


Extending Distortion Viewing from 2D to 3D



M. Sheelagh T. Carpendale,
David J. Cowperthwaite, and F. David Fracchia
Simon Fraser University

Information does not equal knowledge. For information to become knowledge, we need to interpret and understand it. Visualization in general responds directly to this need. However, even after producing a visual representation, we must address issues involving exploration, navigation, and interpretation of the data. This article addresses visual exploration of 3D information layouts.

Comparing 2D and 3D
information layout
adjustment tools leads
directly to a 3D visual access
tool that clears a line of
sight to any region of
interest.

Several visual exploration techniques have been proposed for 2D information layouts. Many of these try to take advantage of humans' natural visual pattern-recognition abilities to understand global relationships while simultaneously integrating this knowledge with local details. This desire for detail-in-context views (also called fisheye, multiscale, and distortion views) has fueled considerable research in the development of distortion viewing tools. Generally, these tools provide space for magnification of local detail by compressing the rest of the

image. In considering a possible detail-in-context view for 3D layouts, we first examine 2D distortion techniques, bearing in mind the particular 3D problem of occlusion.

While our technique can extend to any type of 3D information display, here we focus on graphs. Most previous distortion viewing work dealt only with discrete displays (principally graphs). Also, graphs are structures well suited to information display. A graph can display a set of objects or entities as nodes and relationships between these entities as edges. This basic entity-relationship structure parallels the basic subject-predicate structure of language. It also arises in cognitive science discussions about the distinction between knowledge as declarative or procedural. The same structure has been suggested as forming the basis for the mental models we use to store information internally. Its prevalence across such a variety of areas concerned

with the way humans interact with knowledge or information indicates that related tools and techniques may eventually find quite diverse applications.

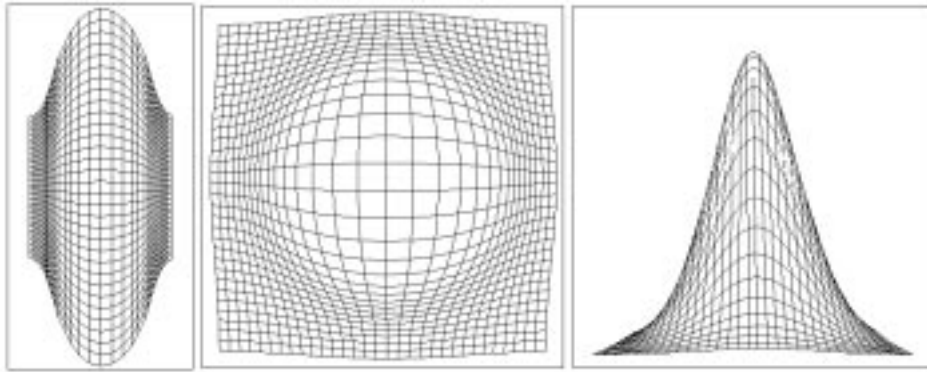
To preserve the representational power inherent in a graph, a distortion-based navigational tool should not disrupt certain properties or relationships in the graph. For example, an entity's spatial position and adjacency relationships may each carry specific meaning. They therefore should be minimally disturbed.

Partly because a graph's structure does not limit the size or complexity of the information it can represent, problems arise, such as a disparity between the information size and display size, and the number of information variables compared to the availability and length of display variables. Three-dimensional layouts offer the possibility of some increase in functional display space and an extra positional display variable.

Ware's studies¹ examined the amount of usable space in a 3D graph display. Discounting the two naive extremes of either an n -fold increase ($3D = n^3$ space or $n \times 2D$ space) or no increase from 2D space, since what is really seen is a 2D projection of the 3D space, Ware's results indicate approximately a three-fold increase in usable space in a 3D display. While encouraging because it seems to indicate that we gain something from our familiarity with 3D space, this result also shows that the fact that we can only see the 2D projection imposes a fundamental limitation. The amount of usable space relates much more closely to the 2D size of n^2 than to a full 3D space.

The fundamental problem remains, just as in the real world, that we cannot see through objects. Something between us and what we want to see blocks our view. While full rotation lets us view the 3D display from all angles, it does not eliminate the fact—inherent to working with 3D information—that some data will be buried within a structure, whether a solid model or a complicated 3D graph layout. Rotational ability definitely improves the situation and was tied closely to the three-fold increase in functional space Ware noted.¹

For discrete information displays, information density can exacerbate occlusion. Given the considerable



1 Distortion dimensionality showing a 3D distortion applied to a 2D surface. In projection 1D applied along the x axis (left); in projection 2D applied in both x and y (center); and side view of the 3D curve (right).

body of work addressing information density in 2D displays, those techniques might well apply to 3D displays. Specifically, can we apply distortion techniques to 3D in a manner that deals with occlusion and preserves context as in the 2D applications?


We believe that studies supporting integrated detail-in-context viewing in 2D displays extrapolate to 3D. For instance, Furnas² indicated that humans store and recall information in great detail for areas of interest and gradually decreasing detail for the related context. This characteristic presumably applies for information in general regardless of the type of display. Certainly more general cognitive support for integrated displays—which permit use of our visual gestalt abilities and minimize cognitive load—applies to our ability to assimilate and interpret information, not to a particular style of display.

Information viewing and dimensionality

Researchers have used both filtering and distorting approaches to create detail-in-context views. Spence and Apperley³ introduced distortion viewing with the Bifocal Display, and Furnas² introduced filtering with Generalized Fisheyes. Unlike filtering, most distortion techniques present all aspects of the image even if very compressed. Some techniques combine the two. For instance, Continuous Zoom⁴ uses both filtering and distortion, and Graphical Fisheye⁵ creates a graphic interpretation of Furnas's filtering method using compression as well as removal. We will discuss primarily spatial reorganization of an existing representation—thus, distortion.

One-, two-, and three-dimensional information representations are common, and a collection of viewing tools exists for each. A distortion can be applied along the x , y , or z dimensions of the computer display or in a combination thereof. Most current distortion techniques use a 2D distortion applied to a 2D information layout.

However, the dimensionality of the information representation and the dimensionality of the viewing technique do not have to match. Figure 1 shows a 3D distortion applied to a 2D surface from 3-Dimensional Pliable Surfaces (3DPS).⁶ The distortion relies on perspective projection to create its reorganized views. The



2 Distortion patterns found in Bifocal Display (left), Perspective Wall (center), and Document Lens (right).

left and center images show the resulting projection. The right image shows the 3D curve from the side. In the leftmost image the distortion is applied along the x axis and not the y axis. In the center image it is applied in both the x and y directions.

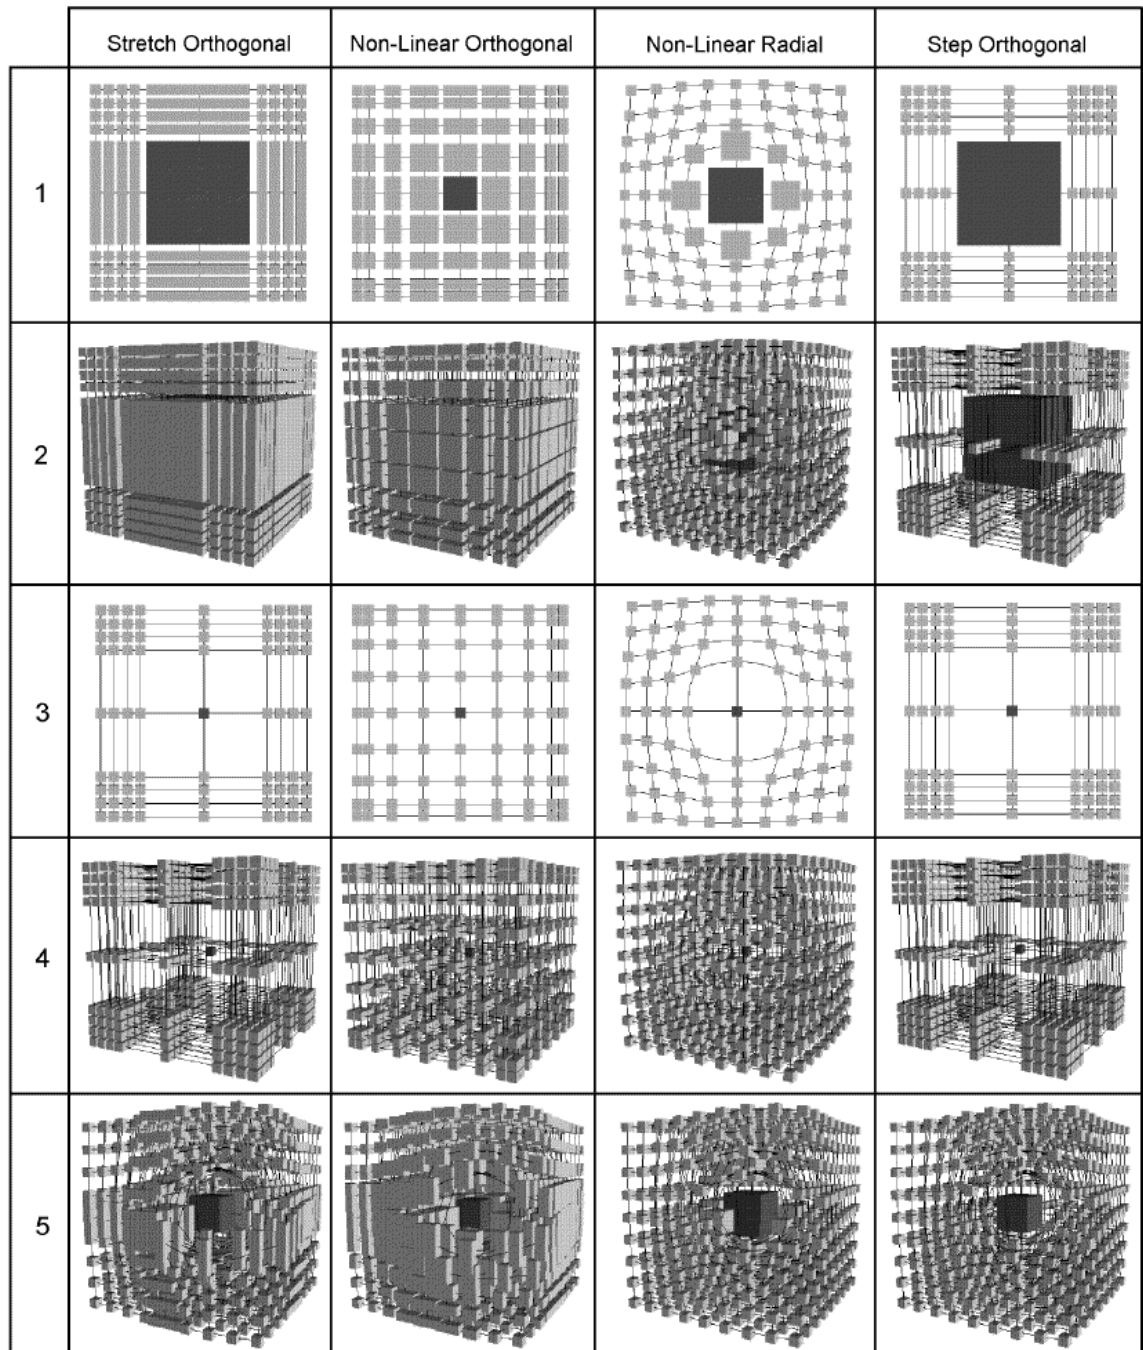
Other examples of discrepancies between dimensionality of the representation and the distortion include Bifocal Display³ (Figure 2, left image) and Perspective Wall⁷ (Figure 2, center image). They apply, respectively, a 1D and a 3D distortion to a linear strip of information that can be thought of as 1D or, since it has width, 2D. Document Lens⁸ (Figure 2, right image) offers a single rectangular focus through a 3D distortion of 2D text fields.

Viewing techniques for 3D data

Currently, the primary methods for accessing 3D space either adjust the viewing angle (rotation) or the viewing position (navigation). Combined, these two would seem to allow all possible views. However, the many problems identified include loss of context when flying through, loss of orientation, and the ever-present problem of occlusion. Nondistortion approaches to accessing the internal details of 3D structures use cutting planes, layer removal, and transparency. Cutting planes and layer removal provide visual access but remove context, while transparency requires some compromise between obtaining visibility and maintaining context.

Previous 3D detail-in-context approaches include Fairchild et al.'s Semnet⁹ and Mitra's aircraft maintenance approach.¹⁰ Semnet included three techniques. One, which uses semantics for positioning, creates an octree to display the focal region in full detail and more remote regions in progressively larger sections. This approach suffers from the sudden changes that occur between boundaries of regions of differing scales. A sec-

Distortion Patterns: A Visual Comparison



3 This chart presents visually the effect of different distortion techniques on spatial organization.

ond approach, based on density, samples more fully around the focus and less frequently as the distance from the focus increases. This approach would increase the congestion and therefore the occlusion problems in the focal region. Third, Fairchild noted the implicit fisheye provided by perspective in a 3D display. A natural single focal point exists for the information in the foreground.

Mitra¹⁰ suggested using linear radial distortion with interactive filters for aircraft maintenance diagrams—3D exploded views of aircraft assembly parts. An adjustable threshold produces a filtered view based on

the function of parts rather than proximity in the diagram. The user could adjust the threshold level to create views with more or less context. In this case exploding and filtering the view does create the space required to see into the structure, but doesn't ensure an unobstructed view. Moreover, progressive filtering removes much of the context, and the overall structure is not apparent in the overall exploded view.

While Fairchild and Mitra focused on providing detail-in-context views, the viewing techniques for 3D data largely concern various types of removal and filtering.

This is probably because in a 3D data display some parts of the display prevent you from seeing other parts—confirming the importance of addressing occlusion.

Examining the distortion viewing techniques developed for 2D data reveals possible extrapolations to 3D data. Rather than critique each technique's usefulness for 2D data, we observe the visual results of spatial reorganization patterns applied to 3D data.

Two-dimensional distortion patterns

Much of the considerable recent work on developing viewing tools for 2D information displays has focused on displaying sufficient detail within the global context. This prompted a general notion of distortion viewing or multiscale diagrams¹¹ where different sections of the information are displayed at different scales. These differing scales of magnified detail and compressed context can be integrated through various distortion functions.


Each approach produces characteristic distortion patterns. We examine displacement separately from magnification, one of several useful distinctions¹² when considering distortion viewing functions. The displacement function adjusts the display to accommodate the increased amount of space the magnified focus requires. Viewing tools generally apply displacement and magnification functions simultaneously. However, Leung and Apperley¹³ introduced the possibility of a distinction between magnification and displacement. They discussed distortion viewing in terms of a transformation or displacement function with a derivative magnification function.

Since our purpose here is exploring what types of distortions might prove useful with 3D data, we examine the 2D distortion functions from the perspective of the resulting visual pattern. The top row in Figure 3 contains a sample of these 2D distortion patterns illustrated on a simple 2D grid graph. While not exhaustive, this set of techniques represents the types of distortion currently used in 2D. All four examples in row 1 show the characteristic patterns created by a traditional application of 2D distortion methods using displacement and magnification simultaneously: stretch orthogonal, nonlinear orthogonal, nonlinear radial, and step orthogonal techniques.


Many 2D techniques offer a choice of one or more of these distinctions. For example, Multi-Viewpoint Perspective,¹⁴ Catgraph,¹⁵ Rubber Sheet,¹⁶ and Shrimp¹⁷ all offer both orthogonal and radial approaches. As a result some of these techniques will appear under more than one of the following headings.

Stretch orthogonal

The first example (Figure 4 and row 1, column 1 in Figure 3) shows a simple orthogonal stretch formed by stretching all data on either of the two axes centered at the focus and compressing the remaining areas uniformly. Bifocal Display by Spence and Apperley³ applied this distortion in one dimension and thus introduced the notion of distortion viewing to computational displays. They created a single-focus detail-in-context view for a personal information space displayed in a long



4 Orthogonal stretch applied to a 2D grid graph.



5 Nonlinear orthogonal distortion applied to a 2D grid graph.


strip. Leung extended this technique to use 2D stretch in 2D Bifocal Display.¹⁸

The image in Figure 3, row 1, column 1 shows a 2D stretch with a single focal point. Subsequently, the orthogonal approach in Rubber Sheet¹⁶ used this distortion to provide multiple focal points. The resulting distorted image uses available screen space well but has entire rows and columns of distorted data. In a multiple foci situation these stretched rows and columns create unrequested or “ghost” foci where they intersect.


Nonlinear orthogonal

The second example (Figure 5 and row 1, column 2 in Figure 3) shows a nonlinear orthogonal approach. Here the focus is magnified to the requested amount, and the magnification decreases according to some function of the orthogonal distance from the focus. This more gradual integration into the foci's immediate surroundings either limits the amount of magnification in the focal region or causes more extreme compression at the edges.

6 Nonlinear radial distortion applied to a 2D grid graph.



7 Step orthogonal distortion applied to 2D grid graph.



Various mathematical functions have been used with the nonlinear orthogonal approach, including arctan in Catgraph¹⁵ and Multi-View Point Perspective,¹⁴ and the hyperbola in Hyperbolic Space.¹⁹ This figure's particular pattern is based on the sine curve.

The nonlinear orthogonal technique supports smooth integration of the focal area into its surrounding context. One problem noted with this approach is that the compression can become extreme at the edges if you allow much magnification in the focal and adjacent regions.

Nonlinear radial

The third example (Figure 6 and row 1, column 3 in Figure 3) is distinct from the first two because of the radial application of the magnification and distortion functions and because it is a constrained distortion. Note that the nonlinear function provides an effect of relative adjacent magnification similar to the image in column 2. However, the radial application causes adjacent edges to curve away from the focus. As a result items directly above

and below or side by side shift slightly. This interferes with the orthogonal relationships in the original grid.

Misue et al.²⁰ asserted the importance of preserving orthogonality, proximity, and topology in creating distorted views that do not interfere with our mental map of the original image. As the other columns of Figure 3 illustrate, orthogonal distortions certainly respect orthogonality, but some have argued that radial distortions best preserve proximity.^{12,17} This distortion pattern is constrained because its effects diminish towards the edge of the image. While the distortion in Figure 3, row 1, column 3 is minimally constrained, you can see that the outer rows of the grid are hardly affected. Most radial distortions suffer from extreme compression and distortion (the image having become virtually circular) at the edges. Constrained distortions were introduced in Pliable Surfaces⁶ and subsequently used in nonlinear transformations.²¹

Creating a distorted view magnification in one place occurs at the expense of compression in another. Avoiding extreme compression at the edges of the image by constraining the distortion does not avoid the compression entirely. Pliable Surfaces⁶ provides user control of both the location and relative rate of compression.

Step orthogonal

The fourth example (Figure 7 and row 1, column 4 in Figure 3) displays a step orthogonal approach. This performs the same distortion as the space filling orthogonal but leaves the data in the rows and columns aligned with the focus unstretched. This basic approach creates less data distortion but leaves more unused space. It also causes a marked grouping of the data not related to the information itself and could lead to misinterpretations. The Zoom family of viewing techniques^{4,22} uses this method, as does the more recent Shrimp Views.¹⁷


Introducing distortion viewing to 3D

Certain problems arise when applying 2D distortion techniques to 3D displays. We intend to demonstrate the distinctive patterns of different types of distortion and show that while they offer considerable advantages in a 2D display, a naive application does little to improve access in 3D. Figure 3 serves as our reference for this discussion.

Figure 3 is organized as a chart. The top row displays four 2D distortion techniques, each used exclusively in the column it heads. We chose to use 2D and 3D grid graphs because these simple structures most clearly reveal the underlying patterns of the distortion functions. We chose a central focal point because we want to explore revealing obscured foci.

Figure 3, row 2

The second row shows a direct naive extrapolation of the 2D schemes to a 3D grid graph. Note that in the first two columns the distortion pattern propagates straight through to the surface. In fact, given focal points on the surface of a 3D structure, we achieve the same benefits that these distortion patterns obtain in 2D. This would continue to apply to any chosen focus visible in the 2D projection. However, simply applying these approaches to a 3D display does more to obscure a central focal



8 Cross-section views illustrate the visual access distortion algorithm: calculating the direction and distance from line of sight (left), calculating the displacement (center), and displacing the occluding objects (right).

point than reveal it. In fact, the usual problem of some objects occluding others in 3D layouts is exacerbated in distortion approaches with space-filling aspects, notably columns 1 and 2.

Applying the radial Gaussian function in 3D best preserves the actual appearance of the 3D grid itself, as the function only minimally extends to the edges. However, the magnification/displacement appears as increased congestion in the center.

The amount of displacement at the edges of the orthogonal step function (row 2, column 4) does provide a view of the internal focal node. While this hints that displacement by itself might be useful, the resulting view is not entirely satisfactory—it still does not allow viewing from all angles, and the artificial groupings are pronounced.

For distortion to help us fully examine the internal aspects of 3D data, we need unrestricted visual access to the chosen focus. Furthermore, if we expect to provide context, it would be preferable to avoid radically reorganizing the data.

Figure 3, row 3

Following the insight provided by the naive application, the third row presents the same set of functions, revealing the displacement-only aspect on the 2D grid. Note that the stretch and step orthogonal (columns 1 and 4) resolve into the same pattern.

Figure 3, row 4

Row 4 applies this displacement-only distortion to the 3D grid. Despite eliminating the obscuring magnification, little improvement results from applying graduated and radial techniques (columns 2 and 3). Note that while the orthogonal approaches had seemed a less efficient use of space in two dimensions, in three dimensions the separation provides partial visual access. However, it creates artificial groupings that can still occlude the focus during rotation. The partial solution provided by the displacement-only patterns indicates the potential usefulness of using distortion to remove occluding objects.

Observations

At this point we have determined that a displacement-only function might best provide visual access. However,

it appears that aligning this function with the data creates artificial groupings of apparent significance. Also, limiting the spread of the distortion produces a much more recognizable exterior, and the objects that concern us lie only between the focus and the viewer.

On the other hand, it seems that the magnification still aligns more appropriately with the data. For instance, the choice between relative local magnification or focal-only magnification depends on the task and information.


These observations led us to apply two techniques first developed in our 2D distortion method, 3DPS⁶—viewer-aligned distortion and constrainable distortions. In 2D we aligned focal regions with the viewer to keep more than one in sight and prevent the focal regions from occluding each other. In 3D we actually apply the displacement distortion radially along the line of sight, permitting interactive displacement of objects that obstruct the view. In 2D we constrained the distortion to maintain as much undisturbed context as possible and to give the user interactive choices on the compression's location and pattern. Applying the constrained distortion in 3D directly parallels this.

Visual access distortion


Visual access distortion²³ is a viewer-aligned, radially constrained, reversible distortion that clears the line of sight to chosen focal regions. We believe that effective 3D detail-in-context viewing requires

- controlling the magnification of a chosen focus or foci to display detail,
- viewing the focus as a 3D object with the usual advantage of rotation (examination from all angles),
- maintaining a clear visual path between the user and the focal point(s), and
- maintaining the surrounding context in a manner that respects the original layout.

Specifically, visual access distortion proceeds as follows. Select a focal point; in Figure 8 (left image) the central point has been selected. Then let L be a line segment extending from the focus to the viewpoint (the line of sight), indicated in the left image of Figure 8



9 This series shows Gaussian visual access distortion applied progressively to the 3D grid.



10 Even when the two foci are in line with the viewpoint, both remain visible because the distortion function from the furthest focus affects all occluding objects, including other foci.

extending from the focus. The vector \mathbf{d} is the shortest vector from an object O in the display and a nearest point P on the line L . In Figure 8 (left image) yellow arrows reach from the focus and the line of sight to the adjacent points. The vector \mathbf{d} defines the direction of the distortion at O , and its length $|\mathbf{d}|$ is used to determine the magnitude M of the distortion.

To achieve smooth integration back into the original data topology, use a Gaussian distribution to determine the displacement's magnitude. The profile of a Gaussian function (Figure 8, center) shows how the $|\mathbf{d}|$ (indicated in yellow) is used to calculate magnitude M of the distortion (indicated in green). For a given value of $|\mathbf{d}|$ you can determine the height of the Gaussian that gives the magnitude M . Figure 8 (right image) shows (in green) using the magnitude M along direction \mathbf{d} to create the displacement. You can control the shape of the Gaussian function, and hence the distribution of the distortion, simply by adjusting the height and standard deviation of the curve. Since the viewing direction is along the line of sight, the distortions will appear to the viewer as radi-

ally symmetrical about the focus, though moderated by the effect of planar perspective projection.

The resulting distortion of the original data provides a clear visual path from the viewer to the focal node. The visibility of the focus persists under rotation of the data or motion of the viewpoint, smoothly deflecting nodes away from the line of sight as they approach it and returning them to their original positions as they move away (Figure 9). The creation of a clear visual path can now be combined with one of the magnification distortions described earlier to permit an unobstructed view of the magnified focus.

Multiple foci

Visual access distortion scales well to multiple focal points. Because each line of sight employs its own access distortion function, you can combine more than one focus in a single view.

In Figure 10 a simple average of the two functions at each point produces clear lines of sight to the two foci. The upper right focus is one layer deep into the $9 \times 9 \times$

9 cube. The lower left focus is eight layers deep, but still visible. As with a single focus, visual access persists during rotation. Figure 10 rotates the lattice (from left to right) until both foci are in line, one above the other, in the rightmost image. Figure 10 proceeds to continue rotating, but from top to bottom instead of left to right. In the left image the closer focal point comes down between the user and the further focal point. However, visual access distortion merely considers it another occluding object and shifts it to one side. The three images in Figure 10 show how close one focal point can come to occluding another as it crosses the line of sight.

Arbitrary graphs


While we chose simple grid graphs to clearly reveal patterns in the distortion techniques, their effectiveness is not limited to this type of 3D grid layout. Figure 11 shows a polar graph layout that positions nodes by randomizing the magnitude of both the radius and angles. This image shows both access displacement and focal magnification. The displacement function applied to the nodes only can leave edges cutting across the focal node. In Figure 11 visual access distortion applied along the length of the edges curves them away from the line of sight, leaving a clear view of the focus.

Figure 3, row 5

The fifth and final row of the visual comparison chart (Figure 3) applies visual access distortion to the four 2D approaches, in each case successfully exposing the focus in context. Here the magnification component from each column's 2D distortion pattern is applied relative to the data, resulting in a range of node shapes and sizes. The displacement is then provided by visual access distortion applied relative to the viewer. Even in cases where the magnification has completely occluded the central focus node, applying the visual access distortion clears a line of sight to the focus.

In row 5, columns 1 and 2, the space-filling orthogonal approach and the graduated sine function had completely occluded the central focus node (see row 2, columns 1 and 2), virtually creating a solid. Similarly, with the radial Gaussian distortion (see row 2, column 3), the central focal node is practically obscured by its neighbors, since they also are magnified, though to a lesser degree. In all these cases visual access distortion provides visibility of the central focus.

In the case of the orthogonal step function, if you don't apply the distortion's displacement aspects, the artificial clusters are not generated (compare row 2, column 4 to row 5, column 4). The actual focus is magnified, while the entire context remains undisturbed. Here, applying visual access distortion achieves the desired focal visibility while minimally disturbing context.



11 Visual access distortion applied to a central node in a random graph. Note how the edges also curve away from the focal node.

Discussion

Distinguishing between data or viewer relative magnification or displacement patterns offers new flexibility in applying these techniques.

Browsing

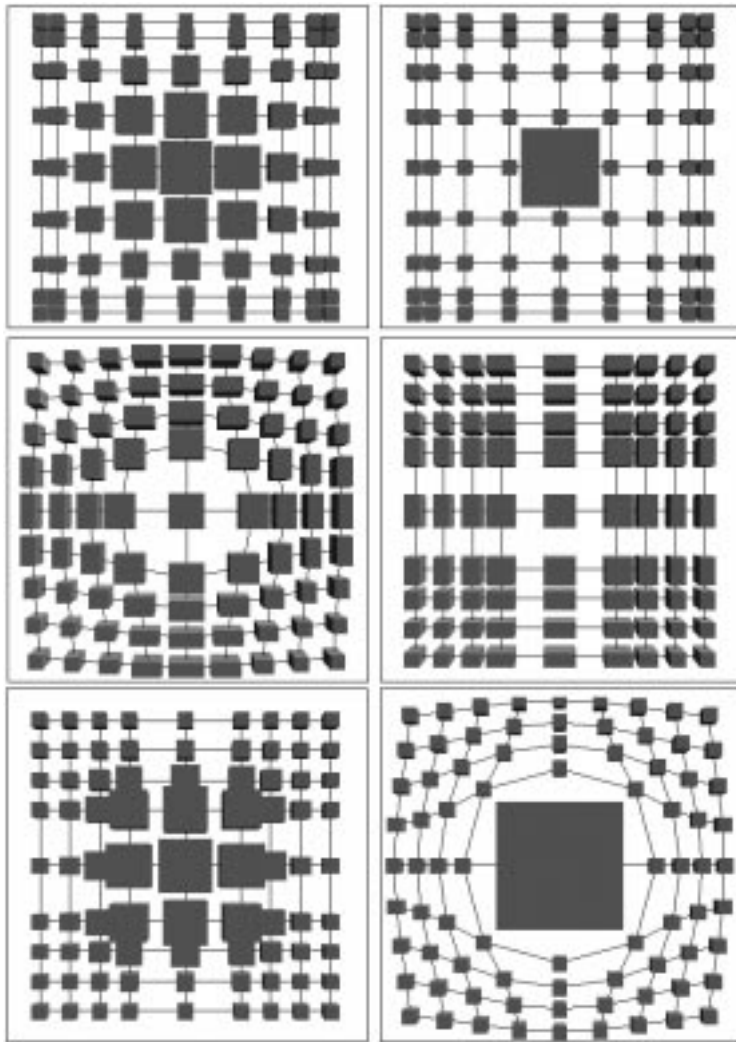
In practice we most frequently use visual access distortion by itself—displacement only. If the focal node requires magnification, we use a simple step function, magnifying the chosen nodes only. By itself visual access distortion allows in-context browsing of a 3D display. With magnification it provides detail-in-context viewing.

A focal point can be either data objects or locations in space. When the focal point is an object, we apply visual access distortion from the viewpoint to the object's center. Browsing can involve sequential selection of objects or nodes. Alternatively, a location in space can set the end point of the line of sight cleared. The user can interactively control this line-segment-of-sight, creating a dynamic probe that moves fluidly through the space.

In browsing a 3D display, the user can select focal type and position as well as which distortion method to use for displacement, magnification, and access. During visual exploration each item is shifted out of the line of sight and then back into its original position. This motion provides effective visual feedback about the context and relative positions of the individual data items.

Other variations

Separating the magnification and displacement functions opens up possibilities for many new distortion viewing variations for 2D data as well. Figure 12 (next page) shows a sampling. The displacement function in the top row is nonlinear orthogonal, as in the second column of Figure 3. In Figure 12 the top row left uses radial Gaussian magnification, and the top row right uses step magnification. In the second row of Figure 12



12 Additional distortion variations.

both images use a nonlinear (sine) magnification curve: radial Gaussian displacement on the left and orthogonal displacement on the right. The bottom row shows radial Gaussian magnification with orthogonal displacement on the left and step magnification with radial displacement on the right.

The last pattern (step-radial) in Figure 12 is currently used as part of Shrimp Views.¹⁷ The others have not yet been explored in actual applications. Using distinct functions for magnification and distortion may improve finding a good match between a distortion viewing approach and the particular information and task at hand.

Conclusions

Future plans for visual access distortion include applying the ideas presented here to both general 3D graph structures and to solid 3D data. We also intend to investigate the potential of perceptual cues (3D grids, color and shading, stereo display) to reveal the nature of the distortions when applied to more general data sets.

Presumably, the availability of this type of access will allow fuller use of the third positional variable. While a

3D computational display will still resolve to its 2D projection when movement stops, being able to interactive-ly shift objects to see behind them will make it possible to plan spatial organization with more freedom. ■

Acknowledgments

Thanks to the editors and reviewers for their comments and direction and to Anne Grbavec for recent help and encouragement. This research was supported by graduate scholarships and research and equipment grants from the Natural Sciences and Engineering Research Council of Canada, Forest Renewal British Columbia, and British Columbia Science Council. Thanks also to the Algorithms Lab, Graphics and Multimedia Research Lab, and School of Computing Science, Simon Fraser University.

References

1. C. Ware, D. Hui, and G. Franck, "Visualizing Object-Oriented Software in Three Dimensions," *Cascon 93 Proc.*, IBM's Center for Advanced Studies, Toronto, Canada, 1993, pp. 612-620.
2. G.W. Furnas, "Generalized Fisheye Views," *Human Factors in Computing Systems: CHI 86 Conf. Proc.*, ACM Press, New York, 1986, pp. 16-23.
3. R. Spence and M. Apperley, "Database Navigation: An Office Environment for the Professional," *Behavior and Information Tech.*, Vol. 1, No. 1, 1982, pp. 43-54.
4. L. Bartram et al., "The Continuous Zoom: A Constrained Fisheye Technique for Viewing and Navigating Large Information Spaces," *UIST 95: Proc. ACM Symp. on User Interface Software and Tech.*, ACM Press, New York, 1995, pp. 207-216.
5. M. Sarkar and M.H. Brown, "Graphical Fisheye Views," *Comm. ACM*, Vol. 37, No. 12, 1994, pp. 73-84.
6. M.S.T. Carpendale, D.J. Cowperthwaite, and F.D. Fracchia, "Three-Dimensional Pliable Surfaces: For Effective Presentation of Visual Information," *UIST: Proc. ACM Symp. on User Interface Software and Tech.*, ACM Press, New York, 1995, pp. 217-226.
7. J.D. Mackinlay, G.G. Robertson, and S.K. Card, "The Perspective Wall: Detail and Context Smoothly Integrated," *CHI 91 Conf. Proc.*, ACM Press, New York, 1991, pp. 173-180.
8. G. Robertson and J.D. Mackinlay, "The Document Lens," *UIST: Proc. ACM Symp. on User Interface Software and Tech.*, ACM Press, New York, 1993, pp. 101-108.
9. K.M. Fairchild et al., "Three-Dimensional Graphic Representations of Large Knowledge Bases," *Cognitive Science and its Applications for Human-Computer Interaction*, Lawrence Erlbaum Assoc., Hillsdale, N.J., 1988, pp. 201-234.
10. D.A. Mitra, "A Fisheye Presentation Strategy: Aircraft Maintenance Data," *Human-Computer Interaction: Interact 90*, Lawrence Erlbaum Associates, Hillsdale, N.J., 1990, pp. 875-880.
11. G.W. Furnas and B.B. Bederson, "Space-Scale Diagrams: Understanding Multiscale Interfaces," *CHI 95: Proc. ACM Conf. on Human-Computer Interaction*, ACM Press, New York, 1995, pp. 234-241.

12. M.S.T. Carpendale et al., *Exploring Distinct Aspects of the Distortion Viewing Paradigm*, Tech. Report TR 97-08, School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada, March 1997.
13. Y.K. Leung and M.D. Apperley, "A Review and Taxonomy of Distortion-Oriented Presentation Techniques," *ACM Trans. on CHI*, Vol. 1, No. 2, 1994, pp. 126-160.
14. K. Misue and K. Sugiyama, "Multi-Viewpoint Perspective Display Methods: Formulation and Application to Compound Digraphs," *Human Aspects in Computing: Design and Use of Interactive Systems and Information Management*, Elsevier Science Publishers, Amsterdam, 1991, pp. 834-838.
15. K. Kaugers, J. Reinfelds, and A. Brazma, "A Simple Algorithm for Drawing Large Graphs on Small Screens," *Graph Drawing 94*, Lecture Notes in Computer Science, Springer-Verlag, Berlin, 1994, pp. 278-282.
16. M. Sarkar et al., "Stretching the Rubber Sheet: A Metaphor for Viewing Large Layouts on Small Screens," *UIST: Proc. ACM Symp. on User Interface Software and Tech.*, ACM Press, New York, 1993, pp. 81-91.
17. M.A. Storey and H.A. Muller, "Graph Layout Adjustment Strategies," *Graph Drawing 95*, Lecture Notes in Computer Science, Springer-Verlag, Berlin, 1995 pp. 487-499.
18. Y.K. Leung, "Human-Computer Interaction Techniques for Map-Based Diagrams," *Designing and Using Human-Computer Interfaces and Knowledge-Based Systems*, G. Salvendy and M. Smith, eds., Elsevier Science Publishers, Amsterdam, 1989, pp. 361-368.
19. J. Lamping, R. Rao, and P. Pirolli, "A Focus and Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies," *Proc. ACM Conf. Computer-Human Interaction (CHI 95)*, ACM Press, New York, 1995, pp. 401-408.
20. K. Misue et al., "Layout Adjustment and the Mental Map," *J. Visual Languages and Computing*, Vol. 6, No. 2, 1995, pp. 183-210.
21. T. Keahey and E. Robertson, "Techniques for Nonlinear Magnification Transformations," *InfoVis 96: Proc. IEEE Conf. on Information Visualization*, IEEE Computer Soc. Press, Los Alamitos, Calif., 1996, pp. 38-45.
22. L. Bartram et al., "Contextual Assistance in User Interfaces to Complex, Time Critical Systems: The Intelligent Zoom," *Proc. Graphics Interface 94*, Canadian Information Processing Society, Toronto, Canada, 1994, pp. 216-224.
23. M.S.T. Carpendale, D.J. Cowperthwaite, and F.D. Fracchia, "Distortion Viewing Techniques for 3D Data," *InfoVis 96: Proc. IEEE Conf. on Information Visualization*, IEEE Computer Soc. Press, Los Alamitos, Calif., 1996, pp. 46-53.



M. Sheelagh T. Carpendale is a doctoral candidate and a research associate in the Graphics and Multimedia Lab in the School of Computing Science at Simon Fraser University. She received her BS in computing science at Simon Fraser University and also has a fine arts background, having attended Emily Carr College of Art and Sheridan School of Design. Her research interests include information representation, visualization, visual languages, user interfaces, and computer graphics.



David J. Cowperthwaite is a PhD candidate and member of the Graphics and Multimedia Research Lab in the School of Computing Science at Simon Fraser University. He received a BS with honors in computer science from York University in Toronto in 1994 in the Space and Communications Sciences program. His research interests include scientific visualization, 3D information presentation and perception, graphics, and animation.



F. David Fracchia is an assistant professor in the School of Computing Science and co-director of the Graphics and Multimedia Lab at Simon Fraser University. He is currently on leave as a senior software developer at Mainframe Entertainment, Vancouver. His research interests include scientific visualization, computer graphics, and mathematical models of biology. Fracchia received an MMath degree in 1988 at the University of Waterloo and BS and PhD degrees in 1986 and 1992, respectively, at the University of Regina. He is a member of ACM, ACM Siggraph, and IEEE.

Contact Carpendale at the School of Computing Science, Faculty of Applied Sciences, Simon Fraser University, Burnaby, B.C. V5A 1S6, Canada, e-mail carpenda@cs.sfu.ca.