ImAxes: Immersive Axes as Embodied Affordances for Interactive Multivariate Data Visualisation

Maxime Cordeil[1]
maxime.cordeil@monash.edu

Andrew Cunningham[2]
andrew.cunningham@unisa.edu.au

Tim Dwyer[1]
tim.dwyer@monash.edu

Bruce H. Thomas[2]
bruce.thomas@unisa.edu.au

Kim Marriott[1]
kim.marriott@monash.edu

[1]Monash University, Melbourne, Australia
[2]University of South Australia, Adelaide, Australia

ABSTRACT
We introduce *ImAxes*, an immersive system for exploring multivariate data using fluid, modeless interaction. The basic interface element is an embodied data axis. The user can manipulate these axes like physical objects in the immersive environment and combine them into sophisticated visualisations. The type of visualisation that appears depends on the proximity and relative orientation of the axes with respect to one another, which we describe with a formal grammar. This straight-forward composability leads to a number of emergent visualisations and interactions, which we review, and then demonstrate with a detailed multivariate data analysis use case.

ACM Classification Keywords
H.5.2 User Interfaces: Information Interfaces and Presentation (e.g., HCI)

INTRODUCTION
The rapid development and commodification of virtual-reality head-mounted display devices in recent years has been largely motivated by obvious opportunities in entertainment. Devices like Microsoft HoloLens represent a similar step-change in the adoption of Augmented Reality for applications like situated architectural walk-throughs [56] and overlays of engineering models on their real-world counterparts [43]. These initial explorations are of applications that are impossible with traditional desktop computing environments. With these types of devices becoming ubiquitous, people are starting to wonder how more traditional computing applications such as data analysis will look in Virtual and Augmented Reality (VR/AR). A new topic of study emerging from the Information Visualisation (InfoVis) and AR/VR research communities is to explore how data analysis can be reimagined with—and benefit from—such emerging display and interaction technologies [9, 1].

UIST 2017, October 22–25, 2017, Quebec City, QC, Canada

DOI: **https://doi.org/10.1145/3126594.3126613**

More fundamentally, in recent years information visualisation researchers have been calling for research into interaction modalities that go beyond Mouse and Keyboard to provide greater freedom of expression [34] and more fluid interaction [16]. In this paper we explore an interaction paradigm built around the notion of data axes as embodied affordances for the construction of a variety of rich, immersive data visuals for exploring multivariate datasets. The axes can be placed anywhere in space and the type of data visual that appears is determined purely by their proximity and orientation with respect to one another.

We introduce *ImAxes* (Immersive Axes), an interactive multi-dimensional visualisation tool that relies on the arrangement of data axes in space. *ImAxes* is built in the Unity game engine, and can be used with most common Virtual Reality headset and tracked controllers. We demonstrate its use with the HTC Vive Virtual Reality headset, which provides a room-sized tracked interaction space that allows users to walk, interact with their full bodies, and provides sufficient space for users to create a whole suite ("gallery") of data visualisations within ImAxes.

ImAxes is a modeless visualisation system. No menus or hotkeys are required; users' actions are all available by direct manipulation of the objects within the virtual environment. We define a formal grammar that allows for the construction of a diversity of data representations based only on the spatial arrangement of axes. We then show how *ImAxes* allows the creation of familiar InfoVis visualisations, and how it allows for the emergence of new types of meaningful visualisations for multivariate data exploration. Then we detail both the designed and emergent interactions of *ImAxes* that we have identified and finish by demonstrating how *ImAxes* can be used in a realistic multivariate data exploration scenario.

RELATED WORK
Interactive Multidimensional Data Visualisation (MDV) is a well-explored subject in Information Visualisation and statistics, yet research is still very active in this domain. One class of approaches use dimension reduction techniques from statistics. For instance, Multidimensional Scaling (MDS) [8] seeks to preserve the distance between the high-dimensional data points when they are projected onto a 2- or 3-dimensional scatterplot, with the aim of showing clusters

and outliers. However, like other dimension reduction techniques, MDS is "lossy" and can be difficult to understand by non-expert users.

Thus, variations of more traditional visual idioms have been developed to show many data dimensions through various types of small-multiple views. This type of relatively straightforward dimension-by-dimension exploration of multivariate data through small-multiple views has proved immensely popular in Business Intelligence (BI) software [44]. Scatterplot matrices (SPLOMs) display all 2D combinations of data dimensions in a matrix form. Users can visually detect patterns between dimension pairs, a process called "Scatterplot Diagnostics" or *Scagnostics* by the Tukeys [58]. However, the usefulness of SPLOM views is limited; as dimensionality increases, a traditional desktop display only offers a limited surface to display data.

Parallel Coordinates Plots (PCPs) are an alternative way to visualise data across dimensions, through axes positioned parallel—rather than orthogonally—to one another [25]. Each data point is represented as a sequence of line segments, intersecting each axis at a position corresponding to the data point's attribute value in the corresponding dimension. Unlike SPLOMs, in PCPs only the relations between adjacent axes are easily visible.

Regardless of the visualisation technique, to avoid cluttered displays when dealing with high-dimensional datasets, interaction is required to filter and interactively explore the data. Focus+context techniques such as brushing and linking of views were introduced to highlight a subset of data points across visualisations [7]. Animations have been used to improve the understanding of correlations between dimensions of a dataset. A compelling example of animation in MDV is the use of 3D rotation to transition between 2D visualisations of dimensional pairs in a scatterplot matrix [15].

Researchers have explored direct interactions with MDV systems using mouse [25, 55] or multitouch [32, 46] and drag-and-drop style interactions to assign data dimensions to visual encodings and are now fairly common in BI tools such as Tableau. However, the affordances for interaction in these standard tools are usually a simple list of attributes in a WIMP interface. The interfaces remain highly constrained by the particular set of visual idioms supported by the tools. Arguably, the state of the art in InfoVis system design is around more flexible systems that use a component model to allow more flexible construction of data visualisations through direct interaction [20, 36, 47], though usually in non-immersive environments that require WIMP interaction.

In this paper, we take inspiration from a proposal for less constrained (yet still mouse- and WIMP-based) interactions with axes components in a 2D canvas style view by Claessen and van Wijk [10]. Their system lets users freely draw lines on the canvas and assign them to data dimensions, link them to other axes, and determine the style of visualisation (PCP or scatterplot) through contextual menus. Using this interface, the authors are able to create a number of highly novel combinations of PCPs and scatterplots. In this paper we also use axes as the fundamental building block, but make the interaction fluid by eliminating the context menu through a grammar based on the positions of the axes that entirely determines the visualisation. By implementing this interface in a virtual environment, the axes become *embodied affordances* for visualisation construction.

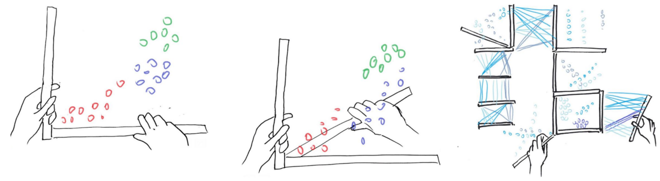

Figure 1. *ImAxes* design concept sketches.

Immersion, Presence and Embodiment In the domain of VR, the terms *immersion*, *presence* and *embodiment* have a range of different definitions. We begin with the notions developed by Slater et al. [17, 31, 50, 51]. *Presence* is a state of consciousness, the (psychological) sense of being in the virtual environment. *Immersion* describes the extent to which technology is capable of delivering a vivid illusion of reality to the senses of a human participant. Thus, an immersive system is capable of producing a sensation of presence in the user. *Embodiment* is when a virtual body is spatially coincident with a user's physical body. The use of VR in this manner enables the user to see through the eyes of that virtual body and the virtual body reacts in concert with the user's actions. In, for example, the HTC Vive VR system, the user's head and hand positions are accurately tracked, supporting a degree of natural proprioception [39].

Dourish takes a broader definition of embodiment relating not-only to the user's body, but to objects in the computing environment having a presence in the user's space that somehow embodies their potential for function. *"What embodied interaction adds to existing representational practice is the understanding that representations are also themselves artifacts. Not only do they allow users to 'reach through' and act upon the entity being represented, but they can also themselves be acted upon—picked up, examined, manipulated, and rearranged."* [14, p. 169] Dourish goes on to imply that a well-designed artefact can transcend a designer's intended use, and be employed by the user to create new meaning. We believe *ImAxes* demonstrates this principle through the set of novel visualisations and interactions that emerge through simple manipulation of the axes as described later in this paper. Our original design concept is illustrated in Fig. 1.

Jacob *et al.* examine *Reality-Based Interactions* [26] as a means of providing a more expressive form of human computer interaction that provides a unification of a set of emerging interaction concepts. In particular, these include four themes: 1) naïve physics, 2) body awareness and skills, 3) environment awareness and skills, and 4) social awareness and skills. *ImAxes* leverages these themes to provide a coherent user experience. The virtual world the user operates in has a consistent set of laws of naïve physics. As with embodied interactions, we will leverage the user's sense of proprioception. *ImAxes* will provide an environment where users will have a sense of their environs and retain proficiencies for negotiating, manipulating, and navigating.

Visualisation in Virtual Reality and 3D InfoVis

VR has been employed for scientific visualisation [6], information visualisation [45] and immersive browsing of images [48]. Virtual environments have also proven effective for other numerous scientific applications such as brain tumour analysis [64], archaeology [33, 52], geographic information systems [4], geosciences [22, 24] or physics [30]. In such so-called *Scientific Visualisation* applications, where the data is inherently spatial, there is little debate about the benefits of 3D rendering and the use of stereoscopic VR and AR for making the best use of people's spatial perception to understand the 3D space of the data.

When visualising *abstract data*—without an inherent spatial arrangement—there is more freedom to adapt the data to the dimensions of the user's display. Researchers in this field, known broadly as *InfoVis*, have been very cautious about the use of 3D representations and mainly focused on 2D visualisation techniques that are well suited to the current desktop display modality. Projecting 3D data visualisations onto a 2D screen suffers from inherent issues due to occlusions, perspective distortion and interaction problems using a 2D input device [41, Ch. 6]. Some previous studies suggest that 3D representations may show high-level structure in 3D shapes and terrain more clearly [53, 54], while others suggest that 2D representations are preferable for precise manipulation or accurate data value measurement or comparison and advocate the use of linked 2D and 3D representations [54, 57]. Other studies showed benefits for 3D network visualisation [59, 60, 21]. There are mixed results on the effectiveness of 3D scatterplots when compared to 2D scatterplots. Studies providing binocular depth cues found benefits [35, 62, 42] while those using only monocular cues were mixed, with one finding benefits [19] but others finding 2D was more effective [61, 49].

VR technology is maturing rapidly. We believe that we need to explore immersive data visualisation to be ready to exploit the technology fully as it becomes mainstream and to support users working in native 3D environments with high-quality stereo rendering, head-tracking and true spatial input devices. While this does not imply that all data visualisations must use a depth dimension, it does raise the questions of, given a good capability to do so, how and when should we take advantage of this capability? In *ImAxes* the third spatial dimension is available for users to use as they please. It is a system for creating and manipulating visualisations (2D, 2.5D and 3D) in a VR immersive workspace. Thus, it has quite a different focus to the studies cited above, each dealing with the efficacy of a single 3D visualisation. While users are free to create 3D scatter plots they are also free to use 2D visualisations and use the third dimension to arrange these in meaningful ways, such as stacking them for comparison [11, 5] or to leverage from so-called Spatio-Data Coordination [12].

IMAXES GRAMMAR AND DEFINITIONS

At a high level, *ImAxes* is a system that allows the creation of a variety of data visualisations based on the direct manipulation of data axes in 3D space. The main steps in this creation are shown in Fig. 2.

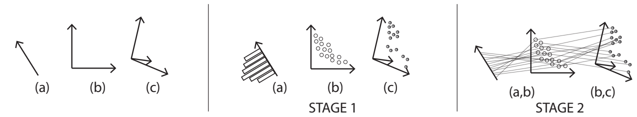

Figure 2. Main steps in *ImAxes* recognition and visualisation. On the left are the symbols in the *ImAxes* grammar (a) S_1, (b) S_2 and (c) S_3. In the middle the corresponding visualisation and on the right facet linking.

Our goals are to define a system that is *expressive* enough to create standard visualisations for understanding multidimensional data: scatterplots, SPLOMs, PCPs and their three-dimensional variants. We want the system to be *generous* in the sense that the user can combine the elements freely, with virtually any combination producing some sort of visualisation, thus encouraging exploration of the data. We want there to be a *physical affordance* so that data visualisations look and behave like physical objects [18]. Finally, we want the system to be *declarative* in the sense that the order of construction should not matter. We feel declarativeness is important in immersive VR because it means that the user can simply imagine something in their mind's eye, and then place the elements in 3D space to create it. The appearance of the object being created gives full information about its construction. This (mostly) accords to our experience in the real world where the appearance of a physical object reveals its components and how they should be placed together to make it.

One of the strengths of *ImAxes* is that it is built upon a simple grammar which formalises the rules for constructing visualisations in the VRE using axes' spatial placement. The formal description of data visualisation using grammars is an active field of research. Notable examples include Mackinlay [37] who formalised the primitive languages and operators that define visual information, the formalisation of two-dimensional visual languages like circuit diagrams [38], the "Grammar of Graphics" [63] which provides a language to create statistical charts, and grammar-based approaches for defining buildings and other objects in virtual worlds [40, 23]. Our contribution is a grammar of graphics based on the spatial agency of *data axes* in a 3D immersive virtual environment. The great strength of a grammatical approach is that grammar rules are compositional, allowing users to combine the basic elements in open-ended and unforeseen ways.

The basic element in the *ImAxes* grammar is the *axis*. An *ImAxes* workspace involves a set of axes A. An axis is represented by the tuple $axis(s,e,d)$ where s and e are the positions in 3D space of the axis start and end points, and d is the data attribute linked to the axis. This is the only *token* in the grammar. The three other symbols in the grammar model are the 1-, 2- and 3-D structures that can be built with axes:

$S_1(\vec{v},A_1)$ – linear concatenation of axes. $\vec{v}$ is the canonical unit vector giving the direction of the axis and $A_1 \subseteq A$

$S_2(\vec{v}_1,\vec{v}_2,A_1,A_2)$ – combination of axes forming a 2D SPLOM. $\vec{v}_1$ and $\vec{v}_2$ are orthogonal vectors, axes in A_1 parallel to $\vec{v}_1$, axes in A_2 parallel to $\vec{v}_2$.

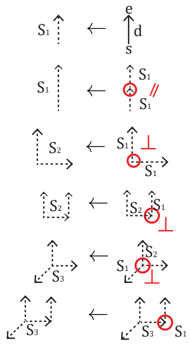

$S_I(dir(s,e),\{axis(s,e,d)\})$	$\leftarrow axis(s,e,d)$		(R_1)
$S_I(\vec{v}_1, A_1 \cup A_2)$	$\leftarrow S_I(\vec{v}_1, A_1), S_I(\vec{v}_2, A_2)$	$s.t. \parallel(\vec{v}_1, \vec{v}_2) \wedge \circledast(A_1 \cup A_2)$	(R_2)
$S_2(\vec{v}_1, \vec{v}_2, A_1, A_2)$	$\leftarrow S_I(\vec{v}_1, A_1), S_I(\vec{v}_2, A_2)$	$s.t. \perp(\vec{v}_1, \vec{v}_2) \wedge \circledast(A_1 \cup A_2)$	(R_3)
$S_2(\vec{v}_1, \vec{v}_2, A_1, A_2 \cup A_3)$	$\leftarrow S_2(\vec{v}_1, \vec{v}_2, A_1, A_2), S_I(\vec{v}_3, A_3)$	$s.t. \parallel(\vec{v}_2, \vec{v}_3) \wedge \circledast(A_1 \cup A_2 \cup A_3)$	(R_4)
$S_3(\vec{v}_1, \vec{v}_2, \vec{v}_3, A_1, A_2, A_3)$	$\leftarrow S_2(\vec{v}_1, \vec{v}_2, A_1, A_2), S_I(\vec{v}_3, A_3)$	$s.t. \perp(\vec{v}_1, \vec{v}_2, \vec{v}_3) \wedge \circledast(A_1 \cup A_2 \cup A_3)$	(R_5)
$S_3(\vec{v}_1, \vec{v}_2, \vec{v}_3, A_1, A_2, A_3 \cup A_4)$	$\leftarrow S_3(\vec{v}_1, \vec{v}_2, \vec{v}_3, A_1, A_2, A_3), S_I(\vec{v}_4, A_4)$	$s.t. \parallel(\vec{v}_3, \vec{v}_4) \wedge \circledast(A_3 \cup A_4)$	(R_6)

Figure 3. Grammar Rules (R_i) for *ImAxes* where $A_j \subseteq A$. There exist symmetric rules for $R_4 : A_1 \cup A_3$ and $R_6 : A_1 \cup A_4, A_2 \cup A_4$.

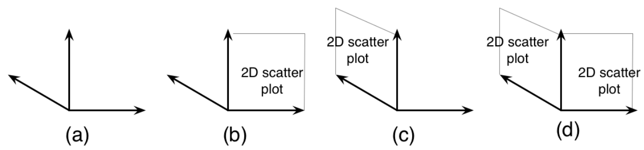

Figure 4. Multi-set grammar reduction vs set grammar reduction. (a) shows 3 axes, not all of which are orthogonal so it is not a 3D scatterplot. In a multi-set grammar this is ambiguous and can lead to either recognising (b) or (c). However in a set grammar (d) there is no ambiguity and two 2D scatterplots are recognised through repeated application of (in this case) R_1, R_3 and R_4.

$S_3(\vec{v}_1, \vec{v}_2, \vec{v}_3, A_1, A_2, A_3)$ – axes forming a 3D SPLOM, axes in A_i parallel to $\vec{v}_i$ and $\vec{v_{1...3}}$ mutually orthogonal.

The grammar rules for constructing these objects are straightforward, Fig. 3. Reflecting the aim for physical affordance we require the axes in each of the various symbols to actually touch the other, so that they form a kind of physical skeleton for the data visualisation. Reflecting the aim for generosity we do not restrict the arrangement beyond this, except that the axes must be parallel or orthogonal to the other axes in the object.

The rules capture that an S_I is either a single axis or two parallel connected S_I; that a S_2 is made up of two orthogonal S_I which can be extended by adding other connected S_Is that are parallel with one the original S_Is; and, that a S_3 is made up of an S_2 and an orthogonal S_I and can be extended by adding other connected S_I. The rules make use of the following functions:

$\parallel(\vec{u}, \vec{v})$ – holds if vectors $\vec{u}$ and $\vec{v}$ are parallel;

$\perp(\vec{u}, \vec{v})$ – holds if $\vec{u}$ and $\vec{v}$ are orthogonal, similarly $\perp(\vec{u}, \vec{v}, \vec{w})$;

$dir(s,e)$ – returns the unit vector along line (s,e); and

$\circledast(A')$ – holds if all axes in $A' \subseteq A$ are spanned by a tree through pairs of axes touching at their ends.

One of the more complicated questions was the underlying semantics of the grammar rules: grammars for visual languages are typically either multi-set grammars which behave like traditional string grammars in that when a symbol is recognised its component symbols are consumed by the reduction, or a set grammar (also formalised as a definite clause logical formula) in which the component symbols are not consumed but can instead be involved in other reductions [38].

We decided to formalise reduction using the set grammar semantics in which symbols are not consumed in a reduction. This was driven by the desire for declarativeness. The example in Fig. 4 shows the problem. With multi-set grammar semantics this will be recognised as a S_I and a S_2, but the result is not deterministic and depends upon the order of reduction. Thus, the grammar is ambiguous, not declarative and two identical collections of axes may not give rise to the same visualisation, with the result depending upon the order of reduction. On the other hand, with set grammar semantics this will be recognised as two S_2. This reading also accords with physical affordance as physical objects can belong to more than one higher order shape: for instance a corner post in a house forms part of the two walls.

We therefore use set grammar semantics in reduction: symbols are not consumed. This ensures that the result of parsing is deterministic and that it is declarative. The result of grammar reduction is the set of *maximal* symbols, as we do not want intermediate results. More exactly, a symbol s_1 contains symbol s_2 if $axes(s_2) \subset axes(s_1)$ where $axes(s)$ returns the set of axes in the symbol. A symbol is *maximal* if no other symbol contains it. Thus, if a 2-D scatterplot forms part of a 3-D scatterplot, only the 3-D scatterplot will be recognised.

After the various grammar symbols have been recognised there are two more steps in the *ImAxes* visualisation pipeline. The next step is to generate the associated data visualisation for each of the symbols (see Fig. 2):

$S_I(v,A)$ – a histogram showing frequency distribution for each axis $a \in A$.

$S_2(v_1, v_2, A_1, A_2)$ – a 2D SPLOM with a scatterplot for each pair of connected orthogonal axes $a_1 \in A_1$ and $a_2 \in A_2$.

$S_3(\vec{v}_1, \vec{v}_2, \vec{v}_3, A_1, A_2, A_3)$ – a 3D SPLOM with a 3D scatterplot for each triple of connected orthogonal axes.

Note that 2D and 3D scatterplots are just simple cases of the 2D and 3D SPLOMs respectively. Also note that the 2D and 3D scatterplots will be overlaid if there is more than one axis for one of the dimensions.

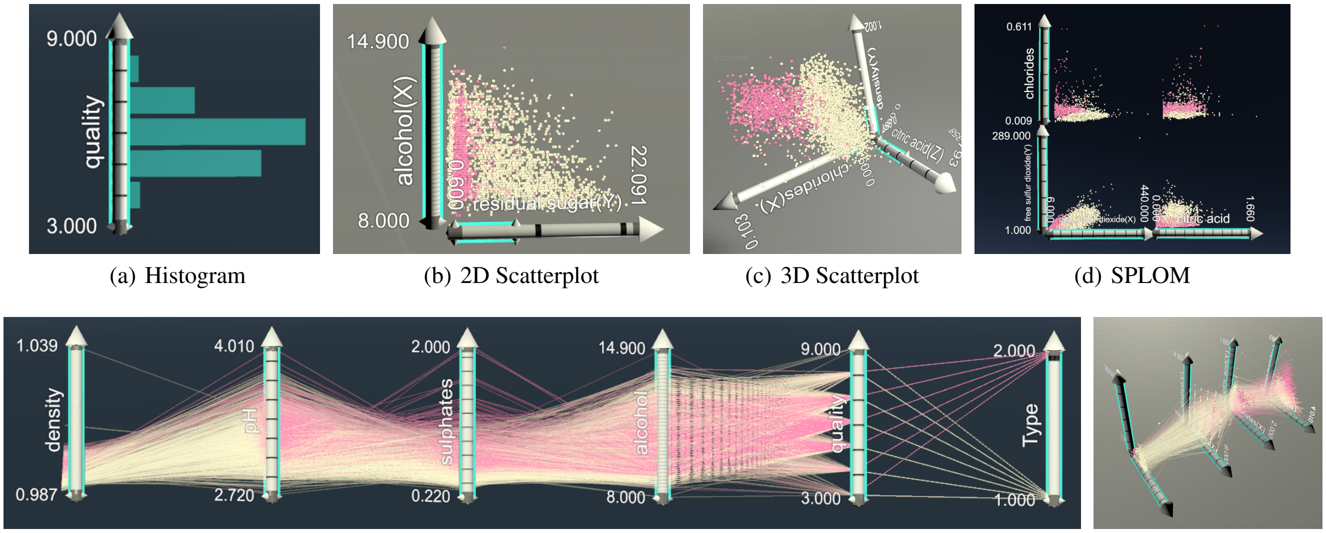

(a) Histogram (b) 2D Scatterplot (c) 3D Scatterplot (d) SPLOM

(e) Parallel Coordinate Plot (PCP) (f) Linked scatterplots

Figure 5. Familiar 2- and 3-D visualisations built with *ImAxes*.

The final step is to visually link elements in these different visualisations, as illustrated in *STAGE* 2 of Fig. 2. Each of the visualisations (histograms, 2D and 3D scatterplot) is called a visualisation *facet*. We use lines to link the same elements in nearby facets. This is inspired by similar linking in 2.5D visualisations. The centres of the facet must be within a threshold distance from each other. A facet can be linked to more than one other facet. Remarkably, as we shall see, such *facet linking* allows us to produce PCPs, 3D PCPs and a wide range of other charts in a very natural way.

IMAXES VISUALISATIONS

In *ImAxes*, specific manipulation and arrangement of axes in 3D space enables the straight-forward construction of familiar InfoVis visualisations in two and three dimensions. However, the user is free to place axes in configurations that correspond to no known InfoVis idiom, but which can still reveal useful properties of the data.

By directly manipulating and positioning axes in space, *ImAxes* uses embodied spatial data mapping, i.e. when a user builds a scatterplot the axis parallel to the ground corresponds to the X axis of a visualisation, the Y axis is the axis perpendicular to the ground, and the Z axis is the axis looking at, or away from, the user's point of view.

Familiar Visualisations

We first present how standard InfoVis idioms can be produced with *ImAxes* (see Fig. 5), and later we show how visualisations that we did not predict are produced.

Histograms – A lone axis (S_1) shows a histogram of the data distribution (Fig. 5(a)).

Scatterplots – Positioning two axes perpendicular to each other creates a 2D scatterplot (a 1×1 S_2), Fig. 5(b). Adding a third axis perpendicular to the first two extends this to a 3D scatterplot (a $1 \times 1 \times 1$ S_3), Fig. 5(c).

Scatterplot Matrices (SPLOM) – 2D scatterplots can be extended to SPLOMs by aligning additional axes along the X and Y axes. Fig. 5(d) shows a 2×2 S_2.

PCPs – Parallel Coordinates Plots are created in *ImAxes* by positioning a series of parallel axes nearby each other. Fig. 5(e) is simply 5 parallel S_1.

Linked Scatterplots – Two or more 2D scatterplots placed near each other create a 2D linked visualisation [11], or 2.5D PCP.

Emergent Visualisations

The spatial manipulation and positioning of axes and visualisation in *ImAxes* lead to the emergence of new and useful types of visualisations that we did not anticipate in designing our grammar and system. See Fig. 6.

3D circular connected Parallel Coordinate Plots – For example, it is possible to arrange axis circularly to form a closed PCP. Placing an axis at the centre of this type of circular PCP creates a one-to-many linked view. In Fig. 6(a) the visualisation shows the relation of the central axis to all the other axes. This type of linked view is a novel approach to perform focus+context operations where the focused object is the central axis.

Connected "plot superposition" – The introduced grammar makes it possible to overlay scatterplots. When positioning three axes in a 'U' shape, two S_2 are recognised. Since they are closer than the linking threshold, links appear between similar points in each plot. We found this type of visualisation is useful when visually inspecting the relations between two visualisations that share same axi(e)s.

Similarly, a 3D connected overlaid plot is produced when placing a fourth axis aligned with one of the three axes of an existing 3D scatterplot. The two 3D scatterplots appear at the same position and get linked. Obviously, such superposition

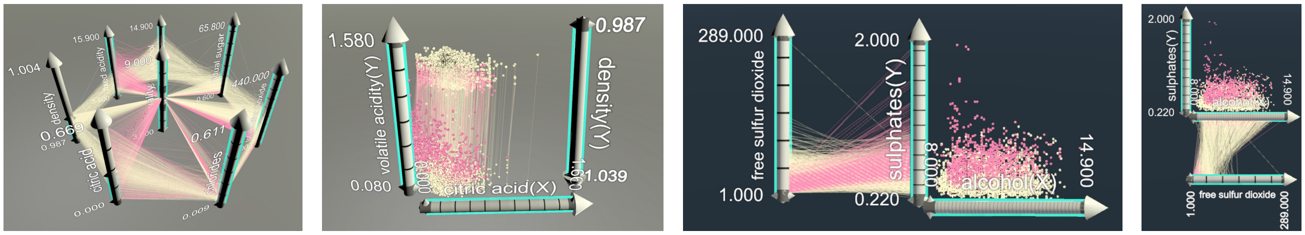

(a) Circular 1 to many PCP (b) Overlaid SPLOM (c) Linking to Y-axis (d) X-axis link

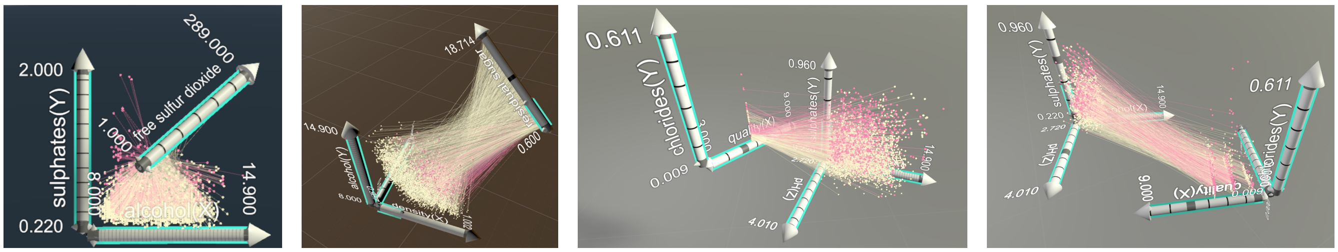

(e) Linking to data points (f) Links to 3D SPLOM (g) Linked 2- and 3-D SPLOMs (h) Linked 3D SPLOMS

Figure 6. Emergent Visualisations.

is only useful when the data in each plot are well separated or sparse.

PCP Brushing – Links appear between two SPLOMs when they are within a given distance. To limit clutter, when a single axis *a* links to a SPLOM *s* with more than one axes, the system considers the distances between the centre of *a*, the centre of *s* and the centre of each axis contained in *s*, and links are shown only between the closest pair. For example, Figs. 6(d),6(c),6(e) show links appearing between a wand axis and the *x*-axis, the *y*-axis and the data points, respectively. Thus, the user can quickly "wave" the single axis around a SPLOM (like a TSA agent) and quickly compare correlations between their wand axis and each of the axes in the SPLOM or with the data points inside the SPLOM. We call this feature PCP Brushing (or Angel Dusting).

Linked 3D Scatterplots – The same rules that cause linking between 1- and 2-D scatterplots also apply to 3D scatterplots. Thus, we can link a 2D plot to a 3D plot (Fig. 6(g), or a 3D plot to another 3D plot (Fig. 6(h)). We can also use a whole plot like a wand for Angel Dusting.

3D SPLOM – We discovered using *ImAxes* that it is possible with our system to form 3D scatterplot matrices from existing 3D scatterplots, Fig. 7(a).

Tree of linked visualisations – As with [10], *ImAxes* offers the freedom of linking plots in many different ways, however in 3D it is practical to link many 2- or 3-D plots to each other. A compelling example is the fabrication of a hierarchy tree linking 2D scatterplots, Fig. 7(b), linking along logical paths defined by categorical data. For example, in this tree of 2D linked visualisations, the user can follow the divergent red and the yellow paths through different pairwise plots.

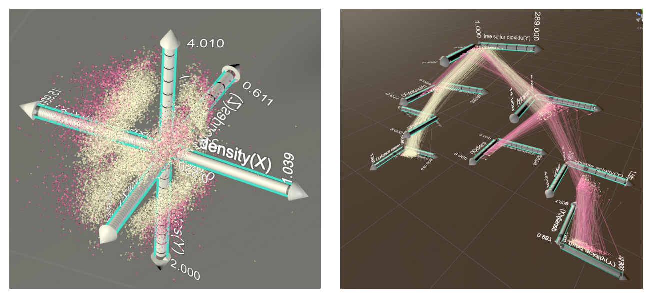

(a) 3D SPLOM (b) Tree linking

Figure 7. Complex emergent visualisations.

IMAXES INTERACTIONS AND METAPHORS

A principle of *ImAxes* is to provide the user with extra spatial degrees of freedom to explore multivariate data in a Virtual Reality Environment (VRE). Building the SPLOM and visual link elements described by our grammar requires the free creation and spatial manipulation of axes and the derived visualisations. Hence it is natural to leverage 6 DOF tracked hand controllers to directly manipulate these elements in 3D space, letting the user create, orient, position and discard elements with natural actions. Similar to a mouse drag-and-drop interaction, ImAxes uses a *grab and place* interaction to manipulate visualisation elements in a VRE. This interaction consists of colliding a hand controller with a visualisation element and pulling the controller's trigger, thereby attaching the element to that controller. Moving the controller in space will directly manipulate the position and orientation of that element. Manipulating the low-level axis elements will stretch and distort associated higher order elements (SPLOMs and linked visualisations) to accommodate the manipulations. Grabbing and manipulating higher-level elements will directly manipulate any associated lower-level elements. For example, grabbing and rotating a SPLOM will rotate the associated axes of that

SPLOM in the expected manner to maintain its appearance. Releasing the controller's trigger will "place" the element in its current position.

Attribute Shelf

Data exploration in *ImAxes* begins with a collection of axes loaded from a dataset, arranged in a series of rows, to invoke a book shelf metaphor. A user can grab any of these axes and—when pulled from their origin beyond a certain threshold—will duplicate the axis in the user's hand and snap the original axis back to the shelf. We leverage haptic feedback in the controllers to provide users with a cue to this threshold.

Axis Interactions

Each axis corresponds to a data attribute from the dataset. An arrowhead at the end of the axis indicates the direction of increasing attribute value, and minimum and maximum value labels display the attribute range. A text label identifies the data attribute. All labels are dynamically reoriented to face the user.

To support the direct manipulation metaphor, filtering and adjusting the range of the axis occurs directly on the axis itself. Axes have two pairs of interactive widgets that allow the user to perform:
scaling – adjust the range $[d_{min}, d_{max}]$ that defines the domain of the attribute represented by the axis, Fig. 8(a) (1).
filtering – adjust the range $[f_{min}, f_{max}]$ that will filter data points outside of this range, Fig. 8(b) (2).
The data points that are not within the specified ranges disappear from the visualisation.

These widgets initially appear in-line with the axis (Fig. 8(a), 8(b), left). When a hand controller approaches, indicating that the user may want to manipulate the axis, these widgets animate out of alignment with their axis to allow better targeting and manipulation (Fig. 8(a), 8(b), centre). This reduces clutter in the VRE, while still affording direct manipulation of the widgets when needed. The direct manipulation of the widgets results in a continuous animation when *scaling* the axis (Fig. 8(a)) or performing a range *filtering* (Fig. 8(b)). The fluid control and feedback of the widgets keep users focused on their task. Additionally, this interaction supports bi-manual manipulation of a single axis as the user can position and orient an axis with one hand, and simultaneously adjust one of the axis range or filter widgets with the other hand.

We implemented a "throw it away" metaphor to discard axes, whereby a user grabs and throws the axis as they would a crumpled piece of paper. This causes the axis to follow the momentum of the throw and shrink to nothing.

Emergent Interactions

The affordances of VREs, coupled with our defined grammar, provide the emergence of new interaction types that enable novel multivariate exploration paths.

Axis Swipe – Dragging an axis across an array of axes or visualisations results in the successive apparition and disappearing of visualisations as defined by our grammar. This provides the user with a fluid and continuous method to quickly explore the relationship between an axis and a set of other axes or visualisations.

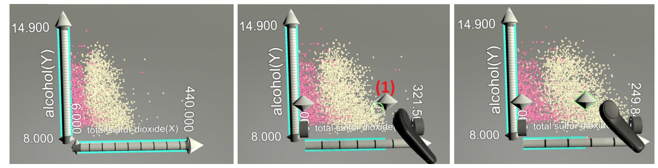

(a) Interactive scaling of an axis

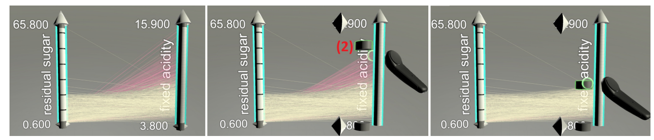

(b) Interactive filtering of an axis

Figure 8. Axes Interactions

Brushing with Motion – Manipulating (even just wiggling) an axis that is part of a composite visualisation causes the points and/or lines that are associated with that axis to move proportionally to the displacement of the axis. This causes these data elements to perceptually "pop" relative to any other points or clutter in the field of view. In our explorations we found this to be a very useful feature, although this style of "brushing with motion" [3] was not explicitly designed or coded. It simply emerged from the positional coupling of points to axes.

Parallel Plot Distortion – Graphic occlusion is a common problem in Information Visualisation. For example, visualising a dense PCP can result in clusters of line segments obscuring other line segments, making the mental task of delineating relationships difficult. ImAxes alleviates this issue as users are able to directly manipulate axis within 3-space and distort the associated visualisations, effectively twisting PCPs to separate overlapping clustered line segments within the visualisation. This is further enhanced through bi-manual input and the stereoscopic nature of VREs, as the user can manipulate both axes of the PCP and make even slight head adjustments to better understand the presented relationships.

Fluid visualisation switching – *ImAxes* enables users to switch between visualisation types using continuous hand and arm gestures. E.g. holding two axes parallel and rotating a wrist fluidly switches the visualisation between a parallel plot and a scatterplot, providing the user agency in perceiving the relationships between the two visualisations through one continuous gesture.

Embodied Queries – We described earlier how entire 2- or 3-D plots or SPLOMs may be grabbed and used as brushes against other plots or SPLOMs, simply by holding them close enough that links appear between either the data points in the two visualisations, or between the axes. By further restricting the set of data by filtering the axes and using the plot as a brushing (or angel dusting) "wand" the object effectively becomes a sophisticated *embodied query* tool.

Gallery/Workspace – In practical data analysis scenarios (such as the one described in the next section) the user typi-

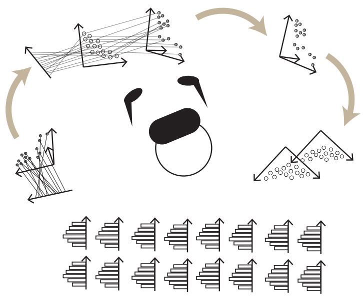

Figure 9. An example *ImAxes* workspace. We have found a useful workflow when exploring data is to construct a sequence of visualisations, each covering a different aspect of the data. The user rotates his body position to begin a new visualisation in a new space, while all the previous visualisations and the Attribute Shelf remain within easy reach.

cally creates a number of independent sets of data visuals as they explore different data dimensions or particular subsets of the data. They are free to place these anywhere in their surroundings. We have noticed that in such scenarios users tend to arrange them in a circular pattern such that the different data visuals all remain within arms reach and the user can shift their attention to a different visual simply by rotating in-place. Thus, the space around them becomes a "workspace" and they are able to arrange the visuals in an order that makes sense to their analysis, Fig. 9. In presentation scenarios, the workspace becomes a "gallery" through which they can guide other users as they explain the results of their analysis.

TECHNICAL IMPLEMENTATION

ImAxes is developed in the Unity game engine and runs on a Windows 10 PC with an Intel i7, and Nvidia GTX 1080. We chose the Unity engine due to its support for VR devices (Oculus VR, HTC Vive) and its support for rapid development of complex interactive VRE applications. At a high level, *ImAxes* is composed of a generalised visualisation toolkit, axis and visualisation facet objects, and the grammar recogniser.

We developed a generalised visualisation toolkit—the Immersive Analytics Toolkit (IATK)—capable of building common visualisations including histograms, PCPs, and scatterplots. This toolkit is responsible for loading multivariate datasets, reading dataset metadata, and generating the geometry to represent the visualisations. The IATK is intended to be application agnostic and as such the visualisation geometry produced by IATK is generated in normalised model coordinates, $-0.5 \leq x, y, z \leq 0.5$. It is the responsibility of the implementing application to project this visualisation geometry into the required coordinate space.

The grammar recogniser is incorporated into the update cycle of ImAxes as presented in Fig. 10. Every update cycle, the *ImAxes* grammar recogniser is responsible for parsing the spatial layout of axis tokens within the scene and creating or destroying the appropriate visualisation facets based on the grammar rules presented previously in this paper. These visualisation facets are composed of the geometry generated by the IATK. Visualisation facets maintain references to their constituent axes and are responsible for positioning and ori-

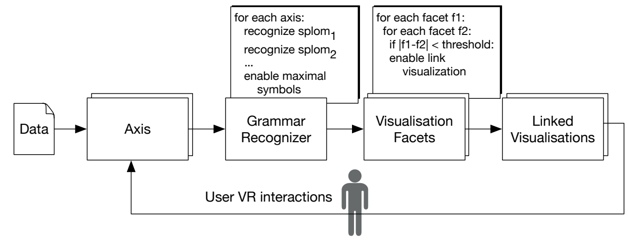

Figure 10. *ImAxes* system diagram showing the grammar recognition process that occurs during the update cycle.

enting themselves relative to their axes. The facets use distortion techniques to stretch and skew visualisations to fit within their axes and to provide immediate, reactive feedback to the user's interactions. With this system design, expensive geometry generation algorithms are infrequently called, relying on shader programs to manipulate, reproject, and even filter the geometry. As a final process in the update cycle, *ImAxes* calculates the Euclidian distance between each visualisation facet and, when below a designated threshold, a linked visualisation is created to connect the two facets visually.

USE CASE

In this section we provide a scenario that demonstrates the use of *ImAxes* as a multivariate data visualisation tool. Luís, a sommelier data scientist, wishes to perfect his knowledge of *vinho verde* wines from Portugal. He has in his possession a wine dataset from this region of the world [13]. The dataset contains 6,497 samples of wine (1,600 *red* and 4,897 *white*). Each wine sample is described by 12 data attributes (Fig. 11).

Overall, Luís wants to gain general knowledge of the differences between white and red wines of this region (*goal 1*), and eventually understand the features that qualify wines as mediocre or excellent (*goal 2*). The data contains a *Quality* rank ranging from 0 (very bad) to 10 (excellent).

The sommelier sets up the color binding to pink for the red wine data type and yellow for white wine data type in the metadata control file of *ImAxes* . When *ImAxes* starts, he views all the data attributes on the *attribute shelf* (Fig. 11).

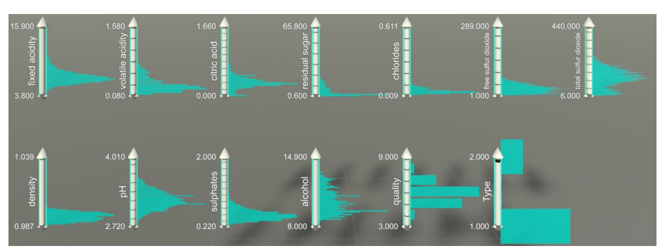

Figure 11. When *ImAxes* starts Luís views all the data dimensions arranged in the attribute shelf.

Luís locates and grabs the *type* axis which contains the two categorical values: *white* (1) and *red* (2).

Goal 1: Red versus White Wine Analysis

To gain a general overview of the differences between white and red wines in the dataset, Luís swipes the *type* axis across the shelf. As the axis is swept across the other data attributes, temporary PCP visualisations appear and Luís views emerging patterns at a glance (Fig. 12).

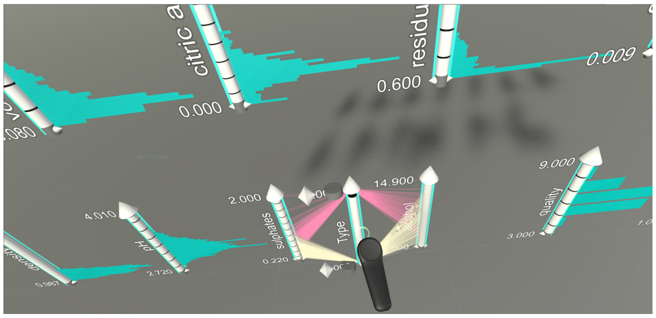

Figure 12. As Luís swipes the *Type* axis across the attribute shelf, some PCPs showing different spread of data may pop-up.

Swiping the attribute shelf and observing some emerging patterns leads Luís to further investigate the *Alcohol*, *Ph*, *Residual Sugars*, and *Sulphate* axes. His primary interest is to perform pairwise comparisons between the *Type* and these four axes.

For the first steps, Luís uses *ImAxes* to verify some initial comparisons between red and white wines. Luís holds the *Type* axis in his right hand, parallel to the *Alcohol* axis in his left hand. A PCP is formed (Fig. 13, left) and by tilting the axes back and forth, the effect of *brushing with motion* makes clear the spread of alcohol values for both types of wine. By examining the PCP, Luís notices a pattern indicating that the red wines have a higher alcohol concentration than the white wines. Luís continues his investigation by inspecting the acidity attribute denoted by the *pH* axis (Fig. 13, right). He twists the PCP to resolve occlusion between the *pH* values spread of the two wine types. The sommelier instantly sees that the red wines from this wine region are less acidic than white wines. Luís just confirmed his general knowledge about degree of alcohol and acidity for red and white wines.

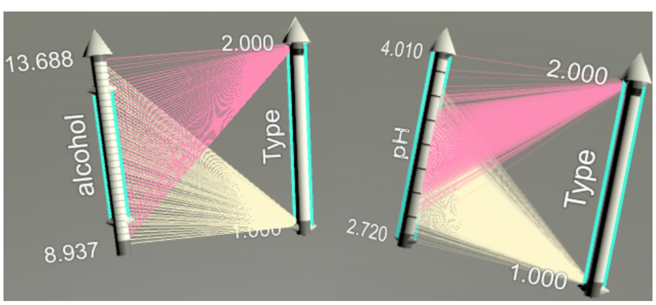

Figure 13. Visual exploration of *Type* against *Alcohol* and *Ph*.

Residual sugars can indicate an incomplete fermentation, or that the winemaker added extra sugar to the wine. Luís wants to refine his analysis and inspect the level of residual sugars for both red and white *vinho verde* wines. He creates a 2D scatterplot between *Residual Sugar* and *Alcohol* to observe trends between the two types of wine. By scaling the *x*-axis, Luís sees that the concentration of alcohol for white wines decreases with the quantity of residual sugars. However, the residual sugar level in red wines is constant with alcohol concentration (Fig. 14 (top left)).

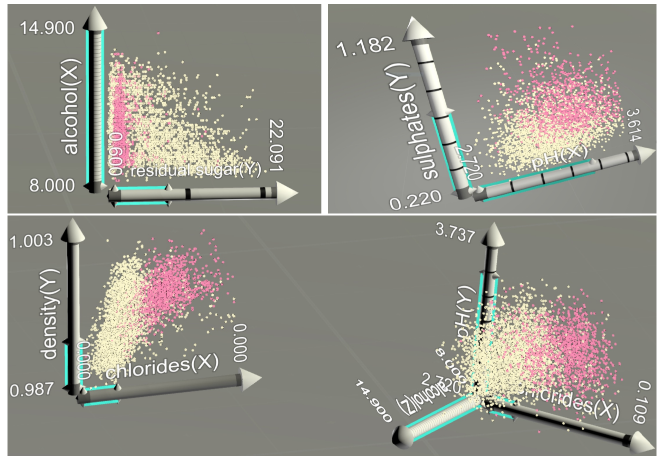

Figure 14. Plot of *Alcohol* against *Residual Sugar* (top left). *(Sulphate, Ph)* 2D scatterplot (top right). *(Chlorides, Density)* plot (bottom left), extended with *Alcohol* (bottom right)

Luís has finished exploring familiar wine attributes. He is now interested in exploring more combinations of attributes that separate red wines and white wines.

Luís examines the *pH* and the *Sulphates* attributes. Through bimanual manipulation of these two axes, a PCP is first formed and then transformed dynamically into a 2D scatterplot. This *(Ph,Sulphates)* 2D scatterplot reveals a clear visual pattern; two clusters emerge between red and white wine: red wines have higher concentration of Sulphates and are more alkaline (i.e. have a higher pH) than white wines (Fig. 14 (top right)). Then, Luís wants to look for three-way relationships between data attributes, starting with the creation of a 2D scatterplot of *Density* and *Chloride*, producing another clear cluster separation between the two wine types (Fig. 14, bottom left). Luís extends this 2D scatterplot to a 3D scatterplot by placing the *Alcohol* attribute on the depth axis (Fig. 14, bottom right). By removing outliers, Luís observes a clear separation between red and white. Luís hence understands that these three attributes are determining factors that differentiate red and white wines.

Goal 2: inspecting wine quality

Luís' second goal is to inspect the ranking of good quality red wines and what makes poor quality red wines from the *vinho verde* region. To do so, he selects three attributes that he is familiar with to determine wine quality, and wants to understand the interplay between them. He selects the *alcohol*, *Volatile acidity*, and *sulphates* attributes from the shelf.

Luís assembles a 3D scatterplot with those attributes of interest. He grabs the *Quality* axis and uses the *PCP Brushing* interaction. However, the resulting connected visualisation between the axis and the 3D scatterplot is too visually cluttered, so he restricts the filter to show only good wines. Luís finds out that he can assemble the *Quality* axis with *Type* in order to formulate his query: 'I want to view the link between the *good red wines*, and my 3D scatterplot' (thus, he has created an *embodied query*). This query allows him to filter out low values on the *Quality* axis, and filter out white wines on the *Type* axis. He creates a similar query visualisation for poor quality red wines.

He uses these two visualisations as brushes on the 3D scatterplot (Fig. 15). The visual links (in blue on the figure to improve print legibility) help him to understand that good red wines have high degrees of alcohol and sulphates, but have few volatile acidity. Conversely, bad red wines have less concentrated alcohol, contain less sulphates and have more volatile acidity.

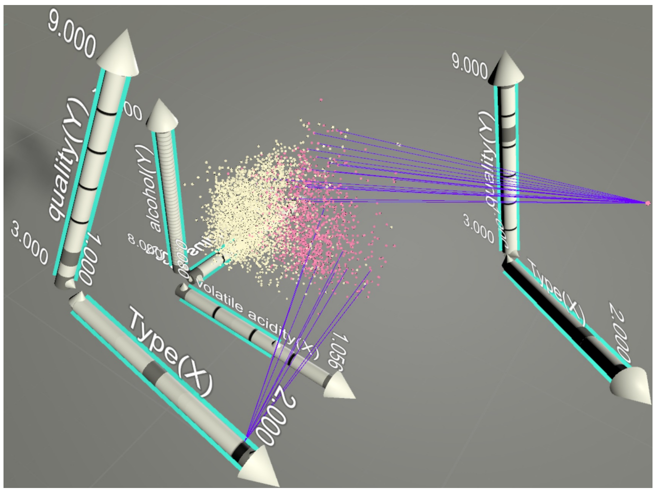

Figure 15. A 3D scatterplot brushed by two 2D filtered scatterplots reveals the attributes of good and bad quality wine

This use-case illustrates the use of *ImAxes* in the context of MDV, demonstrating the traditional visualisations and interactions and the more advanced *embodied query* concept. However, we recognise that this illustrative scenario cannot be considered a *case study* and that in the future we acknowledge that empirical validation of *ImAxes* is needed.

LIMITATIONS

Number of axes. Conceptually, the number of axes that a user can reasonably handle and interact with is limited. Over this limit, usability issues such as the need for the user to walk, crouch or reach to access the axes may occur. For the system to be scalable to extremely large and complex systems of axes, some sort of semantic zoom would be required to allow the user to expand and collapse subsets of axes to focus on different levels of detail.

Number of data points. The size of the dataset is a common factor that impacts the visualisation's legibility and interaction. *ImAxes* takes advantage of hardware graphics accelleration through shader programs and efficient use of geometry such that we can display tens-of-thousands of datapoints with responsive rendering. However, the desire to explore even larger data sets, and the ease with which many plots can be composed with *ImAxes* means that both graphics system and user-perception limitations can still be reached before long. Semantic zoom would also help in this regard.

Interactive Visualisations. Our underlying *Immersive Analytics Toolkit* is at the moment limited to the presentation of table data using histograms, 2D and 3D scatterplots, PCPs, and basic filtering and normalisation. We intend to refine the toolkit to support more multidimensional, multivariate data such as graphs, multimedia data and geovis data, and provide the user with more control over the data binding to visual variables.

Automated analytics. ImAxes currently does not integrate automated presentation algorithms nor analytics tools as part of the system. While we recognise that automation and statistical algorithms play a major role in sensemaking, a balance needs to be achieved between user interaction and system automation. In *ImAxes*, we made the decision to focus on user-centric visual analytics by exploring fluid interaction in a 3D VRE to explore multidimensional data in the VRE. The interactions in this paper do not preclude the use of automation but should be viewed as novel expressive methods to create, brush and query visualisations.

Evaluation. As previously discussed, we consider it necessary to conduct a comprehensive evaluation of *ImAxes* but this validation must go beyond a Likert UX usability validation or base-line comparison. Because of the novelty of immersive visualisations in VREs, established reference evaluation frameworks do not currently exist that we can build upon. Only very recently have InfoVis researchers started evaluating *beyond the desktop* visualisation such as Mixed Reality [2], physical visualisations [28, 29] and proposing increments of the standard models of InfoVis [27]. Consequently, we believe that developing and using such a framework that is theoretically sound would go beyond the scope of this paper.

CONCLUSION AND FUTURE WORK

ImAxes is a multivariate data visualisation tool in a virtual environment that provides highly fluid interaction through an embodied axis metaphor. This is enabled by our declarative spatial grammar which is expressive enough to capture a family of known data visualisation idioms as well as to allow the construction of visualisations that we have not seen before. Furthermore, the fluid interaction that it supports, enables a number of novel interactions. In particular, we find the way a whole visualisation can become an *embodied query* tool that can be applied to brush against other tools using the PCP links to be both novel and compelling.

We are excited by the many directions for extending the *ImAxes* concept. Obviously, it can be extended to other types of visualisations such as time series or graph data. We would like to explore adaptive bundling mechanisms to reduce link clutter. More significantly, there are many possibilities for collaborative data analysis. Compared to desktop systems, the physicality of the VRE allows collaborating analysts to use the usual physical social cues and behaviours, such as gesticulation, passing, personal and shared spaces and so on, to collaborate more naturally and effectively. For example, one user can create an *embodied query* and hand it to another user, for her to try brushing against a visualisation in her personal space.

REFERENCES

1. Bach, B., Dachselt, R., Carpendale, S., Dwyer, T., Collins, C., and Lee, B. Immersive analytics: Exploring future interaction and visualization technologies for data analytics. In *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces*, ISS '16, ACM (New York, NY, USA, 2016), 529–533.
2. Bach, B., Sicat, R., Beyer, J., Cordeil, M., and 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* (2018).
3. Bartram, L., and Ware, C. Filtering and brushing with motion. *Information Visualization 1*, 1 (2002), 66–79.
4. Bennett, R., Zielinski, D. J., and Kopper, R. Comparison of Interactive Environments for the Archaeological Exploration of 3D Landscape Data. In *IEEE VIS International Workshop on 3DVis* (2014).
5. Brandes, U., Dwyer, T., and Schreiber, F. Visual understanding of metabolic pathways across organisms using layout in two and a half dimensions. *Journal of Integrative Bioinformatics (JIB) 1*, 1 (2004), 11–26.
6. Bryson, S. Virtual reality in scientific visualization. *Communications of the ACM 39*, 5 (1996), 62–71.
7. Buja, A., McDonald, J. A., Michalak, J., and Stuetzle, W. Interactive data visualization using focusing and linking. In *Proceedings of the 2Nd Conference on Visualization '91*, VIS '91, IEEE Computer Society Press (Los Alamitos, CA, USA, 1991), 156–163.
8. Buja, A., Swayne, D. F., Littman, M. L., Dean, N., Hofmann, H., and Chen, L. Data visualization with multidimensional scaling. *Journal of Computational and Graphical Statistics 17*, 2 (2008), 444–472.
9. Chandler, T., Cordeil, M., Czauderna, T., Dwyer, T., Glowacki, J., Goncu, C., Klapperstueck, M., Klein, K., Marriott, K., Schreiber, F., et al. Immersive analytics. In *Big Data Visual Analytics (BDVA), 2015*, IEEE (2015), 1–8.
10. Claessen, J. H., and Van Wijk, J. J. Flexible linked axes for multivariate data visualization. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 2310–2316.
11. Collins, C., and Carpendale, S. Vislink: Revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (2007), 1192–1199.
12. Cordeil, M., Bach, B., Li, Y., Wilson, E., and Dwyer, T. A design space for spatio-data coordination: Tangible interaction devices for immersive information visualisation. In *Proceedings of the 10th IEEE Pacific Visualization Symposium (PacificVis)* (2017).
13. Cortez, P., Cerdeira, A., Almeida, F., Matos, T., and Reis, J. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems 47*, 4 (2009), 547–553.
14. Dourish, P. *Where the action is: the foundations of embodied interaction*. MIT press, 2004.
15. Elmqvist, N., Dragicevic, P., and Fekete, J.-D. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE transactions on Visualization and Computer Graphics 14*, 6 (2008), 1539–1148.
16. Elmqvist, N., Moere, A. V., Jetter, H.-C., Cernea, D., Reiterer, H., and Jankun-Kelly, T. Fluid interaction for information visualization. *Information Visualization 10*, 4 (2011), 327–340.
17. Falconer, C. J., Slater, M., Rovira, A., King, J. A., Gilbert, P., Antley, A., and Brewin, C. R. Embodying compassion: A virtual reality paradigm for overcoming excessive self-criticism. *PLOS ONE 9*, 11 (11 2014), 1–7.
18. Gaver, W. W. Technology affordances. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM (1991), 79–84.
19. Gracia, A., González, S., Robles, V., Menasalvas, E., and 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* (2014), 1473871614556393.
20. Gratzl, S., Gehlenborg, N., Lex, A., Pfister, H., and Streit, M. Domino: Extracting, comparing, and manipulating subsets across multiple tabular datasets. *IEEE transactions on visualization and computer graphics 20*, 12 (2014), 2023–2032.
21. Greffard, N., Picarougne, F., and Kuntz, P. Visual community detection: An evaluation of 2d, 3d perspective and 3d stereoscopic displays. In *Graph Drawing*, Springer (2011), 215–225.
22. Helbig, C., Bauer, H.-S., Rink, K., Wulfmeyer, V., Frank, M., and Kolditz, O. Concept and workflow for 3d visualization of atmospheric data in a virtual reality environment for analytical approaches. *Environmental Earth Sciences 72*, 10 (2014), 3767–3780.
23. Hohmann, B., Krispel, U., Havemann, S., and Fellner, D. Cityfit-high-quality urban reconstructions by fitting shape grammars to images and derived textured point clouds. In *Proceedings of the 3rd ISPRS International Workshop 3D-ARCH*, vol. 2009 (2009), 3D.
24. Hsieh, T.-J., Chang, Y.-L., and Huang, B. Visual Analytics of Terrestrial Lidar Data for Cliff Erosion Assessment on Large Displays. In *Proceedings SPIE Satellite Data Compression, Communications, and Processing VII*, vol. 8157, SPIE (2011), 81570D.1–17.
25. Inselberg, A. Multidimensional detective. In *Information Visualization, 1997. Proceedings., IEEE Symposium on*, IEEE (1997), 100–107.

26. Jacob, R. J., Girouard, A., Hirshfield, L. M., Horn, M. S., Shaer, O., Solovey, E. T., and Zigelbaum, J. Reality-based interaction: a framework for post-wimp interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM (2008), 201–210.

27. Jansen, Y., and Dragicevic, P. An interaction model for visualizations beyond the desktop. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2396–2405.

28. Jansen, Y., Dragicevic, P., and Fekete, J.-D. Tangible remote controllers for wall-size displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2012), 2865–2874.

29. Jansen, Y., Dragicevic, P., and Fekete, J.-D. Evaluating the efficiency of physical visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2013), 2593–2602.

30. Kageyama, A., Tamura, Y., and Sato, T. Visualization of Vector Field by Virtual Reality. *Progress of Theoretical Physics Supplement 138* (2000), 665–673.

31. Kilten, K., Groten, R., and Slater, M. The sense of embodiment in virtual reality. *Presence: Teleoperators & Virtual Environments 21*, 4 (2012), 373 – 387.

32. Kosara, R. Indirect multi-touch interaction for brushing in parallel coordinates. In *IS&T/SPIE Electronic Imaging*, International Society for Optics and Photonics (2011), 786809–786809.

33. Kurillo, G., and Forte, M. Telearch – Integrated visual simulation environment for collaborative virtual archaeology. *Mediterranean Archaeology and Archaeometry 12*, 1 (2012), 11–20.

34. Lee, B., Isenberg, P., Riche, N. H., and Carpendale, S. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2689–2698.

35. Lee, J. M., MacLachlan, J., and Wallace, W. A. The effects of 3d imagery on managerial data interpretation. *MIS Quarterly* (1986), 257–269.

36. Loorak, M. H., Perin, C., Collins, C., and Carpendale, S. Exploring the possibilities of embedding heterogeneous data attributes in familiar visualizations. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2017), 581–590.

37. Mackinlay, J. Automating the design of graphical presentations of relational information. *Acm Transactions On Graphics (Tog) 5*, 2 (1986), 110–141.

38. Marriott, K., Meyer, B., and Wittenburg, K. B. A survey of visual language specification and recognition. In *Visual language theory*. Springer, 1998, 5–85.

39. Mine, M. R., Brooks Jr, F. P., and Sequin, C. H. Moving objects in space: exploiting proprioception in virtual-environment interaction. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co. (1997), 19–26.

40. Müller, P., Wonka, P., Haegler, S., Ulmer, A., and Van Gool, L. Procedural modeling of buildings. In *ACM Transactions On Graphics (Tog)*, vol. 25, ACM (2006), 614–623.

41. Munzner, T. *Visualization analysis and design*. CRC Press, 2014.

42. Nelson, L., Cook, D., and 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 (1999), 39–52.

43. Nolle, S., and Klinker, G. Augmented reality as a comparison tool in automotive industry. In *Mixed and Augmented Reality, 2006. ISMAR 2006. IEEE/ACM International Symposium on*, IEEE (2006), 249–250.

44. Oestreich, T. W. Magic quadrant for business intelligence and analytics platforms. *Analyst (s) 501* (2016), G00275847.

45. Ribarsky, W., Bolter, J., den Bosch, A. O., and van Teylingen, R. Visualization and analysis using virtual reality. *IEEE Computer Graphics and Applications 14*, 1 (Jan 1994), 10–12.

46. Sadana, R., and Stasko, J. Expanding selection for information visualization systems on tablet devices. In *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces*, ACM (2016), 149–158.

47. Satyanarayan, A., and Heer, J. Lyra: An interactive visualization design environment. In *Computer Graphics Forum*, vol. 33, Wiley Online Library (2014), 351–360.

48. Schaefer, G., Budnik, M., and Krawczyk, B. Immersive browsing in an image sphere. In *Proceedings of the 11th International Conference on Ubiquitous Information Management and Communication*, IMCOM '17, ACM (New York, NY, USA, 2017), 26:1–26:4.

49. Sedlmair, M., Munzner, T., and Tory, M. Empirical guidance on scatterplot and dimension reduction technique choices. *IEEE transactions on visualization and computer graphics 19*, 12 (2013), 2634–2643.

50. Slater, M. A note on presence terminology. *Presence connect 3*, 3 (2003), 1–5.

51. Slater, M., and 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 (1997), 603–616.

52. Smith, N. G., Knabb, K., DeFanti, C., Weber, P., Schulze, J., Prudhomme, A., Kuester, F., Levy, T. E., and DeFanti, T. A. ArtifactVis2: Managing real-time archaeological data in immersive 3D environments. In *Proceedings Digital Heritage International Congress*, vol. 1, IEEE (2013), 363–370.

53. St. John, M., Cowen, M. B., Smallman, H. S., and 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 (2001), 79–98.
54. St. John, M., Smallman, H. S., Bank, T. E., and 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, SAGE Publications (2001), 1409–1413.
55. Stahnke, J., Dörk, M., Müller, B., and Thom, A. Probing projections: Interaction techniques for interpreting arrangements and errors of dimensionality reductions. *IEEE transactions on visualization and computer graphics 22*, 1 (2016), 629–638.
56. Thomas, B., Piekarski, W., and Gunther, B. Using augmented reality to visualise architecture designs in an outdoor environment. *International Journal of Design Computing Special Issue on Design Computing on the Net (DCNet) 1*, 4.2 (1999).
57. Tory, M., Kirkpatrick, A. E., Atkins, M. S., and Möller, T. Visualization task performance with 2d, 3d, and combination displays. *Visualization and Computer Graphics, IEEE Transactions on 12*, 1 (2006), 2–13.
58. Tukey, J. W., and Tukey, P. A. Computer graphics and exploratory data analysis: An introduction. *The Collected Works of John W. Tukey: Graphics: 1965-1985 5* (1988), 419.
59. Ware, C., and Franck, G. Viewing a graph in a virtual reality display is three times as good as a 2d diagram. In *Visual Languages, 1994. Proceedings., IEEE Symposium on* (Oct 1994), 182–183.
60. Ware, C., and 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*, ACM (2005), 51–58.
61. Westerman, S. J., and Cribbin, T. Mapping semantic information in virtual space: dimensions, variance and individual differences. *International Journal of Human-Computer Studies 53*, 5 (2000), 765–787.
62. Wickens, C. D., Merwin, D. H., and Lin, E. L. Implications of graphics enhancements for the visualization of scientific data: Dimensional integrality, stereopsis, motion, and mesh. *Human Factors 36*, 1 (1994), 44–61.
63. Wilkinson, L. *The grammar of graphics*. Springer Science & Business Media, 2006.
64. Zhang, S., Demiralp, C., Keefe, D., DaSilva, M., Laidlaw, D., Greenberg, B., Basser, P., Pierpaoli, C., Chiocca, E., and Deisboeck, T. An immersive virtual environment for DT-MRI volume visualization applications: a case study. In *Proceedings Visualization 2001*, IEEE (2001), 437–440.