

Focus and Context for Volume Visualization

Marcelo Cohen and Ken Brodlie
School of Computing
University of Leeds
{marcelo,kwb}@comp.leeds.ac.uk

Abstract

Scientific investigation and simulation have been producing increasingly large datasets. Usually that kind of data is visualized employing some Scientific Visualization technique. Another related field, Information Visualization, deals with non-scientific large scale data, using different approaches to achieve an effective understanding. This paper aims to bring these two fields together, proposing a taxonomy for volume visualization, employing Information Visualization techniques, particularly the focus and context idea. A prototype is being developed, which aims to implement and evaluate a number of solutions for this problem, identifying through demonstration which applications would benefit most from focus+context approaches. This paper describes work in progress.

1. Introduction

Increasingly large data sets have been generated by science, engineering and medical applications, either by simulation or direct observation. Direct interpretation of this data is usually impossible, due to two factors: size and complexity. To help solve this problem, the discipline of Scientific Visualization[17] has emerged. This brings together two rather distinct fields: the traditional sciences (Physics, Biology, Chemistry, Medicine, etc) and Computer Graphics. In a similar sense, we have been observing a growth of information being stored electronically, which is mostly a consequence of the Internet's emergence. A considerable amount of that information is not directly related to physical data and maybe is not even numeric, but more abstract data. This kind of data requires a different approach to be effectively visualized, which has led to the creation of a new field: Information Visualization[22]. Making a parallel with Scientific Visualization, we can possibly say that this new field now brings together the social sciences and computer graphics. Would it be possible to achieve gains by bringing these two fields rather closer together ?

Let us first consider the case of scientific data: many techniques have been developed to assist in the visualization of this kind of data. For example, in the case of volume data (three dimensional datasets), direct volume rendering [4] is one of the most used approaches. Medical applications are amongst the most popular ones: data acquired from a scanner (computerized tomography, magnetic resonance, etc) is fed to a volume rendering system, allowing physicians and radiologists to see internal structures and organs with much greater detail than with conventional methods.

However, in some cases there is too much data to be displayed at once on a computer display (or the display's resolution may be insufficient for practical use). A simple and widely used solution is to apply a magnification factor to get closer to a specific region. But by doing so, it is equally easy to get lost in the dataset. This is generally called *loss of context*, because we are no longer able to visualize the entire dataset. When we zoom in, we are *focusing* on a certain feature that is of interest. In the field of Information Visualization this problem is called *focus+context* [22] and a number of successful solutions have emerged. The challenge is to find a way of looking at a high level of detail at this area of focus, without losing the overall context.

This paper describes work in progress to explore the transfer of focus+context ideas from Information Visualization to Volume Visualization. Section 2 gives background to the two fields which we bring together in this research. In section 3 we present a taxonomy of focus and context methods for volume visualization, in terms of nature of focus region, type of focus effect and space in which the effect is applied. Section 4 details two particular methods that we have implemented, and describes their application to a medical dataset - this dataset is of a cerebral aneurysm in the Circle of Willis, requiring focus on the aneurysm without losing context of the surrounding arteries.

2. Related Work

This section provides background on the two main fields of interest to the present research: Volume Visualization and Information Visualization.

2.1. Volume Visualization

Extensive research has been done on the visualization of three-dimensional datasets. When we are trying to visualize this kind of data, there are primarily two main approaches:

- Surface extraction and reconstruction methods: these methods work best when there is a defined structure within the volume. Medical imaging is probably their best application. By applying diverse algorithms to segment and extract surfaces from the image data, this class of methods allow a precise visualization of almost any structure, at the cost of generating a potentially large number of polygons. The classical example among these algorithms is the isosurfacing approach: from a user-specified threshold value, a boundary surface is extracted through the use of some method, like marching cubes[15].
- Direct volume rendering methods: this class of methods deal with the raw 3D data, employing algorithms that build an image without requiring the creation of an underlying geometric structure. Examples of such methods include: ray casting [11, 12], splatting [24], shear-warp [6] and 3D texture mapping [1]. In some situations, volume rendering offers a clear advantage over surface extraction, as it is possible to visualize all the data at once, using different colours and opacities.

The 3D texture mapping method is one of the most efficient methods to do volume rendering, because it uses the graphics hardware capabilities, allowing the quick display of an entire volume. This is essential for many applications, especially in medicine. 3D texture mapping works by loading the original dataset (or part of it) into the graphics card's texture memory. Then a number of planes (or concentric shells, depending on the method) are drawn, sampling this texture in order to render a 3D object. Parameters such as opacity and colour can be changed, affecting the resulting image.

This algorithm is very effective, but needs a lot of texture memory. This is a lesser problem now, as graphics cards are increasingly getting better. But citing medical technology as an example case, resolution of the imaging devices (scanners) is constantly improving, so more texture memory will likely be required to render a 3D dataset. Thus some research has been done on optimizing texture utilization. Levoy [14] proposed a gaze-directed technique, where the quality of the rendering varied according to the user's

gaze (using a head mounted display). LaMar [7] introduced a multiresolution approach to volume rendering, based on a hierarchical texture. Later, Weiler [23] has improved on that, conceiving an algorithm to minimize rendering artifacts due to interpolation. Both rely on the specification of a focus area. In this way, they also deal quite effectively with the texture problem, but at the expense of losing image detail outside their focus area.

2.2. Focus+Context (Information Visualization)

According to Leung's [10] taxonomy of presentation techniques for large graphical data spaces, it is possible to categorize techniques in two different ways: non-distorted and distorted views. Examples of non-distorted approaches include scrolling and zooming, and hierarchical views, where part of the information is hidden. This is adequate for small datasets, but using those techniques it is very easy to lose the context in larger amounts of data.

The primary approach to solve this problem is the use of distortion-oriented displays (where a non-linear magnification is applied). This implies presenting the dataset in a distorted fashion, requiring fast processing speeds (for the computation of the distortions) and high resolution displays (needed to provide enough screen space to show the data, even though the approach itself minimizes screen occupation). Those were the main difficulties in the past, as the technology was not readily available. Now there are a number of very high resolution displays, although they still are rather expensive.

One of the first ideas was the Bifocal Display [21], where an information space was "folded" to allow simultaneous viewing of all data. This concept was implemented as the Perspective Wall [16]: using a three-dimensional perspective effect, it was possible to display all data but at the same time keep the focus on the non-distorted region. A two-dimensional extension of the idea was used earlier [9], to assist in the presentation of topological networks. In a similar way, Furnas [5] introduced the concept of a Fisheye view: by the suppression of irrelevant data it was possible to present a large number of structures. Sarkar and Brown [20] suggested the application of Furnas' text-based Fish-eye view in a graphical fashion, implementing the technique to effectively display large graphs.

Leung [10] classified distortions into cartesian (where each coordinate of each data point is distorted based on the scalar difference between it and the corresponding coordinate of the centre of focus) and polar (where each point of the data is distorted based on the actual (vector) distance between it and the centre of focus). Distortions also can be considered as continuous or non-continuous. A continuous distortion function defines some degree of magnifi-

cation at the centre of focus and a continuous mathematical function that decreases the magnification gradually with distance from the central point (such as a fisheye view). A non-continuous distortion function defines a fixed magnification in a region of a certain size, surrounding the centre of focus, and a mathematical function that decreases magnification of points outside of that region (such as a bifocal display).

Figure 1 presents some examples of continuous and non-continuous distortions and their effects on the data presentation.

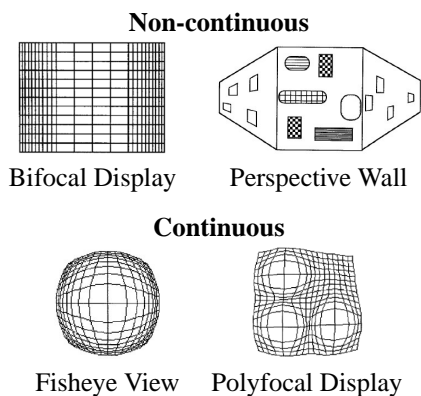


Figure 1. Common distortion techniques [10]

For three-dimensional datasets, the foundations were laid by Carpendale’s [2] research on the effects of distortion techniques on 3D data. Her work indicated that there is an additional complication when viewing 3D data: occlusion of the focus region. She proposed a distortion technique that allowed a clear visual path to this region. Also important is her analysis on the comprehensibility of distortions [3], based on human perceptual skills.

Recently Winch [25] presented the 3D Cartesian Fisheye Display and the 3D Polar Fisheye Display. The 3D Cartesian Fisheye Display produces distortions that are magnified along each axis from the centre of magnification, outwards to the edges. The 3D Polar Fisheye Display uses the radial distance from the centre of magnification to the point under consideration to determine the magnification factor. Figure 2 presents both approaches using a 12 x 12 x 12 dataset.




Figure 2. 3D fisheye distortions proposed by Winch [25]

Later, LaMar [8] proposed the use of a magnification lens to explore volume data, allowing one to selectively enlarge regions inside the volume.

Similarly, McGuffin [18] developed a strategy based on surgical metaphors for distortion of volume data: cutting, peeling, spreading. These metaphors were implemented through the use of direct manipulation techniques in 3D. Rendering only point primitives, he was able to achieve fast speeds.

3. A Framework for Focus and Context in Volume Visualization

Most information visualization techniques create an abstract representation of the original data. For example, a user’s file system can be represented as icons on a flat wall. To help comprehension, the representation can be distorted in order to make better use of screen space. In the wall example, we can bend the two ends of the wall backwards and view in perspective (the so-called “perspective wall”).

By contrast, volume visualization typically creates a representation which reflects a real spatial relationship amongst the data. We can apply distortion techniques but we need to be aware that we are changing those spatial relationships. Therefore this is one of the research goals: to find out how big an issue this is.

In the earlier part of the paper, we have explained how the “Focus + Context” concept in Information Visualization allows a user to study fine detail in a restricted region of a large dataset, without losing the context of the overall data space. In this section, we aim to develop a taxonomy for F+C in volume visualization, building a classification of existing techniques and identifying possible new approaches. We shall concentrate on direct volume rendering and we shall classify an approach in terms of three attributes:

- **Focus:** this defines the nature of the focus region, which can either be a point or a circular or rectangular region (in 2D or 3D depending on the Visualization

Space). In the case of a point, the focus effect is greatest at a single point, and falls off with distance towards the boundary of the visualization space. In the case of a region, the focus effect is constant within the specified region, and then falls off with distance towards the boundary.

- **Effect:** this defines the way in which the detail of the focus is achieved, and can be by refinement or by distortion, or by a combination of both. Refinement is an accuracy attribute, and indicates that the rendering within the focus region is carried out to higher accuracy (but there is no spatial distortion). Distortion on the other hand, achieves the focus effect by spatial deformation of the Visualization Space – the focus region is magnified. It is possible to combine both, where the greater spatial extent of the focus region is accompanied by higher accuracy rendering in that region.
- **Visualization Space:** this defines the space in which the F+C operation is applied, and can be either image space or object space, or indeed both. By object space, we are thinking in terms of the dataset being classified and shaded as though it were a block of inhomogeneous coloured gel – this gel is our object which is rendered.

This taxonomy therefore yields twenty-seven (3 x 3 x 3) different approaches, based on the values of these three attributes. Some approaches are well known, others are relatively new. We look at each in turn, as an instance of the triple {Focus, Effect, Visualization Space}:

1. **{Point, Distortion, Image Space}**: this is the well-known fisheye view of Furnas [5], and later Sarkar and Brown [20], applied to the image created by a volume rendering algorithm. The distortion is a simple magnification process, with pixels being replicated (or removed) from the original image. This can be applied as a post-processing step on the original image.
2. **{Rectangular Region, Distortion, Image Space}**: this is the bifocal display approach of Spence and Apperley [21], with similar effect to the above.
3. **{Circular Region, Distortion, Image Space}**: this is a variant of the Spence and Apperley bifocal display.

Having made this distinction between point and region, for simplicity we now treat them together in what follows, considering point as limiting case of circular region:

4. **{Point/Region, Refinement, Image Space}**: here the accuracy of each pixel in the image is related to its distance from the focus region. This approach has its ori-

gins in the flight simulation area, where high resolution insets were blended in to a low resolution background (see references in Levoy’s gaze paper). We can do exactly the same with volume rendering, simply by generating inset images at higher resolution. This is also related to Levoy’s “Volume Rendering by Adaptive Refinement” paper [13] where areas of greater image complexity are refined by casting more rays.

5. **{Point/Region, Distortion and Refinement, Image Space}**: this combines approaches 1, 2 and 3. As well as carrying out the spatial distortion, we also increase the relative accuracy within the focus region, either by increased sampling or by interpolation from existing samples.
6. **{Point/Region, Distortion, Object Space}**: in this approach, we apply a 3D spatial distortion to the dataset. Within a 3D focus region, the object (that is, the portion of gel within the focus region) is magnified by a certain factor (greater than 1), and the remainder is magnified with decreasing factors as we approach the boundary of the gel. A variation of this approach was presented by LaMar et al in their magnification lens paper [8]. This is also explored further in this paper, for both point and region.
7. **{Point/Region, Refinement, Object Space}**: in this approach, we model the object, i.e. gel, at higher accuracy within the focus region, before the gel is rendered. An example of this method is LaMar’s work on multiresolution volume visualization [7].
8. **{Point/Region, Both, Object Space}**: this would combine 6 and 7 above.
9. **{Point/Region, Distortion, Both}**: this has not been studied, but we could envisage the image distortion being applied as a post-processing step.
10. **{Point/Region, Refinement, Both}**: this is basically the idea in Levoy’s gaze directed volume rendering [14]. Refinement in image space is achieved by higher density of rays cast, and refinement in object space is achieved by higher sampling along the ray. Note that the focus region in object space is a frustum evolved from the focus region in image space.
11. **{Point/Region, Both, Both}**: this would combine 9 and 10 above.

To serve as a testbed for the implementation and evaluation of new methods, a prototype volume rendering system was developed. This prototype currently uses the traditional 3D texture mapping method, as proposed by [1]. But so far no distortion is actually taking place in the prototype: the distortion methods are being done in a stand-alone program, which reads a volume dataset, applies the desired dis-

tortion and creates a new dataset. Only then the new dataset can be visualized on the prototype.

4. Distortion Methods

According to the taxonomy proposed in section 3, it seemed natural to begin the research with existing methods. So this section discusses the implementation of two three-dimensional distortion methods based on two-dimensional counterparts: 3D cartesian bifocal and 3D cartesian fish-eye. Applying the terminology already defined, these methods can be classified as **{Point/Region, Distortion, Object Space}**, as there is no refinement yet.

Normally the distortion can be seen as a mapping function from original coordinates to new (distorted) coordinates. However, in our case we are dealing with discrete data (voxels). If the function was to be applied in the traditional way, the resulting volume possibly would have holes, as happens with image warping. The solution is to use the inverse of the function and work backwards, from distorted coordinates to original ones. This works as follows: first we create a new volume with the same size as the original one. Then we traverse this new volume, applying the inverse mapping function to obtain the original coordinates. Using the computed position, we fetch the corresponding voxel from the original volume and store it on the new one. This process is better illustrated in figure 3, which shows its application to a two-dimensional image.

Furthermore, it is worth noting that although this is essentially a nearest neighbour sampling, there is nothing that prevents a better interpolation method (e.g. trilinear) from being used. In such a case, the methods would be classified as **{Point/Region, Both, Object Space}**.

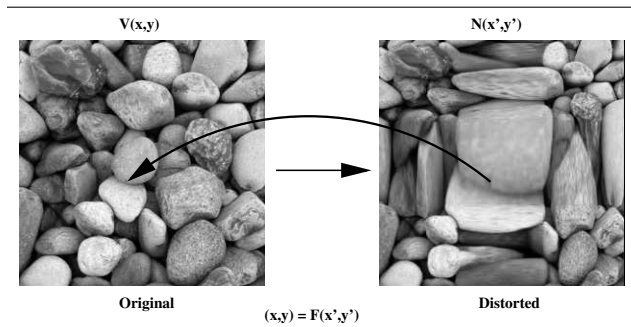


Figure 3. Inverse mapping with an image

4.1. 3D Cartesian Bifocal Display

This method can be seen as an extension to the 2D bifocal display: the focus is a rectangular region which will be uniformly enlarged by a factor and the remaining sections

of the volume will be compressed, keeping the same volume size. As stated in section 4, this is an object space method, as it effectively changes the original data before rendering. For simplicity, we are assuming all coordinates are normalized. The basic distortion itself is directly adapted from Winch's work [25], depending on three parameters:

- Centre of focus (x_f, y_f, z_f)
- Magnification factor (*mag*) - determines how much the focus area will be enlarged
- Size of each half of the focus region (s_x, s_y, s_z) in the original volume

The first step is to compute six scaling factors: for each axis and for each side of the focus region. These factors are needed to determine how much the volume is compressed to each side of the focus region. Figure 4 shows how the factors for the x axis are obtained. Let Δ_{x+} be the distance from the right edge of the original focus region (x_{max}) to the boundary; $\Delta_{x'+}$ be the corresponding distance after the distortion. Using similar notation for the left side, we define scaling factors:

$$scale_{x+} = \frac{\Delta_{x'+}}{\Delta_{x+}}, scale_{x-} = \frac{\Delta_{x'-}}{\Delta_{x-}}$$

(same for y and z)

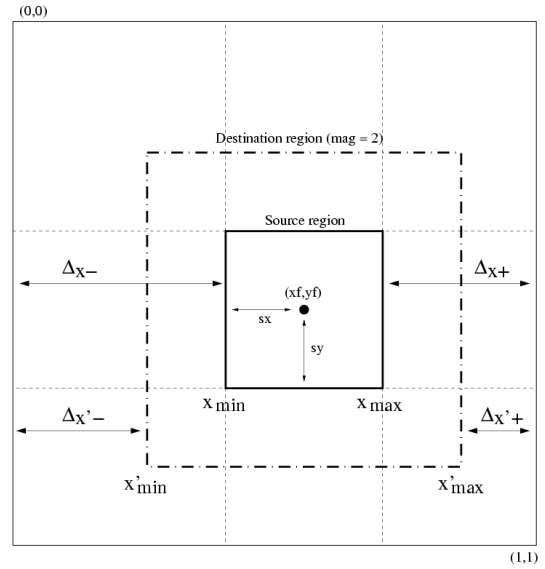


Figure 4. Scaling factors for the 3D bifocal cartesian distortion

Then, for each voxel in the new volume (distorted), the following algorithm calculates the corresponding **original**

voxel coordinate $V(x,y,z)$, using the scaling factors and the limits of both regions:

- if x' is outside focus region:
 - if $x' < x_f - mag \cdot s_x$ (left side):

$$x = \frac{x'}{scale_{x-}}$$
 - if $x' > x_f + mag \cdot s_x$ (right side):

$$x = 1 - \frac{(1-x')}{scale_{x+}}$$
- if x is inside focus region, just apply the magnification factor:

$$x = \frac{x' - x_f + mag \cdot x_f}{mag}$$
 (the formulation is the same for y' and z')

4.2. 3D Cartesian Fisheye Display

Like the previous method, this can be also be seen as an extension of a two dimensional technique, the fisheye display [20]. In this case there is no “focus region”, but a focal point (in this case, a voxel). The magnification factor - here specified by a “distortion” factor - is highest at this focal point, and decreases as we move towards the dataset borders. This produces a good deal of distortion on the borders themselves. Adapted from Sarkar & Furnas work, this method depends on two sets of parameters:

- Centre of focus (x_f, y_f, z_f)
- Distortion factor (d)

Like the cartesian bifocal method, we use the distances between the centre of focus and each dataset boundary. So considering the x axis, there are two distances - right (Δ_{x+}) and left (Δ_{x-}):

$$\Delta_{x+} = 1 - x_f$$

$$\Delta_{x-} = x_f$$

(same for y and z)

Then, by traversing each voxel in the distorted volume $N(x,y,z)$ the following algorithm calculates the voxel coordinate in the **original** volume $V(x,y,z)$:

- if $x' < x_f$

$$x = x_f - \frac{\Delta_{x-}(x_f - x')}{\Delta_{x-} + d(\Delta_{x-} + x' - x_f)} = x_f - \frac{x_f(x_f - x')}{x_f + dx'}$$
- if $x' > x_f$

$$x = x_f + \frac{\Delta_{x+}(x' - x_f)}{\Delta_{x+} + d(\Delta_{x+} - x' + x_f)} = x_f + \frac{(1-x_f)(x' - x_f)}{1-x_f + d(1-x')}$$
 (same for y and z)

4.3. Example of application

To demonstrate how the methods currently perform with real data, a section of a medical dataset has been used, comprising of 512 x 512 x 146 slices. This dataset contains a CTA (Computerized Tomography Angiography) of a patient’s Circle of Willis - a circle of arteries at the base of the brain. Congenital defects may occur in these arteries and result in the formation of aneurysms. An aneurysm is a dilation of an artery caused by the vessel wall yielding and stretching due to the pressure of the blood. The surgeon needs a precise view of the region around the aneurysm so that he/she can apply a clamp, but only needs an overview of the wider picture. Thus the idea is to use the aneurysm as a focal point, in order to assess whether the distortion is providing a better visualization and understanding of the data or not.

First we present the results, viewed on a single slice, in figure 5. The left image shows the original, unaltered slice, where the aneurysm can be seen as a small white ball. The centre image shows the distortion obtained from the 3D cartesian bifocal method. A small region around the aneurysm is used with a magnification factor of two. Finally, the right image presents the results from the 3D cartesian fisheye method. In this image, the focal point is the centre of the aneurysm but here the distortion factor was 1.2, trying to keep the aneurysm roughly with the same size obtained with the bifocal method. From the images, it can be seen how the fisheye method changes to a greater extent the spatial proportions of the data, while the bifocal method retains more of the original balance. This finding is also confirmed by looking at the volume rendered data, as seen in figure 6. In the right image, it is also easier to notice that the borders of the dataset are indeed very compressed. Therefore, although the bifocal method is non-continuous, this preliminary experiment shows that it is more appropriate for this particular medical dataset. Further experiments are to be conducted with other kinds of data.

5. Conclusions and Future Work

Volume rendering is a well-known, established technique to visualize three dimensional datasets. A number of different solutions have already been developed to visualize this kind of data, but most of them suffer from the *focus+context* problem outlined before: usually there is too much data and the most used approach is to zoom in and get closer to the desired feature of interest, at the expense of losing the overall context.

Therefore this study aims to provide a better understanding of how Information Visualization ideas can be applied to Scientific Visualization in an effective way, in particular for the visualization of very large volume datasets. This

is a relevant problem: in a panel held at IEEE Visualization 2002, Zuiderweld [19] mentioned that the ability to efficiently process very large datasets will be essential in the near future.

A contribution of this work is to develop a number of solutions to this problem, and identify through demonstration which applications would benefit most from *focus+context* approaches. We intend to investigate the feasibility of new techniques to help visualizing large scale data and to study how the distortion effects change our views and therefore, our understanding of the data. Some possible applications include (but are not limited to) medical data, CFD, meteorology, failure analysis, molecular modeling, seismic data and oceanography (time-varying) data.

6. Acknowledgments

This research project is being accomplished as part of a PhD research programme sponsored by the Brazilian Research Agency CAPES. We would also like to acknowledge the support of Nick Phillips, a consultant neurosurgeon at Leeds General Infirmary.

References

- [1] Brian Cabral, Nancy Cam, and Jim Foran. Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *Proceedings of the 1994 symposium on Volume visualization*, pages 91–98. ACM Press, 1994.
- [2] M. Sheelagh T. Carpendale, David Cowperthwaite, and F. David Fracchia. Distortion viewing techniques for 3-dimensional data, 1996.
- [3] M. Sheelagh T. Carpendale, David J. Cowperthwaite, and F. David Fracchia. Making distortions comprehensible. In *Visual Languages*, pages 36–45, 1997.
- [4] Robert A. Drebin, Loren Carpenter, and Pat Hanrahan. Volume rendering. In *Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, pages 65–74. ACM Press, 1988.
- [5] G. W. Furnas. Generalized fisheye views. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 16–23. ACM Press, 1986.
- [6] Philippe Lacroute and Marc Levoy. Fast volume rendering using a shear-warp factorization of the viewing transformation. In *SIGGRAPH '94*, pages 451–458. ACM Press, 1994.
- [7] Eric C. LaMar, Bernd Hamann, and Kenneth I. Joy. Multiresolution techniques for interactive texture-based volume visualization. In David Ebert, Markus Gross, and Bernd Hamann, editors, *IEEE Visualization '99*, pages 355–362, San Francisco, 1999.
- [8] Eric C. LaMar, Bernd Hamann, and Kenneth I. Joy. A magnification lens for interactive volume visualization. In *IEEE Pacific Conference on Computer Graphics and Applications*, pages 223–232, 2001.
- [9] Y. K. Leung. Human-computer interaction techniques for map-based diagrams. In G. Salvendy and M. Smith, editors, *Designing and Using Human-Computer Interfaces and Knowledge-based Systems*, pages 361–368, Amsterdam, 1989. Elsevier.
- [10] Y. K. Leung and M. D. Apperley. A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction*, 1(2):126–160, 1994.
- [11] Marc Levoy. Volume rendering: Display of surfaces from volume data. *IEEE Computer Graphics & Applications*, pages 29–37, May 1988.
- [12] Marc Levoy. Efficient ray tracing of volume data. *ACM Transactions on Graphics*, 9(3):245–261, 1990.
- [13] Marc Levoy. Volume rendering by adaptive refinement. *The Visual Computer*, 6(1):2–7, 1990.
- [14] Marc Levoy and Ross Whitaker. Gaze-directed volume rendering. In *Computer Graphics*, volume 24, pages 217–223, March 1990.
- [15] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 163–169. ACM Press, 1987.
- [16] Jock D. Mackinlay, George G. Robertson, and Stuart K. Card. The perspective wall: detail and context smoothly integrated. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 173–176. ACM Press, 1991.
- [17] B. H. McCormick, T. A. DeFanti, and M. D. Brown. Visualization in scientific computing. *Computer Graphics*, 21(6), November 1987.
- [18] Michael J. McGuffin, Liviu Tancau, and Ravin Balakrishnan. Using deformations for browsing volumetric data. In *IEEE Visualization 2003*, pages 401–408, October 2003.
- [19] Michael Meissner, Bill Lorensen, Karel Zuiderveld, Vikram Simha, and Rainer Wegenkittl. Volume rendering in medical applications: we've got pretty images, what's left to do. In *Proceedings of the conference on Visualization '02*, pages 575–578, 2002.
- [20] M. Sarkar and M. H. Brown. Graphical fisheye views of graphs. In *CHI'92*, volume 3:7, pages 83–91. ACM, May 1992.
- [21] R. Spence and M. D. Apperley. Database navigation: an office environment for the professional. *Behaviour and Information Technology*, pages 43–54, 1982.
- [22] Robert Spence. *Information Visualization*. ACM Press, 1st edition, 2001.
- [23] M. Weiler, R. Westermann, C. Hansen, K. Zimmerman, and T. Ertl. Level-of-detail volume rendering via 3d textures, 2000.
- [24] Lee Westover. Footprint evaluation for volume rendering. In *Proceedings of the 17th annual conference on Computer graphics and interactive techniques*, pages 367–376. ACM Press, 1990.
- [25] D. Winch, P. Calder, and R. Smith. (focus+context)3: distortion-oriented displays in three dimensions. In *User Interface Conference, 2000. AUIC 2000*, pages 126–133, 2000.




Figure 5. Non-distorted (left), cartesian bifocal (centre) and cartesian fisheye (right) distortions in a single slice




Figure 6. Non-distorted (left), cartesian bifocal (centre) and cartesian fisheye (right) distortions, volume rendered
