3D InfoVis is Here to Stay: Deal with It

Richard Brath*

Oculus Info Inc. and London South Bank University

ABSTRACT

3D information visualization has existed for more than 100 years. 3D offers intrinsic attributes such as an extra dimension for encoding position and length, meshes and surfaces; lighting and separation. Further 3D can aid mental models for configuration of data within a 3D spatial framework. Perceived issues with 3D are solvable and successful, specialized information visualizations can be built.

Keywords: 3D Information Visualization.

Index Terms: [Systems, man, and cybernetics]: User interfaces – Data visualization. [Engineering management]: Product development – Graphical user interfaces.

1 INTRODUCTION

3D vs. 2D has been debated many times and in many scenarios, academic and otherwise. Papers continue to be published showing negative effects of 3D interfaces (e.g. [1]). Personally, the author has been involved in the creation of visualizations since the late 1980's, initially in domains such as architecture, industrial design, animation – all of which have inherent 3D properties – as well as 3D applications in information visualization since 1990.

3D in infovis is already here and is here to stay. 3D has existed for a long time in information visualization. For example, before OpenGL, before 3D graphics in Excel and Lotus 1-2-3, there are numerous print-based examples of 3D infographics and visualizations, such as Brinton's 1919 book *Graphic Methods for Presenting Facts* [2]. Experimentation with 3D interfaces for abstract data has been proliferous – for example, more than 40 different sphere-based information visualizations were identified in 2013 [3]. The debate should be pivoted towards which criteria and which applications can benefit from 3D vs. applications that will be negatively impacted by 3D. This paper itemizes applications areas identified over the last 25 years where 3D infovis interfaces provide benefits beyond 2D.

2 INTRINSIC ATTRIBUTES OF 3D

InfoVis is different than traditional applications that are inherently 3D, such as architecture, industrial design and some scientific visualizations of 3D phenomena. Based on abstract data, the visual representation can be mapped to any information space including 2D and 3D. A 3D space offers some intrinsic benefits that can be leveraged; while having intrinsic challenges that need to be addressed.

2.1 Position and length are effective encodings

Whether reading Bertin, MacKinlay or Cleveland and McGill [4,5,6], one will find consensus on the effectiveness of using position and length as visual encodings for data. Cleveland shows a low degree of error for length-based encodings compared to other encodings such as hue or brightness. While this error rate will increase with 3D perspective views, there is a potential greater dynamic range available for a 3D length encoding vs. a 2D encoding such as hue or area. Consider the pin map in figure 1. Stacks of pins indicate counts: as a 3D display the heights of various columns can be accurately estimated and compared whether large (e.g. Boston vs New York) or small (e.g. Chicago vs Seattle).

Encoding this same data in 2D presents challenges. Encoding the data using hue or brightness would not offer many levels of discrimination (e.g. some authors suggest brightness only has 2-5 levels of discrimination). Alternatively encoding the data on a 2D map using circles with the area set to the value is feasible, but error estimation for circular areas is higher than length estimation of non-aligned lengths [7]. Furthermore, 2D circles also suffer from occlusion in areas with many overlapping circles. Also the anchor point for the pins in 3D is unambiguous, where as the estimating the center of a circle could be more challenging for larger circles or dense areas of the map.

The counter argument is that issues with occlusion, navigation, etc, will further erode the 3D benefit; however, there are datasets where occlusion can be low – long-tailed datasets will have a few outliers (tall bars) and many small bars, with the user benefiting from being able to discriminate between the differences in the small bars and still having a low degree of occlusion. See figure 1.

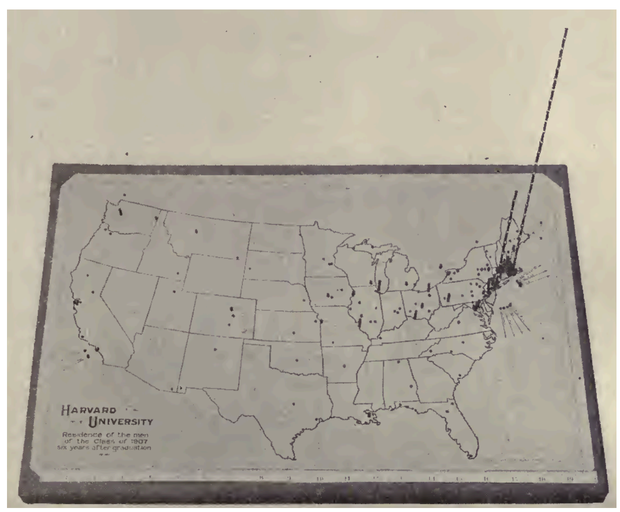

Figure 1: Pins indicate counts on a map. A 3D display easily handles what would be overplotting in 2D (or encoded in a visual attribute such as hue). Height provides accurate estimation of relative sizes both large (e.g. Boston vs. New York) and small (e.g. Seattle vs. Chicago) [2].

2.2 Meshes and Surfaces

Meshes and surfaces are unique 3D representations. These representations can be used to represent a measure using height across two independent variables, e.g. figure 2, as well as other forms, such as a sphere (e.g. commonly used to represent globes

* richard.brath<at>oculusinfo.com; brathr<at>lsbu.ac.uk

and star maps). Examples of real-world applications of 3D surfaces include models of functions (figure 2), financial interest rates (figure 3), financial derivatives (figure 8), periodicity surfaces (e.g. time of day vs. day of week power consumption), bivariate distributions for exploratory data analysis, etc.

Rectangular meshes with cells spaced at regular intervals provide a strong perspective cue when the grid lines are shown explicitly, as shown in both of the surfaces depicted in figure 2.

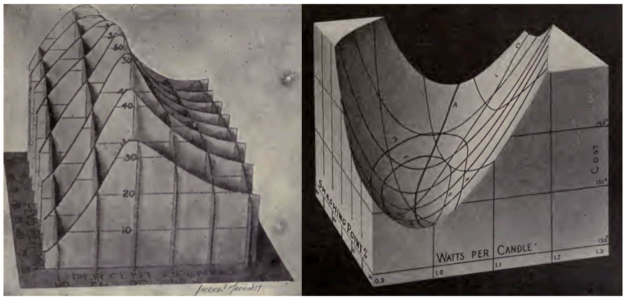

Figure 2: Left: Curves constructed from 2D slices implying a 3D surface. Right: 3D surface with various projection lines. [2].

2.3 Lighting models that highlight anomalies

Intrinsic to many 3D renderers (e.g. OpenGL, Direct3D) are lighting models. While it is easy to create a badly lit scene (e.g. headlight or single directional light), lighting can be used to positive benefit. Shading and/or specular highlights can indicate subtle variations on a surface that may be much more difficult to see using a simple 2D hue and/or brightness based heatmap.

Figure 3 shows a 3D view of interest rates over time. As bars, each surface of the bar has a uniform shade making it difficult to see variations. As a surface, shape is revealed by shading: the macro understanding of the surface as essentially a fairly smooth curve in one dimension with a noisy trend in the other dimension is visible. Subtle local shifts in the day to day movement of curve are visible with highlights and shading resulting from lighting [8]. Both local detail and global form are accommodated by a well lit surface. While a 2D view could show the day to day deviations by calculating a daily difference, the global structure (i.e. the curvature of the broad surface would be lost).

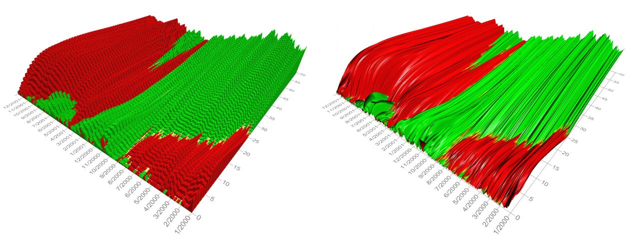

Figure 3: Left: a grid of bars. Right: same data as a surface. Surface shading and highlights reveal more of the local structure than visible in the bar representation.

2.4 Spatial Separation

The extra spatial dimension can be used to create a visual separation between what would otherwise be overlapping items in 2D. Tufte's small multiples [9] utilize 2D spatial separation to indicate different instances indexed by another variable, e.g. in time. Some 2D visualization techniques, such as wormplots or movement on maps, can resemble "hairballs", difficult to decipher to tell what came before and what came after. Space-time cubes, (e.g. [10]) are examples that use the third dimension to make the extra attribute of time more visibly distinguishable than, say, encoding time in 2D using brightness, see figure 4. The 3D space-time cube "GeoTime" is a successful, deployed visualization in use by regional law enforcement officials analysing geotemporal data and securing successful convictions [11].

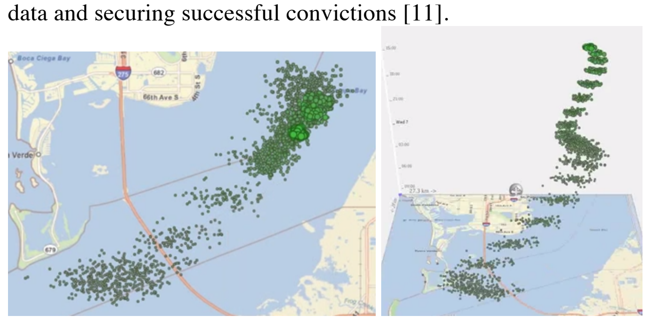

Figure 4: In the 2D view, the spatio-temporal data points overlap, even with brightness it is difficult to perceive the motion. In the 3D view, the vertical axis is used to indicate time, revealing the oscillation in the point cloud over time.

2.5 Perspective is a log transformation

Figure 5 could be dismissed as chart junk with gratuitous perspective, particularly with a strange 3D foreground and 2D background. However, it also hints that perspective may be used as a way of creating a log-like scale – observe the y-axis on the left [12].

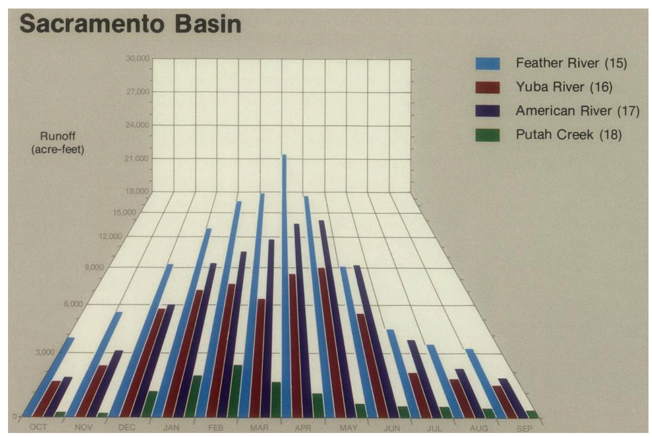

Figure 5: A bar chart tilted in 3D using perspective to provide more area to the smaller numbers in the foreground, forming a log-like chart [12].

Log scales can be useful where data has wide variation (such as the runoff of different rivers in figure 5), a ratio relationship between different levels (e.g. measuring earthquakes, audio volume), or providing focus + context in a single view. Figure 6 shows a timeseries with a focus on recent performance (in the foreground) in the context of the long-term performance (receding in the background). The right side clearly shows the discrete daily movement of the price with more than 30 times the 2D area compared to the start of the timeseries which provides the context where daily movement is not clearly visible but the longer trend and broader vertical range is clearly visible. Note how this effect is accomplished with a single view using perspective thereby potentially improving cognitive performance by eliminating cross-referencing when the focus view and the context view are two different views.

Figure 6: A timeseries chart tilted in 3D to provide more detail to recent performance (right) while providing long term performance (left) as context within a singular view.

3 3D Mental Models and Interaction Techniques

InfoVis datasets are often multi-dimensional. Some of these combinations are effectively translated into 3D forms and 3D mental models. Combined with interaction techniques, these 3D environments may be made effective for data analysis and presentation.

3.1 Cross-tabulations

As shown in the previous section 3D surfaces are effective representations of a dependent variable based on 2 independent variables. The approach also works cognitively with 2 independent categoric variables, using bars instead of surfaces, and adding hue variation to increase differentiation to avoid the problem shown in figure 3 left. 2D bars in a grid can achieve similar results, although potentially with fewer levels of length discrimination, and perceptually, it is easier to compare lengths in one direction (e.g. vertically in figure 7 left) where the bars have a common baseline than in the other direction (where bars float in separate columns). The 3D variant (figure 7 right) offers a common base plane but can suffer from occlusion, particularly in datasets that do not have a long-tail distribution.

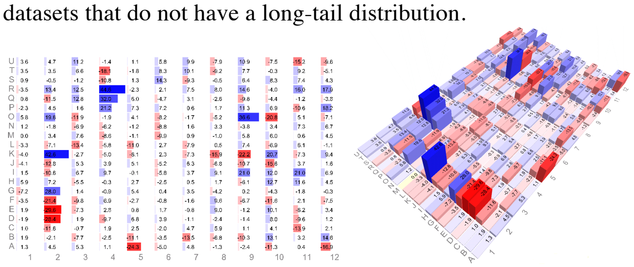

Figure 7: 2D bars (left) facilitate vertical comparison not horizontal comparison. 3D bars facilitate comparison in both directions.

3.2 3D Context + 2D Focus

Deriving from engineering and architectural representations, 3D and 2D linked views can be provided where the 3D view provides the context of the global dataset and associated 2D views represent slices through that data (figure 8). The 2D views can be used to make accurate estimations and comparisons while the 3D view can provide a model to aid understanding the information space. Figure 8 represents a specialized 3D visualization used by financial professionals to assess the valuation of financial derivatives, from an application that is in use by 100,000+ users [13].

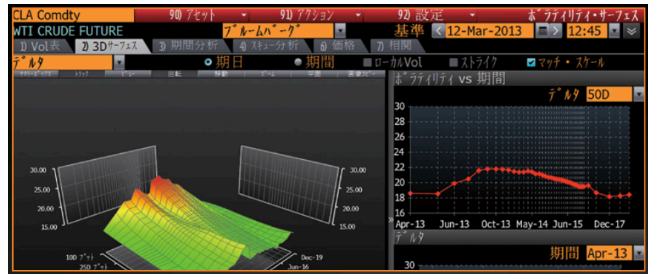

Figure 8: Professional trading application showing valuation of a stock option with a 3D surface (left) and a specific slice through the surface (right) [12].

The approach can also be applied to other 3D representations, e.g. bar, sphere, etc, where the 3D view provides context and also acts as an index to access a specific detailed view (e.g. see [14]).

3.3 Object Constancy

Object constancy makes it easier to follow data through animated transitions that allow data to be visually tracked through transitions between representations. Cognitive burden is lessened by using preattention to motion rather than sequential re-reading and processing of labels [15].

3D transitions can provide object constancy across familiar 2D representations [16]. For example, a multidimensional dataset, such as companies' revenue over time, may be represented as a 2D bar chart – showing revenue for each company; and as a 2D line chart – showing revenue for one or two companies over time. A single view can be used with a 3D transition when moving back and forth between the 2 views. This 3D transition provides object constancy between the two representations (figure 9). A 2D transition could be used, however, Heer and Robertson identify some 2D occlusion issues that are problematic and are partially solved by staggering animations. The 3D transition provides a different means to reduce the occlusion problem and potentially gains the added benefit of exposing a 3D mental model indicating the relationship between these attributes. Using 3D and motion in this way could be considered a specific instance of "structure-from-motion" [17].

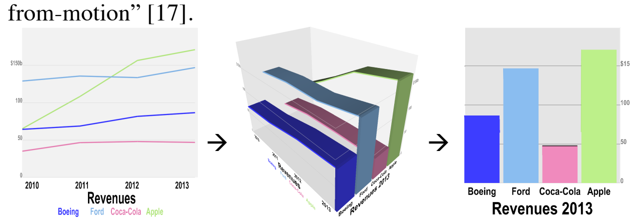

Figure 9: Left: a 2D timeseries chart of revenue. Right: a 2D bar chart of revenue. Middle: a smooth animation from one 2D view to the other can transition through a 3D view providing object constancy and a 3D mental model for these related attributes.

3.4 Mapping Conceptual Models to 3D

Some domain problems may have language or other attributes that convey spatial relationships or may be effectively encoded spatially. Spatial language, such as near, far, up, down, left, right may indicate potential encodings. Other types of relationships may also indicate potential for 3D encoding.

For example, in financial services, correlations are used to inform tasks such as hedging and diversification. A strong correlation (a correlation value approaching 1) is used to find stock prices that tend to move *close together* and can be used to identify alternative stocks with similar price movements. For diversification, non-correlated stocks (a correlation value approaching zero) are of interest to find stocks where price movement is independent of each other or also described as *orthogonal*. In some types of hedging, it is important to find a stock whose performance is *inverse* (correlation approaching -1) to a target stock. These ideas of close, inverse and orthogonal can be used as the basis for a visual layout based on a force-directed layout of stocks based on correlations set out on the surface of a 3D sphere. Inverse correlations will push as far away as possible, i.e. to the opposite side of the sphere; close correlations will be close; and uncorrelated stocks will be in between – i.e. likely orthogonal (figure 10 top [3]). The visual layout in 3D maps well to the mental model of the user through the language that they already use.

The same data and force-directed layout applied to a 2D plane does not result in the same mental model. For an item near the perimeter of the 2D plot and an item on the opposite side the relationship is ambiguous: the correlation could be high but not close due to constraints of the force-directed solution. Interaction can partially address this, but there will still be a perceptual bias for items at the center vs. items at the edge. The notion of opposite becomes more difficult to define in 2D (figure 10 bottom).

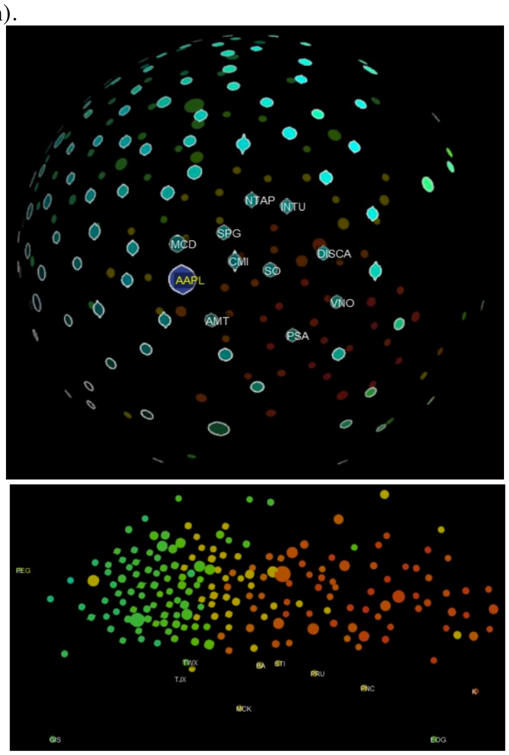

Figure 10: Top - Sphere of correlations shows relationships such as close, inverse and orthogonal for any point in the dataset in 3D. Bottom – same data in 2D. Inverse relationships are no longer intuitive.

3.5 Immersion

In a 2D visualization, the viewer, must be outside of the plane in order to see the visualization. The viewer of a 3D perspective scene is located inside the scene. There are potentially items outside the current viewing frustrum, even behind the viewer. There are items behind other items, partially occluded. While this might be considered a problem, it can also be used as a creative opportunity; and certainly is used in such way in gaming (e.g. clues, fore-shadowing). Such items could be used in 3D visualization possibly when combined with story-telling techniques i.e. hints to the next step in a narrative analytic sequence.

Moving beyond mouse-based navigation opens the potential more intuitive models of navigation. Touch based devices have popularized and familiarized gestures for zoom and pan extensible to 3D although user studies on 3D navigation touch gestures imply additional work may still be required (e.g. [18]).

More immersive interfaces such as immersive display (e.g. virtual reality displays, stereo glasses, parallax barriers, etc.) and input devices (e.g. Kinect, Wii controller, Leap, etc.) can support more intuitive models for 3D navigation and gestures, such as flying or grasping (e.g. [19]).

4 Issues

Beyond the intrinsic properties and mental models, there remain issues with 3D that need to be addressed: Naysayers will continue to refute the use of 3D based on some fundamental challenges that 3D can present.

4.1 Navigation

Navigation is often identified as an issue to be solved with 3D interfaces. Navigation is a critical point of failure when the visual configuration requires navigation to make the scene comprehensible (e.g. a 3D scatterplot), but the particular situation does not allow for interaction, e.g. in published images, in static snapshots in PowerPoint presentations, in collaborative settings where only one person has control over the interaction. As such, reliance on interaction is a liability should any of the above static scenarios occur.

Assuming that interaction is available, 3D navigation can be problematic. In mouse-based environments, the user may desire to do various click-and-drag operations including object selection, scene pan, scene zoom and scene rotation. Creating modes or requiring different contexts to differentiate the interactions makes the system more difficult to use than a 2D scene.

Some 3D game interfaces are extremely easy to operate, e.g. Mario Galaxy has a character running in any direction across small spheres and the navigation of the character and the viewpoint is integrated into a single intuitive control that does not need training nor documentation – young children seem to be able to figure it out quite well. Possibly better or simpler navigation models are required from the 3D visualization community, including both immersive and non-immersive techniques.

4.2 Occlusion

Occlusion may often be claimed as a problem with 3D environments. Occlusion, however, is not unique to 3D. 2D occlusion occurs in scatterplots regularly, with overplotting; or in bubble plots with bubbles of various sizes overlapping.

Furthermore, when a visualization is created for a particular domain, the datasets of interest may result in low occlusion: in some visualizations in some domains, the datasets may typically have normal distributions or long-tail distributions. The behaviour of the data is such that large items and expected patterns are visible and anomalies are still apparent.

4.3 Selection and Manipulation

Navigation and occlusion can impact selection and manipulation. The common 2D click and drag interaction to create a selection bounding box may not be available if drag is used for navigation or may be modal requiring more cognitive effort by the user for managing modes. Picking of small items occluded by other items may be a problem and can be partially addressed in 2D by depth sorting based on size which is not available to spatially anchored 3D objects. Similarly, manipulation of 3D objects can be more difficult given more degrees of freedom. This can be addressed through novel manipulation interfaces (e.g. [20]) or immersive input controllers, discussed earlier.

4.4 Non-anchored points

3D scatterplots are problematic: with no motion or interaction, the location of completely random points in a 3D scene cannot be determined without additional perspective cues (e.g. stereoscopic viewing, shadows, reference lines).

Once again, this may not be a problem with certain types of data. Consider, for example, viewing timeseries data in 3D. Time increments in only one direction. In many datasets, time may update at regular intervals. Some types of measures cannot change dramatically from one interval to the next (e.g. maximum speed and acceleration). Figure 4 right is a 3D scatterplot that is very readable without additional cues – the structure of the data itself provides the cues for decoding the spatial relationships.

Graphical objects, whether discrete (e.g. bars, points) or continuous (e.g. meshes, surfaces) spaced at regular intervals, for example, along a 2D plane or sphere, can be spatially located (e.g. figure 11: WebGL Globe shows bars on a sphere [21]; Walrus Graph Visualization shows a tree with successive levels of depth on sphere [22]). Viewers can adequately assess the relative position in 3D space, even if an underlying reference plane is not visible (e.g. figure 3, 7), although the timeseries in 3D polar coordinates (figure 11 right, by author) is borderline with insufficient cues in the data or scene to aid decoding.

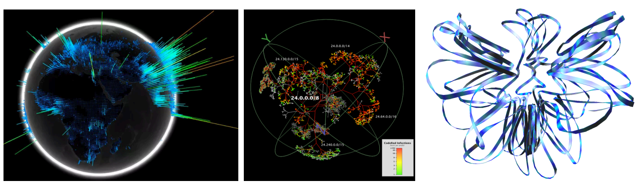

Figure 11:Left: The WebGL Globe showing population [20]. Middle: Walrus visualization showing worm infections [21]. Right: Timeseries in polar coordinates (author).

4.5 Perspective Perception

One complaint about 3D is perspective perception. Perspective is criticized as misleading with numerous examples used to illustrate the potential effect (e.g. see Wikipedia/Misleading_graph or [9]). One example of the problem is shown by considering two objects in a 3D scene by comparing the 2D areas of the objects as presented in the viewing plane, such that an object in the foreground has a larger area than a similar or even larger object in the background. However, to determine that the object in the foreground is smaller or similar, perspective clues are included and these are used to confirm this difference. This indicates that the viewer is quite capable of making this determination as well. The counterargument is that perspective cues aid the decoding of objects as items embedded in a 3D spatial scene, *not* by comparing relative 2D areas, as can be shown by the Ponzo illusion, particularly when the sufficient cues imply a 3D perspective (figure 12).

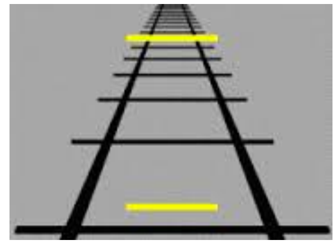

Figure 12: Ponzo illusion. Both yellow bars are the same 2D size, although the yellow bar near top appears much larger given the presence of perspective cues (via en.wikipedia.org/wiki/Ponzo_illusion).

Assuming the visualization has cues to gauge depth, such as linear perspective (e.g. reveal as reference grid), lighting models that reveal shape, stereoscopic viewing, motion parallax, reference planes, and so on, then some degree of visual comparison can be accomplished (see [17] for a discussion on combinations of depth cues). As expressed earlier, the level of accuracy is not as high as a flat 2D projection with objects sharing a common baseline, however, 3D objects encoded by length with appropriate cues may be more effective than other encodings. This would be an excellent future research task to evaluate in a similar fashion to Cleveland and McGill [5].

4.6 3D Fonts are not Readable

Reading text in 3D is sometimes raised as an issue. This may be a historic problem related to low-resolution displays. For two decades the standard screen resolution was 72-96 pixels per inch. A 12-point font may be only 10 pixels high. To improve readability, special 2D sub-pixel font rendering technology was created by Microsoft (ClearType) and Apple (Quartz), which improved legibility for horizontal type. This technology does not apply to 3D type. At 72 ppi, 3D fonts sized similar to their 2D counterparts look terrible and are not particularly readable: figure 13 shows magnified 12 point Times Roman font drawn in 2D and 3D (viewed at similar size) on a 72 PPI monitor. The subpixel rendering optimizations of ClearType manages to keep the 2D text clear spacing between letters, balanced letter forms and clear openings. Simple 3D anti-aliasing results in many muddy letters with poor quality on serifs and seemingly unequal weights on vertical stems.

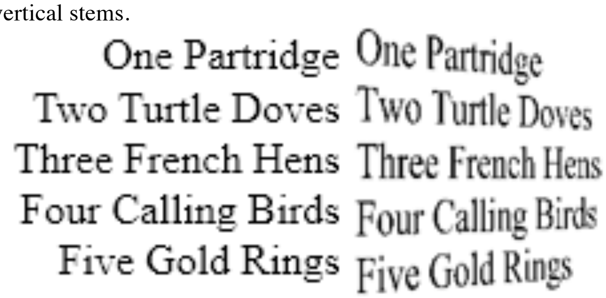

Figure 13: Magnified 12 point Times Roman font drawn in 2D and 3D (viewed at similar size) on a 72 PPI monitor.

3D type is common in the real-world, for example, street signs, store signs, etc. Walking in Paris, it is impossible to see an overhead store sign directly in aligned 2D. Yet we can read these signs readily.

New displays are now available in much higher resolutions. With 250-350 pixels per inch, fonts can be rendered in much greater detail. Fine details on serifs and even subtle changes in weights along a stem are (almost) visible now. Typography in visualization, both 2D and 3D, will have more opportunities to express information through attributes such as bold, italics and small caps; as has long existed on printed maps that support fonts down as small as 3 and 4 point [23].

With higher resolutions, innovative 3D typographic visualizations may become more possible, e.g. Muriel Cooper's pioneering 3D text-based visualizations [24] or Ben Fry's 3D text web browser "tendril" [25] (figure 14).

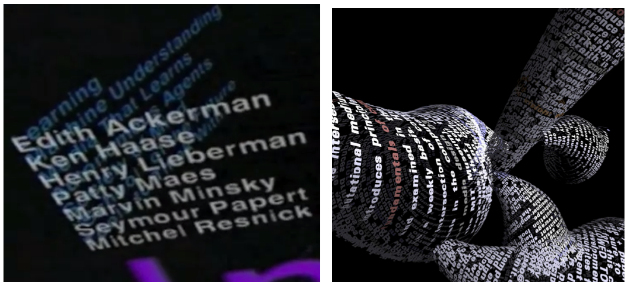

Figure 14: Left: Cooper's Visible Language Workshop visualization snapshot. Right: Fry's tendil: a 3D web browser rendering text applied to 3D forms.

Closely associated with 3D fonts is the issue of labelling 3D scenes. Creating clear labels that work well with whitespace in 2D is a challenging problem and in 3D this challenge is compounded by changing viewpoints which can impact text readability in many ways including loss of contrast between text and its immediate background, occlusion, and orientation (e.g. upside down or perpendicular to the viewer). One highly novel approach to the 3D labelling problem is to view the text on separate coordinated devices [26].

4.7 Bad 3D implementation

Just as it is feasible to create a bad 2D visual encoding, it is perhaps even easier to create a bad 3D implementation, for example, failing with the encoding, failing with the interface interaction paradigms, poorly lit scenes, or use of fonts difficult to read. Extra care must be taken with 3D to ensure that the overall 3D interface provides a benefit.

5 Other Arguments for 3D

Discussions for and against 3D may use other appeals, more difficult to quantify. These include:

- *The real world is 3D therefore 3D works*. This shortcut ignores all the issues discussed above, but also suggests that some stated concerns about 3D visualizations (e.g. fonts are not readable in 3D) should be solvable.
- *Visual literacy of 2D vs. 3D*. One interesting argument for 3D is the potential to extend familiar 2D visualizations to 3D, such as bar charts and line charts into 3D bar and 3D surface charts. This is an intriguing argument – sometimes the acceptance of visualizations can fail if the users have difficulty understanding the encoding. A 3D extension of a known 2D view to a particular community may have greater chance of success than a novel unfamiliar 2D visualization.
- *Visceral appeal*. One argument used for 3D is the immediate emotional attraction to 3D. All the effects available in 3D API's such as glows, transparency, motion, etc can further enhance this. The 3D data visualizations of some artists leverage this attraction to create visually compelling artworks such as Marius Watz (e.g. mariuswatz.com/2008/07/08/trajectories/) and Alex Dragulescu (e.g. sq.ro/malwarez.htm), figure 15.

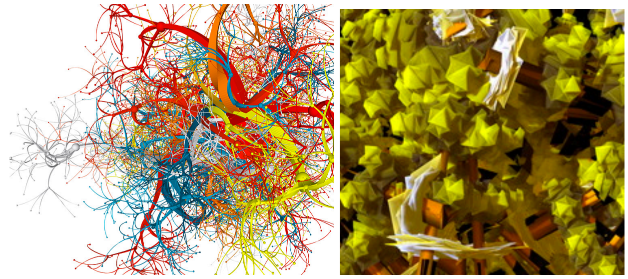

Figure 15:Left: Marius Watz – branching network visualization. Right: Alex Dragulescu – visualization based on code execution of malware.

6 Conclusion

3D information visualization already exists and has successful application in specialized applications such as law enforcement use of space-time cubes (figure 4), or financial professionals visualization of derivative securities as surfaces (figure 8). Unlike solving generic 3D information visualization where data cannot be characterized a priori, when 3D is used for specialized domains, the 3D layout can leverage the structure of the data to reduce or remove the impact of some of the typically limiting factors of 3D such as occlusion or deciphering spatial position. When combined with effective use of other 3D elements (e.g. effective navigation, clearly readable fonts, clear lighting) these 3D information visualizations can be successful.

While there has been a wide variety of 3D evaluation already done (e.g. [16,17,18,19]) there is still more to do. Future work should include more effort to evaluate the effectiveness of encodings in 3D similar to McGill's methodology; evaluate the effectiveness of techniques that use 3D to achieve a spatial separation over 2D counterparts; and evaluate possible cognitive efficiencies of some 3D representations. Successful visualizations may currently rely on regular structures such as reference planes or visualization attributes based on regular intervals – much more dynamic free form 3D information visualizations, such as by artists, provide an interesting starting point for investigating potential novel 3D information visualizations.

References

[1] M. Kyritsis, S. R. Gulliver, S. Morar and R. Stevens. "Issues and Benefits of Using 3D Interfaces: Visual and Verbal Tasks." In *MEDES 13*, ACM. 2013.

[2] W. Brinton. *Graphic Methods for Presenting Facts*. The Engineering Magazine Company, New York, 1919.

[3] R. Brath and P. MacMurchy. "Information Visualization on Spheres." In *Information Visualization: Techniques, Usability and Evaluation*. E. Banissi, F. Marchese, C. Forsell, J. Johannson, ed., Cambridge Scholars Publishing, 2014.

[4] J. Bertin. *Semiology of Graphics*, University of Wisconsin Press, 1983.

[5] W. Cleveland, R. McGill. "Graphical Perception: Theory, experimentation and application to the development of graphical methods" in *Journal American Statistician Association*, 79, 387: 531-554, Sept. 1984.

[6] J. Mackinlay. "Automating the Design of Graphical Presentations", in *ACM Transactions on Graphics*, Vol. 5, No.2, April 1986.

[7] J. Heer, M. Bostock. Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design. *ACM Human Factors in Computing Systems (CHI)*, 203–212, 2010.

[8] R. Brath, M. Peters and R. Senior. Visualization for Communication: The Importance of Aesthetic Sizzle. *Information Visualisation, 2005. Proceedings. Ninth International Conference on*. IEEE, 2005.

[9] E. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT. 1983.

[10] T. Kapler and W. Wright, *GeoTime Information Visualization*, IEEE InfoVis 2004.

[11] K. McDonough. GeoTime Visualization Software. *Law and Order Magazine*. November 2012. URL: lawandordermag.epubxp.com/i/93506/57

[12] W. Kahrl, W. Bowen, S. Brand, et al. *The California Water Atlas*. State of California, 1979. URL: http://www.archive.org/stream/The_California_Water_Atlas/5788000#page/n3/mode/2up Retrieved: 2014/09/28

[13] Bloomberg Commodities Brochure. URL: http://about.bloomberg.co.jp/content/uploads/sites/5/2014/02/web_43536997_JP-CORE-Oil-Brochure.pdf. Retrieved 2014/09/20

[14] S. Feiner and C. Beshers. Worlds within worlds: Metaphors for exploring n-dimensional virtual worlds. *Proceedings of the 3rd annual ACM SIGGRAPH symposium on User interface software and technology*. ACM, 1990.

[15] J. Heer, G. Robertson. Animated Transitions in Statistical Data Graphics. *IEEE Information Visualization (InfoVis)*. 2007.

[16] M. Cordeil, C. Hurter and S. Conversy. "Experimenting and Improving Perception of 3D Rotation-based Transitions between 2D Visualizations" In *Human-Computer Interaction–INTERACT 2011*. Springer Berlin Heidelberg, 2011. 531-534.

[17] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann, Waltham, MA. 2013.

[18] Lingyun Yu, et al. "FI3D: Direct-touch interaction for the exploration of 3D scientific visualization spaces." *Visualization and Computer Graphics, IEEE Transactions on* 16.6 (2010): 1613-1622.

[19] J. McIntire, P. Havig, and E. Geiselman. "What is 3D good for? A review of human performance on stereoscopic 3D displays." *SPIE Defense, Security, and Sensing. International Society for Optics and Photonics*, 2012.

[20] T.T.H. Nguyen and T. Duval. "3-Point++: a new Technique for 3D Manipulation of Virtual Objects." Anatole L´ecuyer and Frank Steinicke and Mark Billinghurst. *3DUI 2013*, IEEE, Mar 2013.

[21] The WebGL Globe. URL: www.chromeexperiments.com/globe. Retreived 2014/09/20.

[22] Walrus – Graph Visualization Tool. URL: www.caida.org/tools/visualization/walrus/ Retrieved 2014/09/20.

[23] R. Brath and E. Banissi. Font Attributes in Knowledge Mapping and Information Retrieval. *Workshop on Knowledge Maps and Information Retrieval (KMIR)*. 2014.

[24] M. Cooper. Visible Languages Workshop. URL: https://www.youtube.com/watch?v=Qn9zCrIJzLs Retrieved: 2014/09/20.

[25] B. Fry. Tendril. URL: benfry.com/tendril/index.html Retreived: 2014/09/20.

[26] C. Roberts, B. Alper, J. Kuchera Morin and T. Hollerer. Augmented Textual Data Viewing in 3D Visualizations Using Tablets. *3D User Interfaces (3DUI), 2012 IEEE Symposium on*. IEEE, 2012.