Chapter 4

EPS: Extending Capabilities

In any detail-in-context presentation, compromises must be made between the amount of magnification in the foci and the amount of compression and/or distortion in the rest of the image. Extreme focal magnification creates a severe space problem causing extreme compression of the remaining context. Also, if smooth visual integration is desired, some type of distortion will be needed to connect the magnified regions to the compressed regions. Greater differences in scale between magnified and compressed regions require a greater degree of distortion to visually connect them.

Controlling the degree of magnification is discussed in Section 3.5. In this Chapter, controlling the degree and location of the distortion is considered. The reasons for distortion control include: providing regions of the representation that are untouched and therefore easy to recognize, preventing the distortion from becoming too extreme, locating the distortion in the regions that are currently of least interest for the current task, and maximizing magnification without losing context.

An important advantage of detail-in-context presentations is that spatially located foci reveal details of current interest within the context of the entire representation, providing information about how sub-regions are embedded in the overall structure. One reason for more than one focus is to facilitate visual comparison of the sub-regions (i.e. detail-to-detail readings, see Chapter 1). However, these sub-regions are spatially separated by their context. Frequently, when carrying out a visual comparison task people place the objects to be compared adjacently. This freedom of re-positioning chosen regions of interest is one advantage of separate views in multiple windows. Previously, this freedom of re-positioning and the advantages of detail-in-context had appeared to be incompatible. Both of these facilities are important when examining information visually, and it is unfortunate to have only presentation paradigms that support one to the exclusion of the other. We introduce *folding*, the freedom of focal re-positioning within detail-in-context presentations. At this point we know of no of other work that examines the repositioning of foci in a detail-in-context view.

This Chapter proceeds as follows. In Section 4.1 we discuss the provision of distortion control allowing the user much more precision in allocation of space. This includes the amount of space a focus uses, the spread of the distortion and the location and nature of the compression. Section 4.2 desrcibes how further control of distortion pattern can be achieved through adjusting the distance metric. In Section 4.3 we use the fact that EPS can explain both detail-in-context presentation (Section 3.2) and re-positioning (Section 5.1) to investigate their incorporation through folding. In Section 4.2 investigates the effect of different distance metrics, and a discussion in Section 4.4 concludes the chapter.

4.1 Distortion Control

The combination of magnification, compression and distortion creates a visual pattern. In this section we discuss this visual pattern, introducing the possibility of distortion control to offer the user considerable choice not only over the amount of compression but where minimum and maximum compression occurs. (This idea is discussed further in Chapter 5).

Since previous methods used *global distortion* there were only two regions: the focus and distorted context. As a result, discussions had only considered the visual integration from focus to context and took for granted that the maximum compression should be at the edges of the context. Limiting the spread of the distortion raises concerns with integration from the focus into the distortion and from the distortion into the untouched context. Moreover, there is the possibility of choosing the nature and location of the maximum compression. The Gaussian drop-off function lends itself to creating constrained distortions, making it possible to affect only as much of the context as necessary to provide the space for the magnified focus.

4.1.1 Controlling the Size of the Lens

There are three aspects to the size of the lens: the lens height, the focal radius, and the lens radius. Adjusting the lens height provides magnification control (see Section 3.5). As the height increases, there is an increase in magnification and a compensating concentration of compression that develops in the distorted area that surrounds the focus (Figure 4.1).

Adjusting the focal radius allows a choice between point or scaled-only foci (Section 3.4) and moves the distorted region towards the edge of the context. Rather than increasing the degree of

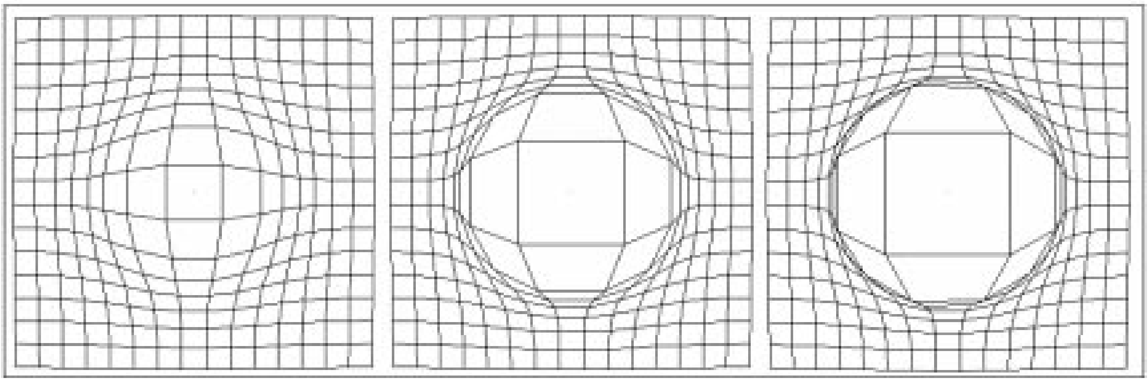

Figure 4.1: Single foci: effects of varying height only

magnification of the focus, a larger portion of the surface is included in the focus. Figure 4.2 shows the side view of a point focus next to a focus with a wider radius.

The lens radius can also be adjusted, controlling the spread of the distortion into the rest of the image. With the Gaussian drop-off function this is achieved by varying the standard deviation. In Figure 4.2 the right image shows how a larger standard deviation affects the profile view of the lens that is in the centre image. In Figure 4.3 the images show how a larger standard deviation affects the range of the distortion.

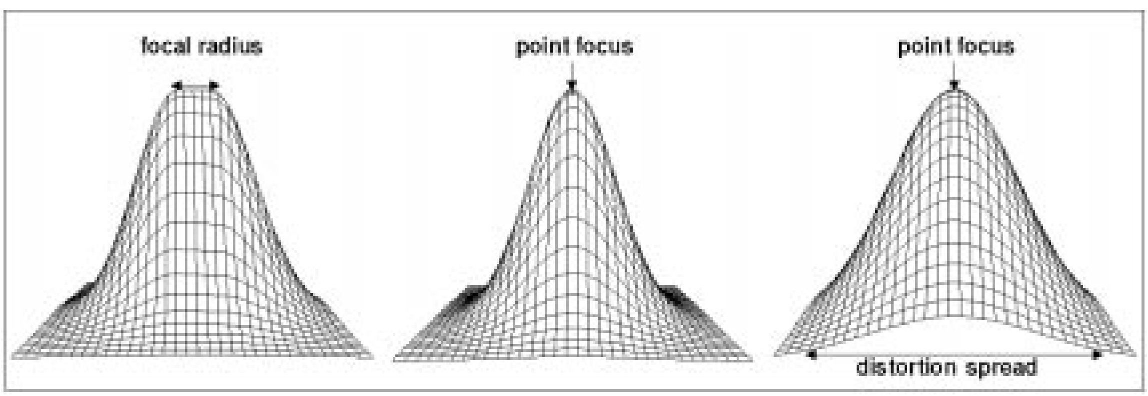

Figure 4.2: Changing the width of the focal radius (left) versus changing the width of the lens (right)

The characteristic Gaussian function has a broad top around the focus where the magnification of the adjacent area lags only slightly behind that of the focus. In some applications this is ideal, providing good local context. However, in other situations this uses too much screen space. If a small standard deviation is used the Gaussian drop-off is quite steep. This creates a small area of

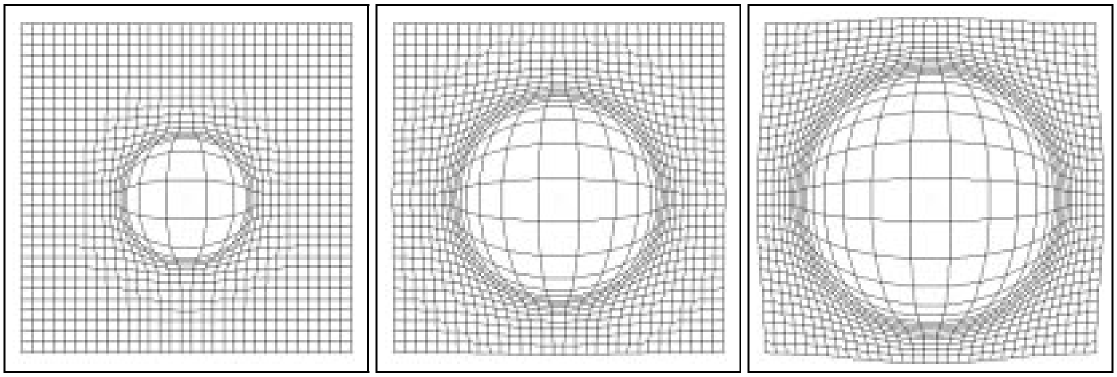

Figure 4.3: Single foci with a Gaussian drop-off function: effects of varying the standard deviation with fixed height

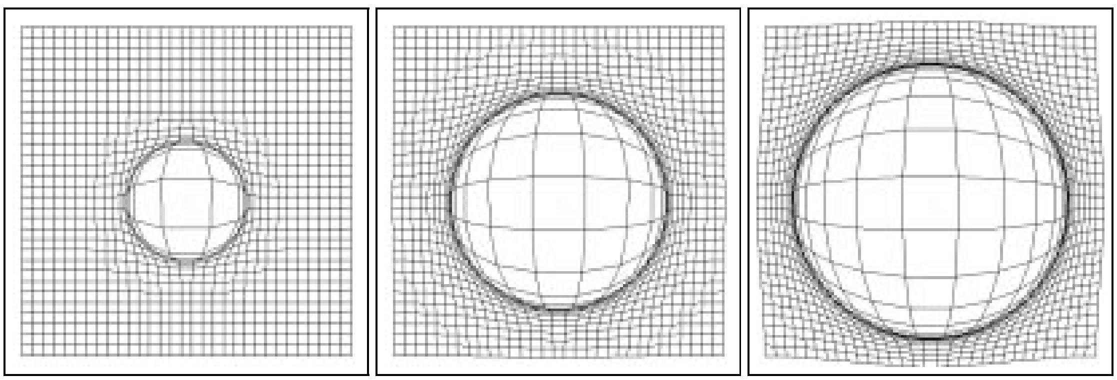

Figure 4.4: The effects of varying deviation once a ring of compression has formed

contextual magnification adjacent to the focus leading into the area of maximum compression before blending into the undistorted context. Increasing the standard deviation increases the over-all lens size, and moves the area of maximum compression further from the focus. Expanding the distance between the focus and the area of maximum compression allows more space for magnification of the context adjacent to the focus (Figure 4.3). Note that adjusting the standard deviation does not change the degree of magnification in the focus.

In this manner it is possible to choose a more gradual integration for the region of context immediately adjacent to the focus, control the location of the maximum compression, and limit the extent of the distortion. This allows choices to be made between the amount of magnification adjacent to the focus and the amount of undistorted context. However, as shown in Figure 4.4, if a ring of maximum compression already exists, adjusting the standard deviation merely re-positions it.

4.1.2 Relating Compression to Curvature

While parallel planes preserve geometric relationships, sloped sections do not. When a region of the surface is not normal to the central axis, the slope of the surface will affect the visual pattern in the projected image. The compression needed to create the space for the increased magnification can become concentrated (Figure 4.4). All curves have characteristic profiles and resulting patterns of compression; in particular, the Gaussian curve's familiar bell shape tends to result in a ring of high compression at its point of inflection.

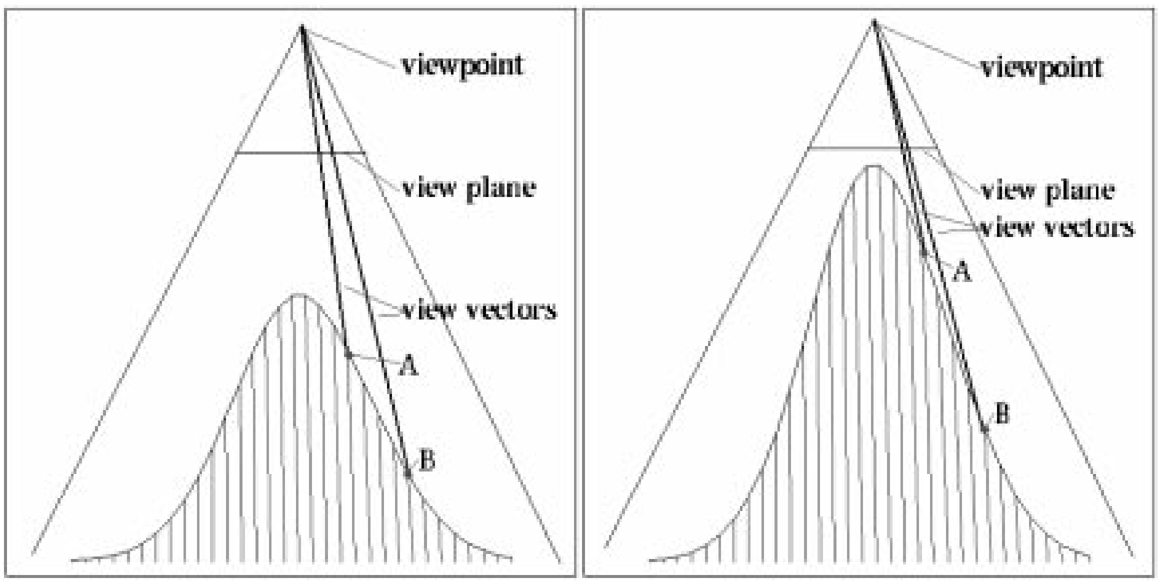

Figure 4.5: As the focal magnification increases the region between points A and B becomes increasingly compressed

Figure 4.5 shows two cross-section views of Gaussian lenses with their reference viewpoints and view planes. The two points A and B are drawn with *view vectors* (vectors from the point on the surface to the reference viewpoint). The intersection of point A's view vector and the view plane is point A's projected location (similarly for point B). The portion of the surface between the points A and B will be projected onto the region of the view plane between the intersection points of the view vectors of these points. Within this region of the view plane the magnification will not be uniform. The change in magnification will vary with both the slope and the distance. When the two view vectors become coincident the entire portion between them will be projected as a line. Figure 4.6 shows the cross-section view of a lens where the view vectors of points A and B are coincident. The projected view (Figure 4.6, right image), shows the resultant line forming the dark compressed ring

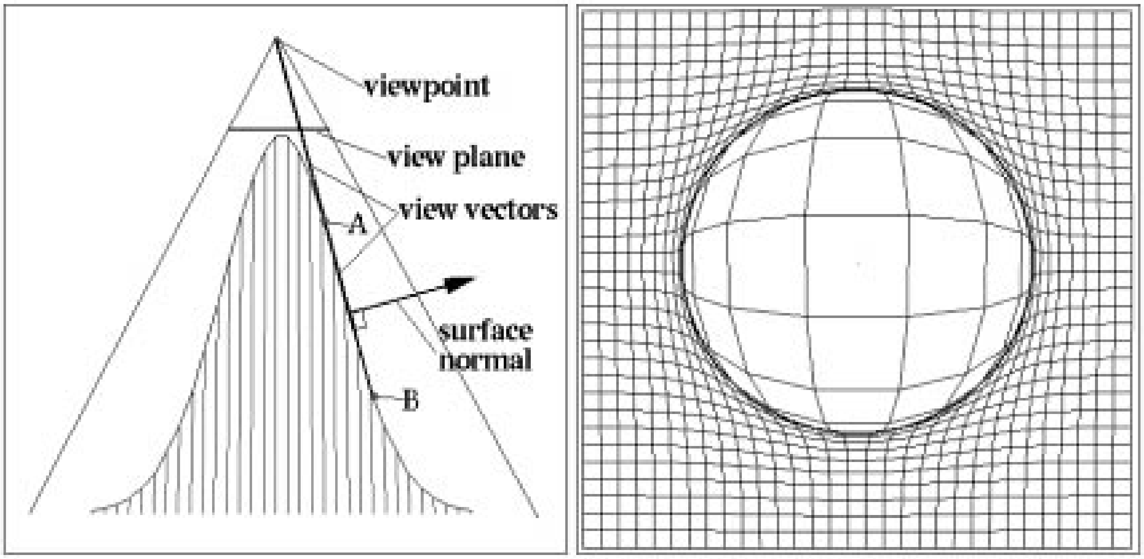

Figure 4.6: Relationship between compression and angle of surface to viewer. Maximum compression occurs when the surface normal is orthogonal to the view vector

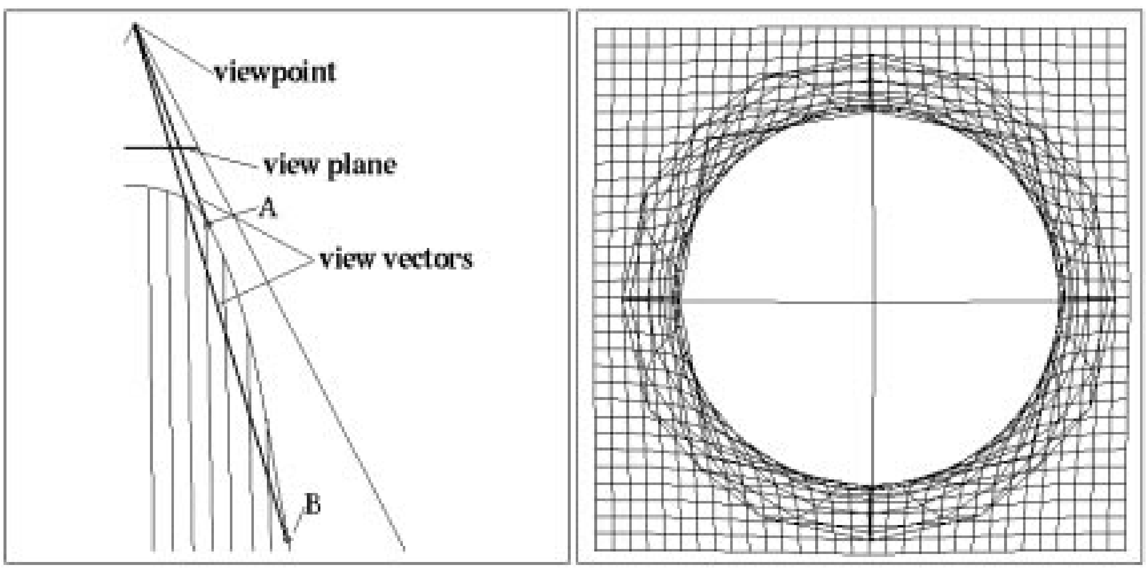

Figure 4.7: When a point's view vector passes through the surface it will either be occluded or result in information reversal in the projected view

surrounding the focus. This type of maximum compression occurs when the normal to the surface is orthogonal to the view vector (Figure 4.6, left image).

When the curvature becomes even steeper, that is, when the angle between the view vector and the surface normal is greater than 90 degrees, then the projected view will no longer give full consecutive context (Figure 4.7). In discrete representations (point and line drawings, graph layouts and network diagrams etc.) this will appear as information reversal. In opaque representation it will result in occlusion. Figure 4.7 shows point B's view vector passing through the surface before it reaches the view plane. As a result it reaches the view plane to the left of point A's view vector even though on the surface, point B is to the right of A. The grid in Figure 4.7, right image, is not opaque and therefore shows the inside of the curve between points A and B.

4.1.3 Using an Auxiliary Function

Increasing the focal magnifications eventually leads to a ring of extreme compression and then occlusion or reversal. The chief effect of controlling the standard deviation is to locate the position of the maximum curvature. This allows choices to be made between the amount of magnification adjacent to the focus and the amount of undistorted context. However, as shown in Figure 4.4 if a ring of maximum compression already exists, adjusting the standard deviation merely re-positions it. To change the rate of the compression, we need to affect the curvature itself. The basis curve may be modified by subtracting from it a second function, in this case a half sine wave (Figure 4.8). The domain of the half sine $(0, \pi)$ is normalized across the domain of the Gaussian curve.

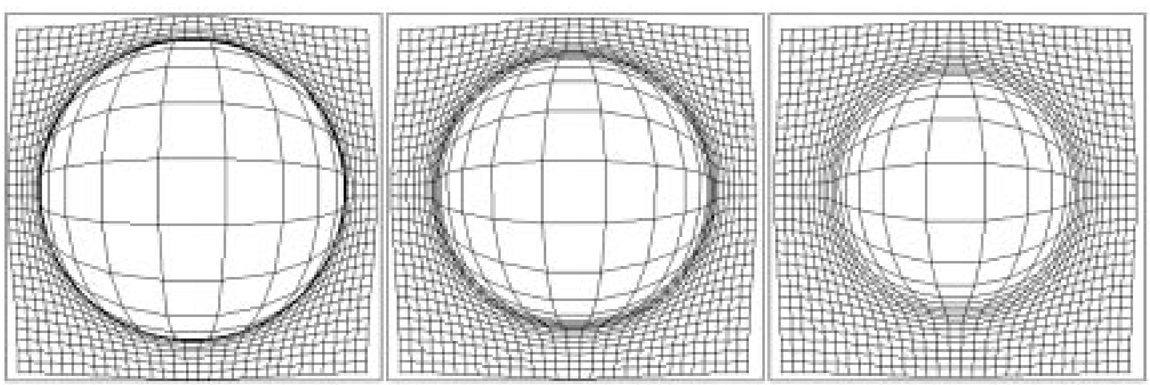

Figure 4.8: Using an auxiliary curve removes the ring of maximum compression.

In summary we allow control of the curve's profile through the height (Figure 4.1), standard deviation (Figure 4.3) and use of an auxiliary function (Figure 4.8). In this manner it is possible to chose a more gradual integration from focus to context or limit the extent of the distortion, causing more compression in the distorted region. As all of these distortion controls are left up to the user, it is possible to extend the distortion to a point that causes some areas to be compressed beyond visibility. However, as the slope and curvature are adjustable and reversible, it is possible to interactively redistribute the context in non-focal areas without losing focal magnification.

4.2 Changing the Concept of Distance

Distance is one of the factors that work together to create the organization of the presentation; the drop-off function takes a point at a specific distance from the chosen focus and computes a new location for that point. Many types of distance concept can be used. The notion of distance could be tied to the structure of the information representation; for example, path length in a graph structure. Alternatively the distance concept could be based on a combination of the task at hand and Euclidean distance. For example, in looking at a road map a police officer may be interested in the distance representing travel time when in pursuit.

One concept of distance is L_p-metrics [124]. For two-dimensional distances between points $p_1(x_1, y_1)$ and $p_2(x_2, y_2)$ an L_p metric is defined as:

$$L(p) = \sqrt[p]{|x_1 - x_2|^p + |y_1 - y_2|^p}$$

Changing the distance metric changes the shape of the unit circle and consequently changes the corresponding shape of the lens.

The L-one metric:

$$L(1) = \sqrt[1]{|x_1 - x_2|^1 + |y_1 - y_2|^1}$$

resolves to the distance in x plus the distance in y:

$$L(1) = |x_1 - x_2| + |y_1 - y_2|$$

As this corresponds to distances on city streets (assuming regularly laid out blocks) it is referred to as *Manhattan distance*. An L-one unit circle is diamond shaped (Figure 4.9).

The L-two metric:

$$L(2) = \sqrt[2]{|x_1 - x_2|^2 + |y_1 - y_2|^2}$$

is Euclidean distance. Lenses using L-two metrics have radial distortion (Figure 4.10).

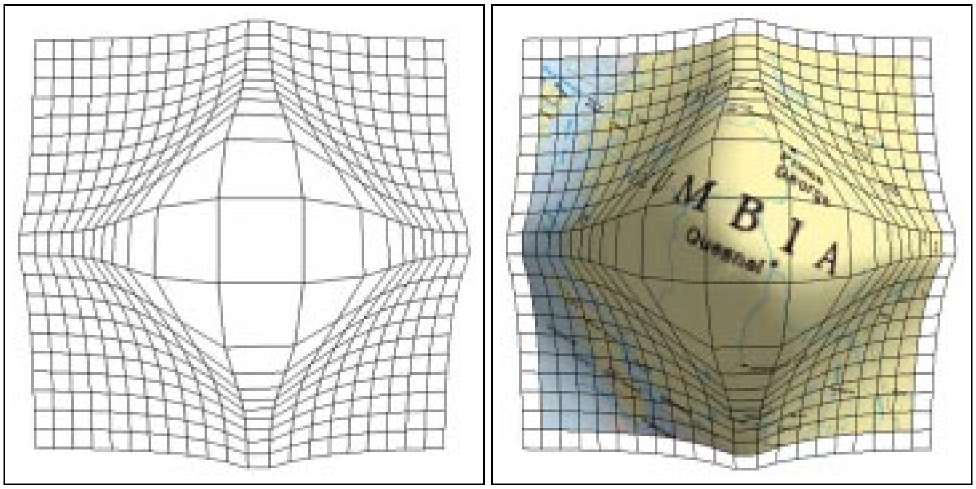

Figure 4.9: L-one metric (for image credit for the map of British Columbia see Appendix C.2)

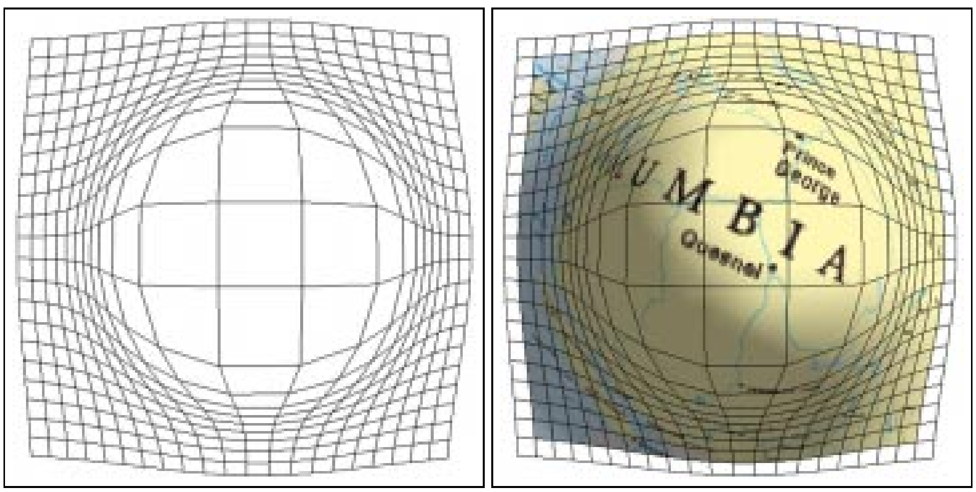

Figure 4.10: L-two metric

The L-three metric (Figure 4.12) and L-four metrics (Figure 4.11) are:

$$L(3) = \sqrt[3]{|x_1 - x_2|^3 + |y_1 - y_2|^3}$$

$$L(4) = \sqrt[4]{|x_1 - x_2|^4 + |y_1 - y_2|^4}$$

Notice how the circle in Figure 4.10 becomes gradually more ‘square’ in Figures 4.11 and 4.12. This trend increases as the L-metric approaches ∞.

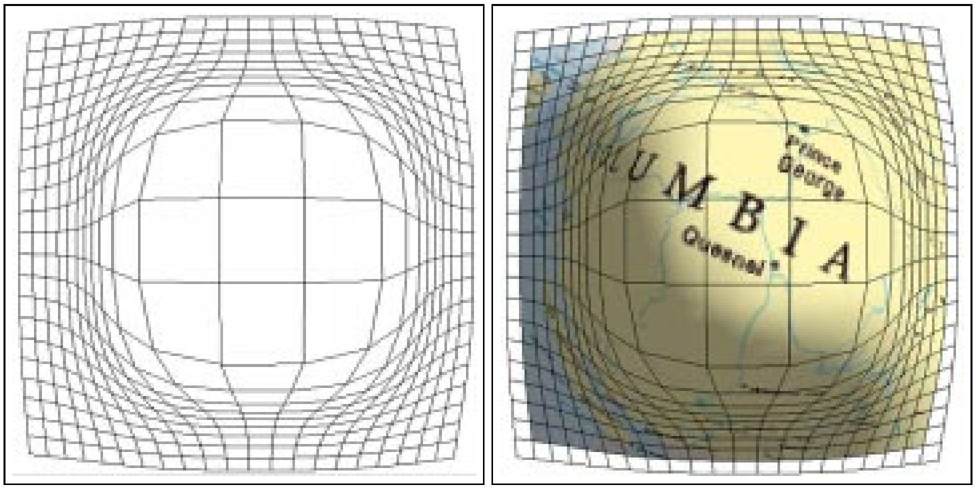

Figure 4.11: L-three metric

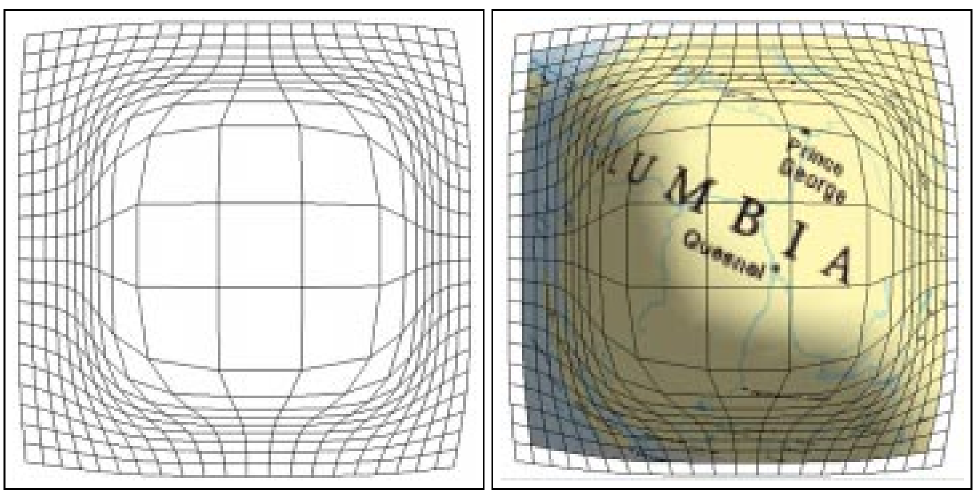

Figure 4.12: L-four metric

The L-∞ metric (Figure 4.13) is:

$$L(\infty) = \sqrt[\infty]{|x_1 - x_2|^\infty + |y_1 - y_2|^\infty}$$

which resolves to:

$$L(\infty) = \max(|x_1 - x_2|, |y_1 - y_2|)$$

The unit circle for this distance metric is a square and the resulting lenses are orthogonal.

Any given distortion function can be applied either orthogonally or radially creating significantly different layouts (Figures 4.10 and 4.13). This distinction has been frequently noted [108, 138, 150].

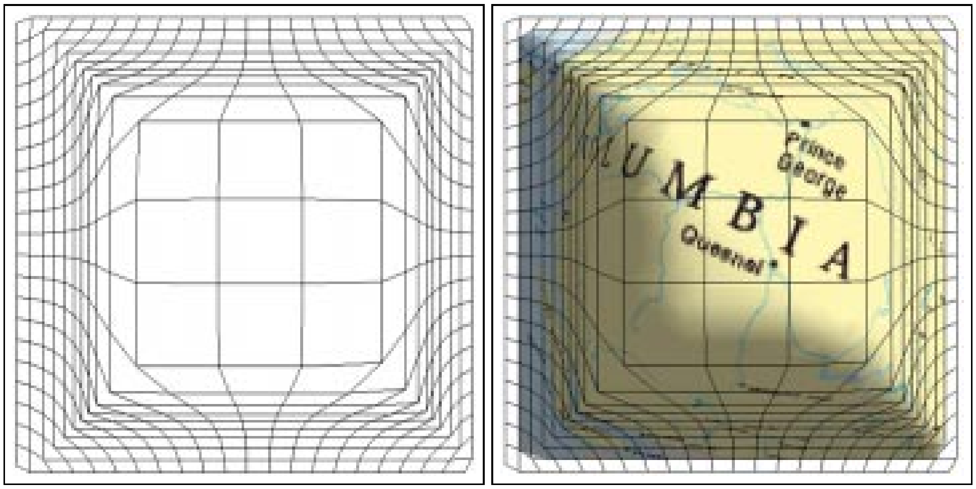

Figure 4.13: L-∞ metric

However, previously only two distinct options have been offered: radial distortions created with polar coordinates, and orthogonal distortions created with Cartesian coordinates. The difference in layout is considered important in terms of preservation of the user's mental map [107] and in terms of finding appropriate presentations for specific applications [150, 166]. For further discussion see Chapter 2, Section 2.2.1 and Chapter 6, Section 6.5.4.

Varying the distance metric offers this choice (radial L-two and orthogonal L-∞), and a full range in between is possible. Furthermore, L-one offers an interesting possibility in that its diamond shape provides slower integration for those aspects of the representation that are orthogonal to the focal centre (Figures 4.9 to 4.13 vary only the distance metric).

The discussion thus far has focused on spatial manipulation. However, many representations have a concept of distance that is integral to their structure while not necessarily spatial. For example, graphs can be thought of as having a notion of distance such as path length. Noik [116] discusses the use of conceptual distance in creating his layouts. Conceptual distance depends on how "far apart" ideas are and depends on domain knowledge.

4.3 Folding: Re-positioning Established Foci

We want to combine the visual gestalt advantages of detail-in-context presentations with the ability to bring magnified regions of spatially separated lenses together in order to facilitate visual comparisons. These seem mutually exclusive because in detail-in-context presentations foci are separated

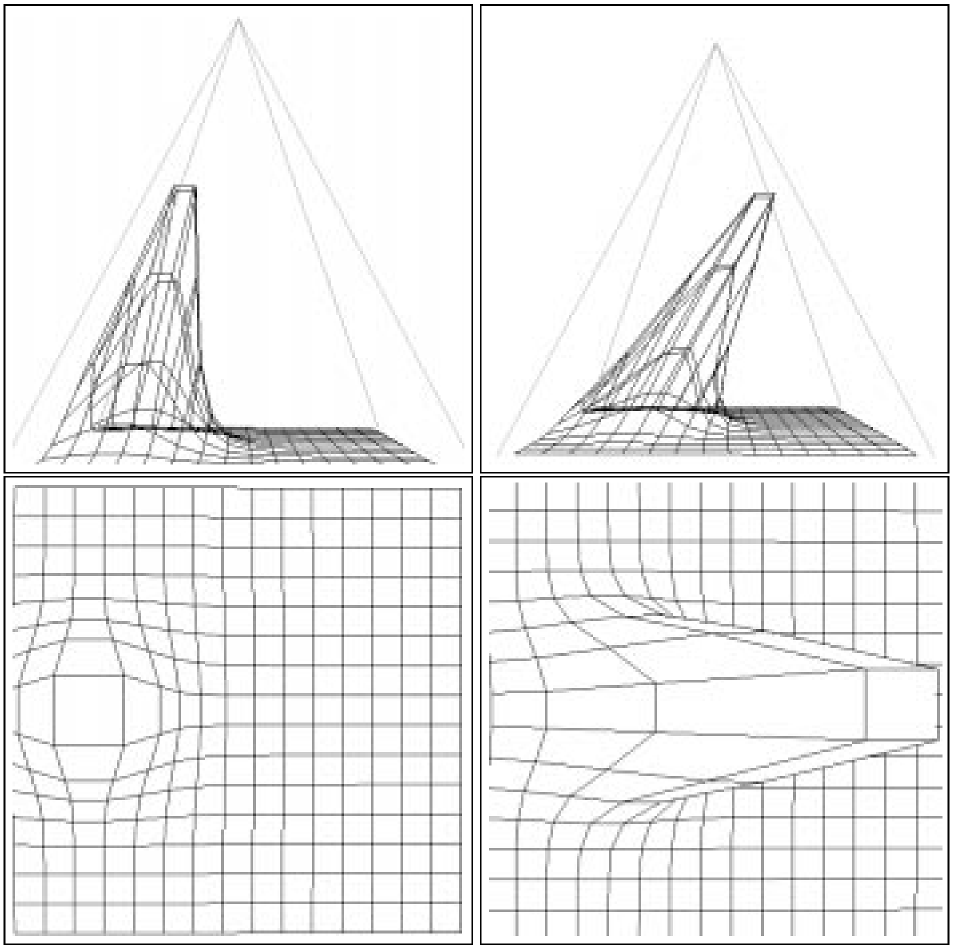

Figure 4.14: Single off-center focus; left, pulled towards viewpoint; right, folded across the image

by regions of compressed context, and preservation of this context imprisons the foci, preventing re-positioning or re-aligning. However, the EPS framework has allowed us to re-position foci within a detail-in-context presentation.

We introduce *folding*, the re-positioning of established foci. There is an important distinction between folding a focus and performing a *roving* search with a focus. Both involve moving a focus. When using a focus for a roving search it passes across the representation and magnifies a new region of the representation according to its new location. When folding a focus the region of the representation that is magnified in the focus remains constant. It is the x, y position of the focus that changes, re-positioning the region of the surface in the focus, causing it to 'fold' over other regions of the surface. Folding allows freedom to reposition magnified regions without detaching them from

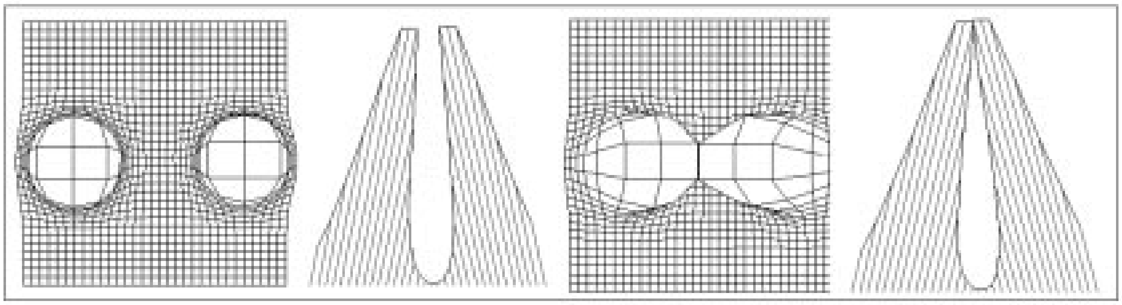

Figure 4.15: A pair of foci, repositioned to be adjacent

the rest of the image.

Figure 4.14 shows side and projected views of a viewer-aligned focus (left) and a folded focus (right). The folded focus has been pulled across from the left side of the image to the right side of the image. While it is apparent that the focus is still part of the surface, it is also clear that due to the folding a region of the surface is occluded. If the surface is sufficiently visually convincing as an object, then it will be assumed to be complete. If the surface is perceived as complete, then it can be stretched, folded, and warped without portions of it appearing to cease to exist. Furthermore, it is interactively possible to unfold the focus or view it from different angles exposing temporarily obscured sections. In introducing folding we have moved from preservation of full context into a variation of preservation of sufficient context. The question of whether folded context is sufficient for a user to retain the sense of the representation being intact and of the location of the focal region within it remains open.

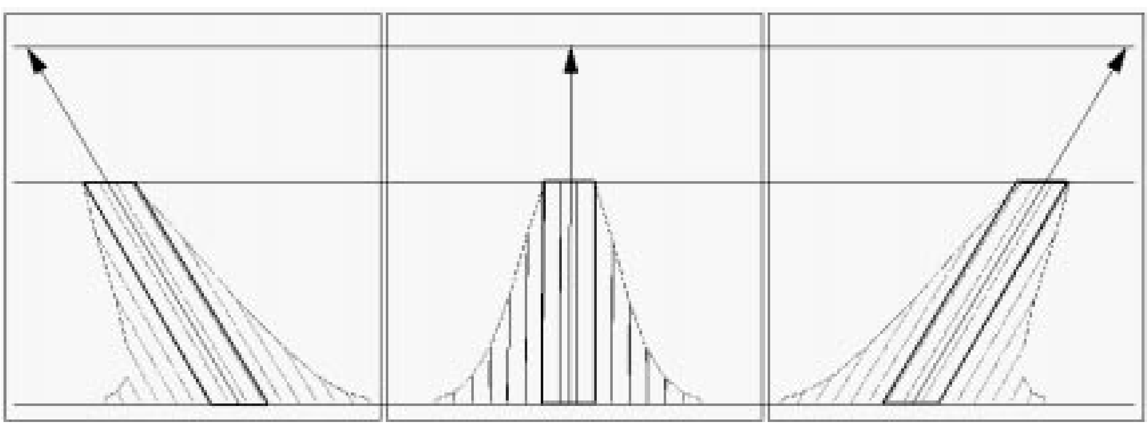

Figure 4.16: A central lens with the viewer-aligned vector sheared to the left (left) and to the right (right) folds the focus accordingly

Repositioning a focus while retaining the effect of a detail-in-context presentation is possible in part because of the three-dimensional nature of the distortions. A focus, or magnified section, is the top of a ‘hill’. The steepness of the sides of this ‘hill focus’ can be adjusted to minimize interference

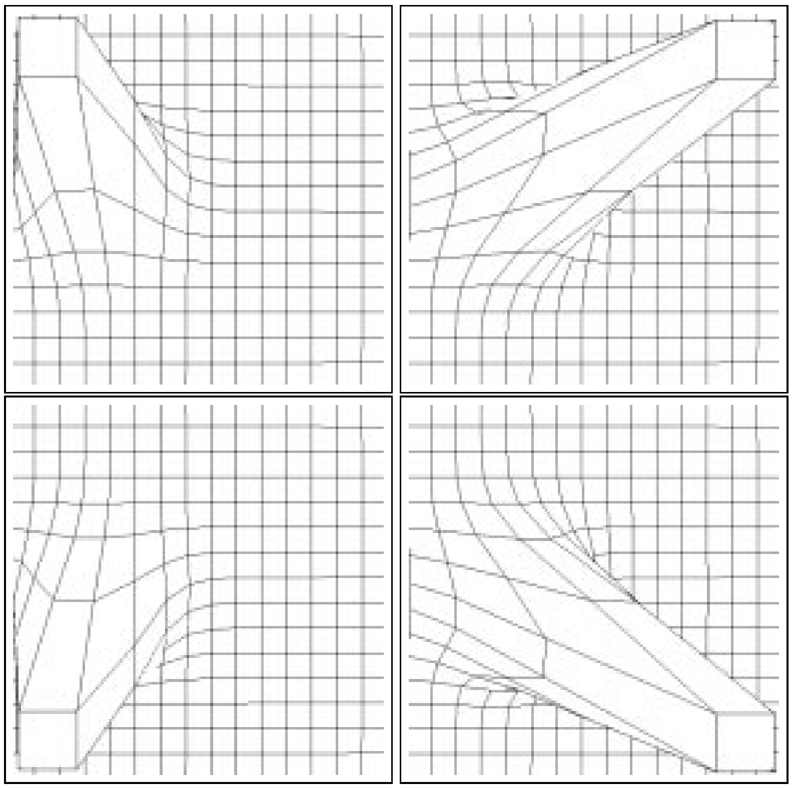

Figure 4.17: Four views of a single focus, folded to position it over each of the four corners of the surface.

with other foci. The top, or focus can then be moved without changing the location of the base of the hill. The stable base maintains the same section of the representation within focus, and keeping the focus at the same height retains the degree of magnification. It is the sides of the hill that are stretched and bent. The context is maintained over the sides of the hills and across the valleys. This allows foci from different locations on the surface to be repositioned adjacently, without losing the sense of context. Figure 4.15 show a pair of foci repositioned to be adjacent.

Surface folding is achieved by shearing the viewer-aligned vectors. The centre image of Figure 4.16 shows a lens in the centre of the field of view with the perpendicular viewer-aligned vector as it is directed towards the RVP. The lens’ viewer-aligned vector starts from the centre of the focus at the base plane, and passes through the centre of the elevated region of the surface (the magnified focus). Just like viewer-aligned foci, folded foci have one central translation vector that determines their orientation. Figure 4.16 shows these central *orientation vectors*. The position of the focus can be readily shifted by pointing the orientation vector elsewhere. The central orientation vector is directed at any point on the plane which contains the RVP and is parallel to the base plane. The vector is then z-normalized. The parallelogram formed by connecting the selected focal region on the base plane and the elevated focal region on the surface is sheared when the angle of the orientation vector is changed (Figure 4.16). The properties of height, magnification and scaling for folded foci remain constant, while their position in x and y changes. This allows spatially separate areas to be positioned adjacently while maintaining a continuous surface between them. At any moment the representation on this surface can be viewed by rotating and adjusting the three-dimensional image or by unfolding the surface. Figure 4.17 shows that a single focus from an arbitrary location can be placed over any region of the surface through folding. Note how the magnification stays constant.

Folding was introduced to allow spatially separate focal regions to be placed adjacently. In use, one frequently tips a focus to obtain a better view of the its adjacent compressed context. For example, Figure 4.19 left, shows an H-curve [61, 92] (an H-curve is a DNA walk in which the orientation of each line segment represents one of A,C,G,T) in normal position. The centre image shows a lens magnifying the start of the h-curve. In the right image the lens is folded slightly, keeping the focal magnification and providing a better view of the rest of the h-curve.

Figure 4.18 shows an iterated K12 graph. The outer subgraphs are so densely packed that one cannot tell for sure what they are. Two of these have been magnified and compared to verify that they are indeed K12's.

It is possible to think of folding as a method for hiding sections of the representation that are not of immediate concern. The result is a method that can be used analogously to folding a printed map to hide and/or expose the areas of interest. Since folding is directly reversible, regions that have temporarily been hidden are readily retrievable. In fact, one can unfold to allow closer examination of the connections between the expanded detail in the focus and the rest of the representation.

4.4 Discussion

While detail-in-context distortion presentations hold much promise as a method for interacting with large representations, the amount of magnification is limited by the preservation of context, as was the freedom to re-position and align regions of interest. We introduce distortion control and folding.

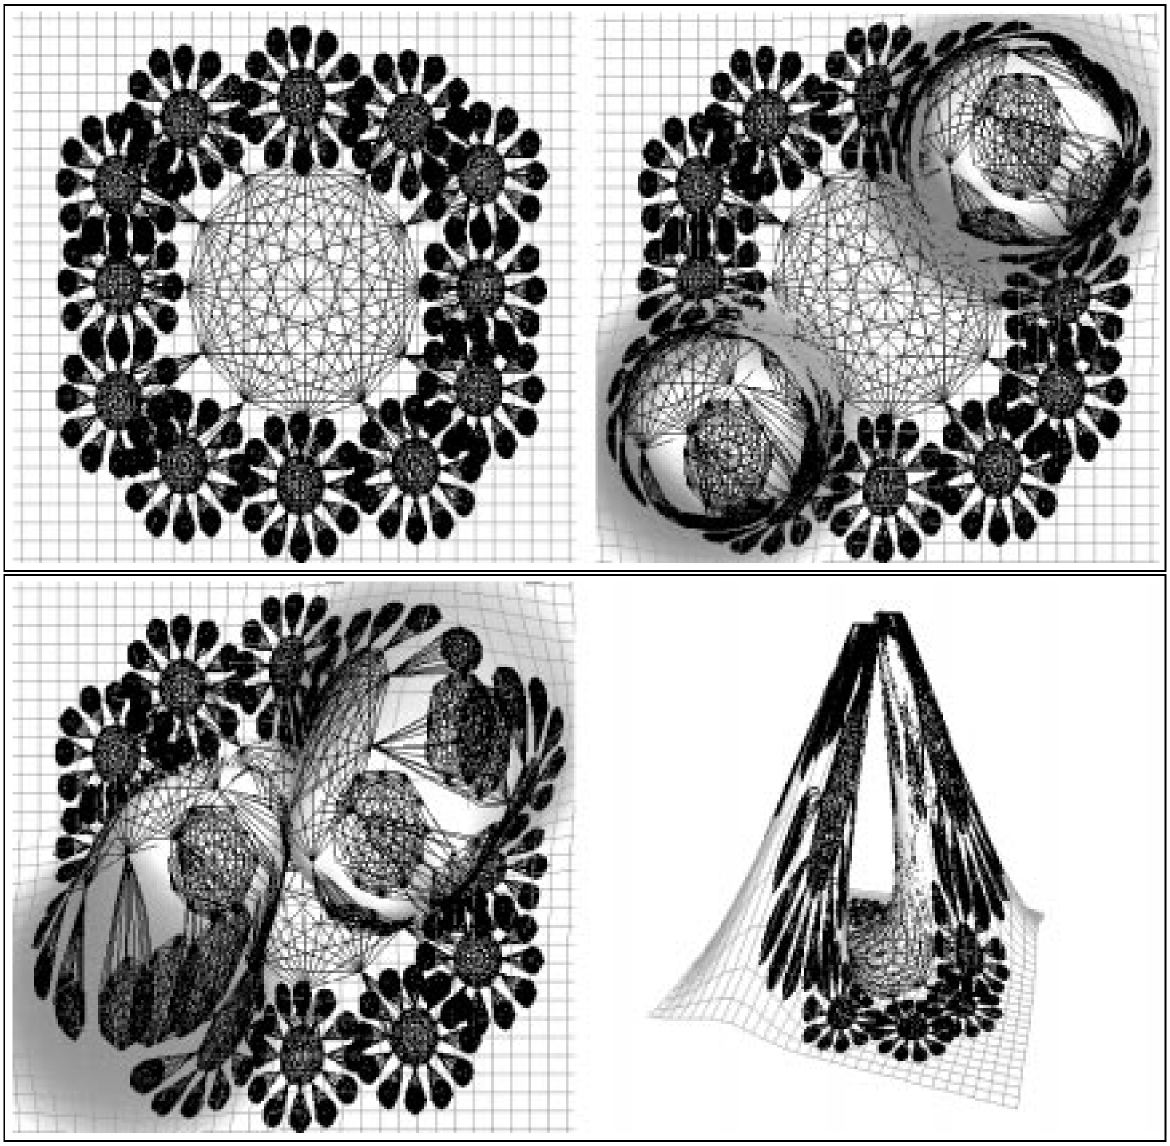

Figure 4.18: An iterated K12: top left, normal view; top right, two foci magnified to reveal subgraphs of K12; bottom left, iterated K12 folded; bottom right; a side view of the folded iterated K12

By adjusting the lens' curvature parameters the user can determine the distribution of compression. For instance, distortion can be contained within a relatively small region of the surrounding representation leaving most of the context undistorted. In other techniques the pattern of distortion was controlled by the system, user choices being limited to such things as a global choice between Cartesian or polar transformations [83, 138]. In our approach, user access to the parameters that

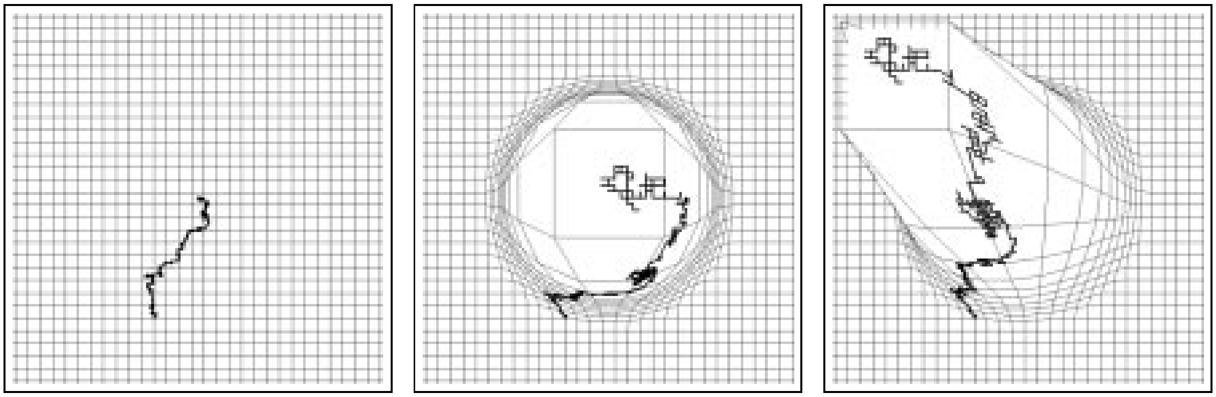

(a) A DNA walk in normal position

(b) The beginning of the walk is magnified

(c) Folding the lens provides an improved view of the rest of the DNA walk

Figure 4.19: Taking advantage of folding with a single focus

affect distortion patterns is unconstrained, therefore it is possible to create curves that obscure some context. However, just what has been obscured is always evident and the actions are readily reversible. Further control of the distortion is possible through use of different distance metrics.

We also introduced folding as a novel concept which combines the advantages of detail-in-context viewing with the freedom of movement provided by magnified views in separate windows. Other approaches had made the repositioning of focal regions difficult if not impossible.