Comparing ExoVis, Orientation Icon, and In-Place 3D Visualization Techniques

Melanie Tory

School of Computing Science
Simon Fraser University
mktory@cs.sfu.ca

Colin Swindells

Computer Science Department
University of British Columbia
swindell@cs.ubc.ca

Abstract

With large volume data sets, it can be difficult to visualize the data all at once. Multiple views can address this problem by displaying details in areas of interest while still keeping track of the global overview. Many "detail and context" techniques exist for volume data, but it is unclear when to use each one. We introduce a new class of methods called ExoVis, an alternative that balances trade-offs of existing techniques. We then heuristically compare ExoVis to existing methods to provide insight into when each technique is appropriate.

Keywords: volume visualization, detail and context, direct volume rendering, isosurface, slice, subvolume.

1 Introduction

Gaining insight from volume data sets can be challenging because high data density makes it difficult to view all the data simultaneously. For example, a medical scan might generate a volume with 256^3 voxels. Since we cannot perceive objects occluded in depth, we typically display 256 two-dimensional (2D) slices if we want to see every detail. Here the third inherent spatial dimension is lost, forcing viewers to mentally reconstruct a three-dimensional (3D) model of the data. Although some professionals (e.g., radiologists) have extensive training in mental reconstruction, other people who use visualization tools (e.g., family physicians and patients) can find this task difficult. Furthermore, on a typical computer monitor, only a few slices can be displayed at full resolution, forcing us to either view the slices a few at a time or drastically reduce the image resolution. Displaying multivariate volume data is even more challenging. For instance, the amount of data in the medical example could double if the patient had two types of medical scans, such as Magnetic Resonance Imaging (MRI) and computed tomography (CT).

Methods to visualize volume data in 3D have been extensively studied – primarily isosurface extraction [11] and direct volume rendering (DVR) [15]. These methods successfully portray the 3D nature of the data set, but each voxel contributes only a small amount to the final image, such that details are lacking.

A common solution is to allow users to interact with a data set to create multiple views and select only the important data. Many multiple-view approaches allow users to simultaneously view details in areas of interest and a global overview showing the location of all detail views. This "detail and context" approach allows users to zoom in without losing track of their overall location.

We present a new detail and context technique for volume data. Our method combines a 3D overview with surrounding views of slices or subvolumes (the details). Examples are shown in Fig. 1d and Fig. 3. Because the details are "outside" or "surrounding" the overview, we call these views "ExoVis widgets" (from the Greek "exo-", meaning "outside" or "external"[1]). ExoVis widgets are motivated by "callouts" in technical drawings (e.g., medical textbooks), and are similar to callout ideas in the calendar visualizer [12]. Because the details surround the overview, ExoVis widgets are similar in concept to Stasko and Zhang's "Detail Outside" method for hierarchy visualization [18].

After introducing ExoVis, we heuristically compare it to existing detail and context approaches, specifically "in-place" and "orientation icon" methods (defined in Section 2). This comparison provides insight regarding when each technique is most useful.

---

1. *Webster's Encyclopedic Unabridged Dictionary of the English Language* (New York: Portland House, 1989), p. 500.

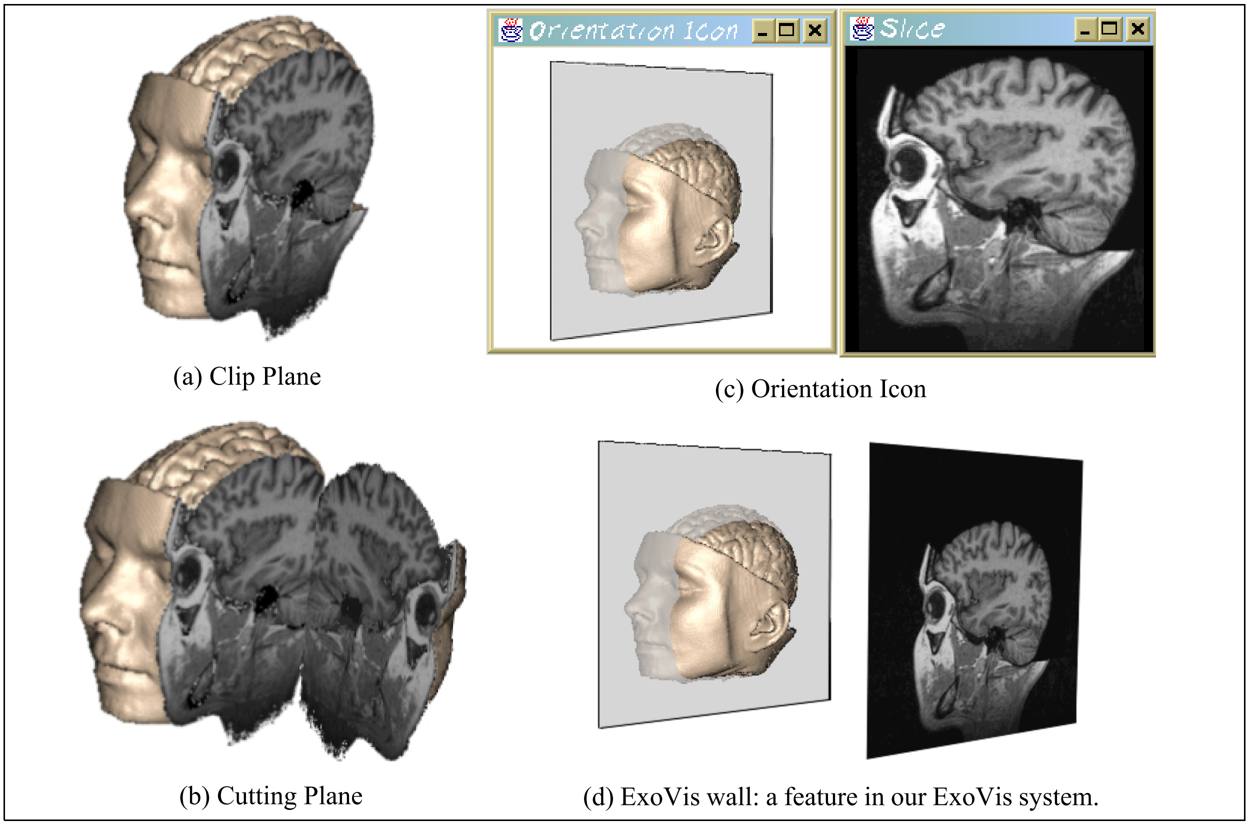

(a) Clip Plane
(c) Orientation Icon
(b) Cutting Plane
(d) ExoVis wall: a feature in our ExoVis system.

Figure 1: *Methods to combine a 2D slice with a 3D overview. Examples show a slice through an MRI brain scan.*

In the "Task by Data Type Taxonomy", Shneiderman introduces guidelines for developing effective visualizations [17]. He suggests that visualization tools should support the following generic tasks:

- **Overview:** Gain an overall picture of the data and keep track of the global location of items.
- **Zoom:** Adjust the size of items of interest.
- **Filter:** Remove uninteresting items.
- **Details-on-demand:** See details of an interesting area or item when desired.
- **Relate:** Identify relationships between items.
- **History:** Easily undo and redo actions.
- **Extract:** Extract a subset of the data set for separate analysis.

We use Shneiderman's guidelines as the basis of our comparison criteria.

2 Existing Techniques

2.1 Slice Details with 3D Context

To combine 2D slice details with a 3D overview, three main approaches exist: clip planes, cutting planes, and orientation icons. Differences between these methods are illustrated in Fig. 1 (a, b, and c).

Clip planes are common in many domains (e.g., geological visualizations). They show slice details in their exact relative position to the 3D context (i.e. the slice is "in-place"), but adding a clip plane removes all data between the plane and the viewer (see Fig. 1 a). Showing a slice deep within the volume removes almost all of the 3D context information. A similar method, the "planar brush", shows a cross-section "in-place" [20]. Although the planar brush does not clip away 3D information, 3D context is limited to a simple outline or semi-transparent surface. Another "in-place" alternative is to cut open the volume using a book or cutting metaphor [2, 4, 8], as in Fig. 1 b. Thus the overview information is not removed, but simply pushed aside.

With an orientation icon, the overview and details are in separate windows, and positions of the detail views are indicated in the overview (see Fig. 1 c). Orientation icons are common in medical imaging. Medical images have been traditionally viewed as sets of 2D slices; an orientation icon simply adds an extra window (containing a 3D overview) to any 2D slice viewer. Gerlach and Hersch describe an example [5].

2.2 Subvolume Details with 3D Context

Analogous detail and context techniques exist for subvolumes. In the orientation icon approach, each subvolume has its own window. A separate window shows a global overview and indicates subvolume positions via geometric primitives, as shown in Fig. 2.

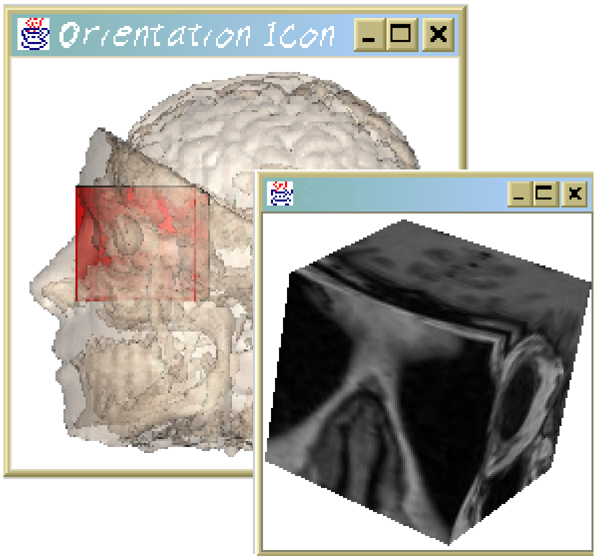

Figure 2: *Orientation Icon with a 3D subvolume. Subvolume position is indicated by the shaded red box.*

"In-place" techniques keep subvolume details in their original positions within the volume, and enhance subvolume details by magnifying them and/or changing their view properties. Most magnification techniques use image distortion. Fish-eye lens techniques greatly magnify objects at the centre of the field of view with a continuous fall-off in magnification towards the edges. Such lens techniques are commonly applied to 2D images, where the image is modelled as a pliable surface that can be stretched to emphasize specific areas [3, 14].

Several extensions of the fish-eye lens for discrete 3D data sets exist [4, 6, 13]. This distortion method is most effective when a 3D grid is drawn along with the data set. Bends in the grid provide perceptual cues to help users understand what type of distortion has occurred. 3D distortion displays for volume data have not been explored in much detail, but some examples do exist:

- LaMar et al. integrated a 3D distortion lens with a texture-based volume renderer [9], and
- Kurzion and Yagel developed a 3D distortion method [8]; however, it was not evaluated as a focus + context technique.

For volume data, it is difficult to draw a 3D grid that will not be occluded; thus, lens distortions are harder to understand. Interacting with the lens may help clarify the distortion, but angles may still be unclear.

Several techniques change display properties to make in-place subvolumes more visible. "Volume brushes" show subvolumes within an overview drawn as a simple outline, similar to planar brushes [20]. In direct volume rendering, transfer functions make voxels semi-transparent (so that voxels behind them may be seen). Thus a transfer function determines how much each voxel contributes to the image. Boyles and Fang developed a system that allows users to specify different transfer functions for different regions of a volume [1]. Similarly, Shaw et al. allow users to render regions of interest in various styles by moving 2D "plates" through the data set [16]. Data under a plate is rendered with that plate's style, while other data is rendered normally. These techniques can make subvolumes more visible but do not increase the screen space allotted to them.

3 Our New Technique: ExoVis

We propose two types of ExoVis widgets for viewing overview and detail components of volume data: a "wall" displays a 2D slice of the volume, and a "callout" displays a 3D subvolume. The global context object (also called the overview object) represents the entire data set and can be rendered using either isosurface extraction or direct volume rendering. Fig. 3 illustrates this terminology and shows sample ExoVis displays.

ExoVis walls and callouts provide visual cues to help users mentally relate the overview and detail views, as shown in Fig. 3. Each ExoVis wall or callout has an associated "placeholder" that illustrates its position within the overview. Each widget/placeholder pair is assigned a unique colour to distinguish it from other widget/placeholder pairs. Callouts and walls can be connected to their placeholders via connecting lines. Furthermore, widgets and their placeholders are never rotated relative to one another.

In our ExoVis system, walls and callouts can be combined in very flexible ways. Interaction with walls and callouts is also very flexible. Users can reposition:

- Walls and callouts, to change the position or orientation of individual detail areas,
- The camera, to rotate the entire scene, and
- The 3D overview object, to simultaneously change all detail views.

3.1 2D Walls

Lantin and Carpendale [10] and König et al. [7] display information on planes surrounding a 3D object.

ExoVis walls extend this idea to display slices of volume data, as shown in Fig. 1d and Fig. 3 (top). Each wall may be assigned unique display properties (e.g., colour scale). Slices may be obtained in any orientation. Walls are initially oriented along one of the three major axes of the volume; the orientation can then be modified by picking and rotating the wall using the mouse.

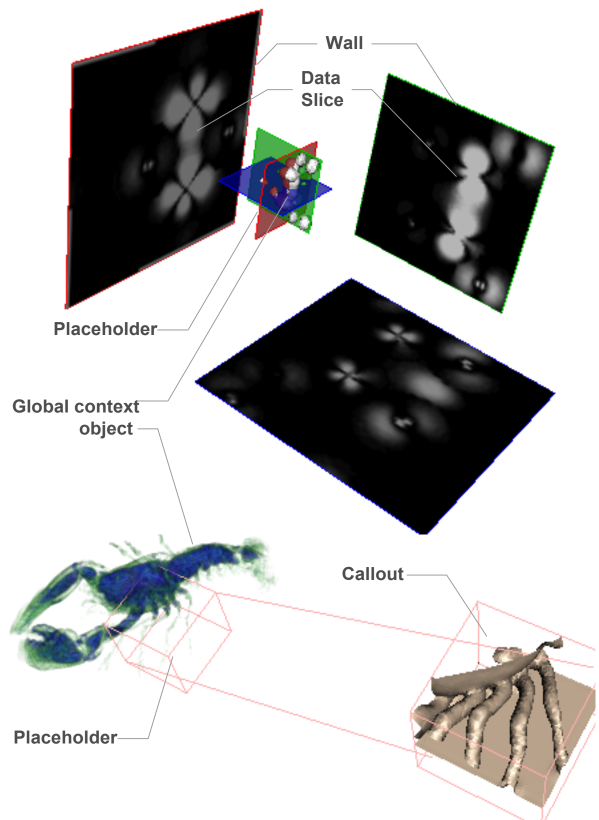

Figure 3: *ExoVis widgets. Walls show 2D slices of a protein data set (top) and a 3D callout shows subvolume details of a lobster's legs (bottom).*

A placeholder illustrates the position and orientation of the data slice within the 3D global context object (i.e., the volume rendered object or isosurface, as defined in Fig. 3). Placeholders are particularly important in the case of arbitrary slice angles, to ensure users can orient walls accurately and understand slice positions. Moving the placeholder along its normal changes the slice displayed on the wall. Rapid movement back and forth in this direction provides an animation that may be considered an overview and detail version of cine mode. (Cine mode is a term used in radiology to describe rapidly scrolling back and forth through a set of 2D images to gain a better sense of the depth dimension.)

Furthermore, grabbing and rotating the global context object automatically updates slices on the walls, producing a rotational version of cine mode. The walls remain in their original positions and orientations while content updates according to the new global context object orientation. Alternatively, rotating the world as a whole changes the camera view but does not change the orientation of objects in the scene relative to each other.

3.2 3D Callouts

3D callouts show volumes of interest (VOIs or subvolumes). VOIs may be rendered using either isosurfacing or direct volume rendering, similar to the global context object. Subvolumes and the global context object need not, however, share the same rendering style or transfer function. Each callout may be assigned unique rendering properties. For example, the lobster in Fig. 3 is direct volume rendered, and the callout shows an isosurface of the leg area.

Interacting with callouts and walls is similar. Users can change callout position, size, and orientation using the mouse. Video clips in the supplementary material and on our website (http://www.cs.sfu.ca/~mktory/personal/volvis) demonstrate these interactions.

3.3 How are ExoVis Walls and Callouts Different from Existing Methods?

ExoVis widgets have features that differentiate them from existing techniques. For example, an ExoVis wall does the reverse of a cutting plane: the overview remains in the centre of the display and slices are shown in the surroundings (compare Fig. 1 a and d).

Orientation icons and ExoVis widgets both separate details from the overview (i.e. the details are "out of place"). However, ExoVis slices and subvolumes remain in their correct orientation relative to the overview object (i.e. they are simply translated from their original position). By contrast, orientation icon slices are translated and rotated from their original location (so the slice is viewed straight-on). This difference is illustrated by Fig. 1 c and d. Similarly, orientation icon subvolumes may be rotated relative to the placeholder, whereas ExoVis subvolumes are not. Restricting slices and subvolumes to their original orientation decreases flexibility, but should make it easier for people to relate the detail and context views since aligning them does not require mental rotation. This is especially useful for non-standard orientations (e.g., when viewing several non-orthogonal slices through a brain tumor), since understanding the spatial relationship between views is more difficult. Keeping the views in their original orientations thus allows easier context switching between views.

Overall, ExoVis walls and callouts are new alternatives that fit somewhere between existing “in-place” and “out of place” techniques. Walls and callouts provide similar flexibility to orientation icons, but the details and context are more closely integrated to make their spatial relationships easier to understand.

4 Comparison of Techniques

Each detail and context technique has a unique set of benefits and drawbacks, so different techniques will be useful for different tasks. For example, “in-place” techniques may be useful when only one variable is visualized (e.g., in a surgery planning application). By contrast, “out of place” methods such as ExoVis walls and callouts and orientation icons may be more valuable for comparing multiple variables (e.g., in a fuel cell data set containing temperature, pressure, and concentrations of several different gases). “Out of place” methods allow multiple copies of each slice or subvolume to be displayed simultaneously without re-displaying the overview, facilitating comparison between different data components or rendering styles (see Fig. 4 and 5).

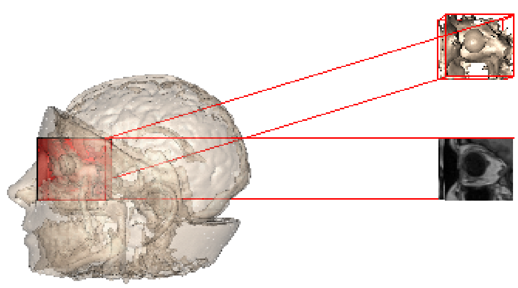

Figure 4: *Two copies of a subvolume show different display properties. An ExoVis VOI is displayed using an isosurface (top) and direct volume rendering (bottom).*

Based on Shneiderman’s general visualization tasks [17], we define the following criteria for comparing and evaluating detail and context methods for volume data:

- Overview and details:
  - **Minimize occlusion** of views by other views or placeholders.
  - **Minimize spatial distortion** to simplify interpretation. (E.g., slices are not shown obliquely and views are not cut apart or otherwise deformed.)
  - **Relationship among views is clear**. It is easy to understand relative positions/orientations of subvolumes or slices. Visual cues include connecting lines, visual landmarks, overlap between views, colour, etc. In other words, there is “visual momentum” [19].

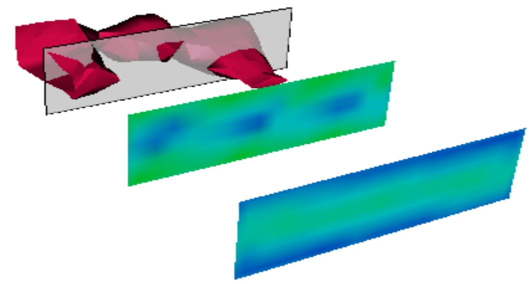

Figure 5: *Visualization of a fuel cell. The overview shows an isosurface of temperature. Two copies of a slice show concentrations of hydrogen (plane closest to overview) and oxygen (furthest from overview). The colour scale shows increasing concentration from green to blue (light to dark in greyscale). Connecting lines and coloured outlines could be added to clarify that the two walls represent the same slice.*

- Zoom and filter:
  - **Maximize ability to selectively magnify** individual views.
  - Make **multiple detail views** showing different spatial areas possible to reduce demands on visual working memory.
- Relate, history, and extract
  - Make **multiple copies** of a detail area possible. This makes it easy to (1) relate different data sets (e.g., several medical imaging modalities, such as CT, MRI, and Ultrasound) or rendering styles (e.g., transfer functions) at the same time and (2) save rendering styles in one copy and create new copies to try new styles (a history function).
- **Minimize screen space**.

Results of our comparison are summarized in Tables 1 and 2 for 2D slices and 3D subvolumes respectively.

5 Discussion

Our comparison helps define when each technique is most useful. Tables 1 and 2 show how ExoVis walls and callouts balance trade-offs of existing techniques. In-place methods have advantages in terms of screen space and ease of relating different views, but this comes at the cost of limited flexibility. Thus, in-place methods are best for simple displays where flexibility is not important (e.g., to compare only a few detail areas in one data set). Orientation icons offer flexibility (e.g., they can display any combination of detail areas as well as multiple data sets and rendering styles). They are also valuable if the task prohibits viewing slices obliquely, or if there are many detail views (since a large number of detail views may be displayed without occlusion).

Table 1: *Comparison of Techniques for Combining a 2D Slice with a 3D Overview* [a]

	Clip/Cutting Plane	***Orientation Icon***	***ExoVis Walls***
Minimize occlusion: 2D slices do not occlude each other or the 3D view.	– Cutting planes are collocated and may occlude each other. Clip planes hide objects in front of them.	+ Each slice is displayed in a separate window, but placeholders partially occlude the 3D view.	– Consecutive slices are stacked and may occlude each other. Placeholders partially occlude the 3D view.
Minimize spatial distortion:			
The 3D view is not cut open, and orthogonal 2D slices do not cut each other.	–	++	++
Slices are not distorted by being viewed obliquely.	–	++	–
Relationship among views is clear:			
Position and orientation of a slice is easy to see.	++ Slices are not moved from their original positions.	– Slices are both translated and rotated from their original positions.	+ Slices are translated but not rotated from their original positions.
When several slices are present, it is easy to determine which slice corresponds to each slice position/orientation in the 3D overview.	++ Discrimination is trivial since all slices are "in-place".	– Specific cues are needed (colour, standard layouts, interactive highlighting, etc.).	+ Can utilize the same cues as orientation icons, plus slices are shown in their original orientations. Connecting lines may also be added.
Selective magnification of slices is possible.	– Slices are collocated.	++	++
Multiple detail views of different spatial areas (slices) are possible.	+ Slices may be occluded by other clip planes. Several cutting planes will produce a "chopped up" image.	++ Since slices are shown "out of place", they will not influence each other and any combination of slices is possible.	++
Multiple copies of the same slice can be shown to support relationship and history tasks.	– Two copies of a plane cannot be shown in the same location.	++ Since slices are shown "out of place", several copies of a single slice may be displayed without creating several copies of the 3D view.	++
Minimize screen space.	++	+	+

a. Legend: ++ (strongly satisfies criterion), + (weakly satisfies criterion), – (does not satisfy criterion).

Table 2: Comparison of Techniques for Combining a 3D Subvolume with a 3D Overview [a]

	In-Place Subvolumes	*Orientation Icon (OI)*	*ExoVis Callouts*
Minimize occlusion: Subvolumes and the overview do not occlude each other.	– Subvolumes cannot be moved and may be occluded by the overview or each other.	++ The overview does not occlude VOIs. ExoVis callouts and overview may occlude each other more than with OI, but subvolumes can be repositioned by the user to minimize occlusion. An automated layout algorithm could also help.	+
Minimize spatial distortion: Magnifying a subvolume does not distort the overview.	– "In-place" subvolumes cannot be selectively magnified without distortion.	++ Since the overview and subvolumes are physically separated, subvolumes can be magnified relative to the overview without spatial distortion.	++
Relationship among views is clear:			
Position and orientation of a subvolume is easy to see.	++ Subvolumes are not moved from their original positions	– Subvolumes can be both translated and rotated from their original positions	+ Subvolumes are translated but not rotated from their original positions
When several subvolumes are present, it is easy to determine which subvolume corresponds to each position/orientation in the 3D overview.	++ Discrimination is trivial since all subvolumes are "in-place".	– Specific cues must be added (colour, interactive highlighting, etc.). Connecting lines are hard to draw.	+ Can use OI cues plus connecting lines (since subvolumes are in the same window as the overview).
Selective magnification of subvolumes is possible.	+ Possible if subvolumes do not overlap.	++ Since subvolumes are "out of place", they can be magnified independently.	++
Multiple detail views of different spatial areas (subvolumes) are possible	++ Overlapping subvolumes merge.	++ Since subvolumes are "out of place", any combination of subvolumes is possible.	++
Multiple copies of the same subvolume can be shown to allow comparison of related data sets or display styles.	– Two copies of a subvolume cannot be in the same location.	++ Since subvolumes are "out of place", several copies of a single subvolume may be displayed without creating several copies of the 3D overview.	++
Minimize screen space.	++	+	+

a. Legend: ++ (strongly satisfies criterion), + (weakly satisfies criterion), – (does not satisfy criterion).

Mentally relating different views with orientation icons can be challenging. ExoVis walls and callouts offer similar flexibility to orientation icons but provide better integration between views. Allowing users to toggle back and forth between ExoVis walls and orientation icon displays provides benefits of ExoVis walls and also allows users to view images straight-on.

With our current ExoVis system, occlusion can be a problem. Manually manipulating objects to reduce occlusion can be time consuming. An automatic layout algorithm that adjusts positions of walls and callouts based on the current view direction could help to resolve this issue. However, such an algorithm should be under user control so users can switch it off if they have particular positioning requirements (e.g., some objects may be more important than others).

6 Conclusion and Future Work

We introduced a new class of display techniques for volumes and compared these methods to existing approaches. Besides being new display types, ExoVis walls and callouts provide a framework for further analysis. ExoVis widgets fall between orientation icons and “in-place” methods, forming a continuum. Studying this continuum should more clearly identify which *characteristics* of the methods are valuable. We are currently planning empirical studies to compare detail and context approaches for several visualization tasks. This should provide additional insight and determine the relative importance of our evaluation criteria. The evaluation criteria themselves may also be valuable for future comparisons. Additional details on our work and video clips demonstrating ExoVis may be found in the supplementary material and on our website: http://www.cs.sfu.ca/~mktory/personal/volvis/.

Acknowledgements

We thank Torsten Möller, Lyn Bartram, Stella Atkins, and Ted Kirkpatrick for their suggestions and acknowledge funding support from the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

[1] M. Boyles and S. Fang. 3DIVE: An Immersive Environment for Interactive Volume Data Exploration. *Int’l Conf. Computer Aided Design and Computer Graphics*, pages 573-580, 2001.

[2] M.S.T. Carpendale, D.J. Cowperthwaite, M. Tigges, A. Fall, and F.D. Fracchia. The tardis: A Visual Exploration Environment for Landscape Dynamics. *Visual Data Exploration and Analysis VI*, SPIE, vol. 3643, 1999.

[3] M.S.T. Carpendale, D.J. Cowperthwaite, and F.D. Fracchia. 3-Dimensional Pliable Surfaces: for the Effective Presentation of Visual Information. *ACM Symp. User Interface Software and Technology*, pages 217-226, 1995.

[4] D.J. Cowperthwaite. *Occlusion Resolution Operators for Three-Dimensional Detail-In-Context*, doctoral dissertation, Computing Science Dept., Simon Fraser Univ., Burnaby, B.C., 2000.

[5] S. Gerlach and R.D. Hersch. A Real-time Navigator for the Visible Human. *IEEE Internet Computing*, pages 27-33, March/April 2002.

[6] T.A. Keahey. Visualization of High-Dimensional Clusters Using Nonlinear Magnification. *Technical Report LA-UR-98-2776*, Los Alamos National Laboratory, 1998.

[7] A. König, H. Doleisch, and E. Gröller. Multiple Views and Magic Mirrors - fMRI Visualization of the Human Brain. Institute of Computer Graphics and Algorithms, Vienna University of Technology, *Technical Report TR-186-2-99-08*, February 1999.

[8] Y. Kurzion and R. Yagel. Interactive Space Deformation with Hardware-Assisted Rendering. *IEEE Computer Graphics and Applications*, vol. 17, no. 5, pages 66-77, 1997.

[9] E. LaMar, B. Hamann, and K.I. Joy. A Magnification Lens for Interactive Volume Visualization. *9th Pacific Conf. on Computer Graphics and Applications*, pages 223 -232, 2001.

[10]M.L. Lantin and M.S.T. Carpendale. Supporting Detail-in-Context for the DNA Representation, H-Curves. *IEEE Visualization*, pages 443-446, 1998.

[11]W. Lorensen and H. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. *Computer Graphics*, vol. 21, no. 3, pages 163-169, 1987.

[12]J.D. Mackinlay, G.G. Robertson, and R. DeLine. Developing calendar visualizers for the information visualizer. *Symp. User Interface Software and Technology*, pages 109 - 118, 1994.

[13]A. Raab and M. Rüger, 3D-Zoom: Interactive Visualisation of Structures and Relations in Complex Graphics. In *3D Image Analysis and Synthesis*, B. Girod, H. Niemann, and H.P. Seidel, eds., pages 125-132, 1996.

[14]M. Sarkar et al. Stretching the Rubber Sheet: a Metaphor for Viewing Large Layouts on Small Screens. *ACM Symp. User Interface Software and Technology*, pages 81-91, 1993.

[15]W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, 2nd ed., Prentice Hall, Upper Saddle River, N.J., 1998.

[16]C.D. Shaw et al. Interactive Lens Visualization Techniques. *IEEE Visualization*, pages 155-160, 1999.

[17]B. Shneiderman. *Designing the User Interface*, 3rd ed., Addison-Wesley Publishers, Reading, Mass., 1997, Chapter 15.

[18]J. Stasko and E. Zhang. Focus + Context Display and Navigation Techniques for Enhancing Radial, Space-Filling Hierarchy Visualizations. *IEEE Symp. Information Visualization*, pages 57-65, 2000.

[19]D.D. Woods. Visual Momentum: a Concept to Improve the Cognitive Coupling of Person and Computer. *Int’l J. Man-Machine Studies*, vol. 21, pages 229-244, 1984.

[20]P.C. Wong and D. Bergeron. Brushing Techniques for Exploring Volume Datasets. In *Proc. IEEE Visualization (Vis)*, pages 429 - 432, 1997.