Chapter 7

Beyond 2D

While a great variety of visual representations (VReps) are two-dimensional, there are, of course, VReps of other dimensionality. Presenting VReps of any dimensionality on a computer involves many of the same issues. The size of the screen tends to be very much smaller than the VRep at a reasonable resolution. Therefore readability and navigation issues that gave rise to detail-in-context research will also affect VReps of other spatial dimensionality. Furthermore, with only a few expections [43, 99, 111, 123, 147], detail-in-context presentation is relatively unexplored for VReps of other dimensionality. In this chapter we lift our restriction of considering presentation space only for two-dimensional VReps and use EPS to explore the presentation variations that arise from varying the spatial dimensionality of the VRep and the dimensionality of the distortion.

In this exploration there are three different types of spatial dimensionality to consider: that of our framework EPS, that of the distortion and that of the VRep. The first is the three-dimensional space used by EPS as described in Chapter 3. The dimensionality of EPS will not be changed in this discussion as the intention is to test its versatility and extensibility by changing the dimensionality of the VRep and the distortion.

The second is the dimensionality of the distortion. A distortion operates in the dimensions of the mathematical space in which it exists. While we recognize that this is by no means an inherent limitation of distortion possibilities, we restrict the distortion dimensionality to the three dimensions of our framework. A distortion can be applied along any of the three dimensions or in any combination.

The third is the *spatially dimensionality* of the VRep. The VRep's spatial dimensionality relates to the number of variables used to create its layout, or in Bertin's [13] terms, its use of *positional variables*: a line has one positional variable, an area has two, and a volume has three. Various discussions [160, 161, 13] consider visual dimensionality to include such things as colour, size, value, texture, and orientation. However, spatially a VRep must contend with the limitations of what can be displayed on a computer screen and perceived visually. Even multi-dimensional information will have a VRep whose spatial dimensionality is limited by three dimensions. This discussion explores the effect of changing the spatial dimensions of the VRep.

This discussion is organized according to a simple classification of VReps according to their utilization of spatial dimensions. If one, two, and three dimensional VReps are defined as using one, two, and three axes to create their spatial organization respectively, then there are many types of representations that seem to exist somewhere between these three categories. In other words, they make primary use of one or two dimensions and partial use of another. These VReps have been referred to as being of 1.5 and 2.5 dimensions [49], but we find this notation too precise to capture a broad category that exists between either one and two dimensions or between two and three dimensions. Instead, we use five categories: one (1D), one-plus (1D+), two (2D), two-plus (2D+) and three dimensional (3D). We recognize three-plus dimensional VReps exist but consider that discussion beyond the scope of this thesis. These apply only to the spatial dimensionality of the VRep, not the information that is being represented.

This chapter proceeds by examining, for each category of VRep, the presentation variations that arise when the dimensions of the distortion are changed. Sections 7.1 to 7.5 look at 1D, 1D+, 2D, 2D+ and 3D VReps respectively.

7.1 One-Dimensional Visual Representations

A *one-dimensional visual representation* (1D VRep) is a representation whose layout makes use of only one positional variable. The individual symbols that compose the VRep may have area or volume but the relationships between the symbols is linear, like beads on a string. Any changes in information are indicated only linearly. For example, text symbols are two-dimensional, but as text itself consists of a series of symbols in a row, it is considered to be a 1D VRep. Primary sources for one-dimensional information are single output streams. Monitoring equipment often produces 1D VReps.

The nine node graph in Figure 7.1 is a simple formal example of a 1D VRep in normal presentation. We assume a standard x, y, z space with x denoting the horizontal axis, y the vertical and z the perpendicular. The focus is the slightly darker central node.

Though the VRep is one-dimensional the distortions can be one, two or three-dimensional. The

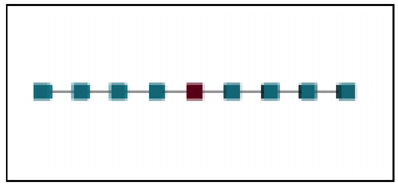

Figure 7.1: 1D VRep: A nine node graph with a linear layout along the x axis in normal presentation

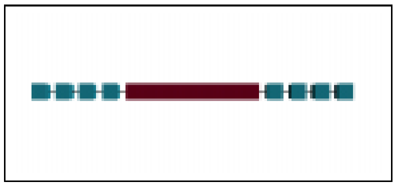

(a) 1D VRep: 1D distortion applied along the x axis

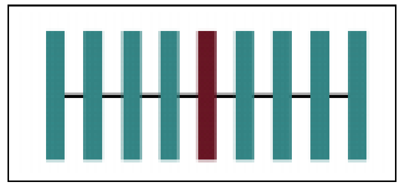

(b) 1D VRep: 1D distortion aligned orthogonally to the VRep's dimension

Figure 7.2: 1D VRep with 1D distortions

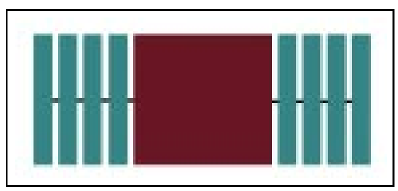

Figure 7.3: 1D VRep: 2D distortion in x and y

series of Figures 7.2 to 7.5 illustrate different presentation patterns that result when the dimensionality of the distortion is changed. Figure 7.2(a) shows a 1D graph with a one-dimensional step distortion applied along the x axis to align it with the single dimension of the VRep. The effect is to elongate the focal node in x.

Figure 7.2(b) shows the linear graph with a one-dimensional distortion applied along the y axis. This distortion stretches the graph nodes rather than changing their spacing. Figure 7.3 shows the same graph with the a two-dimensional distortion. The 2D distortion preserves the proportions of the focal node, compresses the placement of the context nodes along the x axis, and stretches them along the y axis.

The distortion patterns shown in Figure 7.4(a) and Figure 7.4(b) make use of z displacement. As the distance from the focus is calculated from x only they are 2D x, z distortions. Figure 7.4(a) shows a step drop-off function and Figure 7.4(b) a linear drop-off function. With the step function

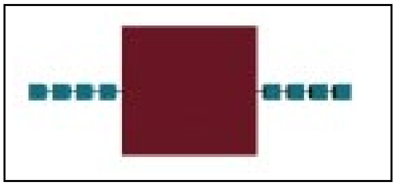

(a) 1D VRep: x and z step distortion

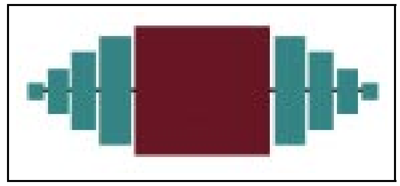

(b) 1D VRep: x and z linear distortion

Figure 7.4: 1D VRep: 2D distortion in x and z

all nodes maintain their proportions but considerable white space is introduced around the context nodes. With the linear function the context nodes have gradually decreasing magnification.

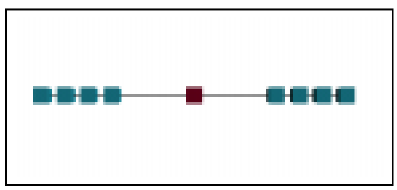

Figure 7.5: 1D VRep with 1D distortions and no magnification, aligned with the visual representation's single dimension

Figure 7.5 shows the same graph with a distortion aligned as in Figure 7.2(a) but used only for displacement. This creates space around the focal node.

7.2 One-Plus Dimensional Visual Representations

A *one-plus dimensional visual representation* (1D+ VRep) is a representation whose layout makes primary use of one positional variable and auxiliary use of another. Figure 7.6 shows three simplified examples of 1D+ VRep types. Figure 7.6(a) shows elementary use of an auxiliary dimension where extra symbols are attached to the primary 1D VRep. Figure 7.6(b) shows a type of a 1D+ VRep that is created when a few related linear output streams are displayed adjacently. Perspective Wall [99] was designed for this type of VRep. Another type of 1D+ VRep organizes its different components in a strip that is predominately linear but has some width (Figure 7.6(c)). Bifocal Display [147] was designed for this type of VRep. Both of these examples are usually referred to in the literature as 1D but clearly make some use of a second dimension. Many common visual information representations fall into this category. Mathematical and musical notation are common examples of 1D+ VReps.

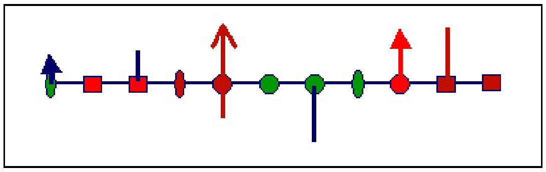

(a) 1D+ VRep: a primary 1D VRep with symbolic extensions into the second dimension

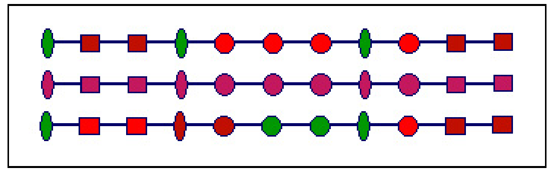

(b) 1D+ VRep: layered 1D VReps

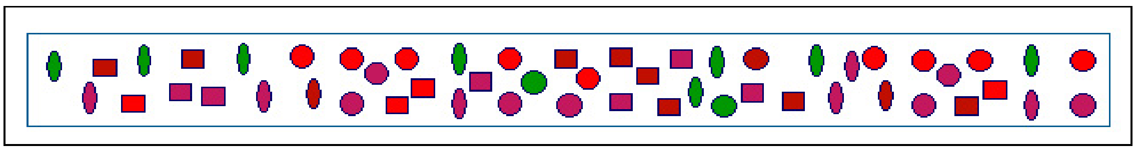

(c) 1D+ VRep: a strip of information with one extremely dominate dimension

Figure 7.6: 1D+ VRep

Visual information is seldom truly one-dimensional. In fact, as the 1D VRep used to illustrate the changing presentation patterns in Figures 7.2 to 7.4 does have some width in y, the illustrations Figures 7.2 to 7.4 also indicate the presentation patterns for 1D+ VRs. Even those examples, that refer to their application as linear information, such as Bifocal Display [147] and Perspective Wall [99], would more appropriately be placed in the 1D+ category. Figure 7.7 shows the presen-

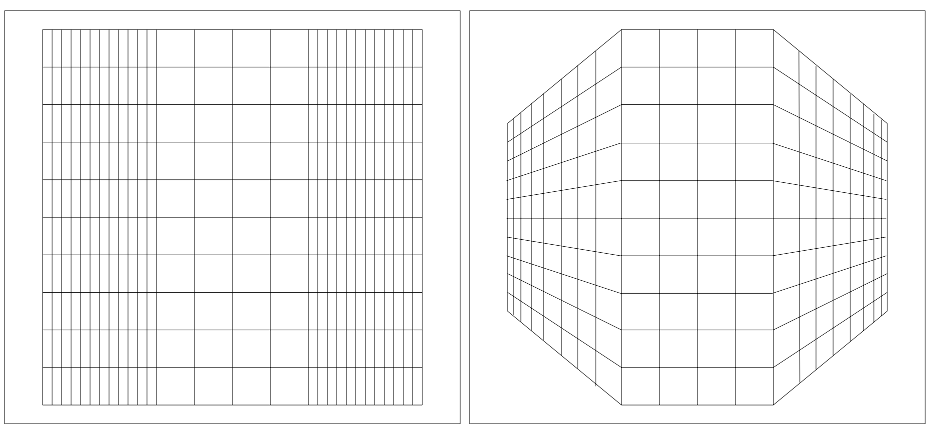

Figure 7.7: 1D+ VReps: the presentation patterns for Bifocal Display [147] and Perspective Wall [99]

tation patterns for Bifocal Display and Perspective Wall. Bifocal Display compresses the context on each side of the focal region, creating the presentation pattern shown in Figure 7.2(a). Perspective Wall uses an x plus z distortion as in Figure 7.4(b). Both pay attention to the predominately 1D characteristic of the VRep by creating a focal region that extends across the information strip, providing magnification for the full width.

7.3 Two-Dimensional Visual Representations

A *two-dimensional visual representation* (2D VRep) is a representation whose layout makes use of two positional variables. A 2D VRep has two predominant axes and thus lies on a plane. We have discussed using 3D distortions to vary the presentation possibilities for 2D VReps throughout this thesis. We now examine varying the dimensions used by the distortion.

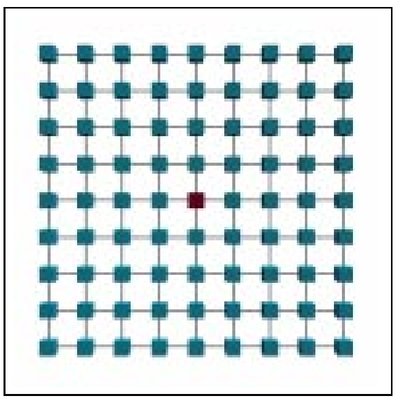

Figure 7.8: 2D VRep: a nine by nine grid graph in normal presentation

Figure 7.8 shows a 9×9 2D graph in normal presentation. Figure 7.9 illustrates the application of one-dimensional distortions to magnify the central node of the 9×9 graph. If the distortion is applied in x, all the nodes that have the same x location as the central focal node are elongated, creating a stretched column. Similarly, the distortion applied in y (Figure 7.9(b)) results in a stretched row.

Figure 7.10 illustrates the application of two-dimensional distortions to magnify the central node of a nine by nine grid graph. Figure 7.10(a) shows the stretched row and column that results from applying the distortions in Figures 7.9(a) and 7.9(b) together. A step magnification applied with a two-dimensional stretch is shown in Figure 7.10(b).

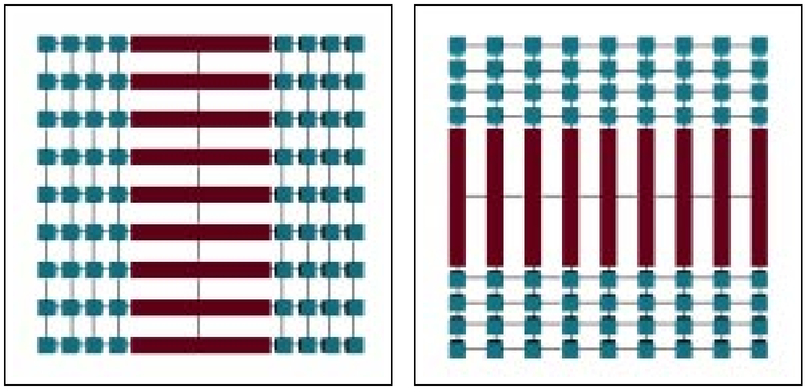

(a) 2D VRep: 1D distortion in x

(b) 2D VRep: 1D distortion in y

Figure 7.9: 2D VRep with one dimensional distortions. All the affected nodes are slightly darker

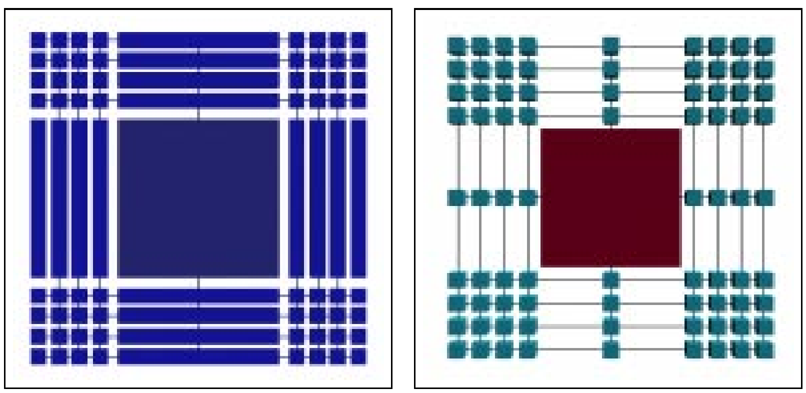

(a) 2D VRep: 2D distortion in x and y, with orthogonal stretch

(b) 2D VRep: 2D distortion in x and y, with step magnification

Figure 7.10: 2D VRep with two dimensional distortions

This type of uneven stretch in x and y is not a natural result of an EPS three-dimensional distortion. However, the same presentation pattern as in Figure 7.10(b) can be obtained with a three-dimensional step distortion. Chapters 3, 4 and 5 have extensively explored the variations in presentation patterns when applying three-dimensional distortions to 2D VRs. We now examine the effect of varying the dimensionality of the distortion within EPS. The surface containing the 2D VRep is still manipulated in a three-dimensional presentation space, however, the distortion may be applied in less than three dimensions. The resulting manipulated surface will still have a 3D form

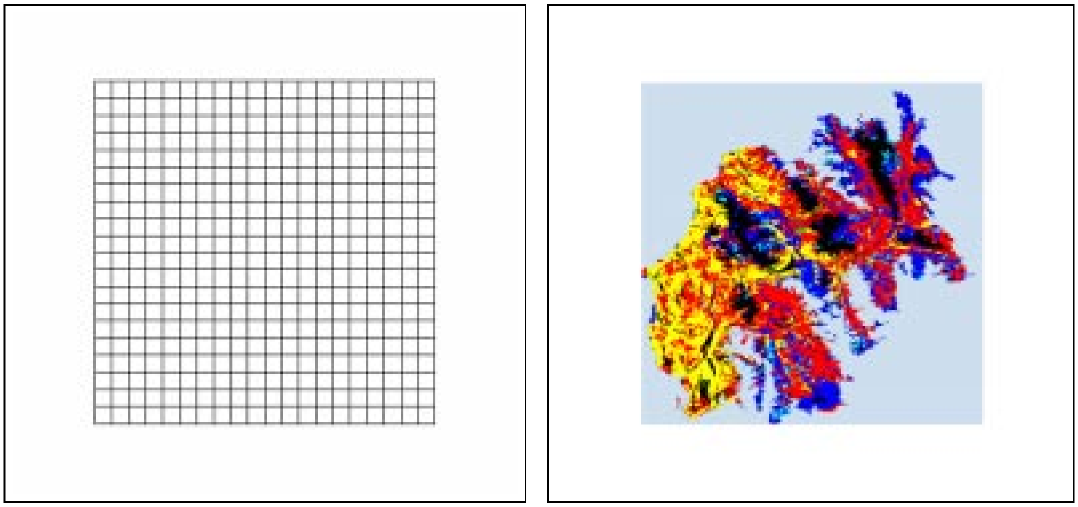

(a) 2D VR: grid in normal presentation

(b) 2D VR: a map showing dominant forest Species of part fo Rocky Mountain Trench in British Columbia in normal presentation (for map credit see Appendix C.5)

Figure 7.11: 2D VR: normal presentations

but choosing in which dimensions the distortions will be performed provides considerable freedom. The discussion will be illustrated with both a grid (Figure 7.11(a)) and an image (Figure 7.11(b)). The image is a thermal map of a volcano.

Figure 7.12 shows a three-dimensional distortion with one central focus. This focus has a small central region. Note that the four central squares of the grid (Figure 7.12(a)) are magnified but not distorted. Figure 7.12(b) uses the same distortion as Figure 7.12(a). The Gaussian drop-off function that creates this lens is calculated from a measure of distance. We call a lens three-dimensional if it is either calculated on the basis of or performs translations in three-dimensions. The lens in Figure 7.12 is three-dimensional because the distance it uses to calculate the z translation is based on both x and y. If (x_f, y_f) are the coordinates of the focal point and (x_1, y_1) are the coordinates of a point p, the distance d_p from p to f_c is:

$$d_p = \sqrt{(x_1 - x_f)^2 + (y_1 - y_f)^2}$$

If d_p is based solely on y, $(d_p = abs(y_1 - y_f))$, then any portions of the VR that have the same y coordinate as those in the focal region will be translated in z as if they were in the focal region. The result is that all portions with the same y coordinate will have the same magnification as the focal

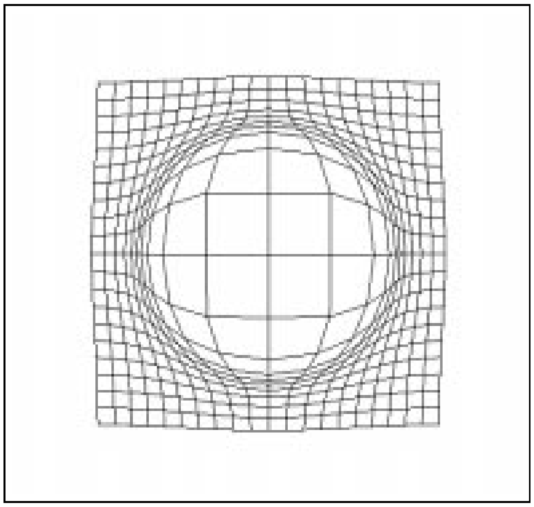

(a) 2D VR: an x, y and z distortion with the grid to show the pattern

(b) 2D VR: an x, y and z distortion

Figure 7.12: 2D VR: This lens with a small scaled-only focal region (center four squares in the grid) this used as the basis for the next series of images (Figures 7.12 to 7.19

region. Note that the region of scaled-only magnification in the centre of the lens in Figure 7.12 extends the width (Figures 7.13) of the magnified region creating the visual effect of a *scroll*.

Since the distortion is only affected by the distance in y, this is a y, z distortion. Figures 7.13(a) and 7.13(b) shown the same lens as in Figure 7.12 with no distortion in x. The same is true for x. If the distortion is based solely on x, that is $(d_p = abs(x_1 - x_f))$, then the magnified strip or scroll will extend from top to bottom of the image (Figures 7.14(a) and 7.14(b)).

Figures 7.13 and 7.14 show the effect of a y, z and x, z distortion, respectively. As the distortion has been applied to a 2D surface the resulting surfaces have three-dimensional form. However, both of these are 2D distortions. Note that the proportions of the squares of the grid are preserved along the length of the focal regions in Figures 7.14(a) and 7.13(a). This is due to the fact that magnification is a function of distance from the reference viewpoint (or height in z) and this height is the same for the entire length of the scroll.

Note that if the surface is kept at the same distance from the reference viewpoint to maintain the same degree of magnification, it is not possible to see the entire magnified region, as part of it is beyond the edge of the frame (Figures 7.13(a), 7.13(b) 7.14(a), and 7.14(b)). This is because for one of the dimensions of the VR, x in Figures 7.14(a) and 7.14(b) and y in Figures 7.13(a) and 7.13(b),

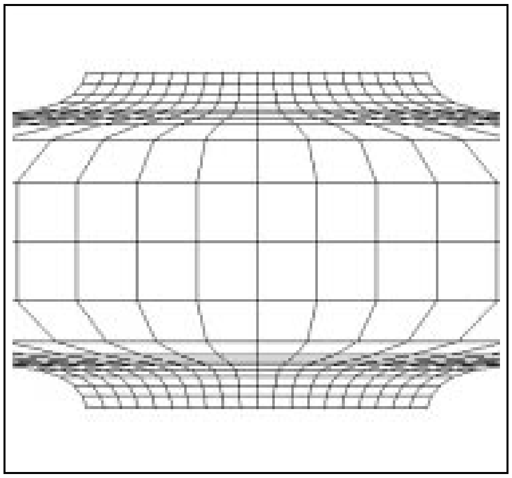

(a) 2D VRep: y and z distortion, with the same magnification in the focal region of the scroll as in the lens in Figure 7.12

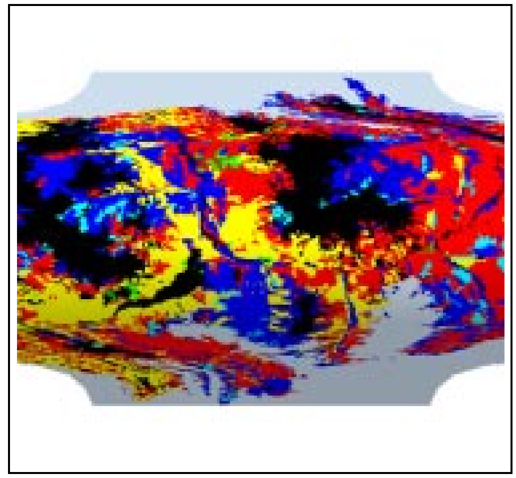

(b) 2D VRep: y and z distortion as in Figure 7.13(a) with the image

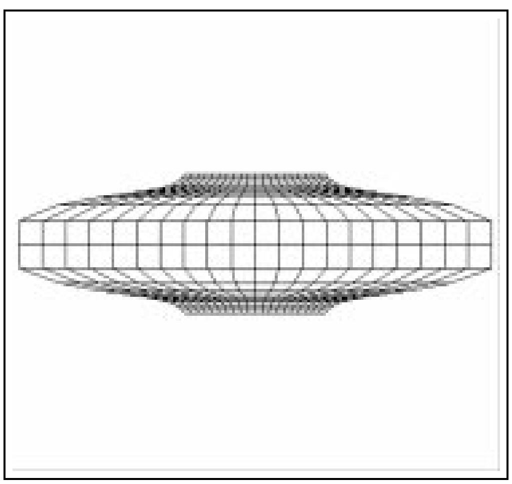

(c) 2D VRep: the same y and z distortion as in Figures 7.13(a) but translated back away from the reference viewpoint sufficiently to see both ends of the scroll

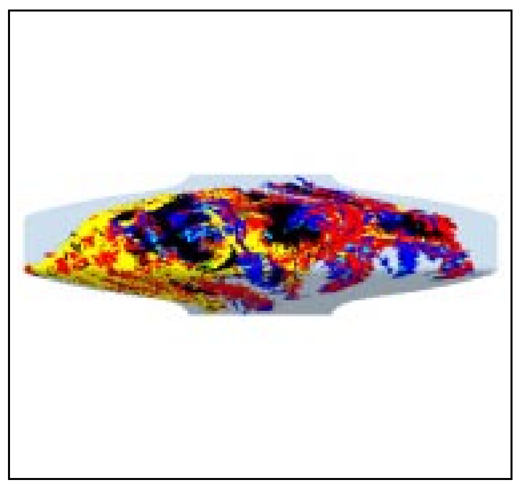

(d) 2D VRep: as in Figure 7.13(c) with the image

Figure 7.13: 2D VRep: y and z distortion

no distortion is used to maintain context. Therefore the edges of the image move out of the viewing frustum. To fit the entire magnified region into the frame the whole manipulated surface can be

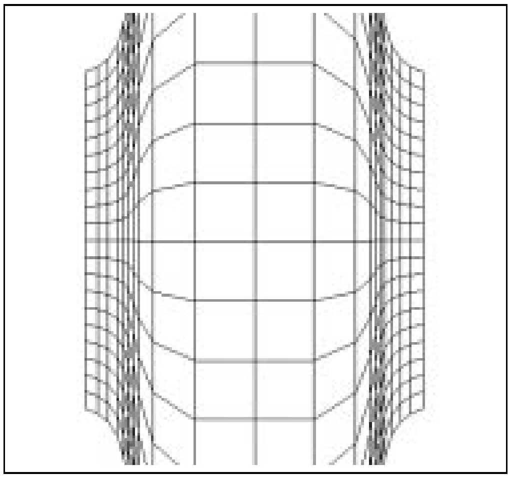

(a) 2D VRep: x and z distortion, with the same magnification in the focal region of the scroll as in the lens in Figure 7.12

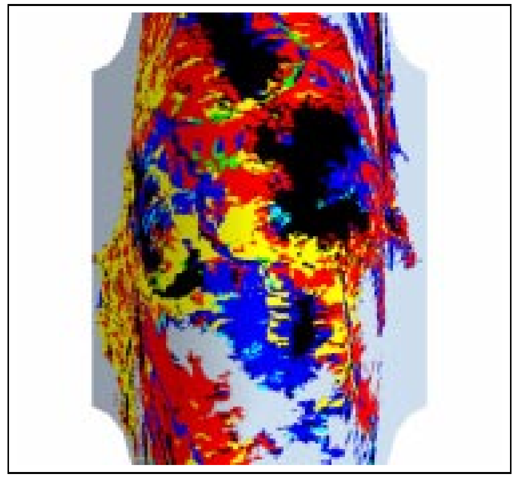

(b) 2D VRep: x and z distortion as in Figure 7.14(a)

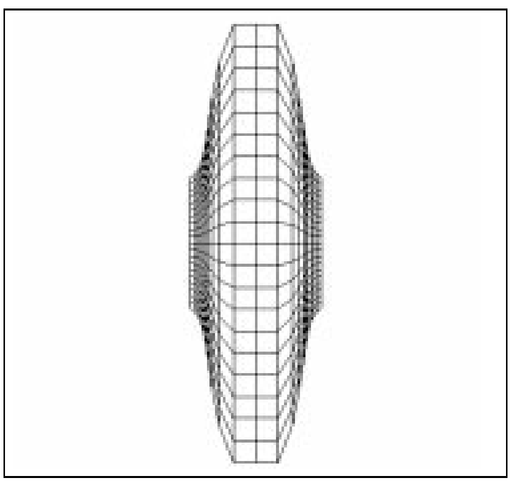

(c) 2D VRep: the same x and z distortion as in Figure 7.14(a) but translated back away from the reference viewpoint sufficiently to see both ends of the scroll

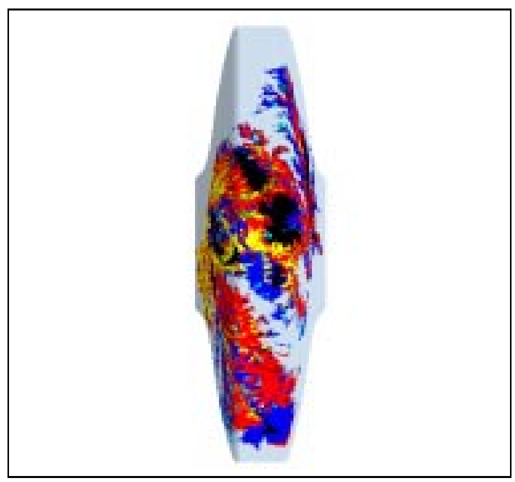

(d) 2D VRep: as in Figure 7.14(c) with the image applied

Figure 7.14: 2D VRep: x and z distortion

moved away from the reference viewpoint in z (Figures 7.14(c) and 7.14(d) and Figures 7.13(c) and 7.13(d)).

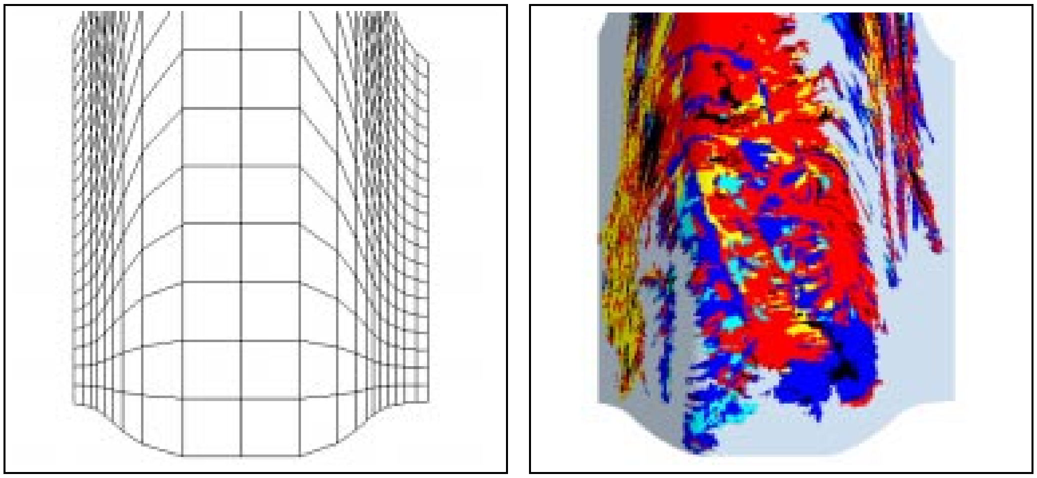

(a) (b)

Figure 7.15: 2D VRep: moving the focal centre in order to bring the lower end of the scroll into view

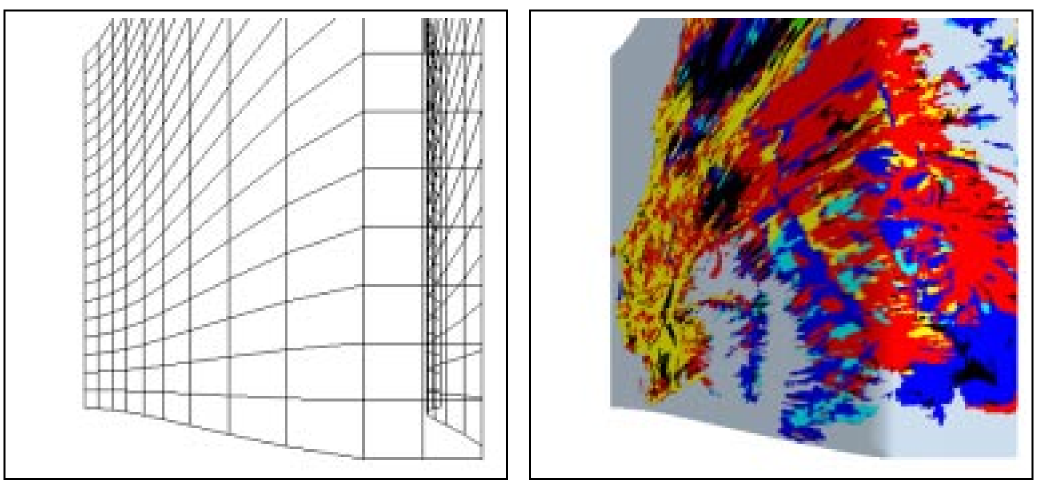

(a) (b)

Figure 7.16: 2D VRep: Folding the scroll towards the right shows a clear view of the left hand context

Moving the entire surface away and losing the corresponding magnification is not a satisfying solution. However, this scroll is still a lens as described in Chapter 3, which means that it has a

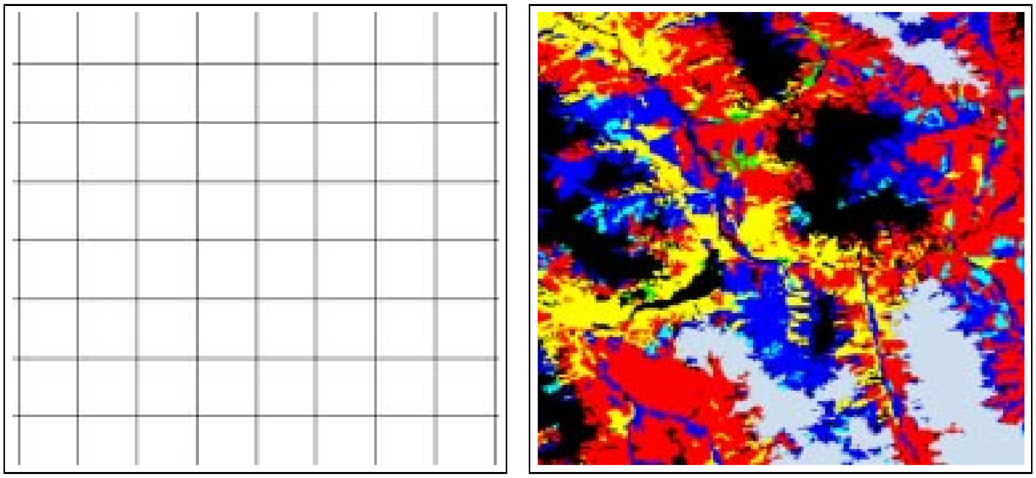

(a) (b)

Figure 7.17: 2D VRep: a z only distortion

viewer-aligned focal centre. The centre of a viewer-aligned focus is translated towards the reference viewpoint. One of the effects of a viewer-aligned focus is that its centre will always be in the same relative x, y location to the edge of the image as it was in normal presentation. The centre of the lenses in Figures 7.13 and 7.14 is in the centre of the image in normal presentation. Moving the centre of the lens towards either end of the scroll will bring that end of the scroll into view. Figure 7.15 shows the effect of moving the focal centre towards the bottom of the image: the bottom of the scroll comes into view.

Just as a roving search is possible with these scrolls, so is folding. In Figure 7.16 the same scroll as in Figure 7.15 is folded to temporarily improve the view of the context on one side of the scroll. In Figure 7.16(a) one can see the fold, as no surface is applied to the grid. In Figure 7.15(b) because of the opaque surface that holds the image one can see just the focal region and the extended side on the left.

The regions of scaled-only magnification for the scrolls extends the width in Figure 7.14 or the length in Figure 7.13 of the full image. Therefore, if both x and y are not used to calculate the distance, that $(d_p = 0)$ always, the whole image will be treated as if it was in the focal region. Figure 7.17 shows the same lens as used in all the figures in this series with distances calculated from neither x nor y. This is a both a lens and a full zoom. Because it is a lens, moving the mouse as in a roving search will adjust the position of the image due to viewer-alignment, moving the mouse as in folding will also re-position the image. Also, it has full, precise and semi-infinite magnification and because this is a lens that has merely been extended in both x and y, the context is interactively retrievable.

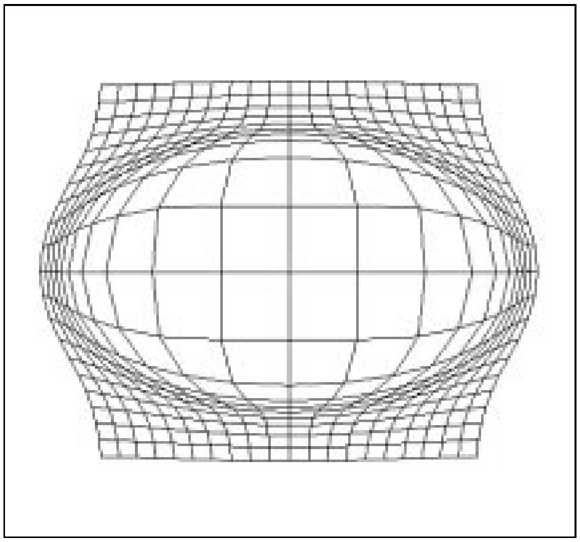

(a)

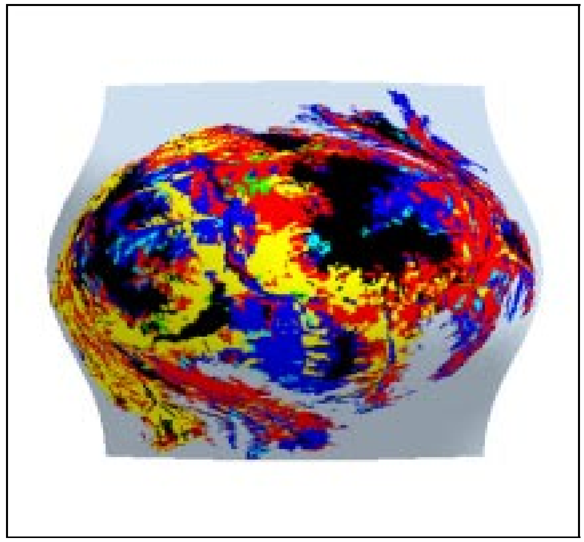

(b)

Figure 7.18: 2D VRep: a partial x, y and z distortion

To allow interactive control of these features, distances in both x and y are multiplied by independent controllable factors (x_{fac}, y_{fac}) which range between zero and one.

$$d_p = \sqrt{(x_{fac}(x_1 - x_f))^2 + (y_{fac}(y_1 - y_f))^2}$$

Figures 7.18 shows a y and z distortion with partial x distortion (x_{fac} of 0.6) and Figure 7.19 show a partial y distortion (y_{fac} of 0.6). The ability to interactively select whether to use distortion in x or y or either or both, partially allows for rapid change from a full zoom environment to a detail-in-context environment.

Keeping a recognizable detail-in-context presentation places limits on the amount of space that is available for magnification. There is a trade-off between the degree of compression and degree of magnification. With this type of control a lens can become a scroll and a full zoom and a lens again. One can locate a region of interest with a detail-in-context lens, change to full zoom to be able to use the full frame for magnification, and return to a lens to relocate oneself in the context again.

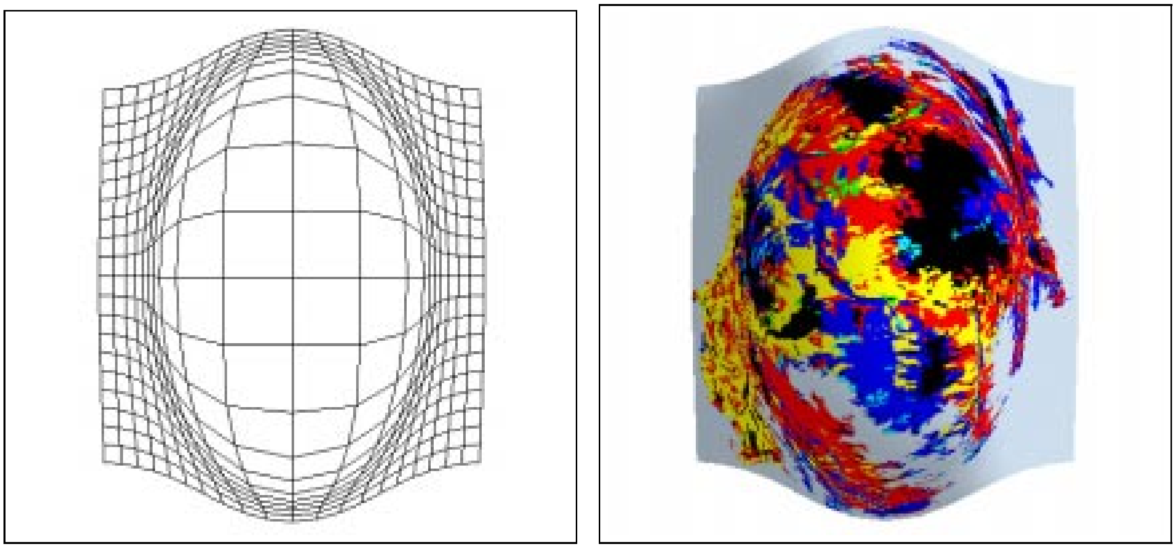

(a) (b)

Figure 7.19: 2D VRep: an x, partial y and z distortion

7.4 Two-Plus Dimensional Visual Representations

A *two-plus dimensional visual representation* (2D+ VRep) is a representation whose layout makes use of two positional variables and auxiliary use of another. A 2D+ VRep is still predominately two-dimensional. The use made of the third dimension varies from VReps that simply allow a three-dimensional interpretation to those that make partial use of the third spatial dimension.

Some types of images are minimumly 2D+ VReps. For instance, photographs still lie on a two-dimensional plane but they often portray three-dimensional information and include enough visual support for a three-dimensional interpretation. Perspective images are similar. A map can include topological information which implies 3D information but is still a 2D representation. Figure 7.20(a) shows a photograph of the surface of Mars. The image is actually 2D but the shape of the land forms can be read as 3D. Figure 7.20(b) is a distorted view of the same region of Mars. Note how as the ravines are magnified the 3D interpretation persists. It appears that one can see further into the ravines. We include these VReps in this category because the fact that they appear 3D has consequences for the interpretation of presentation (see Section 6.8).

A map can use three-dimensional relief to reveal land formations. Height fields in scientific visualization make use of the third dimension to represent some aspect of the information (Figure 7.21).

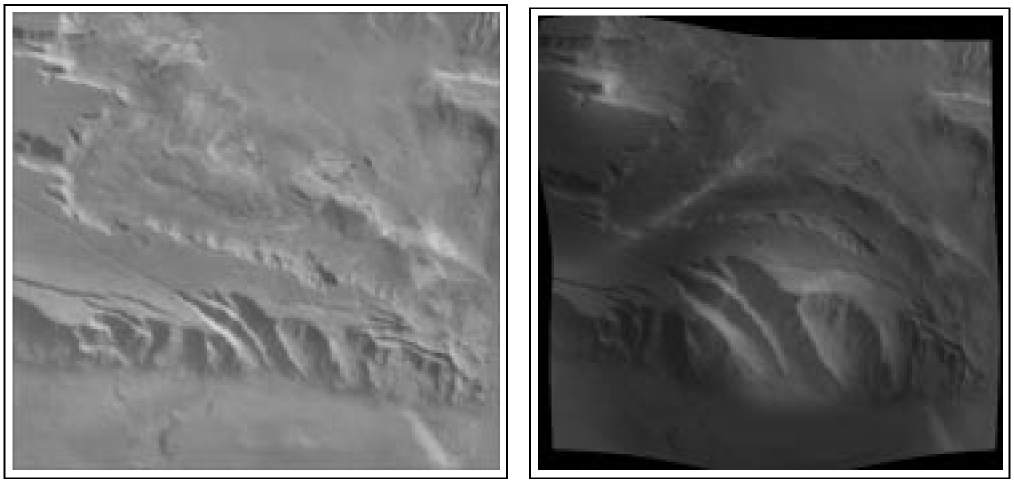

(a) The photograph of the surface of Mars in normal presentation

(b) The photograph of the surface of Mars with two focal regions

Figure 7.20: 2D+ VRep: The surface of Mars can be interpreted as 3D. The detail-in-context presentation does not interfere with this interpretation

These representations make use of the third dimension while still being predominately 2D. These 2D+ VReps form a narrow layer or blanket of 3D information in much the same manner that musical notation formed a narrow band or strip of 2D information.

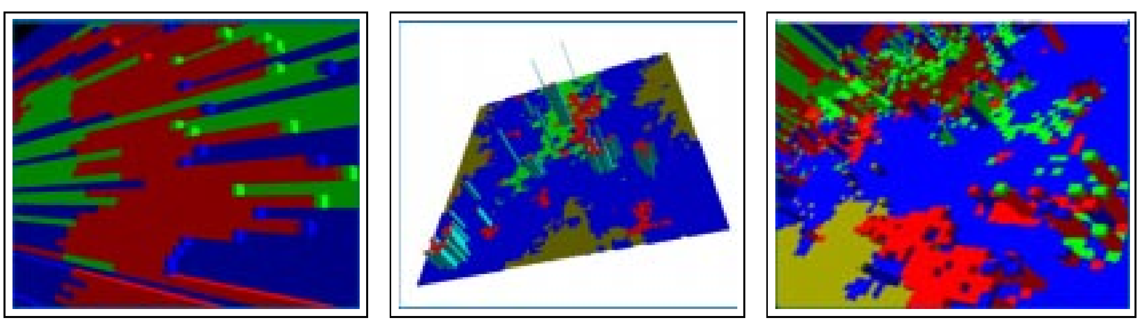

(a) 2D+ VRep: using the third dimension to emphasize changes in a 2D VRep

(b) 2D+ VRep: using the third dimension to express an additional variable

(c) 2D+ VRep: using the third dimension to progressive changes in an atribute

Figure 7.21: 2D+ VRep examples from the SEED project showing data from simulated landscape dynamics [27]

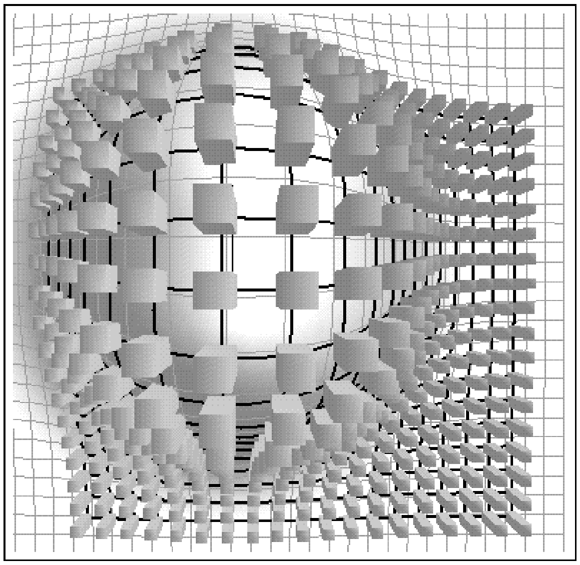

(a) 2D+ VRep: A single focus with the height field aligned perpendicularly to the base plane

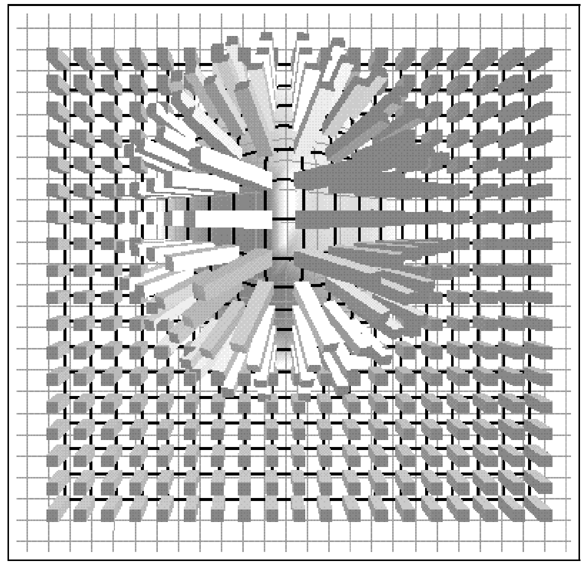

(b) 2D+ VRep: A single focus with the height field aligned with the surface normals

Figure 7.22: 2D+ VRep: Placing a height field on a pliable surface

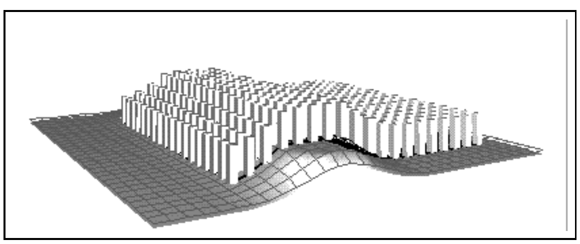

(a) 2D+ VRep: with perpendicular height fields the pillars do not collide

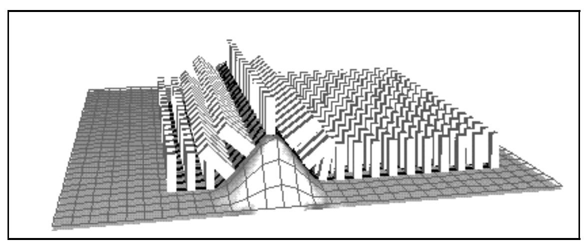

(b) 2D+ VRep: aligning the height fields with surface normals opens up the space between the pillars but causes collisions

Figure 7.23: 2D+ VRep: looking at the differences between perpendicular and surface normal aligned height fields from the side

Figures 7.22 to 7.25 show an initial exploration of the possibilities of using EPS concepts with 2D+ VReps. In particular, we look at applying surface manipulations to a simplified representation of a height field – a grid graph with uniform pillars as nodes.

Figure 7.22 shows the uniform regular height field with single focus. However, the height field in

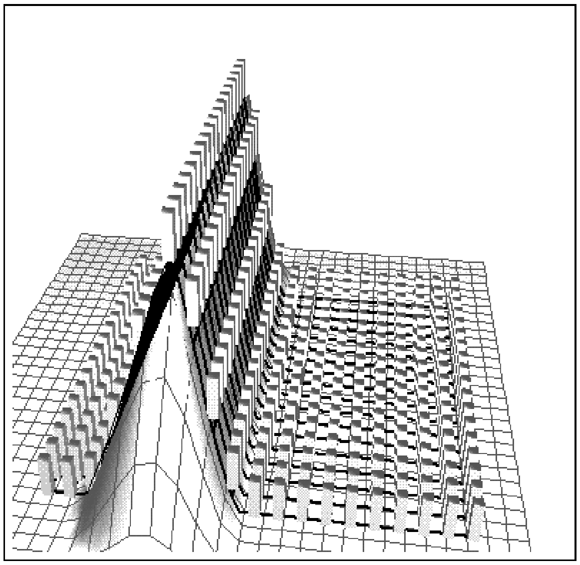

(a) 2D+ VRep: one row of height fields is raised

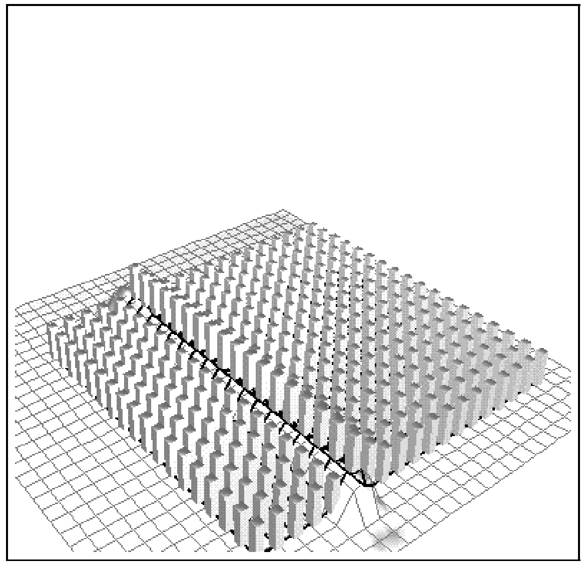

(b) 2D+ VRep: a bank of three rows of height fields are raised for possible comparison

Figure 7.24: 2D+ VRep: perpendicular height fields allow comparisons of aligned height fields

Figure 7.22(a) has been treated as if the centre of the base of each pillar was attached to the surface. Therefore when the surface is raised so is the pillar. Other than this the pillar is not adjusted. As a result the raised pillars are still perpendicular to the base plane. In the magnified region of the focus the pillars are also magnified and they are presented, as before the magnification, so that the most visible part of them is their top.

In contrast, the pillars in the focus in Figure 7.22(b) are treated as if they were fully attached to the surface. That is, the pillars are aligned to the surface normals. The effect is to spread the pillars across the curve the surface of the magnified focal region. Here it is the sides of the pillars in the focus that are most visible in the resulting presentation.

Figures 7.23(a) and 7.23(b) show side views of linear foci that extend across the graph, illustrating the different effects of aligning the pillars either by maintaining their perpendicular orientation (Figure 7.23(a)) or with the surface normals (Figure 7.23(b)).

In Figure 7.24(a) a single linear focus with steep sides has been used to raise an entire row of pillars without disconnecting them from their context. In Figure 7.24(b) similar linear focus is used to present three rows of pillars, connected to their context yet visually accessible.

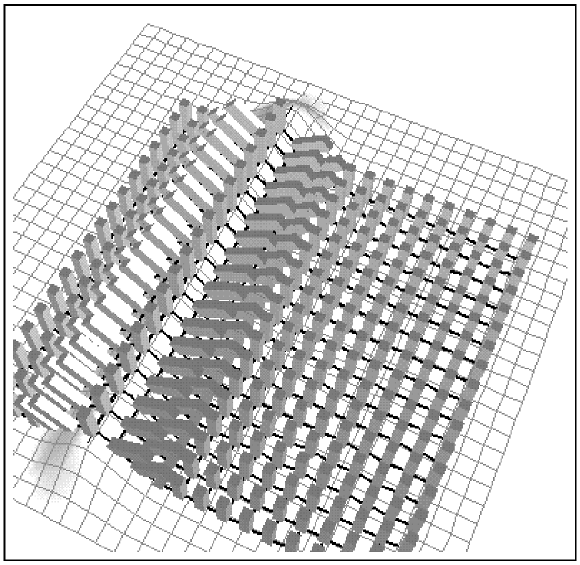

(a) 2D+ VRep: surface normal aligned height fields from the side

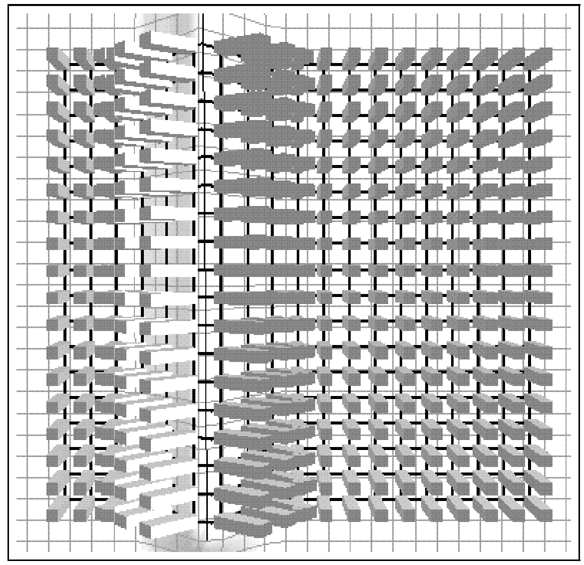

(b) 2D+ VRep: surface normal aligned height fields from above

Figure 7.25: 2D+ VRep: surface normal aligned height fields open like pages in a book

Figures 7.25(a) and 7.25(b) show two views of a single linear focus with steep side and the pillars aligned by surface normals. If the only pertinent information was represented by the height of the pillar these views might not be interesting. However, if information was represented along the pillars or on the surface to which the pillars are connected this type of presentation allows one to open the spaces between the pillars, and see their sides and the surface to which they are attached. If information is represented along the sides of the pillars, opening to see between them is useful. Also it allows visual access to the connection of the pillars to the surface.

7.5 Three-Dimensional Visual Representations

A *three-dimensional visual representation* (3D VRep) is a representation whose layout makes use of three positional variables. Frequently a 3D VRep is created for information that is perceived as three-dimensional, such as a building or the human body. More abstract concepts such as graphs can also be laid out in three dimensions, for example, cone trees [133], fractal trees [87] and Disc Trees [72]. There is considerable contention about the potential of 3D VReps for information presentation. Those in favour [171] of using 3D VReps claim that they offer the potential of an extra positional display variable, allow relational readings between three variables and may create some increase in usable display space. Those opposed [98] claim that 3D VReps add visual complexity with little or no gain, resolve into 2D when displayed on the computer screen, and place objects in front of one another causing occlusion.

The amount of usable space in a 3D graph layout has been examined by Ware et al. [172]. 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, Ware's results indicate that there is approximately a 3-fold increase in usable space in a 3D display (ie $3n^2$). While this result is encouraging because it seems to indicate that we gain something from our familiarity with 3D space, it also illustrates the fact that the 2D screen projection is a fundamental limitation. The amount of usable space relates much more closely a 2D rather than a full 3D space.

While 3D VReps may have increased visual complexity over 2D VReps, this could be due to the fact that as of yet there is little experience to draw from while creating 3D VReps. In contrast, over the centuries a lot has been learnt about how to make effective 2D VReps [160, 13]. It is to be hoped that 3D VReps will improve as the community gains familiarity with them and that this improvement will include both the ability to create understandable 3D VReps and to read them.

The fact that a computer's display can only show 3D VReps on a 2D screen adds an additional presentation problem, that of occlusion. When creating a presentation of a 3D VRep, access to its interior regions must be considered, along with the presentation problems of providing context, adequate magnification, the detail-in-context readings and the freedom to reposition (as discussed for 2D VReps). Currently the primary methods for accessing 3D VReps are either by adjusting the viewing angle (rotation) or the viewing position (navigation). In combination it would seem that these two allow all possible views. While the possibility of full rotation allows us to view the 3D display from all angles this does not eliminate the fact that it is inherent to working with 3D information that some data will be buried within a structure, whether it is a solid model or a complicated 3D graph layout. This rotational ability definitely improves the situation and was tied closely to the 3-fold increase in functional space noted by Ware [172]. However, there are many identified problems including: loss of context, loss of orientation, and the ever present problem of occlusion. Non-distortion approaches to providing access to the internal details of 3D structures make use of 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.

The fundamental problem remains, just as in the real world, that we cannot see through objects.

If something is placed between us and what we want to see it will block our view. For discrete information displays the occlusion problem can be exacerbated by information density. Noting the considerable body of work addressing the issues of information density in 2D VReps, we investigate possible application of these techniques to 3D VReps. We are particularly interested in whether distortion techniques can be applied to 3D in a manner that deals with occlusion as well as preserving context as in the 2D applications. We believe that the studies supporting the concept of integrated detail-in-context viewing in 2D displays [52, 70, 141] extrapolate to 3D. For instance, Furnas' [52] observations that humans store and recall information in great detail for areas of interest and gradually decreasing detail for the related context, presumably apply for information in general. Certainly the general cognitive support for the importance of integrated displays, to allow for utilization of our visual gestalt capabilities and minimize cognitive load, applies to our ability to assimilate and interpret information, not to a particular style of display.

Recent years have seen considerable research towards developing detail-in-context methods for 2D VReps, but relatively few have explored these issue for 3D VReps. Previous 3D detail-in-context approaches include Semnet [43] and Mitra [111]. Semnet presents three techniques. One creates an octree and uses this to display a 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. The second approach is based on density, and samples more fully around the focus and less as the distance from the focus increases. This approach increases the congestion and therefore causes more occlusion problems in the focal region. Semnet's third method is merely the implicit fisheye provided by perspective in a 3D display. A natural single focal point exists for the information that is in the foreground.

Mitra [111] employs a linear radial distortion approach with interactive filters for aircraft maintenance diagrams. These diagrams are 3D exploded views of aircraft assembly parts. An adjustable threshold is used to produce a filtered view based on the function of parts rather than proximity in the diagram. The user can adjust the threshold level creating views with more or less context. In this case both techniques of exploding and filtering the view create the space required to see into the structure. However, an unobstructed view is not ensured, progressive filtering removes much of the context, and the overall structure is not very apparent in the exploded view.

The next section examines the application of distortion techniques developed for 2D VReps to 3D VReps.

7.5.1 Applying Distortion to 3D VReps

We observe a naive application of four representative 2D distortions to 3D VReps. Each 2D (x and y) distortion that was intended for a 2D (x and y) VRep is applied as a 3D distortion (x, y and z) to a 3D VRep. As in their original 2D applications their dimensions of the distortion are aligned with the dimensions of the VRep. The observations made in this manner coupled with our framework lead directly to a 3D visual access method that deals effectively with occlusion and is capable of providing detail-in-context presentations for 3D VReps.

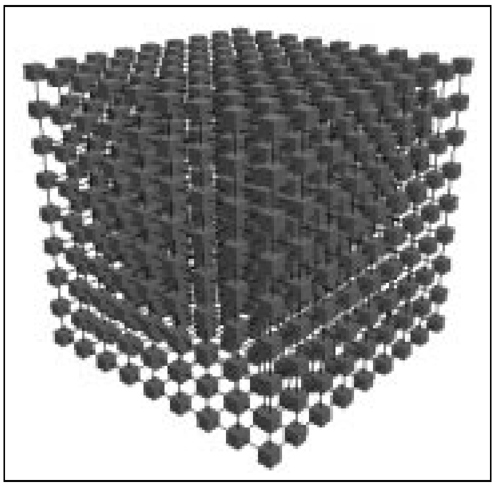

Figure 7.26: 3D VRep: the grid graph

To most clearly reveal the variations in presentation patterns we use a regular 9×9 2D grid graph (Figure 7.8) and a regular $9 \times 9 \times 9$ 3D grid graph (Figure 7.26).

Figures 7.27(a), 7.28(a), 7.29(a) and 7.30(a) show four representative distortions from 2D VRep methods that reveal the different possible presentations, providing a magnified focus for a 2D VRep. While this set of example techniques is not exhaustive, it is representative of the differing types of distortion.

The space-filling orthogonal approach Figure 7.27(a), similar to stretch tools [138], is formed by stretching all data that lies on either of the two axes centered at the focus and compressing the remaining areas uniformly. The resulting distorted image makes good use of available screen space but has entire rows and columns of distorted data.

The second example, Figure 7.28(a), and in the chart, Table 7.1 (column 2, row 1)) shows a nonlinear orthogonal approach. Here the focus is magnified to the requested amount and this magnification decreases according to some function (in this case sine) of the orthogonal distance

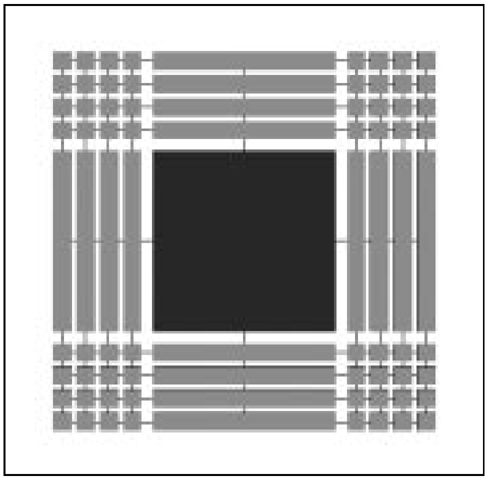

(a) Space-filling orthogonal distortion applied in x and y to a 2D grid

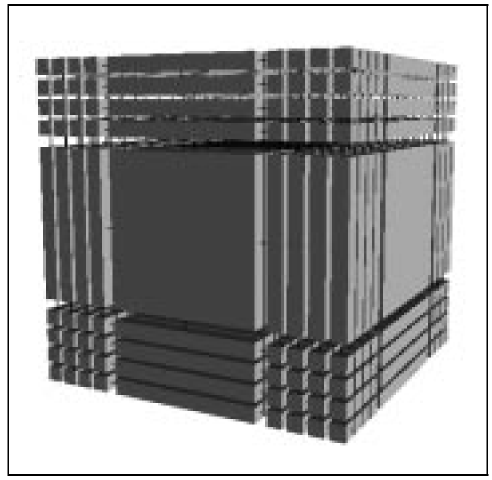

(b) Space-filling orthogonal distortion applied in x, y and z to a 3D grid

Figure 7.27: Space-filling orthogonal distortion applied to a 2D grid graph and a 3D grid graph

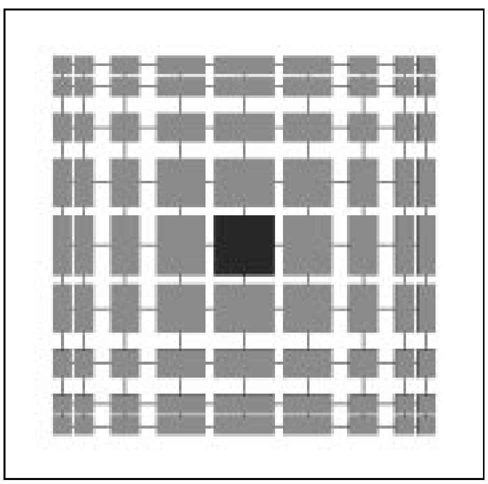

(a) Nonlinear distortion applied in x and y to a 2D grid

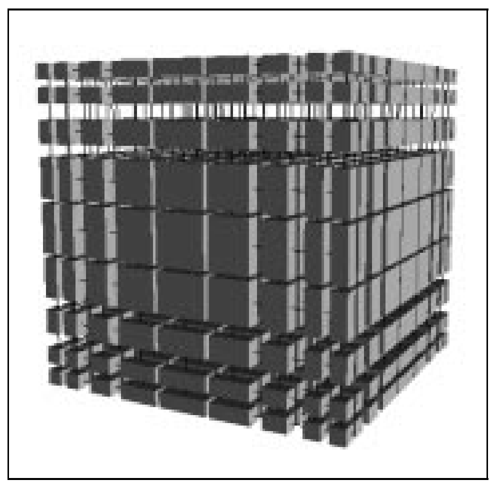

(b) Nonlinear distortion applied in x, y and z to a 3D grid

Figure 7.28: Nonlinear distortion applied to a 2D grid graph and a 3D grid graph

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.

A variety of mathematical functions have been used with this approach including arctan [82, 108], hemisphere [137], and hyperbola [90] (see Chapter 5). This type of technique supports the smooth integration of the focal area into its surrounding context.

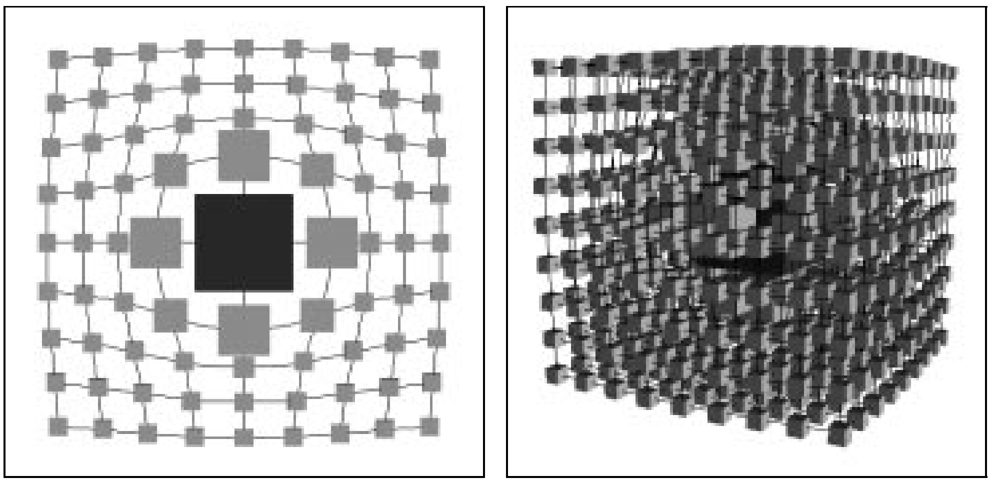

(a) Gaussian radial distortion applied in X and Y to a 2D grid

(b) Gaussian radial distortion applied in X, Y and Z to a 3D grid

Figure 7.29: Gaussian radial distortion applied to a 2D grid graph and a 3D grid graph

The third example, (Figure 7.29(a), and in the chart Table 7.1(column 3, row 1)) 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 non-linear function provides a similar effect of relative adjacent magnification as in 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 are shifted slightly. This interferes with the orthogonal relationships in the original grid. The principle of nonlinear radial distortion is common to several approaches [20, 90, 138].

The fourth example (Figure 7.30 and in the chart Table 7.1 (column 4, row 1) 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 and leaves more unused space. It also causes a marked grouping of the data that is not related to the information itself and could lead to misinterpretations. This technique has been used in the Zoom family of viewing techniques [7, 8] and more recently in Shrimp Views [150].

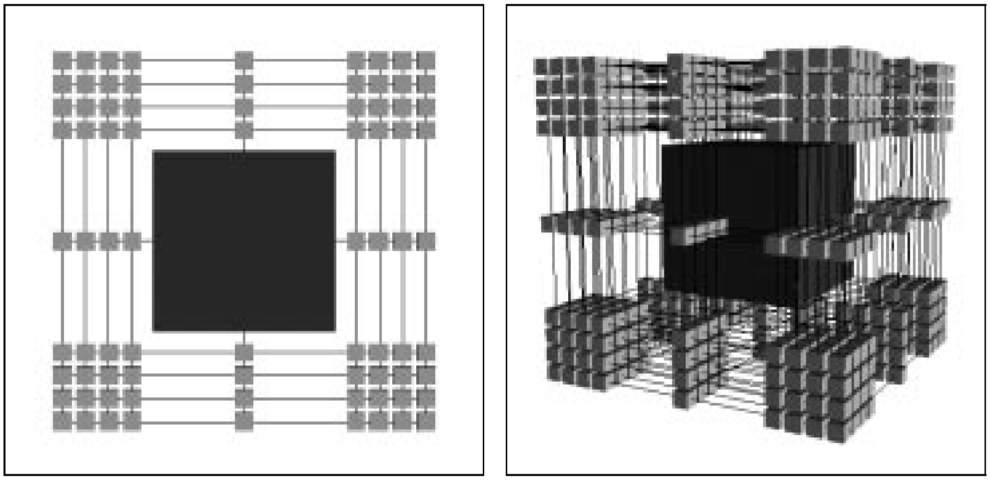

(a) Step orthogonal distortion applied in X and Y to a 2D grid

(b) Step orthogonal distortion applied in X, Y and Z to a 3D grid

Figure 7.30: Step orthogonal distortion applied to a 2D grid graph and a 3D grid graph

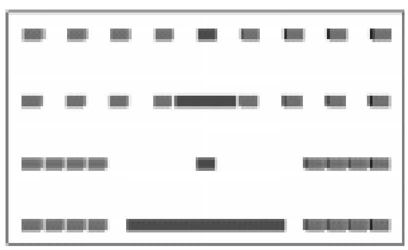

Figure 7.31: For a 1D VRep; top to bottom, normal presentation, magnification without displacement, displacement without magnification, magnification and displacement combined

Table 7.1 illustrates the application of 2D distortion techniques to 3D VReps. The top row displays the four 2D distortion techniques just discussed. Each of these techniques is used exclusively in the column it heads.

The second row shows a direct naive extrapolation of the 2D distortion methods to a 3D grid graph. Note that in the first two columns the distortion pattern propagates straight through to the surface. In fact, for focal points on the surface of a 3D structure, the same benefits that these distortion patterns obtain in 2D are achieved here. This would continue to apply to any chosen focus that was visible in the 2D projection. However, simply applying these approaches to a 3D VRep does more to obscure a central focal point than reveal it. In fact, the usual problem of some

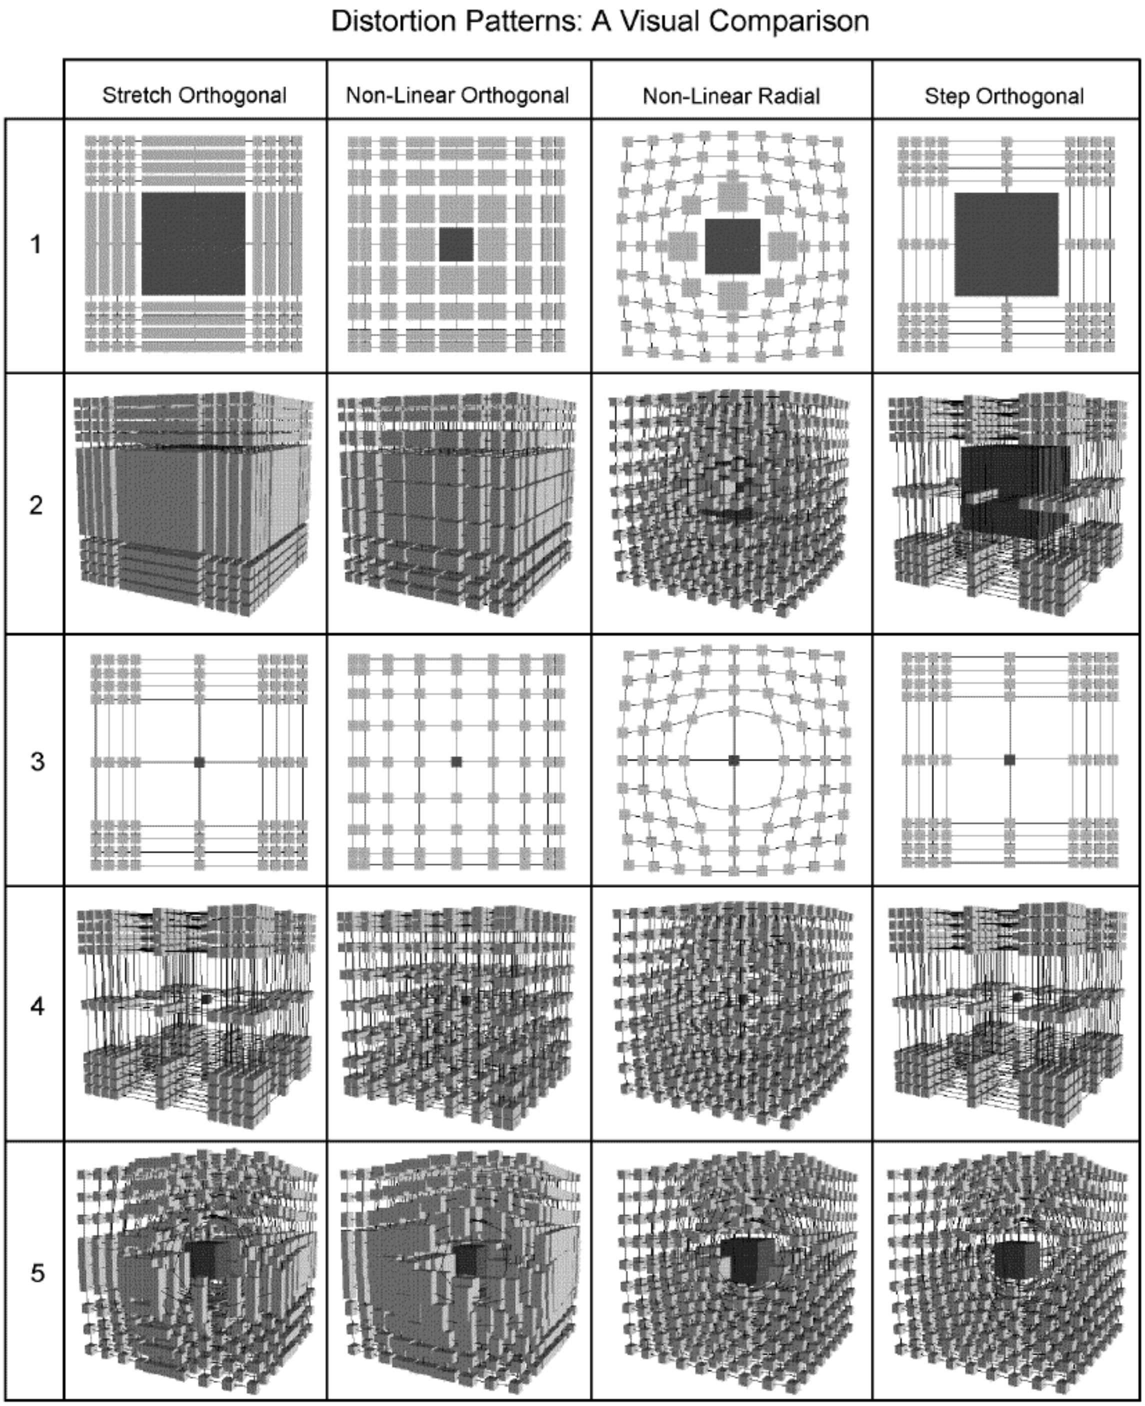

Table 7.1: This table illustrates the application of 2D distortion techniques to 3D VReps

objects occluding others in 3D layouts has been exacerbated in 2D distortion approaches with space-filling aspects, notably columns 1 and 2. The application of 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 just appears as increased congestion in the center. The amount of displacement at the edges of orthogonal step function (Table 7.1, row 2, column 4) does provide a view of the internal focal node. While this offers an initial clue that displacement by itself may be useful, the resulting view is not entirely satisfactory because it still does not allow viewing from all angles, and the artificial groupings are very pronounced. If distortion is to aid in full examination of internal aspects of 3D data, it is desirable that the user actually have unrestricted visual access to the chosen focus. Furthermore if we are endeavoring to provide context it would be preferable to not reorganize the data so radically .

From this naive application, we observe that the most effective method at providing visual access to the central focal node has a pattern that creates a lot of unused space. As we have discussed in Chapter 3.5 magnification and displacement are separate functions and can be applied separately. Leung and Apperly [94] discussed distortion viewing in terms of a transformation or displacement function with a derivative magnification function. As the 2D examples reveal, displacement provides the space that allows for the magnification. The usual application creating 2D detail-in-context presentations is achieved by applying both magnification and displacement distortions simultaneously. However, this does not have to be the case; these ideas are distinct and can be applied alone or in conjunction. Figure 7.31 shows at the top an undistorted one-dimensional array of cubes, second from the top is the same array with a central focus magnified only (no nodes are displaced), third the distortion is applied to displace nodes to either side of the focus (no magnification), and the last example makes use of both magnification and distortion.

Following the insight provided by the naive application, the third row (Table 7.1) presents the same set of functions revealing the displacement-only aspect on the 2D grid. Note that the stretch and step orthogonal approaches (Table 7.1, row 3, columns 1 and 4) resolve into the the same pattern.

For the four approaches we have been examining, the effects of applying displacement-only distortion to the 3D grid graph can be seen in Table 7.1, row 4. Despite the elimination of the obscuring magnification, there is little improvement with the application of graduated and radial techniques (Table 7.1, row 4, columns 2 and 3). Notice that without the magnification both orthogonal techniques have the same appearance (Table 7.1, row 4, columns 1 and 4). 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 which can still occlude the focus during rotation. The partial solution provided by the displacement-only patterns gives an indication that using distortion to remove occluding objects is potentially useful.

7.5.2 Observations

At this point we have identified that a displacement only function may be the most useful in providing visual access. However, it would appear that aligning this function with the data creates artificial groupings that appear to have significance. We have also noted that limiting the spread of the distortion provides a much more recognizable exterior, and that the objects that concern us lie only between the focus and the viewer. On the other hand it would appear that the magnification is still more appropriately aligned with the data. The choice, for instance, between relative local magnification or focal-only magnification, is task and information dependent. These observations led to the application of two techniques first developed in our 2D distortion method 3DPS [20], namely viewer-aligned distortion and constrainable distortions.

In our framework, for 2D VReps, focal regions are aligned with the reference viewpoint to keep more than one in sight and prevent the focal regions from occluding each other. In this manner the 2D VRep presentations are organized primarily to be viewed from the reference viewpoint. From this vantage point the focal regions are perceived as if they are being brought closer to the viewer as they are magnified. With 2D VReps, we have not reorganized the presentation during rotation. Instead the focal regions stay aligned to the reference viewpoint, allowing the viewer to see the 3D form of the manipulated surface. For 3D VReps we will use the same idea of a line of sight from the chosen focus to the viewpoint and apply the displacement distortion radially along this line of sight. This will displace objects that obstruct the view. However, we do not use the idea of a reference viewpoint and allow the line of sight to be re-positioned during rotation. This provides interactive visual access to the chosen focus from all viewpoints. Furthermore, for 2D VReps the distortion is constrained to maintain as much undisturbed context as possible and allow the user interactive choice about the location and pattern of the compression. The application of the constrained distortion for 3D VReps directly parallels this.

7.5.3 Using Distortion to Provide Visual Access

We have extended traditional magnification and displacement distortion viewing to include a visual access distortion [36]. This is applied radially along the line of sight to displace data items in gradually decreasing amounts as the distance from the line of sight increases. This creates an opening directly in front of the viewpoint, giving visual access to the focus while minimally affecting the rest of the data. Figure 7.33 shows progressive opening of the 3D grid graph.

Visual access distortion [21, 36] is a viewer aligned, radially constrained, reversible distortion that clears the line of sight to chosen focal regions. We believe that for effective three-dimensional detail-in-context viewing it is important to:

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

Specifically, visual access distortion proceeds as follows. A focal point is selected, for example in Figure 7.32(a), the central point has been selected. Let L be a line segment extending from the focus to the viewpoint (the line of sight). We compute the vector $\vec{D}$ as the shortest vector from each object O in the display (in this case a grid node) and a nearest point P on the line L. These are illustrated in Figure 7.32(b) with yellow arrows from the focus and the line of sight to the adjacent points. Vector $\vec{D}$ defines the direction of the distortion at O and its length $|\vec{D}|$ is used to determine the magnitude M of the distortion. To achieve our goal of smooth integration back into the original data topology we use a Gaussian distribution to determine the magnitude of the displacement. Figure 7.32(c) shows with the profile of a Gaussian function how the $|\vec{D}|$ (indicated in yellow) is used to calculate the magnitude M of the distortion (indicated in green). For a given value of $|\vec{D}|$, we can determine the height of the Gaussian, which gives the magnitude M. Figure 7.32(d) shows (in green) the use of the magnitude M along direction $\vec{D}$ to create the displacement. The shape of the Gaussian function, and hence the distribution of the distortion, can be controlled simply by adjusting the height and standard deviation of the curve. Since our viewing direction is along the line of sight, the distortions will appear to the viewer as radially symmetrical about the focus, though it is moderated by the effect of planar perspective projection.

The result is a distortion of the original data that provides a clear visual path from the viewer to the focal node. The visibility of the focus will be maintained 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 smoothly to their original positions as they move away (Figure 7.33). The creation of a clear visual path is combined with step magnification.

Having described the basic idea behind visual access distortion we return to the comparisons

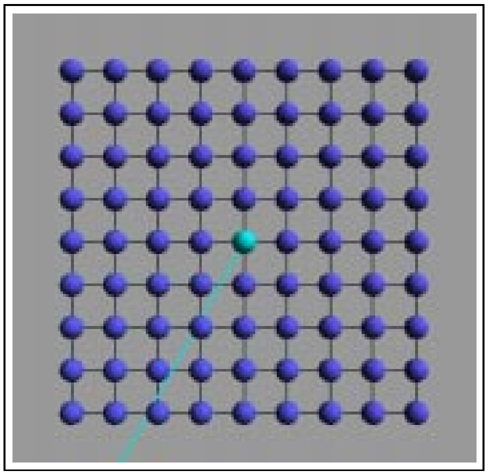

(a) The focal point is selected

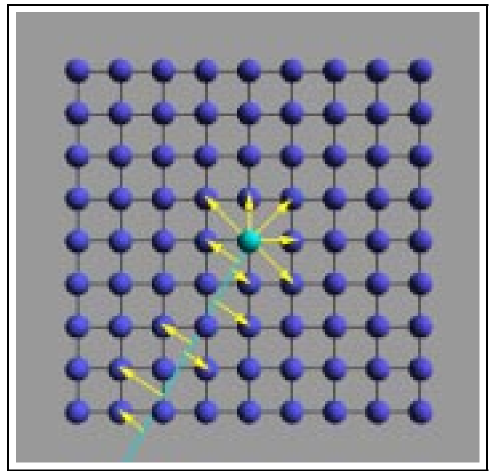

(b) The vectors to the nearest points on the line of sight

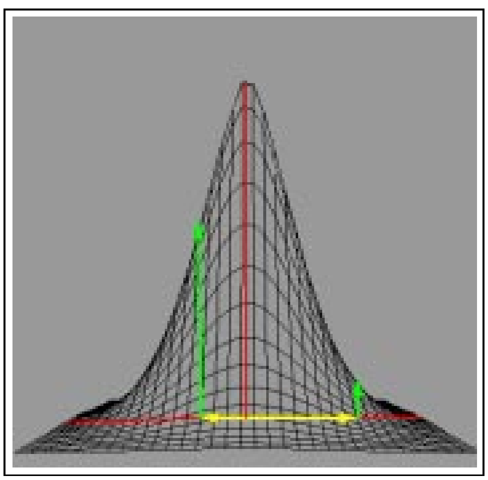

(c) The displacement is obtain for the drop-off function

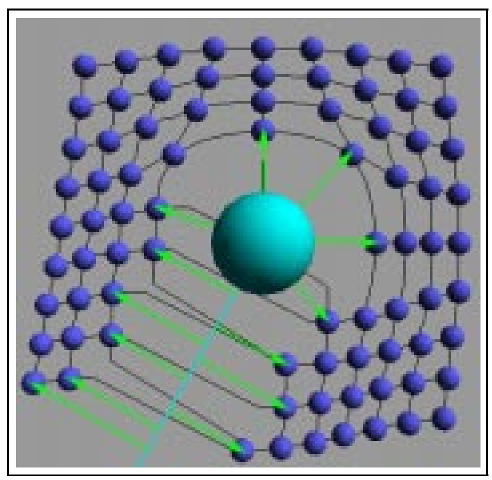

(d) The nodes are displaced from the line of sight

Figure 7.32: Visual access distortion

in Table 7.1. Applying a visual access distortion to any of the four examined 2D approaches successfully exposes the focus in context. This can be seen in Figure 7.34 where Figure 7.34(a) shows space-filling orthogonal magnification, Figure 7.34(b) shows non-linear orthogonal magnification, Figure 7.34(c) shows radial Gaussian magnification, and Figure 7.34(d) the step approach or in

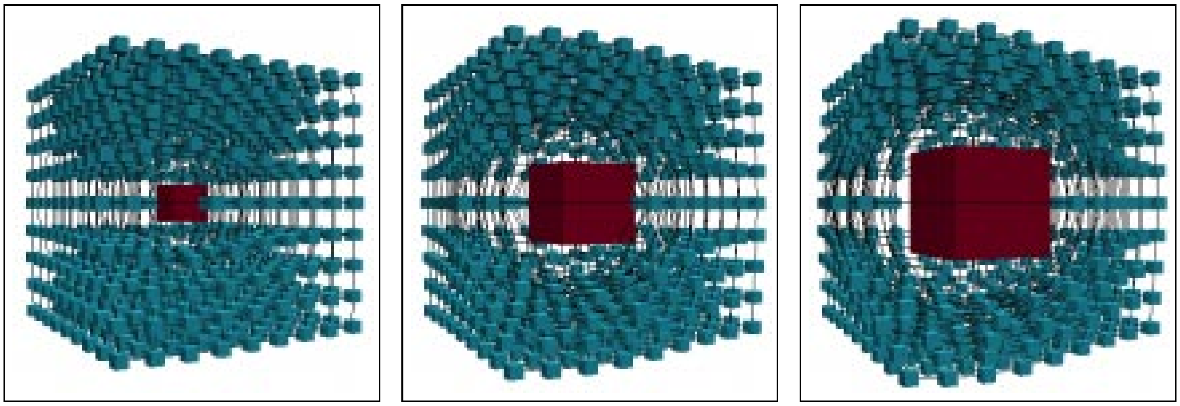

(a) (b) (c)

Figure 7.33: This series shows Gaussian visual access distortion applied progressively to the 3D grid

Table 7.1, row 5. In these images 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 viewer-aligned visual access distortion. Even in cases where the magnification has completely occluded the central focal node, applying the visual access distortion clears a line of sight to the focus. This is particularly evident in the images Figure 7.34(a) and 7.34(b), (Table 7.1 row 5, columns 1 and 2). Here the space filling orthogonal approach and the graduated sine function had completely occluded the central focus node (see Table 7.1, row 2, columns 1 and 2), virtually creating a solid. Similarly, with the radial Gaussian distortion (see Table 7.1 row 2, column 3), the central focal node is practically obscured by its neighbors as they are magnified as well, though to a lesser degree. In all these cases visual access distortion provides visibility of the central focus (Figure 7.34) (Table 7.1 row 5). In the case of the orthogonal step function (see Table 7.1 row 2, column 4), without the application of the displacement aspects of the distortion the artificial clusters are not generated (Figure 7.34(a) or Table 7.1 row 5, column 4). The actual focus is magnified, while the entire context is undisturbed. In this situation applying visual access distortion achieves the desired focal visibility in a situation where context is minimally disturbed.

Visual access distortion [36] can be applied by itself or in combination with another VRep aligned distortion. These distortions can use either or both magnification and displacement. In practice we find that most frequently we use visual access distortion by itself; i.e. displacement only. If the focal node requires magnification, we use a simple step function, only magnifying the

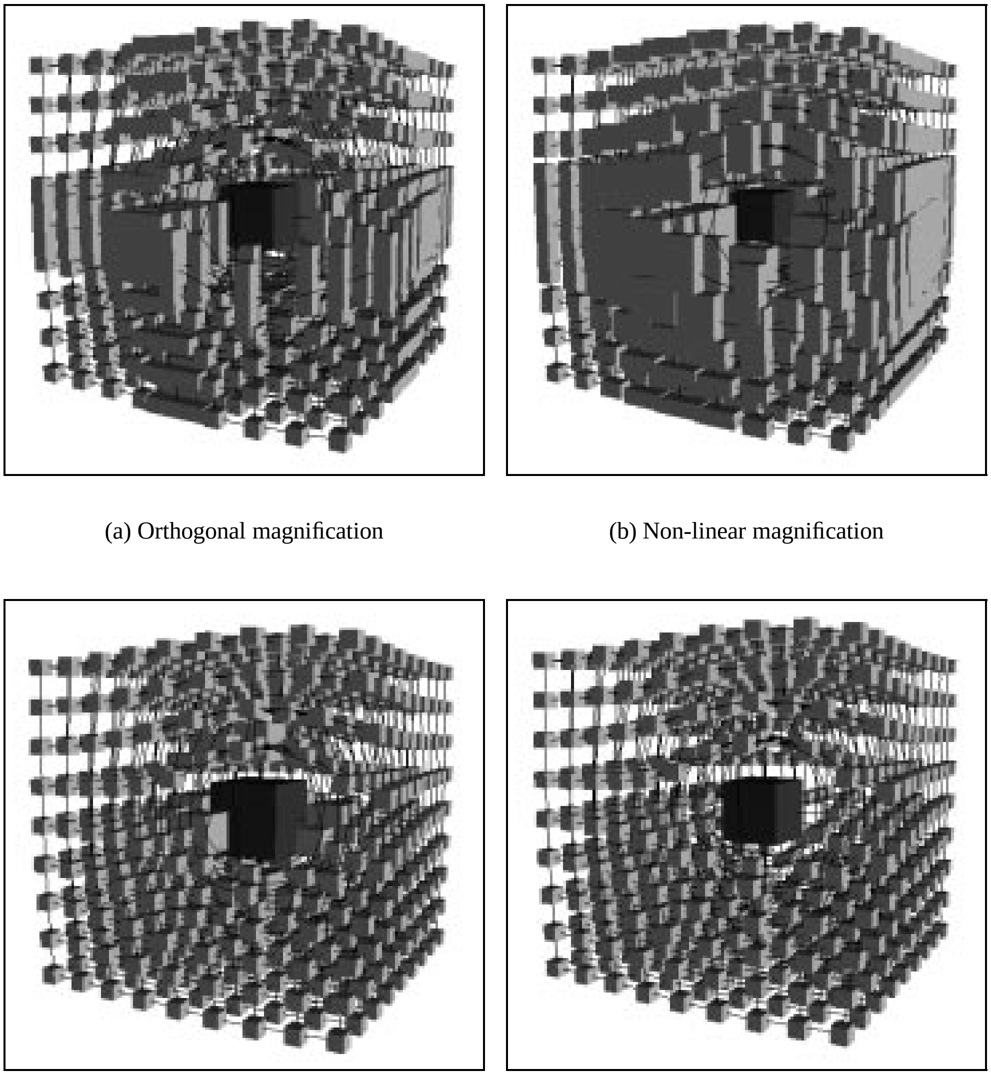

(a) Orthogonal magnification

(b) Non-linear magnification

(c) Gaussian magnification

(d) Step magnification

Figure 7.34: Applying visual access distortion to the 3D grid graph with different magnification patterns

chosen nodes. By itself it allows in-context browsing of a 3D display. With magnification it provides detail-in-context viewing.

7.5.4 Application to Arbitrary Graphs

While simple 3D grid graphs have been chosen to clearly reveal the patterns in the distortion techniques, the effectiveness of this approach is not limited to this type of 3D grid layout. The Figure 7.35 shows a polar graph layout that positions nodes by randomizing the magnitude of both the radius and angles with visual access displacement and focal magnification applied for a single node. In Figure 7.36 increasing the displacement and magnification bring edges into conflict. In Figure 7.37 visual access distortion has also been applied along the length of the edges, curving them away from the line of sight, leaving a clear view of the focus.

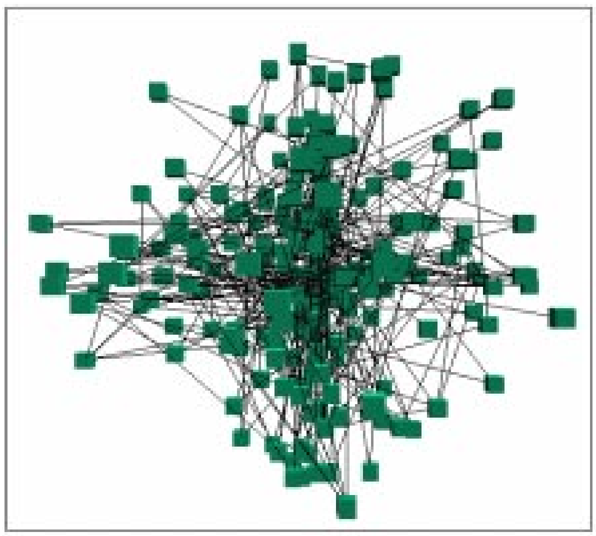

Figure 7.35: Undistorted random polar graph layout

7.5.5 Providing Multiple Foci

Visual access distortion scales well to multiple focal points. Because each line of sight utilizes its own access distortion function, it is possible to combine more than one focus in one view. For each object in the volume the distortion is now a function of each line of sight. The contributions from each focus are combined [20]. The displacement of any object O relative to a set of n line segments L_i, where each segment L_i begins at an Focus F_i and terminates at the viewpoint V, may be summarized as follows:

$$D = \frac{\sum_{i=0}^{n-1}(O - N_i) \times g(|O - N_i|)}{n}$$

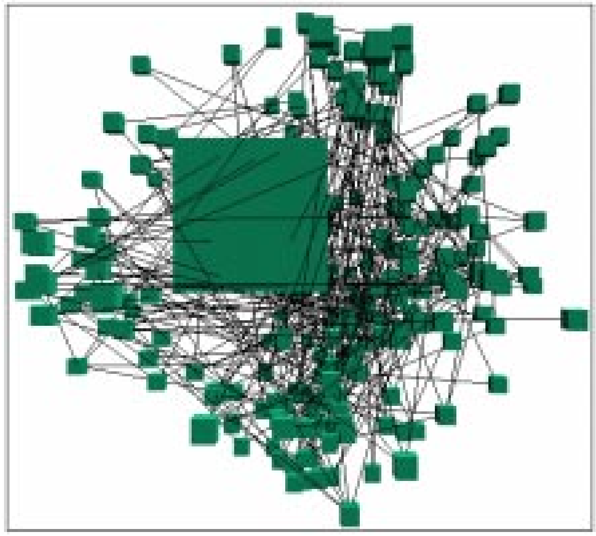

Figure 7.36: Increased magnification brings crossing edges into conflict

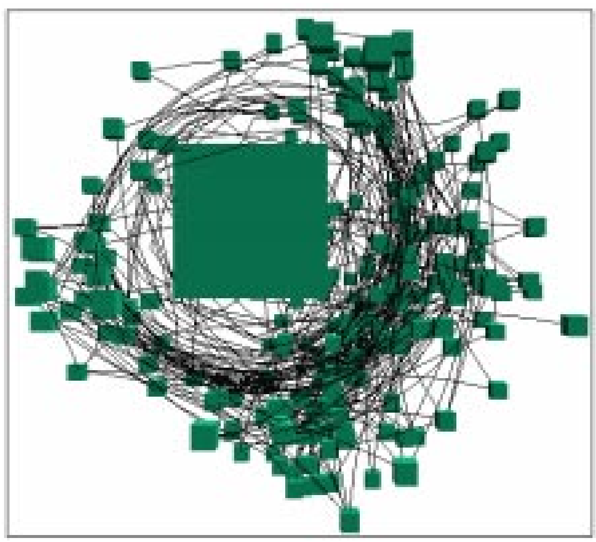

Figure 7.37: Visual access distortion also applied to edges

where N_i is the point on the i^{th} line segment nearest the object O, and $g()$ is the function that returns the height of a Gaussian given the distance from the object O to the line L_i. The resulting displacement is the average of several independent distortions, each in a radial direction away from the line segment L_i. Thus it is possible to clear lines of sight to several objects simultaneously. In Figure 7.38 a simple average at each point produces clear lines of sight to the two foci. The upper

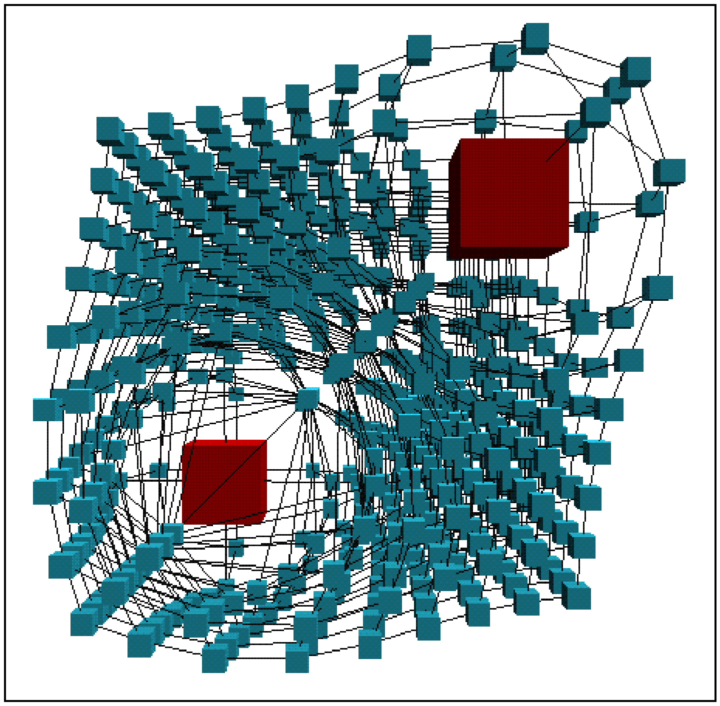

Figure 7.38: The visual access function may be applied simultaneously to more than one focus

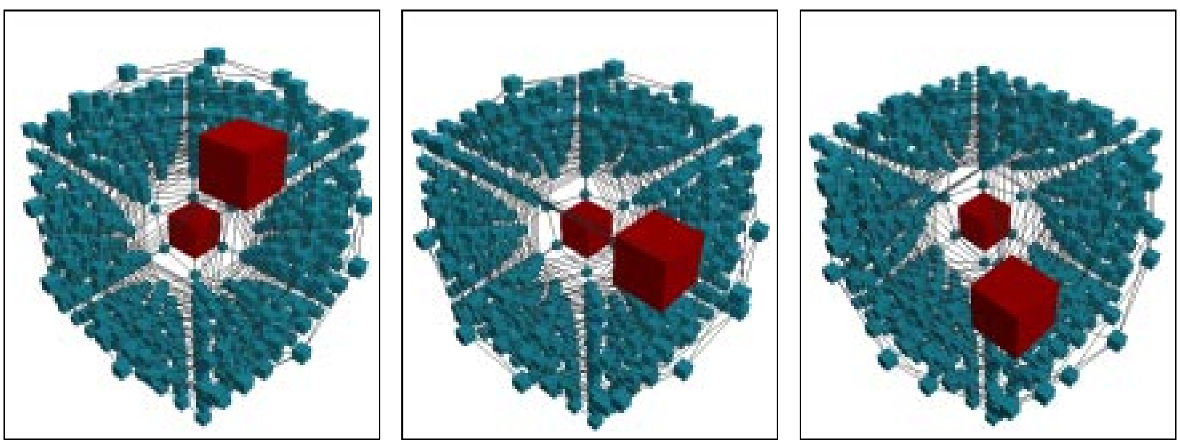

(a) (b) (c)

Figure 7.39: Even when the two foci are in line with the viewpoint both foci remain visible because the distortion function from the furthest focus effects all occluding objects including other foci

right focus is one layer deep into the $9 \times 9 \times 9$ cube. The lower left focus is 8 layers deep, but still visible.

As with a single focus visual access is maintained during rotation. Figure 7.39(a) shows the lattice positioned with both foci in line one above the other. Figures 7.39(b) and 7.39(c) illustrate rotating the lattice from top to bottom. The closer focal point is rotated to come in between the viewer and the further focal point. However, visual access distortion merely considers the closer focal point as another occluding object and shifts it to one side. The series of three images in Figure 7.39 shows how close one focal point can come to occluding another as it crosses through the line of sight.

In browsing a 3D VRep the user can select focal type and position, as well as which the 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 very effective visual feedback about the context and relative positions of the individual data items.

7.5.6 Using Distortion for Visual Search

Visual access distortion [36] is distinct from the other distortions discussed in that it is not aligned with any of the dimensions of the VRep. This separation from the VRep allows it to be used as a dynamic probe with which the inner regions of a 3D VRep can be searched.

Note that for the images in Figure 7.34 or in Table 7.1, row 5, two distinct transformation functions are applied at the same time, with respect to the same focus. One is the 3D x, y and z distortion that is aligned with the x, y and z of the grid graph. The other is the viewer-aligned visual access distortion. The *central-line-segment* of the visual access distortion is aligned with the line-of-sight which is independent of the orientation of the VRep. This independence can be taken further. As discussed thus far, the central-line-segment's end points are the chosen focus and the viewpoint. The focal endpoint of the central-line-segment can be either a data object or a location in space. When it is a data object, the visual access distortion is applied from the viewpoint to the object's center. Browsing can involve sequential selection of objects or nodes. Alternatively, the focal endpoint of the distortion's central-line-segment can be a location in space. Providing the user with interactive control of this line-segment-of-sight creates a dynamic probe, fluidly movable through the space.

The next series of images show use of visual access distortion as a search probe. The information in these images is spatio-temporal data generated by SELES [44], a spatially explicit landscape event

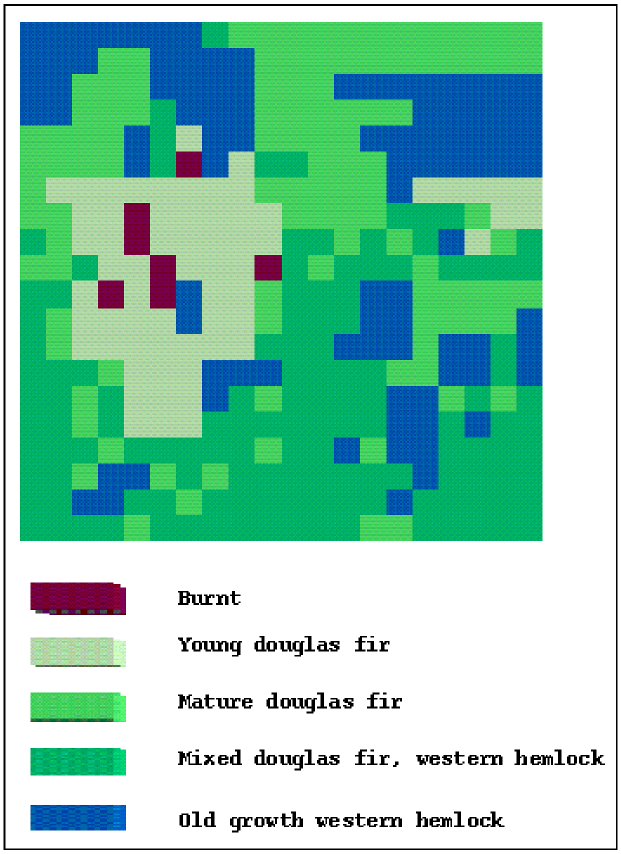

Figure 7.40: Land coverage image and legend

simulator. We are using this data to illustrate the use of a visual access probe because it has a simple spatial organization (still a 3D lattice of cube cells) and variations in data can be seen through changes in colour. The following is a very brief description of the data (for further explanation see [27, 44]).

In SELES landscape structures, each grid cell indicates the dominant aspect of vegetation cover of the area it represents. The current information in Figure 7.40 includes British Columbia forest cover in a system that has one disturbance type, namely fire. The type of vegetation cover includes young douglas fir, mature douglas fir, mixed douglas fir, western hemlock, old growth western hemlock and burned. Figure 7.41 shows the spatio-temporal cube of landscape data as it changes over time.

Figure 7.42 uses a visual access probe to look through the information space along the temporal axis. This reveals separate columns of contiguous temporal data. Increasing the displacement of the

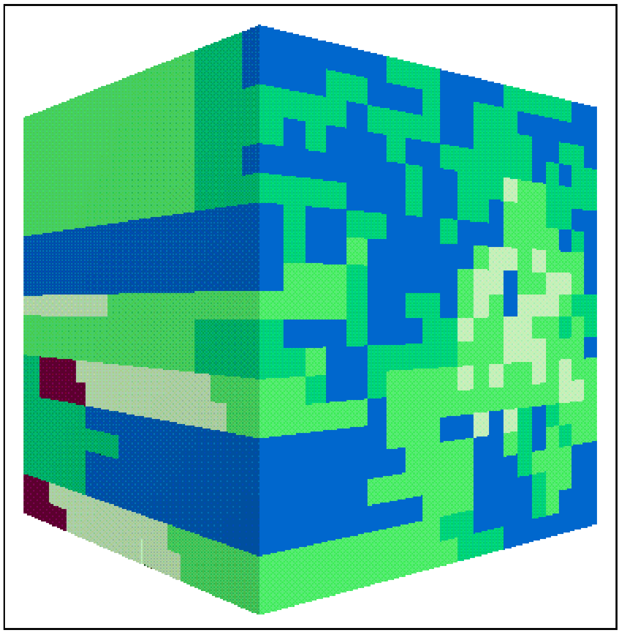

Figure 7.41: The temporal block of landscape coverage information

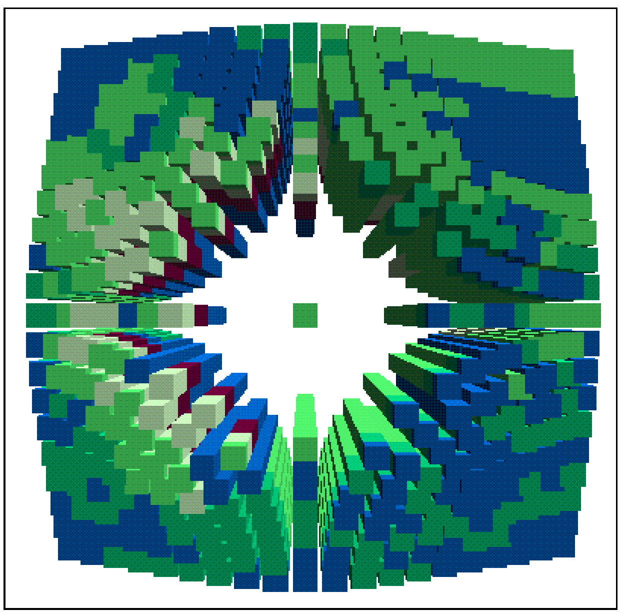

Figure 7.42: Peering through the temporal axis

access distortion can separate these columns enough to see along their full length. However, it still preserves spatial adjacencies sufficiently to be able to recognize neighbors.

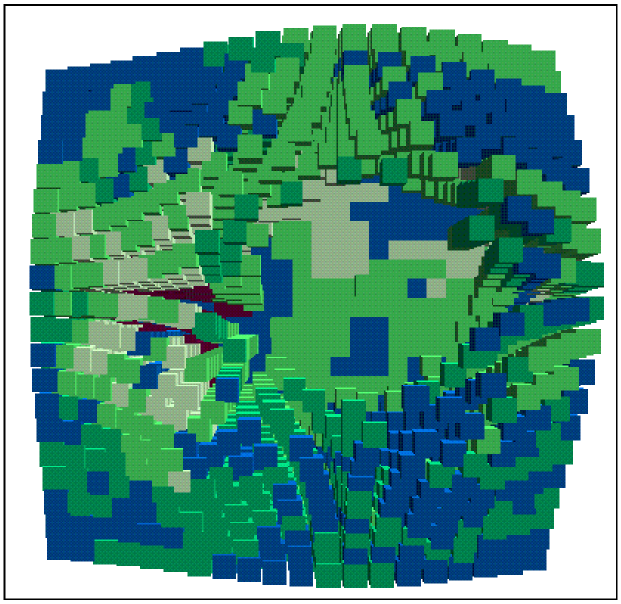

Figure 7.43: Browsing through temporal space to a spatial layer

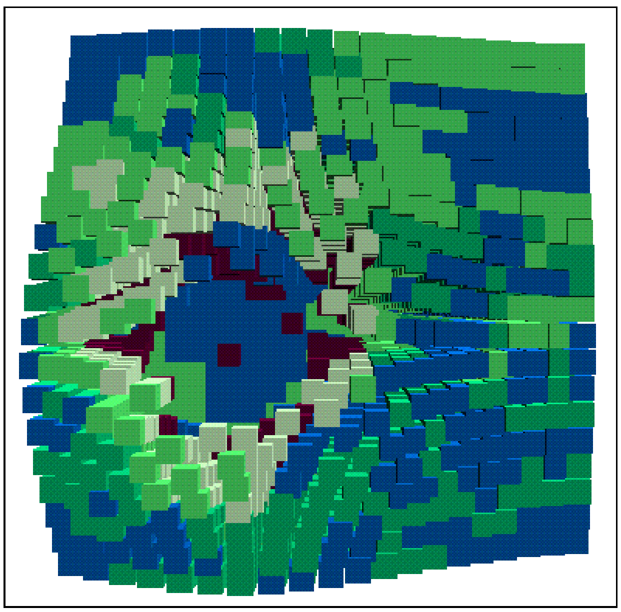

Figure 7.44: Continued browsing reveals fire start

If visual access is being used as a search probe with no specific focus, it can be useful to leave chosen section of the data undisturbed. In Figures 7.43 and 7.44 a particular spatial layer has been chosen for exploration. The first opening in Figure 7.43 shows a typical forest cover pattern. Moving across the data Figure 7.44 shows the beginning of a forest fire, one red cell indicating fire starting in a section of old growth. The distortion access allows one to open a ‘portal’ through time to the event of interest. One can browse across the space to see the initial extent of the fire and see the degree of spread to the subsequent time layers. It is also possible to rotate the model and examine up to the chosen layer, revealing the landscape conditions at the out-break of the fire.

7.5.7 Discussion

The distinction between data or viewer relative magnification or displacement patterns offers new flexibility in the application of these techniques through various combinations.

We have examined the issues that arise when applying 2D detail-in-context distortion methods that were designed to improve visual access of 2D VReps to 3D VReps. These 2D methods are insufficient to derive much benefit for 3D VReps as the recurring problem of interior, or hidden, data becoming inaccessible nullifies any advantages of a straightforward extrapolation.

Visual access distortion provides a solution to one of the fundamental problems in 3D viewing, namely occlusion. This is achieved by clearing a line-of-sight to any chosen internal region of 3D VRep. This distortion uses viewer rather than data alignment, allowing the user to retain sight of their focus during rotation. The distortion is interactively constrained and distinct from any magnification function. It can also be used with a chosen magnification pattern to create detail-in-context views. This distortion also maintains visual conformation of inter-node connections and relative positions.

Future work with this concept is being carried out by D. Cowperthwaite and includes the application of these ideas to volumetric 3D data.