The Perspective Tunnel:
An Inside View on Smoothly Integrating Detail and Context

Kenneth Mitchell *Jessie Kennedy*

Napier University
Canal Court, 42 Craiglockhart Avenue, Edinburgh, EH14 1LT
Scotland

Abstract. The perspective tunnel, a general kind of information visualisation artefact, embodies a visual form which exploits natural human visual perception. Perspective tunnels map information on to the floor, ceiling and walls of a tunnel, so that both every item of data is visible and the visual field is used effectively. This maximises the effective use of screen real estate, with the potential of representing an infinite amount of data. In real terms, the visualisation of large amounts of information is achieved. A justification of the perspective tunnel's appropriateness is sought from the study of human visual perception. In addition, a selection of perspective tunnel visualisation techniques are presented.

Introduction

Comprehending large quantities of data is crucial in order to make informed decisions from the mass of information resulting from technology advances. To a large extent, visualisations can aid the analysis and understanding of such information sources [27]. This work investigates the effectiveness of 3D visualisations in providing comprehensive, comprehensible representations of information.

Many problems exist in visualisation including, occlusion of overlapping items, ambiguity of symbols in representation, effective layout of structured information, and physical display limitations. These problems are potentially exacerbated in three dimensional (3D) visualisation, where the distribution of objects, geometry of shapes and structures, and required display technology increase in complexity. However, with care, it is precisely this capacity that 3D visualisations can exploit to improve upon 2D visualisation counterparts.

The fundamental difference between 2D and 3D visualisations is that of perspective projection. The techniques outlined here use perspective projections in representations which are familiar, efficient, effective, and contextual:

Familiar As we are used to seeing objects and surfaces in 3D, the perspective tunnel suggests a more familiar form than the use of complex geometrical mappings,

Efficient Where the entire visual display is filled with information,

Effective Through the use of perspective, the entire dataset is visible, and

Contextual Both details and overview are smoothly integrated in a focus+context manner.

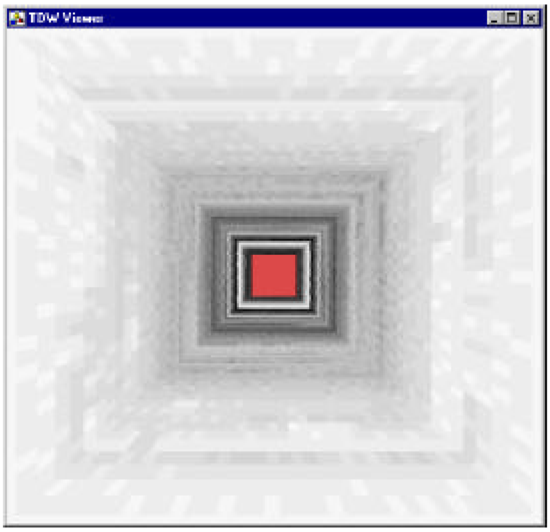

Figure 1: A perspective tunnel visualisation of daily aggregate share prices on the Wall Street stock exchange from 1907 (far) to 1992 (near).

1. Perspective

The celebrated Italian scientist and artist, Leonardo da Vinci (1452-1519) once remarked, *'perspective is the rein and rudder of painting'*. The concept of perspective, which makes distant objects appear smaller than closer ones, was first developed in the 15th century by two Florence architects, Brunelleschi (1377-1446) and Alberti (1404-72). Alberti's *Della Pittura* [1] (meaning, on painting) identifies a method for systematically constructing paintings in perspective. With an art historian's eye, Loretta Staples [32] describes perspective as a visual device,

'One in which near objects are differentiated visually from those afar, and in which those objects' relativity to one another is emphasised'

In the context of calculating 3D perspective projections for computer graphics, Foley *et al* [10] state,

'In 3D, the parallel lines meet only at infinity, so the vanishing point can be thought of as the projection of a point at infinity.'

It is implied here that the concept of a point at infinity is visually perceptible. At the centre of figures 1 and 2 one can imagine such points simply through the suggestion of the converging patterns in perspective. As one's focus moves towards the centre of these images progressively more information is represented. Details are gradually merged revealing patterns at multiple granularities simultaneously.

2. Visual Perception

Human visual perception and cognition can be used to detect patterns in data suggested through a graphical presentation. Indeed the symbol, ¥, is a representation of an infinitely complex entity. In a similar way that an icon can suggest an arbitrary complex entity in a desktop environment, a single pixel can be used to represent multiple samples from an image [28]. The difference is that the colour or luminance of a pixel is determined through a mathematical calculation, whereas, icons are determined by what is deemed to be an appropriate symbol for their referent entity typically covering 16x16 pixels each. When placed in an appropriate visual context (perhaps using a legend), a pixel may represent the same concept as an icon, but at the cost of additional difficulty in its interpretation.

In the retina of the human eye there exists a mass over 100 million of receptor cells which respond to varying intensities of light. In the visual cortex at the rear of the brain there is evidence to suggest that the inputs from these cells are processed by neurons in the detection of bars and edges to build up a raw primal sketch [19]. A possible justification for this theory is that humans require these low level processes to be able to navigate and interact adeptly with distinguishable objects in the world. This is also a possible justification for the human ability to read graphs, because of the low level disposition towards edge detection.

In human vision, the mass of cells on the retina are classified as rod and cones, with the vast majority of being concerned with monochromatic light intensity (rods) and about 5% centred on the fovea respond to one of three particular wavelengths of light, i.e. red, green and blue, for colour vision (cones). All these cells are excited through changes in light intensity, firing more rapidly as light increases. In combination they can be used to enable the visual perception of texture.

As these cells are located on the retina, where our image of the world enters the brain, this perception seems to happen at a very early stage in human visual processing. This suggests the primal nature of dynamic query interfaces [30] which rapidly change the display of scatter plots upon user manipulation.

Throughout our life, we tune our visual perception abilities to seeing the world in 3D. There is evidence to suggest that neurons in the visual cortex respond to light from both eyes, possibly directly involving the perception of depth through stereo disparity.

Although 3D depth perception is not our primary vision ability [8], dealing with 3D surfaces at any orientation is an intrinsic ability of sighted people. Marr’s 2½D sketch [19] is a modern day analogue of recognising 3D objects through edge detection and surface orientation. Visualisations which exploit our primal abilities of visual perception through multi-grained textures and surfaces may yield effective representations of information which are tuned to the limits of our optical system.

3. The Perspective Tunnel

Many methods exist for the visualisation of information providing integrations of focal and contextual views, such as zoomable displays [3] and lenses [11, 16]. In retrospect, such visualisations can be seen to exhibit the perceived benefits of Shneiderman’s visual information seeking mantra, ‘overview, zoom and filter and details on demand’ [29].

The Perspective Wall visualisation [17] makes use of this principle, where a central panel in the foreground is sided by two panels extending into the distance in false perspective. The images on these two panels provide the context for the central panel which provides the detailed section.

If a large number of objects are randomly distributed in 3D, then occlusion is inevitable. The Perspective Wall avoids this by only using the surface of the wall which is entirely visible in a false perspective projection onto the view plane. However, one possible criticism of this is that much space surrounding the wall is not utilised.

The Perspective Tunnel attempts to maximise the effective use of screen space in 3D perspective projections (figures 1 to 6). The image is formed by looking along the major axis of a prism whose cross section is square. The natural perspective projection permits all sides to be visible at once, extending into the distance. When combined with a texture mapping technique [24], a large amount of information may be represented. Indeed, an infinite number of data values may be perceived if the prism extends to a point at infinity.

Of course, the largest amount of information that can be shown on a screen is limited by its resolution, where the colour of each pixel on the screen is used to represent one data value. However, this physical limit may be broken by exploiting human visual perception abilities [19]. More precisely, through the three dimensional (3D) perspective projection of texture maps [10] extending towards the visible horizon, the visualisation of potentially infinite quantities of data is suggested.

3.1. The Time Tunnel: A Helical Layout Technique

In experimenting with texture mapping techniques an obvious step would be to apply existing space filling paths of pixel oriented visualisation techniques [14]. However, its is suggested that the clarity of visualisations may be enhanced through simplified uniform paths. Figure 1 shows a visualisation of stock market data arranged around the perspective tunnel in a helix. This is achieved by a simple using a raster scan texture map image with a period of 128 texels (texture elements). Each texel corresponds to a single daily value, aggregated over the entire Wall Street stock exchange. This gives an indication of the historical trading trends of the overall market. Each revolution of the helix corresponds to about half a year of data beginning in 1907 up to 1992. The farthest light coloured square with a dark square immediately following this shows the Wall Street crash of 1929, whose effects can be seen for a number of years after. In the foreground, individual values can be seen with minor fluctuations occurring on a day to day basis.

The Spiral Calendar [18] technique also maps temporal information in a visual spiral. However, where the Spiral Calendar shows successive levels of time periods in 2D panels facing the user positioned in a spiral, the helical technique shows time as the smooth progression along the floor, ceiling and walls of the perspective tunnel.

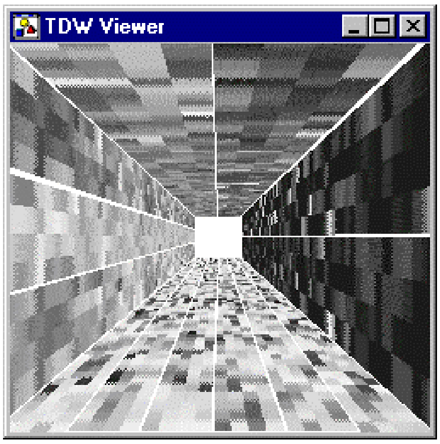

Figure 2: A visualisation of 4 years of bi-weekly lottery results

3.2. The Lotto Road: A Parallel Channels Technique

A related technique to the helical layout technique is achieved through aligning multiple data dimensions in parallel channels (figure 2). Figure 2 shows the history of a lottery competition. The bottom six channels show the six actual ball numbers. Reassuringly, the effect is one of random noise signifying the non-determinism of lottery balls. On the right wall of the tunnel the number of 1st prize winners is shown with the associated prize money. The roof of the tunnel shows the 2nd prize winners and their prize amounts. The left wall show 3rd prize winners along with the total prize money at the lower left. Although the balls seem random, there is an interesting recurrence of high payouts roughly every 8 weeks. Perhaps indicating additional incentives induced by the lottery organisation or some social factors.

In a way similar to the circle segments pixel oriented technique [14], correlations between many data dimensions can be viewed simultaneously. However, as the circle segments technique is rigorously pixel oriented the number of pixels at each concentric circle of data dimensions decreases towards the centre. Whereas the parallel channels technique maintains a uniform number of data values all along the tunnel. This makes the task of visually quantifying areas of the data and identifying periods easier, because of the perceptual uniformity of data distribution.

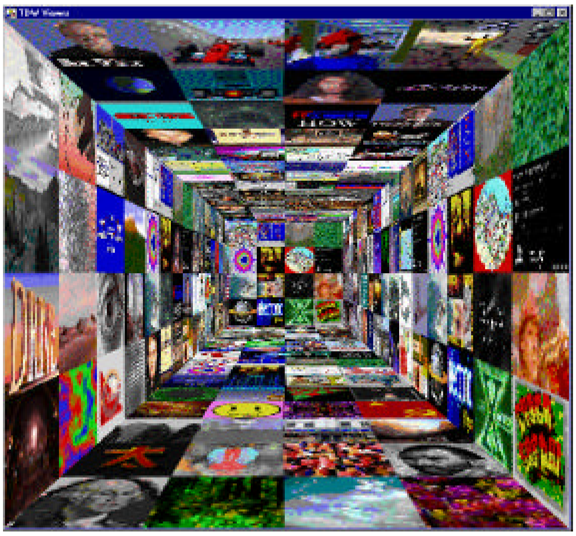

Figure 3: The video tunnel prototype simulating digital video-on-demand channels

3.3. The Video Tunnel: A Surface Subdivision Technique

As the number of available television channels increases, there is an increasing problem in obtaining an informed decision of programmes which are of interest. In analogy with the video grid channel commonly provided by cable service providers, which consists of a 2D 4x4 array of channels, the video tunnel sub-divides the space on each surface of the perspective tunnel into a set of video grids. The end of the tunnel is sealed off with another video grid to make complete use of the entire image, resulting in almost 300 channels being visible.

The ubiquitous use of texture mapped images in the perspective tunnel yields a large range of possible applications. The quad-tree sub-division of surfaces makes full use of the space in a uniform manner. General surface sub-division schemes can be applied in the perspective tunnel information visualisation artefact, e.g. Tree-Maps [13] and CocktailMaps [2]. In summary, the advantages of smoothly integrating focus and context in the perspective tunnel can be obtained for any existing 2D visualisation through the mapping of images to the surfaces of the perspective tunnel.

3.4. The Warren: A Recursive Projection Technique

The recursive projection technique (figure 4) takes the sub-division of panels a step further. As the perspective tunnel is a projection of a 3D space onto a 2D plane, each panel on the sides of the perspective tunnel can itself show a perspective tunnel image. The depths of each tunnel can be seen, because the entire contents of each tunnel are projected onto a flat plane. In general, a mapping of directed acyclic graphs to the perspective tunnel is possible with the recursive projection technique.

General graph structures may be obtained in a manner based on that used in the hyperbolic web browser [16] or simply by providing two views corresponding to branches in and out of the current node. For such structures Bertin suggests [4] that nodes take up progressively less space according to the length of their graph path distance from the root node. Thus, catering for an exponential growth in the cumulative fan-out from graph nodes, by exponentially reducing their allocated screen space through natural perspective.

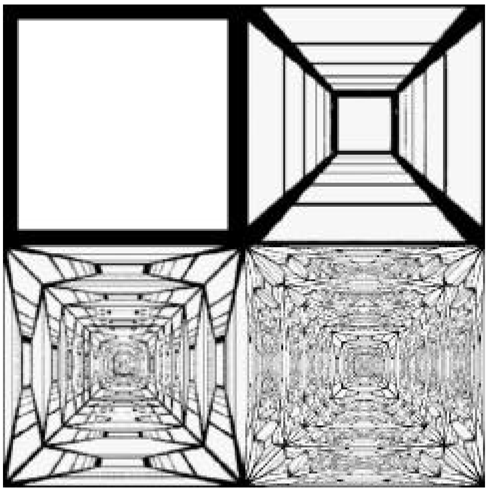

Figure 4: Recursive perspective projection technique for visualising graph structures shown at successive levels of recursion.

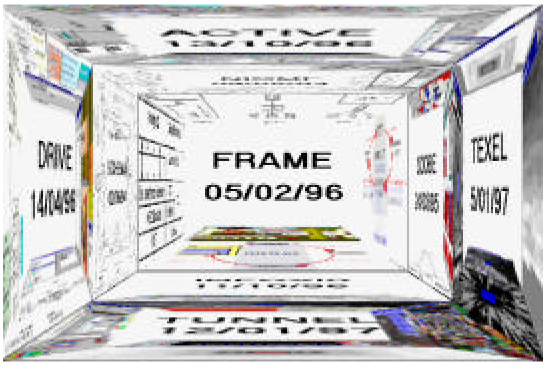

Figure 5: Recursive projection of images from previous papers.

Figure 5 shows the recursive projection technique applied to the visualisations presented in the author's previous papers in a tree structure, with a hierarchy of root, papers and images. For each node of the tree, the far panel is used to represent the data of that node in the tree (in this case a name with date or an image). The other walls show the children of the node. If there are no children, then the data of that node fills the projection.

Clearly, hierarchical structures of arbitrary depth can be suggested in a self contained area, without the necessity for collapsing or expanding branches of the tree. This technique avoids the occlusion problems of Cone Trees [26] and makes more complete use of space than the hyperbolic lens [16].

Recursively projecting from a 3D perspective tunnel onto a 2D surface (in effect, seeing round corners) provides the potential for visualising complex spatial structures. In relation to the n-Vision system [9], the visualisation of *n*-dimensional virtual worlds is possible. However, where n-Vision's worlds within worlds metaphor encourages the user to think of nested worlds it can only show a few nested 3D graphs at any one time, the recursive projection technique may represent such nesting of dimensions simultaneously without occlusion. For example mapping each data dimension to helii in a recursive version of the perspective tunnel permits the visualisation of correlations between data dimensions nested within each subtunnel

Clearly, as the domain of each dimension increases in breadth (corresponding directly to the depth of a helical perspective tunnel), the limitations of the display resolution may degrade the effective number of visible dimensions. In such cases, texel oriented techniques (figures 1 and 2) may yield effective results.

An analysis of the information content versus display resolution of texel oriented techniques is given in [24].

In terms of human visual perception, it may be argued that the use of recursive projections potentially yields an unfamiliar representation of space. However, in contrast with distortion based techniques, which visually perturb the image, the technique presented here is more a topographical distortion where the natural aspects of perspective projection is ultimately retained in the final projected image.

3.5. Win3D: A Prototype 3D Operating Environment

The perspective tunnel is an ideal candidate for managing space contention on today's desktop screen displays. The occlusion problems associated with multiple overlapping windows is one which should be addressed by information visualisation practitioners. The Rooms [12] and 3D/Rooms [27] metaphor effectively enlarges the screen space by permitting fast access to a number of virtual workspaces. However, the overall view of the workspace is presented as a 2D grid in a separate view, thus losing the visual context of the virtual workspace. The Elastic Windows system [31] uses the TreeMaps [13] technique for organising hierarchical arrangements of windows, where non-focal windows are shrunk around windows of more importance to the user. Varying the size of such windows, although visually pleasing, detracts from interacting with the window's content. Therefore, the automatic management of focus control is desirable.

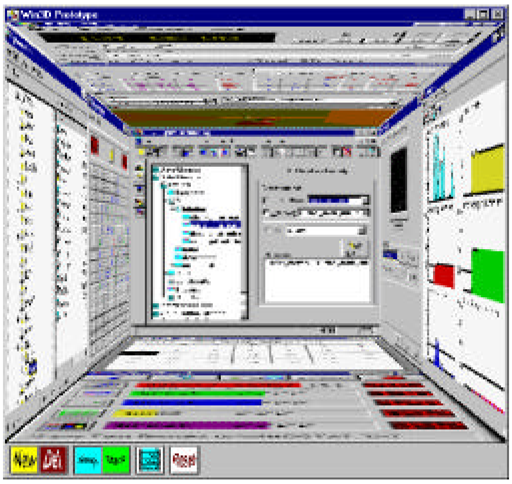

Figure 6: The Win3D prototype for investigating focus control.

Win3D (figure 6) constitutes a brief investigation into these issues. Images captured from various windows applications are laid out in the perspective tunnel according to two schemes, swap and stack. When swapping, window panels are laid out in a helix round the tunnel as they are added to the workspace. Selecting a panel on the tunnel's sides makes it fly to the centre panel. If a panel currently exists in the centre, it is swapped with the newly selected panel. Whilst the stacking technique organises the panels in a helical stack, i.e. the most recently used panel is pushed on the helical stack at the near end of the tunnel. Over time this effectively allocates more screen space to frequently used panels, thereby mirroring and therefore relieving the user's short term memory cache of windows.

Acknowledging the perhaps excessive amount of screen space devoted to contextual information, the centre panel may be maximised to cover the full screen and restored to its original size. However, the consistency of the workspace is retained through smooth animated transitions between these states.

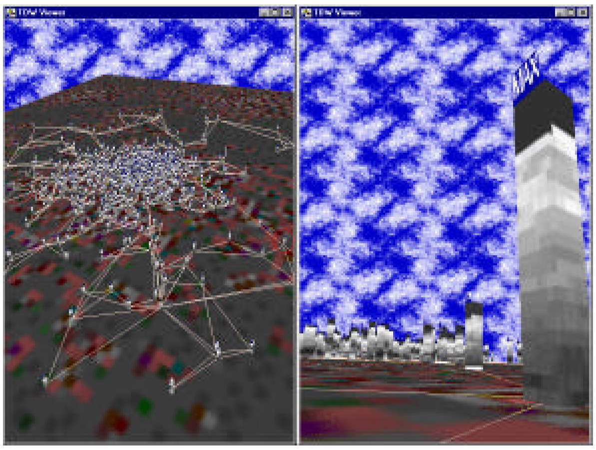

Figure 7: The City 3D database environment.

3.6. The City: A Historical Stock Market 3D Database Environment

A 3D database environment is a visual database interface in which the user may navigate and manipulate data structures in 3D [23]. Figure 7 shows a 3D database environment concerning a megabyte of simulated stock market information. Each share's history is represented as a tower, which uses the helical layout technique of figure 1.

Share towers are positioned relative to each other according to a similarity metric based on the work of Chalmers [6] with a downward force applied to improve the legibility of the layout and reduce the effects of occlusion. A simple triangulation of the resulting tower distribution, based on nearest neighbours, is added to reinforce the visual effect of clustering and cohesion.

As the user navigates the 3D database environment s/he can obtain an overview by flying around the towers of the city. Flying close to towers reveals its share name abbreviation and entering a tower and looking up or down offers a perspective tunnel view. In this way, multi-level scales of focus and context has been achieved in a manner consistent and familiar with real world spatial properties.

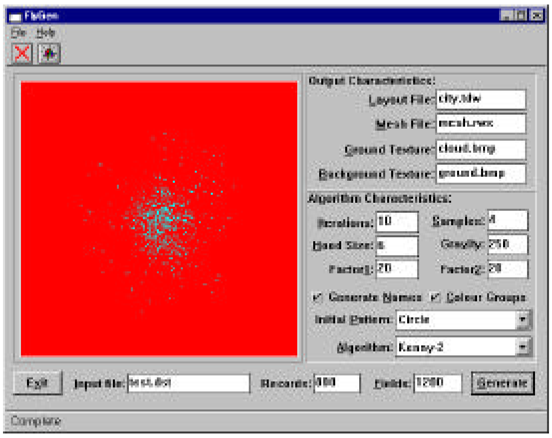

Figure 8: The FlyGen layout generation testbed

With the development of linear iteration time algorithms for computing such complex layouts [6], the potential exists for interactively manipulating layout parameters with quick responses. As the generation of such layouts can be enhanced through user discretion, the ability to view the layout algorithm, as it progresses, is invaluable. The FlyGen test bed, animates this layout process according to user-defined parameters and generates the resulting 3D database environment. Figure 8 shows the configuration and final state of the layout shown in figure 7.

4. Systematic Support for Visualisation with the Perspective Tunnel

In order to experiment with perspective tunnel visualisations, a development framework is required. In previous work [20, 21, 22, 5], a framework for user-interfaces to databases with advanced visualisation support has been developed. This includes the ability to construct user-interfaces to databases which integrate 3D graphical widgets (TDW) [5] and standard desktop widgets. The framework's relevance to information visualisation has been identified in [15].

In providing support for perspective tunnel visualisation in the context of this framework an object oriented 3D perspective tunnel widget has been implemented. This widget simply renders the geometry of the perspective tunnel in a simple perspective projection. A single texture map is wrapped round the four sides of the tunnel to provide the visualisation of texels. This texture map is held as a bitmap image in memory for rapid manipulation of both texels and mapping parameters. A full specification of this widget is detailed in [23].

The use of texture mapping in 3D information visualisation indicates a recent trend in computer graphics. The realisation that texture maps are an efficient way of increasing the complexity of real time 3D graphics is permeating to user-interface developers [25]. This moves the load in graphics algorithms from polygon processing to texture mapping. A few well chosen texture maps can replace thousands of polygons in a given scene, i.e. a texture map is worth a thousand polygons. This results in surprisingly undemanding computational needs, with ~5 frames per second on a DX-2 50MHz PC at a resolution of 512×512 without graphics card acceleration.

5. Conclusion

The perspective tunnel, which maximises the effective use of screen real estate, with the potential of representing an infinite amount of data, has been presented. The visualisation of large amounts of information has been demonstrated and a possible justification of its appropriateness through visual perception has been discussed. In addition, a selection of perspective tunnel visualisation techniques have been presented.

As texture mapping has been of major interest to computer graphics researchers in the quest for visual realism, there now exists many advanced methods. Such results may be utilised in improving the quality of 3D information visualisations. Further, as a 3D database environment, the perspective tunnel can be navigated and viewed from any point. Animated transitions through graph structures represented using the recursive projection technique should provide an interesting sense of both immersion and overview.

6. References

1. L. Alberti. Della Pittura, Florence, Italy, 1435.
2. C. Ahlberg. Spotfire: An Information Exploration Environment, ACM SIGMOD Record Special Issue on Information Visualisation, 24(3), 1996.
3. B. Bederson, J. Hollan, D. Perlin, J. Meyer, D. Bacon & G. Furnas. Pad++: A Zoomable Graphical Sketchpad for Exploring Alternative Interface Physics, Journal of Visual Languages and Computing, 7(1):3-32, 1996.
4. J. Bertin. Semiology of Graphics, University of Wisconsin Press, 1983.
5. J. Boyle & K. Mitchell. Embedding Three Dimensional Graphics Inside a User-Interface Development Framework, Technical Report, Robert-Gordon University, Aberdeen, 1996.
6. M. Chalmers. A Linear Iteration Time Layout Algorithm for Visualising High-Dimensional Data, IEEE Visualisation, 1996.
7. A. Dix, J. Finlay, G. Abowd & R. Beale. Human Computer Interaction, Prentice-Hall, 1993.
8. A. Dix. Time, Space and Interaction, FADIVA Workshop, 3. Gubbi, Italy, 1996.

9. S. Feiner & C. Beshers. Worlds Within Worlds: Metaphors for Exploring N-Dimenional Virtual Worlds, ACM Symposium on User Interface Software & Technology 6:76-83, 1990.
10. J. Foley, A. van Dam, S. Feiner & S. Hughes. Computer Graphics: Principles and Practice, Addison-Wesley.2nd Edition, 1990.
11. G. Furnas. Generalised Fisheye Views, ACM Special Interest Group on Computer Human Interaction, 16-23, 1986.
12. D. Henderson & S. Card. Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention in a Window-Based Graphical User Interface, ACM Transactions on Graphics 5(3):211-245, 1986.
13. B. Johnson & B. Shneiderman. Treemaps: A Space-Filling approach to the Visualisation of Hierarchical Information Structures, IEEE Visualisation 2:284-291, 1991.
14. D. Keim. Pixel-oriented Database Visualisation, ACM SIGMOD Record Special Issue on Information Visualisation 24(5), 1996.
15. J. Kennedy, K. Mitchell & P. Barclay. A Framework for Information Visualisation, ACM SIGMOD Record Special Issue on Information Visualisation 24(4), 1996.
16. J. Lamping & R. Rao. The Hyperbolic Browser: A Focus+Context Technique for Visualising Hierarchies, Journal of Visual Languages and Computing 7(1):33-56, 1996.
17. J. Mackinlay, J. Card & G. Robertson. Perspective Wall: Detail and Context Smoothly Integrated, ACM Special Interest Group on Computer Human Interaction 173-179, 1991.
18. J. Mackinlay, G. Robertson & R. Deline. Developing Calendar Visualisers for the Information Visualiser, ACM Symposium on User Interface Software and Technology 10:109-118, 1994.
19. D. Marr. Vision, Freeman, 1982.
20. K. Mitchell, J. Kennedy & P. Barclay. Using a Conceptual Data Language to Describe a Database and its Interface, British National Conference on Databases 13:101-119, 1995.
21. K. Mitchell, J. Kennedy & P. Barclay. A Framework for User Interfaces to Databases, ACM International Workshop on Advanced Visual Interfaces, 1996.
22. K. Mitchell & J. Kennedy. DRIVE: An Environment for the Organised Construction of User-Interfaces to Databases, International Workshop on Interfaces to Databases, 3, 1996.
23. K. Mitchell. 3D Database Environments, PhD Thesis, Napier University, Edinburgh, 1997.
24. K. Mitchell. To Infinity and Beyond: Texel Oriented Techniques for Visualising Large Databases, to be submitted for publication, 1997.
25. R. Pausch. User Interface Design in the Unreal World: What I Learnt at Disney, ACM Symposium on User-Interface Software & Technology, Banquet Speech, 1996.
26. G. Robertson, J. Mackinlay & S. Card. Cone Trees: Animated 3D Visualisations of Hier-archical Information, ACM Special Interest Group on Computer Human Interaction 189-194, 1991.
27. G. Robertson, S. Card & J. Mackinlay. Information Visualisation using 3D Interactive Animation, Communications of the ACM 36(4):37-71, 1993.
28. A. Ray Smith. A Pixel is NOTa Little Square, A Pixel is NOT a Little Square, A Pixel is NOT a Little Square..., Technical Memo-6, Microsoft, 1995.
29. B. Shneiderman. Designing the User Interface: Strategies for Effective Human-Computer Interaction, Addison-Wesley, 1987.
30. B. Shneiderman. Dynamic Queries for Visual Information Seeking, IEEE Software 11(6):70-77, 1994.
31. B. Shneiderman & E. Kandogan. Elastic Windows: Improved Spatial Layout and Rapid Multiple Window Operations, ACM Workshop on Advanced Visual Interfaces 23-28, 1996.
32. L. Staples. Representation in Virtual Space: Visual Convention in the Graphical User Interface, ACM INTERCHI 348-354, 1993.