Chapter 2

The Screen Real Estate Problem

The screen real estate problem is ubiquitous. The fact that we can store and manipulate vast amounts of information in a computer and have only a comparatively small screen on which it can be viewed is an issue in almost all aspects of computing. Consequently screen usage discussions appear in many branches of the literature. In some areas, such as visualization, interface design and visual languages, there is research directly focused on presentation. In other areas, such as databases, software engineering, and GIS systems, applications exist that have made significant contributions to screen usage. There are also other subject areas outside of computing in which relevant matters such as representation, effective use of display space, human perception and comprehension are discussed.

A full review of all the relevant literature would cover many areas in computing as well as several other disciplines. Since we are considering computational presentation, this review proceeds from a computational perspective, mentioning related work in other fields only where it seems particularly relevant. We review the computing literature that explores presentation issues, focusing on detail-in-context presentations that are intended for two-dimensional representations. Since our results extend to three-dimensional representations, we also include a review of related three-dimensional research.

The chapter is organized as follows: Sections 2.1 and 2.2 discuss presentation methods designed for two-dimensional representations, focusing on the manner in which they provide context, and their resulting visual effects. Section 2.1 describes three early computational presentation methods: Windowing [76], Bifocal Display [147] and Generalized Fisheyes [52]. These three presentation approaches consider the problem of effective screen usage and make different suggestions concerning the maintenance of context. Section 2.2 discusses subsequent presentation research, organized into six groups according to which ideas regarding context have motivated them. Section 2.3 compares these methods according to the functionality they provide. Section 2.4 reviews detail-in-context presentations designed for three-dimensional representations. Section 2.5 reviews taxonomies that have been suggested in regards to this literature.

2.1 Initial Presentation Approaches

This section discusses the initial computational presentation approaches which have proved seminal in presentation research. These are the Windowing [76, 97], Bifocal Display [147] and Generalized Fisheyes [52]. The ideas embodied in these approaches have generated several themes and have been the basis for most presentation research. The description of these three approaches focuses on the ideas that motivated them.

2.1.1 Windowing Paradigm

Xerox Star [76, 145] introduced several presentation advances over command line access. These include the separation of the presentation space into discrete views or windows, scrolling, panning and zooming within these windows, and a free-form two-level-of-detail presentation that uses icons to represent windows or applications that are not of current interest. These methods have become widespread and are frequently referred to as WIMP (Windows, Icons, Menus, and Pointing) interfaces.

During the last fifteen years the ideas in Xerox Star have become the dominant paradigm, though no one is claiming, at least in their current manifestations, that they are the perfect solution. Even their general acronym WIMP is pejorative. As discussed in Chapter 1 some of the shortcomings of the windowing paradigm are difficulties: in navigation, in recognizing relationships between items in separate windows, and in providing adequate context. The use of windows in WIMP interfaces increases usable space by providing multiple display spaces but does not address the issue of detail-in-context. The main presentation ideas in windows are:

- the partitioning of space,
- the tiling of regions,
- the free-form overlapping of regions, supporting freedom to reposition,
- the replacing of large objects with a smaller symbols,
- the zooming, panning and scrolling capabilities within a window, and

- the inclusion of pointing using the mouse as a movable indication device.

2.1.2 Bifocal Display

In 1982 Spence and Apperley [147] noted the frequency of both crowding and navigation problems with WIMP interfaces and introduced several divergent ideas. They motivate their work from the perspective of the search problem and declare that humans often have a less than precise recollection of exactly what they are looking for. Someone may not know the exact name of a book but may know the subject or the location it was last seen. This imprecision leads to a need for preliminary browsing and searching activities before actual retrieval. They discuss how searching in physical space is supported by spatial memory, memory of previous actions, and visual and verbal clues. Through subsequent physical searches all of these clues are continually reinforced by a reasonable degree of constancy. They proposed a presentation method, Bifocal Display, to support the use of spatial memory and visual scan in computational search.

Bifocal Display is a two-level presentation method that places a single focal area within a unified presentation as opposed to the creation of separate views. The central rectangular focal area extends from top to bottom across a long strip of information and is presented at suitable magnification for the information. The regions on each side of this focus are compressed uniformly in length to fit in the remaining display space. This compressed *full context* maintains at least some symbolic vestige of the remaining information in the strip. They suggest that maintaining a compressed version of the full context in which spatial constancy is preserved will support use of visual scan and spatial memory. Maintaining *spatial constancy* requires ensuring that all parts or regions of the representation stay in the same positions relative to each other. That is, the locations of the items in the context are considered important for preserving spatial memory and therefore should be respected and moved as little as possible. Bifocal Display's incorporation of two levels of detail within a single unified presentation is the first computational detail-in-context presentation.

The main presentation ideas in Bifocal Display are:

- the maintenance of spatial constancy to support spatial memory,
- the introduction of more than one scale in a unified presentation, and
- the compression of a full context.

2.1.3 Generalized Fisheyes

Furnas [52, 54] observed that a fisheye, or very wide angle, lens shows a world view with a detailed central focus and the surrounding context in gradually decreasing detail. He noted that an artistic presentation of this type of view not only used the fisheye distortion but also kept a few selected objects in disproportionate detail considering their distance from the focus. He mentions a poster, the New Yorker's view of the USA, that shows Manhattan street by street while just indicating the Hudson River and New Jersey and showing only Chicago, the Rockies, and California from the rest of the country. The display includes detail around the focus and only those aspects that are considered important elsewhere.

He also studied how people retain and present information in various subject areas and workplaces such as geography, history and newspapers. These studies revealed that people usually know the details about their own interests set in enough domain knowledge to provide context. These studies confirm the importance of maintaining context and of setting a focus within its context. He suggested that the widespread evidence shown in his studies indicates that detail-in-context may be a useful and intuitive way to present information.

Furnas' studies suggest that, when enough is known about the domain to ensure preservation of crucial aspects, a filtered context may be sufficient. The basic idea behind Generalized Fisheye Views [52] is that there is a centre of interest or focus about which detail is important and that with the exception of some domain specific items, interest decreases as distance from the focus increases. Furnas suggested achieving this through filtering the context using a *degree of interest* (DOI) function. A DOI is based upon the distance from the current focus and an *a priori importance* (API) that is domain specific and known for each item. Using a *domain specific* function requires that the person creating the presentation have knowledge about the information or data in the representation. If this type of knowledge is available, it is very useful in creating appropriate presentations; however, visual exploration of unknown or little know information spaces is frequently desired.

The main presentation ideas in Generalized Fisheye Views are:

- the maintenance of a sufficient context (this requires domain specific knowledge),
- variation in the size of the area of interest around the focus (this is again domain, task and user specific), and
- the possibility of more than one area of interest.

2.2 Methods of Maintaining Context

There is agreement on the importance of maintaining context. However, the exact manner in which it is to be preserved is still under discussion. While the methods vary, the central purpose, which is to preserve the essential features of the context in less space, remains constant. This section

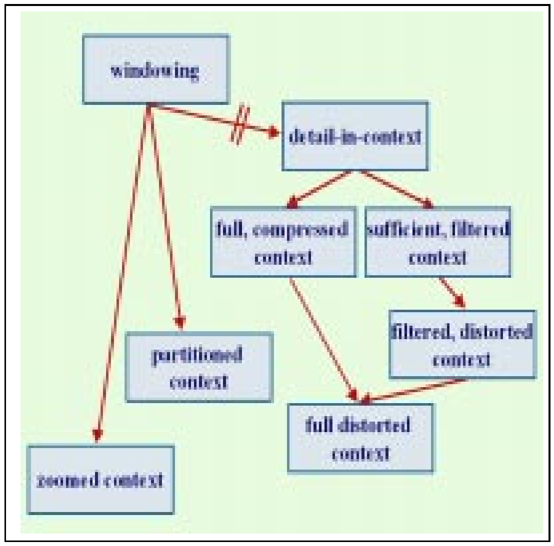

Figure 2.1: The relationships between the different ideas concerning maintaining context. The cross-hatch over the connection between windowing and detail-in-context indicates that detail-in-context ideas developed to provide a feature not available with windowing.

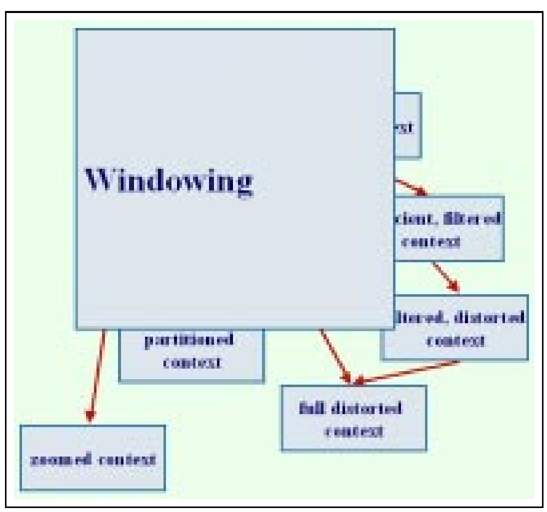

Figure 2.2: A diagram of the free-form overlapping allowed by most windowing approaches

discusses six different methods of preserving context: compression, filtering, filtering and distorting, distortion, partition, and zooming (Figure 2.1). These methods place presentation approaches into groups which often have characteristic visual patterns.

Figure 2.1 is used through this discussion on the different ideas behind maintaining context to provide a simple illustration of the basic types of visual organization. For example, Figure 2.2 shows the expansion of the windowing node covers parts of the rest of the diagram.

One important feature of visual presentation patterns is the nature of the visual connection between different regions in the presentation. Methods used to maintain context often result in presentations that have regions that differ in scale or distortion. The visual relationships between these regions can be classified by the nature of their visual connection.

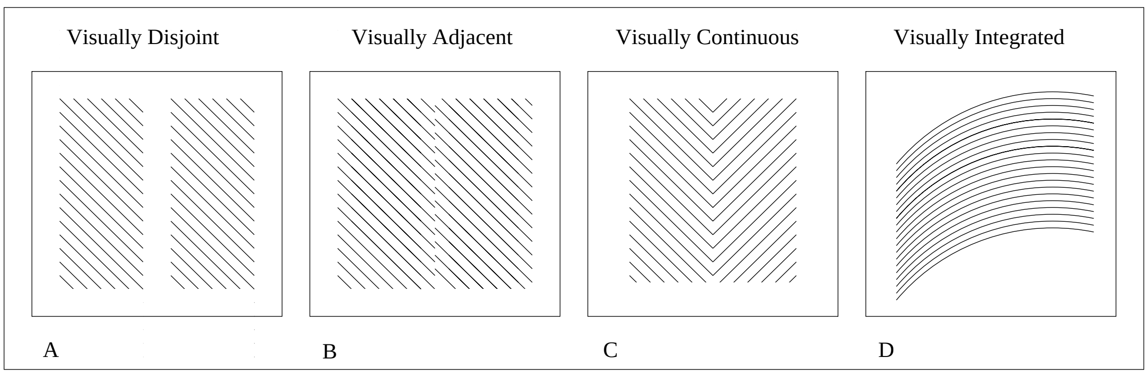

Figure 2.3: Types of visual connections

The visual connection pattern between two regions is *visually disjoint* when the regions of the image are presented with a gap between them (Figure 2.3 A). This separation can be in any direction. In this type of presentation, relationships between regions are not visually apparent and therefore must be deduced from domain knowledge or assumed based on other information.

Two regions that are placed next to each other to compose a single image but have contents or sub-regions that do not align, are called *visually adjacent* (Figure 2.3 B). In this type of presentation it is clear that there is only one image, but the relationships between contents or sub-regions in the image can be unclear. Such things as domain knowledge and visual cues are needed to clarify actual relationships. For example, Figure 2.3 B gives no information as to whether a line in the left-hand region should be connected to the line above or below it in the right-hand region.

The connection pattern is *visually continuous* when two regions are combined in a manner that also connects all their components and sub-regions (Figure 2.3 C). This visual connection can be minimal, such as G_0 continuity. The lines in the components and sub-regions are continuous but not necessarily differentiable. The connection may include a change in direction; components and/or sub-regions in adjacent regions will meet but not smoothly. These images provide enough visual information to indicate whether a sub-region or line is a continuation of one in the adjacent section or not. They also provide visual information about the position of the join.

The connection pattern is *visually integrated* when two adjacent image regions and all the components and sub-regions in these regions are visually connected and when the connection is sufficiently smooth for the resulting image to appear unified (Figure 2.3 D). This visual integration can be as minimal as G_1 continuity. The connection may include a change in direction but the change will not be abrupt. Visual integration supports visual gestalt but makes recognizing location of the boundary between regions difficult.

The six sub-sections that follow each focus on the ideas behind their methods of maintaining context and the affect this has on visual connections. The first group follows and extends the ideas about context expressed in Bifocal Display. These approaches provide a single unified presentation, consider it important that all items in the context are visible in the presentation, and try to preserve spatial constancy. As the primary method used to achieve this is to uniformly compress the context, these approaches are discussed under Compressed Contexts in Section 2.2.1.

The ideas put forward in Generalized Fisheyes have fostered three ways of handling context. One is to filter the context. This stream of presentation research is discussed under Filtered Contexts in Section 2.2.2. Next, DOI functions were extended to include both filtering and distorting. These are discussed in Section 2.2.3. In response to the difficulty experienced in reading contexts that were both filtered and distorted, research emerged that combined the idea of gradually diminishing interest as the distance from the focus increases, with the idea of preserving full contexts. This group uses only distortion to maintain context and is discussed under Distorted Contexts in Section 2.2.4.

Two of the ideas in the windowing paradigm have also been explored further. These are partitioning and zooming. The approaches that explore uses of partitioning are discussed in Section 2.2.5, and those that have extended the idea of zooming to create a unified presentation approach are discussed under Zoomed Contexts in Section 2.2.6.

2.2.1 Compressed Contexts

Methods that attempt to use compression-only to adjust their context follow, in principle, the directions introduced by Spence and Apperley in 1982 [147]. Research in this group endeavours to provide a unified detail-in-context presentation that maintains full context and preserves spatial constancy. The primary method has been to provide space required by the focus by uniformly compressing the context. However, in order to create a unified presentation the context is sometimes compressed in x and y separately, resulting in regions where the compression in x does not equal the compression in y. An essential feature of a compressed context is that every item in the context is visible in the adjusted presentation even if it is either very small or has been replaced by a smaller symbol. One of the important contributions from this group is the refinement of what is meant by spatial constancy.

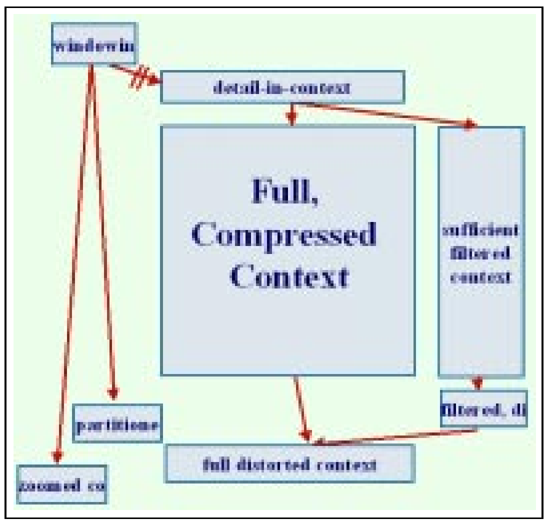

Figure 2.4: A diagram of the characteristic compression pattern for compressed contexts

These presentations have a characteristic visual pattern (Figure 2.4). The resulting foci and regions of compression are typically square or rectangular (Figure 2.5) and there are abrupt changes in scale between regions. These regions are usually visually connected but sometimes are merely visually adjacent. The transformation of these methods are *step* functions in that there are clearly defined regions that are compressed or magnified differently.

Bifocal Display [147], as discussed in Section 2.1, uses uniform compression in x only as it is intended for a linear strip of information. When browsing, the information is moved across the single, central, stationary focus. Figure 2.5, left image, shows the one-dimensional compression pattern for Bifocal Display.

Leung, Apperley and Spence [95, 146] extend the ideas in Bifocal Display to a two-dimensional method for a two-dimensional representation (Figure 2.5, central image). The selected focus is

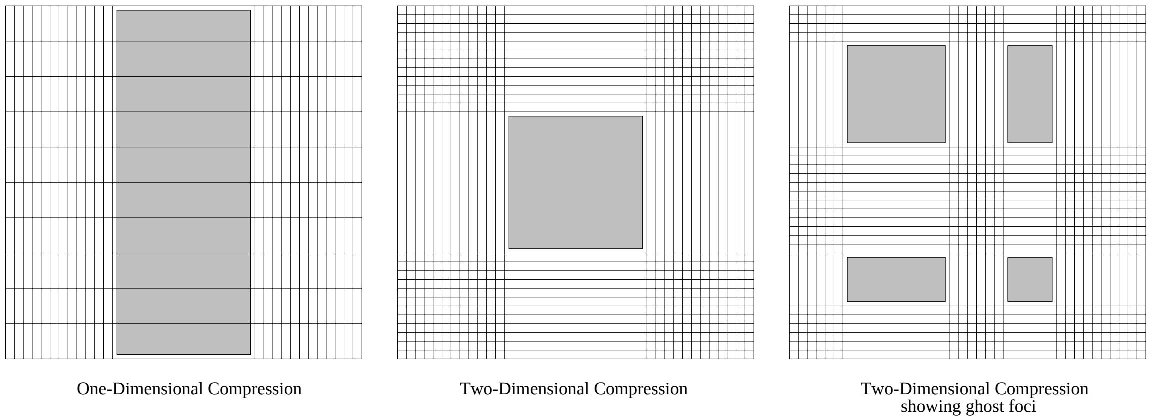

Figure 2.5: Characteristic visual presentation patterns of using compressed contexts. The darker regions are areas of magnification. The selected foci expand uniformly. Selected foci of differing sizes cause non-uniform magnification in the ghost foci.

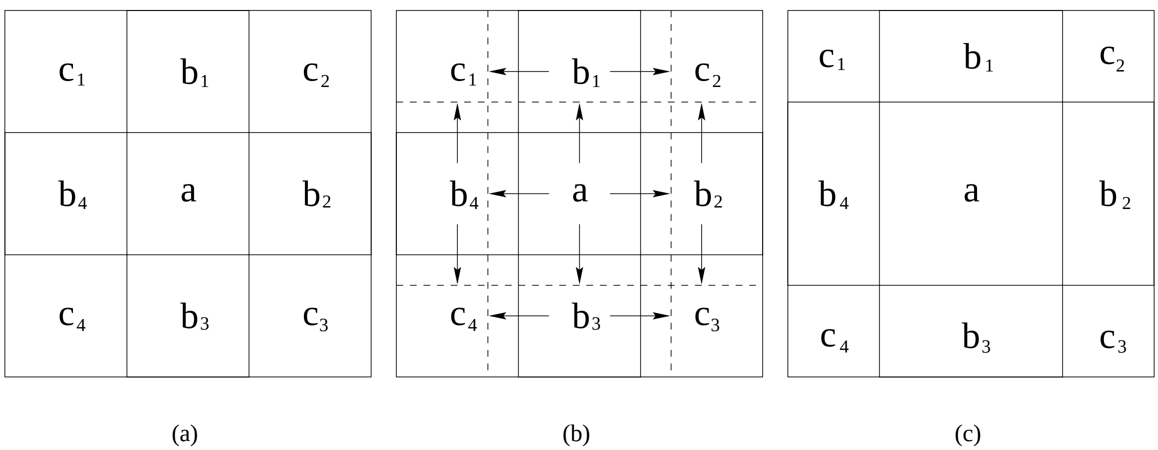

Figure 2.6: The dynamics of an orthogonal stretch

magnified, then the context is compressed uniformly but separately in x and y, creating a visual pattern with nine distinct regions. Figure 2.6 illustrates how this pattern is formed. Enlarging a selected focal region a causes distinct regions of context. In Figure 2.6(b) the focus a is uniformly magnified, the regions adjacent to the focus b_1 to b_4 are magnified in one direction and compressed in the other, and the remaining regions c_1 to c_4 are uniformly compressed.

Misue et al. [42, 107] refine the idea of spatial constancy further. They introduce the possibility of aiding user's ability to recognize distorted views by preserving their mental map. They suggest maintaining three spatial properties:

- *orthogonality* - objects maintain relative right/left up/down positioning,
- *proximity* - adjacent objects remain adjacent, and
- *topology* - containment relationships are preserved.

The thought is that limiting the pattern of changes will provide better support for the mental map that the user has formed of the information. However, adhering to these ideas limits the nature of the distortion. This is further discussed in [138] where it is noted that an orthogonal distortion is required to maintain orthogonal relationships. The resulting pattern is called an *orthogonal stretch* because it preserves orthogonality and has stretched sections that lie adjacent to the focus which introduce distortion into the context. Figures 2.5, 2.6, and 2.7(a) all show examples of orthogonal stretch.

There has been a stream in the literature that has favored these considerations [109, 131, 138, 139, 147, 158] and furthered them with the introduction of multiple focal points. Multiple focal points have in turn introduced even more variation in scale and distortion. Figure 2.5 C shows orthogonal stretching with two foci. Notice how the expansion of each focal point creates two stretched strips, one vertical and one horizontal due the the stretch in x and y respectively. Where two of these strips intersect unrequested *ghost foci* result. While the intentional foci have uniform magnification the ghost foci's magnification can vary in x and y. Both the stretched strips and the ghost foci create magnified regions for which no interest has been directly specified. However, full context, orthogonality and topology are maintained.

Leung, Spence and Apperley [95] conducted a comparative study of five map presentation methods: a paper map, Bifocal Display [147] with scrolling access, Bifocal Display with point-and-shoot access, a single scrolling view, and a two-view split-screen approach. Point-and-shoot access allows the user to point and click to select a new focal position and the two-view presentation offers one detailed view and one overview. Their results indicate that the scrolling version of Bifocal Display is better than the other three computer presentations. They noted that this result may well be task specific.

Table Lens [131] makes use of the basic orthogonal stretch to provide a multi-focal distortion view specifically for tables. Without explicitly mentioning the mental map, the authors discuss maintaining orthogonality and straight lines in order to preserve the table's integrity while enlarging individual cells. The resulting distortion is very similar to Biform Display [42, 108] as well as the rectilinear approach in both Rubber Sheet [138] and CATGraph [82, 83]. As tables are laid out in rows and columns this type of distortion matches very well with the representation.

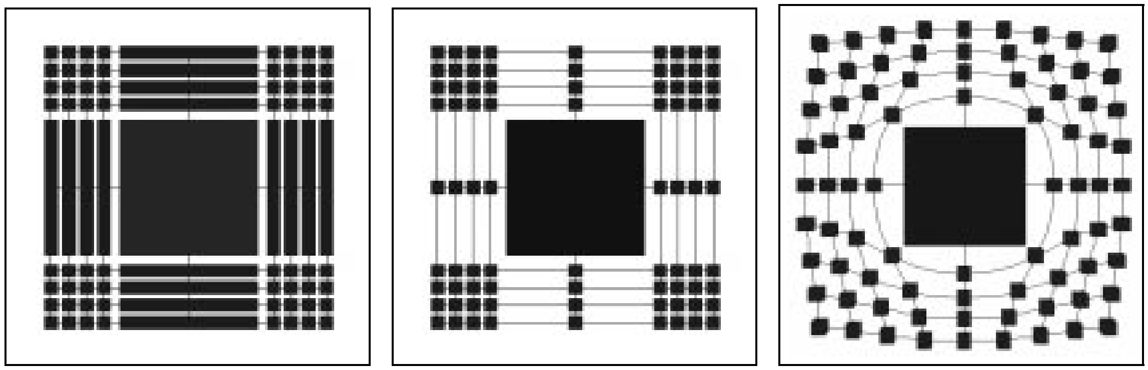

(a) An orthogonal stretch

(b) The pattern of SHriMP's variation of orthogonal distortion

(c) The pattern of SHriMP's variation of radial distortion

Figure 2.7: Variations of compressed contexts

SHriMP [150, 151] introduces variations of the same basic idea of magnifying the focus and then uniformly compressing the full context. The issue of unrequested non-uniform magnification, as in ghost foci, is avoided by allowing the space adjacent to the foci to be stretched but not applying this stretch to the items in the space. This avoids the variant distortion in x and y for the nodes but introduces considerable unused or *white* space. Figure 2.7(a) shows a typical orthogonal stretch and Figure 2.7(b) shows SHriMP's variation.

SHriMP also comes closer to providing solutions that support the users mental map. Storey and Müller [150] note that different aspects of the representation are preserved when different distortion methods are used. In particular, the orthogonal stretch method better preserves orthogonality (Figure 2.7(b)), while applying radial compression better preserves proximity (Figure 2.7(c)). SHriMP algorithm expands the focal node, pushing other nodes out of the display area to make space, then scales all nodes to once again fit into the area. Repeating this procedure over sequentially selected focal nodes results in each subsequent selection reducing the magnification of the previously selected focal nodes.

SHriMP was designed to address users' needs in a software maintenance environment and studies [152, 153] have been conducted to compare it with other software maintenance tools. In these studies SHriMP included a single focus and full-zoom capabilities analogous to those described in Pad++ [10, 11, 120] (see Section 2.2.6). These studies were designed to discover appropriate ways to support the cognitive processes involved with software maintenance. From their observations they conjecture that a supportive environment would offer a combination of working strategies and the ability to switch easily between them.

Van der Hedyen et al. [166] extend the SHriMP approach to include various alternatives to the layout adjustments with regards to minimizing variations in scale and the amount of white space while still maintaining orthogonality and parallelism. These presentation variations are designed specifically to suit the needs of radiologists working with MR images.

2.2.2 Filtered Contexts

A filtered context provides more space for the focus by removing parts of the context that are not currently important (Figure 2.8). There is a significant distinction, as noted by Price [125, 126], to be made between simple filtering which merely removes objects from the presentation, and *eliding* which removes objects in a manner that leaves some indication of their absence. Ellipses (...) elide text and icons elide windows on a computer screen. As simple filtering is directly tied to the representation and requires no spatial reorganization, there are no typical visual characteristics with these approaches.

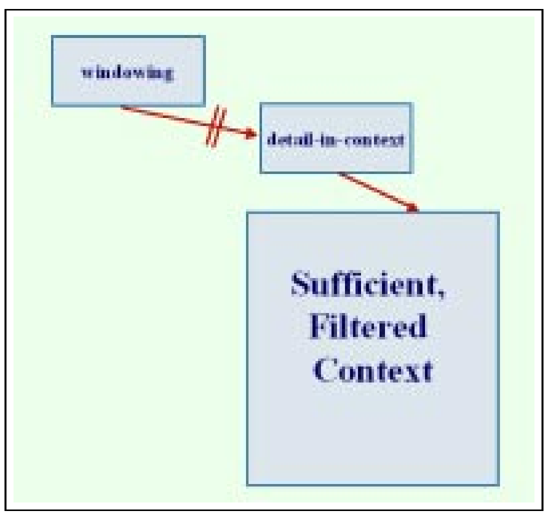

Figure 2.8: A diagram of a filtered context

The possibility of creating detail-in-context views through selective filtering with a DOI was introduced by Furnas [52] in 1986. This filtering is implemented by applying a threshold to the DOI function below which objects are not displayed. His examples included program code, calendar and a tree graph. When information is filtered the program code and the calendar naturally provide some indication of what is missing. With the code, this information is supplied by the line numbers. If the

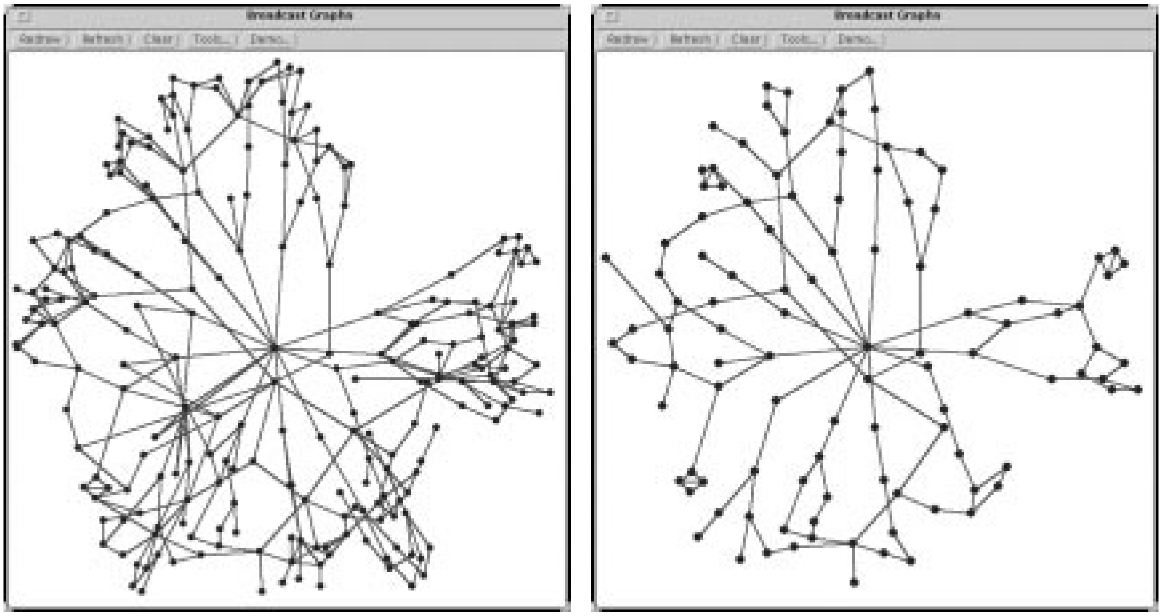

(a) A random graph with a spring layout [69]

(b) Filtering a graph can make some details easier to detect by removing others

Figure 2.9: An example of a filtered context

numbering in the filtered view skips from fifteen to twenty-one then it is clear that the lines sixteen to twenty have been elided. In the calendar example as the nodes get smaller they still display as much information as possible, even if this is only a few letters. However, even the first few letters of a word may trigger a memory and indicate that something has been entered for that day. However, the tree graph example gives no indication that anything has been filtered or that the filtered tree is only part of a larger tree. Figure 2.9(a) shows a graph and a filtered version of the same graph (Figure 2.9(b)). While the filtering has removed a lot of visual clutter, making it easier to to see some features (such as three four-cycles with a diagonal edge) it could be simply perceived as a different graph.

Graphlog [34] is a graph based visualization of databases that makes extensive use of filtering and features a visual database query language. As the scale of the database becomes too large to make reasonable visual displays, two filtering presentation choices are offered. One is to provide an overview of dots, where each dot represents a node in the graph (with no edges) and allow the user to select a sub-region. The other is to use a query to set up a view with only the relevant information.

Seenet [9] is a visual analysis tool for network data that employs three tactics to cope with visual clutter. Two are geographic, one displaying link information on the map and the other displaying node information on the map. The third is a diagram that displays link information in a matrix representation. However, any one of these displays can get too cluttered to be readable. To combat this the user can dynamically filter information. For instance, the user can dynamically shorten edge lengths, creating edges that operate as pointers. These edge pointers still give reasonable visual indication of their connections and the effect of shortening long edges the reach across the network prevents occlusion of shorter intermediate ones.

Starfields [1, 73] display multi-dimensional information on a two-dimensional scatter plot. Each item is represented by a small mark whose position is determined by two of its ordinal attributes. Keeping the marks very small allows for many items to be on the screen at once. Dynamic filtering makes this interface remarkably usable for fairly large data sets. It appears that the type of interaction is particularly important in creating effective filtered contexts.

2.2.3 Filtered and Distorted Contexts

Graphic interpretations of the concepts in Generalized Fisheyes include some distortion to make better use of the space provided by filtering. Hollands et al. [70] created a detail-in-context presentation that combined distortion with filtering in 1989. They use a DOI to decide what items are displayed (giving a filtered context) and also to decide the size of the displayed items (creating a distorted context) (Figure 2.10). The combination of domain specific filtering and distorting creates presentations that differ quite radically from the normal presentation and provides no characteristic visual patterns by which the adjusted presentations can be recognized. Figure 2.11 shows a graph in normal presentation on the left and a filtered and distorted presentation of the same graph on the right.

Topographic Networks [70] provide an approach based on Furnas's work that includes both filtering and distortion. A DOI function based on a built-in API and the distance from the focus is used to both filter and distort a single focus view. Nodes are adjusted in terms of level of detail displayed as well as size and position. A new focus can be set interactively, however, it is possible that every location on the display will change. There is no animation between views.

Hollands' study [70] which compared Topographic Networks with the ability to pan and scroll gave fairly inconclusive results. The subjects performed three tasks: simple route finding given a

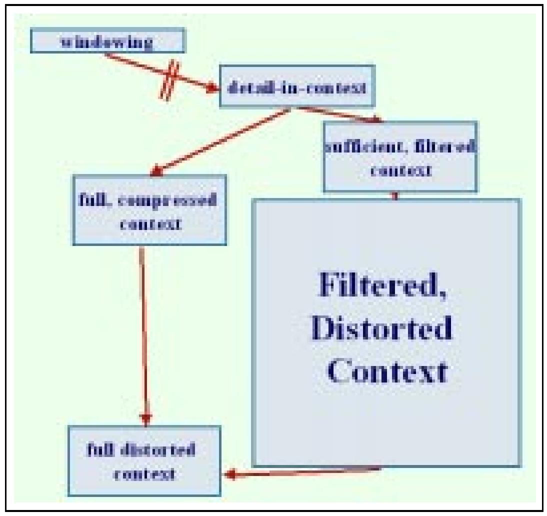

Figure 2.10: A diagram of a slightly filtered context that has also been distorted

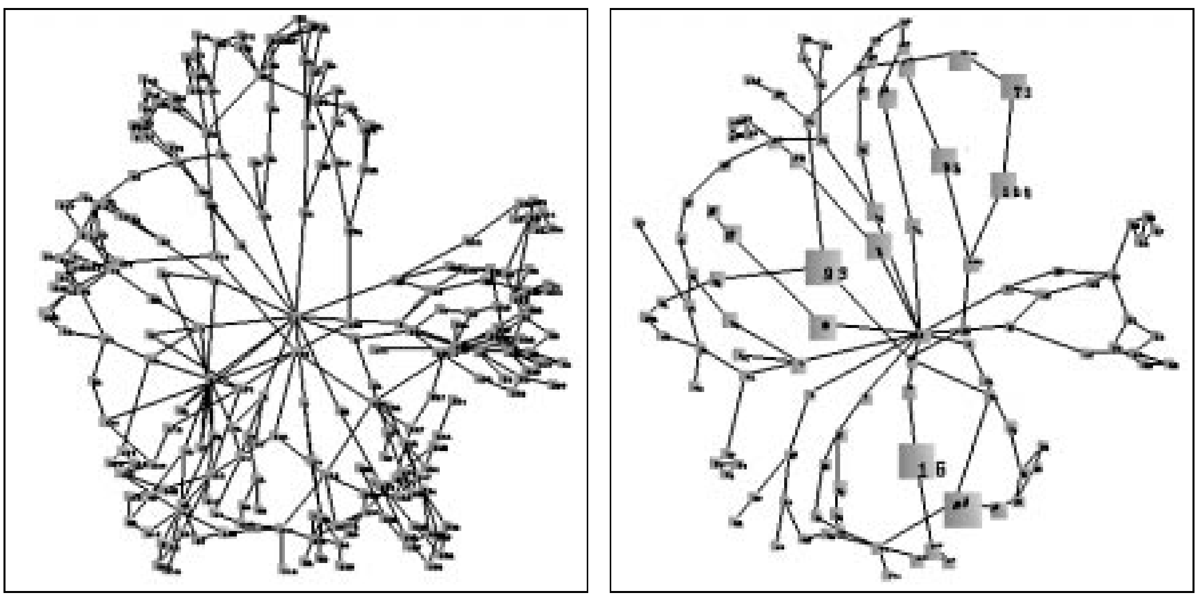

(a) The graph in normal presentation

(b) The same graph filtered and distorted

Figure 2.11: An example of a filtered and distorted context

start and end point, route finding which included locating one of the end points which was not initially visible, and itinerary planning of a route with either three or five stops. Topographic Networks were only better when both the start and end points could not be seen on a single view. This result raises the question of whether one would use a distorted view if all the relevant information could be seen. Itinerary planning is indicated as being faster with the distorted presentation than with scrolling.

Also, they acknowledge that their distortion may be less than ideal as subjects reported that when the entire presentation changed completely with change of focus it was jarring. It would seem that there would be considerable cognitive effort reconciling the new view with remembered ones. Their study raises questions about whether filtering and distortion together might be too difficult to read.

Sarkar and Brown [137] also employ a DOI calculation and the distance from the focal point to establish the position, size, and amount of detail to display for each vertex. This approach offers a single focal point and maintains constancy of focal position in a transformation that effects the entire presentation. A threshold function can be incorporated so that nodes whose DOI is sufficiently low are not displayed. There is no indication in the filtered image of what or how much information has been filtered. Of particular interest are their comments about user disorientation when using this transformation. They respond by providing a hemispherical distortion for maps, postulating that this may seem more familiar. They also note that filtering images may make them less recognizable and suggest that their distortion-only variations may be preferred.

A group of related approaches: Variable Zoom [40, 141], Continuous Zoom [7], Intelligent Zoom [8], and CZ Web [33, 45], also use both filtering and distortion. However, they use the style of distortion discussed under compressed contexts (Section 2.2.1) and elide instead of simply filtering. These multiple foci detail-in-context approaches support maintenance of orthogonality. Though nodes change radically in size and proportion they stay in the same spatial orientation to each other. Nodes are nested hierarchically and effect each other's size hierarchically. As a node grows, its siblings are the first to adjust, then parents or children follow. Relative size is decided by user request and a balancing factor. The factors used by the DOI are dynamic semantics derived from monitoring the network, user behavior and balancing usable space. While node size is in direct control of the user in that it will immediately expand or contract as requested by mouse action, the rest of the nodes will adjust themselves as needed, balancing the DOI with network information and orthogonality constraints. This approach was designed for a situation where time is of essence and minimizing the complexity of the user's task is important. However, it has limitations. For instance, one cannot return to an exact previous setting because of the balancing factors. An aspect of this system that has not been emphasized is that their use of hierarchies, continuous compression, and dynamic closing of cluster nodes provides a very comprehensible elision. In fact, it would appear that much of their capability to handle very large networks derives from this.

Schaffer et al. [140, 141] compared the Variable Zoom [40], with a full-zoom approach. Their task was to find a break in a network and than re-route the connection. The full zoom was implemented within the same interface as the distortion zoom, so that other than the zooming technique, the interfaces were the same. The task was performed with more ease and accuracy when the Variable Zoom was used. They note that while these results indicate the usefulness of the Variable Zoom approach for the task tested, they do not necessarily extend to all tasks. One recommendation from this study, which has sub-sequently been included in almost all distortion methods, is the importance of smoothly animated transitions.

Fisher et al. [45] studied CZWeb [33], the most recent version of the Zoom family described above. CZWeb supports web browsing by building up a hierarchical graph overview of the sites visited as the user navigates the web. The user can dynamically adjust the overview using the Continuous Zoom technique [7]. The study found that CZWeb was easy to use, useful in moving from site to site (particularly backward more than five sites) and generally made the web structure more understandable. However, the users did not think it provided support for the way they thought about the web spatially.

Other than the relatively successful combination of elision and distortion in the Zoom family of approaches, the idea of using filtering and distortion together has not been pursued recently.

2.2.4 Distorted Contexts

Hollands et al. [70] and Sarkar and Brown [137] both independently used a graphic interpretation of a DOI which combined the use of distortion with filtering. However, while these types of presentations seemed promising as far as task performance, users complained that they were jarring [70] and disorienting [137]. In response to these comments Sarkar and Brown [137] created presentations that used a simple mathematical distortion instead of a complex DOI and maintained full context.

Most distortion-only approaches vary the context's degree of compression as the distance from the foci increases (Figure 2.12). These methods endeavor to maintain full context, consider spatial constancy, and attempt to maintain topology but often fall short of maintaining orthogonality. Instead proximity and clustering are maintained. The general appearance of these types of approaches is very similar. Figure 2.13 shows three variations from this stream of research.

These approaches are quite distinct from those that focus on compressed contexts and those that filter and distort with a DOI. For distortion-only approaches the selected focus is magnified as

Figure 2.12: A diagram of a full distorted context. Note how the distortion continues right across the focus and note the extreme compression at the edges.

desired and then the size and location of the rest of the information is positioned to create curved distortions with no ghost foci that smoothly integrate a magnified region with its surrounding context. There are, however, two mathematical schools of thought within the distortion-only approaches. One uses a 2D-to-2D transformation to create the distorted presentations [59, 80, 82, 84, 137, 108], and the other uses a 2D-to-3D manipulation and incorporates perspective projection to create the final presentation [20, 99, 132].

In 1978 cartography literature, Kadmon and Shlomi describe a 2D-to-2D distortion-only transformation which they call Polyfocal Projections [80]. While their approach was not interactive it did describe the mathematics for a 2D-to-2D transformation that provides a multiple foci presentation.

Several 2D-to-2D distortion approaches based on different mathematical functions were proposed. Multi-Viewpoint Perspective (MVP) [108] offered an approach based on the arctan function. The arctan distortion function can be applied in two ways: either using polar coordinates to give a radial distortion and mapping the representation unto a circle, or using Cartesian coordinates to give an orthogonal distortion and creating a rectangular presentation. In both of these approaches the distortion is *global* in that it is spread over the entire representation. As compression increases with the distance from the focus the edges of the representation can become extremely compressed. In fact, MVP's authors, Eades et al. [42] rejected their own arctan approaches because although they achieve their goal of being able to map an infinite domain into a finite space, in the process they crush the outer regions of the representation beyond recognition. Also these presentations create magnified but distorted foci. In subsequent research they describe goals for maintaining a user's

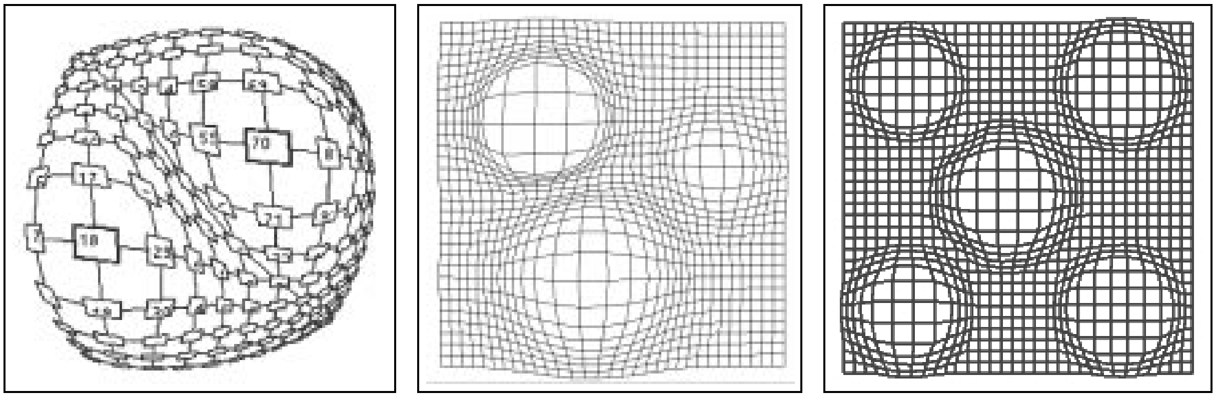

(a) Arctan Fisheye from CATGraph [82] (for image credit see Appendix C.7)

(b) 3-Dimensional Pliable Surfaces (3DPS) [20]

(c) Non-Linear Views [84] (for image credit see Appendix C.8)

Figure 2.13: Distorted contexts

mental map (discussed in Section 2.2.1), and propose two orthogonal compression methods, Biform Display [158] and Abductor [109]. CATGraph [82, 83] offers an arctan function based approach very similar to MVP. Figure 2.13(a) shows a radial CATGraph distortion with two focal regions.

Rubber Sheet [138] features two distinct approaches: an orthogonal stretching method (as discussed in Section 2.2.1) and a method based on morphing [12]. The morphing-based approach, Polygon Handle, provides multiple foci as convex polygons and does not cause areas of unrequested magnification. The resulting presentation comes closest, at the date it was published, to providing a smooth unwasteful multi-focal integration of detail-in-context. However, it has limitations that cause the authors to prefer their orthogonal approach. In regions of high compression between two foci or between a focus and the edge of the image, the ordering of the information in the representation can become reversed. For instance, two vertices that were positioned left to right may be repositioned right to left in the distorted presentation. We call this problem *information reversal*. As information reversal is an extreme violation of spatial constancy, the authors place limits on the size and location of the foci. This interferes with freedom of focal location and user control of level of detail. The resulting image lacks visual information about the original undistorted topology of the image. Also, more than one iteration of the transformation may be required to provide the right balance between detail and context.

Hyperbolic display [90, 91] uses the hyperbola to provide magnification and compression. It places a graph on a hyperbolic plane and then projects this plane onto a circular display, providing a central single focus. The focus is changed by dragging the newly selected node to the center of the presentation. This method is intended for the presentation of hierarchical structures. There is no possibility of creating a focus that spans an area of the graph, nor of allowing for magnification of a focus to scale in a manner that maintains distance relationships. Also, just as the arctan functions described above, the function is asymptotic and the distortion it creates can cause extreme compression at the edges.

All the approaches discussed this far use 2D-to-2D transformations. Leung and Apperley [94] review these distortion approaches explaining how the 2D-to-2D distortion transformations are linked mathematically. This review is discussed further in Chapter 5.

Perspective Wall [99] and Document Lens [132] use the third dimension and perspective compression. Perspective Wall [99] creates a detail-in-context view for one-dimensional information, placing it on a wall which is then bent to display three sides, the central one providing the focus. Not only is context provided, but it is readable. One can see the information on each side and understand, at least approximately, relative importance because one can interpret actual size from position in perspective. Document Lens [132] provides a single movable focus or *lens* for viewing text files. It is a two-dimensional interpretation of Perspective Wall and as such retains most the advantages. The perspective information is provided by the shape of the lens, a truncated pyramid, and the text itself.

Our approach, 3-Dimensional Pliable Surfaces (3DPS) [20], also makes use of the third dimension and perspective projection. 3DPS combines and extends the power and flexibility suggested by the smoothly integrated multi-focal presentation of Polyfocal Projection [80] and the readability of the 2D-to-3D perspective approaches Perspective Wall [99] and Document Lens [132]. 3DPS offers multiple, arbitrarily shaped foci and the possibility of controlling the location and degree of magnification and/or compression. Figure 2.13(b) shows 3DPS with a three focal distortion (see Chapter 3 for a full explanation of 3DPS).

Subsequently, Keahey and Robertson [84] offer similar freedoms in Non-Linear Views with a 2D-to-2D based approach. Figure 2.13(c) shows five foci created with Non-Linear Views. However, for those approaches based on 2D-to-2D transformations obtaining an appropriate transformation given a chosen magnification is non-trivial [85, 94]. Keahey and Robertson [85] have developed Magnification Fields an approximate iterative approach that allows them to derive a transformation given a magnification field. While use of Magnification Fields is not dependent on the number of foci it is an iterative approach and can can require several iterations to achieve a suitable presentation. Also in order to prevent the creation of degenerate presentations (which include information reversal) the number of iterations is dependent on the degree of magnification requested. With sufficiently powerful machines this approach can still be interactive. However, in a 2D-to-3D framework the relationship between magnification and transformation is much simpler (see Chapter 3).

Focus Line [59, 106] incorporates the use of domain knowledge to create a distortion-only approach that can extend its distortion into chosen directions in the representation. For instance, with maps a given road may have several important factors such as a city or airport in one direction and relatively open space in another direction. Focus Line will attempt to locate the compression in the less used areas of the representation.

2.2.5 Partitioned Contexts

Partitioned contexts create more usable space by dividing up the available space. These partitioned regions can be created and organized in different ways. Figure 2.15 shows various familiar organizations of windows and Figure 2.14 shows a heirarchical partioning of context.

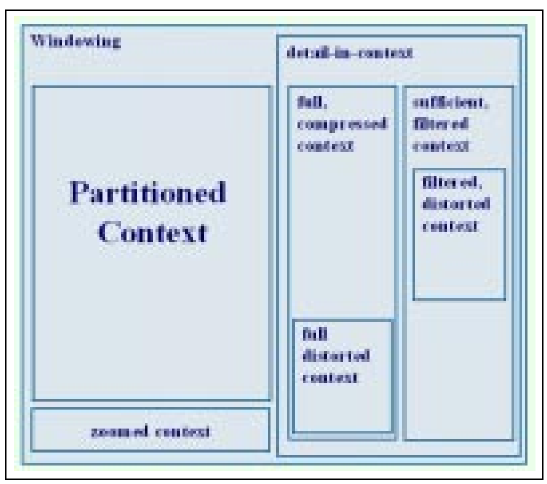

Figure 2.14: A diagram of a partitioned context

Rooms [66] extends the partitioning used in WIMP interfaces to include division of virtual spaces. A user can establish several different screen configurations that have each been set up for a particular task. An overview map allows the user to switch between these screen configurations. In 3D rooms [134] a user can move about the room, zoom in on objects of interest and open doors to other rooms. This approach gives the user control of creating virtual partitions, which are then accessed through visual symbols. There are now many available variations on this idea, such as the virtual desktop.

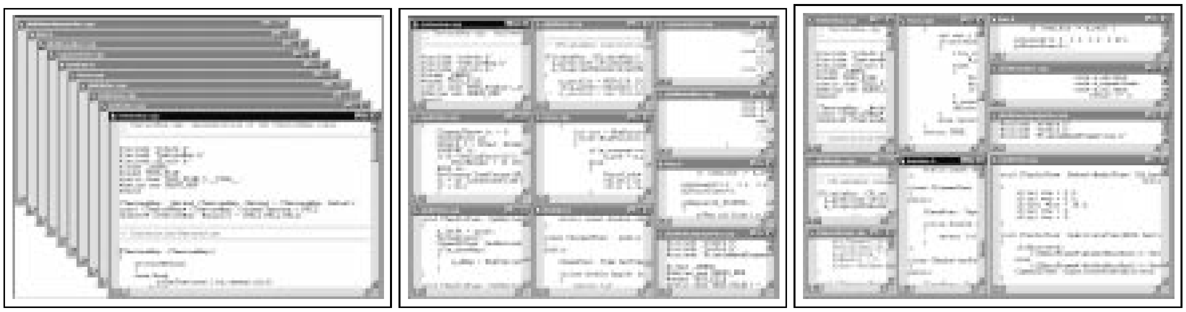

Figure 2.15: Partitioned contexts

A Treemap [74, 75, 144, 163] is created by hierarchically partitioning the representation. Each node's children share their parent's space according to a chosen information-specific weighting factor such as size or age. Advantages claimed by Treemaps include the complete usage of space, and a considerable increase in the number of nodes that can be displayed versus node and edge layouts. This presentation lends itself to easy comparisons of the weighting factor used to affect the node size. Since a user can interactively set the attribute used as the weight, this in effect creates a distortion view based on a single importance function. Treemaps also degrade gracefully, providing filtering of sibling nodes whose aspect ratio is extreme. The latter point is also clearly a disadvantage; there are times when losing nodes, no matter how small, is a problem. Placing a limit on how small a node can get may prevent this. However, as the number of the nodes to be displayed surpasses the number of pixels available, nodes will still disappear.

Interactive graph layout [67, 68] allows the nodes in a selected subgraph to be sized and positioned either interactively or with simple layout algorithms. Subsequent layout of the rest of the graph will respect the bounding boxes of previously laid out sections. In this way a user can personalize a layout using different algorithms for different sections. While this semi-automatic approach takes considerable user input the results can be effective. The basic premise of allowing the user to interact with large graphs in conveniently sized pieces is an innovative use of partitions. While not a detail-in-context approach it actively supports creation of layouts that contain user selected foci.

Layout-independent fisheye graphs [114, 115, 116] provide a fisheye view without requiring a previous layout as a starting point. Space is allocated in a bottom-up manner and focal point size is calculated from a DOI function and sibling and parent size. This allows for different layouts for separate branches of the tree. As these fisheye views are developed independently from notions of physical distance, adjusting focal points can cause a total reorganization of the image.

Lens methods are closely related to detail-in-context methods. In fact, a detail-in-context focus can be explained as a non-occluding magnification lens. Lenses generally include broader possibilities for focus adjustment than simply magnification. They also do not attempt to visually integrate the focus with its context. As a result the regions are at best merely visually adjacent and will cause loss of local context if the adjustment to the focus requires more display space than the initial focal presentation.

Toolglass [15] and Magic Lenses [149] are transparent widgets [14, 64] that operate between the presentation and the cursor. They incorporate lens filters [15, 149] through which the representation will be altered according to the functions embodied in the filters. By providing the interactive ability to change the representation for the region covered by the lens, these widgets allow users to explore multiple options without requiring multiple views. Various types of focus-and-context presentations are also possible; for instance, edge separation can be exaggerated through a filter for a particular region of the image. One important contribution these approaches make to effective use of screen space is that they create widgets that require no additional screen space.

2.2.6 Zoomed Contexts

While *full zooming* (magnifying and compressing the entire representation to scale) was initially considered as just part of a windowing environment the zooming idea has been extensively revisited. There is now a family of approaches that have combined the goals of maintaining spatial constancy with the ability to zoom. Interfaces based on zooming are starting to be referred to as ZUIs. Figure 2.16 illustrates a zooming appraoch. With ZUIs context is provided through interactivity.

Pad/Pad++ [11, 10, 120] is a highly zoomable interface that operates on a theoretically infinite two-dimensional space that can be shared among users. As an aid for spatial memory the distance relationships within the information space are maintained. Their version of semantic zooming uses techniques similar to cartographic generalization to create alternate symbolic mappings for eliding information as representations become small. Pad++ offers rapid interactive zooming instead of detail-in-context presentations. Spreading information out and not re-organizing it spatially could interfere with the ability to compare information that was placed in separate locations and the ability to look at context and detail simultaneously. “Portals” and “sticky” objects address these needs. A portal is like a “hole” that allows the user to look through the Pad++ surface onto other regions on the same surface or other Pad++ surfaces. A portal is created by selecting a region and relocating

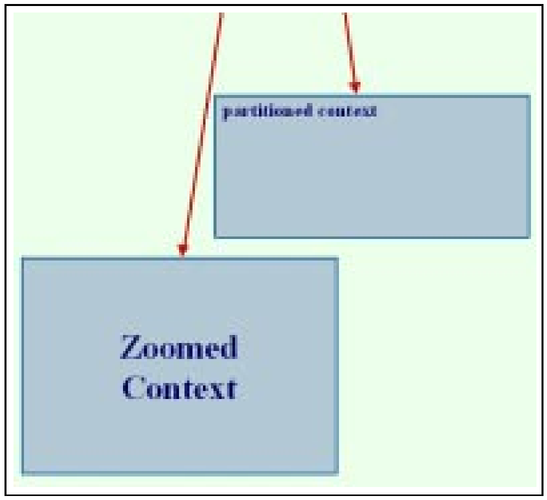

Figure 2.16: A diagram of a zooming appraoch

it. Portals are a visual copy and share the same information. Editing through either the portal or the original region has the same effect and is visible in both. A sticky object sticks, metaphorically, to the glass of the computer screen instead of the the Pad++ surface. Therefore when the Pad++ surface moves the sticky object remains in view.

A ZUI provides not just an infinitely (within the bounds of a computer's capabilities) large space on which one can spatially organize information, but also an infinite amount of space between any items. This new spatial freedom intensifies navigation problems [55]. Jul and Furnas [78, 79] call the problems of navigating in an electronic environment that is devoid of cues, critical zones in a desert fog. Their suggestion, ZTracker [78], indicates where the information is with visual cues in the form of outlines. However, it does not yet address what the information is.

In MuSE [56] Furnas and Zhang explore questions of authoring in ZUI environments. The premise is that issues of operating in multi-scale environments extend to authoring needs. For example, a snap to grid exists in z as well as x and y.

A different use of zooming is Powers of Ten Thousand [96] which creates a layered interface. Each layer has a different magnification factor, with the top layer being the most magnified. All but the base layer are partially transparent. These layers are independent and zoomed and panned over the original image. The idea is that the base image will provide context and the zoomed image overlays are the focal points. This idea is interesting in that in the physical world people use the ability to focus at different depths of field. However, without motion these presentations are very difficult to decipher.

2.2.7 Summary

While actually being able to see the context is important, the type of visual connection between the focus and its context is also significant. There are two principle directions in the literature in regards to the maintenance of a unified presentation. One group [7, 109, 131, 138, 139, 147, 150, 158, 166] favours preserving scaled only regions. However, to do this requires abrupt changes in scale resulting in the regions being simply visually connected. The others [20, 50, 59, 80, 82, 84, 85, 99, 108, 132, 137] prefer smooth visual integration which is usually achieved through a distortion function.

There are trade-offs that arise. Context is compressed to provide space for a magnified focus. Eventually, continuing to compress the context makes it unreadable and therefore no longer providing visual context. The desire to maintain complete context therefore limits the degree of magnification. Also, maintaining scaled-only regions results in abrupt visual changes between regions. The introduction of distortion can visually integrate regions but literally distorts the information.

2.3 Comparing Visual Capabilities of Detail-in-Context Methods

The last section discussed related work in terms of the way that context was maintained. Another major concern of presentation research is the flexibility of the focus. This is essentially how the user can examine the details of the information representation and concerns such things as the number and type of foci and the nature of the interaction with the foci. These visual capabilities are the focus of this section.

2.3.1 Freedom of Focal Shape

Any notion of ideal focal shape is tied to nature of the information, the chosen representation, the task at hand and the current user's preferences. As none of these items are fixed factors, the freedom to choose the shape of the focus is considered important. A user may be interested in a focal point or region depending on the nature of their data and their task. A region of interest could be any arbitrary shape, for instance a river's drainage area. Furthermore, it is useful to be able to switch between differing focal types.

Table 2.2 compares detail-in-context methods by the choice of focal shapes offered and Table 2.1 lists the meanings of the abbreviations used point (Pt), node (Nd), rectangular (Rt), convex (Cx), concave (Cv) and poly-line (Ln). A point focus can be a vertex in a graph or a point in an image or a point in the representation space. For Table 2.2 a focus is considered to be a point focus if it is

Pt	point focus
Nd	Node focus
Rt	Rectangular focus
Cx	Convex focus
Cv	Concave focus
S	Single focus
M	Multiple foci

Table 2.1: These are the meanings of the abbreviations used in Table 2.2

treated as a single point in an image, whether it is located on a graph vertex or otherwise. A focus is considered to be a node focus (Nd) if the node in the focus is treated as an object and magnified as a unit. Rectangular (Rt), convex (Cx), concave (Cv) foci are all region foci. A region focus is magnified to scale within the region that has been specified as the focus. A poly-line (Ln) focus maintains uniform magnification only along its focal line. Compression and/or distortion starts orthogonal to and immediately adjacent to the poly-line.

MVP [108], Rubber Sheet [138], and CatGraph [82, 83] offer two distinct methods; an orthogonally based distortion provides square or rectangular foci and a radially based approach provides point foci. However, this is not the same as having a single distortion method that incorporates freedom of choice in focal shape. Sarkar et al. [138] first made the call for the inclusion of arbitrarily shaped foci. Their morphing-based [138] approach approximates arbitrarily shaped foci with convex polygons and does offer user choice between point or convex polygon foci.

Our approach, 3DPS [20], offers user choice between point, line, poly-line, and arbitrarily shaped polygons. Users can draw the shape of the focus interactively. Subsequently, Non-linear Views [85] and Focal Line [59] offer the same freedom of choice.

2.3.2 Multiple foci

Most of the methods that initally included multiple foci had caveats. Polyfocal Display [80] was not interactive at the time, though with current computers it could be. Several of the orthogonal compressed-context approaches offer multiple foci [138, 139] but as discussed in Section 2.2.1 cause unrequested ghost foci.

The Zoom family [7, 8, 140, 141] of approaches are also based on orthogonal distortion but avoid

	Focal Shape and Number								
	Date	Pt	Nd	Rt	Cx	Cv	Ln	S	M
PolyFocal Display [80]	78	○							○
Bifocal Display [147]	82			○				○	
Generalized Fisheyes [52]	86	○						○	
2D Bifocal [93]	89			○				○	
Topographic Networks [70]	89	○	○					○	
Perspective Wall [99]	91			○				○	
MVP (Fisheye, Biform) [109]	91	○	○						○
Treemaps [74]	91		○					○	
Graphical Fisheyes [137]	92	○	○					○	
CatGraph [82]	92	○	○						○
Stretch Tools [139]	92		○	○					○
Document Lens [132]	93			○				○	
Variable Zoom [140]	93	○	○						○
Layout-Independent [114]	93	○							○
Pad++ [10]	93	○	○	○				○	
Magic Lens [149]	93			○				○	
Rubber Sheet [138]	93	○	○	○	○				○
Table Lens [131]	94		○	○			○		○
Hyperbolic Display [91]	94	○						○	
3DPS [20]	95	○	○	○	○	○	○		○
Continuous Zoom [7]	95	○	○						○
SHriMP [150]	95	○	○						○
Non-linear Views [84]	96	○	○	○	○	○	○		○
Focus Line [59]	96	○	○	○	○	○	○		○
Magnification Fields [85]	97	○	○	○	○	○	○		○

Table 2.2: This table compares the choice of focal shape and indicates whether or not multiple foci are supported. The meanings of the abbreviations are listed in Table 2.1

Fx	Fixed focal point
DOI	Degree og interest function
User	User selection
Sc	The option of scaled-only or not
Sx	the ability to re-size a selected focus
Mg	the ability to change the amount of magnification
Rs	change location by re-selecting
If	stationary focus, information moves
Rv	drag a focus through the information

Table 2.3: These are the meanings of the abbreviations used in Table 2.4

the ghost foci issue through a looser interpretation of the maintenance of orthogonality. Nodes stay within a range defined by their neighbours. However, the algorithm is global therefore each subsequent selection of a focus will influence existing focal choices. This prevents achieving equivalent magnification in serially selected foci and makes returning to previous configurations difficult. SHriMP [150] a closely related approach has similar results. As nodes are uniformly displaced and then scaled to fit, to obtain several foci of the same magnification one has to ask for them at the same time. If working sequentially, each subsequent request for a new foci will reduce somewhat the scale of the previously established foci. Recently Van Heyden et al. [166] have been investigating to what extent multiple levels of magnification in foci pose a problem for viewing series of medical MR images. Various alternatives to the layout adjustments in regards to orthogonality, parallelism and uniformity of scale are suggested.

Sarkar and Brown's [137] morphing-based approach also achieved multiple foci with caveats. If large foci are too close to each other or at the edge of the image they cause local information reversal. As this information reversal is an extreme violation of preserving overall shape, the foci are constrained in order to avoid this problem. This constraint interferes with freedom of focal location, and user control of level of detail. Due to these caveats they recommend the orthogonal approach as described in the same paper.

	Focal Interaction									
	Date	Fx	DOI	Us	Sc	Sz	Mg	Rs	IF	Rv
PolyFocal Display [80]	78			○			○	○		
Bifocal Display [147]	82	○							○	
Generalized Fisheyes [52]	86		○				○	○		
2D Bifocal [93]	89		○	○				○		
Topographic Networks [70]	89		○	○			○	○		
Perspective Wall [99]	91	○							○	
MVP (Fisheye, Biform) [109]	91		○	○			○	○		
Treemaps [74]	91		○				○	○		
Graphical Fisheyes [137]	92		○	○			○	○		
CATGraph [82]	92		○				○	○		
Stretch Tools [139]	92		○	○			○	○		
Document Lens [132]	93			○			○	○		○
Variable Zoom [140]	93		○	○			○	○		
Layout-Independent [114]	93		○	○			○	○		
Pad++ [10]	93			○	○	○	○	○		○
Magic Lens [15]	93			○			○	○		○
Rubber Sheet [138]	93			○	○	○	○	○		
Table Lens [131]	94			○			○	○		○
Hyperbolic Display [91]	94	○							○	
3DPS [20]	95			○	○	○	○	○		○
Continuous Zoom [7]	95		○	○			○	○		
SHriMP [150]	95			○			○	○		
Non-linear Views [84]	96			○	○	○	○	○		○
Focus Line [59]	96		○	○	○	○	○	○		○
Magnification Fields [85]	97			○	○	○	○	○		○

Table 2.4: This table compares the modes of selection and interaction with foci. The meanings of the abbreviations are listed in Table 2.3

2.3.3 Focal Interaction

Table 2.4 compares the same detail-in-context methods by the type of focal interaction offered. Table 2.3 explains the meanings of the abbreviations used. Table 2.4 compares three issues: selection, detail and content. Selection is concerned with how the location of the focus is specified. Detail looks at what type of control the user has over the level of detail presented in the focus. Content examines the methods by which a user can change the content of the focus.

There have been three basic methods by which a user can select a focus: One method provides no choice and the location of the focus is fixed (Fx). Another uses a DOI, where the user expresses interest in some aspect of the information and the algorithm makes the adjustments (DOI). Alternatively. the user selects the location of the focus in the image (Us). The user selected location can be used create a focus either in conjunction with a DOI or solely with a distance based function.

Issues compared under control of level of detail in the focus include: the whether the user has the option of choosing between uniform scale or not (Sc), the possibility of interactively changing the size of a focus once it has been selected (Sz) and the user interactive control of degree of magnification (Mg). There are situations in which it is important that the focus is magnified to scale, for instance, when distances in the representation have meaning such as in a map or architectural blueprint. There are representations based on connectivity such as graphs where ensuring scaling might not provide advantages and may use up valuable space unnecessarily. Many approaches allocate the amount of space a focus uses and consequently the level of detail it displays based on a composite DOI function. These methods can provide advantages by creating balanced views and allowing a system to incorporate information, such as alarms, from input other than the user [7, 8]. However, the user does not have full control. Other approaches allow the user to expand the focus until sufficient detail is revealed [108, 138]. These allow user control of relative magnification but not to specify the degree of magnification.

Sarkar and Brown's [138] morphing-based [12] approach offers the choice of whether or not the scaling is preserved in the focal region. 3DPS [20] also offers this choice and can allow the user to specify the exact degree of magnification [28]. For those approaches based on 2D transformations obtaining an appropriate transformation given a chosen magnification is not trivial [85, 94]. Keahey and Robertson [85] have developed an iterative approach that allows them to derive a transformation given a magnification field.

The last three columns in Table 2.4 compare the methods by which a user can change the content of a focus. The most common method of changing the content of a focus is to have the user select a new focus and create a new presentation (Rs). This method has caused user disorientation between presentations and given rise to discussions around supporting a users mental map and the creation of animated transitions. Another method is to move the information into a stationary focus (If), basically allowing the user to stretch the information so that the chosen portion is in the focus. Alternatively, roving search is supported by allowing a focus to be moved across the information (Rv).

2.4 Beyond 2D

The previous section reviewed literature was concerned with detail-in-context presentations that adjusted existing two-dimensional representations. This section reviews related research for three-dimensional representations.

SemNet [43] was the first exploration of detail-in-context presentation for three-dimensional graphs. In this paper the authors discuss three different possibilities. The fist method notes the implicit fisheye provided by perspective projection of the 3-dimensional layout. This creates a single focal point for the information that is in the foreground. The second method uses semantics for node positioning and creates an octree. With this octree a focal region is displayed in full detail and more remote sections of the octree are displayed in progressively larger chunks. Changes in the display occur when the focus is move across an octree boundary; small changes occur as a focus crosses a low-level boundary and larger changes occur at higher level boundaries. These sudden changes can be disconcerting particularly when the user is unaware that a boundary is approaching. To give some indication of where these boundaries are, nodes are highlighted as a boundary is approached. The third method uses a sampling density approach that samples fully around the focus and less as the distance from the focus increases.

Mitra's [110, 111] fisheye approach for aircraft maintenance diagrams employs an interactive filtering approach. These diagrams are three-dimensional radially-exploded views of aircraft assembly parts. A filtered view is based on the function of parts rather than proximity in the diagram. The user can adjust the threshold level creating views with more or less context.

Ware et al. [4, 170, 171] performed studies to explore the usability of space in a three-dimensional layout as opposed to a two-dimensional one. They discovered that people's accuracy in finding paths in networks was the same for 3D graphs of size $3X$ as it was for 2D graphs of size X. They also examined the problems in creating comprehensible 3D graph displays. In GraphVisualizer3D [51]

(GV3D) they attempt to express the informations semantics through semiotic[1] principles. This work is significant in its studies of the usability of 3D networks and the exploration of the use of semiotics in a 3D computer space.

Cone trees [133] layout hierarchical information in a three-dimensional tree structure. All of a nodes' children (displayed as cards) are hung from a partially transparent cone like a carousel. The transparency allows detection of the nodes that are behind other nodes. These cones rotate allowing selected nodes to be brought to the front. The rotating cones together with perspective provide a natural single focus detail-in-context view. Fractal trees [87] extend the three-dimensional cone layout by using fractal dimension to provide views of huge hierarchies.

One result of our framework, Visual Access Distortion [21, 26, 36], uses distortion techniques to provide visual access to the interior regions of a three-dimensional layout. This method is discussed in Chapter 7.

3D Magic Lenses [168] extend the ideas in ToolGlass [15] and Magic Lenses [149] into three-dimensions, allowing selected regions of a 3D layout to present alternate representations concurrently within the existing presentation. These lenses can be either flat or volumetric. A flat lens exists between the viewer and the representation and effects all of the representation that lies behind it, while a volumetric lens affects only a selected volume. As in 2D Magic Lenses, while the lens is situated in context, the lens region and the context region of the presentations are visually adjacent, changing abruptly from one representation to another. If the representation in the lens region uses more or less space than the representation in the context region, there will be visual discontinuities between the regions.

3D Zoom [123, 129] applies the Continuous Zoom [7] algorithm to three-dimensional models. Similarly to the Variable Zoom this approach is tied to the data model. The model is considered as having distinct parts; each part fits in its own bounding box. As chosen parts of the model are enlarged other parts are compressed to compensate. The differences between enlarging and compressing while maintaining the part's aspect ratio are discussed. This approach was developed to support visual exploration of medical images. The authors note that the ability to adjust the relative sizes of separate parts of the model in context seemed useful in revealing the relationships between these parts.

Kurzion and Yagel [88, 89] introduce distortion to volume visualization, describing a method, the deflector, that can be used to deform any object that is ray traceable. During ray casting, their

[1] Semiotics is the study of signs and the way that they are used [2, 3, 13]

method curves the sight rays according to a space deformation. These curved rays create an illusion of a deformed screen without actually adjusting the object.

2.5 Analysis Literature

There are an increasing number and variety of techniques that have been developed to address the spatial or screen real estate aspect of presentation space. However, the proliferation of terms and differences in approach have made it difficult to understand how these individual methods relate to each other. Several taxonomies have been suggested, which themselves vary in scope and focus. This variability may well stem from a lack of common ground regarding computational presentation. Presentations can be characterized by many different aspects such as; the data they can handle, the representation that is used, the tasks they support, the type of interactivity they offer, and their spatial organization.

Taxonomies are developed to help us understand differences and relationships within a field of research. A taxonomic characterization illuminates the field from a particular prespective, tending to emphasize the importance of certain aspects over others. Blackwell and Englehart [16] in their taxonomy of diagram taxnomies note that, rather than trying to create a “super-taxonomy” from which to judge other taxonomies their intention is to make the range of options more familiar and to encourage diversity. We also see the importance of understanding this multi-disciplinary field from many perspectives. Blackwell and Englehart’s [16] are classifying other taxonomies on the basis of how people think about diagrams Their six dimensions of comparison are;

1. The representation or the visual graphic display. This is considered to have two aspects: the graphic domain or vocabulary and the graphic structure or the visual spatial relations.
2. The message or the represented information (the information domain and structure).
3. The relation between the representation and the message (pictorial, realistic or abstract, and analogical correspondence).
4. Task and process (interpreting and modifying representations).
5. Context and convention (culture and communicative context).
6. Mental representation (diagrams in the mind).

Item 4 incorporates the ability to reorganize a representation while interpreting it, or in our terms, presentation space. However, it also includes tasks to be performed and interactivity issues. This section examines previous taxonomies that classify screen real estate research within this category and proposes a taxonomy based on presentation dimensions.

2.5.1 Examining Previous Taxonomies

A taxonomy can classify the appearance of existing presentations. Early presentation taxonomies such as Leung [93] in 1989 and Misue and Sugiyama [108] in 1991 fall into this category. Leung [93] divided presentation possibilities into: two screen (separate regions for detail and context), split screen (separate but visually adjacent regions), bifocal (compressed uniform peripheral context, as discussed in Section 2.1), and fisheye (compression that increases with distance from focus, as discussed in Section 2.2.4). Misue and Sugiyama [108] classify visual presentations as: detailed (displays full details always), image-switching (single window, switches views), multi-view (multiple views in separate windows) and fisheye (which in their taxonomy includes all presentations discussed under Sections 2.2.1 through 2.2.4). They note that only the fisheye views satisfy the first four of their requirements and that the fifth is domain specific.

This type of taxonomy divides existing individual approaches into categories according to the main presentation transformation supported. Our division of Section 2.2 into: compressed full contexts, filtered sufficient contexts, filtered and distorted sufficient contexts, distorted full contexts, partitioned contexts, and zoomed contexts. can be considered as such a classification. It groups research into families that create similar visual presentation patterns. However, a taxonomy that is written to create groupings for existing approaches may well break down as research progresses.

In contrast Misue and Sugiyama's [108] presentation requirement specifications and Plaisant et al.'s [122] task taxonomy which base classification on human tasks and capabilities, may remain useful for a longer period of time as humans tend to change less rapidly than computational methods. Misue and Sugiyama's [108] suggested set of requirements are: detailness (all of the details in the representation are accessible visually), wholeness (the entire representation is visually accessible), simultaneity (parts of the representation do not have to be accessed sequentially), image-singleness (supports the interpretation of the whole), and appropriateness (the presentation is suitable for the representation).

Plaisant et al. [122] describe two taxonomies for 2D image browsing: a task taxonomy and a method taxonomy. Their task taxonomy is: generation (creating, editing, annotating), exploration (examine an unknown information space), diagnostic (scanning, looking for patterns, making comparisons), navigation (locating objects or regions), and monitoring (surveying systems, keeping track of security, safety, and system failure). Their method taxonomy falls into the first category in that it classifies existing approaches. It makes an initial division into presentation and operation then divides presentation into static and dynamic. The static categories relate most closely to the other taxonomies listed here and are: multiple views, which in turn is divided into various organizational and tiling methods, and single views which is divided into detail-only, zoom-and-replace and fish-eye.

Alternatively, a taxonomy can classify the transformations that can be used to develop presentation methods. A classification of this type tends to remain useful in that while new transformations may be developed, those transformations that have been used will still exist as part of the space. Both Spence's [146] and Noik's [117] taxonomies are based on how the transformation from one presentation to the next is achieved. Furthermore, as these are classifications of transformations that can be used to create a presentation method, they can be considered generative in that it is possible to try new combinations of transformations.

Spence's [146] taxonomy for presentation techniques describes four orthogonal transformations: two distortion transformations (one in x and one in y), one encoding transformation (replacing objects with a symbol), and one threshold transformation (filtering according to a user defined value). A given presentation method may then make use of more than one of these transformations. For example, using Spence's terms, Zoom [7] uses both x and y transformations and a threshold transformation.

Noik [117] narrows the field from computational presentation in general to those types of presentation that use a single image, providing what he terms a presentation emphasis taxonomy. He defines an emphasized presentation as one that contains one or more regions of interest that are distinguished visually from the rest of the image in some manner. Emphasized presentations are divided into implicit (achieved without distortion through perspective viewing), filtered (displays a subset of the elements), distorted (distorts size, shape and position), and adorned (emphasis is provided through visual additions such as colour or motion).

Leung and Apperley [94] describe a taxonomy of presentation techniques in general and a more specific taxonomy for distortion-oriented techniques. The distortion technique taxonomy is described under the category of distortion presentation dimension in our taxonomy to folow Section 2.5.2. Their presentation taxonomy is a hierarchy and includes some aspects of the information visualization pipeline [30] in that it starts by considering large volumes of data (Figure 2.17). They

Data	Graphical	distorted (detail-in-context)	encoding spatial transformation
		non-distorted (detail without context)	zooming windowing
	Non-Graphical	distorted (detail-in-context)	data suppression
		non-distorted (detail without context)	clipping paging

Figure 2.17: This figure is derived from Leung and Apperley's Fig 1 in [94]. It shows their taxonomy of presentation techniques for large graphical data spaces. The arrow indicates that non-graphical representations are sometimes re-interpreted as graphic representations

consider data as either inherently graphical or non-graphical and recognize that sometimes non-graphical data is represented graphically. They continue to consider graphical and non-graphical information spaces separately. Both of these spaces are divided into distorted and non-distorted views. Graphical distorted views include encoding and spatial distortion, and non-graphical distorted views include data suppression. Graphical non-distorted views include zooming and windowing, and non-graphical non-distorted views include paging and clipping.

2.5.2 Our Taxonomic Dimensions for Presentation Space

Presentation space consists of possible transformations from one presentation to the next. These possible transformations can be grouped into *presentation dimensions*. A presentation dimension is a continuum with a domain of transformations causing range of effects. We identify seven presentation dimensions and note their range.

Presentation Dimensions

1. **Partition** is the act of dividing a given presentation into components or regions. The degree of partition affects the construction of a presentation. The construction of the presentation is the arrangement and connection of items or groups of items within the presentation. The construction of a presentation ranges from each of the components being presented individually (as clearly separated items) to the other extreme where all the components are integrated into a single unified image.

A common method for organizing space is to create divisions. In our homes and offices many variations of partition are used to organize space. Items are placed in boxes, files, shelves, etc. There are analogous methods for dividing up a computer's display space. Lines, planes and volumes can be used to partition display space and these can be further sub-divided; for instance, a plane can be divided into rectangles as in Treemaps [74]. A space is usually partitioned into sub-spaces by creating borders or frames around regions of the space. Partitioning the components of the presentation involves keeping them in or out of these defined regions. Windows are the most common computational presentation use of partition.

Use of partition includes Leung's [93] two-screen and spilt-screen, Misue and Sugiyama's detailed, and multi-view, and Plaisant et al.'s [122] various tiling methods and detail-only categories. It is not mentioned in either Spence [146] or Noik's [118] taxonomies.

2. **Aggregation** is the act of changing a given presentation by collecting components or regions of a presentation into a groups. Again in the extreme case all items are unified into an integrated single item. The degree of integration of a presentation is the degree to which all the components of the presentation are visually coherent. Aggregation is considered a distinct from partition in that they can be used independantly within a given approach. For example, Windows is a use of partition and tilings are groupings of windows [81].

   Objects can be combined or grouped; lined up in rows, piled up, spread evenly, or packed together. Items can be organized in rows on toolbars, in columns in menus, and in various tilings and organizations. Further aggregation can place items in situ, combining two or more items into a unified region.

3. **Distortion** is a spatial re-organization that changes the original proportions. A distortion will change at least one of the geometric relationships of size, angle, parallelism, orthogonality, proximity or topology. Changes in size only or scaling is at one end of the distortion continuum. Evidence of this can be taken from the fact that some researchers in presentation space do not consider changes in scale to distort information while others do [166]. This causes us to consider scaling as the most minimal form of distortion. At the other end of the distortion continuum are complex transformations that affect all of the presentation's geometric properties. In the range inbetween many variations of distortion exist with changing degrees of visual and interpretive complexity.

- **Scale** An object that changes in scale changes in size while maintaining the other geometric relationships. Magnification and compression are simple forms of distortion in that they are merely changes in scale.

- **Complex Distortions** There are many ways of introducing complexity into a distortion. For instance, a distortion transformation can be applied along any combination of the x, y, or z dimensions of the representation, or dimensions of the presentation space, or independently of either of these (see Chapter 7). Examples include; a one-dimensional distortion (Bifocal Display [147]), a three-dimensional distortion for linear information (Perspective Wall [99]) and a three-dimensional distortion for two-dimensional information (Document Lens [132], 3DPS [20]). Furthermore, distortions do not have to operate in Euclidean space, for example they can use of hyperbolic space [112]. Increased distortion complexity of does not necessarily involve creating unreadable presentations. Instead it is possible that investigations in this direction may lead to new methods of visual access [23].

  Leung and Apperley [94] discuss distortion as a transformation function. The slope of this distortion transformation function indicates the degree of magnification. Their distortion taxonomy is based on the derivative or magnification function. They divide distortions into piecewise-continuous and continuous, then further divide piecewise-continuous into constant and varying. Piecewise-continuous distortions are a relatively simple distortion based on step functions. Constant piecewise-continuous results in the same grouping as compressed contexts in Section 2.2.1. Varying piecewise-continuous includes approaches such as Perspective Wall [99] and Document Lens [132]. Continuous distortions are those discussed in Section 2.2.4 and provide visually integrated presentations.

  A type of subtle but complex distortion, called *displacement* [136] in cartography, involves a transformation that preserves distinctions in the information. For example, if two roads have no intersection it is important to keep the visual separation even if at the current scale the distance between them is less than the line thickness with which they are drawn. With careful distortion this spatial distinction can be maintained even at a reduced scale. These types of adjustments, while preserving information integrity, are a form of distortions; their computational application is being explored [59, 123, 157].

Distortion includes Leung's bifocal [93], and Spence's [146] x and y transformations, and

Noik's [118] implicit and distorted categories.

4. **Simplification** is the act of removing some aspect of the representation. Not all of the representation need be presented simultaneously. If the information presented is too cluttered one option is to remove components that are not currently pertinent. Simplification can be performed through thresholding via a numeric value such as those used in a DOI (Generalized Fisheye [52]), or by filtering via semantic relevance (G+/Graphlog [69]). Components of the representation can be fully removed [52] or partially removed [9, 96], or individual details can be simplified.

   This idea relates closely to the cartographic notion of *generalization* [136] which is concerned with how best to maintain information integrity at any chosen scale. Insufficient space can make it difficult or impossible to display all the details of a representation. However, what is presented should have integrity. That is, it should appear to be a smaller or larger version of the same representation. Sometimes it is possible to maintain the important aspects of a representation while removing considerable detail. For instance, a line on a map, which signifies a road with several minor deviations, can be usefully shown as a straight line at a smaller scale.

   Simplication includes Spence's [146] threshold category and Noik's [118] filtered category.

5. **Augmentation** is the selective visual addition that is intended to increase the clarity of the presentation. A given presentation of a given representation can be augmented by making any aspect or component of the representation more visible. This includes re-introducing filtered or elided components, displaying details which had been simplified, and enhancing aspects of the representation to make them clearer, more noticeable or in otherwise visually enhanced. Examples of augmentation include the use of colour as a highlight [117], and increasing detail in selected regions as in SemNet [43].

   Augmentation includes Noik's [118] adorned category.

6. **Motion** As presentations can be changed over time, presentation space includes the possibility of motion. Access to the temporal axis is one of the more significant characteristics of computational information presentation in contrast to printed information graphics. Motion can be incorporated as an integral part of the presentation [5, 6]. The use of motion ranges from an abrupt transition to the use of motion that in itself adds meaning to the presentation.

- **Transitions** The simplest transitions are a direct change from one presentation to the the next. In this case it is the static presentations that are intended to be interpreted. Motion only exists as a means of changing from one presentation to another. Animating these transitions is a small step along the continuum of possible uses of motion in that it makes the new presentation more understandable because its relationship to the previous presentation is clearer to the user. However, the motion itself is not being used to convey meaning.

  Inbetween these two extremes use of motion includes: moving an image (scrolling, panning, or zooming), moving around partitioned sections, and moving through time as in animation. These different types of motion can be applied to either the whole image or subsections of the image.

- **Semantic use of motion** Motion can be added to a presentation in a symbolic manner. For example, it can signify the onset of an alarm, it can indicate a grouping [35] and can draw paths [148] and make connections.

Motion includes Misue and Sugiyama's [108] image-switching, Plaisant et al.'s [122] zoom and replace categories.

7. **Illumination** Lighting has been perhaps one of the most under-utilized computational presentation variables. While its importance is well noted in theatre at this point in time it is sparingly used in computational presentations. This is at least in part due to the computational expense of lighting. The illumination continuum extends from no use or uniform lighting through complex use of lighting to highlight and emphasize regions of interest to use of lighting that has semantic value.

   - **Uniform** A computational presentation can create uniformly lit presentations as do many current presentations.

   - **Semantic use of lighting** Minimal computational uses include such things as shading buttons and window frames to give a slight three-dimensional appearance.

     Lighting is essential to creating convincing three-dimensional graphics. We use lighting to disambiguate distortion patterns (see Chapter 6).

The only category mentioned in earlier taxonomies not included in these seven presentation dimensions is Spence's [146] and Leung and Apperley's [94] encoding. Encoding is a process of abstraction and as such involves changing the representation. A different abstraction of the information involves a change in the concept that is being represented. For example, Cox and Roman [38] note that computer programs can be thought about in several levels of abstraction: direct (statement to statement), structural (code structure as in nesting), synthesized (the result of the process), analytical (loop invariants), and explanatory (meta-level understanding). As discussed in Chapter 1, we make a distinction between representation and presentation. Presentation variables vary the type of access to a given representation. They do not change the representation itself. Abstraction as representation adjustment is discussed thoroughly in Strothotte [157].

Table 2.5 indicates the use that many of the approaches discussed in this chapter make of these presentation dimensions. As these dimensions can be used to greater or lesser extent a small symbolic slider is used to indicate approximate extent of utilization.

	Presentation Dimensions in Use							
	Dt	Part	Agg	Dist	Simp	Aug	Mot	Ill
Star [76]	78	———o–	o———	o———	–o———	o———	———o–	o———
Bifocal [147]	82	–o———	———o–	–o———	–o———	o———	–o———	o———
Gen FE [52]	86	o———	———o–	o———	———o–	o———	o———	o———
Topo Net [70]	89	o———	———o–	———o–	———o–	o———	o———	o———
P Wall [99]	91	–o———	———o–	—-o—-	o———	o———	–o———	o———
Treemaps [74]	91	———o	o———	o———	–o———	o———	o———	o———
Grph FE [137]	92	o———	———o–	———o	–o———	o———	o———	o———
D Lens [132]	93	–o———	———o–	———o–	o———	–o———	—-o—-	o———
Zoom [7]	93	———o–	———o–	—-o—-	—-o—-	–o———	–o———	o———
Pad++ [10]	93	o———	———o	o———	o———	o———	———o–	o———
M Lens [149]	93	———o–	—-o—-	o———	o———	———o–	–o———	o———
R Sheet [138]	93	o———	———o–	———o–	o———	o———	o———	o———
T Lens [131]	94	–o———	–o———	—-o—-	o———	o———	–o———	o———
HB Dply [91]	94	o———	———o	———o	o———	o———	–o———	o———

Table 2.5: Some examples of the use of presentation dimensions. The amount of use is approximately indicated from very little or none (o———) to extensive use (———o)

This notion of distinguishable dimensions which can be combined in various ways, parallels Bertin's *visual variables* [13]. A visual variable is a basic visual characteristic that can be used alone or in combination to create visual representations. For two-dimensional printed graphics Bertin notes these visual variables as: the mark (a point, line or an area), the position of the mark (location in x and y), and the perceptual qualities of the mark. The perceptual (or retinal) qualities of the mark are; colour, size, value, orientation, texture and shape. Understanding these visual variables supports the creation of useful visual representations. Bertin's visual variables are representation space variables in that they describe basic building blocks creating representations for two-dimensional printed graphics.

The transformations that compose these presentation dimensions can be used alone or in combination to create presentations. They are the basic building blocks from which presentations can be created and adjusted. Identifying basic building blocks and understanding their constraints and potential contributes to the development of a design paradigm. However, as the capabilities of computational presentation are still in a state of flux this set of variables may well require future revision.

2.6 Discussion

There has been considerable research in computational presentation in general and in developing a detail-in-context solution in particular. Even though there is general agreement as to what features a complete detail-in-context approach would include, none of the methods thus far had offered a full set. We provide a framework that have supported the development of a fully functional detail-in-context presentation method, explains existing presentation methods and suggests new possible methods and directions.