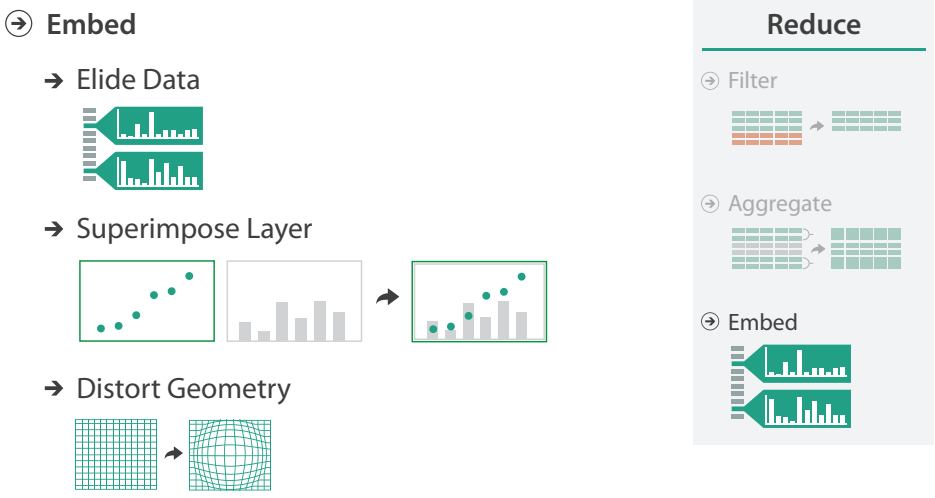

Figure 14.1. Design choices for embedding focus information within context.

Chapter 14

Embed: Focus+Context

14.1 The Big Picture

The family of idioms known as **focus+context** are based on the design choice to **embed** detailed information about a selected set—the **focus**—within a single view that also contains overview information about more of the data—the **context**. These idioms reduce the amount of data to show in the view through sophisticated combinations of filtering and aggregation. The design considerations are sufficiently complex that embedding is addressed separately in the analysis framework as a special case. A very large family of specific idioms that use some form of focus+context embedding has been proposed.*

One design choice for embedding is to elide items, where some items are filtered out completely while others are summarized using dynamic aggregation for context; only the focus items are shown in detail. Another choice is to superimpose layers, where a local region of focus information can be moved against the background layer of contextual information. A third choice is to distort the geometry, where context regions are compressed to make room for magnified focus regions. In all of these cases, there is a choice of single or multiple regions of focus. With geometric distortion, there are choices for region shape, region extent, and interaction metaphor. Figure 14.1 summarizes this set of design choices.

★ Many names are essentially synonyms for or special cases of *focus+context*: **bifocal displays**, **degree-of-interest models**, **detail in context**, **distortion-oriented presentations**, **distortion viewing**, **elastic presentation spaces**, **fisheye lens**, **generalized fisheye views**, **hyperbolic geometry**, **multifocal displays**, **nonlinear magnification fields**, **pliable surfaces**, **polyfocal projections**, **rubber sheet navigation**, and **stretch and squish navigation**.

14.2 Why Embed?

The goal of embedding focus and context together is to mitigate the potential for disorientation that comes with standard navigation techniques such as geometric zooming. With realistic camera motion, only a small part of world space is visible in the image when the camera is zoomed in to see details for a small region. With

► Zooming is discussed in Section 11.5.1.

geometric navigation and a single view that changes over time, the only way to maintain orientation is to internally remember one's own navigation history. Focus+context idioms attempt to support orientation by providing contextual information intended to act as recognizable landmarks, using external memory to reduce internal cognitive load.

Focus+context idioms are thus an example of nonliteral navigation, in the same spirit as semantic zooming. The shared idea with all of them is that a subset of the dataset items are interactively chosen by the user to be the focus and are drawn in detail. The visual encoding also includes information about some or all of the rest of the dataset shown for context, integrated into and embedded within the same view that shows the focus items. Some idioms achieve this selective presentation without the use of geometric distortion, but many use carefully chosen distortion to combine magnified focus regions and minimized context regions into a unified view.

Embedding idioms cannot be fully understood when considered purely from the visual encoding point of view *or* purely from the interaction point of view; they are fundamentally a synthesis of both. The key idea of focus+context is that the focus set changes dynamically as the user interacts with the system, and thus the visual representation also changes dynamically. Many of the idioms involve indirect control, where the focus set is inferred via the combination of the user's navigation choices and the inherent structure of the dataset.

14.3 Elide

One design choice for embedding is **elision**: some items are omitted from the view completely, in a form of dynamic filtering. Other items are summarized using dynamic aggregation for context, and only the focus items are shown in detail.

A general framework for reasoning about these three sets of items is a **degree of interest (DOI)** function: $DOI = I(x) - D(x, y)$. In this equation, I is an interest function; D is the distance, either semantic or spatial; x is the location of an item; y is the current focus point [Furnas 86]. There could be only one focus point, or multiple independent foci. The DOI function can be thought of as a continuous function that does not explicitly distinguish between focus items to show in detail, context items to aggregate, and com-

pletely elided items to filter out. Those interpretations are made by algorithms that use the function, often based on threshold values. These interest functions typically exploit knowledge about dataset structure, especially hierarchical relationships. For example, if a few subsections of a document were selected to be the foci, then a good context would be their enclosing sections.

Example: DOITrees Revisited

The DOITrees Revisited system shown in Figure 14.2 uses multiple foci to show an elided version of a 600,000 node tree. The shaded triangles provide an aggregate representation showing the size of the elided subtrees. The context in which to show them is computed using tree traversal from the many focus nodes up toward their common ancestors and the tree root. In this case, distance is computed topologically based on hops through the tree, rather than geometrically through Euclidean space. The focus nodes can be chosen explicitly by clicking, or indirectly through searching.

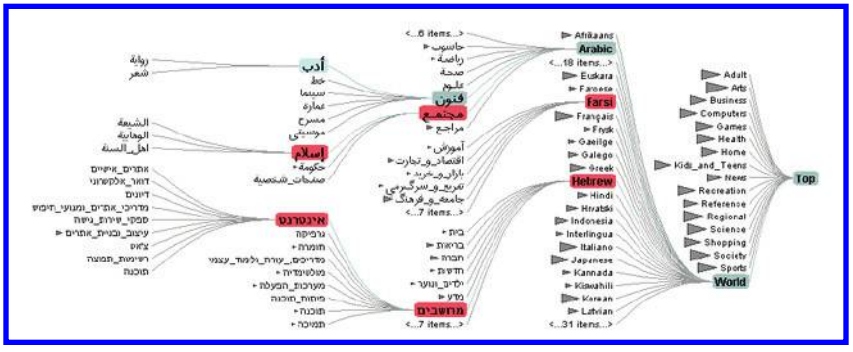

Figure 14.2. DOITrees Revisited uses elision to show multiple focus nodes within context in a 600,000 node tree. From [Heer and Card 04, Figure 1].

System	**DOITrees Revisited**
What: Data	Tree.
How: Encode	Node–link layout.
How: Reduce	Embed: elide, multiple foci.
Scale	Nodes: hundreds of thousands.

14.4 Superimpose

► Superimposing layers globally is discussed in Section 12.5.

Another choice for integrating focus and context is the use of superimposed layers. In this case, the focus layer is limited to a local region, rather than being a global layer that stretches across the entire view to cover everything.

Example: Toolglass and Magic Lenses

The Toolglass and Magic Lenses system shown in Figure 14.3 uses a see-through lens to show color-coded Gaussian curvature in a foreground layer, atop the background layer consisting of the rest of the 3D scene. Within the lens, details are shown, and the unchanged remainder of the other view provides context. The lens layer occludes the region beneath it. The system handled many different kinds of data with different visual encodings of it; this example shows 3D spatial data. The curvature lens shows that the object in the scene that appears to be a perfect sphere when rendered with standard computer graphics techniques is in fact a faceted object made from multiple patches.

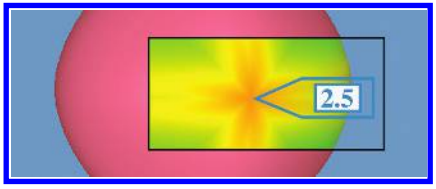

Figure 14.3. The Toolglass and Magic Lenses idiom provides focus and context through a superimposed local layer: the see-through lens color codes the patchwork sphere with Gaussian curvature information and provides a numeric value for the point at the center. From [Bier et al. 93, Figure 12].

System	Toolglass and Magic Lenses
What: Data	Spatial, quantitative curvature attribute across surface.
How: Encode	Use given, color by curvature.
How: Reduce	Embed: superimpose.

14.5 Distort

In contrast to using elision or layers, many focus+context idioms solve the problem of integrating focus and context into a single view using geometric distortion of the contextual regions to make room for the details in the focus regions.

There are several major design choices that underlie all geometric distortion idioms. As with the other two choices, there is the choice of the number of focus regions: is there only a single region of focus, or does the idiom allow multiple foci? Another choice is the shape of the focus: is it a radial, rectangular, or a completely arbitrary shape? A third choice is the extent of the focus: is it global across the entire image, or constrained to just a local region? A fourth choice is the interaction metaphor. One possibility is constrained geometric navigation. Another is moveable lenses, evocative of the real-world use of a magnifying glass lens. A third is stretching and squishing a rubber sheet. A fourth is working with vector fields.

These choices are now illustrated through five examples of distortion idioms: 3D perspective, fisheye lenses, hyperbolic geometry, stretch and squish navigation, and magnification fields.

Example: 3D Perspective

Several early idioms used 3D perspective to provide a global distortion region with a single focus point. The interaction metaphor was constrained geometric navigation. The perspective distortion arising from

Figure 14.4. The Cone Tree system used 3D perspective for focus+context, providing a global distortion region with a single focus point, and using standard geometric navigation for interaction. From [Card and Mackinlay 99, Figure 10].

the standard 3D computer graphics transformations is a very intuitive and familiar effect. It was used with the explicit intention of providing a distortion-based focus+context user experience in many early vis systems, such as the influential Cone Tree system shown in Figure 14.4 [Robertson et al. 91].

► The costs and benefits of 3D are discussed in Section 6.3.

Although many people found it compelling and expressed a strong preference for it on their first encounter, this approach lost popularity as the trade-offs between the costs and benefits of 3D spatial layout for abstract information became more understood.

System	**Cone Trees**
What: Data	Tree.
How: Encode	3D node–link layout.
How: Reduce	Embed: distort with 3D perspective; single global focus region.
Scale	Nodes: thousands.

Example: Fisheye Lens

The **fisheye lens** distortion idiom uses a single focus with local extent and radial shape and the interaction metaphor of a draggable lens on top of the main view. The fisheye idiom provides the same radial distortion effect as the physical optical lenses used with cameras and for door peepholes. The lens interaction provides a foreground layer that completely replaces what is beneath it, like the magic lens idiom, rather than preserving what is beneath and superimposing additional information on top of it, like the superimposing layer design choice for faceting data between views. The lens can be moved with the standard navigation approach of 2D translation. The mathematics of fisheye distortion is straightforward; modern graphics hardware supports high performance fisheye lenses using vertex shaders [Lambert et al. 10].

Figure 14.5 shows two examples of a fisheye lens used with an online poker player dataset. The scatterplot in Figure 14.5(a) shows the percentage of time that a player goes to showdown (playing until people have to show all of their cards) versus the flop (playing until the point where three cards are placed face-up on the board). In the dense matrix view of Figure 14.5(b), blocks representing players are color coded according to their winning rate, and a space-filling curve is used to lay out these blocks in order of a specific derived attribute; in this case, a particular betting strategy. In the parts of the scene under the fisheye lens, the labels are large enough to read; that focus region remains embedded within

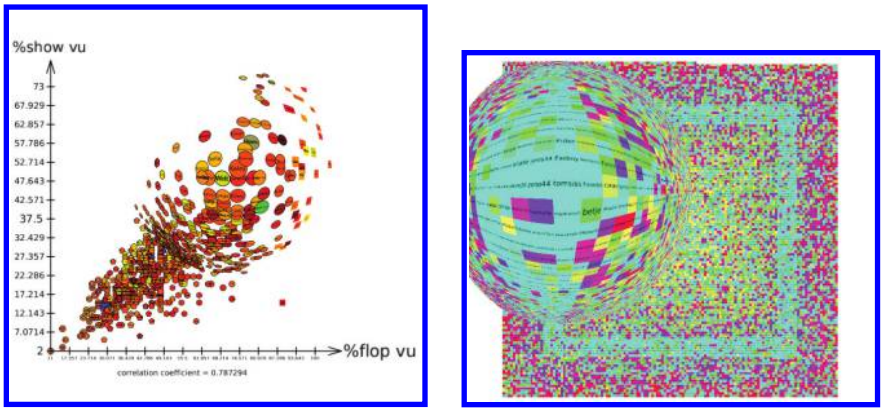

(a) (b)

Figure 14.5. Focus+context with interactive fisheye lens, with poker player dataset. (a) Scatterplot showing correlation between two strategies. (b) Dense matrix view showing correlation between a specific complex strategy and the player's winning rate, encoded by color.

the surrounding context, showing the global pattern within the rest of the dataset.

Idiom	**Fisheye Lens**
What: Data	Any data.
How: Encode	Any layout.
How: Reduce	Embed: distort with fisheye; single focus, local radial region, moveable lens interaction metaphor.

Example: Hyperbolic Geometry

The distortion idiom of hyperbolic geometry uses a single radial global focus with the interaction metaphor of hyperbolic translation. This approach exploits the mathematics of non-Euclidean geometry to elegantly accommodate structures such as trees that grow by an exponential factor, in contrast to standard Euclidean geometry where there is only a polynomial amount of space available for placing items. An infinite non-Euclidean plane can be mapped to a finite Euclidean circle, and similarly an infinite non-Euclidean volume can be mapped to a finite sphere in Euclidean space. The interaction metaphor is hyperbolic translation, which corresponds to changing the focus point of the projection; the visual effect

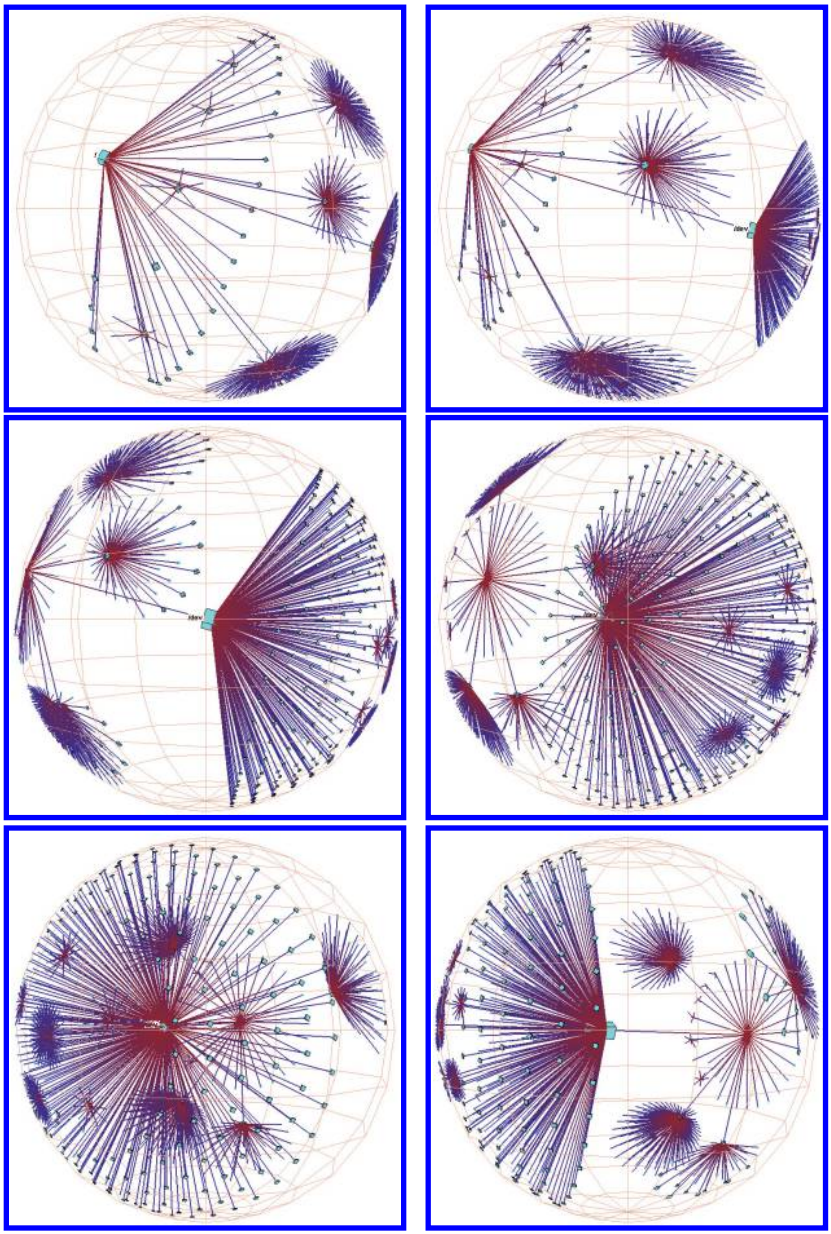

Figure 14.6. Animated transition showing navigation through 3D hyperbolic geometry for a file system tree laid out with the H3 idiom, where the first three frames show hyperbolic translation changing the focus point and the last three show standard 3D rotation spinning the structure around. From [Munzner 98, Figure 3].

is changing which items are magnified at the center, versus minimized at the periphery, for a global effect with similarities to using a fisheye lens that extends across the entire scene.

Figure 14.6 shows a 3D hyperbolic node–link tree representing the structure of a file system laid out with the H3 idiom [Munzner 98], through a sequence of frames from an animated transition as the view changes over time. The first three frames show hyperbolic translation to change what part of the tree is magnified, where the subtree on the right side gets larger as it moves toward the center of projection. The last three frames show standard rotation to clarify the 3D structure. The rationale for using 3D rather than 2D was to achieve greater information density, but at the cost that any single frame is partially occluded.

► The costs and benefits of using 3D for abstract data are covered in Section 6.3.

Idiom	**Hyperbolic Geometry**
What: Data	Tree or network.
How: Encode	Hyperbolic layout.
How: Reduce	Embed: distort by projecting from hyperbolic to Euclidean space; single global radial focus; hyperbolic translation interaction metaphor.

Example: Stretch and Squish Navigation

The **stretch and squish navigation** idiom uses multiple rectangular foci of global extent for distortion, and the interaction metaphor where enlarging some regions causes others to shrink. In this metaphor, the entire scene is considered to be drawn on a rubber sheet where stretching one region squishes the rest, as shown in Figures 11.7, 14.7, and 14.8. Figure 14.7 shows stretch and squish navigation with the TreeJuxtaposer system [Munzner et al. 03], where Figure 14.7(a) shows two small trees juxtaposed with linked highlighting and navigation, and Figure 14.7(b) shows the result of multiple successive stretches within a single large tree. The borders of the sheet stay fixed so that all items stay visible within the viewport, although they may be projected to arbitrarily small regions of the image. The user can choose to separately stretch the rectangular focal regions in the horizontal and vertical directions.

These figures also illustrate the visual encoding idiom of **guaranteed visibility** that ensures that important objects are always visible within the scene, even if they are very small. Guaranteed visibility is an example of aggregation that operates at the subpixel level and takes the importance attribute of each item into account. Standard graphics systems use assumptions that work well when drawing realistic scenes but are not necessarily true for abstract data. In reality, distant objects are not visually

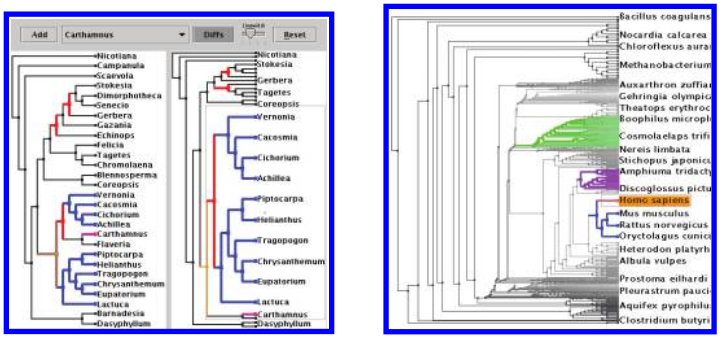

(a) (b)

Figure 14.7. TreeJuxtaposer uses stretch and squish navigation with multiple rectangular foci for exploring phylogenetic trees. (a) Stretching a single region when comparing two small trees. (b) Stretching multiple regions within a large tree. From [Munzner et al. 03, Figures 5 and 1].

Figure 14.8. PRISequenceJuxtaposer supports comparing gene sequences using the stretch and squish navigation idiom with the guaranteed visibility of marks representing items with a high importance value, via a rendering algorithm with custom subpixel aggregation. From [Slack et al. 06, Figure 3].

salient, so it is a reasonable optimization to simply not draw items that are sufficiently far away. If the viewpoint is moved closer to these objects, they will become larger on the image plane and will be drawn. However, in abstract scenes the distance from the camera is a poor stand-in for the **importance** value for an object; often an original or derived attribute is used instead of or in combination with geometric distance. The example in Figure 14.8 is a collection of gene sequences that are over 16,000 nucleotides in width displayed in a frame of less than 700 pixels wide [Slack et al. 06]. The red marks that indicate differences between gene sequences stay visible at all times because they are given a high importance value, even in very squished regions where hundreds or thousands of items may fall within a single pixel. Figure 11.7 also illustrates this idiom: the value

used for the box color coding also indicates importance, so the boxes representing alerts that are colored red are always visible.

Idiom	**Stretch and Squish Navigation**
What: Data	Any data.
How: Encode	Any layout.
How: Reduce	Embed: distort with stretch and squish; multiple foci, global rectangular regions, stretch and squish navigation interaction metaphor.

Example: Nonlinear Magnification Fields

The nonlinear magnification fields idiom relies on a general computational framework featuring multiple foci of arbitrary magnification levels and shapes, whose scope can be constrained to affect only local regions. The underlying mathematical framework supports calculations of the implicit

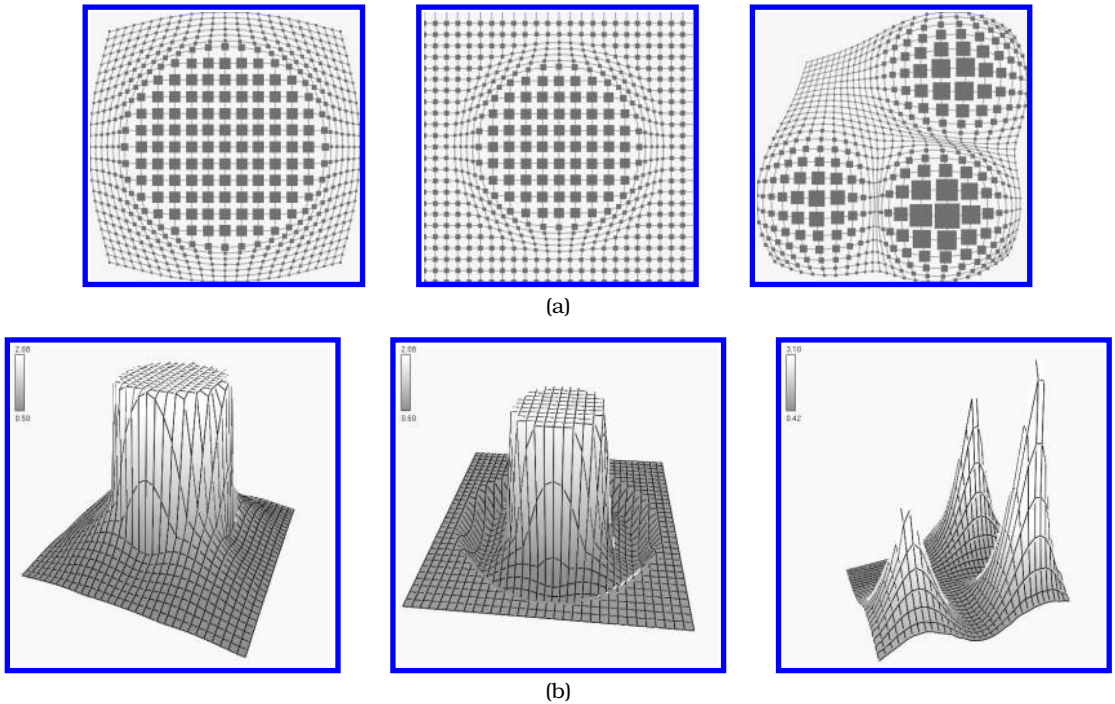

Figure 14.9. General frameworks calculate the magnification and minimization fields needed to achieve desired transformations in the image. (a) Desired transformations. (b) Calculated magnification fields. From [Keahey 98, Figure 3].

magnification field required to achieve a desired transformation effect, as shown in Figure 14.9(a). The framework supports many possible interaction metaphors including lenses and stretchable surfaces. It can also expose the magnification fields shown in Figure 14.9(b) directly to the user, for example, to support data-driven magnification trails of moving objects.

Idiom	**Nonlinear Magnification Fields**
What: Data	Any data.
How: Encode	Any layout.
How: Reduce	Embed: distort with magnification fields; multiple foci, local arbitrary regions, lens or stretch or data-driven interaction metaphors.

14.6 Costs and Benefits: Distortion

Embedding focus information within surrounding context in a single view is one of five major alternatives to showing complex information. The other four choices are deriving new data, manipulating a single changing view, faceting into multiple views, and reducing the amount of data to show. The trade-offs of cost and benefits between these five approaches are still not fully understood.

What has gradually become clear is that distortion-based focus+context in particular has measurable costs, in addition to whatever benefits it may provide. One cost is the problem that distance or length judgements are severely impaired, so distortion is a poor match with any tasks that require such comparisons. Thus, one of the most successful use cases for geometric distortion is with exploration of node–link layouts for networks and trees. The task of understanding the topological structure of the network is likely to be robust to distortion when that structure is shown using lines to connect between nodes, or containment to show nested parent–child node relationships, because precise angle and length judgements are not necessary.

One potential cost is that users may not be aware of the distortion, and thus misunderstand the underlying object structure. This risk is highest when the user is exploring an unfamiliar or sparse structure, and many idioms incorporate explicit indications of distortion to lessen this risk. Hyperbolic views typically show the enclosing circle or sphere, magnification fields often show a

superimposed grid or shading to imply the height of the stretched surface.

Even when users do understand the nature of the distortion, another cost is the internal overhead of maintaining **object constancy**, which is the understanding that an item seen in two different frames represents the same object, just seen from a different viewpoint. Understanding the underlying shape of a complex structure could require mentally subtracting the effect of the transformation in order to recognize the relationship between the components of an image before and after the transformation. Although in most cases we do this calculation almost effortlessly for standard 3D perspective distortion, the cost of mentally tracking general distortions increases as the amount of distortion increases [Lam et al. 06]. Some empirical evidence shows that constrained and predictable distortion is better tolerated than more drastic distortion [Lam and Munzner 10].

The originator of the generalized fisheye view approach has expressed surprise about the enthusiasm with which others have embraced distortion and suggests that the question *what* is being shown in terms of selective filtering is more central than that of *how* it is shown with any specific distortion idiom [Furnas 06]. For example, the fisheye metaphor is not limited to a geometric lens used after spatial layout; it can be used directly on structured data, such as a hierarchical document where some sections are collapsed while others are left expanded.

Figure 14.10 illustrates four different approaches on the same node–link graph [Lambert et al. 10]: a fisheye lens in Figure 14.10(a), an ordinary magnifying lens in Figure 14.10(b), a neighborhood highlighting idiom using only superimposed layers in Figure 14.10(c), and a combination of that layering with the Bring and Go interaction idiom in Figure 14.10(d). Discussing these examples in detail will shed some light on the costs and benefits of distortion versus occlusion versus other interaction.

The local fisheye distortion has a small circle region of very high magnification at the center of the lens surrounded by a larger intermediate region that continuously varies from medium magnification to very high compression, returning to low compression in the outer periphery. Although fisheye lenses were developed with the goal of reducing the viewer's disorientation, unfortunately they can be quite disorienting. The continuous magnification change introduces some amount of cognitive load to untangle the underlying shape from the imposed distortion. Distortion is less problematic with familiar shapes, like geographic maps of known places, be-

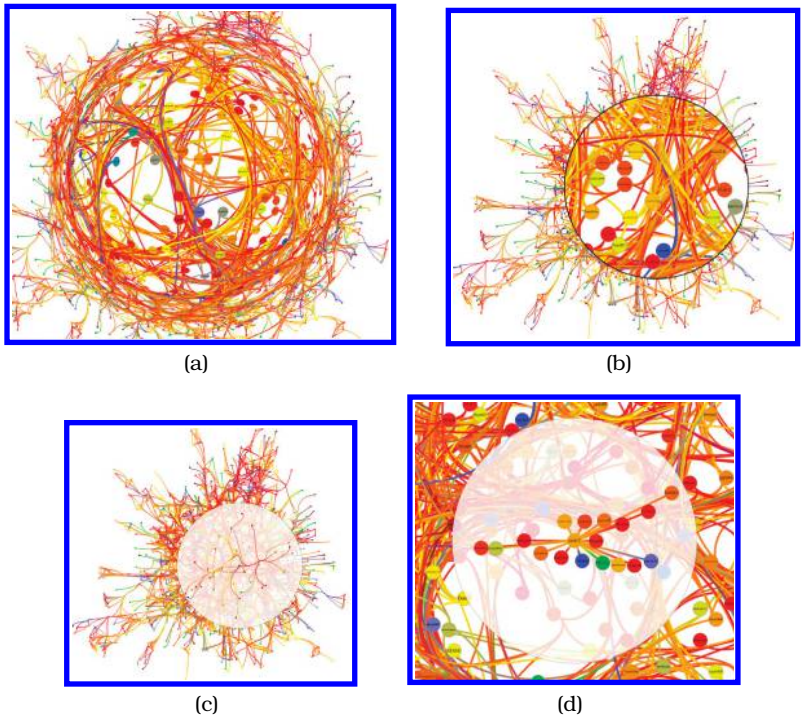

Figure 14.10. Four approaches to graph exploration. (a) Fisheye lens. (b) Magnifying lens. (c) Neighborhood highlighting with layering. (d) Neighborhood highlighting with both layering and Bring and Go interaction. From [Lambert et al. 10, Figures 2a, 2b, 3b, and 4b].

cause people can interpret the distortion as a change to a known baseline. With unfamiliar structures, it can be difficult to reconstruct the basic shape by mentally removing the distortion effects. While these idioms are designed to be used in an interactive setting, where the user quickly moves the lens around and compares the undistorted to the distorted states, there is still some cognitive load.

In contrast, the superimposed local magnifying lens has just two discrete levels: a highly magnified circle of medium size, and the low-compression periphery of medium size. There is a discontinuous jump between these two levels, where the lens occludes the immediately surrounding region. In this particular example, the fisheye lens may be more difficult to use than the magnifying lens; it is not clear that the benefit of avoiding occlusion is worth the cost of interpreting the continuous magnification change.

The last two approaches show a specific region of interest, in this case a local topological neighborhood of nodes reachable within one or two hops from a chosen target node. The neighborhood highlighting lens does not distort spatial position at all; it uses layering by reducing the opacity of items not in that neighborhood, automatically calculating a lens diameter to accommodate the entire neighborhood. While this approach would not help for tasks where magnification is critical, such as reading node labels, it does a good job of supporting path tracing.

A fisheye lens can be interpreted as a temporary warping that affects the location of all objects within the active region. The **Bring and Go** interaction idiom for network data [Moscovich et al. 09] is also temporary, but selectively changes the location of only specific objects of interest, by bringing the one-hop neighbors close to the target node. The layout is designed to simplify the configuration as much as possible while still preserving direction and relative distance information, in hopes of minimizing potential disorientation. This interaction idiom exploits topological structure information to reduce the cognitive load cost of tracking moving objects: only the one-hop neighbors move during animated transitions, in contrast to fisheye lenses that affect the positions of all items within their span.

14.7 Further Reading

The Big Picture Two extensive surveys discuss a broad set of idioms that use the choices of changing the viewpoint through navigation, partitioning into multiple views, and embedding focus into context, including an assessment of the empirical evidence on the strengths and weaknesses of these three approaches and guidelines for design [Cockburn et al. 08, Lam and Munzner 10].

Early Work Early proposals for focus+context interfaces were the Bifocal Display [Spence and Apperley 82] and polyfocal cartography [Kadmon and Shlomi 78]. An early taxonomy of distortion-based interfaces introduced the unifying vocabulary of magnification and transformation functions [Leung and Apperley 94].

Fisheye Views The fundamental idea of generalized fisheye views [Furnas 82, Furnas 86] was followed up 20 years later with a paper

questioning the overwhelming emphasis on geometric distortion in the work of many others in the intervening decades [Furnas 06].

3D Perspective Influential 3D focus+context interfaces from Xerox PARC included the Perspective Wall [Mackinlay et al. 91] and Cone Trees [Robertson et al. 91].

Frameworks Two general frameworks for focus+context magnification and minimization are elastic presentation spaces [Carpendale et al. 95, Carpendale et al. 96] and nonlinear magnification fields [Keahey and Robertson 97].

Hyperbolic Geometry Hyperbolic 2D trees were proposed at Xerox PARC [Lamping et al. 95] and 3D hyperbolic networks were investigated at Stanford [Munzner 98].

Stretch and Squish Navigation The TreeJuxtaposer system proposed the guaranteed visibility idiom and presented algorithms for stretch and squish navigation of large trees [Munzner et al. 03], followed by the PRISAD framework that provided further scalability and handled several data types [Slack et al. 06].