Chapter 1

Introduction

Use of the computer as an information medium is expanding daily. The rapid proliferation of the World Wide Web is creating an unprecedented universal information exchange. Information visualization is playing an increasing role, in that businesses, schools and research are using it to communicate, teach and promote insight. Creating visualizations of information involves developing visual representations and presenting them within the capabilities and limitations of the selected medium. For instance, a given set of information can be presented very differently on film, on a billboard, or on the radio. A good understanding of the chosen medium is an important factor in the creation of effective presentations [13, 41, 47, 160]. With continuing advances in hardware and software technology the possibilities inherent in presenting information on a computer are in a constant state of flux. A full understanding of the potential of a computer as an information medium will, of necessity, evolve with the development of the computer itself.

Historically, a new medium usually undergoes several developmental stages. The computer has proved no exception. The first of these stages, sometimes called the 'dancing-bear stage' [119], refers to the effect of bringing a dancing bear into a small village for the first time. The novelty of the act alone will bring a lot of attention. With a computer, dancing-bear effect is a recurring phenomena. During the second or 'copy-cat' stage the new medium is used to mimic an existing more familiar medium. For instance, film was initially used as taped theatre and the camera as instant painting. During this stage it is common for discussions to erupt as to the ability of the new medium to replace the old. Eventually, as understanding of the unique potential of the new medium develops, it receives recognition in its own right.

An interesting twist in this history as it applies to the computer is that it has been possible to extend the 'copy-cat' stage by using the computer not just to mimic one previous medium but to perform the functions of many media that have previously existed separately. For example, a computer can operate as a calculator, a typewriter, and a filing cabinet. Through this, we are learning that one of the computer's distinguishing features is its versatility. However, perhaps there has been a tendency to dwell in this stage, encouraging the development of what have been termed metaphors. In contrast, this thesis is directed towards increasing our understanding of what is unique about the visual presentation of information on a computer.

Contributions to the development of a new medium can be made by discovering some technique that can only be done with that particular medium. For example, morphing [12] is a uniquely computational ability. Also, understanding the limitations in comparison with other media is significant [51, 170]. Alternatively, developing a meta-understanding of the basic components of a medium and the manner in which they relate will provide creative freedom. For example, a knowledge of vocabulary and grammar is needed to write. For visual media, this type of analysis is less common, perhaps more difficult, and certainly more controversial. However, there are some examples such as Bertin's [13] analysis of printed information graphics, Dondis' [41] book on visual literacy and McCloud's [104] explanation of comics.

The research in this thesis addresses one specific aspect of information visualization: that of developing an understanding of the variations possible for computational information presentation. *Presentation* is the act of spatially organizing a given display appropriately for the task at hand. It is defined in terms of the information visualization field in Section 1.2 and in terms of presentation problems in Section 1.3. Limiting the scope in this manner has allowed us to develop a fuller understanding of one step in the information visualization process. As a contribution to the emergence of the computer as a medium in its own right, we develop an improved understanding of computational presentation space. For instance, before using a new tool, even one as simple as a pencil or a brush, an artist will test it to gain a knowledge of the characteristic range of marks that can be made. Just as an artist benefits from knowledge of the tools they are using, a person creating an information visualization for a computer will benefit from fuller understanding of the possible presentation choices.

This introduction proceeds as follows. Section 1.1 notes that a framework can be important in developing a fuller understanding of presentation possibilities and that this is an appropriate moment in information visualization research to develop a unified presentation framework. Section 1.2 uses the information visualization pipeline to position this research in terms of other research in the field. Section 1.3 explains the recognized issues that exist in a computer's presentation space and the use of the term *elastic* to describe a computer's presentation space. Section 1.4 covers common terminology. Section 1.5 outlines the concepts behind our framework. Section 1.6 states the contributions of this research and Section 1.7 explains the organization of the thesis.

1.1 The Motivation for Developing a Framework

An information visualization application is usually designed for a specific set of information and a specific set of tasks. The process of designing a visualization involves understanding the nature of the information and the nature of the tasks to ensure that specific needs are met. Most visualizations, both information and scientific, have been approached in this manner, and the research community has learnt a lot from this process. New methods have been discovered and refined, expanding our awareness of the scope of what is possible. Presently, within this space of possible visualizations, many successful "point" techniques exist. In fact, the number of individual successes has led to an interest in developing an understanding of the space itself.

A good theoretical framework will provide a structure that helps to define the space, supports description of possibilities and provides guidance in developing variations. To date, most information visualization research has been concerned with producing visualizations for specific data. This work has led to several individual successes and a variety of independent visualization techniques, resulting in an increasing number of point solutions. While these point solutions may not yet span the full space, they are sufficient in number to consider an analysis of the information visualization space. The benefits of developing a meta-understanding of the information visualization space include:

- classification of existing visualizations through the recognition of the differences and relationships between visualizations;
- knowledge of the component parts that have been used to build visualizations and how they can be combined to create new visualizations; and
- recognition of the process involved in creating a visualization which facilitates more in-depth study of individual sections of the design space.

Research towards developing an information visualization analysis drew initially from information graphics and scientific visualization. Tufte [160, 161, 162] shows how using critical assessment can reveal potential traps and pitfalls. Bertin's [13] semiotic analysis of information graphics describes how relatively few components, or in his terms visual variables, are used to create printed presentations, and how an understanding of their characteristics enables the creation of more readily understood information graphics. Ware [169, 172] and Mackinlay [100] have examined adjusting Bertin's analysis to cope with the differences between printed graphics and computer graphics.

More recently Card and Mackinlay [19] proposed a framework for information visualization design space, and Chi and Riedl [30] developed an operator interaction framework that is based on an extension of Card and Mackinlay's work. Both of these proposed frameworks tackle the difficult task of describing the full information visualization space. In contrast, Tweedie [165] and Chuah and Roth [31] examine a sub-region and develop frameworks for the interactive aspects of visual information exploration. In this thesis, we also examine a sub-region of information visualization design space, developing a framework for visual presentation space.

Narrowing our focus this way has allowed us to describe a geometric framework which is both explanatory and generative. This framework provides a structure that can explain existing point presentation solutions. It allows extrapolation between solutions and in doing so describes novel presentation possibilities.

1.2 Positioning Presentation Space in the Information Visualization Pipeline

Basically, information visualization is involved with creating visual displays on a computer that are intended to provide support for problem-solving by externalizing information. There are a considerable number of steps before the information or data becomes a visualization. Research in visual presentation space forms only a small component of information visualization design process. Chi's [30] information visualization pipeline (Figure 1.1) is expanded to position presentation space in the context of other information visualization research. There are two non-exclusive modes of interacting with an information visualization: authoring or the interaction of creating it, and the users interaction of interpreting the display. Figure 1.1 draws an information visualization pipeline from the perspective of authoring.

The first step in the information visualization process is to create a representation of the raw data that can be stored in the computer. The term representation is used as defined by Marr [101] to mean a formal system by which the data can be specified. For instance as in Marr's example, Arabic, Roman and binary representations can be provided for the number thirty-four, giving 34, XXXIV, and 100010 respectively. Defined in this way a given representation provides specific information about the data and differing representations more readily reveal particular aspects of the

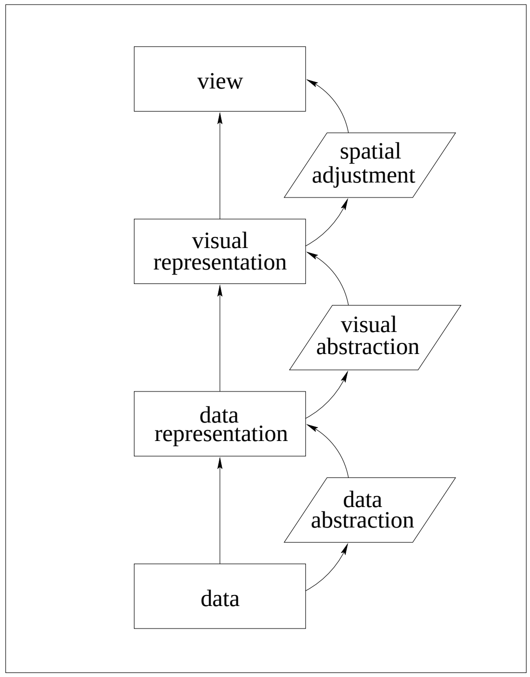

Figure 1.1: The information visualization pipeline, varying only in terminology from that published by Chi and Riedl [30]

data. For example, arabic numerals reveal information about powers of ten while binary representations reveal information about powers of two. This data representation is often arrived at through a process of abstraction. For example, it is the process of abstraction that selects which aspects of the data are to be the most accessible. The data representation will not necessarily have a visual form; therefore a second process of abstraction may be required to create a visual representation. This meta-representation may have several visual forms. For example, a hierarchical graph can be displayed in many ways including treemaps [74, 75], cone trees [133], disc trees [72] and several node and edge layout configurations [39]. We expand the next step, called the visual mapping transformation by Chi and Reidl [30], making a distinction between representation and presentation.

At this point in the process the visual representation can be displayed as a basic image. This basic image directly corresponds to the information, for example a drawing of a graph. However, in displaying the visual representation, factors such as the size limitations of the physical display are encountered. Very often a visual representation of real world data is bigger than can be fully displayed on a computer screen. It may be necessary to change the presentation, taking into consideration such things as the size of the display, the areas of interest and the task at hand. For example, a map of Vancouver may be presented with a work route magnified to reveal street names. The techniques that allow adjustment of a visual presentation, without interfering with the information content of the representation, form a *presentation space*. These transformations re-organize the representation spatially to allow users greater freedom in visual exploration. Figure 1.2 positions presentation space between visual representation and the final view.

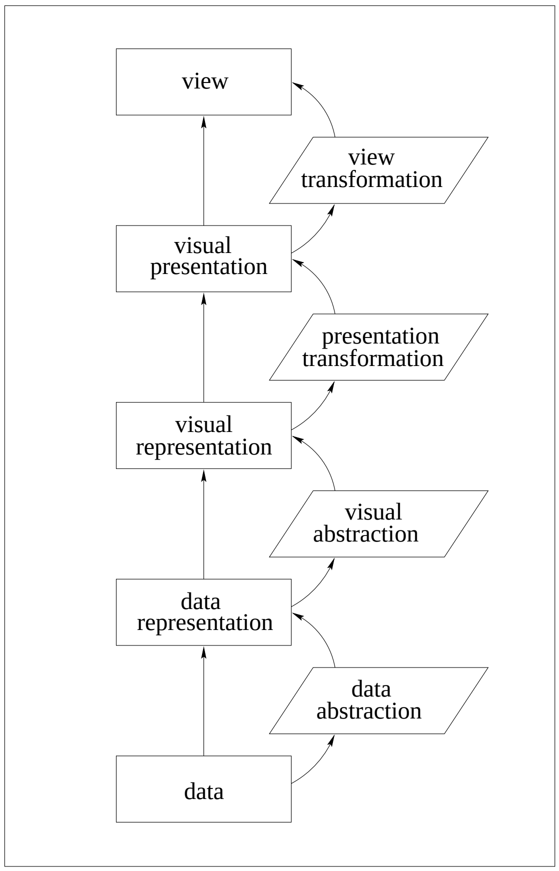

Figure 1.2: The expanded information visualization pipeline, including presentation space

To discuss this from the perspective of the user, the distinction between view and value [30] operations is useful. A value operation changes the information that is stored in the data representation and a view operation merely changes what is currently displayed. For example, removing data is a *value operation* and rotating a graph is a *view operation*. However, this separation is not always distinct nor is it necessarily obvious to the user. The extent to which the different operations affect either the information or the visual display varies. For instance, filtering maybe done simply to temporarily clear visual clutter in a display or actually be applied to the data itself, removing certain aspects of the information. As defined here, presentation operations are view operations.

Users readily consider rotation, pan, scroll, and zoom as view operations. This may be in part because they relate to real world view operations of turning and moving (from side to side, up and down or further and closer). While presentations transformations are intended to merely present the user with a more useful view for the current task, the visual changes do not always relate to view changes in the real world. This can leave users uncertain about the change that has occurred and about how to interpret a given presentation. Creating new presentation techniques carries with it a responsibility to ensure that it is possible to read them as such. The user should be able to use various presentation transformations and remain convinced the information has not been changed in the process.

1.3 Issues in Presentation Space

As the primary metaphor for computer use shifts from an extension of one's personal desktop to a form of access into a vast information space, viewing this expanding information space through relatively small computer screens becomes increasingly problematic.

A wide variety of observed and simulated data has resulted in many visual representations. Discrete representations such as graphs and vector drawings are used in a great variety of fields. In computing alone they are used to express such things as visual languages, software, hypertext, natural language parsing, and databases. High resolution raster information is being generated scientifically and artistically as well as being collected from many disparate sources such as satellites and radar. A growing library of analysis and manipulation tools make it advantageous to work with these representations on computers. However, these representations are seldom small enough to fit on a computer's display space without compression. This discrepancy between a computer's display space and its information space has been associated with problems in navigation, interpretation and recognition of relationships between items in the representations. Making the best use of this display space has been an important issue in user interface design since the introduction of video display terminals. While there is research into alternate display technologies [155, 159], video display terminals are still the primary interface to the computer. The necessity for effective solutions to this problem has intensified as technology has advanced, with the ability to produce visual data continuing to outstrip the rate at which display technology has developed. This issue is referred to as the *screen real estate problem*.

The introduction of windows [76] was the first notable advance in presentation solutions. This overlapping partitioning of two-dimensional space has greatly increased the amount of usable display space. However, even with the inclusion of scrolling, panning and zooming it has become apparent that the ability to examine details of the representation often conflicts with the ability to maintain global context. Zooming-out, or compressing the data to fit within the space of the screen, can result in an image that is too dense to discern detail. Figure 1.3(a) shows a land usage map of Champaign, Illinois compressed uniformly to fit a given frame. This image shows the full map but much of the detail is difficult to see as it is too compressed. Zooming-in, or magnifying the whole image, provides a detailed reading but results in the loss of context because only a sub-region will fit in the available display space.

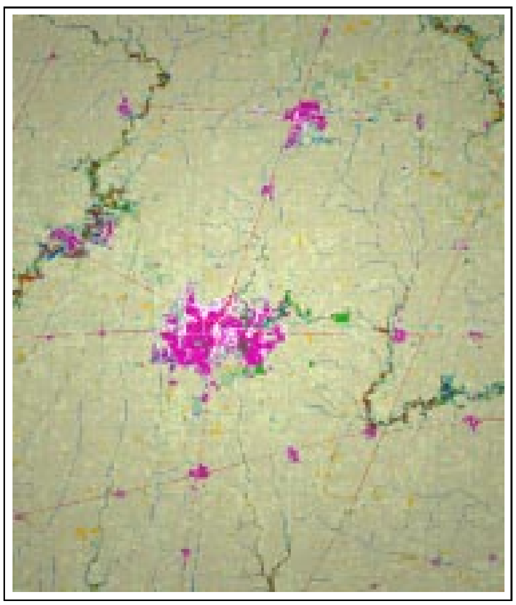

(a) A single scale view of a land usage map of Champaign, Illinois (for image credit see Appendix C.1)

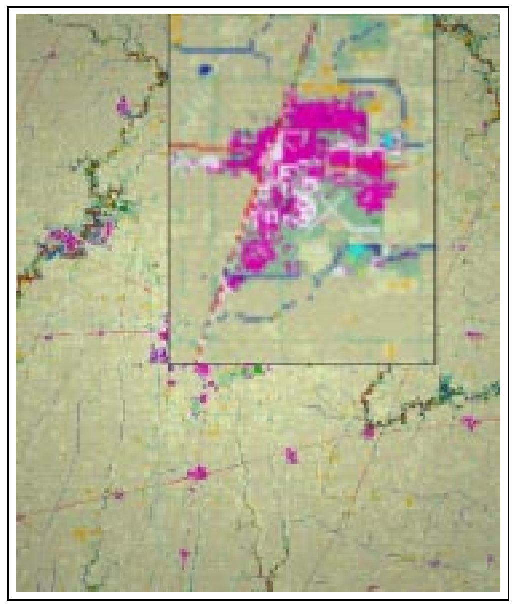

(b) The Champaign, Illinois map with a magnified inset

Figure 1.3: Two different presentations of a land usage map of Champaign, Illinois

Examining the representation with panning and scrolling has been compared to peering through a keyhole onto a vast display of information. Panning and scrolling allow movement of the information across the keyhole but require the user to keep track of their location. This is one of the factors that has led to discussions of being *lost in computer space* [105, 143].

Creating an inset, by zooming-in or magnifying a sub-region in place, obscures local context.

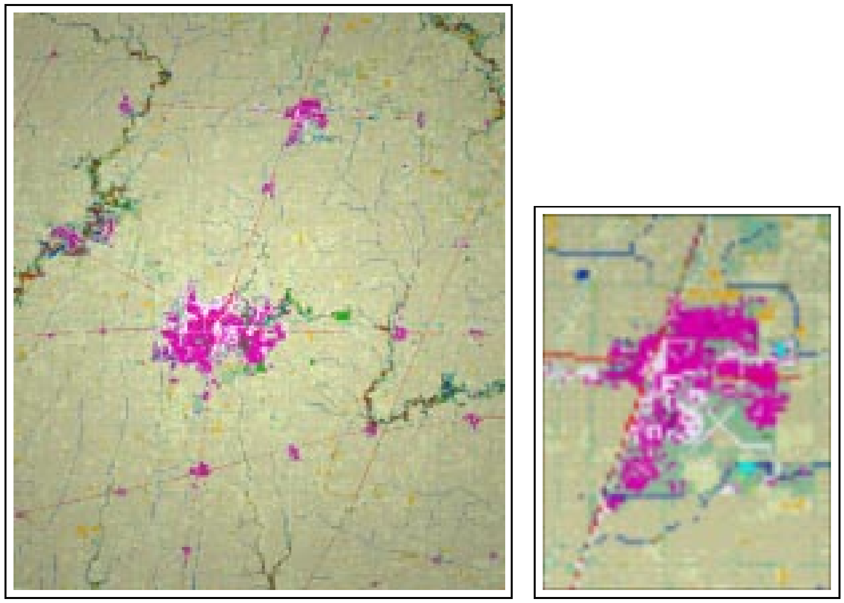

(a) The single scale view of a land usage map of Champaign, Illinois

(b) A separate magnified view of a selected region of the map

Figure 1.4: A presentation containing two distinct views, showing the map of Champaign, Illinois in one and a selected region magnified in the other

Figure 1.3(b) shows a view with a magnified inset. The inset provides detail for the selected region but the space required for magnification causes the adjacent regions to be covered, making it impossible to see how the details, for instance roads, in the inset connect to the roads in the rest of the map.

Multiple views in separate windows allow global structure to be displayed in one view and the required detail in another. Figure 1.4 displays a magnified sub-region separately in its own frame. This solution removes the occlusion in Figure 1.3(b), however the connections between the two images are not necessarily obvious and must be performed consciously by the user. The only situation where detail can be viewed within its context is when the entire image will fit without compression into the display space.

The phrase *detail-in-context* is defined as the ability to see a *focus*, or chosen region of the representation, in sufficient local detail while it is still set in its global context. The difficulty with supporting detail-in-context readings in a windowing environment has led to several techniques that

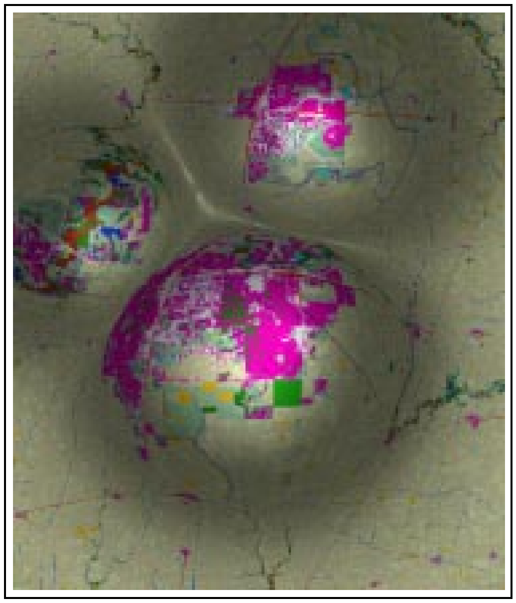

Figure 1.5: A three foci detail-in-context view of Champaign, Illinois. Shading is used to reveal the distortion (see Chapter 6)

combine the advantages of zooming-in with those of zooming-out. Essentially these techniques allow a user to magnify chosen sections to reveal the desired detail and compensate for the extra space this magnification requires by various types of compression in the rest of the image. Figure 1.5 shows a detail-in-context view with three foci. The advantages attributed to these techniques include:

- increases in the amount of information that can usefully be presented on a computer screen;
- human preference for remembering and presenting information in this manner [52];
- utilization of visual gestalt by retaining the perception of the information space as a single event; and
- increased user performance in path finding tasks [70, 140].

These detail-in-context techniques usually create a new presentation that includes several spatial variations from the initial presentation. Consider the fact that objects in the presentation can change in size, or appear to change in size, by the use of view/value operations. In our physical world we are familiar with the fact that a magnifying lens simply makes an object appear larger without altering the object itself. An object is said to be *magnified* if it is enlarged in appearance. This distinction has become less clear with use of the word magnification in presentation on the computer. This is partly because discussions often blur the distinction between altering the computational representation, a value operation, and altering the computational presentation, a view operation.

For instance, compression is a frequent topic in computational literature where it usually is concerned with the reduction of actual size or physical storage space. A key point is whether or not it will cause information that was in the representation to be lost. In contrast an object is visually *compressed* if its apparent size or volume is reduced, therefore information will not be lost, it will merely appear smaller. An object is *scaled* if it is magnified or compressed in a manner that maintains its proportions. A change in scale maintains geometric relationships of angles, parallelism, orthogonality, relative proximity and topology. An object is *distorted* if it is spatially adjusted in a manner that changes at least one of the geometric relationships.

Given these definitions a change in scale is not a distortion, while changes in magnification and compression can be distortions or simply changes in scale. There is a more general definition of distortion that declares that any change in an object is a type of distortion [166]. We consider changes in scale to be a minimum distortion (see Section 2.5).

Humans have the potential for *visual gestalt*. That is, our visual perception allows us to develop an understanding of the whole that is more than the sum of the parts. This factor, which makes integrated visual access so appealing, has added a new twist to the screen real estate issue. Evidence from a number of studies in experimental psychology points out an intriguing paradox as to how humans combine sensory information. It appears that while monitoring multiple sensory events seems to demand a great deal of cognitive resources, monitoring sensory input that is precieved as pertaining to a single event is relatively simple [46, 102, 103]. In other words we are capable of detecting changes in pattern, texture, and overall structure in an integrated manner, developing an understanding not derivable by summation of the parts. For example, plotting statistics on a graph reveals patterns that would be arduous to pick out from a list of numbers. However, to gain this advantage we need to see the whole image, preferably as a unit. If the desired detail view can be provided in a manner that smoothly integrates it into the global context, then it preserves the possibility of visual gestalt.

In cognitive science there is considerable discussion of *cognitive load* or the relative degree of effort involved in a cognitive task. It appears that integration of separately obtained pieces of information is cognitively very difficult, often causing people to make mistakes, including saying one thing while doing another. This brings one to wonder if high-level conscious mechanisms are effective for maintaining space constancy (i.e. the mental model of perceptual space surrounding the perceiver). The answer, according to many psychologists, is that once low-level mechanisms have broken down, secondary cognitive efforts are insufficient to restore normal function [17, 48, 127].

This supports the idea that if detail and context are provided in separate multiple views then increased cognitive effort is needed for integration. While the user may be cognitively aware that views in multiple windows pertain to a single information space, perceptually they remain distinct. The effort of maintaining which detail belongs where has to be performed consciously by the user. The effort in synthesizing separately obtained pieces of information compromises the benefits offered by a global image.

However, one advantage of separate views in multiple windows which has previously been overlooked within detail-in-context research is the ability to move and reposition individual views. This is often used to align images, or selected sub-sections of the image, in separate views to facilitate visual comparisons. Traditionally this has meant the use of magnified views in sub-windows which are moved independently of the original image and hence have no direct visual connection to the rest of the image. The framework developed in this thesis can describe both movable separate views and unified detail-in-context presentations. This has allowed us to provide the freedom to reposition magnified regions without detaching them from the rest of the image (see Chapter 4).

A number of techniques have been developed that try to simultaneously address the detail-in-context problem and allow presentation of larger amounts of information. Though some of these techniques do achieve their goals, there has not been wide-spread acceptance. While this current lack of acceptance may merely be a factor of time, there are issues with the detail-in-context approaches in general. First, to date they have lost one of the great advantages of windowing, the ability to reposition chosen sections freely. Second, they violate many of the information design guidelines that have been developed over the centuries - i.e. they distort information. Third, windowing and detail-in-context have seemed mutually exclusive. While it is possible to place a detail-in-context approach in a window or to organize windows with a detail-in-context method, it has not seemed possible to integrate the advantages of one with the other.

1.4 Terminology

While detail-in-context is one of the more commonly used terms, it is by no means the only term that has been used to describe research in this area. No single term has yet emerged as identifying the whole field. Furthermore, there are several names for varying types of detail-in-context views.

The term *fisheye view* [52] is an analogy to a fish-eye camera lens which is highly curved, allowing light to enter the lens from a wide angle. Therefore a fisheye view has central magnification set in a background that becomes increasingly compressed as the distance from the centre increases.

This term is sometimes used as a general term and as such has been used to describe many presentations that diverge quite widely from the original camera analogy. For example, *Bifocal views* [147] have two levels of magnification, creating a clear visual division into regions of two separate scales.

A *filtered view* removes some aspects of the representation to provide more space for the presentation of detail in the focus. Filtering maintains only a partial context. A *presentation emphasis view* [118] uses varying visual effects, possibly including magnification, to draw attention to a selected focus. For instance, a section could be emphasized by changing the colour or through use of motion [5, 6]. The term *layout adjustment* [107, 150] explicitly states that what is being done is the re-organization of an existing layout.

A *multi-scale view* incorporates several scales in a single image. Fisheye views and bifocal views are multi-scale views. Detail-in-context views and presentation emphasis views are not necessarily multi-scale views. A *distortion view* makes use of distortion to create multi-scale views. A distorted view will contain at least one region in which the scale is not uniform. A *non-linear* distorted view is a distorted view in which the rate of distortion itself changes.

1.5 The Concept of Elastic Presentation Space

In order to develop a general framework for elastic presentation space, we examine the presentation problem independently from the application. The term *elastic* is used because it implies both the ability to be stretched and deformed and the ability to return to its original shape. With a computer, a space of presentation possibilities exist, including the ability to dynamically adjust a presentation. Aspects of the computer's elastic facility have been utilized in the creation of several existing techniques, for instance, Stretch Tools [139], Rubber Sheet [138], Malleable Graphics [32], Pliable Surfaces [20] and, more recently, Elastic Labels [71] and Elastic Windows [81]. The term elastic reflects the resilient deformability that appears to be one of the distinguishing characteristics of a computer's presentation space.

To limit the scope of the problem, we restrict our exploration of elastic presentation space to two-dimensional representations. There are several reasons for this decision. Two-dimensional representations are very plentiful, they are the most commonly used in computational display, and there is a lot of research about the creation of two-dimensional representations [13, 160, 161, 162]. Symbols and the two-dimensional positional relationships between them are a powerful method of encoding information. Charts reveal relationship between two components, diagrams can display concepts and processes, and graphs portray relationships among many entities. Two-dimensional methods for displaying graphs include a variety of positional organizations that use small marks for the entities or nodes and connecting lines for the relationships between the entities. Graphs can also be represented topologically by using an area to represent the entity and containment to represent the relationships. Maps are common two-dimensional representations where the Euclidean distance in the representation has symbolic, positional and relationship interpretations. Also, we first encountered the screen real estate problem when considering the presentation of two-dimensional representations and often extrapolating results from lower to higher dimensions is more amenable than the reverse.

Most research into screen real estate has concentrated on specific applications, exploiting their underlying information structure to obtain reasonable displays. In this work we take a different approach, examining the display problem independent of the application. In particular, we divide the display problem into two components, representation and presentation. Since representation is inherently dependent on the information, we do not consider this part of the problem. We instead concentrate on presentation and assume that we start with a valid 2D representation. Our approach to the presentation problem is to map the representation onto a surface in three dimensions and use perspective projection to create the final display. Manipulating the surface transforms the two-dimensional flat surface into a three-dimensional curved surface, allowing control of magnification and organization of image details.

The three-dimensional nature of this deformation approach has several advantages. Using perspective projection to view the three-dimensional surface provides the possibility of maintaining magnification to scale in chosen areas and of controlling the organization of the compression and distortion in the context. The process of magnification in elastic presentation space by what appears to be the action of pulling a selected region towards oneself in order to see it more clearly is directly analogous to a natural physical action, providing a useful metaphor for one's actions. Further, if sufficient visual support is provided about the nature of the manipulated surface's three-dimensional form, the relative magnification and compression should be readable. It has been established that humans can discern three-dimensional shape from shading alone [130, 138], and there is considerable evidence to support the fact that this is a low level pre-attentive skill [86]. Such a low level visual routine will interfere less with conscious processing and may even provide an aspect of the interface that requires little or no learning [169]. Our challenge in this regard is to provide appropriate visual cues that will access these pre-attentive abilities [24].

1.6 Contributions

The framework presented in this thesis describes existing presentation methods, identifies new presentation variations, and provides simple methods for combining them. This framework removes some of the current either/or choices, allowing a designer of a new information visualization to chose a combination of presentation methods that best suit the information and task needs. This framework:

- Allows for the creation of a detail-in-context solution that fulfills the list of functionality that has been suggested in literature as desirable. Previous research had recognized the detail-in-context problem and delineated the functionality that might exist in an ideal detail-in-context approach. In comparison with the suggested list all previous solutions had caveats.

- Creates detail-in-context methods independent of the information's characteristics. Previously, considerable effort had been placed on creating displays appropriate for particular types of information. The resulting solutions can be applied to any two-dimensional visual representation. In particular, this includes raster image data. Previous methods for viewing raster image data did not include the possibility of detail-in-context viewing.

- Explains previous presentation approaches and their inter-relationships. Existing presentation approaches such as windows, full zooming environments, and various distortion approaches, create visual displays that vary considerably visually and algorithmically. This framework provides a way of understanding how these seemingly distinct approaches relate to each other. Furthermore, it provides a method of relating them algorithmically, allowing the inclusion of more than one presentation approach in a single interface.

- Allows extrapolation between the presentation approaches it describes. Being able to relate previously distinct presentations methods has allowed creation of approaches that exist in the spaces between the current point approaches. Of particular interest are the approaches existing between insets and detail-in-context views, between full-zoom and detail-in-context views and between separate views and detail-in-context views.

- Extends distortion viewing to include re-positioning of foci, which we call *folding*. Folding allows spatially separated focal regions to be repositioned while maintaining their information contents and without disconnecting them from their context. We know of no other work that examines the repositioning of foci within a detail-in-context view.

- Extends the advantages of distortion viewing into three-dimensional representations. The ideas from this framework have been applied to 3D representations creating a novel visual access method that deals effectively with occlusion and is capable of providing three-dimensional detail-in-context views.

Identifying the components that comprise the distortion paradigm and understanding their effects provides a basis for developing new techniques and for more readily tailoring existing techniques. Through our framework we describe several new presentation paradigms, for instance, 3DPS [20], Folding [25] and Manhattan Lenses. We do not make any claims about one being better than the other. This framework was intentionally developed independently of the information's characteristics in order to form an understanding of presentation space in general. This thesis offers a description of elastic presentation space, creating a palette of techniques from which choices can be made. We believe these choices need to be made in careful consideration of the information, the visual representation and the intend users and their tasks. We are involved with such application work, for instance, providing detail-in-context for the DNA representation H-curves [92], exploring screen usage for MR images [166] and creating a visualization environment for landscape dynamic simulations, Tardis [27]. These applications are being developed with user involvement and are beyond the scope of this thesis.

1.7 Thesis Organization

This section describes the organization of the thesis. Chapter 1 introduces the presentation problem and positions presentation space in terms of other research in information visualization. Chapter 2 examines research into effective screen usage as that is the problem we address through our framework for *elastic presentation space* (EPS). Chapter 3 describes our geometric framework EPS and explains how this framework was used to develop a detail-in-context approach, 3DPS [20]. Chapter 4 introduces three new EPS concepts: distortion control, folding and the use of adjustable distance metrics. Chapter 5 steps through existing and new presentation possibilities describing them in EPS terms. Chapter 6 discusses the comprehension issues involved with distortion viewing. Chapter 7 applies EPS to visual representations of different dimensions. Chapter 8 concludes this thesis by outlining the contributions and discussing future directions.

Throughout this thesis the ideas discussed are illustrated with examples from several prototypes. Appendix A lists these prototypes and the people involved in creating them. Many but not all of the variations discussed in this thesis have been implemented. Most of these implementations are simple prototypes built to provide proof of concept or to resolve some issues visually.