Chapter 5

Lens Library

Existing presentation solutions such as full zooming environments, insets and various distortion approaches, create visual displays that vary considerably, visually and algorithmically. EPS provides a way of understanding how these seemingly distinct solutions relate to each other. Furthermore, EPS provides a method of relating them algorithmically, allowing the inclusion of more than one presentation solution in a single interface.

We explain how each presentation type can be thought of as an EPS lens and then show how EPS functionality follows. An EPS lens has the following parts (see Figure 5.1):

1. a *focus* is the selected point or region,
2. a *focal height* is the distance from the focus to the base plane,
3. the *focal magnification* is the degree of magnification of the focal region,
4. a *focal centre* is the centre of the focus and locates the lens on the surface,
5. *central alignment vector* sets the angle of the focal translation,
6. a *focal edge* is the outline of the selected focal region,
7. a *distortion region* is the compensating connection that exists between the focus and the context,
8. a *focal connection* is the join between the focal region and the distortion region, and
9. a *context connection* is the join between the region of distortion and the context.

The lens itself includes all of these parts, though some of them may be degenerate. For instance, with a point focus, the focus, focal centre and focal edge are the same.

The effect of a lens on the presentation is a result of all of these parts. In Section 3.5 we discussed the interaction between focal magnification and focal height, allowing specific degree of

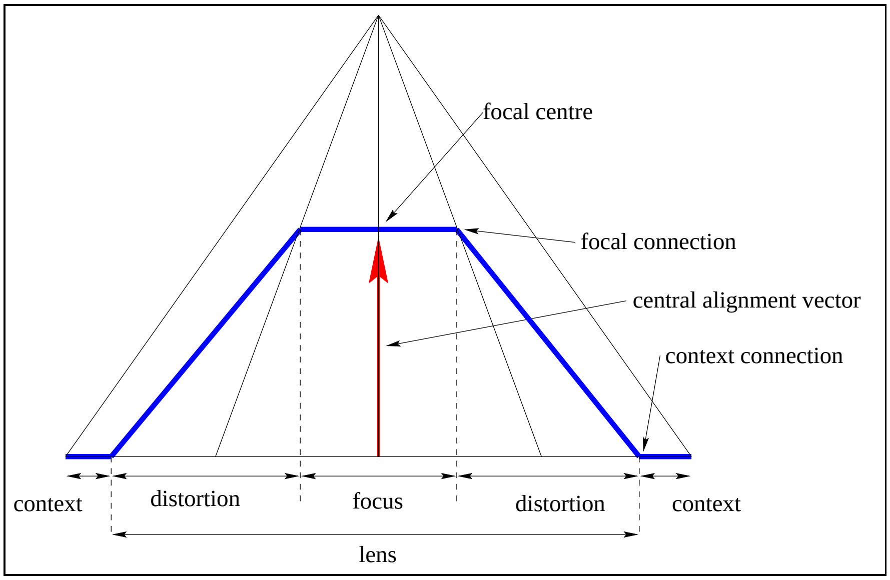

Figure 5.1: The components of a lens

magnification to be used as input, and in Section 3.4 we discussed the provision of arbitrary focal shapes. In this chapter we discuss how varying the drop-off function affects the region of distortion and the nature of the focal connection and the context connection. Though the first lens developed with EPS used the Gaussian drop-off function to create the region of distortion, EPS is not tied to any particular function. In this chapter we step through a variety of drop-off functions discussing the variations in visual patterns that arise.

Creating a presentation involves finding a balance between the magnification required and some compensation. This compensation can take the form of loss of context, compression, distortion, or other visual discontinuities. Different drop-offs create characteristic curvatures and result in different presentation patterns. Whether these characteristic patterns have advantages or disadvantages is probably not absolute but dependent on the information, the task and the preferences of the user. The purpose of this chapter is not to judge these lenses but to extend the variety of lenses. Because these lenses can all co-exist in a single presentation environment we think of this as developing a lens library. We leave the choice of which lens(es) to use up to the user or the developer of an application. We hope that having a lens library to choose from will allow an application developer to better match specific information and task needs. The purpose of this exploration is to develop an understanding of interrelationships between existing presentation methods and to extend the scope of presentation space, creating a library of lenses that exist within one framework.

Section 5.1 examines the simple situation in which the slope of the drop-off function is zero. Section 5.2 and Section 5.3 look at step drop-off functions. Section 5.4 looks at linear drop-off functions and Section 5.5 at non-linear functions. Section 5.6 compares the distortion effects of previous methods with EPS. Section 5.7 discusses different functions in combination, and Section 5.8 concludes the chapter.

5.1 Zooming, Panning and Scrolling

This section describes the effect of replacing the Gaussian drop-off function (see Chapter 3.2) with a function where there is *zero* drop-off from the focus. With a zero drop-off function, the height h_p, for any point p is the same as the focal height h_f, ($h_p = h_f$). The surface that holds the 2D representation is kept parallel to the base plane and the view plane. Recall from Section 3.2 that raising the entire surface parallel to the base plane towards the viewpoint corresponds to zooming, and moving the surface laterally in x and y corresponds to panning and scrolling respectively (Figure 5.2).

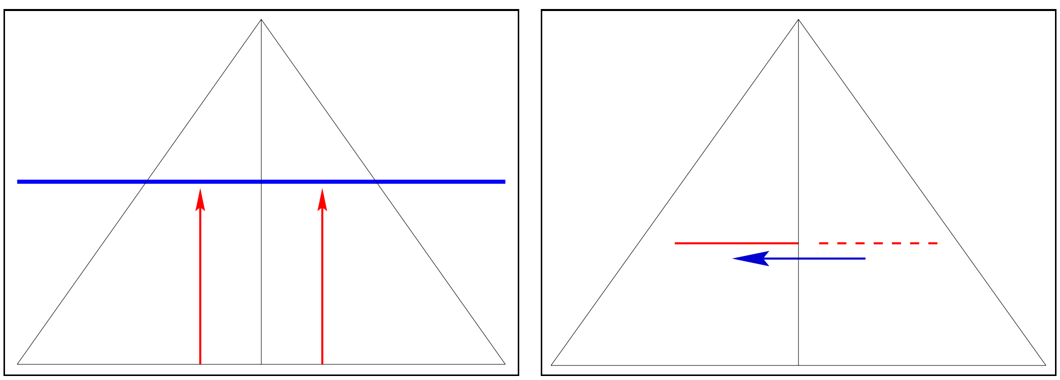

(a) Cross-section diagram of zooming

(b) Cross-section diagram of panning

Figure 5.2: Diagrams of zooming and panning

Zooming within EPS can have the functionality of a lens. This includes precise control of magnification, viewer-alignment and folding. Precise control of magnification offers the possibility to zoom in infinitely with fine control at high magnification. Figure 5.3 illustrates a zooming into a land usage map of Champaign, Illinois. In this series the magnification ranges from one (Figure 5.3(a)) to one hundred and fifty (Figure 5.3(f)).

In EPS a full-zoom lens has a focal centre and an alignment vector. If the translation vector is

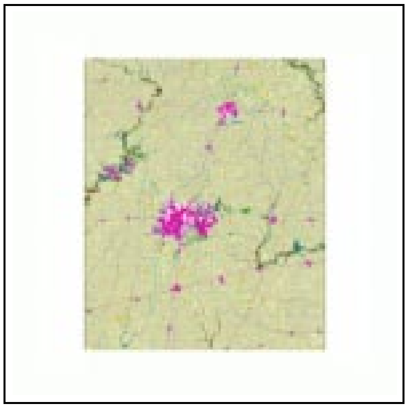

(a) Magnification of 1 (for image credit see Appendix C.1)

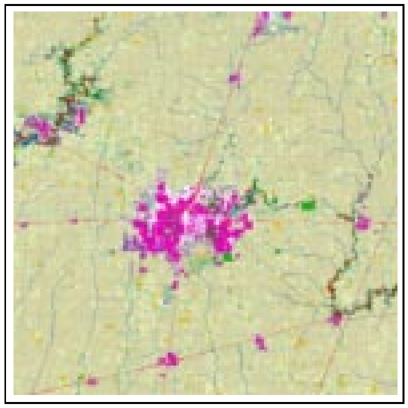

(b) Magnification of 2

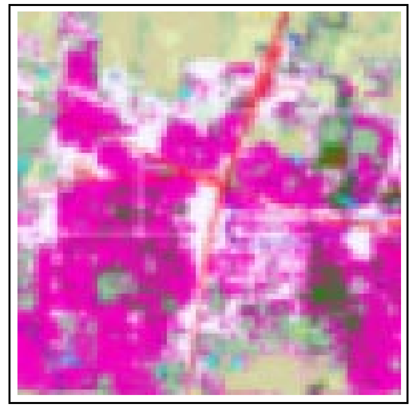

(c) Magnification of 10

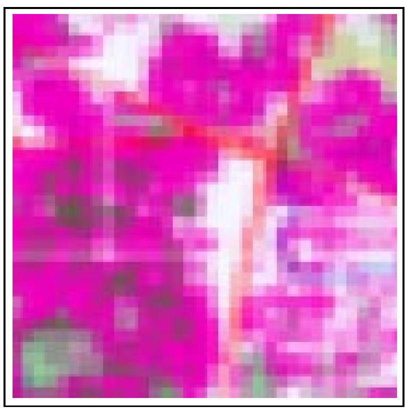

(d) Magnification of 20

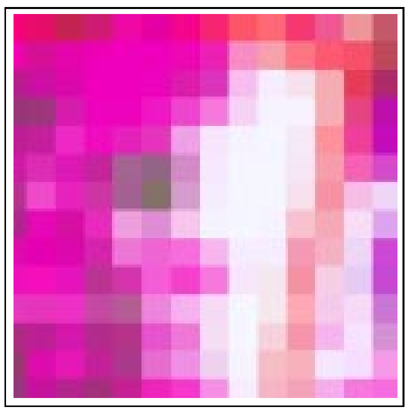

(e) Magnification of 50

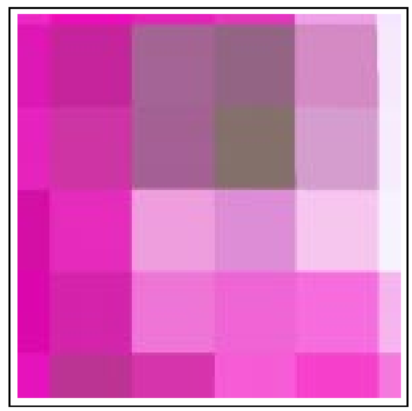

(f) Magnification of 150

Figure 5.3: Zooming in to fat pixels

used perpendicularly to the base plane (Figure 5.2(a)) then no matter where the focal centre is in the image, the centre of the image will stay in the centre of the view plane. If the full-zoom lens' central alignment vector is viewer-aligned (Figure 5.4(a)) then its current focal centre will retain its position in the view plane. For instance, if the current focal centre is in the upper left corner it will still be in the upper left corner when zoomed. The series in Figure 5.5 shows a viewer-aligned zoom. The arrow in the upper left corner indicates the focal centre before zooming. Viewer-aligned zooming keeps the region of interest in view without affecting the scaling-only advantages of the simple zoom. The focal region will always remain in view as it is magnified, keeping the same position

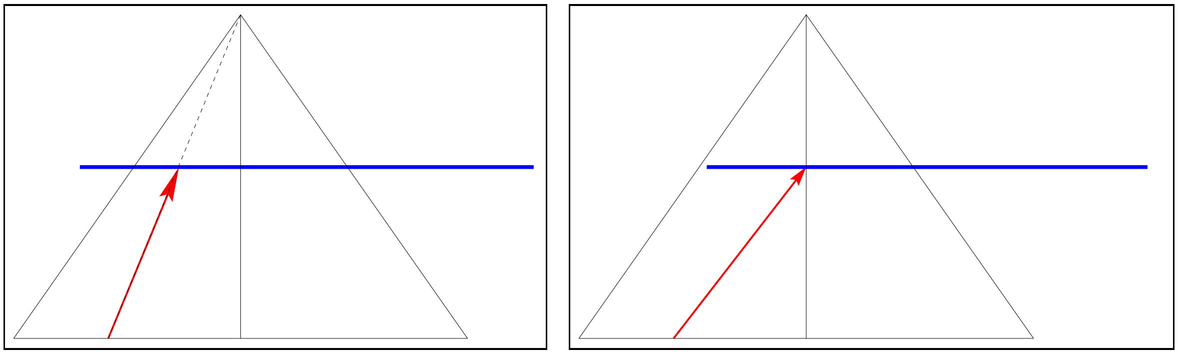

(a) Viewer aligned zoom

(b) Centred zoom

Figure 5.4: Using the central alignment vector to position the zoom

relative to the edges of the display. This provides an elegant solution to choosing an appropriate translation vector that incorporates motion in x, y and z in an intuitive manner. Alternatively the translation vector can be set to place the centre of the lens in any chosen location. Figure 5.4(b) is a diagram with the translation vector set to place the lens centre in the centre of the field of view after zooming.

Folding a full-zoom lens increases the possible modes of lateral movement (allows for panning and scrolling). The cursor can operate as though it is 'sticky', moving the whole surface directly when it moves. Moving the cursor can correspond to a roving search lens. Recall from Chapter 4 that during a roving search the base of the lens moves, changing the region of the representation that is in the lens centre. For instance, if the desired detail were not visible in the current presentation and one knew that the details were in a particular region, one could use that knowledge to locate them.

There are advantages to being able to zoom in (magnify) until all details of interest are displayed to full resolution. There are also advantages to being able to zoom in much further, as has been demonstrated by the success of Pad++ [10], to fat pixels and beyond. For instance in Pad++, one can place annotations at a sufficiently enlarged level so that they will not be visible at the intended resolution. The surface can also be zoomed out (compressed) until the entire representation is visible in the current display. This provides complete albeit compressed context. The surface can be zoomed out even further until it is like an icon of itself, or even has disappeared entirely. Again Pad++ [10] makes use of the variations within this, allowing items to get slowly smaller until they effectively disappear from sight.

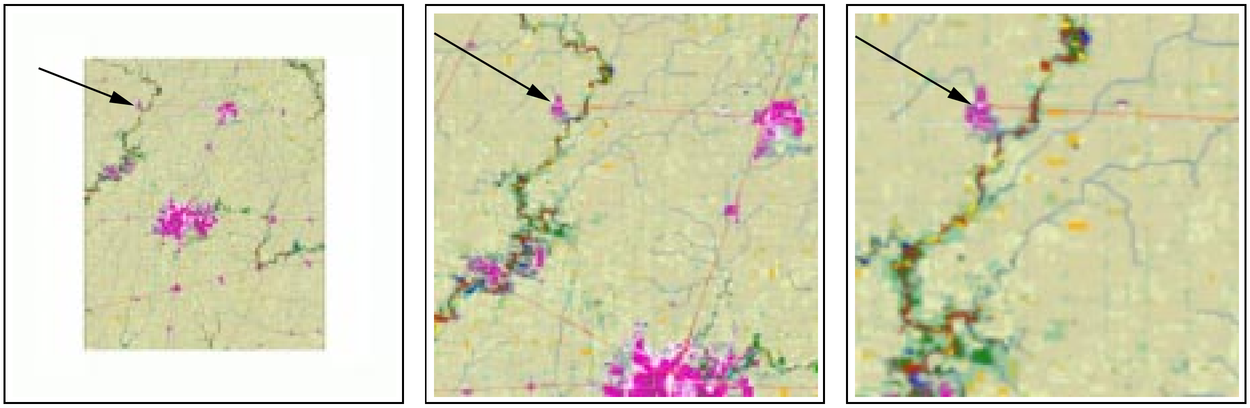

(a) The focal point has been selected (for image credit see Appendix C.1)

(b) Magnification of 2.5

(c) Magnification of 5

Figure 5.5: Viewer aligned zooming keeps the focal point in the same relative position

At any level of magnification one can pan or scroll, allowing lateral exploration of the representation. For large lateral distances relative to available display, keeping track of one's location in the representation can be difficult. If the section of interest is visible from the zoomed out position, it is possible to centre it by shifting laterally and then zooming perpendicularly as before. However, the section of interest may be so compressed when zoomed out that recognition is difficult or impossible. With Space Scale Diagrams [53] different notions of optimality of zooming paths have been explored. The authors show that for sufficient lateral distance, the optimal path is to zoom out and then zoom back in. Their user observations indicate that this may also be a more intuitive navigation method.

This full-zoom lens capability relates to the zooming capability in Pad++ [10]. The significance of this is that in EPS zooming exists within a paradigm that covers many other presentation variations.

5.2 Step Drop-off Functions

The other extreme from using a zero drop-off from the focus to the context is to use a complete drop-off, stepping abruptly down from the focus to the context on the base plane. For a step drop-off function the height h_p of all points p not in focal region is $h_p = d_b$, where d_b is the location of the base plane.

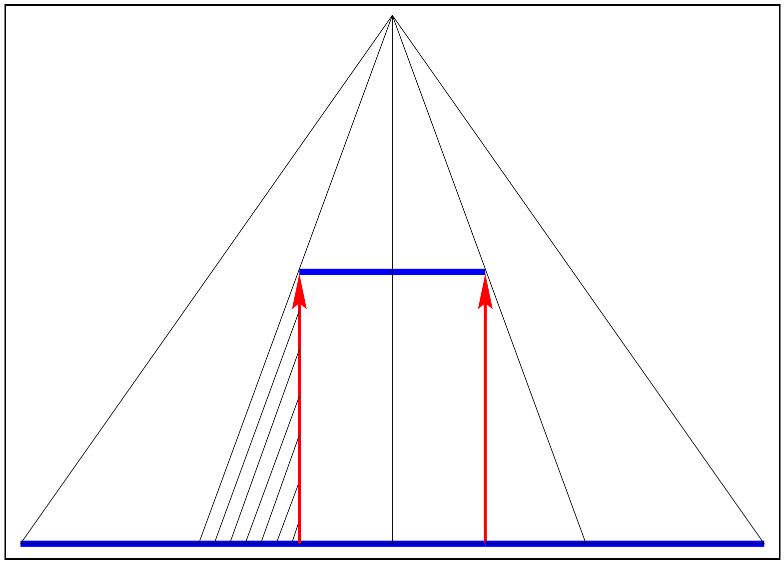

Figure 5.6: A diagram of an inset lens indicating the occluded region

A common example of a step drop-off lens is an inset. An *inset* is a selected sub-region of the representation that is magnified in place. In EPS terms an inset is a detached lens with a described focal region. As it is a detached sub-region of the original surface or surface patch, an inset can move independently from the rest of the surface. This is a slight departure from conceiving of the surface as a unit. With insets, magnification is achieved at the cost of local context. Drawing the projection rays from RVP through the edges of the focus to the base plane indicates the occluded region (Figure 5.6). Insets maintain partial context in that usually some context is still visible but the adjacent context is occluded, causing visual separation between the focus and its context. Figure 5.7 shows the effect of insets from RVP. While the issues of occlusion and separation remain significant, allowing freedom of lateral translation through folding can provide the ability to see the region that was occluded. However, either a new region will be occluded or the inset and the context will be completely separate. Folding can be used to re-align foci for ease of visual comparisons between previously separated sections.

5.2.1 Manhattan Lenses

While insets provide some context, there is always some occlusion or separation that makes the focus and its context perceptually distinct. A *Manhattan lens* is still a complete step function but the

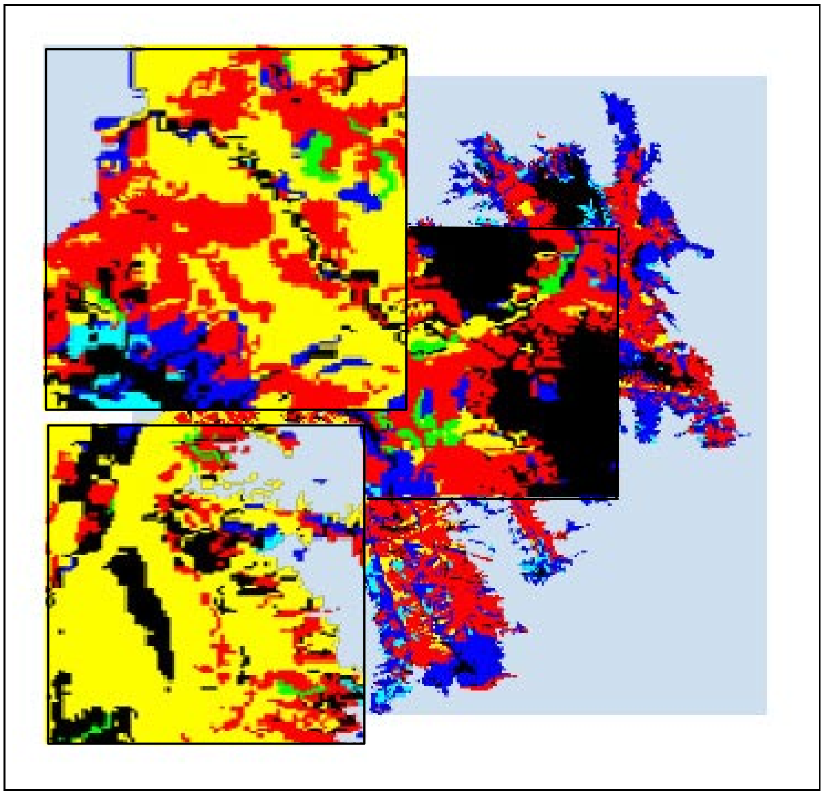

Figure 5.7: Insets provide magnification to scale but cause local occlusion (for image credit see Appendix C.5)

surface is stretched to keep the focal region attached to its context. Figure 5.8 is a cross-section diagram showing the edge of the focus stretched to connect to the surface at the position the focus had before magnification. When viewed from RVP, this presentation is identical to an inset. Figure 5.7 has the appearance of insets magnified in place. However, when the focal regions are attached to the surface, folding provides a visual connection (Figure 5.10). Manhattan lenses provide interactive access to a modified detail-in-context reading. The name Manhattan comes from their appearance. In the profile view they look a little like skyscrapers, albeit, with more than a little influence from the Tower of Pisa.

Magnification is provided to scale. The region of distortion connecting the focal region to its context is extreme. The line that forms the edge of the focus is stretched to make the full connection. This one directional stretch is an extreme distortion, however, its purpose is solely to provide visual support for cognitive integration. Any actual reading of the representation can still be done on the scaled only surface sections. It is to be hoped that navigational issues will be even less significant

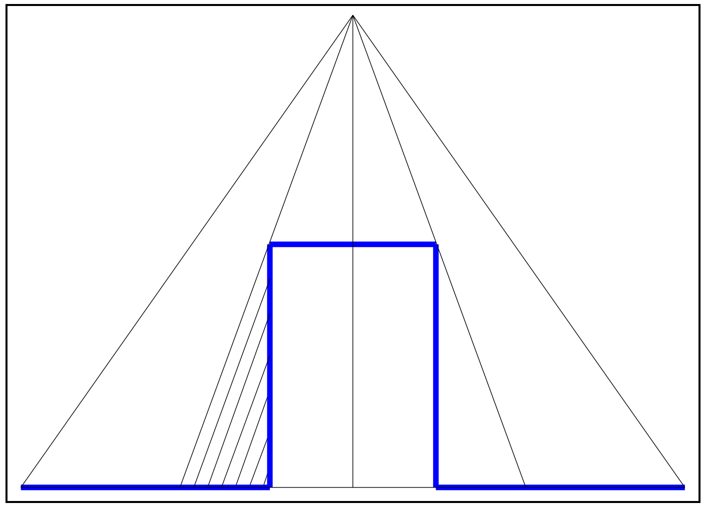

Figure 5.8: Cross-section diagram of a Manhattan lens. The surface connects the base plane to the magnified focal region

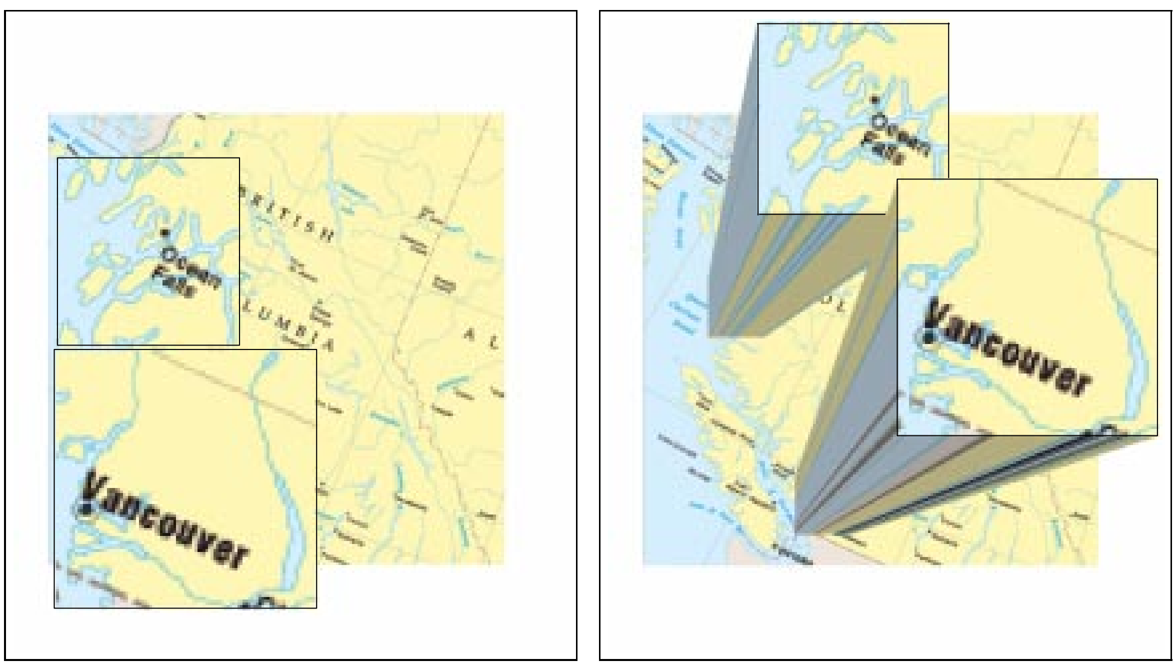

(a) Two Manhattan lenses, in place appearing as insets (for image credit see Appendix C.2)

(b) Folding the Manhattan lenses slightly shows the magnified regions are located in the map.

Figure 5.9: Using Manhattan lenses to magnify the coast line around Ocean Falls and Vancouver

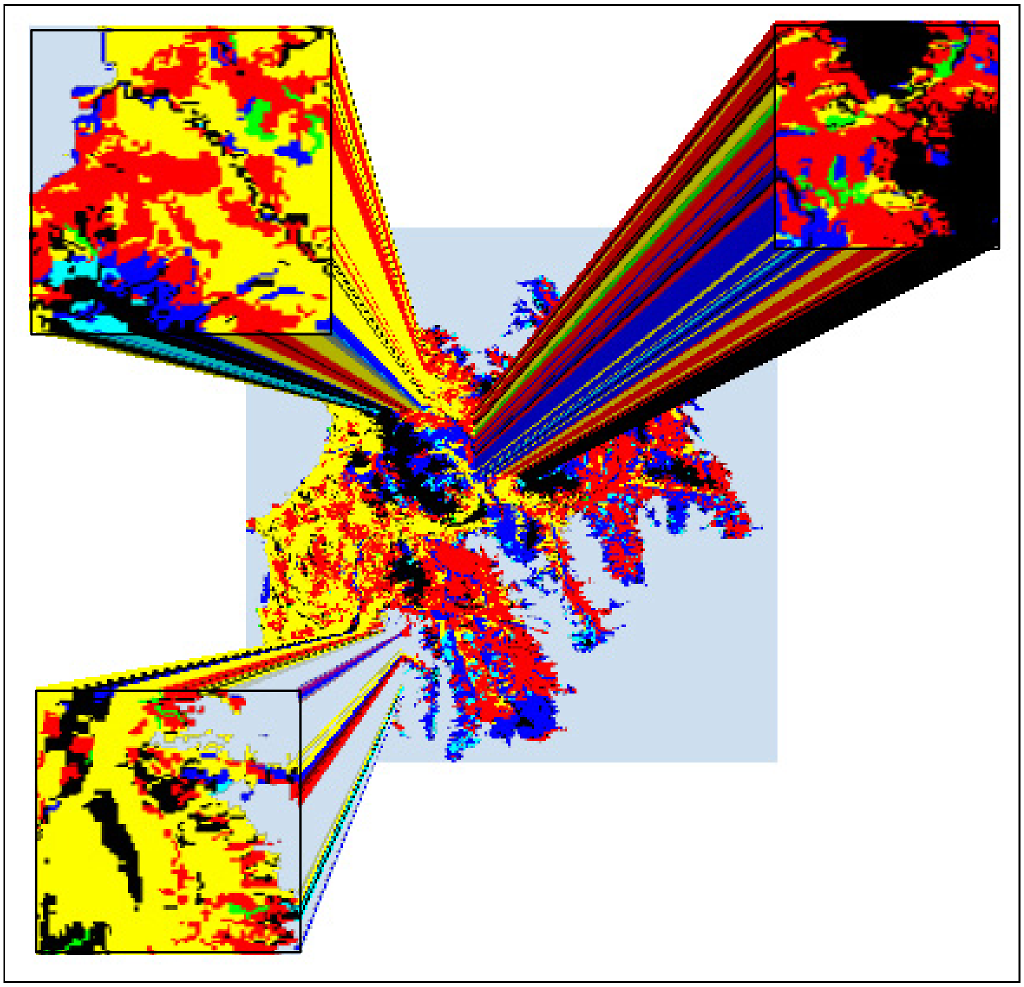

Figure 5.10: A folded view of the Manhattan lenses in Figure 5.7

than with simple insets in that visually connected views are possible. As roving search is possible while a lens is folded, the Manhattan lens can be tipped slightly, showing the connection on one or two sides and moved in that orientation. Figure 5.9(b) shows the two Manhattan lenses in Figure 5.9(a) folded slightly to reveal their connection to the rest of the map. They can be moved across the map in this orientation.

Since magnification is accomplished without effort to maintain context, the focal section can be magnified to fill the entire available display space. Here the limitations on magnification will depend on the initial resolution of the representation, the size of the available display space and the amount of the representation that has been chosen for the focus.

The advantages of Manhattan lenses are: magnification of focal region is to scale, freedom of re-positioning is provided, interactive visual connection is available. However, while adjacent context can be seen in a visually connected manner, not all of it can be seen simultaneously.

5.3 Non-Occluding Step Functions

One difficulty with insets and Manhattan lenses is local occlusion. Using EPS the projection rays from the RVP through the edge of the raised focus can be used to define the visible regions (Figure 5.11). Intersecting the projection rays with the base plane defines the edge between the occluded section and the visible area.

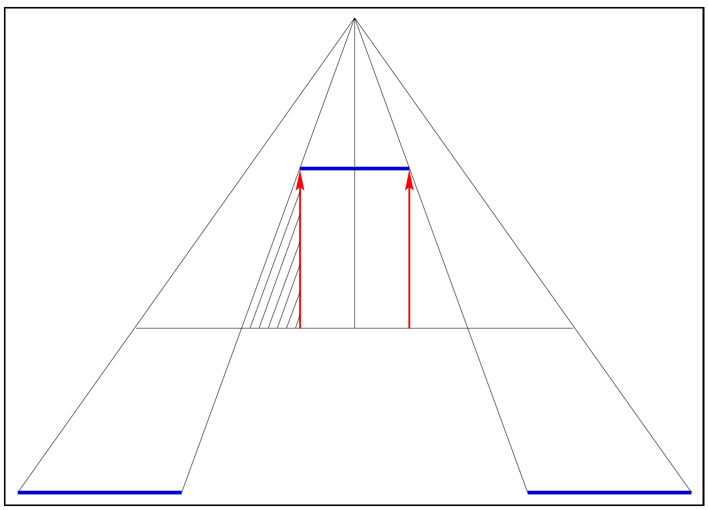

Figure 5.11: The profile view of a non-occluding disjoint step function

Occlusion can be eliminated by avoiding the regions that cannot be seen from the RVP. Once a selected focus is magnified the rest of the representation is compressed to fit within the non-occluded regions that remain. In EPS to compress the context, the distance between the context and the RVP is increased. As the focal regions and the context remain parallel to the view plane both magnification and compression are to scale. These are step functions in that different regions of the surface are treated differently.

However, a problem remains in creating a detail-in-context presentation using this approach. The focus and the context were initially of the same scale and size. As the focus is now larger and the context is now smaller they will no longer fit (Figure 5.12). Various layout approaches have been suggested in this regard.

- An *orthogonal stretch* as described in Chapter 2 has been used frequently in the literature (Bifocal Lens [147], Stretch Tools [139], Rubber Sheet (orthogonal method) [138]). The regions adjacent to the focus are stretched in order to fill the space (Figure 5.12(a)).

- In an *orthogonal separation*, gaps are allowed to form in the context rather than the additional distortion of the orthogonal stretch (Figure 5.12(b)). Disparate distortion in x and y is avoided, but the regions are visually disjoint. Some of the presentation space is used for maintaining positional organization of the separated regions instead of for presenting information. This introduces considerable white space. SHRiMP [150] uses this approach.

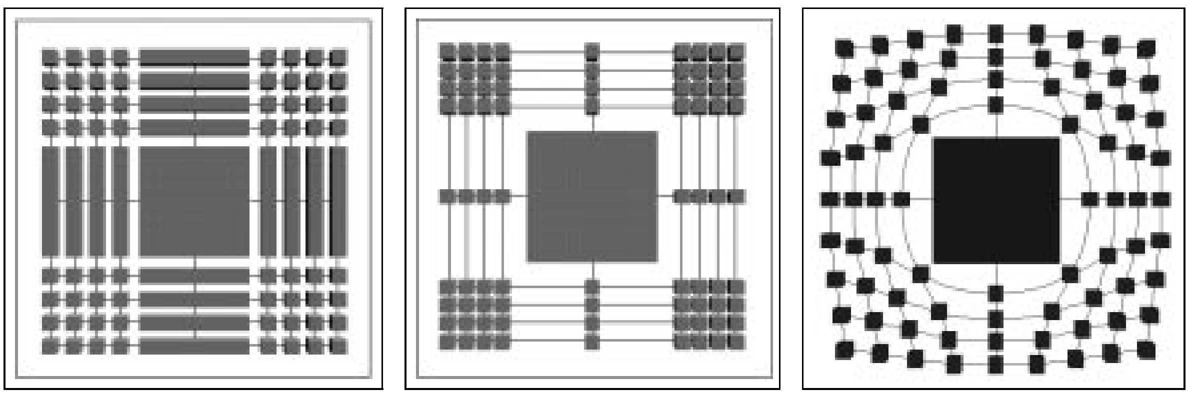

(a) Orthogonal stretch (b) Orthogonal separation (c) Radial separation

Figure 5.12: Different approaches to the focus to context fit problem

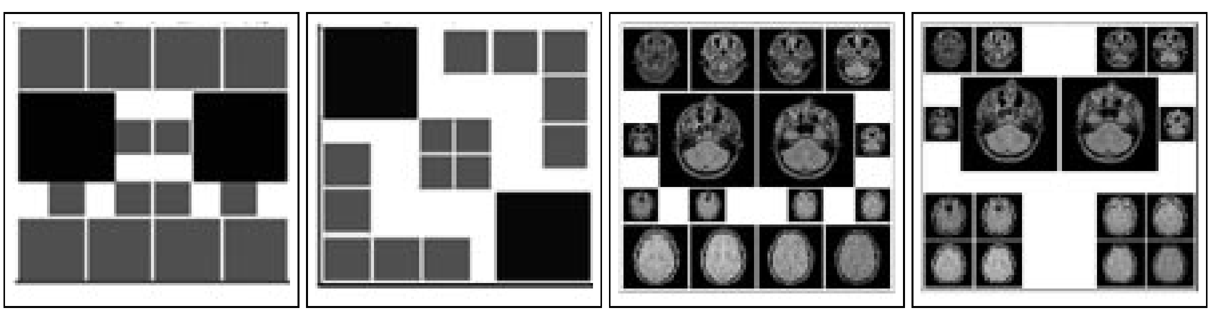

Figure 5.13: Orthogonal variations developed in response to the particular needs of presenting MRI images [166]

- Several *orthogonal variation* methods have been developed. These methods work similarly to an orthogonal separation but allow for more variations in the positioning of the context. The looser interpretation of orthogonal ordering reduces the white space slightly and often creates a more complex presentation. There are now several variations of this approach (Zoom family [7, 8, 45], MRI Presentations [166, 167]. MRI Presentations developed several variations to discover an appropriate balance between: keeping all the magnified foci at the same scale, minimizing the amount of white space, minimizing the variations in scale in the context, and maintaining orthogonality (Figure 5.13).

- A *radial separation* displaces the context radially (Figure 5.12(c)). While this pattern does not preserve the orthogonal relationships it does a better job of preserving proximity and does not cause such pronounced artificial clusterings [150].

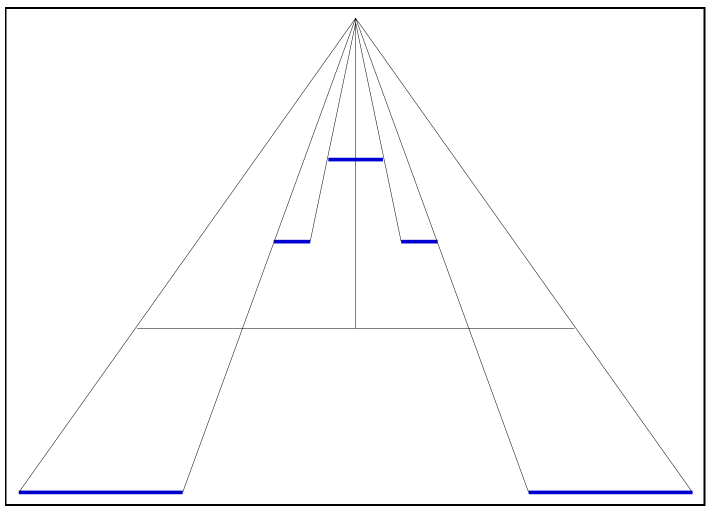

Figure 5.14: The profile view of a non-occluding disjoint step function

All these variations provide detail-in-context presentations with scaled-only foci and none of the context either occluded or removed. Along with the usual trade-off between the amount of magnification and compression, there is a trade-off between degree of distortion and amount of white space. There is another trade-off between orthogonal distortions, which maintain orthogonality and parallelism, and radial distortions, which better preserve proximity (for further discussion see Chapter 6). There is the possibility of extending the step function further as suggested in Figure 5.14 (also noted in [94]).

These step functions provide detail-in-context presentations from RVP but when viewed from the side, the foci are visually disjoint (just like insets). One solution is to stretch the surface along the connecting vector from the foci to the context as in Manhattan lenses (Figure 5.15). This has

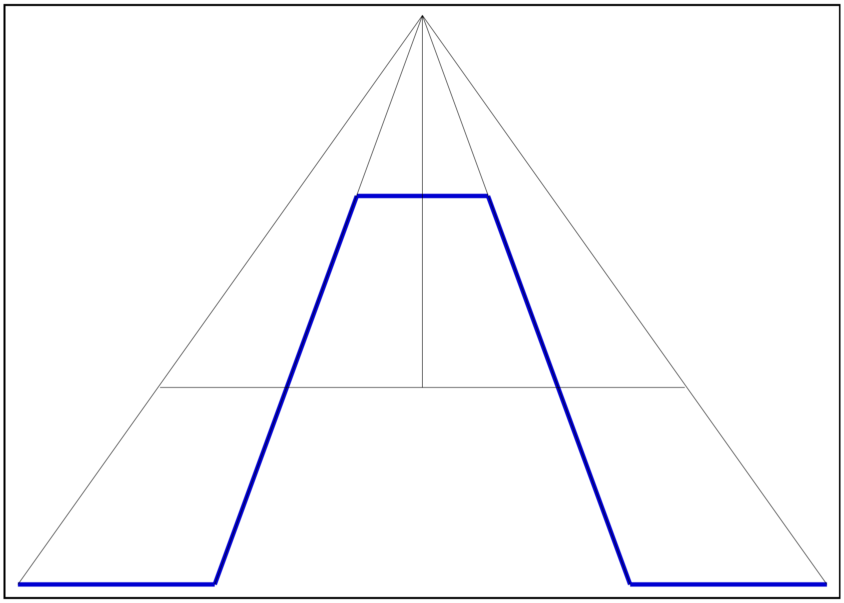

Figure 5.15: The profile view of a connected non-occluding step function

not been tried for at least two reasons. First, the stretched connection, as in the Manhattan lenses, will be seen when folding only. Second, the step function issue of focus-to-context fit exists for this method as well.

The advantages of these non-occluding step approaches are: focal regions are scaled only, all variations provide detail-in-context presentation, and there are several spatial organization options. The disadvantages are: providing a detail-in-context presentation places limits on the degree of magnification, there are several trade-offs between the various options in that compressed only contexts create white space and cost visual connectivity, and maintaining connectivity creates distortion.

5.4 Linear Drop-off Functions

While the non-occluding step functions discussed above provide detail-in-context views, they suffer from significant visual discontinuities as regions of scale change abruptly from focus to context. A smoother visual integration between focal regions and their context can be a desirable attribute [138], in that it is thought to provide a presentation that is more convincingly unified. It is thought that this will provide a better visual gestalt [35].

The foci and context can be visually connected by ensuring that the drop-off function is in the region that can be seen from the RVP. One approach is to use a linear drop-off function that connects

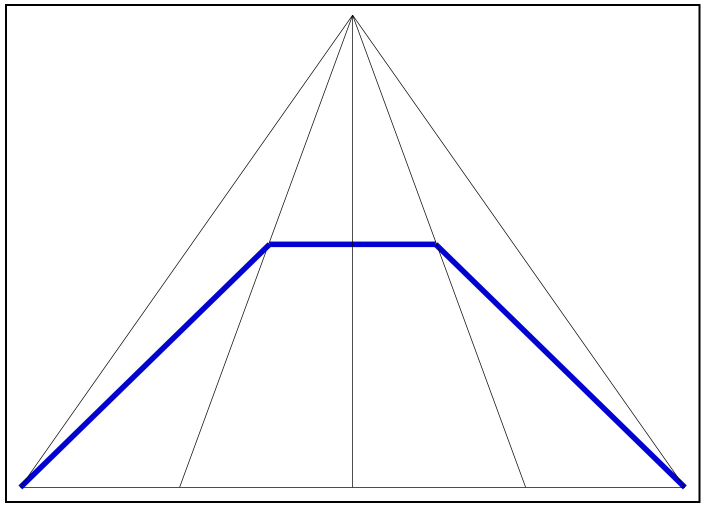

Figure 5.16: The profile view of a global linear drop-off function

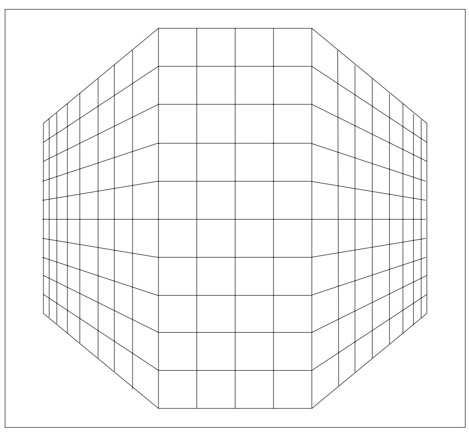

(a) Perspective Wall

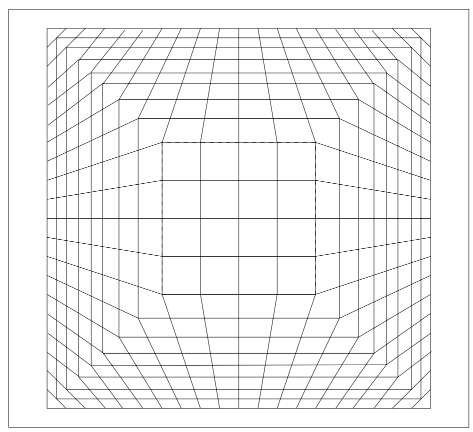

(b) Document Lens

Figure 5.17: Perspective Wall and Document Lens use orthogonal global linear drop-off functions

the edge of the focal region to the edge of the context (Figure 5.16). This creates a *global* detail-in-context presentation, in that the distortion is spread throughout the context. In EPS terms both Perspective Wall [99] and Document Lens [132] use orthogonal global linear drop-off functions (Figure 5.16). Alternatively the effect of the drop-off function can be *constrained*[1], keeping some

[1] Although lenses that did not extend globally were first introduced in 3DPS [20], it was in Non-Linear Views [84] that the word constrained was first used to describe them

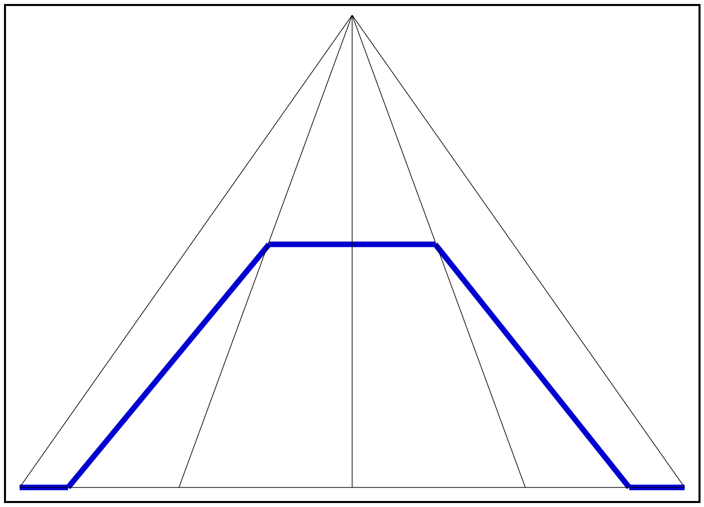

Figure 5.18: The profile view of a constrained linear drop-off function

of the context undistorted. Figure 5.18 shows the profile of a constrained linear lens with a scaled focal region and Figure 5.19 shows top views of a constrained linear lens with a point focus. In Figure 5.19(a) the grid shows the structure of the gradually increasing compression in the distorted region and the sharp connection to the context and Figure 5.19(b) uses this lens to magnify Prince George, BC.

In an EPS linear drop-off function the surface height h_p of a point p, that is a distance d_p from the focal centre f_c with a lens radius l_r and a focal height h_f is calculated by:

$$h_p = h_f \cdot (1 - (d_p/l_r)) \tag{5.1}$$

Varying the lens radius l_r affects the limits of the lens and the resulting slope. Figure 5.20 shows side views of a variety of possible slopes and Figure 5.21 shows the corresponding top views. Keeping the lens radius constant and varying the focal height h_f changes the magnification (Figure 5.22). Note how increasing magnification causes increasing compression just before the connection to the context. Figure 5.23 illustrates limiting the focal height to provide scaled-only focal regions.

These lenses have all the functionality described in Chapters 3 and 4 with a linear drop-off function. Again choice of which linear slope to use involves trade-offs. Keeping the slope as steep as possible provides much un-distorted context but a very compressed visual connection. Extending the linear function to the edges of the surface minimizes the distortion and compression but leaves less or no untouched context.

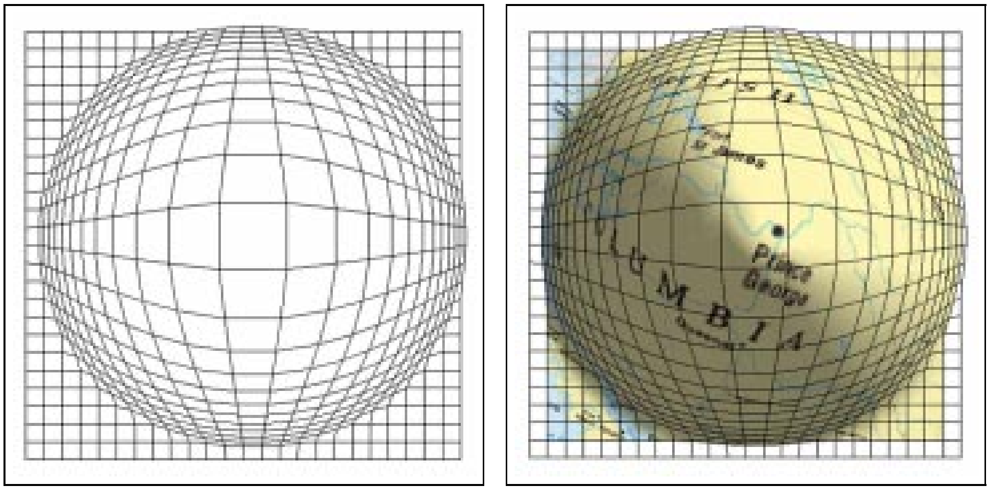

(a) The view from the RVP of a grid showing the pattern of compression of a linear lens

(b) Prince George, BC magnified with point focus and constrained linear drop-off

Figure 5.19: The constrained linear drop-off function

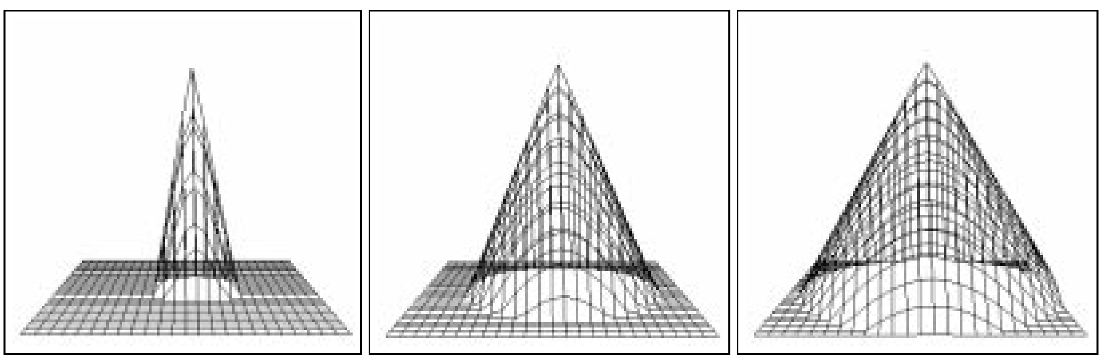

Figure 5.20: Linear slope sides views

The linear drop-off function provides both compression necessary to make space for the magnified focus and the distortion necessary to make the visual connection between the focus and the context. While the different regions of the representation will be connected visually, these connections tend to be angular and as a result quite noticeable. This can be an advantage in that they define visually regions of differing scale and distortion. Noticeable visual connections can also be a disadvantage in that these are abrupt changes. The distinction between radial and orthogonal patterns is available by varying the distance metric (see Chapter 4). However, as the compression varies with the distance from the focal point, maintaining a context that does not cause regions of extreme

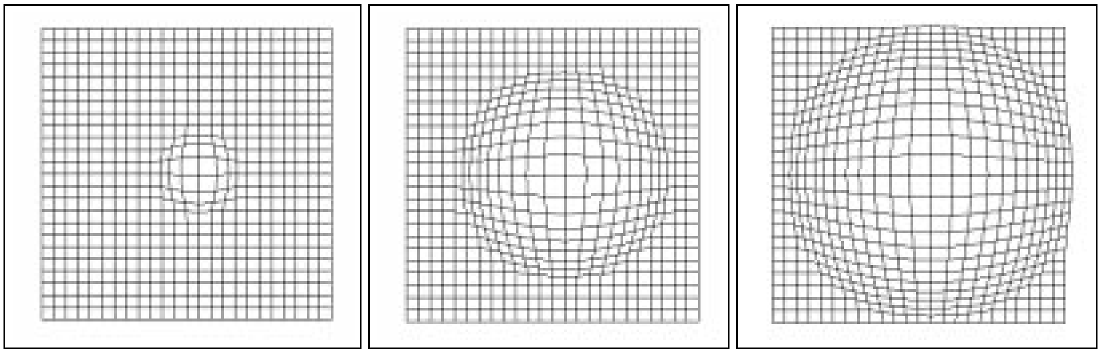

Figure 5.21: Linear slope top views with varying lens radius

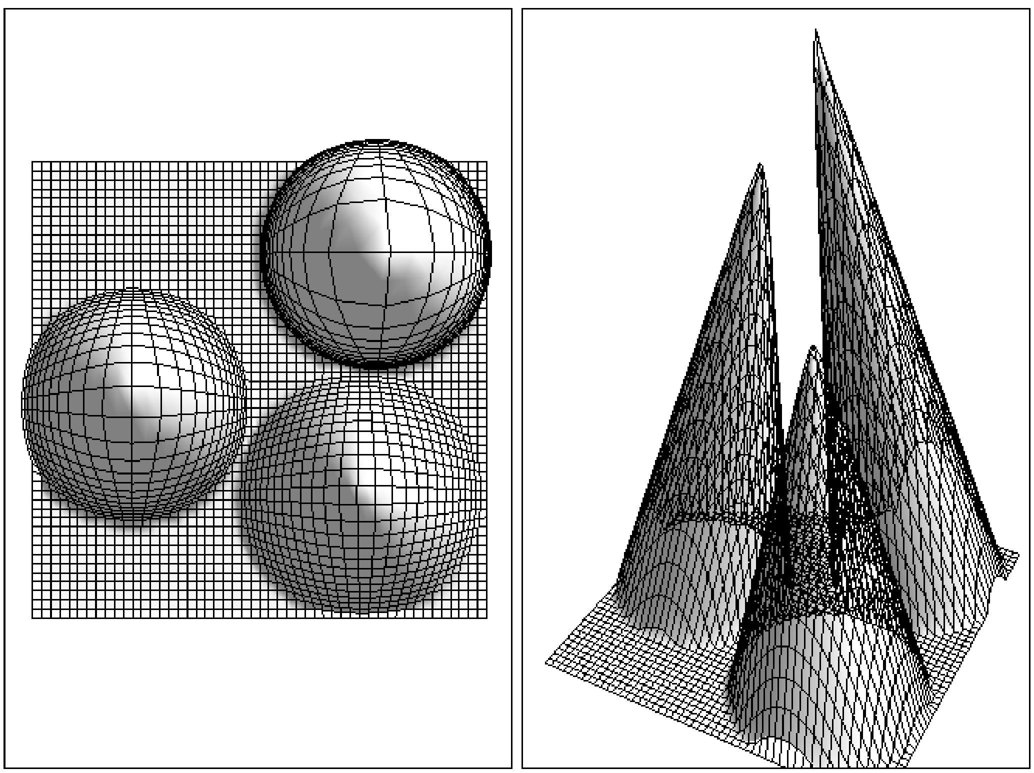

Figure 5.22: Constrained linear drop-off functions with the same lens radii and varying magnification

compression places more limits on the amount of magnification that can be provided.

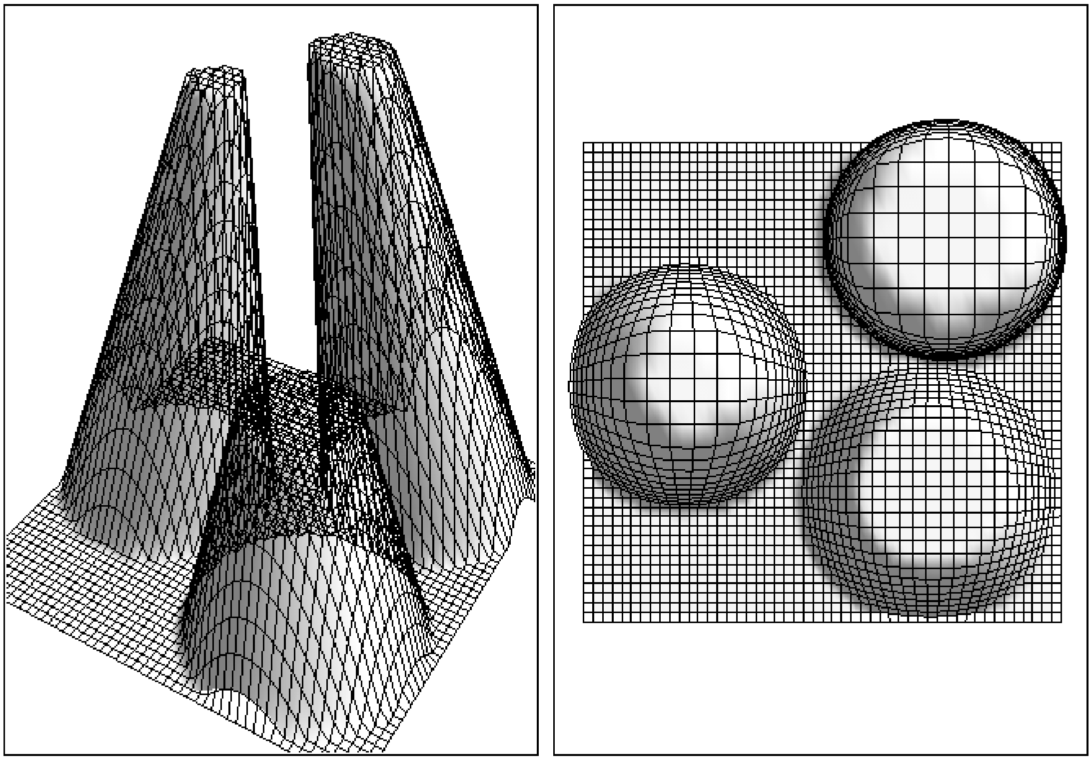

Figure 5.23: Constrained linear drop-off functions with scaled-only focal regions

5.5 Non-Linear Drop-off Functions

While linear drop-off functions provide a visual connection between the focus and its context they also create sharp visual transitions. Non-linear drop-off functions make it possible to combine focal magnification with gradual integration into the context. A great variety of mathematical functions can be used in this manner, each with its own characteristic pattern of distortion. The following discussion is illustrated with a few of the possible functions: hemisphere (Figure 5.24), cosine (Figure 5.25), hyperbola (Figure 5.26) and Gaussian (Figure 5.27).

A hemisphere drop-off function (Figure 5.24) has a very gradual initial drop-off that increases rapidly towards the edge of the lens and meets the context perpendicularly. This causes the context adjacent to the focus to be almost as magnified as the focus, and results in some occlusion at the connection of the lens to the context. Minimizing the occlusion severely limits the amount of focal magnification. The characteristics of hemisphere are: limited focal magnification, good visual integration from the focus into its immediate surroundings, when constrained there is an abrupt visual transition or discontinuity where the lens meets the context, and the edges of the region of distortion may be occluded or reversed. It has been suggested that the familiarity of the hemisphere may aid in readability [137].

The cosine function provides a simple drop-off function (Figure 5.25) with slightly more moderate magnification of regions adjacent to the focus and a more gradual connection to the context. The cosine has moderate focal magnification and good visual integration from the focus into its immediate surroundings, although less magnification of the regions adjacent to the focus than the hemisphere. The slope of the curve towards the edges of the distorted region is more gradual than the hemisphere, spreading the compensating compression more throughout the distorted region. However, as magnification is increased the compression builds at the connection to the context. If the cosine is constrained then there is still an abrupt visual transition where the distortion meets the context.

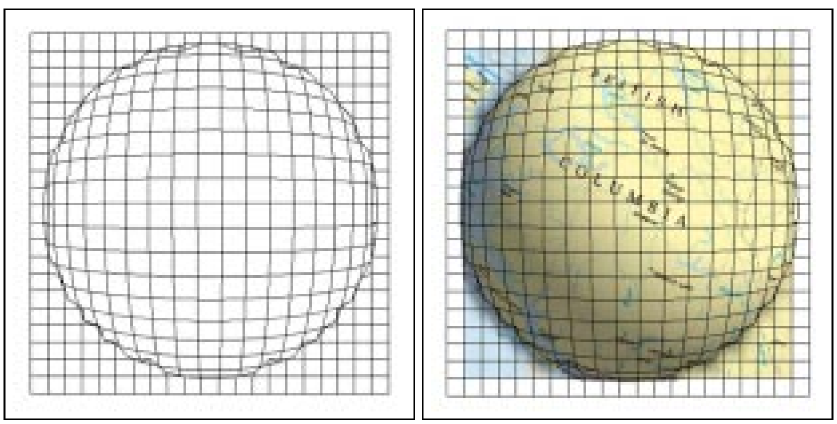

Figure 5.24: The hemisphere drop-off function

The characteristics of the hyperbolic drop-off function (Figure 5.26) are similar to those of the cosine drop-off namely: moderate focal magnification, good visual integration from the focus into its immediate surroundings, and a fairly gradual connection to the context. With a global hyperbola it would be possible to adjust the asymptotes with respect to the view volume in order to spread the compression more evenly in the distorted region. However, the compression would still become extreme at the edges. As the hyperbola is constrained, there is an abrupt visual transition where the distortion meets the context.

The characteristics of the Gaussian drop-off (Figure 5.27) are discussed in detail in Chapter 4. The Gaussian with its characteristic bell shape provides a good basis for constrained lenses. It combines the advantages of gentle focal integration with those of gradual integration into the remaining

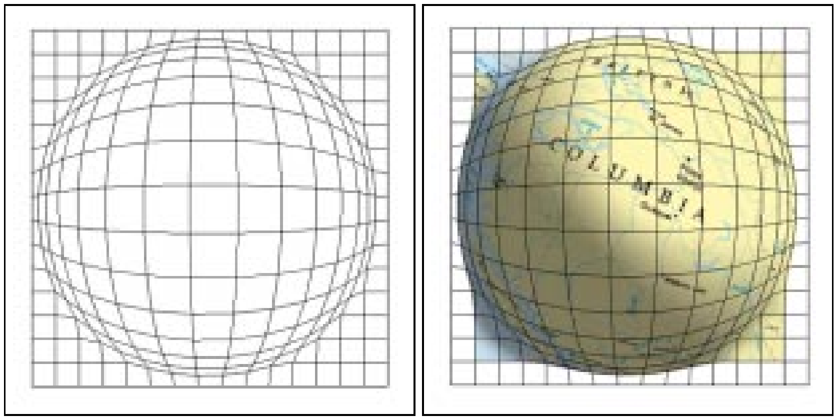

Figure 5.25: The cosine drop-off function

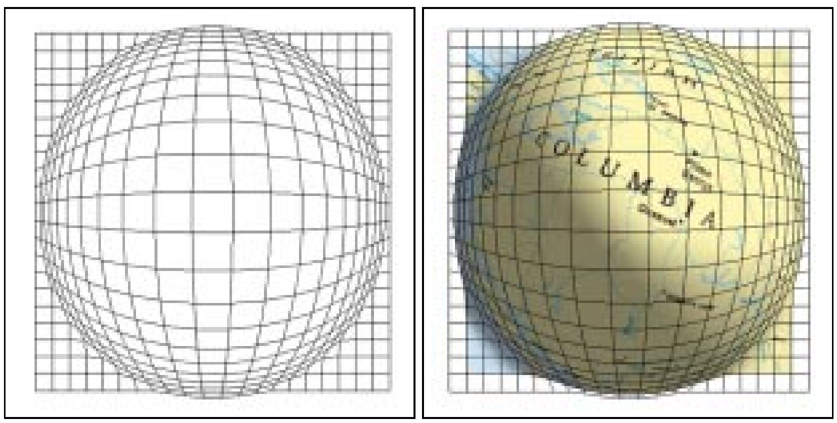

Figure 5.26: The hyperbolic drop-off function

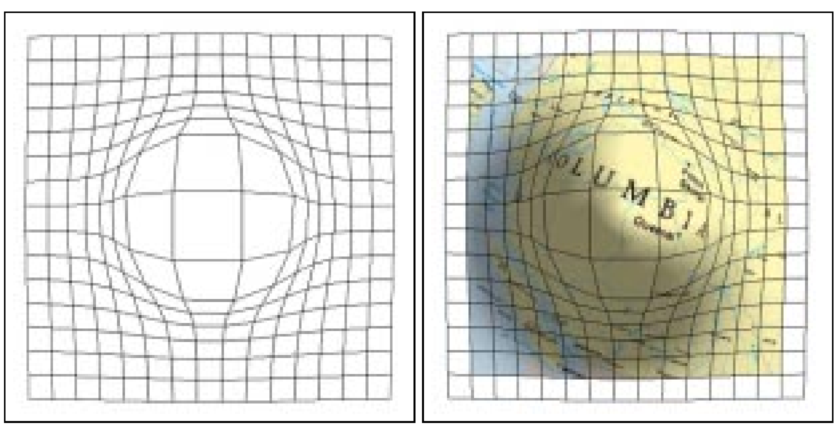

Figure 5.27: The Gaussian drop-off function

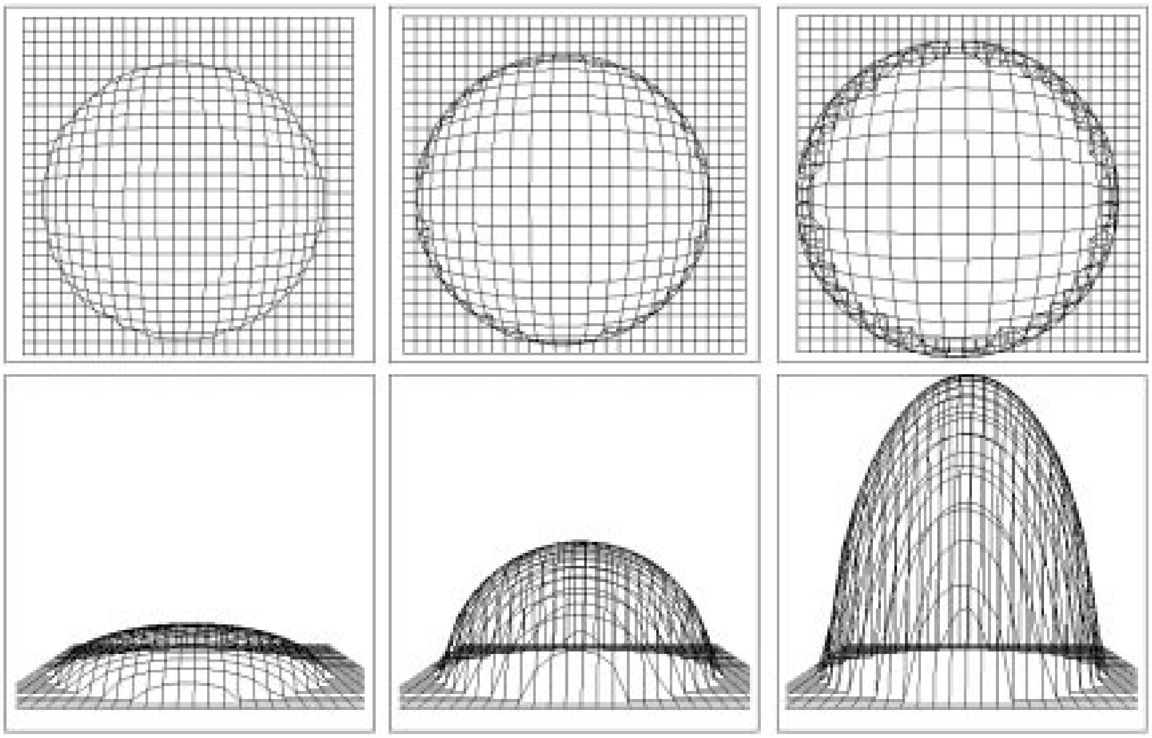

Figure 5.28: The profile views of hemispherical functions

context. However, like all of the functions just discussed, it also has an area of maximum compression at the point of inflection. Its characteristics are: great focal magnification, good visual integration from the focus into its immediate surroundings, good magnification of adjacent context, and good visual integration from the distorted region into the context. The area of maximum compression is located between the visual integration of the focus and visual integration into the context. The exact location of the area of maximum compression can be shifted be adjusting the standard deviation. Even if the Gaussian is used globally, the edges of the context are more preserved than with the other curves discussed. If the Gaussian is constrained, there are no abrupt visual transitions.

For all of these drop-off functions, scaling the distance d_p locates the edge of the lens, and adjusting the focal height h_f sets the degree of magnification, providing a range of related lenses for each drop-off function. For example, Figure 5.28 adjusts the hemisphere by changing the focal height. The basic characteristics hold throughout this range of lenses, though as the series shows increasing the focal magnification increases the reversal at the connection to the context.

For comparison the basic characteristics are tabulated in Table 5.1 and the top views of the distortion patterns are presented (Figure 5.30) below their drop-off functions (Figure 5.29). The

	Linear	Hemisphere	Cosine	Hyperbola	Gauss
focal magnification	good	minimal	fair	fair	great
adjacent focal magnification	minimal	great	good	good	moderate
focal integration	sharp	good	good	good	good
location of max compression	context-connection	context-connection	context-connection	context-connection	mid-distortion
context integration	abrupt	poor	abrupt	abrupt	good

Table 5.1: Comparing drop-off functions

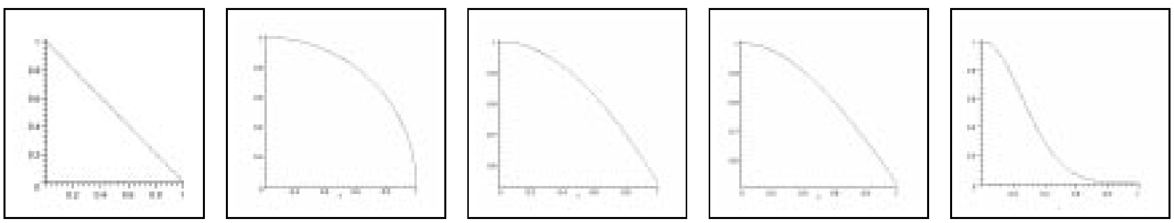

(a) Linear drop-off function (b) Hemisphere drop-off function (c) Cosine drop-off function (d) Hyperbolic drop-off function (e) Gaussian drop-off function

Figure 5.29: Graphs of the drop-off functions

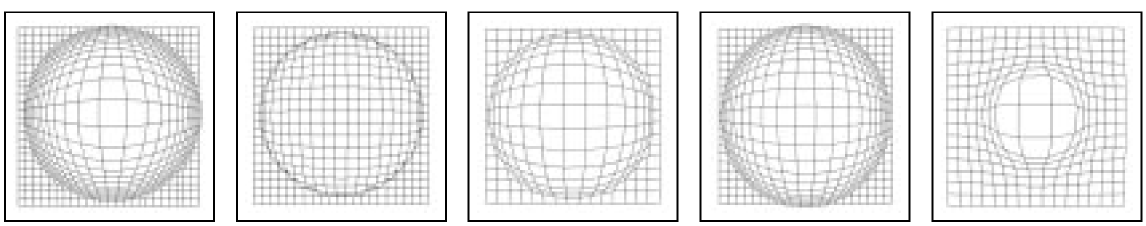

(a) Linear drop-off function (b) Hemisphere drop-off function (c) Cosine drop-off function (d) Hyperbolic drop-off function (e) Gaussian drop-off function

Figure 5.30: For comparison, top views of the different lenses

linear drop-off function of Section 5.4 in included as it results in non-linear magnification due to perspective projection. The general characteristics of non-linear drop-off functions are a magnification/compression trade-off and a trade-off between the spread of the distortion and the degree of compression.

Any mathematical function can be used, however as we develop an understanding of the relationship between curvature and presentation pattern we may be able to make more appropriate choices to suit a particular representation or task. In the literature a lot of attention has been paid to the importance of local context and the nature of the focal connection. Other factors such as the nature of the compression, the location of maximum compression, and the nature of the context connection may also be important in different situations.

5.6 Comparing EPS Lenses with Other Methods

There are many methods for adjusting the presentation of two-dimensional representations to provide detail-in-context presentations (see Chapter 2). There is a broad distinction between 2D-based transformations and transformations that make use of 3D to create 2D presentations. Most methods [80, 84, 108, 137, 138] are 2D-based in that they make use of transformation functions that are applied in the two-dimensional plane of the representation. A few methods [20, 99, 132] are 3D based in that they manipulate the two-dimensional representation in three-dimensions and then apply perspective projection. EPS is a 3D based method that includes the presentation variations offered by the other 3D based methods.

The 2D based methods create a new presentation by spatially adjusting a given two-dimensional presentation to create another two-dimensional presentation. A 2D transformation function $T_{2D}(d_p)$ where d_p is a distance in the x, y plane between a point p and the focal centre f_c, performs the adjustments. The resulting magnification pattern $M_{2D}(d_p)$, is the derivative of $T_{2D}(d_p)$.

EPS as a 3D based method is quite different algorithmically. The plane or surface that holds the two-dimensional representation is manipulated in three dimensions, then viewed through perspective projection. This separates the transformation function T_{EPS} into two distinct steps; a surface manipulation function and perspective projection.

First we show that the EPS approach can produce 2D results. Then we compare two typical 2D-to-2D transformation methods: Sarkar and Brown's Graphical Fisheye [137], and Keahey and Robertson's Non-Linear Views [84] with two EPS lenses: the linear drop-off lens and the Gaussian lens of 3DPS [20]. This comparison shows that the variations in resulting presentations are like the variations between lenses with different drop-off functions. Then we discuss the advantages of using EPS.

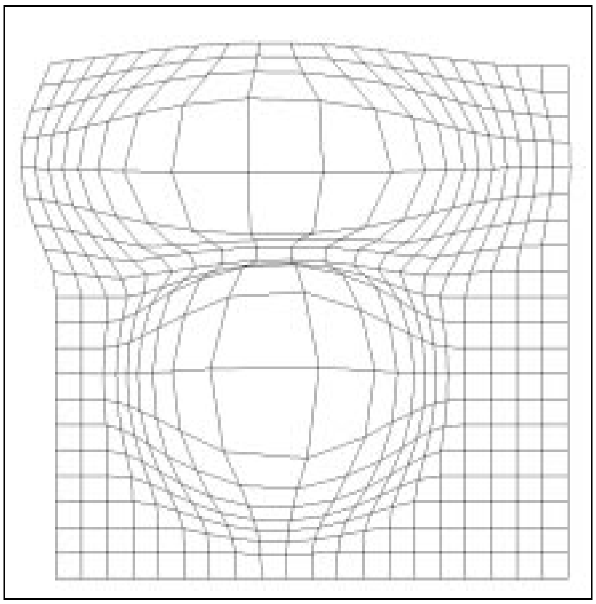

(a) The top view of a two focal grid

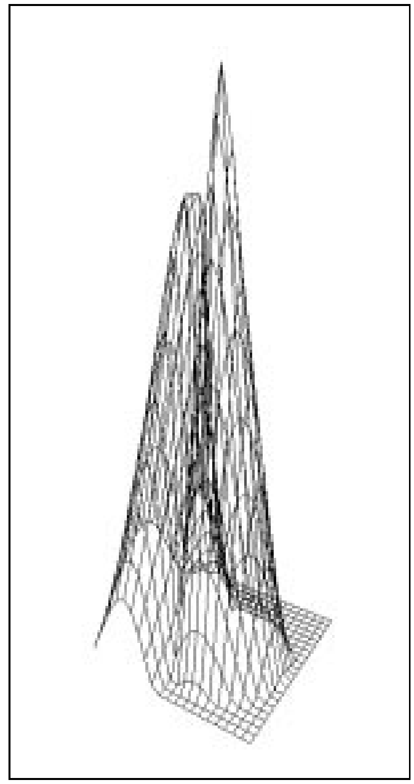

(b) The 3D form of this grid

Figure 5.31: The transformed grid

5.6.1 Using EPS to Obtain 2D Results

EPS can be used to generate 2D-to-2D transformation functions. Recall from Section 3.5 that given a magnification factor we can obtain the focal height h_f and from the focal height and the drop-off function we can calculate the surface manipulation. For a point p that is a distance d_p from the focus and a standard deviation σ of a Gaussian drop-off function, the surface height h_p is given by:

$$h_p = h_f \cdot \exp^{-\frac{(d_p)^2}{\sigma}} \tag{5.2}$$

From the manipulated surface location of a point p (x_i, y_i, h_p) and the distance d_b from the RVP to the base plane we can obtain the apparent x_m and y_m location on the base plane:

$$x_m = d_p \cdot \frac{x_i}{(d_b - h_p)} \qquad and \qquad y_m = y_i \cdot \frac{d_b}{(d_b - h_p)} \tag{5.3}$$

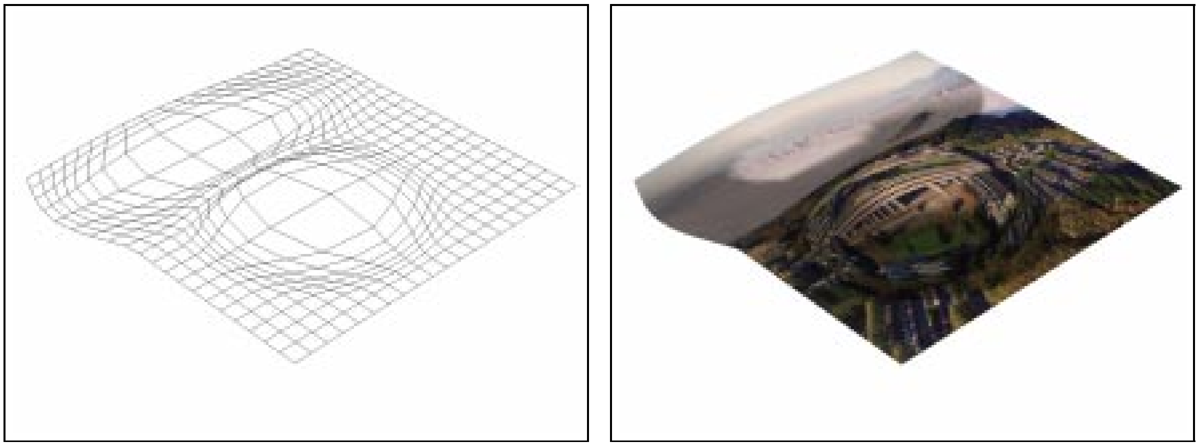

(a) The grid in 2D

(b) The representation applied to the grid in 2D (aerial photograph, R. Long, IMC SFU)

Figure 5.32: The applying the surface to the 2D transformed grid

Since from d_b and the RVP we know the z location of the base plane we can place the point on the base plane z_b at (x_m, y_m, z_b).

This gives two equivalent presentations within EPS. One is the view from RVP of the 3D manipulated surface and the other is a transformed 2D presentation. Figure 5.31(a) shows a manipulated grid with two focal regions. Figure 5.31(b) shows this grid from the side, revealing its 3D form. However, this grid can be transformed to 2D using Equation 5.3 (Figure 5.32(a)) as can the surface (Figure 5.32(b)). Figure 5.33 shows the transformed presentation. For implementation the surface normals can be preserved to maintain the shading in 2D. When viewed from RVP the 2D or 3D images appear equivalent.

This two step process can be combined into a single step by substituting h_p in Equation 5.3 with the Equation 5.2. This gives 2D-to-2D transformation function that is based on EPS's Gaussian drop-off:

$$T_{EPS:Gaus} = x_i \cdot (d_b/(d_b - (h_f \cdot \exp^{-\frac{(d_p)^2}{\sigma}})) \tag{5.4}$$

and a corresponding magnification function:

$$M_{EPS:Gaus} = \frac{d_b}{d_b - h_f \cdot \exp^{-(\frac{x^2}{\sigma})}} - 2 \cdot \frac{x^2 d_b h_f \exp^{-(\frac{x^2}{\sigma})}}{(d_b - h_f \cdot \exp^{-(\frac{x^2}{\sigma})})^2 \cdot \sigma} \tag{5.5}$$

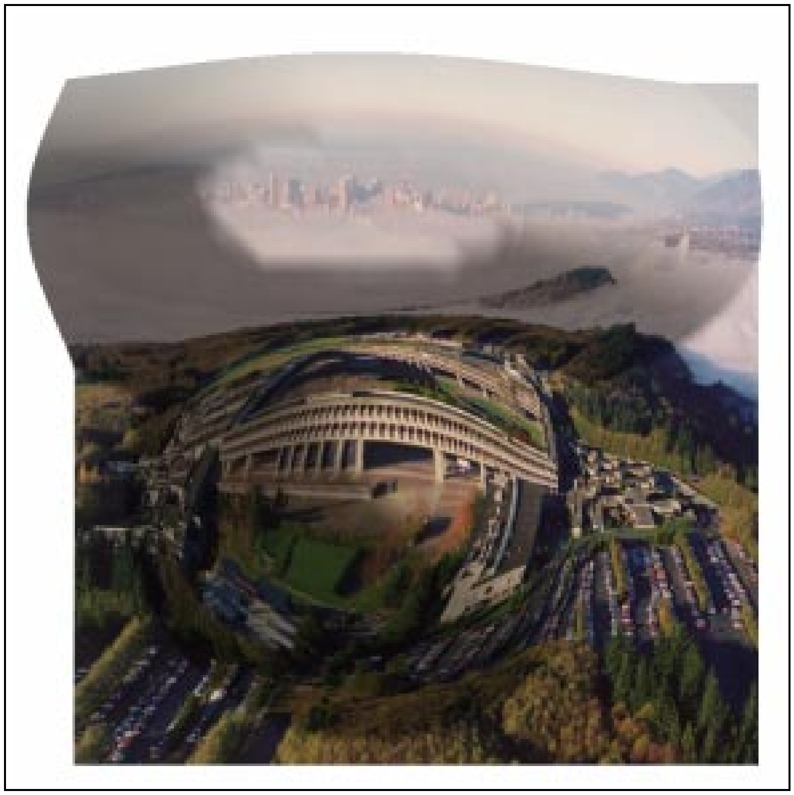

Figure 5.33: Example of a multi-scale view. The three-dimensional surface manipulation that organizes the presentation has been projected back into the base plane with equation 5.3. The foci zoom in on downtown Vancouver and the academic quadrangle of the SFU campus.

For the values $d_b = 2$, $h_f = 1$ and $\sigma = 0.1$, Figure 5.34(a) shows the graph of the Gaussian drop-off function, Figure 5.34(b) shows the EPS derived Gaussian transformation function, and Figure 5.34(c) shows the resulting magnification function.

Replacing the Gaussian drop-off function with any other drop-off function can generate 2D-to-2D transformation functions. This can be useful in that these types of presentations can be created without implementing a full 3D environment.

5.6.2 Comparing Results

We compare an EPS linear drop-off and an EPS Gaussian drop-off with two typical 2D-to-2D transformation functions, Sarkar and Brown's Graphical Fisheye [137] (*SB:fisheye*) and Keahey and Robertson's Non-Linear Views [84](*KR:nonlinear*).

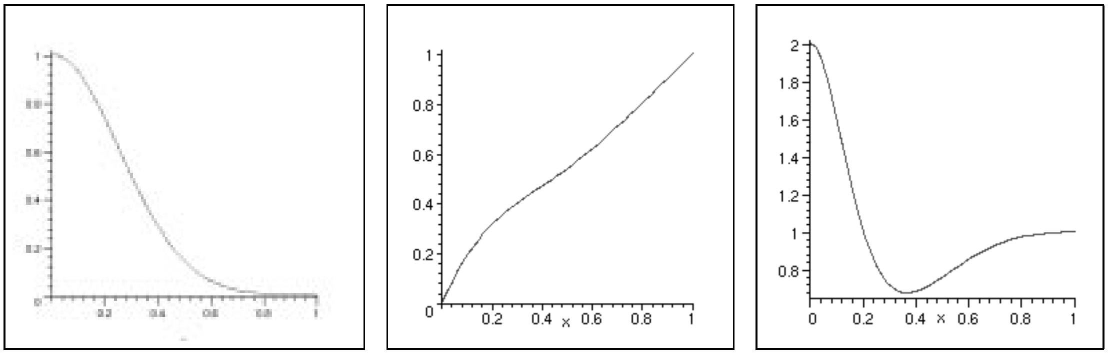

(a) Gaussian drop-off function (b) $T_{EPS:Gaus}$ (c) $M_{EPS:Gaus}$

Figure 5.34: The Gaussian drop-off function and the related Gaussian 2D-to-2D functions

Sarkar and Brown's Graphical Fisheye is:

$$T_{2D:SB:fisheye}(x) = \frac{(m+1)\cdot x}{((m\cdot x)+1)} \tag{5.6}$$

In this fisheye approach adjusting m changes the magnification. Keahey and Robertson's transformation function is:

$$T_{KR:nonlinear}(x) = tanh(m\cdot x) \tag{5.7}$$

where adjusting m also changes the magnification.

The EPS linear drop-off function is:

$$h_p = h_f \cdot \left(1 - \frac{d_p}{l_r}\right) \tag{5.8}$$

where varying h_f adjusts the magnification, varying l_r adjusts the lens radius, and d_p corresponds to the distance from the focus. For comparison d_p is set to x and h_f and l_r are set to 1. Combining surface manipulation of the drop-off function (Equation 5.8) and perspective projection function (Equation 5.3) creates a 2D-to-2D EPS transformation function:

$$T_{EPS:linear}(x) = x\cdot(d_b/(d_b-(1-x))) \tag{5.9}$$

The 2D-to-2D Gaussian transformation function is shown in Equation 5.4. The graphs of these four transformation functions are shown in Figure 5.35. They have all been set to have the same initial magnification factor for ease of comparison. Note that $T_{SB:fisheye}$ and $T_{EPS:linear}$ are coincident.

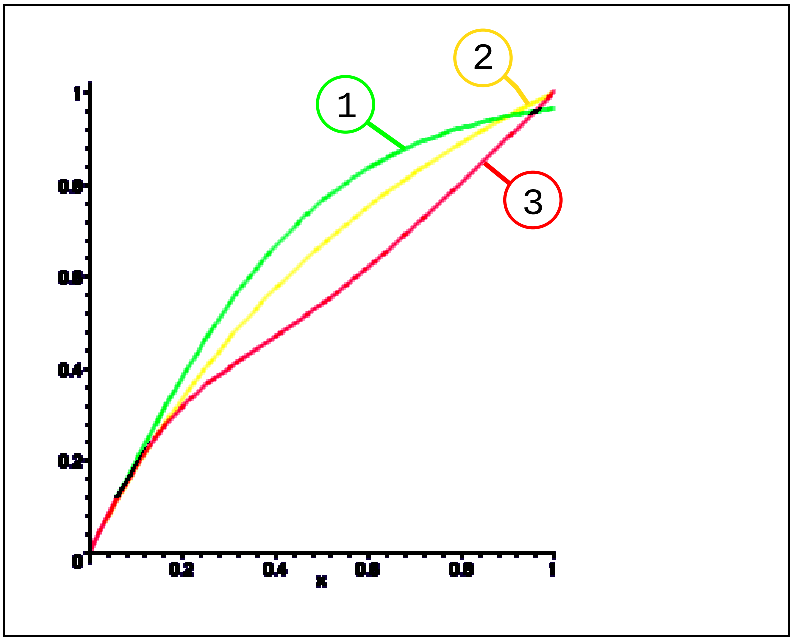

Figure 5.35: The graph of the transformation functions: (1) $T_{KR:nonlinear}$, (2) $T_{SB:fisheye}$ and $T_{EPS:linear}$, and (3) $T_{EPS:gaus}$

Sarkar and Brown's magnification function is:

$$M_{SB:fisheye}(x) = \frac{(m+1)}{((m \cdot x)+1)} + (m \cdot x)\frac{(m+1)}{((m \cdot x)+1)^2} \tag{5.10}$$

Keahey and Robertson's magnification function is:

$$M_{KB:nonlinear}(x) = m \cdot (1 - \tanh(m \cdot x)^2) \tag{5.11}$$

The EPS linear magnification function is:

$$M_{EPS:linear}(x) = \left(\frac{d_b}{(d_b - 1 + x)}\right) - \left(\frac{d_b x}{(d_b - 1 + x)^2}\right) \tag{5.12}$$

The EPS Gaussian magnification function is given in Equation 5.5.

Figure 5.36 shows the graphs of these magnification functions all set to have a a magnification factor of 2. The magnification functions of SB:fisheye and EPS:linear are the same, as are their transformation functions.

Similar visual results can be achieved in either 2D-to-2D or EPS approaches. However, there are several advantages with EPS. As noted in Section 3.5, given a desired degree of magnification, it is

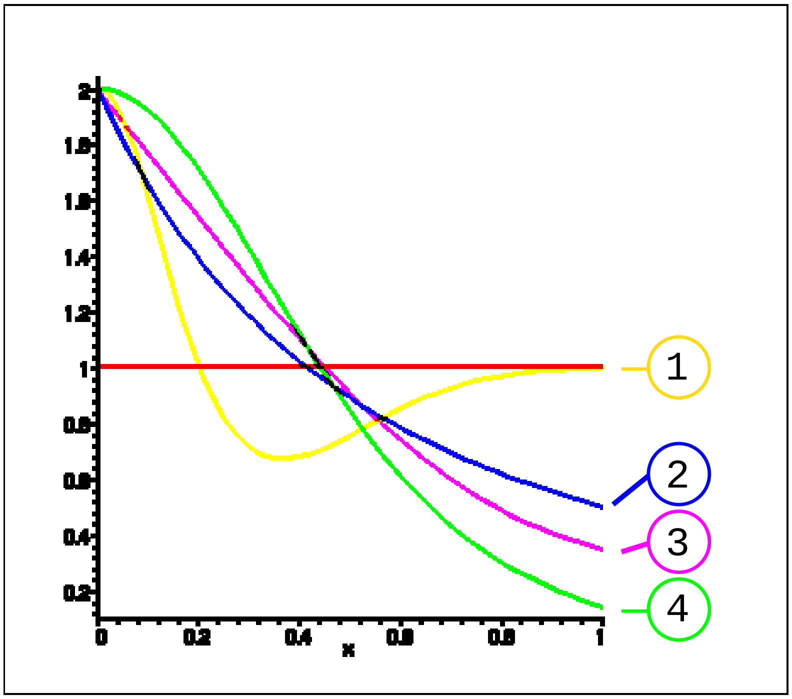

Figure 5.36: The graph of the magnification functions; (1) $M_{EPS:gaus}$, (2) $M_{SB:fisheye}$ and $M_{EPS:linear}$, (3) $M_{EPS:cosine}$, and (4) $T_{KR:nonlinear}$

difficult in the 2D-to-2D paradigm to obtain an appropriate transformation function. This has been extensively discussed in [85, 94]. Keahey and Robertson [85] developed an approximate transformation. They start with a grid and a set of desired magnification amounts. The grid is then adjusted iteratively, ensuring no grid points overlap until the difference between the magnification provided by the adjusted grid and the desired magnification is sufficiently small. In contrast, in EPS using a specified degree of focal magnification as input is simple and precise (Section 3.5).

Furthermore, finding a reverse mapping in a 2D-to-2D method is difficult [138]. In EPS with the surface position h_p one can obtain the original location from the transformed location (see Section 3.5). If EPS is used to generate a 2D-to-2D method and this method is used as a one step method bypassing obtaining h_p, then the difficulty of reversal also arises. Blending (see Section 3.7) complicates the situation and has not been investigated. However, as the drop-off functions are associated with a lens, reversing the changes in the focal magnification also reverses the effect of the distortion. This provides the ability to return to previous configurations.

Separating the transformation function into two distinct steps (surface manipulation and perspective projection) simplifies the mathematics. The magnification factor, the height of the surface, and the apparent transformation have a simple mathematical relationship based on similar triangles. Changing between two- and three-dimensional presentations is achieved by applying Equation 5.3. Note that if our method only supplied us with the transformed coordinates x_m and y_m, as with other 2D transformations, the same situation would result. The magnification factor would be equivalently difficult to retrieve. However, as h (or the z coordinate) of each point is known, the relationships remain as described above.

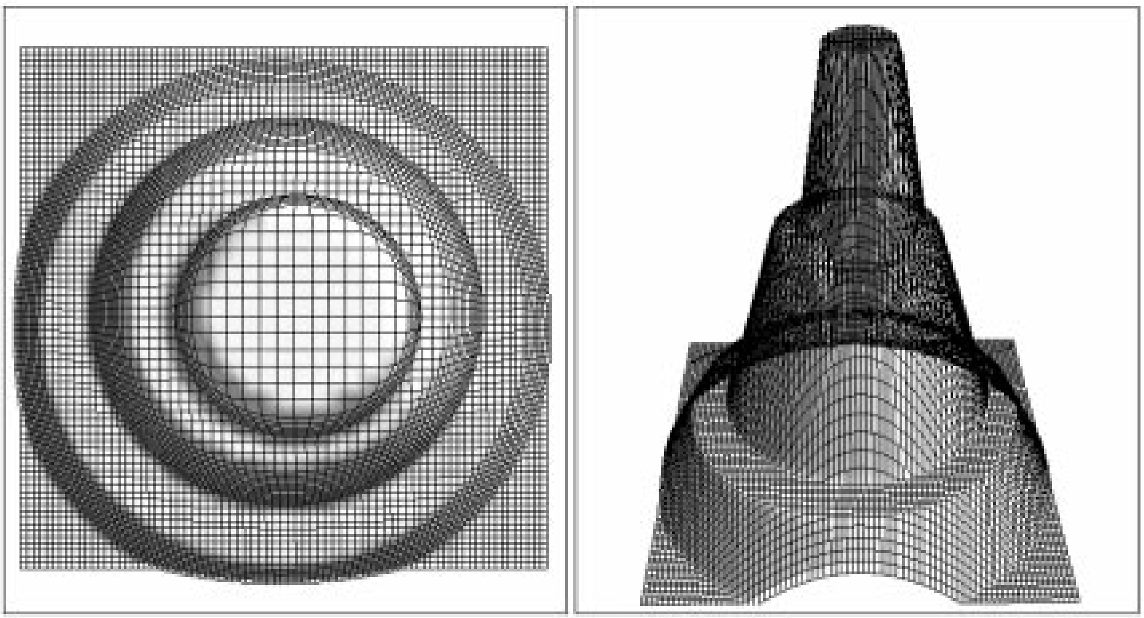

Figure 5.37: Step pyramid

5.7 Drop-off Functions in Combination

Individual drop-off functions have characteristics patterns of magnification and compression. However, there is no reason to use any one drop-off function exclusively. In fact, there are many ways in which they can be usefully combined.

Piecewise Combinations. Many of the lenses discussed in Section 5.4 combined a region of zero drop-off around the focus with a linear drop-off function in the region of distortion. Without this zero drop-off region point foci would literally be points. Combining linear functions with additional regions of zero-drop-off can create step pyramid presentations (Figure 5.37 and Figure 5.14). Two

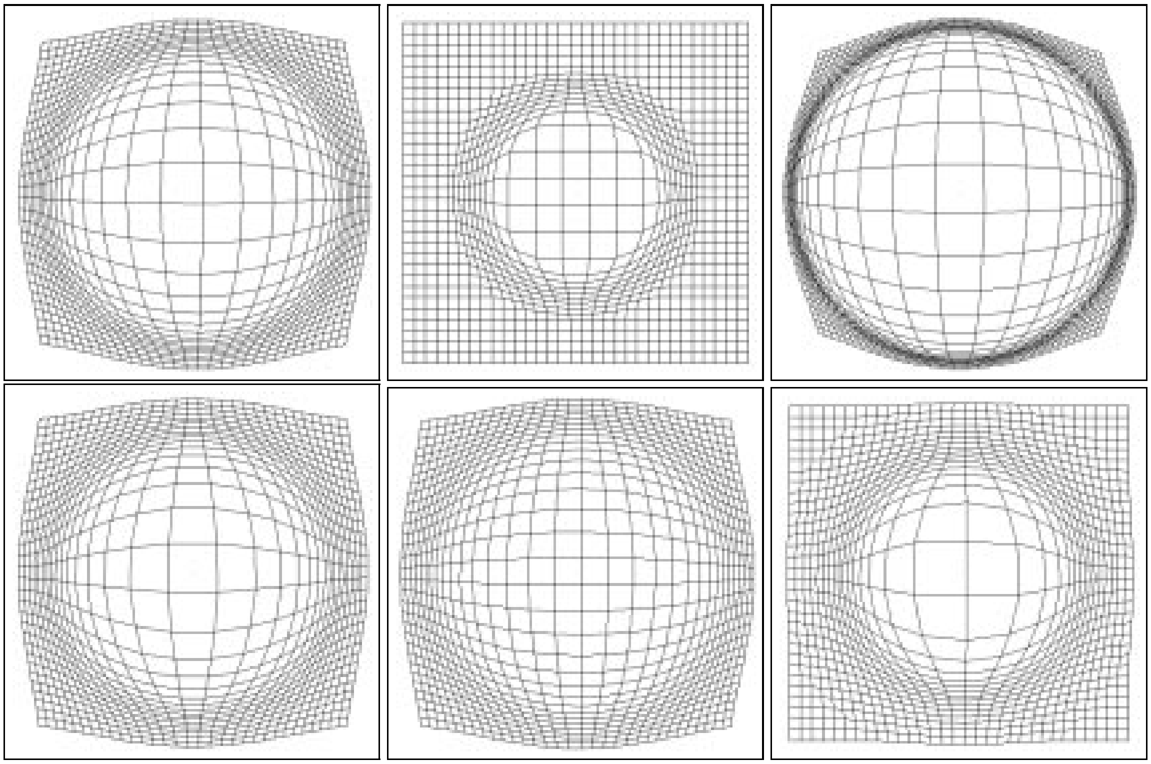

Figure 5.38: A variety of lenses resulting from a Gaussian basis function modified by a half sine function

methods of combining zero drop-off with non-linear drop-offs are discussed in Sections 3.4 and 3.5 as either limiting the function to some maximum height to provide a scaled-only focal region or creating an arbitrary polygon focal region which is entirely projected to the given focal height (also creating a scaled-only focal region). It is possible to create connected piecewise linear functions with different characteristics in different regions.

Additive Combinations. All drop-off functions discussed have at least one less than desirable feature. Rather than look for a new function that might eliminate this feature it is possible to use separate basis and auxiliary curves. The auxiliary curve can be used to modify the problems in the basis curve. For instance, in the discussion in Section 4.1 it was mentioned that the Gaussian's bell shape causes a waist of maximum compression approximately midway between the focus and its context. A half sine function can be subtracted from the Gaussian to straighten out this dip (see Section 4.1). In this manner it is possible to obtain smooth integration of the Gaussian, ensuring that the compression pattern is not locally constricted. Figure 5.38 shows several curves that are the

Figure 5.39: Lenses in combination

result of adjusting the Gaussian and sine functions.

Distinct Lens Combinations. Using constrained distortions opens up the possibility of using different drop-off functions for different lenses. For instance, one lens could use a step function and another a modified Gaussian and still another a truncated hyperbola (Figures 5.39 and 5.40).

5.8 Discussion

Much of this discussion has focused on the type of visual continuity a presentation pattern provides. There has been a general tendency to label as preferable the more visually integrated patterns [85, 137, 138]. However, in a visually integrated distortion, focal areas blend into context. While this provides perception of the image as a single event it can lead to interpretation questions about whether there are any areas that are scaled only, and if so where such areas start and end. Simple visual continuity provides this information more readily.

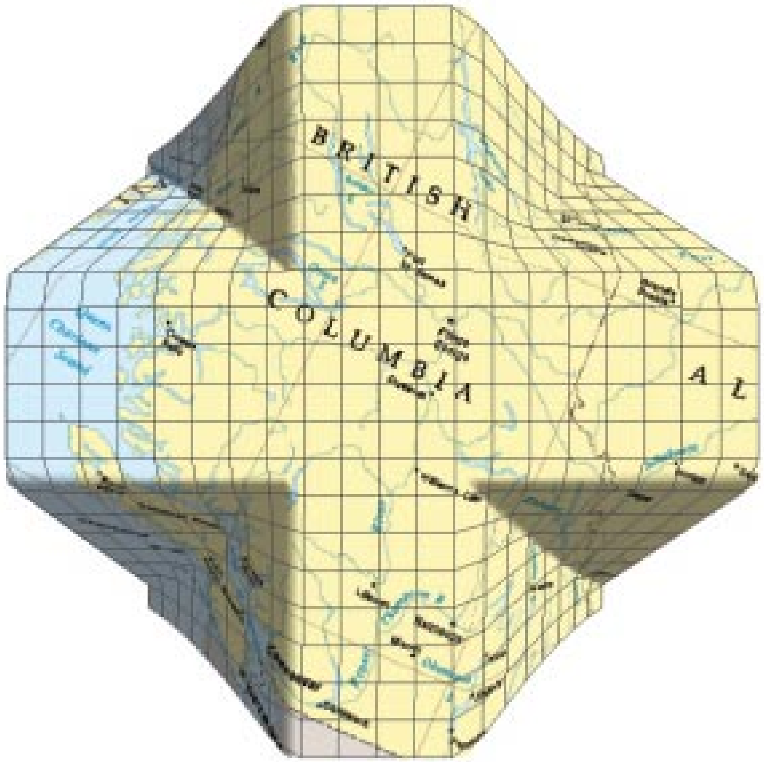

Figure 5.40: Lenses in combination

focal connection					
slope					
context connection					

Table 5.2: Profiles of possible drop-off variations at visually critical points

One may think that small variations in distortion pattern like those between the hyperbola and the cosine are probably not visually signification. However, there may be both critical issues and critical zones in a lens' distortion pattern. The critical issues concern the degree of magnification possible

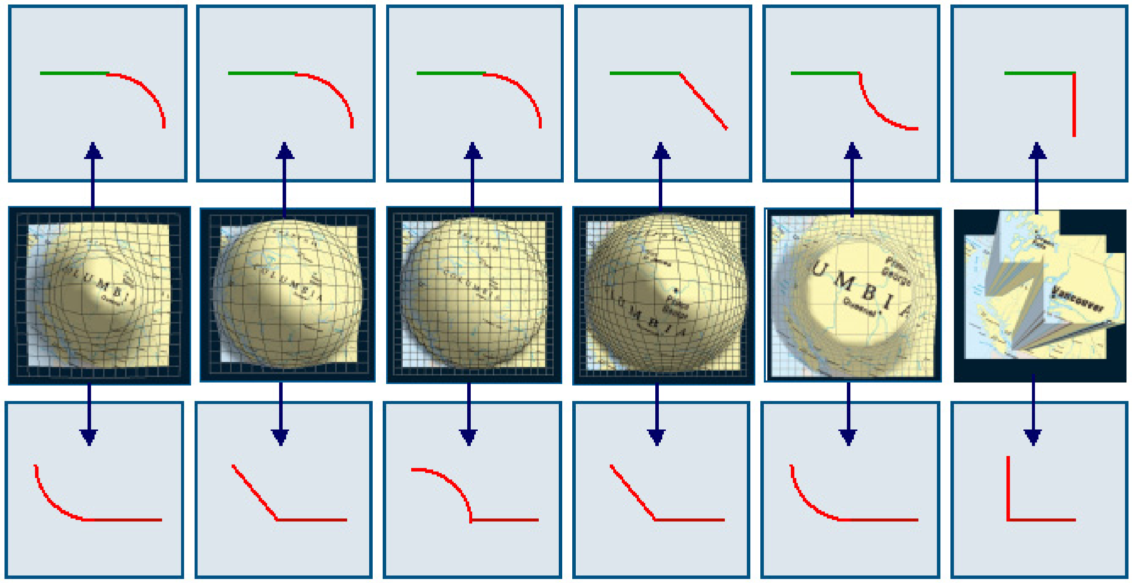

Table 5.3: Choosing the type of visual integration at the focal connection and at the context connection affects the lens' pattern of compression

before the compression becomes too extreme, whether any occlusion is tolerable, and the location of maximum compression. Critical zones include the focal connection, the region of distortion, and the context connection. Table 5.2 indicates the range of drop-off possibilities for these critical zones. A full lens library would offer this range of choice. Since any decreasing mathematical function can be used in this regard, being able to interpret resulting visual patterns from the curve's profile may allow for more appropriate choices between curves for a particular information representation or task. Table 5.3 shows some of the lens possibilities.

While the resulting presentation patterns of other examples from the literature can be explained through EPS, most of them make use of two-dimensional adjustment for their presentation. As discussed in Chapter 2 many cannot provide all the functionality discussed in Chapter 3.2, and none have the functionality of distortion control and folding discussed in Chapter 4. One advantage of EPS is that changing the drop-off function used merely changes the resulting distortion pattern without affecting the functionality. Therefore while the discussion has largely looked at the patterns created by a single lens for ease of comparison, all EPS lenses discussed have full functionality presented in Chapters 3 and 4.