3-Dimensional Pliable Surfaces: For the Effective Presentation of Visual Information

M. Sheelagh T. Carpendale, David J. Cowperthwaite, F. David Fracchia
School of Computing Science, Simon Fraser University
Burnaby, B.C. V5A 1S6 CANADA
carpenda@cs.sfu.ca

ABSTRACT

A fundamental issue in user interface design is the effective use of available screen space, commonly referred to as the screen real estate problem. This paper presents a new distortion-based viewing tool for exploring large information spaces through the use of a three-dimensional pliable surface. Arbitrarily-shaped regions (foci) on the surface may be selected and pulled towards or pushed away from the viewer thereby increasing or decreasing the level of detail contained within each region. Furthermore, multiple foci are smoothly blended together such that there is no loss of context. The manipulation and blending of foci is accomplished using a fairly simple mathematical model based on gaussian curves. The significance of this approach is that it utilizes precognitive perceptual cues about the three-dimensional surface to make the distortions comprehensible, and allows the user to interactively control the location, shape, and extent of the distortion in very large graphs or maps.

KEYWORDS: Distortion viewing, screen layout, 3D interactions, information visualization, interface metaphors, interface design issues

INTRODUCTION

Since the advent of video display terminals as the primary interface to the computer, how to best make use of the available screen space has been a fundamental issue in user interface design. In fact, the necessity for effective solutions to this problem has intensified as technology has advanced, with the ability to produce visual data in greater volumes continuing to outstrip the rate at which display technology has developed. This issue is traditionally referred to as the *screen real estate problem*.

A significant advance in the screen real estate problem was the leap from command line access to a windowing environment. Even with the addition of the now familiar features of panning and zooming, the desire to examine detail often conflicts with the ability to maintain global context. Zooming out of or compressing the data to fit within the space of the screen can result in its becoming too dense to discern detail. Zooming into or magnifying a region will result in the loss of context.

Multiple views allow for the simultaneous display of detail and global structure, however the integration of these distinct views must be performed consciously by the user. Evidence as to how we combine information from multiple sensory channels has arisen from a number of studies in experimental psychology [5, 12, 13]. Information perceived as a single event is integrated automatically, however that perceived as distinct events requires a more strenuous reintegration. While the user may be cognitively aware that views in multiple windows pertain to a single information space, perceptually they remain distinct. For example, the effort of maintaining which detail belongs where and of current locations has to be performed consciously by the user. If the desired detail view can be provided in a manner that smoothly integrates it into the global context then it preserves the possibility of visual gestalt.

This work presents the development of a tool for exploring large information spaces which:

- Increases the amount of information that can usefully be presented on a computer screen.
- Creates a situation that can utilize visual gestalt by retaining the perception of the information space as a single event. This avoids creating situations that are known to be cognitively strenuous.
- Keeps the resulting images comprehensible.
- Encodes as much information as possible in a manner that will access low level perceptual abilities.

This paper focuses on two broad classes of information spaces: graphs and maps. By graphs we refer to visual representations where information is expressed as entities (nodes) and relationships between them (edges), and by maps to representations which also contain distance information.

The next section discusses distortion viewing techniques most closely related to this work. The concept behind 3D pliable surfaces is then described. This is follow by an explanation that explains the development of the 3D pliable surface to provide a detail within context viewing tool. Then the visual cues used to provide comprehension of the surface distortion are presented. Finally, the advantages and limitations of this approach are discussed.

RELATED WORK

There have been several different approaches to the screen real estate problem in the literature. Some simplify the task by applying various filtering techniques [3, 6]. Others investigate the advantages of using 3D to display the information [4, 17, 24]. Finally, there are those that distort the viewing space to provide unified views that contain selected detail within context [1, 8, 10, 11, 14, 18, 19, 20]. It is the latter group, in particular those that use perspective to provide distortion [11, 18], that pertain most closely to the approach presented here.

Initially Furnas [6] observed that a fisheye, or very wide angle lens, provides a world view by showing a focus in great detail and gradually decreasing this detail as distance from the focus increases. His studies in various subject areas (geography, workplaces, history, and newspapers) reveal that people naturally retain and present information in this manner. While this work has laid the foundations for most of the distortion viewing techniques, his examples threshold a degree of interest function, creating filtered fisheye views with gaps in context. This could be problematic if, for instance, the next section of interest was in one of these gaps.

Sakar and Brown [19] expand upon Furnas' approach by using a visual representation to express fisheye views of graphs. Given a current focal point, they use a trade off between a node's assigned *a priori importance* and its distance to the focus point to establish the position, size, and amount of detail to display for each node. However, this approach offers only single focal points and is interactive for fairly small graphs (approximately 100 nodes).

Perspective wall [11] and Document lens [18] both use perspective to provide magnification of the focus and the resulting 3D image to provide visual information about the context and how it has been distorted. However, they only provide a single focal point and were designed for particular data. Perspective Wall is realistically limited to linear information and Document Lens has only been applied to displaying text. Conceivably the latter could be used to display other types of 2D data, such as maps or graphs, although much of the perspective information provided by the regular patterns in text would be lost.

Hyperbolic display [10] and CATGRAPH [8] use a simple mathematical function, hyperbola and arctan respectively, for their magnification and distortion. The Hyperbolic display provides an interactive single focus viewing tool, while CATGRAPH allows for multiple foci and uses both rectangular and polar transformations. Both of these functions are asymptotic and as a result spread the distortion across the entire image and cause extreme compression at the edges. Neither of these approaches provide for the possibility of creating a focus that spans an area of the graph nor allow for magnification of such a focus to scale in a manner that maintains distance relationships.

The rubber sheet approach [20] based on morphing [2] provides multiple foci as convex polygons, does not cause areas of unrequested magnification, and maintains contextual frame. However, several problems are mentioned with regard to this method. The transformational technique does not have a general inverse mapping which makes editing across the entire distorted image non-trivial. Also, this tool provides real time response for graphs of up to only a few hundred nodes and a similar number of links, and sometimes more than one iteration is required to provide the right balance between detail and context. In terms of large information spaces this is quite limiting.

Much of the desired functionality exists but not within a single approach. Reasonable response time is available from the mathematically-based approaches, distortion comprehension from the 3D perspective approaches, and a multi-focal smooth integration display from the morphing approach. Our 3D pliable surface *3DPS* effectively combines these advantages.

CONCEPT

In creating our viewing tool for a two-dimensional information space we make a distinction between the graph or map as the image encoding the information and the surface on which it is displayed. The resulting tool will not be tied to any particular kind of image. As the visual cues are provided about the surface, distortions will still be readable even when there are gaps in the image. Current distortions can be quite readable when applied to regularly spaced information, particularly grids or text; unfortunately not all information can be laid out so regularly.

Choice of the distortion transformation is crucial as it will affect both performance and the visual result. Ideally, one would like a simple mathematical function that will provide smooth integration from the focus through the distorted section and into the context; preferably it would have no discontinuities in curvature. We chose the three-dimensional gaussian curve as its bell shape curves away from the focus at its apex and inflects to curve gently back into the surface (Figure 1). These gaussian curves transform the two-dimensional flat surface into a three-dimensional curved surface. The three-dimensional nature of this distortion approach offers several advantages:

1. Using single point perspective to view the three-dimensional surface from above provides detail with magnification to scale and a readily controllable context.

2. It provides a useful metaphor for the actions performed to

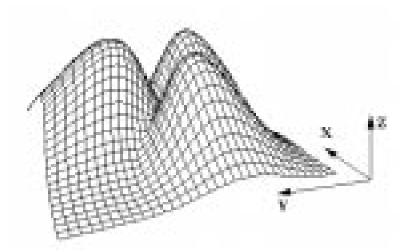

Figure 1: 3-D Surface of Blended Gaussian Curves

create the distortions. Pulling a section towards oneself to see it better, or in this case magnify it, appears to be a natural response.

3. It also provides a metaphor for the overall appearance; the end result of a multi-focal view can be seen as a softly curved 'landscape' with hills and valleys.

4. While there is on-going discussion regarding the amount of information that humans can comprehend when presented three-dimensionally, Ware's studies [23] imply that there is some truth to the notion that because humans evolved in a 3D world they are better at retrieving information about it.

While we feel that being able to understand the resulting distortion is significant for both graphs and maps there is a particular point to be made for the latter. In our culture, interpreting maps assumes that distance is to scale and that scale is normally consistent across the map and is clearly indicated. A distortion view will create an image of the map containing sections of varying scale. User disorientation when viewing a distorted map has been reported [20] particularly when the map was familiar. We suggest that this disorientation results from the discrepancy between the information provided by the distorted map and what the users feel they know to be true about the original map. This effect parallels what Tufte [21] discusses extensively as 'lie factors'. In this case previous knowledge is protecting the user from assimilating false information. In unfamiliar information spaces there is a greater chance of being misled.

To dispel the possibilities of misleading users the form of the 3D pliable surface should be clearly displayed. This will create comprehensible distortion allowing the user to understand at a glance which sections are magnified and which are compressed with an intuitive notion of the extent.

One choice for revealing form is to employ shading. It has been well established that humans can discern three-dimensional shape from shading alone [16, 20], and there is considerable evidence to support the fact that this is a low level precognitive skill [9]. Such a low level visual routine will interfere less with conscious processing and may even provide an aspect of the interface that requires no learning [22].

A different choice is the use of perspective to provide distortion information. However, understanding three-dimensions from perspective appears to be a learned skill and demonstratably culturally tied [7]. Also, perspective has often been indicated with the outlines of a three-dimensional shape [11] or by the visual pattern of the data [18]. The choice of smooth curves for distortion and allowing for data with irregular layouts means neither outlines nor patterns in the data will reveal the nature of the distortions. However, using a regular grid reveals the shape of the distortions by accessing two human depth cues: it provides perspective information without requiring edges and serves as a texture gradient.

3D PLIABLE SURFACE

This section steps through the creation and manipulation of a 3D pliable surface. We start from the simplest case where a single focus is in the center of the field of view and progress through to the interaction of multiple foci.

Single focus at the center of the field of view

Here the action of pulling the region of interest up, perpendicular to the surface, produces the desired magnification. The surrounding region is stretched over a three-dimensional gaussian curve connecting the magnified region to the original plane of the surface (Figure 2). This controls the compression of the surrounding context and the integration of the magnified region back into the original image.

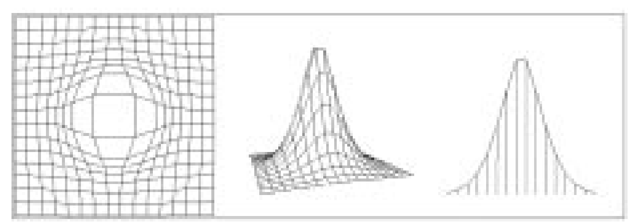

Figure 2: Single focus in the center of the field of view: left; top view, center; 3D view, right; profile with ToEye vectors

Magnification of single focus to scale

The center of the 3D gaussian curve is projected up to the height h_c, but to provide a flat region where only scaling occurs, the curve may be truncated; limited to a fraction f of this maximum height (Figure 3). The points of the graph/map in the central magnified region are all projected up to the same height, $h_c f$. The height h_p of all other points on the curve is a simple relationship of distance d_p to the center of the region, the height h_c and its standard deviation s_c:

$$h_p = h_c \mathrm{exp}^{-\frac{s_c}{d_p}}$$

As in Figure 3 projecting all points perpendicular to the plane provides the desired magnification and compression in the appropriate regions.

Single focus anywhere in the field of view

If, however, the region of interest is located at a point other than that directly below the viewpoint, projection perpendic-

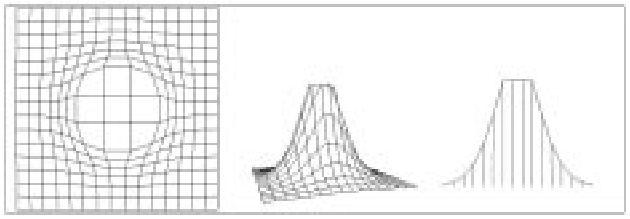

Figure 3: Single focus with flattened top

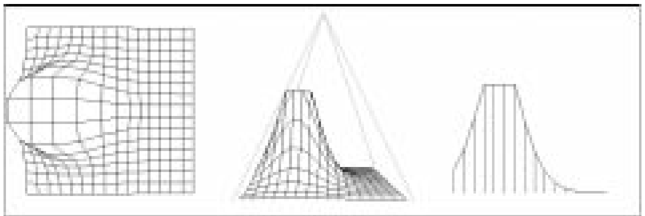

Figure 4: Single off-center focus, projected perpendicular to the plane. Dotted lines denote the viewing frustrum

ular to the plane of the surface causes the desired region to move out of the viewing region as it increases in magnification (Figure 4). This is a fundamental geometric limitation of the configuration of the viewing frustrum used in single-point perspective projection. Document lens [18] solves this problem by translating the viewpoint so that it remains directly above the focus.

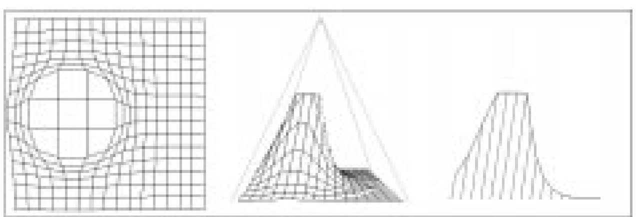

Figure 5: Off-center curve projected to the viewpoint

Since it is our goal to have multiple foci, the solution is to replace the vector perpendicular to the plane with one directed to the eye (Figure 5). This *ToEye* vector is derived from the center of the magnified region on the plane to the viewpoint and is used for all points within the domain of the curve. Simply projecting each point towards the viewpoint would result in all points converging along these vectors at a rate that cancels out the effects of perspective transformation (Figure 6).

Also it is desirable to keep the systems response to a user's actions independent of the location of the focus on the viewing plane. If a unit length vector were used as the basis of the distortions at each point, the *z*-component (which we assume to be normal to the undistorted surface) would differ for each location. What is required is that not the length but the *z*-component of each vector be constant. This is accomplished by dividing each *ToEye* vector by the same constant (Figure

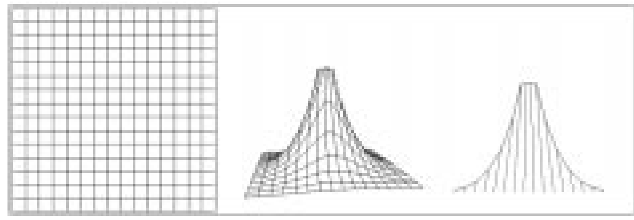

Figure 6: All points projected directly towards the viewpoint and resulting lack of distortion

7). Within the focus region, where all points are projected to the same height, scaling is still preserved. In fact this solu-

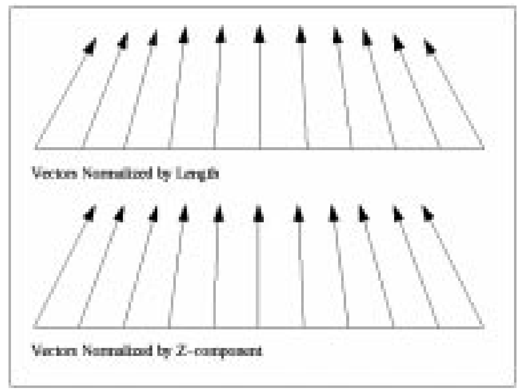

Figure 7: projection vectors: top; normalized to unit length, bottom; normalized to unit z-component

tion has the added benefit of allowing a user to specify more than one focus at a time (Figures 8 and 10).

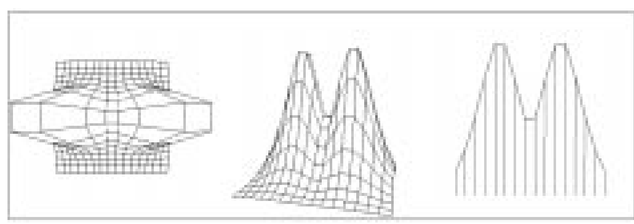

Figure 8: Two foci, vectors perpendicular to plane

Multiple foci

In the case of overlapping multiple foci the end result should be a smoothly curving surface, allowing the entire image to remain visible. A point under multiple curves will have a projection vector associated with each curve. Simply using the vector of the highest curve will result in a discontinuity where the dominance switches from one curve to another. Because all curves have vectors that converge as they approach the viewpoint, points to either side of this discontinuity may, if the horizontal displacements are sizeable, reverse their ordering across the 'seam' (Figure 9).

To prevent this buckling adjoining curves are blended across their seams. From these curves both the height and the direction to which the point is to be projected are calculated as

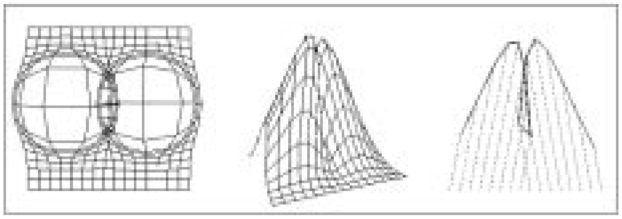

Figure 9: Two foci colliding

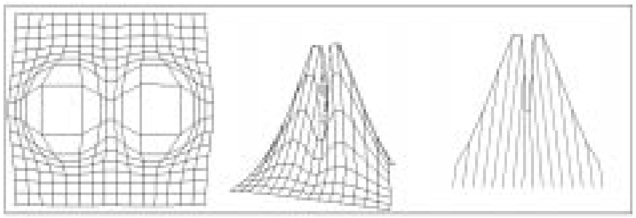

Figure 10: Two foci, now resolved using blending with each foci's ToEye vector

follows:

- The height will be *max*(h_p) of the curves.
- The *ToEye* vector of each curve is weighted by the curve's height at the points location.
- These weighted *ToEye* vectors are then averaged.
- This new averaged *ToEye* vector it then renormalized such that its z component = 1.
- The point is then projected a distance *max*(h_p) along the direction of the averaged *ToEye* vector.

This blending (Figure 11) creates a gradual shift in the direction of the projection vectors between curves allowing for larger distortions to interact more closely, while still maintaining a continuous smooth (unwrinkled) surface.

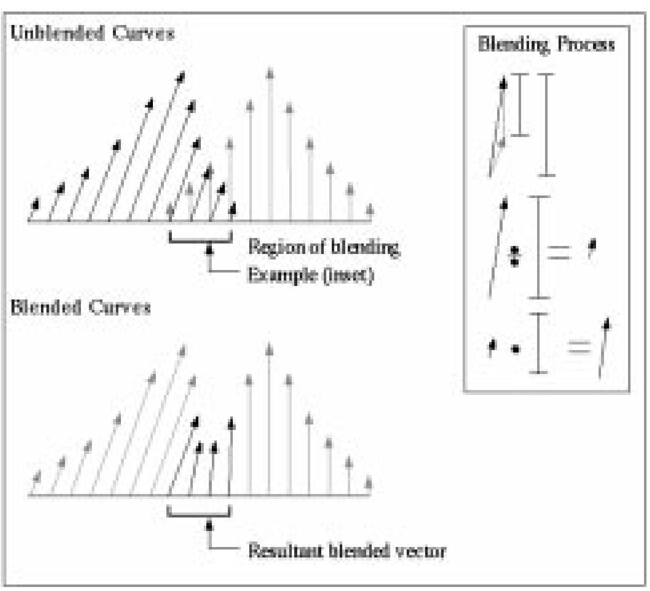

Figure 11: Method of blending vectors

Foci with arbitrary shapes

The single point function foci discussed so far can be extended to provide foci of other arbitrary shapes as well, for example lines or polygons (Figure 12). Now the height h_p of a point outside of a focus but within a region is determined not by its distance from the center of the region but by its distance to the edge of the defined focus. If the point is either on the line or within the polygonal focus it is projected to the full height h_c of the curve. The center of the arbitrary region is still used to determine the vector to the eye.

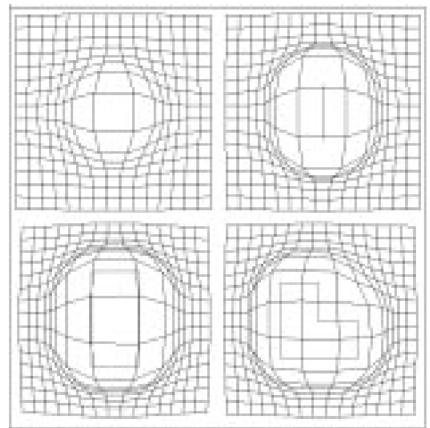

Figure 12: Various shapes of foci: clockwise from the top left: point, line, concave and convex polygons

Distortion control

In any distortion viewing tool compromises are made between the amount of magnification in each foci and the amount of compression in the rest of the image. Our model offers the user considerable control not only over how much compression there is but where minimum and maximum compression occurs. The pattern of compression is a direct result of the slope of the curve. We allow control of the curve's profile through the height h_c (Figure 13) and standard deviation s_c (Figure 14). In this manner it is possible to choose a more gradual integration from focus to context or to limit the extent of the distortion, causing more compression in the distorted region.

All curves have their characteristic profiles and resulting patterns of compression; the gaussian in particular tends to have a broad top around the focus where the magnification of the adjacent area lags only slightly behind that of the focus. In some applications the accompanying magnification of the section around the focus is ideal, providing good local context. However, in other situations this uses too much screen space.

Another characteristic of the gaussian curve is its familiar bell shape that tends to result in a ring of high compression where the tangent to the surface of the curve nears coincidence with the vector from the surface to the viewpoint (Figure 15). In order to adjust this distribution of compression the profile of the basis curve may be modified by subtract-

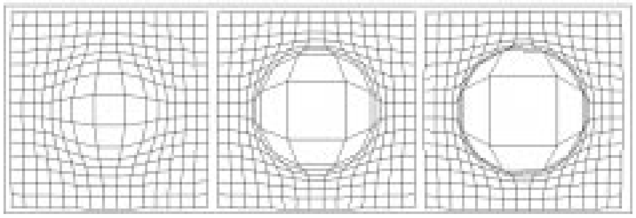

Figure 13: Single foci: effects of varying height with fixed deviation

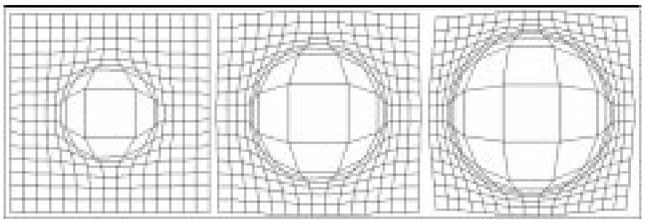

Figure 14: Single foci: effects of varying deviation with fixed height

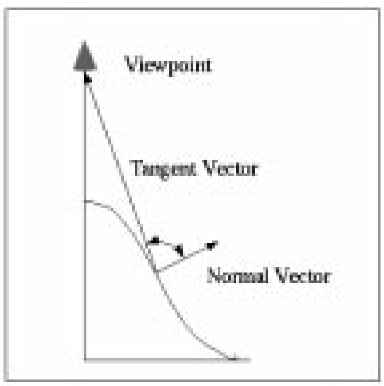

Figure 15: Relationship between compression and angle of surface to viewer

ing from it a second function, in this case a simple half sine wave (Figure 16). The domain of the sine $(0, \pi)$ is normalized across the domain of the gaussian curve (defined to extend to a distance of three standard deviations, beyond which the result of the height calculation is negligibly small).

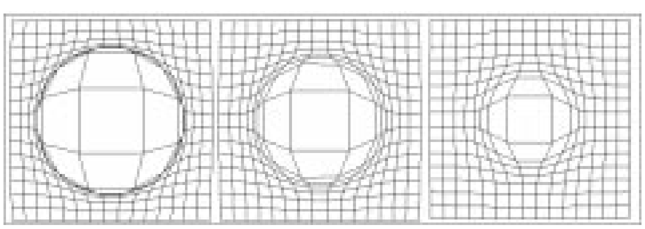

Figure 16: Single foci: effects of varying the secondary function

As all of these distortion controls are left up to the user, it is possible to extend the distortion to a point that causes some areas to be compressed beyond visibility. However, as the slope and curvature are adjustable and reversable directly by the user, it is possible to interactively redistribute the context in non-focal areas without losing focal magnification.

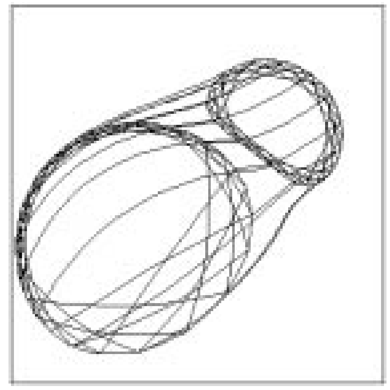

Figure 17: Role of long lines in revealing distortion

Displaying nodes and edges

At this point all of the images we are displaying with 3DPS are stored as graphs; that is they consist of nodes and edges. Nodes are currently displayed either as single points or as squares. Edges are segmented small enough so that they will lie snuggly against the surface. This actually provides additional information about the surface. Information spaces that contains long lines now aid in the description of the surface (Figure 17).

COMPREHENSION

A primary goal of this work has been to make the distortion visually comprehensible. The separation of the image from the distortions of the surface means that the original topology of the image is maintained across the surface. Once the surface is manipulated the image is dropped onto it. Displaying a surface in such a manner as to reveal its three-dimensional form provides the perceptual information that describes the distortion.

Seeing the distortion

A map of the greater Vancouver area (Figure 18) is used to illustrate the visual cues provided. Figure 19 shows the same map with a single focus containing Stanley Park.

A simple cue is to outline the edges of the distorted areas. These outlines show at a glance the size and shape of the magnified scaled areas and the extent of the distorted regions. This is computationally simple and visually minimal, and does not provide degree of distortion information. For this it is necessary to reveal more detail about the shape of the surface. A grid can be displayed over the entire surface, providing both curve and perspective information (Figure 20). Its lines indicate relative magnification as well as serving as a texture gradient. Alternately (Figure 21), or simultaneously (Figure 22), the entire surface can be shaded by placing control points for a NURB surface at the grid intersections and rendering the NURB with a simple lighting model.

All of these visual cues are optional and are displayed in

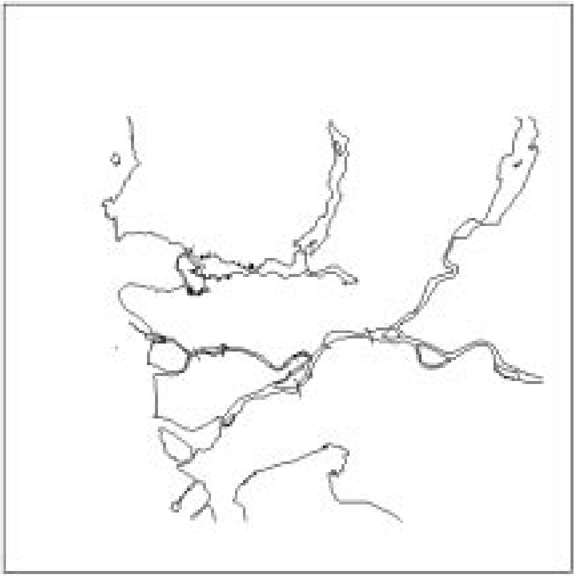

Figure 18: Undistorted map of greater Vancouver area

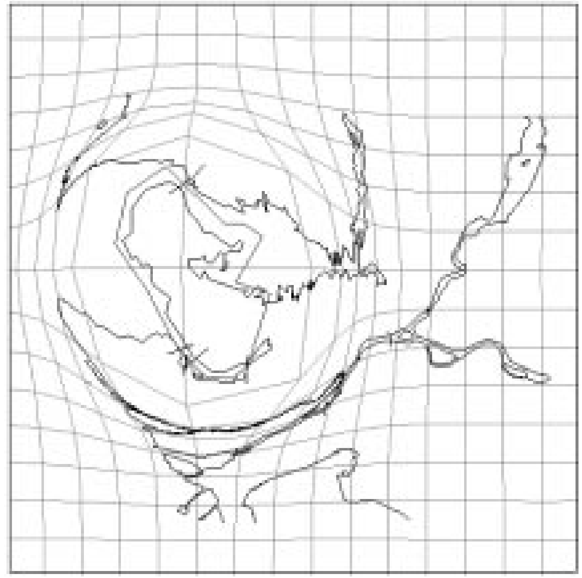

Figure 20: Distorted map with orthogonal grid overlay

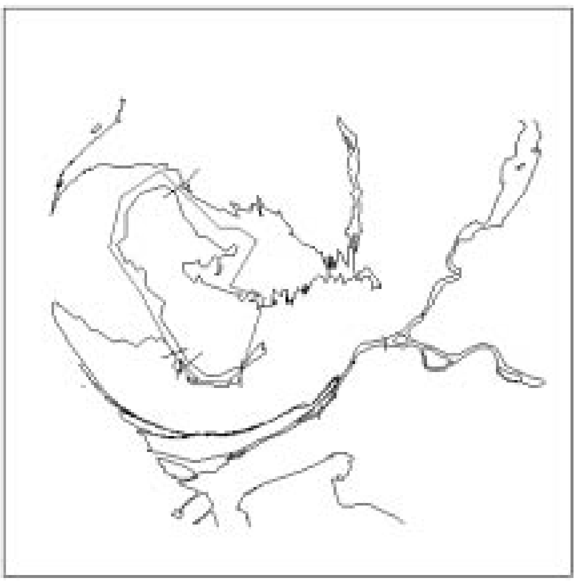

Figure 19: Distorted map of greater Vancouver area, focusing on Stanley Park

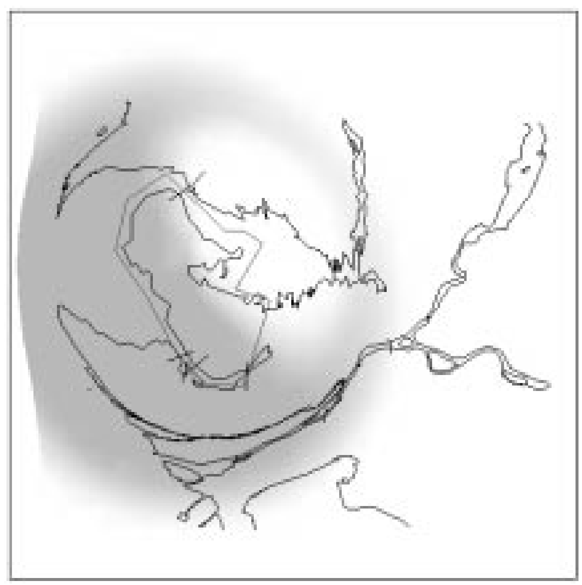

Figure 21: Distorted map with presentation surface shaded to reveal form

shades of grey, so that while they are readily visible apart or in unison they do not dominate the image. This is crucial for the shading since the interpretation of it being the result of luminance has much to do with its use in the establishment of 3D shape.

Maintaining Visual Frame

In maintaining context it is important that the edges of the image stay in sight and are recognizable. A good deal of context information is provided by the fact that large portions of the image remain undistorted and that each distorted area is itself bounded and blends back into the undistorted background. However, as the vector used for the whole curve is the *ToEye* vector for its center it is possible for the edge of a curve as it is pulled up to slide out of the viewing frustrum. One solution is to add a ring of *ToEye* vectors around the edge of the viewing surface such that the vectors are averaged in to counteract the sliding effect. While this does help in keeping the edges of the surface in the viewing frustrum, it also compresses the area available for undistorted foci. Taking advantage of a three-dimensional viewing environment we can simply expand the field of view. While slightly compressing the whole image this will maintain the current distortion and bring all edges back into view.

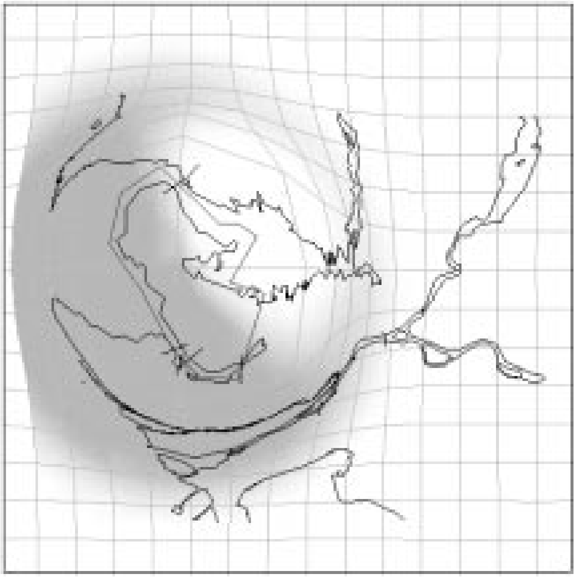

Figure 22: Distorted map with shading and grid

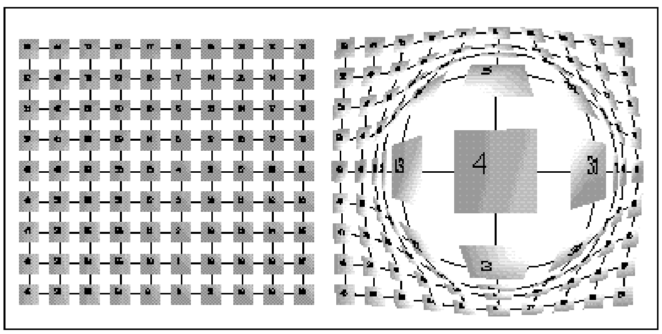

Figure 23: Illustrating degree of magnification

Reversibility

Another important factor in comprehending the distortion is reversibility [15]. If one is making use of this tool as a browser one does not want the undistorted original image to be unrecoverable, in part or in total. Also, if the distortions were quickly and easily reversible the user could use this to reference the original topology. Therefore pulling and pushing should have equal and opposite effects.

Magnification

To indicate the amount of magnification that can be provided we expand a single node in a grid graph. Figure 23 shows the undistorted graph on the left, and on the right shows the same graph with a central node expanded.

Not only is a high degree of magnification available but because the image itself has not been distorted very densely packed information will expand to scale. For example, we scaled the map of the greater Vancouver area (Figure 18) into place in our map of North America (Figure 24). Displaying this combined map undistorted just shows a slightly larger than usual black spot where Vancouver resides. Pulling this

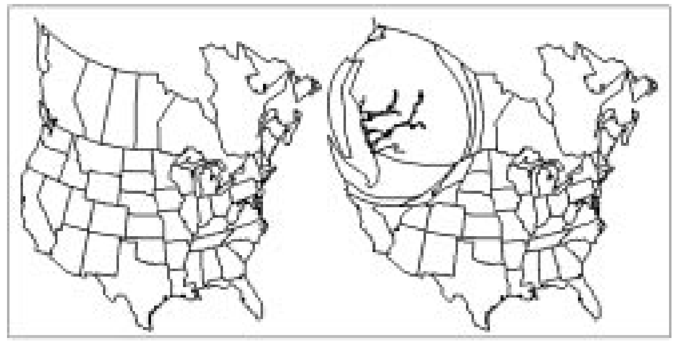

Figure 24: On the left, a map of North America; on the right Vancouver coastline revealed in the map of North America

section up reveals its detail.

3DPS is capable of a great variety of focal shapes and with many patterns of magnification and compression. To illustrate just a few of the possibilities we have included a two foci image of the Vancouver area (Figure 25), and a three foci image of North America (Figure 26). The last series shows the Vancouver area with four focal points (Figure 27, 28 and 29). The contrast between the images with and without distortion clues demonstrates how either the shading or the grid can disambiguate the distortions.

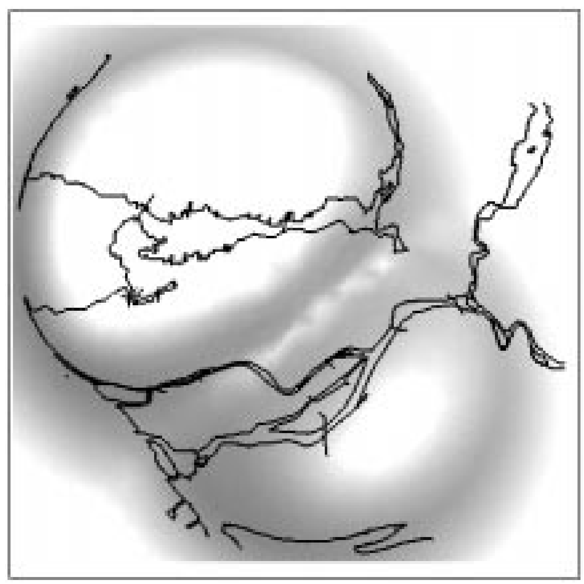

Figure 25: Vancouver area with two focal points; Burrard Inlet and Annacis Island

CONCLUSIONS

This paper presents a three-dimensional pliable surface as a tool for addressing the screen real estate problem. The approach integrates the desirable properties of previous methods to considerably increase the amount of information that can usefully be presented on a computer screen. Through several examples we have shown that the viewing aspect of this tool can handle very dense information.

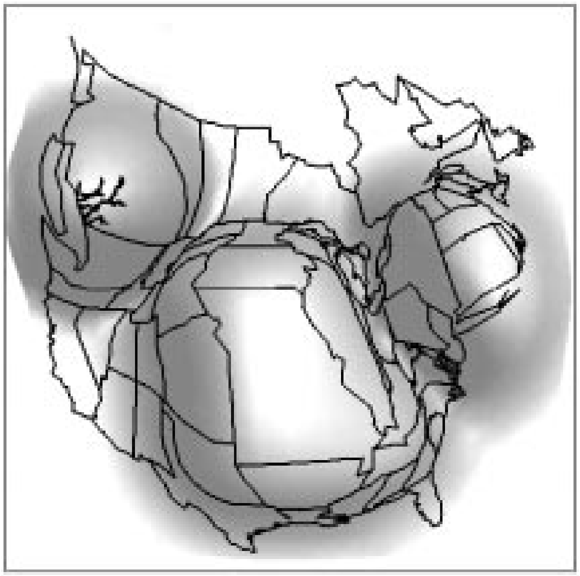

Figure 26: North America with three focal points; Vancouver, Missouri, and Connecticut

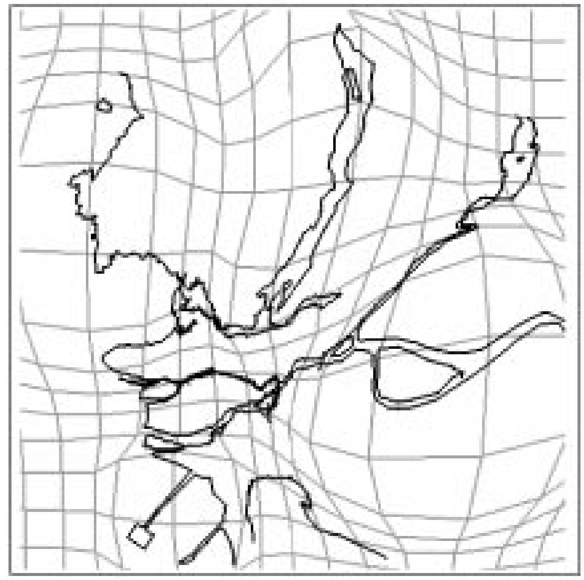

Figure 28: Vancouver area with four focal points, distortion revealed by the grid

Figure 27: Vancouver area with four focal points; Atkinson Point, Indian Arm, Barnston Island, and Tsawwassen (for an undistorted map see Figure 18)

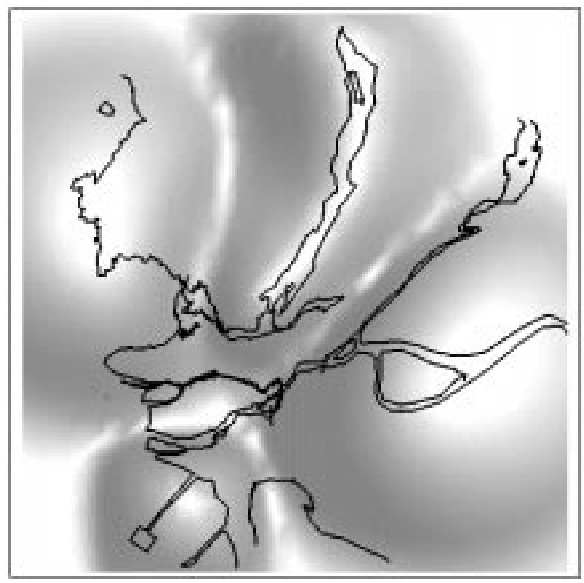

Figure 29: Vancouver area with four focal points, distortion revealed by shading

Our approach makes use of a distortion transformation based on a simple mathematical function (gaussian curve). However, although gaussian curves were chosen because of their gentle curvature out from the focus and back into the context they have a tendency to be radially compressed about halfway up the slope. We addressed this problem by allowing for modification of the gaussian with an auxilliary curve, a half sine wave. To further explore the applicability of other types of curves we propose to build a curve library.

Performance comparisons are difficult across varying platforms and immplementations. However, in terms of complexity this algorithm is comparable to the morphing based approach [20]. Both are $O(nfp)$, where n is the number of points to be projected, f is the number of focal points. The only difference is that we define p as the average number of focal polygon points, while for their algorithm p is the average number of vector pairs per focal polygon. As these vector pairs define the edges of the focal polygons the difference is negligible. However, their algorithm may need to be applied iteratively to keep the resulting image within its window. Furthermore, to prevent change in relative ordering of nodes along x and y dimensions, they limit the possible degree of focal stretch or magnification. Instead of applying similar limits we carefully blend interfocal compression to retain relative ordering of nodes and allow for curvature adjustment to counteract compression that has become too extreme. As a result, while similar to the morphing based approach [20] in complexity and in providing views of multi-focal detail within context, 3DPS offers: a continuous smooth surface between foci where spatial relationships amongst the data points do not transpose, freedom of foci size, and foci positioning. Also, the choice of focus shape is extended to include concave polygons.

3DPS extends the three-dimensional perspective approaches used in [11, 18] for a single focus into a multi-focal tool. Also, rather than rely on the specific shape of the distortion or any characteristics of the information, precognitive perceptual cues are used to reveal the nature of the distortions. We intend to further extend these cues with the addition of aerial perspective. Being able to understand the distortion provides knowledge about the degree of compression, information about the original undistorted topology of the graph, and the cumulative result of the history of the user's actions.

3DPS also extends user control of the distortion through the height, standard deviation, and curvature. Presently these parameters are unconstrained, therefore it is possible to create curves that obscure some context. However, just what has been obscured is always evident and the actions are reversible.

While the use of shading provides instant recognition of the patterns of distortion, it causes some problems. Finding the right balance between light and dark intensities is difficult to achieve, especially if one wants to have convincing shading both on the screen and in print. Also, shading is computationally expensive and noticeably slows interaction.

There is considerable interest in the application of cognitive science knowledge to interfaces. This stems from a desire to make more accessible interfaces but moreover, by creating and testing practical applications such as this one, we extend our awareness of the utilization of low level perceptual skills to offset cognitive load.

ACKNOWLEDGEMENTS

The authors would like to thank the UIST'95 reviewers for helpful comments that have improved the presentation of this paper. Thanks also to Brian Fisher for verifying our cognitive references, James Strickland and Charles Henrich for the map data, Tom Shermer and Art Liestman for ongoing support, and John Dill for helpful comments. This research was supported by graduate scholarships and research and equipment grants from the Natural Sciences and Engineering Research Council of Canada. Thanks also to the Algorithms Lab, Graphics and Multimedia Lab and the School of Computing Science, Simon Fraser University.

REFERENCES

1. L. Bartram, R. Ovans, J. Dill, M. Dyck, A. Ho, and W. S. Havens. Contextual assistance in user interfaces to complex, time critical systems: The intelligent zoom. In *Proceedings of Graphics Interface'94*, pages 216–224, 1994.
2. T. Beier and S. Neely. Feature-based image metamorphosis. *Computer Graphics*, 26(2):35–42, 1992.
3. M. P. Consens. Visualizing queries and querying visualizations. *SIGMOD RECORD*, 21(1):39–46, 1992.
4. K. M. Fairchild, S. E. Poltrock, and G. W. Furnas. Semnet: Three-dimensional graphic representations of large knowledge bases. In *Cognitive Science and its Applications for Human-Computer Interaction*, pages 201–234. Erlbaum Associates, 1988.
5. B. D. Fisher and Z. W. Pylyshyn. The cognitive architecture of bimodal event perception: A commentary and addendum to Radeau (1994). *Cahiers de Psychologie Cognitive/Current Psychology of Cognition*, 13(1):92–96, Feb. 1994.
6. G. W. Furnas. Generalized fisheye views. In *Human Factors in Computing Systems: Proceedings of CHI'86*, pages 16–23, 1986.
7. N. Goodman. *Languages of Art; An Approach to a Theory of Symbols*. Bobbs-Merrill, 1968.
8. K. Kaugers, J. Reinfelds, and A. Brazma. A simple algorithm for drawing large graphs on small screens. In *Lecture Notes in Computer Science: Graph Drawing*, pages 278–282, 1995.
9. D. A. Kleffner and V. S. Ramachandran. On the perception of shape from shading. *In Perception & Psychophysics*, 52(1):18–36, 1992.
10. J. Lamping and R. Rao. Laying out and visualizing large trees using a hyperbolic space. In *UIST: Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 13–14, 1994.
11. J. D. Mackinlay, G. G. Robertson, and S. K. Card. The perspective wall: Detail and context smoothly integrated. In *Proceedings of CHI'91*, pages 173–180, 1991.
12. D. W. Massaro. *Speech Perception by Ear and Eye: A paradigm for Psychological Inquiry*. Hillsdale, N.J.: Erlbaum Associates, 1987.
13. D. W. Massaro. Attention and perception: An information integration perspective. *Acta Psychologica, Special Issue: Action, attention and automaticity.*, 60(2-3):211–243, Dec. 1985.
14. E. G. Noik. Layout-independent fisheye views of nested graphs. In *Proceedings of the 1993 IEEE Symposium on Visual Languages*, pages 336–341, 1993.

15. J. Piaget. Piaget's theory. In P. H. Mussen, editor, *Carmichael's Manual of Child Psychology*. N.Y. Wiley, 1970.

16. V. S. Ramachandran. Perception of shape from shading. *Nature*, 331(14):163–166, 1988.

17. S. P. Reiss. 3-D visualization of program information. In *Lecture Notes in Computer Science: Graph Drawing*, pages 12–24, 1995.

18. G. Robertson and J. D. Mackinlay. The document lens. In *UIST: Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 173–179, 1993.

19. M. Sakar and M. H. Brown. Graphical fisheye views. *Communications of the ACM*, 37(12):73–84, 1994.

20. M. Sakar, S. Snibbe, O. J. Tversky, and S. P. Reiss. Stretching the rubber sheet: A metaphor for viewing large layouts on small screens. In *UIST: Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 81–91, 1993.

21. E. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 1983.

22. C. Ware. The foundations of experimental semiotics: A theory of sensory and conventional representation. *Journal of Visual Languages and Computing*, 4:91–100, 1993.

23. C. Ware and G. Franck. Viewing a graph in a virtual reality display is three times as good as a 2D diagram. In *IEEE Conference on Visual Languages*, pages 182–183, 1994.

24. C. Ware, D. Hui, and G. Franck. Visualizing object oriented software in three dimensions. In *Proceedings of CASCON'93*, pages 612–620, 1993.