

Technical Report

Taxonomy For Discrete Lenses

Petra Neumann^{*}

University of Magdeburg

Department of Simulation and Graphics

Sheelagh Carpendale[†]

University of Calgary

Department of Computer Science

December 16, 2003

^{*}pneumann@cs.uni-magdeburg.de

[†]sheelagh@cpsc.ucalgary.ca

Abstract:

In this report we introduce new ways of thinking about different types of visual data distortion methods. In particular, we clarify terminology that can be used to distinguish between two visual distortion methods, discrete and continuous lenses. The set of presentation techniques that result from application of these distortion methods is extended by this exploration. We describe a theoretical framework for use in future research on these presentation techniques. This involves the exploration of general distortion effects and how these effects are utilized by the visualization tool designer. Finally, the usability of the concept of discrete lenses for developing new interaction techniques is demonstrated by example applications.

Contents

1	Introduction	6
1.1	Motivation	6
1.2	Problem Statement	6
1.3	Context	7
1.3.1	Innovations in Visualization Laboratory	7
1.3.2	Resources	7
1.4	Methods	7
1.5	Evaluation	8
1.6	Overview of this Essay	8
2	Concepts and Definitions	8
2.1	Information Concepts	8
2.2	Information Visualization	9
2.2.1	Visualization Techniques	10
2.2.2	Information Visualization Pipeline	10
2.3	Presentation Space	11
2.3.1	Screen Real Estate—a Presentation Problem	11
2.3.2	Elastic Presentation Space	12
3	Elastic Presentation Framework	13
3.1	Geometric Setup	13
3.2	Lens Setup	14
3.2.1	Focal Point	14
3.2.2	Focal Shape	14
3.2.3	Drop-off Function	14
3.2.4	Magnification	15
3.2.5	Distance Metrics	16
3.2.6	Multiple Foci	17
3.2.7	Blending	18
3.2.8	Additional Information	18
4	A Taxonomy for Discrete Lenses	18
4.1	Discrete Lens vs. Continuous Lens	19
4.1.1	Continuous Lens	19
4.1.2	Discrete Lens	22
4.2	Discrete Representations	22
4.2.1	Discrete Objects	23
4.3	Distortion Transformations	24
4.4	Distortion Calculation Within the Elastic Presentation Framework	25
4.5	Why Develop a Taxonomy	27
4.6	Development of the Taxonomy	27
4.6.1	An Approach Based on Effects	27

4.6.2	An Object-Oriented Approach	28
4.6.3	A Mathematical Approach	28
5	How to Create a Discrete Lens	28
5.1	Step 1: Object Properties	29
5.1.1	Distorting Geometric Properties	29
5.1.2	Distorting other Discrete Properties	34
5.2	Step 2: Reference Points	35
5.3	Step 3: Distortion Properties	36
5.4	Step 4: Distortion Output	36
5.5	Step 5: Mappings	38
5.6	Step 6: Apply Lens and Present	38
5.7	General Discrete Distortion Effects	38
5.8	Problems of Discrete Distortion	39
5.9	Student Ideas	40
5.9.1	Discrete-Image-Lens	40
5.9.2	Ansel-Adams-Lens	40
5.9.3	Discrete-Word-Synonym-Lens	41
6	Related Work	41
6.1	Bifocal Display and Perspective Wall	42
6.2	Generalized and Graphical Fish-Eyes	43
6.3	Nonlinear Magnification	45
6.4	Flip Zooming	46
6.5	Interaction with Distorted Views	47
6.6	Ongoing Work with the Elastic Presentation Framework	48
6.6.1	Bubble Lens	48
6.6.2	Semantic Zoom	48
6.6.3	Discrete Distortion in 3D	49
7	Conclusion	49
	Appendix	51
	List of Figures	54
	References	56

1 Introduction

This section will introduce the motivation and problem statement for the work presented in this thesis. It will also show the place and resources for the internship during which the ideas presented here were developed.

1.1 Motivation

Much of the research in *detail-in-context* presentation has dealt with the *magnification* and *distortion* of seemingly continuous structures like images or maps. SARKAR et al. proposed a *rubber-sheet metaphor* for describing these techniques [SSOR93]. However, much of the research conducted described continuous rubber-sheet techniques but implemented the distortion in such a way that not all aspects of the data were distorted according to the chosen distortion function. A good example is SARKAR and BROWN's implementation of a graphical fish-eye lens [SB92]. The transformation functions were applied to the nodes of a graph but not to its edges. Recently, a popular operating system Mac OS X implemented a discrete technique in its dock-bar to distort the discrete icons representing a menu (see Figure 1). This has lead to new popularity of the use of distortion techniques in HCI and information visualization. A distinction must be made between distortions which apply a technique to all graphical aspects of the data and those that apply the distortion to certain aspects of the data since each has different visual effects. This thesis introduces a distinction between continuous lenses and discrete lenses. Many specific examples for discrete lens distortion exist; however, it is possible to bring these examples together into a single, comprehensive taxonomy. Such a taxonomy would be beneficial to those using these techniques in HCI and information visualization.



Figure 1: Menu bar in Mac OS X.

1.2 Problem Statement

The problem at hand was to define a taxonomy for discrete lenses. The set of existing presentation techniques can be extended by such an exploration and create new ways of thinking about presentation with discrete lenses. The taxonomy was to be based on the Elastic Presentation Framework as presented by CARPENDALE in [Car99] and placed in a two-dimensional representation space. Working with the Elastic Presentation Framework which nicely integrates many detail-in-context techniques, our goal was to come to a general understanding of how the provided lenses could be used discretely on the underlying representations and to place

those techniques in a general context. Unlike other methods for discrete visual structures, the Elastic Presentation Framework lens library handles all representations identically, not taking into account certain aspects which prior work has considered. Orthogonality, proximity or topology have been identified to be necessary for the user to preserve a mental map when graphs are distorted, for example [SM95]. In this sense, the library provides a general way for distorting data points, other aspects are left for the developer to integrate.

1.3 Context

This research is based largely on previous research by CARPENDALE. In her PhD-thesis [Car99] she describes a framework for possible variations in information presentation, and therefore an analysis of visual presentation space. Her framework combined many previously existing ‘point’ solutions in detail-in-context presentation and integrated them into a geometric framework. The main contribution of this work was to unite these methods (and many new ones) algorithmically. Providing them through the application of one geometric setup and one calculation method an extrapolation is supported between the presented methods. Many different distortion effects can be created in a given interaction at the same time.

1.3.1 Innovations in Visualization Laboratory

The Innovations in Visualization Laboratory is headed by PROF. DR. SHEELAGH CARPENDALE. Graduate students, research assistants, and interns work on tasks that involve interactive spatial organization like information visualization with the EPF library and Interface Design. We will shortly describe the lab’s ongoing work with the EPF library in Section 6.

1.3.2 Resources

CARPENDALE and MONTAGENESE have developed an EPF library which is used for research in the area of detail-in-context data visualization. This C++ library implements the EPF data distortion methodology in [Car99]. The library is available for use under Linux and Microsoft-Windows. For our research, we took advantage of the EPF library and implemented the examples found in this essay with the help of C++, OpenGL, Glut and QT.

1.4 Methods

The work presented in this thesis was developed after a period of exploration of the given methods and possibilities in the used EPF library. An implementation and re-implementation phase followed during which the examples presented here were developed. As an assignment for the 581 “Human-Computer-Interaction II” lecture at the University of Calgary the students

had to find example uses for the presented discrete techniques with the help of an early version of the work presented here. Some results will be shown later in Section 5.9.

1.5 Evaluation

This thesis introduces new terminology to distinguish between two distortion methods, discrete and continuous lenses. This terminology is useful since the two methods are visually distinct. Methods to create discrete lenses are introduced which describe the most important aspects of discrete lenses. Some ideas about the space to which lenses can be applied had to be considered for future work since it was beyond the scope of this thesis.

1.6 Overview of this Essay

After this introduction of the problems, motivation and context of the work presented here Section 2 will show concepts and definitions necessary to understand the work presented later and will embed this work in a scientific context. In Section 3, the reader will learn the basic aspects of the Elastic Presentation Framework our work is based on. Section 4 will introduce the ideas on which the taxonomy is based and the approaches towards the creation of this taxonomy. Section 5 will show, how based on the concepts and definitions presented, a discrete lens can be built and what kind of choices are offered to the visualization tool designer. It will also show some examples of students' work. In Section 6, related work will be presented and set into relation to the given taxonomy. Finally, Section 7 will give a concluding statement and ideas about future work.

2 Concepts and Definitions

In this section we will describe relevant terminology and information necessary to understand the concepts and ideas behind our work. We will place our research in the context of the definitions given.

2.1 Information Concepts

No general definition for the concept of *information* exists. According to SHANNON [SW63], information corresponds to redundancy in a message, non-redundant data contains information while redundant data does not. Taking this definition one can say that information is a stimulus that is meaningful for the user and can so assist in decision making or forming of opinion. Therefore, there is a distinction between data and information. Information is to be gained e. g. by visualizing data.

NEWBY defines *Information Space* as the set of objects and relations among them held by a system in contrast to cognitive space which is the set of concepts and relations among them held by a human [New01]. An information space is, therefore, a type of information design in which information objects are located in a space in which location and direction have a meaning, so that mapping and navigation become possible. Components of an information space are usually documents and parts of documents (paragraphs, sentences etc.) and the relations between both. Classes of information spaces may include graphs, maps, charts, tables etc. which are the subjects of information visualization.

An *Information System* is defined as an organized set of documents or a technology aimed at realizing information processes [Eng02]. Therefore, it represents an entire infrastructure or organization including components for the collection, processing, storage, transmission, display, dissemination, and disposition of information and so is a part of an information space. A geographic information system, for example, enables the user to envision the geographic aspects of a collection of data. Basically, it enables the user to query or analyse the data and receive the results in the form of some kind of map.

2.2 Information Visualization

Information Visualization deals with creating visual or graphical aids for data we want to access, distribute or explain.

CARD et al. [CMS99] give a general definition for information visualization:

Information Visualization:

“The use of computer-supported, interactive, visual representations of abstract data to amplify cognition.”

In the literature, there exists a difference between visualization of data which is scientific, mostly physically based (*scientific visualization*) on the one side and the visualization of abstract, non-physically based and therefore not inherently spatial data (*information visualization*) on the other side. Our work is placed in the field of information visualization since the data we deal with is not necessarily acquired from scientific investigations or experiments. However the methods developed in this work are certainly applicable to scientific data as well. Since, inevitably, there is overlap between the two different types of visualization we will use *information visualization* as a term which describes means of creating visual aids which lead to insight in the underlying data sets. Information visualization in this sense is not about producing nice pictures but about making data understandable and explorable so that the visualization helps us gain knowledge about the data. It is the process of forming a mental model for the acquired data and so helping the viewer to understand underlying concepts, patterns, and connections within the data [Spe00]. Data sets in information visualization typically come from large information spaces or information systems which need to be made accessible and understandable to the user.

2.2.1 Visualization Techniques

The evolution of computers in our society has lead to many new means of producing visualizations for information we have gathered. Much of the work in this field focuses on creating graphical interfaces for complicated datasets, databases for example. Visualization techniques include concealment of data, using a three-dimensional space, layering data, scaling techniques to provide more space for certain information (e. g. using lenses), detail-in-context techniques, and taking advantage of psychological principles of layout, such as proximity, alignment, and shared visual properties (e. g., color). Often such visualizations are combined with means for interactive exploration of the given visualization or presentation as mentioned in the given definition.

2.2.2 Information Visualization Pipeline

In information visualization there is a strong separation between the value of data and the visualization of data. Several steps have to be taken to turn acquired data into a visualization. Figure 2 shows a visualization pipeline as seen in [Car99] which adapts a pipeline presented by CHI and RIEDL in [hCR98].




Figure 2: Visualization pipeline as seen in [Car99].

As a first step the acquired data has to be turned into a data representation. This pipeline draws a distinction between the representation and the presentation of data. Representation of data is defined by MARR as follows:

“A representation is a formal system for making explicit certain entities or types of information, together with a specification of how the system does this.” [Mar82]

Therefore, a representation is a formal system by which the information can be specified. Representations portray things or the relationships between them and can be found in many different forms. For instance, words and diagrams can be used to describe quantities, or graphs

and equations to portray the relationships between certain entities. MARR’S example in Figure 3 shows different representations of the concept of the number thirty-seven. The Arabic, Roman, and binary numeral systems are all representation methods for numbers. They reveal information about how the numbers are decomposed into powers of ten, five, and two. This example shows the different mathematical aspects each representation reveals.

37	XXXVII	100101
----	--------	--------

Figure 3: Representations for the concept of the number thirty-seven.

Representations are often created by an abstraction process of the raw data during which important aspects of the data can be selected. This transformation step creates meta-data, or descriptive information about data. This meta-data can later be helpful to choose appropriate visualizations [CMS99]. To arrive at a *visual* representation a *visual* abstraction step follows. For example, a table of Web-page hits per day can be turned into a visualization in the form of a pie-chart, column-chart, or bar-chart. The following presentation transformation step turns the visual representation into a visual presentation. A presentation of data is concerned with how the data is displayed on the medium of choice. In this step it is necessary to consider the possibilities and limitations of the medium at hand and the task that is to be encountered. Our work is to be placed at this and the following point in the pipeline in the field of creating presentations for given data representations. The last step in the pipeline is the view operation which merely changes the way the presentation is currently displayed. It is here where most of the user’s interaction takes place. The user has the possibilities of several different view operations like rotation, scroll, pan, zoom, or filtering.

2.3 Presentation Space

To understand the concept of a *presentation space* one has to take a look at a visualization pipeline as discussed in Section 2.2.2. A presentation space is formed by all the possible forms of changing a visual presentation without changing the information given through the representation of the data. The presentation space is positioned between the visual representation and the view in the given pipeline in Figure 2.

2.3.1 Screen Real Estate—a Presentation Problem

Today’s computers have the ability to store vast amounts of data. Many different types of data use up huge amounts of storage space. One just has to think about databases storing information about e-commerce transactions, satellite imagery, geological information systems, or medical information spaces. Representing such vast information spaces becomes more and more difficult since the size of our computer screens has not increased proportionally over

the past years. Indeed many devices like PDA's or cell-phones support even *smaller* displays. Such technologies try to give the user access to functionality similarly available on a regular PC through a much smaller screen. The issue of how to best use the available display space is one major issue in presenting information and is known as the *Screen Real Estate Problem*. As TUFTE mentioned, it is not important how much information is displayed but rather *how* it is displayed [Tuf90]. Problems to be tackled are:

- How to present information with sufficient detail
- How to present information with sufficient context

In the literature these two problems have often been called the *Detail-in-Context Problem*. Detail-in-context is a presentation technique which uses the ability to view part of the information up close while still seeing the remaining information. This includes the ability to present detail and related contextual information at the same time without enlarging the space occupied by the specified data. The detail-in-context approach is an alternative to approaches found in most conventional applications. The scrolling presentation method, for example, presents a small part of the data in detail while the other parts are hidden outside the viewing area.

Detail-and-context techniques are grouped into:

- *Windowing*, where context is provided in one window and detail in another. This makes it difficult for the user to integrate both mentally into one model.
- *Full Compressed Context*, where detail and compressed context are shown in one view. Spatial constancy is preserved.
- *Sufficient, Filtered Context*, where parts of the context are filtered out to provide more space for the focal area.
- *Full Distorted Context*, where detail and non-uniformly distorted context are presented in one view.
- *Partitioned Context*, where the screen space is divided into several regions in which detail or context are displayed.

2.3.2 Elastic Presentation Space

Elastic Presentation Space is a term defined by CARPENDALE in [Car99]. It is the notion of a space that is elastic and can, therefore, after certain transformations have been applied interactively, always revert to its original shape.

3 Elastic Presentation Framework

The main features of the Elastic Presentation Framework presented by CARPENDALE in [Car99] is that a two-dimensional representation is placed in an Elastic Presentation Space. The representation is put onto a surface which is itself placed in a three dimensional space. This surface is viewed using perspective projection and can be manipulated to change to different presentation styles. The presentation transformation mentioned in Section 2.2.2 is split in two steps: manipulation of the surface containing the representation and perspective projection. This section will introduce the basic geometric framework in which the elastic presentation is placed and show the different aspects and features contained in the Elastic Presentation Framework.

3.1 Geometric Setup

The geometric setup for the work presented here follows the setup as described in [Car99]. A perspective projection view frustum is set up by defining a viewpoint, a view plane and a base plane which holds the discrete representation to be presented. The setup can be seen in Figure 4. The picture is slightly adapted from the one in [Car99] to fit the actual implementation in the EPF library. The viewpoint and viewplane stay fixed once defined. They are referred by *reference viewpoint* and *reference viewplane*. The reference viewplane stays parallel to the x - and y - axis preserving parallelism, angles, and proximity within the representation. Moving the viewplane in z -direction scales it in size. A decrease in z produces a magnification effect while an increase produces a zoom-out effect.




Figure 4: Geometric setup.

3.2 Lens Setup

The detail-in-context approach provided by EPF incorporates several features which have been described as desirable to have in detail-in-context presentations. Figure 5 shows an image of the most important parts of an EPF lens each of which will be described further in the following sections.




Figure 5: Parts of an EPF lens.

3.2.1 Focal Point

The focal point identifies the point of maximum magnification and can be described as the center of the applied ‘lens’ for most focal shapes, although this might be different when the polyline focal shape is applied. Figure 6 shows the focal point identified by a circle for the different focal shapes supplied by the current implementation of the EPF library.

3.2.2 Focal Shape

Several different focal shapes have been implemented. The focal shape specifies the shape and size of the area of maximum magnification, also called the focus or focal area. It can be adapted to the specified task at hand. Figure 6 shows the five focal shapes supplied by the current implementation of the EPF library with a linear drop-off-function applied. The focal shape can be changed further by blending together several lenses as describe in Section 3.2.7.

3.2.3 Drop-off Function

Focal point and focal shape both select the place and area of detail. The drop-off function specifies how this area of detail is integrated in its context. Several different drop-off functions are supplied by the EPF library, each having its advantages and disadvantages of how it integrates the focus into the context. Figure 7 shows the different drop-off functions applied to a 2D grid in 3D space as seen from the side.




Figure 6: Different focal shapes and their focal points.

The *Gaussian Lens* has a characteristic bell curve. It integrates the magnified focus smoothly into its context. The standard deviation describes the exact location of the area of maximum compression.

The *Linear Lens* integrates the focus linearly into its context. This produces harsh transitions between the focal region and the drop-off and where the drop-off meets the context.

The *Cosine Lens* shows a more gradual integration of the focus into its context but compared to the Gaussian a steeper integration of the slope into the undistorted part of the presentation.

The *Hemisphere Lens* slowly integrates the focal point into its surrounding but the drop-off meets the context almost perpendicularly which causes an ring of high compression.

The *Manhattan Lens* does not distort any information that is not inside the focal region. It has a step drop-off function.




Figure 7: Different drop-off functions within EPF.

3.2.4 Magnification

Magnification within EPF is calculated through the use of perspective projection (see Figure 8). Setting the maximum magnification for a selected lens changes the height h_f by which the data point x_i is distorted. The point x_m on the base plane is the point where the distorted point x_h is perceived to be located. The EPF library calculates the distorted points in 3D yielding x_h or in 2D x_m . A simple mathematical calculation makes it possible to calculate h_f :

$$h_f = d_b - d_s \quad (1)$$

$$mag = \frac{d_b}{d_s} \quad (2)$$

$$(2) \text{ in } (1): h_f = d_b - \frac{d_b}{mag}$$

This function reduces the amount of translation towards the viewpoint with higher magnification leading to the desirable effect that points cannot be distorted beyond the viewpoint. The apparent lateral translations are calculated as follows:

$$x_m = x_i \cdot \frac{d_b}{(d_b - h_f)}$$

$$y_m = y_i \cdot \frac{d_b}{(d_b - h_f)}$$




Figure 8: Magnification calculation within EPF (image courtesy of Sheelagh Carpendale).

3.2.5 Distance Metrics

Within the EPF library points are distorted according to the distance from the focus and the selected drop-off function. Since the library takes discrete data points as input several so-called distance metrics are possible which often give different answers for the distance between the same pair of points. L_p -metrics provide a continuum between radial and orthogonal layout. The general formula for L_p -metrics is given as:

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

The four provided L_p -metrics within the EPF library are the common:

Chess Board Distance which assumes that you can make moves on a grid just like a king moves in chess. A diagonal move counts the same as a horizontal move. This L_∞ metric is given by:

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

City Block Distance also known as *Manhattan* distance. This metric assumes that to go from one point to the other one has to travel directly along the coordinate lines of the grid. No diagonal moves are allowed. This L_1 metric is given by:

$$L_1 = |x_1 - x_2| + |y_1 - y_2|$$

Euclidean Distance is the most commonly used distance metric which people are probably also most familiar with. The Euclidean distance is given by the L_2 metric:

$$L_2 = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$$

L_3 metric is given in the library as an interpolation between the L_2 and L_∞ -metric being a compromise between computational effort and a visual interpolation between the orthogonal L_∞ and the radial L_2 -metric:

$$L_3 = \sqrt[3]{(x_1 - x_2)^3 + (y_1 - y_2)^3}$$




Figure 9: The four supported distance metrics: L_1 , L_2 , L_3 and L_∞ .

3.2.6 Multiple Foci

The library supplies the means of adding more than one lens to the representation to be distorted. These lenses can interact in several different ways as will be seen in Section 3.2.7. This makes it possible to create new forms of lenses which have not been supported by default. Figure 10 shows an example of a fish-lens created by combining a lens with a circular and a polyline focal area.




Figure 10: A new EPF lens created by combining a polyline lens and a circular lens.

3.2.7 Blending

Multiple foci can be blended together to avoid visual discontinuities where they meet. A new translation vector is calculated where the drop-off functions lie over each other. The blending can be either additive or an interpolation between both functions. The two lenses in Figure 10 are blended using an interpolating calculation.

3.2.8 Additional Information

For reasons of completeness, we are going to include other notable features of the EPF library. They are not mentioned again in this thesis but serve to get a general overview of the capabilities EPF presents:

Focus Alignment:

The calculation of the focus has been implemented nicely within EPF. To avoid that the focal point moves out of the view frustum in perspective projection the focal point is aligned to the viewpoint making it impossible to move out of sight as long as the full lens is on the visible part of the base plane.

Folding:

Foci can be repositioned by the application of folding. The region of the representation that is magnified stays the same but the position of the focus changes and new information moves into the focus.

4 A Taxonomy for Discrete Lenses

This and the following section presents the taxonomy for the discrete usage of EPF lenses that has been developed during the internship. Discrete lenses, and other related terminology will be defined in this section. These and other definitions can be found in the Appendix. Most lens techniques, and certainly all within the Elastic Presentation Framework, apply a

distortion function to the representation that is to be distorted. The distortion functions - or field functions, as they will be called - currently implemented within EPF are all continuous functions. This means that they are defined for a continuous domain such that for any given input point a displacement vector can be calculated. Discrete and continuous lenses will be defined visually using the application of the distortion technique on the underlying data as a criterion. Examples will show the application of the described methods at the end of the section.

4.1 Discrete Lens vs. Continuous Lens

The greatest challenge of the internship task and, consequently, this thesis is to work out and describe the difference between discrete and continuous lenses. Many researchers have used these types of lenses or described techniques for their application but no good or complete definition is given in the literature. How researchers have used or understood these terms in relation to their work is discussed in Section 6. For further reference please refer to the definition given in the Appendix.

4.1.1 Continuous Lens

A continuous lens on the computer is a visually continuous lens. On the computer as arguably in the real world there is no actual continuity if you think of its mathematical definition. However, you can intentionally create the effect of a visually continuous lens. Many graphics methods have been developed to mimic visual continuity, and while they may achieve convincing results, there will be a limit. For a continuous lens the data has to be accessed so that each visible point on the surface of a representation is changed according to the applied distortion function and the actual distortion is done above display resolution. This ensures that for each point on the display one or more data-points can be found that have been changed according to the applied distortion.

Since lens functions used for the distortion process always take individual points or vectors as input data the representation has to be discretized into data-points prior to distortion. When distorting a visual structure continuously, the discretization of the representation has to be so high as to ensure that the resolution criterion is met. Mathematically speaking, real continuous distortion can only be applied to visual representations which are described by functions or procedures so that a continuous distortion can be applied to the whole visual structure without having to discretize it. Since this is rarely the case distortion will be called continuous in this thesis as long as the visual criteria is met. Visual continuity within the distorted representation ensures for example adjacencies and order.

Researchers described a rubber-sheet view for such distortion-oriented presentation techniques [SSOR93]. This metaphor describes the underlying representation space as a rubber-sheet mounted on a frame. When the rubber-sheet is dragged towards the viewer some parts of it become enlarged and some shrink in size. Figure 11 shows such a rubber-sheet as viewed

from the side and the achieved distortion effect when the sheet is viewed from the top. This rubber-sheet represents a coherent space on which a representation resides. A continuous distortion effect occurs when every part of the representation embedded on the rubber-sheet changes according to the applied “stretching”.




Figure 11: The distorted “rubber-sheet”.

Continuous detail-in-context magnification lenses mimic continuous levels of detail. Of course, on a computer as in the real world, no unlimited detail can be displayed. However, to a certain extent continuous levels of detail can be imitated. As can be seen in Figure 12(b) the level for continuous magnification of the image has been reached, while a different distortion technique in Figure 12(c) shows a still continuous lens applied.




Figure 12: Distortions to an image. Pictures courtesy of Eric Pattison, University of Calgary, original image: Orthoshop.

As with discrete lenses, which will be described later, there are several possible techniques for continuous lenses, of which some will be briefly mentioned here to give the reader a feeling for the field of continuous lenses. While working in the two-dimensional domain, continuous lenses can be applied in one dimension or two dimensions, for example. In Figure 13(a), the

representation is distorted with a two-dimensional distortion applied while in Figure 13(b) a one-dimension distortion is used which merely changed the representation horizontally. As can be seen, the representation stays visually continuous.




Figure 13: A 1D- and 2D-continuous lens.

The major application of lenses using visually continuous distortion is the assessment of detail in a representation while keeping the context visible. The rubber-sheet metaphor is suitable for describing such detail-in-context magnification lenses. As has been described previously by BIER et. al [BSP⁺93], lenses do not necessarily have to be used for magnification purposes: they can act as visual filters on the underlying data. Figure 14 shows such a lens used continuously on a visual structure. The two-dimensional plane is the R-G baseplane of the RGB color cube. If the plane is stretched in the third dimension along the B-axis, the plane's colors adjust accordingly. The effect is that the lens filters the blue part of the RGB-values coloring the plane.




Figure 14: A continuous color lens.

4.1.2 Discrete Lens

A discrete lens makes choices about the use of data and/or representation and/or distortion. This means that at least some points on the representation that will be displayed are not adjusted according to the chosen distortion function. The distortion for a discrete lens is so applied below the resolution of the output medium. A choice then has to be made on how to interpolate between distorted points.

Therefore, in contrast to a continuous lens, a discrete lens does not require data to be discretized above the resolution of the output medium or data to be mathematically described by functions or procedures. It does not have to ensure a distortion coherency to the underlying distorted representation space. Referring to the rubber-sheet metaphor described above, the visual structures comprising a representation do not have to be embedded in the rubber-sheet but can be loosely connected to it, for example by being attached to it at certain corner points. This means that only parts or certain properties of a visual structure are distorted and the remaining parts can be interpolated according to the preferences of the visualization tool designer. As with continuous lenses, all or just discrete parts of the lens's field function can be used for the distortion transformations. In a geometric sense, implying a linear interpolation between distorted data-points, a discrete lens so maps piecewise linear structures to piecewise linear structures while continuous lenses map piecewise linear structures to curves. Figure 15(c) shows one type of discrete lens applied to the grid while Figure 15(b) shows a continuous lens. A more distinct analysis of discrete distortion will be shown in the following sections. For further reference to used definitions please refer to the Appendix. Application for discrete lenses will be shown in Section 6.




Figure 15: Discrete vs. continuous lenses.

4.2 Discrete Representations

A representation of data can be seen as continuous or as being comprised of parts. These parts can be naturally defined by an inherent discrete nature of a representation or can be artificially created. A representation of menu items as icons, for example, is discrete in nature since there are only so many distinct menu options available within a certain program. An image on the other hand as a representation of a certain place in the world at a given time is

inherently discrete only on the pixel-level, it is perceived as representing a continuous image of the world. Artificial parts of an image can be created by cutting an image into several pieces and then dealing with every part separately. An example of a discrete image lens can be seen in Section 5.9. The discrete representations used exemplarily in this research are point sets which comprise complex geometric objects. Discrete objects used within EPF have to be geometrical objects. They will be further described in Section 4.2.1. The big difference between discrete and continuous lenses used on visual entities is that for discrete distortion object properties can be chosen freely for distortion out of the range of properties a certain object offers. It is also possible to distort several properties at the same time, for example by magnifying the area of a square and translating it by its center point in 2d space, so that position and size of the object change as can be seen in Figure 16. Discrete lens transformations will be further investigated in Section 4.3.




Figure 16: Discrete squares moved by their center point in 2D with a magnified area.

4.2.1 Discrete Objects

Discrete visual entities can be spatially isolated, i. e. separated in space, having independent locations or some kind of conceived isolation. One can see a graph, for example, as a unit or as being comprised of nodes and edges. These nodes and edges can be conceived as isolated discrete objects in space which can be accessed and changed.

Discrete lenses in this research have dealt with objects up to a three-dimensional geometry. Since the objects are laid out on the baseplane of the viewfrustum as described in Section 3.1 two-dimensional representations have been investigated further. Some examples will show the application of the presented method to 3D objects laid out according to the baseplane. For more precise interaction with 3D data, revised field functions should be used as described by CARPENDALE et al. [CCF96].

As mentioned, complex geometric objects are made up of parts and can have discrete as well as continuous object properties. Each polygon has these (and more) geometric parts:

Points: like corner-points or center-points, which are considered to be zero-dimensional,

Lines: like border-lines or diagonals that are one-dimensional, and an

Area: surrounded by the corner-points and border-lines which is considered to be two-dimensional.

Figure 17 shows an example of one particular object, which will be used as an example throughout this thesis, a square. The figure also shows the identified geometric object properties. A square was chosen as an object because of its use in HCI and visualizations. Buttons, windows, labels, nodes in graphs, etc. usually are rectangular objects.




Figure 17: One discrete object and its geometric properties.

The decision to use a square has been made because of the easiness to order the objects evenly in a grid and the possibility of comparing distortion techniques used on the objects in the grid. From the specialization of a square it should be easy to generalize shown ideas to general rectangles or other discrete objects and representations. It is also possible to take less-dimensional objects than general polygons. Lines and points can themselves be seen as discrete objects. In this case the objects only have the properties and parts up to and including their own dimension. A line, as an example is defined as not having an area. According to EUCLID a line is a widthless length and, therefore, it consists of one line (itself) and the points on the line. Points being zero-dimensional obviously only consist of one point (itself). Discrete objects can have other object properties as well like color, texture, size, shape, orientation or position. Size, shape, orientation, and position are seen as geometric properties of the object while other properties like color, or texture are optional and do not necessarily have to come with a discrete object laid out in space. In this research only visual object properties were considered but an example at the end of the thesis will show the application of distortion to non-visual object properties.

4.3 Distortion Transformations

Generally within the Elastic Presentation Framework, there exist the following distortion transformations when applying a lens to a two-dimensional representation:

- three-dimensional displacement,
- two-dimensional displacement,

- one-dimensional displacement,
- magnification, and
- hybrids.

One-, two-, and three-dimensional transformations are carried out in the respective geometric dimension. Magnification is a value defined for an output data-point which can be used for a transformation. Hybrid transformations are combinations of the other four. Table 1 shows three identified general object-properties and the transformations defined on them given by EPF. The calculated magnification factor for objects within the focal area corresponds to the perceived magnification of an object in the focal region within EPF. This is the case if the corner points of an object within the focus were distorted or if the magnification factor has been applied to calculate the correct size of the magnified object in some other way, like scaling the borderlines, for example. This is not the case if the magnification factor is used to scale the area of an object. The term enlargement of the selected area would be mathematically correct in this sense. However, for reasons of simplicity the term magnification will be used throughout in this thesis.


Transformation/Objects	points	lines	area
	0D	1D	2D
1D displacement	x	x	x
2D displacement	x	x	x
3D displacement	x	x	x
Magnification	-	x	x

Table 1: Transformation on object properties.


The table could be easily extended to *volumes* in three dimensional representations but this extension is beyond the scope of this thesis. The extension to n -dimensional objects is also possible however, in connection with presentation issues not useful. Points cannot be magnified since they are mathematically defined as zero-dimensional. Lines can be magnified by scaling their length and moved by translating each point on the line. Areas can be magnified by scaling in two dimensions and translated by moving each point in the area. Figure 18 shows two examples of a discrete lens applied to the discrete objects in the grid. Figure 18(a) shows the object's corners distorted with two-dimensional displacement applied. Figure 18(b) shows the object being magnified by applying the magnification factor to its area.

4.4 Distortion Calculation Within the Elastic Presentation Framework

Within the Elastic Presentation Framework, a *field function* is described as a point in a seven-dimensional feature space. Figure 19 shows this feature space with continuous axis. Although within the EPF library only the six default drop-off-functions and focal shapes described in Section 3 are implemented more drop-off functions and focal shapes can be created. It is also

**Figure 18:** Discrete lenses.

possible to introduce new monotonous functions within EPF and create general polygons as focal shapes although they are currently not implemented within the EPF library. The resulting distortion calculation function is influenced by which field functions were chosen with these factors. A lens distortion consists of a field or a piecewise group of selected field functions. A choice of connections can be made for a group of field functions. Two-dimensional spatial coordinates of given data points are the input values for the lens's field function. As described later, the factors influencing the field function, the input parameter, and the output parameter of the lens can be chosen freely for use in discrete lenses since no visual coherency has to be kept. The EPF lens library is a good tool for exploring discrete lens distortion since it handles all input data points the same way and the visualization tool designer has to find for him/herself possibilities to work with the acquired data.

**Figure 19:** The EPF feature space.

4.5 Why Develop a Taxonomy

“An early step towards understanding a phenomena is [...] to develop a taxonomy” [Sim96]

We already have a very well defined understanding on how discrete data can be laid out and presented. Graphs, charts, and trees all have rules for an effective layout and research has also been conducted into what types of relationships have to stay unchanged for a user to understand a certain presentation of data. Orthogonality, proximity or topology have been identified to be necessary for the user to preserve a mental map when graphs are distorted for example [SM95]. Regardless of any such relationships, a taxonomy of discrete presentation techniques will help us to understand new concepts of distortion within detail-in-context presentation and might lead to new visualization techniques previously unthought of.

The taxonomy developed in this work aims at creating an understanding of the wide range of possibilities presented within EPF for discrete lenses and a way of showing the connections and relationships between the methods presented.

4.6 Development of the Taxonomy

Several possibilities were considered for developing a categorization of the discrete presentation techniques made possible by EPF.

4.6.1 An Approach Based on Effects

As a first approach, several effects were considered and categorized according to aspects like overlap, displacement, or magnification. Although this seemed to be a reasonable approach at first, it became apparent that too many possibilities and constraints on them had to be thought of to make a clearly distinct categorization. There are already nearly endless possibilities of choosing discrete object properties to distort as seen in Section 4.2.1. The problem can be shown with the example of overlap. Overlap within a selected distortion depends on the size of the objects, the chosen type of distortion, the placement of the objects in the space, and the selected lens properties. Figure 20 shows the change from no-overlap to increasing overlap with changing focal size. If the focal size would be further increased to be bigger than viewplane the effect would be a zoom-in showing no overlap again. However, it is important especially for the visualization tool designer to know how certain distortion effects can be deliberately created or avoided. Section 5.7 will describe some general effects, which occur during the different types of distortion.




Figure 20: Notice the difference in overlap with changing focal size.

4.6.2 An Object-Oriented Approach

A second approach dealt with categorizing the distortions according to what aspect of an object would be affected. The two main categories were object-oriented distortions and vertex-oriented distortions. Object-oriented distortions changed properties of the considered objects like its size, orientation, or shape. Vertex-oriented distortion merely changed the corner points of objects leading to some effects which showed partial embedding of the object in the underlying elastic space. This approach extended the taxonomy enormously and lead to new insight into EPF but the two main classes were not disjoint. The insight gained from this approach was used in identifying the different object properties and geometric aspects of distortion with the Elastic Presentation Framework as will be seen in Section 5.1.

4.6.3 A Mathematical Approach

The last approach on categorizing the distortion possibilities within EPF discrete space were organized according to the mathematical distortion possibilities offered within EPF and the mathematical or geometric properties offered by the data to be distorted. The following section will put these three approached together and explain the use of the taxonomy on the practical example of how to build a discrete lens.

5 How to Create a Discrete Lens

This section will show a proposed way to create a discrete lens and at the same time introduce the five identified aspects of a discrete lens. The shown steps will also explain ways and techniques related to discrete lenses. Some will be closely related to the perspective projection technique used within EPF but most will be general enough so that developers using different distortion approaches will be able to exploit the ideas shown here as well. As mentioned in Section 4.1.2 and Section 4.2, a discrete lens distorts object properties on a level that is above display resolution so that an interpolation has to take place. Since lens distortion is usually understood as a visualization technique this thesis concentrates on visual lenses. However, other types of lenses can be thought of as will be seen in Section 5.9.

5.1 Step 1: Object Properties

To create a discrete lens for a certain data set it is important to attend to the way this data is represented. The first step in creating a discrete lens, therefore, is to identify the properties of the visual entities we want to distort. We assume to have distinct, usually more than one, visual entities to distort. All the shown techniques can also be applied to a single visual structure but will prove to be visually less powerful. The increase in size of a single visual object, for example, will not show the richness of the shown techniques with the application of an EPF lens as will a whole group of objects. For this reason we concentrate on the case of several discrete visual entities for distortion.

5.1.1 Distorting Geometric Properties

Since the objects within the Elastic Presentation Framework are all laid out on the baseplane of the view-frustum they all have geometry as their basic layout principle. The lens's field function takes geometric coordinates as input data and produces geometric data as output. Therefore, geometric distortion of objects has a strong influence on the location, size, orientation, and shape of the visual entities to be distorted. The following sections will show example images of distortions to the three main identified geometric object properties: points, lines, and the area. They will be explained in detail in the following subsection since influencing these three basic object properties also influences higher-level properties like shape, orientation, size, and location. The chosen properties of geometric objects have also been stated as *classes of representations* by Bertin [Ber83].

Points

According to EUCLID [Byr47, Book 1, Def. 1] a point is defined as that which has no parts. Over the years, several mathematical definitions for a point were developed. Points can be seen as dots, an exact location, an ordered pair, or as a node (or vertex). In this thesis points are seen as zero-dimensional, having no height, length, or width, just a position as defined by BERTIN [Ber83]:

“A POINT represents a location on the plane, that has no theoretical length or area. This signification is independent of the size and character of the mark which renders it visible”

The coordinates of a point are defined by the discrete object or the place in space at which a point is located. Points can be translated but due to the lack of a dimension they cannot be magnified. Therefore, points can be moved in 3d space with perspective projection applied but will never change in size.

Table 2 shows discrete distortion applied to the corner points of the 2D object. You can see that while distorting only points in 2D- or 3D-space other geometric properties of the objects change as well. The size and shape of the objects change, as well might their orientation compared to the baseplane.




Table 2: Sets of distortions to corner points.

An interesting application for such distortions could be a fast artistic transformation of computer generated line drawings, produced with NPR silhouette techniques, for example. The images in Figure 21 show the distortion of 2D-silhouettes produced with the Magdeburg Open-NPAR system ¹.

Lines

EUCLID [Byr47, Book I, Definition II] defines a line as a length without a breadth. The extremities of a line are points [Byr47, Book I, Definition III]. BERTIN again gives a similar definition taking into account the rendering of a line [Ber83]:

“A LINE signifies a phenomenon on the plane, which has measurable length but no area. This signification is independent of the width and characteristics of the mark which renders it visible.”

¹data sets courtesy of Tobias Isenberg




Figure 21: Artistic distortion of an NPR image achieving caricature-like renditions.

Taking these definitions, we can say that a line has a length and a position on the plane but no width. Therefore, a line can be translated and magnified in length.

Similar to point distortions other geometric properties of the object are affected when its lines are distorted as can be seen in the figures in Table 5.1.1.

Area

EUCLID calls a surface what will here be described as an area. He gives the following definitions [Byr47]:

- A surface is that which has length and breadth only,
- The extremities of a surface are lines, and
- a plane surface is that which lies evenly between its extremities.

Therefore, an area is defined as being two-dimensional, having a size and a position as well as bounding lines. Areas can be visually represented by points and lines. Areas can be translated as well as magnified. Examples are shown in Table 5.1.1.














Different Line Distortions			
1D			
2D			
3D			
Mag			
Hybrid			

Table 3: Sets of distortions to borderlines.
















Different Area Distortions			
1D			
2D			
3D			
Mag			
Hybrid			

Table 4: Sets of distortions to the area of the object.

5.1.2 Distorting other Discrete Properties

The work for this thesis concentrates on the discrete distortion of geometric properties. However, some other notable examples will be shown in this section which exemplify the discrete distortion to other object properties. One can think of other visual properties like color, texture, shading, subdivision, etc. or even non-visual properties, like sound.

Raytracing

The following images show an example of raytracing with EPF lenses. The lenses used are continuous lenses since raytracing is done at display resolution but are described here to show the possibility to use other aspects for distortion. The lens is applied to the rays emitted from the view-port. Examples can be seen in Figure 22.




Figure 22: Continuous-raytracing-lens.

Subdivision and Shading

Figure 23 shows how a discrete lens can be applied to other than geometric properties. In Figure 23(a), the lens's magnification value has been applied to define a subdivision value for the displaced spheres. The more a sphere is translated towards the viewport and so increased in size the more subdivisions steps are applied to its mesh. The spheres in focus, therefore, seem to be smooth and can be viewed in more detail than can and need spheres in the context. Figure 23(b) shows an additional shading applied to spheres with magnification > 1 .

Sound

FAZEL an undergraduate student at the University of Calgary, used the idea of influencing sound volume with a discrete lens. In his project, all instruments initially play at the same volume. The lens's magnification factor calculated for the visual representation of the instruments increases or decreases the volume when the mouse moves closer to a particular instrument or away from it. The visual representation also gets bigger and darker. When the user clicks on




Figure 23: Discrete lens applied to subdivision and shading

an instrument its volume gets the loudest and the volume of the other instruments reduces to zero (mute).



Figure 24: Discrete lens applied to sound volume.

We are quite aware that in this case it is debatable if distorting sound can be called discrete since discrete distortion here has been defined visually. However, since the distortion is done visually discrete to the representations of the instruments and the volume of the attached sound-file is seen as a single object property we will categorize it as a discrete distortion.

5.2 Step 2: Reference Points

When choosing points for distortion one has to identify where on the baseplane the parameter for the distortion are calculated. Figure 25 shows all four corner points distorted according to distortion parameters calculated along the horizontal and vertical axis through the focal point. When translating a line on the plane, a reference point has to be defined which determines the amount of translation for all other points on the line. For an area, a reference point has to be determined to define the amount of translation for all points in the area. When magnifying lines or areas, a magnification factor has to be determined for the object at one reference point

in space. Usually the reference points lie on or in the object itself but interesting effects can be achieved when putting them outside the objects as seen in Figure 25.




Figure 25: Reference points laid out along major axis.

5.3 Step 3: Distortion Properties

Distortion functions have to be created by choosing field functions with certain properties as discussed in Section 3.2. The chosen field functions can then be connected to create a final distortion function. They can be disjoint (having undefined parts), adjacent (being discontinuous but having no undefined parts) or connected.

5.4 Step 4: Distortion Output

The taxonomy takes a mathematical approach of categorizing distortion possibilities within EPF for discrete objects. As described in Section 3, the EPF takes discrete points as input data and outputs them distorted in either 2D or 3D including the magnification factor.

The output from the distortion function is always a magnification factor for this particular data-point and displacement factors which depend on the chosen lens properties. These displacement factors can be either two-dimensional or three-dimensional. So the output parameters are categorized as seen in Figure 26.




Figure 26: The EPF distortion parameter.

As mentioned in Section 4.3, these output parameter can be applied to the object parameter chosen. These output parameter can be combined in $2^6 - 1 = 61$ permutations of which not everyone can be applied since it is not possible to apply the x -translation factor for the 2D- and 3D-functions {2D:X,3D:X} to one point at the same time, for example. A point can be translated in x -direction by only one factor. As a simplification, the following distortion parameters in Figure 27 will suffice and give a more general understanding of the mentioned distortion techniques.




Figure 27: The EPF distortion parameter.

Table 5 shows the $2^4 - 1 = 15$ transformation combinations supported by this simplification. It assumes that combinations of translations will be handled either in 2D or in 3D. Either the translation factors will be taken from the two dimensional function of the EPF library or all translations will be done according to the three-dimensional function. The user can choose for her/himself which one to use. The third column shows the functions used for the examples shown in this thesis.

Combinations	Type	Function
{X}	1D	2D
{Y}	1D	2D
{Z}	1D	3D
{X,Y}	2D	2D
{X,Z}	2D	3D
{Y,Z}	2D	3D
{X,mag}	hybrid	2D
{Y,mag}	hybrid	2D
{Z,mag}	hybrid	3D
{X,Y,Z}	3D	3D
{X,Y,mag}	hybrid	2D
{X,Z,mag}	hybrid	3D
{Y,Z,mag}	hybrid	3D
{X,Y,Z,mag}	hybrid	3D

Table 5: Combinations, dimensionality, and used function for the distortion parameter.

5.5 Step 5: Mappings

Each output factor can be applied to the original data-points to produce the final distortion. A mapping of these output factors to the input point describes the actual distortion of the input data-point. We assume to have the point $\{x,y\}$ as input for our distortion and receive $\{x',y',mag\}$ as output parameter. Table 6 shows some of the many possible mappings for the x -parameter of the input data-point.

Input		
x	x'	y'
x	x	y
x	mag	
x	$x'*mag$	$y'*mag$
x	$x*mag$	$y*mag$
x	$x*y'$	$x'*y$
x	...	...

Table 6: Some possible mappings for x.

As will be seen later the output data does not have to be mapped directly back to the input point. The calculated magnification factor, for example, can be applied to the length of an object, its width, height, or color values.

5.6 Step 6: Apply Lens and Present

The lens can be applied to a representation once or interactively while the user moves the focal center or moves the representation under the lens. Interactive presentation with discrete lenses usually allows higher frame rates compared to similar presentations with continuous lenses since only a few points have to be distorted and others can be interpolated by the Graphics API, for example.

5.7 General Discrete Distortion Effects

Due to the number of possibilities in distorting objects and to the number of object properties that can be chosen it is impossible to predict which effects the distortions will have in a general case although some effects can be identified.

Effect Point Distortion:

We assume that our objects are laid out in a grid with no overlap. When the corner points of the objects are distorted, there will be overlap only when the focal size becomes increasingly large compared to the drop-off width. Then the slope of the drop-off function is very steep and can lead to information reversal.

Effect Line Distortion:

We assume that our objects are layed out in a grid with no overlap. When borderlines are distorted, there will be overlap as soon as the objects' size becomes larger than the gap between the objects in the grid.

Effect Area Distortion:

When the area of the objects is displaced, or magnified the orientation and shape of the object does not change. Instead, the size and location in space is changed.

Effects due to Perspective Projection:

When displaced in z objects closer to the viewpoint will increase in size. This is due to perspective projection. This assumption does not hold if the objects are points. Objects moved only in z will also seem to move sideways due to perspective projection as they change size. The focus is then misaligned from the viewpoint. Figure 28 shows the effect of displacement in z -only in the perspective view frustum used by EPF. It is, therefore, advised to use all three coordinates for displacement in 3D.

Effects of Continuous Distortion:

When objects are distorted continuously objects that lie just out of the focal area but not on the baseplane will be compressed. They will most likely be distorted in size, shape, and orientation.




Figure 28: Effect of Z-displacement in perspective projection.

5.8 Problems of Discrete Distortion

Certain effects of discrete distortion can be problematic in some application scenarios. Overlap can lead to reduction of insight in cases where full context is necessary. It creates visual discontinuities and might be confusing for the user. As in continuous distortion, objects might become distorted so that it is difficult for the user to identify or use in a particular task. Such problems have to be taken into account by the visualization tool designer.

5.9 Student Ideas

This section includes examples of discrete lenses as they were implemented by students taking the *CPSC 581 Human-Computer-Interaction II - Interaction Design* class at the University of Calgary. Their task was given as follows:

You are a junior member of this interaction design team and you have been asked to create a novel elastic interaction that will operate like a portal into the information space. You want to create a combination of interaction dynamics that fits well with the images and information in the museum.


The examples are just described to show the many creative ways in which the shown methods can be applied in a presentation.

5.9.1 Discrete-Image-Lens

Mike Larke, an undergraduate at the University of Calgary, developed this lens for a discretized image of San Francisco. The image was cut in several pieces, which were distorted separately. The image pieces contain further information about this particular area of San Francisco which can be accessed by clicking on it.



(a) A discretized and distorted image.



(b) Revealed information.

Figure 29: A discrete image lens, images courtesy of Mike Larke.

5.9.2 Ansel-Adams-Lens

This lens was developed by Kathryn Elliot, an undergraduate student at the University of Calgary. An interesting approach for using one-dimensional line displacement is shown here. She uses this technique to open up text-boxes containing information on Ansel Adams.




Figure 30: The Ansel-Adams-Lens, image courtesy of Kathryn Elliot.

5.9.3 Discrete-Word-Synonym-Lens

This lens was created by Michael Boyle, a PhD student at the Interactions Lab at the University of Calgary. He used a lens to magnify the length of a word in a sentence. As the magnification for a word increases the word is replaced by longer synonym of itself.




Figure 31: A Word Synonym Lens, images courtesy of Michael Boyle.

6 Related Work

This section will show innovative ideas in the field of detail-in-context presentation for two-dimensional representations and research that has used and/or presented ideas in the field of continuous vs. discrete lenses. The shown techniques will be set in context to the approaches introduced in this thesis. At the end of the section ongoing research with the Elastic Presentation Framework will be presented. It has to be noted that techniques which include what we have called detail-in-context presentation so far have many other descriptions in the literature. The term fish-eye lens, for example, is used somewhat confusingly in the literature. It is sometimes used to describe detail-in-context methods in general. On the other hand, it most commonly describes a distortion technique very similar to the polar transformation technique described in Section 6.2. One of these lenses applied to a picture can be seen in Figure 36. Detail-in-context techniques have also been called *distortion oriented presentation* [LA94], *focus-and-context*, *polyfocal display* [NE78], or, as mentioned, *fish-eye-views*.

6.1 Bifocal Display and Perspective Wall

One of the earliest detail-in-context techniques was presented by SPENCE and APPERLEY in their paper *Data Base Navigation: An Office Environment for the Professional* [SA82]. The *bifocal display* was introduced as a method which makes use of our spatial memory when searching for data on the computer. The technique was developed for information that is subdivided into a hierarchy of journals, volumes, issues, and articles. A single focus region and two adjacent regions of distorted context create an environment of detail integrated into its context. Figure 32 shows how such a display is created. We assume an undistorted representation which is too large to be displayed on a single screen (Figure 32(a)). We then fold the sides of the representation so that all of it is still in view (Figure 32(b)). This folding is projected back onto the screen (Figure 32(c)) leading to one central region with undistorted detail and two uniformly distorted regions of full context. An encoding for the distorted items has to be made available so that viewer easily gets an idea about what the content of the distorted regions might be and if it they are possibly of interest for him. In the detail view SPENCE and APPERLEY show pages out of a journal, for example, while the distorted view groups pages into issues or articles.




Figure 32: Techniques used in creating the original bifocal display.

With the Elastic Presentation Framework a discrete lens was created to imitate a bifocal display as seen in Figure 33. A discrete bifocal epf lens uses a rectangular or square focal shape and a linear drop-off-function and out-of-object reference points to create spatial constancy between regions. In contrast to the original version this implementation scales the objects according to their right and left borderlines in x -direction and according to their top and bottom borderlines in y -direction. This increases the size of the object in the focal region and compresses objects in the context part of the lens. The original implementation of the bifocal display kept the objects in focus in their original size and left no possibility for viewing the central objects in greater detail.

In 1991, MACKINLAY et al. presented a variation of the bifocal display called the *Perspective Wall* [MRC91], as seen Figure 34(a). Perspective wall is a detail-in-context technique which uses the third dimension. The data distorted in the perspective wall approach usually has some linear ordering, for example a chronological or alphabetical ordering. Such linearly structured data usually covers a wide area on a two-dimensional layout which is difficult to be displayed on a single screen. Perspective wall creates a central region for viewing detail and two perspective regions, one on each side, for viewing context by folding a two-dimensional layout into a three-dimensional wall. The central region is parallel to the screen and so cor-





Figure 33: A discrete bifocal EPF lens.

responds to our reading-habit. The x -axis corresponds to the linear structure of the displayed information while the y -axis can be used to structure the information by other attributes. An advantage of this approach over the bifocal display is that this technique does not require the context data to be grouped or specially scaled. Also, context adjacent to the focus area has a larger scale than context which is further away. This technique presents three times as much information as a flat wall with detail of the same size but wastes screen real-estate at the sides of the wall. Interaction is possible in moving along a linear direction. Perspective wall has been introduced as a continuous lens technique. Items bend around the border of the focal region as can be seen in Figure 34(a). A discrete variation of this lens was introduced by LAM and SPENCE in 1997 [LS97]. It can be seen in Figure 34(b).



(a) Continuous.



(b) Discrete.

Figure 34: Perspective Wall.

6.2 Generalized and Graphical Fish-Eyes

FURNAS introduced a technique based on a very wide angle lens, the *fish-eye lens* [Fur86]. He notes that such a lens can show parts of the world in great detail and the remaining parts with successively less detail. The main aspect of a fish-eye lens is to have local detail within global context. In his 1986 paper, FURNAS analyzes naturally occurring fish-eye views. He concludes that many natural occurring views of the world have a fish-eye character and that

the application of such views to the computer, therefore, might prove to be a good presentation technique. The basic idea behind FURNAS' fish-eye lens was a *degree-of-interest* (DOI) function which is assigned to the objects to be displayed. The items are filtered according to the applied DOI. Relevant information will be displayed in great detail, while less relevant information will be suppressed. A point of focus is defined in which detail is of importance. The DOI depends on a global a priori importance (API) which has been set for an item and the distance of an item from the current focus. The API for an item requires domain specific knowledge about the data. FURNAS suggests the application of the algorithm for lists, trees, graphs, and Euclidean spaces. FURNAS' algorithm is general enough that it does not require spatial data or graphical output. However the term *fish-eye lens* more frequently describes distortion techniques for graphical representations.

In 1992, SARKAR and BROWN described a system for viewing planar graphs using a graphical fish-eye lens [SB92]. FURNAS work is extended to include layout formalisms for the position, size, and level of detail of objects to be displayed. A fish-eye view for graphs is generated in two steps: First, a geometric transformation which repositions vertices and magnifies or demagnifies areas according to their distance from the focus. Secondly, an API function is used to compute the final size, level of detail, and a visual worth factor for the nodes of a graph. SARKAR and BROWN introduce two transformation techniques, the Cartesian transformation and the polar transformation. They lead to different fish-eye effects as can be seen in Figure 35. The technique has some drawbacks compared to techniques, which have been developed later, like the Elastic Presentation Framework. A focus cannot be placed arbitrarily on the display space, a single item has to be chosen. Also, users cannot choose the size of the focal region.




Figure 35: Transformation functions used in [SB92].

According to the definitions presented here SARKAR and BROWN introduce a discrete fish-eye lens since only the vertices of the graphs are distorted according to their distortion function. Edges are mapped at their endpoints or occasionally at intermediate bend points. The graphical representation of the nodes of the graphs are just scaled in size and not distorted and additional detail is displayed within the nodes when they reach their full size. However, the paper mentions the possibility of a continuous lens by mapping many intermediate points and

border points of the nodes. However, the authors chose not to do so in order to achieve real time performance.



Figure 36: What most people understand as a fisheye-lens.

6.3 Nonlinear Magnification

In their 1996 paper, [KR96b] KEAHEY and ROBERTSON introduced the term *non-linear magnification* to describe effects common to many current detail-in-context visualization techniques, local detail and global context. He describes general non-linear transformations:

- One-dimensional transformations,
- Two-dimensional transformations to two dimensional coordinate space,
- Orthogonal transformation which is a one-dimensional function applied to two dimensions individually,
- Radial (fish-eye) transformations, which use polar-coordinates, and
- bi-radial transformations, which are a combination of radial-and-orthogonal transformations.

In [KR96a] KEAHEY and ROBERTSON propose a general formulation of the “detail-in-context problem”. The detail-in-context problem as seen by Keahey does not only include how to create a focal and context area but also how to provide means for placing more detail within the focus. He examines two areas. First, the information is composed of discrete objects placed on a map of Scotland. The problem is how to render the objects so that they are consistent with the underlying spatial information. He examines two methods for increasing detail, by increasing the size of the objects, which may lead to overlap or by embedding the objects in the underlying space. These two techniques have been described here as discrete and continuous lenses respectively. When increasing the size of an object, a certain level of detail can be blended in. Figure 37 shows two examples where a discrete and a continuous lens have been combined with a level of detail function. For the continuous lens, the LOD is applied to a

first version of the representation to change it to the final representation to which the distortion is applied.




Figure 37: Two methods for transforming discrete objects, images courtesy if Keahey.

For the continuous technique, he identifies certain problems. The layout of the objects on the undistorted representation becomes more challenging because it has to be ensured that the objects do not initially overlap. Also, a certain distortion is introduced in the objects to be rendered.

6.4 Flip Zooming

In Chapter 2 of his PhD thesis, HOLMQUIST [Hol00] presents ideas related to the concepts presented in this thesis. He identifies two types of representations: *continuous visual structures* (maps or images) and *discrete visual structures* (trees or graphs). When a view transformation is applied to a continuous representation the whole visual representation is distorted regardless of what it represents. He compares this technique to the rubber-sheet view defined by SARKAR and BROWN. Discrete detail-in-context techniques are developed specifically for structures made of discrete visual entities. The view transformation is applied to these entities and the overall presentation is then adjusted to retain structural information. HOLMQUIST identifies key ideas. However, his terms *continuous*- and *discrete representation* have been quite undefined. Continuous lenses can also be used on discrete visual structures on purpose and they can keep structural information as well. On the other hand, they might not have structural information at all besides their place on a two-dimensional surface. Discrete lenses can be applied to images and maps when an image is discretized into pieces on purpose. This can be seen in the students examples presented in Section 5.9.

HOLMQUIST visualizes discrete visual structures, which are placed in the context of a certain linear ordering. The flip zooming technique is a detail-and-context technique for documents consisting of several pages. His main idea is to let the user zoom in a particular page while

showing the other document pages in context. He first places all his documents on a two-dimensional surface, as thumbnails for example. Figure 38 was taken from his thesis. It shows how he lays out the elements left-to-right and top-to-bottom to preserve ordering. The limits of this approach are closely connected to his constraint to preserve ordering. The focus and context elements shift in space, the number of rows and columns has to change, and there is a limited amount of information that can be displayed.




Figure 38: Left unzoned view. Right flip zoomed view; image courtesy of HOLMQUIST.

6.5 Interaction with Distorted Views

Lately, usability studies have been conducted to show the interaction issues with distorted representations. GUTWIN and SKOPIK showed that navigation and efficiency of user interaction was much improved by the use of detail-in-context techniques in representation where magnification was required. [GS02]. One interesting discrete fish-eye lens which involves interaction has been introduced by BEDERSON in 2000 [B.B00]. It distorts linear menus by changing the size and place of menu items according to a traditional fish-eye technique (see Figure 39). A study showed that users preferred the fish-eye technique for browsing and hierarchical menus for goal-directed tasks. MCGUFFIN and BALAKRISHNAN explored several expansion techniques in selection tasks [MB02]. They conclude that expanding widgets can be used without sacrificing performance. They found that such tasks can be accurately modelled by Fitt's Law² even if the target only expands after 90% of the distance towards the target has been traversed. Two discrete lens prototypes were developed for the study, as can be seen in Figure 40.

²Fitt's Law: The time to acquire a target is a function of the distance to and size of the target.




Figure 39: A fisheye menu.




Figure 40: Interaction prototypes, images courtesy of MCGUFFIN.

6.6 Ongoing Work with the Elastic Presentation Framework

6.6.1 Bubble Lens

Nelson Wong, a graduate student at the Innovations in Visualization Laboratory currently researches the use of interactive distortion of graph structures. He developed a discrete bubble lens, which influences the display of edges in graphs. His lens curves edges according to the magnification factor calculated for a specific point on the edge. Some results are shown in Figure 41. The University of Calgary Computer Science Webpage is displayed as a graph once undistorted and once with a bubble lens applied.

6.6.2 Semantic Zoom

Eric Pattison, a graduate student at the the Innovations in Visualization Laboratory, researches *semantic zoom* with multi-resolution images using the Elastic Presentation Framework. Semantic zoom involves not only the increase of objects or representations in size but the display of additional detail and/or information. He developed an algorithm for semantic zoom into low resolution images based on underlying high-resolution images of which parts are displayed at the focal region. An example can be seen in Figure 42.




Figure 41: Graph distortion. Pictures courtesy of Nelson Wong, University of Calgary.




Figure 42: Image distortion. Pictures courtesy of Eric Pattison, University of Calgary, Original Image: Orthoshop.

6.6.3 Discrete Distortion in 3D

Julia Schliebenow an intern student at the Innovations in Visualization Laboratory, conducted research on the interaction with and the distortion of three-dimensional data. The concepts developed in EPF have proven to be extensible to three-dimensional data representations. Her work extends the concept from an over-all data distortion to a user controlled probe for 3D data.

7 Conclusion

In this thesis, new ways of thinking about different types of visual data distortion methods were developed. Terminology to distinguish between two visual distortion methods, *discrete* and *continuous lenses*. This terminology can be used regardless of the used distortion algorithms.




Figure 43: Distortion in 3D data, images courtesy of Julia Schliebenow, University of Magdeburg.

It is useful since the two methods are visually distinct. The Elastic Presentation Framework has been used as a basis for the exploration of the shown methods and the description of how to build a discrete lens. The ideas, concepts, and definitions about discrete and continuous lenses have been verified and tested in example implementations of which many can be found in this thesis. All examples have been implemented by the authors unless stated otherwise. This work does not include the description of the different types of spaces that are involved in data presentation. The representation can reside in a certain space which can also be distorted, divided, transformed or left unchanged during the creation of the final presentation. These thoughts had to be left out because it would have been too complex to describe in this thesis. It was felt that the other topics needed to be covered in the presented detail to be understandable. The work on the distortion and use of spaces or projections will have to be considered future work on the basis of this thesis. This thesis is meant to be a theoretical framework for future research on discrete and continuous lenses. It has to be explored whether more general distortion effects can be found and how they can be related to by the visualization tool designer. The student assignments presented in Section 5.9 have already shown sample uses of discrete lenses on the basis of an early version of the work presented here. As mentioned in Section 6, previous well-known lens techniques have been implemented discretely. However it will have to be seen if new methods will be found on the basis of the thoughts and ideas presented here.

Appendix

This appendix lists terminology needed to understand the concepts of this work. For further reference to unlisted definitions, please refer to the work presented by CARPENDALE in [Car99].

Lens

A lens combines the set of choices that can be made to create a presentation. It includes the selection of type(s) of distortion and making deliberate choices on how to apply these to the given representation and/or data. Therefore, to create a lens choices have to be made about the data you want to use, the type and parts of a representation you want to affect for the given dataset, the space in which you want to put your representation, and the type(s) of distortion you want to use.

Continuous Lens

A continuous lens on the computer is a visually continuous lens. On the computer, and arguably in the real world, there is no actual continuity if you think of the mathematical definition for continuity. However, you can intentionally create the effect of a visually continuous lens. Many graphics methods have been developed to mimic visual continuity, and while they may achieve convincing results, there will be a limit. Similarly distortion methods have been developed to mimic visually continuous distortion for which there, too, will always be a limit. The data has to be accessed so that each visible point (i. e., pixel) on the surface of a representation is changed according to the applied distortion function. The actual distortion is done above display resolution so that for each point on the display one or more data-points can be found that have been changed according to the applied distortion.

Discrete Lens

A discrete lens makes choices about use of the data and/or representation and/or distortion so that at least some points on the representation that will be displayed are not adjusted according to the chosen distortion function. The distortion for a discrete lens is so applied below the resolution of the output medium. A choice then has to be made on how to “fill in the gaps”.

Discrete-complete-continuous

Discrete: Defined for a finite or countable set of values [The00].

Continuous: For any two numbers there exists a number inbetween.

Complete: Format which is discrete by the above definition but not separable (e. g., rasters).

Distortion

A lens distortion consists of a field or a piecewise group of selected field functions which, in our case of lens distortions, are applied to a representation in a certain way. A choice of connections can be made for a group of field functions. They can be disjoint (having undefined parts), adjacent (being discontinuous but having no undefined parts) ,or connected.

Field

A field, or field function, is a mathematically continuous function of a certain dimension which can be part of a lens distortion.

Magnification

A mathematical term which can be calculated out of a certain distortion of data. In Elastic Presentation Framework magnification is defined by the height by which a data point is distorted.

Representation

A representation is a formal system by which information can be specified. Representations portray things or the relationships between them and can be found in many different forms. It consists of at least one object of one or more types. It has to have a defined relationship to the underlying data.

Object

An object is a geometric realization or manifestation of data. It can be composed of other objects as well.

List of Figures

1	Menu bar in Mac OS X.	6
2	Visualization pipeline as seen in [Car99].	10
3	Representations for the concept of the number thirty-seven.	11
4	Geometric setup.	13
5	Parts of an EPF lens.	14
6	Different focal shapes and their focal points.	15
7	Different drop-off functions within EPF.	15
8	Magnification calculation within EPF (image courtesy of Sheelagh Carpendale).	16
9	The four supported distance metrics: L_1 , L_2 , L_3 and L_∞	17
10	A new EPF lens created by combining a polyline lens and a circular lens.	18
11	The distorted “rubber-sheet”.	20
12	Distortions to an image. Pictures courtesy of Eric Pattison, University of Calgary, original image: Orthoshop.	20
13	A 1D- and 2D-continuous lens.	21
14	A continuous color lens.	21
15	Discrete vs. continuous lenses.	22
16	Discrete squares moved by their center point in 2D with a magnified area.	23
17	One discrete object and its geometric properties.	24
18	Discrete lenses.	26
19	The EPF feature space.	26
20	Notice the difference in overlap with changing focal size.	28
21	Artistic distortion of an NPR image achieving caricature-like renditions.	31
22	Continuous-raytracing-lens.	34
23	Discrete lens applied to subdivision and shading	35
24	Discrete lens applied to sound volume.	35
25	Reference points laid out along major axis.	36
26	The EPF distortion parameter.	36
27	The EPF distortion parameter.	37
28	Effect of Z-displacement in perspective projection.	39
29	A discrete image lens, images courtesy of Mike Larke.	40
30	The Ansel-Adams-Lens, image courtesy of Kathryn Elliot.	41
31	A Word Synonym Lens, images courtesy of Michael Boyle.	41
32	Techniques used in creating the original bifocal display.	42
33	A discrete bifocal EPF lens.	43
34	Perspective Wall.	43
35	Transformation functions used in [SB92].	44
36	What most people understand as a fisheye-lens.	45
37	Two methods for transforming discrete objects, images courtesy if Keahey.	46
38	Left unzoomed view. Right flip zoomed view; image courtesy of HOLMQUIST.	47
39	A fisheye menu.	48
40	Interaction prototypes, images courtesy of MCGUFFIN.	48
41	Graph distortion. Pictures courtesy of Nelson Wong, University of Calgary.	49

42	Image distortion. Pictures courtesy of Eric Pattison, University of Calgary, Original Image: Orthoshop.	49
43	Distortion in 3D data, images courtesy of Julia Schliebenow, University of Magdeburg.	50

References

- [B.B00] B.B. Bederson. Fisheye Menus. In *Proceedings of ACM Conference on User Interface Software and Technology: UIST 2000*, pages 217–226. ACM Press, 2000.
- [Ber83] Jacques Bertin. *Semiology of Graphics*. University of Wisconsin Press, 1983.
- [BSP⁺93] Eric A. Bier, Maureen C. Stone, Ken Pier, William Buxton, and Tony D. DeRose. Toolglass and Magic Lenses: The See-Through Interface. In *Proceedings of SIGGRAPH '93*, volume 27 of *Computer Graphics Annual Conference Series*, pages 73–80, 1993.
- [Byr47] Oliver Byrne. *The first six books of the Elements of Euclid*. William Pickering, 1847.
- [Car99] Sheelagh Carpendale. *A Framework for Elastic Presentation Space*. PhD thesis, Simon Fraser University, March 1999.
- [CCF96] M. S. T. Carpendale, D.J. Cowperwaite, and F.D. Fracchia. D Viewing Techniques for 3d Data. In *Vis'96: Proceedings of the IEEE Conference on Information Visualization*, pages 46–53. IEEE Computer Society, 1996.
- [CMS99] Stuart Card, Jock Mackinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufman, San Francisco, CA, 1999.
- [Eng02] English-Russian Glossary on Information Society. Web page: <http://www.iis.ru/glossary/index.en.html>, November 26, 2002.
- [Fur86] G.W. Furnas. Generalized Fisheye Views. In *Proceedings of the ACM Conference on Human Computer Interaction: CHI'86*, pages 16–23, 1986.
- [GS02] Carl Gutwin and Amy Skopik. Fisheye Views are Good for Large Steering Tasks. 2002. accepted to CHI'03.
- [hCR98] Ed Huai hsin Chi and John T. Riedl. An Operator Interaction Framework for Visualization Systems. In *Proceedings IEEE Symposium on Information Visualization 1998*, pages 63–70, 1998.
- [Hol00] Lars Erik Holmquist. *Breaking the Screen Barrier*. PhD thesis, Göteborg University, Dept. of Informatics, 2000.
- [KR96a] T. Keahey and E. Robertson. Non-linear image magnification, 1996.
- [KR96b] T. A. Keahey and E. L. Robertson. Techniques for non-linear magnification transformations. In *Proceedings IEEE Symposium on Information Visualization 1996*, pages 38–45, 1996.
- [LA94] Y.K. Leung and M.D. Apperley. A Review and Taxonomy of Distortion-Oriented Presentation Techniques. In *ACM Transactions on Computer-Human Interaction*, volume 1, pages 126–160. ACM Press, 1994.

- [LS97] K. Lam and R. Spence. Image Browsing - a space-time trade-off. In *Proceedings INTERACT '97*, pages 611–612, 1997.
- [Mar82] David Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman & Company, 1982.
- [MB02] M. McGuffin and R. Balakrishnan. Acquisition of Expanding Targets. In *Proceedings of Conference on Human Factors in Computing System: CHI '02*. ACM Press, 2002.
- [MRC91] Jock D. Macklinay, George G. Robertson, and Stuart K. Card. The Perspective Wall: Detail and Context Smoothly Integrated. In S. Robertson, G. Olson, and J. Olson, editors, *Proceedings of the ACM Conference on Human Factors in Computing Systems: CHI91*, pages 173–180, New Orleans, Louisiana, 1991.
- [NE78] N.Kadmon and E.Shlomi. A Polyfocal Projection for Statistical Surfaces. In *The Cartographic Journal*, volume 15(1), pages 36–41, 1978.
- [New01] G.B. Newby. Cognitive Space and Information Space. *Journal of the American Society for Information Science and Technology*, 52(12):1026–1048, 2001.
- [SA82] R. Spence and M. Apperley. Data Base Navigation: An Office Environment for the Professional. In *Behaviour and Information Technology*, volume 1(1), pages 43–54, 1982.
- [SB92] M. Sarkar and M.H. Brown. Graphical Fisheye Views of Graphs. In *Conference Proceedings on Human Factors in Computing Systems: CHI'92*, pages 83–91, New York, USA, 1992. ACM Press.
- [Sim96] Herbert Alexander Simon. *The Sciences of the Artificial*. M.I.T. Press/Trilateral, 1996.
- [SM95] Margret-Anne D. Sotrey and Hausi A. Müller. Graph Layout Adjustment Strategies. In *Proceedings of Graph Drawing*, Passau, Germany, September 20–25 1995.
- [Spe00] Robert Spence. *Information Visualization*. Addison-Wesley Pub Co, 2000.
- [SSOR93] M. Sarkar, S. Snibbe, O.J.Tversky, and S.P. Reiss. Stretching the Rubber Sheet: a metaphor for viewing large layouts on small screens. In *Proceedings of the ACM Symposium on User Interface Software Technology: UIST '93*, Atlanta, Georgia, USA, 1993. ACM Press.
- [SW63] Claude E. Shannon and Warren Weaver. *A Mathematical Theory of Communication*. Univ of Illinois Pr (Pro Ref), 1963.
- [The00] *The American Heritage Dictionary of the English Language*, volume Fourth Edition. Houghton Mifflin Company, 2000.
- [Tuf90] E.R. Tufte. *Envisioning Information*. Graphics Press, 1990.