Information Visualization
Focus and Context

DIVA Research Group - University of Fribourg - Switzerland

Lorenzo Clementi
lorenzo.clementi@unifr.ch

Winter semester 2006 - 2007

Abstract

Displaying a large amount of data on a small display (for example the display of a handled device such a mobile phone) in a clear, readable way is a difficult task. Even on bigger devices such as desktop computers, for which the amount of memory and the CPU speed have grown very fast, the screen size has not increased as fast as the amount of data that needs to be displayed. In order to produce a comprehensible visualization, **focus-and-context** techniques can be used: these methods present a small amount of data in great detail (the focus), while the rest of the data is rendered with less detail (the context).
This article summarizes six papers[1] either proposing concrete solutions to specific problems, or addressing the focus-and-context methodology from a theoretical point of view. We aim to provide a snapshot illustrating the state of the art of the focus-and-context domain.

Introduction

Reproducing large amount of information on a small space has been an issue since the first georaphical maps were to be drawn. Nowadays, the physical support on which the information is displayed has changed, but the problem still exists: we need to visualize more and more data on a small surface, keeping it readable. Since the 1970s, many researchers have addressed this problem producing several solutions that can be classified in two groups: **distortion-oriented** and **non distortion-oriented** techniques. The non distortion-oriented approach provides all the information at the same detail level: since not all the data can be displayed at once, scrolling is used to browse the whole information. Despite their frequent use, non distortion-oriented methods have a major weakness: the information displayed leaks of context, which makes its interpretation difficult. Conversely, distortion-oriented techniques provide a visualization of a relatively small amount of data with great detail, while the remaining information is rendered with less detail.

This paper is structured as follows: we first describe two mathematical concepts relatad to distortion-oriented presentation methods. Next, an overview of the main techniques both distortion and non distortion-oriented is provided. Then, we describe the frameworks proposed by Leung and Carpendale, introducing two methodologies to compare different visualization methods ([Leung], [Carpendale]). The last section is dedicated to the conclusions of our work.

Transformation and magnification functions

The information *representation* is the action of creating an image representing some kind of data (for example a graph, a map or a picture), whereas the information *presentation* is the action of displaying the image. Suppose that the information you want to display already has a graphical representation.
Distortion based techniques transform an undistorted image into a distorted image by applying a **transformation function**, as shown in in Figure 1. This transformation function can be applied either on one dimension (as in Figure 1) or on two dimensions.
The derivative of the transformation function is called **magnification function** and it provides the profile of the magnification factors applied to the undistorted image. Figure 2 shows the magnification function corresponding to the transformation function of Figure 1. In this example, the magnification function reaches its

[1] See References

maximum for x being close to zero, while it is almost zero when x is close to 1 and -1. As a matter of fact, the most magnified ("stretched") region of the distorted image in Figure 1 is situated in the middle (where x is close to zero), while near the borders the image is compressed.
The transformation and the magnification functions have two main practical applications: first, they are used to describe the behaviour of different distortion-based presentation methods. Second, some authors have used their features to classify distortion-based visualizations.

Focus-and-context presentation methods

Non distortion-oriented techniques use only one operation: **magnification or zooming**. Zooming allows two options: full-zoom or inset. The full-zoom magnifies the whole image and its result is a big image that does not fit into the screen area; for this reason dragging or scrolling are used to browse the hidden regions of the image. The inset technique, on the other hand, magnifies only a section of the image (which is called the *inset*). The inset method has a major drawback: the detailed region is displayed with little or no context (the result is called a *detail-in-context* image).

With respect to non-distortion based techniques, distortion based methods can display a larger amount of data on the screen, because not everything is rendered at the same detail level. The user can interactively select where the focus should be (for example using a mouse click) and the representation is consequently updated. All the methods we are going to describe in the next section are distortion-based.

Polyfocal display

This technique was first introduced to represent statistical data on a cartographic map: the principle behind this method is to "stretch" the part of the image where the focus is located, while compressing the remaining areas, as if you were looking at the image through a lens. The magnification function corresponding to a polyfocal display has one peak for each focus in the distorted image.

Bifocal display

The bifocal display in its two-dimensional form produces a distorted-image in which the focus is a rectangle positioned in the middle of the image and the context is situated at the edges of the image. Between the focus and the context areas there is a discontinuity of magnification (i.e. the zoom factor changes suddenly). The context region is not distorted uniformly: the focus line and the focus column are distorted only over one dimension (horizontal and vertical respectively), while the four corners are distorted on both dimensions. Later we will see how Rao and Card have improved this technique merging it with a fisheye distortion [Rao].

Fisheye view

There exists several kind of fisheye views. The original one is due to George W. Furnas, who proposed a method to reduce the information to be displayed, trading-off its *a priory* importance and the distance from the focus [Furnas]. In his paper, Furnas intruduces a **degree of interest** function which is given by the difference of the *a priori* importance of the item and its distance from the focus, as follows:

$$DOI_{fisheye}(x|. = y) = API(x) - D(x, y)$$

where DOI is the degree of interest function, the API is the *a priori* importance function and D is the distance between the current point of focus y and a point x.
Further, a threshold can be used: the data whose DOI value is less than the threshold is not displayed on the distorted image. It is important to note that the DOI function is an information-suppression function, which makes Furnas' fisheye method rather different from other distortion-oriented methods. As a case study, Furnas' paper illustrates how his fisheye view can be applied to tree structures such as, for example, structured text files or a code written in some programming language. In the context of a tree structure, the *a priori* importance is given by the inverse of the distance from the root – the more a point is far from the root, the smaller its importance is. Similarly, the distance between two points is the length of the path that connects them. Furnas states that this technique allows *a logarithmically compressed display of the original tree* [Frunas].

Fisheye menu: an example of fisheye view

Selecting an item from a long list has become a common task; a very concrete example is the selection of a website from a web browser bookmarks list. To address this problem, Bederson proposes a fisheye menu – a classical drop-down menu to which a fisheye distorted view is applied [Bederson]. The menu items have different font sizes: those close to the mouse pointer are bigger, while the size of the other entries is smaller. The minimum and maximun font sizes are computed when the

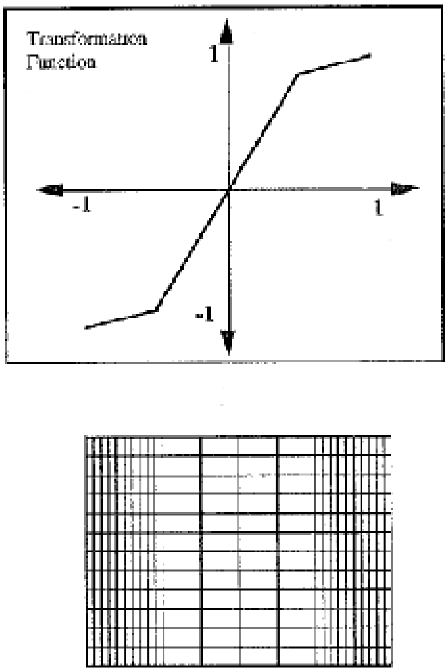

Figure 1: A transformation function and the result of its one dimensional application on a grid (source: [Leung]).

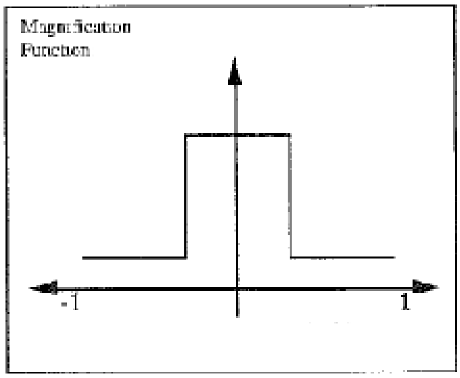

Figure 2: The mangification function corresponding to the magnification function of Figure 1 (source: [Leung]).

menu is instantiated and this allows the whole menu to fit into the screen area. In order to make the selection easier, the items are ordered alphabetically; on the left-hand side of the menu there is an alphabetical index helping the user to quickly reach the region containing the required item. Further, by moving the mouse pointer to the righthand side of the menu, the focus area is blocked, allowing to select an item appearing within the focus region.
Great attention has been given to testing the fisheye menu. A group of 10 users has been selected, half of them having programming experience and the other half without any programming knowledge. The test proposed a comparision between the fisheye menu and other kind of menus. Interestingly, the overall appreciation for the fisheye menu was rather high (6.4/10); this result was even better among programmers (7.0/10) who rated the fisheye solution as the best choice. On the other hand, non programmers' appreciation was slightly lower (5.8/10).
While a fisheye menu requires some time to get acquainted with, it provides an effective way for selecting items from a long, non hierarchically-ordered list.

Table lens: another example of fisheye view

Rao and Card have extended the concept of Bifocal display by applying it to a spreadsheet treatment application [Rao]. As a matter of fact, humans are able to spot patterns much better than what machines can do; however, our ability to observe and understand data structures critically relies on how the information is displayed.
The table lens applies to a conventional relational table the concept of the bifocal display, defining in this way four regions on the resulting presentation: focal, row focal, column focal and context. Further, two DOI functions - one for each dimension (horizontal and vertical) - are used to compute the size of each cell. The DOI functions work independently and compute a cell size according to its distance from the focus area. The user can interact with the table lens application by using three main operations:

- *Zoom* increases or decreases the space allocated to the focal area.
- *Adjust* changes the number of cells within the focal area (without modifying its dimension).
- *Slide* moves the focal area to another location.

The table lens uses additional graphical features such as color and shading to improve the rendering of

the data. These features depend on several factors, for example the cell position with respect to the focus or its content value and type.
The table lens has been tested on six different data tables and it has proven to be a very effective technique to display large spreadsheets and to explore the data they hold. Furthermore, its functionalities have made it possible for the user to extract correlations and patterns that would otherwise have been ignored.

Perspective wall

The perspective wall provides a central detailed view and two distorted sideviews. This technique is similar to a one-dimensional Bifocal display, except for the distortion of the sideviews that is constant in the Bifocal method while it increases close to the borders in the Perspective wall method.

Radial space filling

The paper by Stasko and Zhang illustrates a visualization technique used to display information hierarchies [Stasko]. A hierarchy (for example a file system) can be represented using a radial space filling technique in which the inner circle is the root element and the outer circle arcs represent the files or the directories located in the first hierarchy level. Next, the directories are again subdivided into outer circle arcs, and so on.
The main weakness of this representation is that, if the hierarchy is deep, the elements close to the leafs are represented by tiny circle arcs, making the comprehension of the resulting image very difficult. Stasko and Zhang have tried to solve this problem, keeping the basic idea unchanged: the full circular view must be maintained, but there must be the possibility to magnify peripheral items, keeping the global view as context information. The result of their work is the proposal of three different visualization techniques based on the Radial space filling idea. These three versions differ only in the way the detailed view is displayed; in each of them an animated transition is used to switch from the global view to the focus-and-context view, and vice versa.

Two unified frameworks

So far we have seen several focus-and-context presentation methods: each of them provides a solution for a particular visualization problem. In order to compare these solutions, Lenung and Carpendale provide two frameworks that uniformly describe different focus-and-context presentation methods ([Leung], [Carpendale]).

The Elastic Presentation Framework (EPF)

Carpendale and Montagnese have developed an Elastic Presentation Framework describing and integrating several presentation techniques, both distortion-based and non distortion-based. The framework relies on the distinction between information representation and information presentation. The framework furnishes a set of tools that can be applied to the data representation in order to display it in a focus-and-context fashion.
The basic elements behind the EPF are a three dimensional space and a lens. The image representing the data is projected on the base plane (dimensions x and y) and - at some distance d_b - a reference viewpoint is set (dimension z). Magnification is achieved by reducing the distance between the base plane and the viewpoint (d_b). To produce multi-scale presentations (presentations having a focus-and-context appearance) the lens technique is used. A lens has a focus region, a compression region and a context region. A *drop-off* function determines how the compression zone is deformed in order to shade out from the focus to the context region. This concept is applied to the base plane as follows: for each point x, y the distance from the focus is computed and, according to the drop-off function, the current point is set to its new position.
The last operation allowed by the EPF is folding, which is the displacement of a focus region with respect to the reference viewpoint. In this case the magnified region covers part of the context, thus only part of it is preserved in the resulting image.

Both non distortion-based and distortion-based presentation methods can be reproduced by modifying the drop-off function that defines the transition between the magnified region and its context. Further, the possibility to use several distance metrics - in addition to the ability of folding focus region(s) - make this framework highly parametrizable. These options make the EPF a good tool to describe and compare focus-and-context presentation methods.

The rubber sheet

Leung proposes a more intuitive idea to describe distortion-based presentation methods. His framework is based on the idea of a rubber sheet fixed on a rigid frame. Suppose that the data representation, in the form of an image, is printed on the sheet, which can be stretched or shrunken: *any stretching in one part of the sheet results in an equivalent amount of "shrinkage" in other areas* [Leung].
Applying a zoom factor to a region of the image corresponds to stretching the rubber sheet, while at the same time the other regions are shrunken. The amount of stretching and shrinking is modelized by the use of

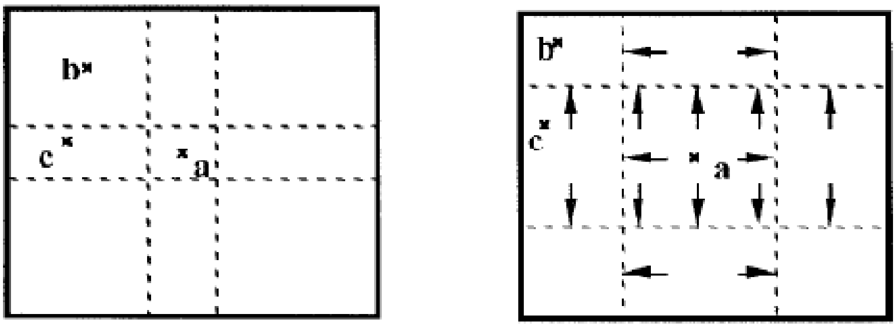

Figure 3: Illustration of the rubber sheet framework (source: [Leung]).

vectors; the functioning of the rubber sheet framework is illustrated in Figure 3. The magnified area contains a certain number of outbound vectors that represent the stretching deformation. Similarly, these vectors must be balanced by other vectors in the remaining areas of the image (the focus line, the focus column and the four corners). The overall resulting vector must be null, but in every single section of the image there can be a resulting vector not null; this vector indicates how the points in that region are displaced. This method works in the same way also with multiple focus areas and it provides a rather intuitive way for describing distortion-based presentation methods.

Conclusion

The most common techniques to display large amount of data on a computer screen are non distortion-based. These methods, although very popular, are far from being perfect. Distortion-based techniques provide an interesting alternative that is often ignored.

In this article we have gone through the main distortion based presentation techniques, pointing out their features and providing references to more specific documentation. The most important focus-and-context presentation method is the fisheye view. This technique relies on a degree-of-interest function that selects the data to be displayed according to two parameters: the *a priori* importance and the distance from the focus region. Fisheye views are not the only focus-and-context methods: hierarchically structured data, for example, can be represented by a radial space-filling approach.

The existence of many focus-and-context techniques generates the need for a theoretical base on which these methods can be described and compared. As we have seen in the second half of this article, there exist two frameworks whose aim is to provide a unified environment to compare different focus-and-context approaches. The EPF is more technical and it constitutes a good toolkit to describe both non distortion-oriented and distortion-oriented methods. The rubber sheet framework, on the other hand, is more intuitive, but it can be applied to distortion-based techniques only.

Distortion-based presentation methods are seldom used in commercial applications; however, several interesting techniques exist in this field. We think that their use should be seriously considered in situations where a large amount of data has to be graphically displayed. However, some of the currently existing techniques still require improvements in order to increase their ease of use, without which they will remain inaccessible to unspecialized users.

References

[Bederson] Bederson, Benjamin, *Fisheye Menus*, Proceedings of UIST 2000, San Diego, CA, Nov. 2000, pp. 217-225.

[Carpendale] Carpendale, M.S.T and Montagnese, Catherine, "A Framework for Unifying Presentation Space", Proceedings of UIST 2001, Orlando, FL, Nov. 2001, pp. 61-70.

[Furnas] G. Furnas, *Generalized Fisheye Views*, Proceedings of SIGCHI '86, New York, pp. 16-23, 1986.

[Leung] Leung, Y. and Apperley, *A review and taxonomy of distortion-oriented presentation techniques*, ACM Transactions on Computer-Human Interaction, Vol. 1, No. 2, June 1994, pp. 126-160.

[Rao] Rao, Ramana and Card, Stuart K., *The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus+Context Visualization for Tabluar Information*, pp. 318-322, Proceedings of the ACM SIGCHI '94 Conference on Human Factors in Computing Systems, Apr. 1992, Boston, MA.

[Stasko] Stasko, John T. and Zhang, Eugene, *Focus+Context Display and Navigation Techniques for Enhancing Radial, Space-Filling Hierarchy Visualizations*, Proceedings of IEEE Information Visualization 2000, Salt Lake City, UT, October 2000, pp. 57-65.