

Visualization Viewpoints

Editors: Theresa-Marie Rhyne and
Lloyd Treinish

The Transfer Function Bake-Off

Hanspeter Pfister,
Bill Lorensen,
Chandrajit Bajaj,
Gordon Kindlmann,
Will Schroeder,
Lisa Sobierajski
Avila,
Ken Martin,
Raghu Machiraju, and
Jinho Lee

Direct volume rendering is a key technology for visualizing large 3D data sets from scientific or medical applications. Transfer functions are particularly important to the quality of direct volume-rendered images. A transfer function assigns optical properties, such as color and opacity, to original values of the data set being visualized. Unfortunately, finding good transfer functions proves difficult. Pat Hanrahan called it one of the top 10 problems in volume visualization in his inspiring keynote address at the 1992 Symposium on Volume Visualization. And it seems that today, almost a decade later, there are still no good solutions at hand. Or are there?

In a panel discussion at the Visualization 2000 conference, we pitched four of the currently most promising approaches to transfer function design against each other. The four approaches and their advocates are

- trial and error, with minimum computer aid (Will Schroeder);
- data-centric, with no underlying assumed model (Chandrajit Bajaj);
- data-centric, using an underlying data model (Gordon Kindlmann); and
- image-centric, using organized sampling (Hanspeter Pfister).

Ahead of time, each of the four panelists received three volume data sets from Bill Lorensen. The data are static 3D scalar volumes sampled on rectilinear grids. The panelists' task was to create meaningful volume renderings using their respective approaches to transfer


function design. During the panel session, each panelist presented a summary of the method and results of the visualization, including visual results (images and animations), performance (timings and memory use), and observations (how easy or hard it was, what the findings were, and so on). At the end of the panel session, Bill Lorensen discussed the content of the volume data, what an experienced visualization practitioner would have hoped to find, and how well the panelists' methods achieved this goal. Bill also announced a winner.

This was a unique event: alternative approaches to a pressing research problem went head-to-head, on multiple real-world data sets, and with an objective quality metric (Bill Lorensen). The panel took place in an atmosphere of lighthearted fun, but with a serious goal, namely to emphasize the importance of further research in transfer function design. This article presents the four methods in more detail and answers such questions as: How well did they do? Which method works best? And who won the bake-off?


Data sets

Bill Lorensen chose three volume data sets to represent a variety of challenges for the bake-off. Figures 1 through 3 show marching-cube isosurface renderings of the data for comparison to the direct volume rendering images presented by the panelists. The data is available for noncommercial use at <http://visual.nlm.nih.gov/>.


The first data set was generated at GE Aircraft Engines in Evendale, Ohio. The data is industrial x-ray computed tomography (CT) data of a human tooth. The axial



1 The tooth data set modeled with isosurfaces.



2 The sheep heart modeled with two isosurfaces.



3 The segmented MRI knee.

slices are ordered from bottom to top, one slice per file. The pixel samples are spaced 1 mm within each slice, and the slices are 1 mm apart. This data set was the easiest of the three to work with. Figure 1 shows the two materials in the tooth extracted as isosurfaces.

The second data set is magnetic resonance image (MRI) data of a sheep heart generated at the Center for In-Vivo Microscopy, Duke University, North Carolina (<http://www.civm.mc.duke.edu/>). These axial slices are ordered from top to bottom, one slice per file. The pixel samples are spaced 1 mm within each slice, and the slices are 1 mm apart. The heart data is a bit more challenging to visualize than the tooth, because the heart has a variety of tissues including some damaged tissue that was caused by blocking circulation to part of the heart. Figure 2 shows the normal and damaged (yellow) tissue in the sheep heart.

The final data set was generated at Brigham and Women's Hospital Surgical Planning Laboratory (<http://splweb.bwh.harvard.edu:8000/>). The data is clinical MRI data of the knee. These sagittal slices are ordered from left to right. The pixel samples are spaced .25 mm within each slice, and the slices are 1.5 mm apart. This was the most challenging of the three data sets to visualize. Meaningful visualizations of this knee data set are only possible using sophisticated segmentation techniques. Figure 3 shows segmentation performed at Brigham and Women's Hospital, Surgical Planning Lab.

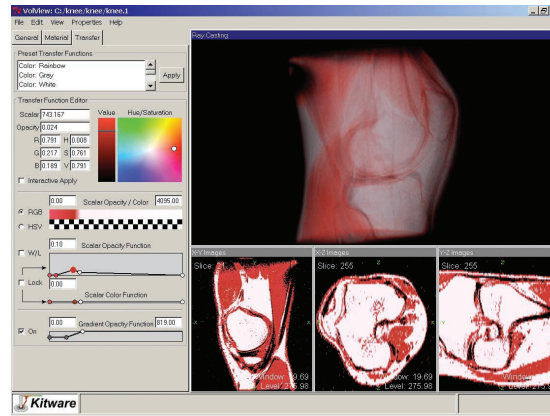
Trial and error

William Schroeder, Lisa Sobierajski Avila, and Ken Martin
Kitware


The widespread use of volume rendering has been hampered by the difficulty of creating effective transfer functions. The complexity of the transfer function is further exacerbated by the blending effects along the depth direction. As a result, recent research has focused on automatic and semiautomatic techniques for creating transfer functions.

Such methods are potentially dangerous because the techniques remove the human from the visualization process. Visualization isn't just about generating pretty pictures. It's also a vehicle of exploration by which the observer comes to understand the data. You can easily imagine semiautomatic and automatic techniques that generate images that fulfill the observer's expectations, but aren't necessarily true to the nature of the data. Thus, we believe that creating a transfer function is a necessary part of the visualization (that is, data exploration) process. Methods that assist the user in creating transfer functions—and thus improve the efficiency of data exploration—are beneficial. Methods that eliminate the human from the exploration process are dangerous and should be avoided.


Figures 4 through 6 demonstrate these ideas. We used Kitware's VolView volume rendering system and the RTViz VolumePro volume-rendering hardware to generate the images quickly. VolView allows interactive, intuitive creation of transfer functions, while the VolumePro board enables maximum interactive response (up to 30 frames per second). For example, we



4 Rendering of the knee data set through trial and error.



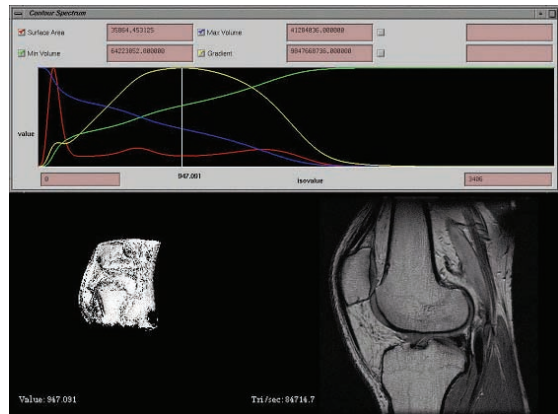
5 Rendering of the tooth data set through trial and error.




6 Rendering of the sheep data set through trial and error.

created both the knee (see Figure 4) and tooth images (see Figure 5) in less than five minutes from start-up to image capture.

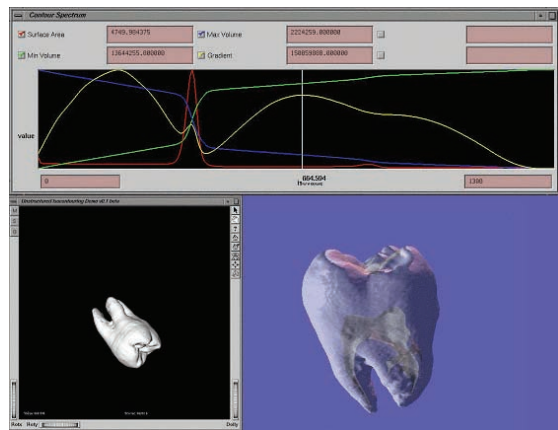
The sheep heart (see Figure 6) was much more challenging to render, requiring approximately 20 minutes to create the final image. However, the time to generate the image was essential: the exploratory process of adjusting the transfer functions taught us much about the data set. For example, after exploring the tooth for five minutes, we felt that we fully understood the important structures within the data. In contrast, the five min-



7 While the maximum of the gradient integral function (yellow signature curve) determined the average scalar intensity for the knee bone, the multiple peaks in the surface area corresponded to the multiple muscular tissues captured by the imaging data set.



8 The isocontour surface automatically selected by the maximum of the gradient integral signature function (left). A volume rendering of the sheep data set. The data set's primary isocontour value indexes the color map range as it centers on purple (right).



9 The gradient integral function again shows four distinctive peaks for the four material types present in the the tooth data set. The left and right renderings were obtained by determining the intensity values for the peaks from the contour spectrum and using that to assign a white color with no transparency (for the left rendering) and purple, pink, and white colors with varied transparency assignments (for the right rendering).

utes spent visualizing the knee taught us that the data was fairly complex, requiring additional segmentation processing for more effective visualization.

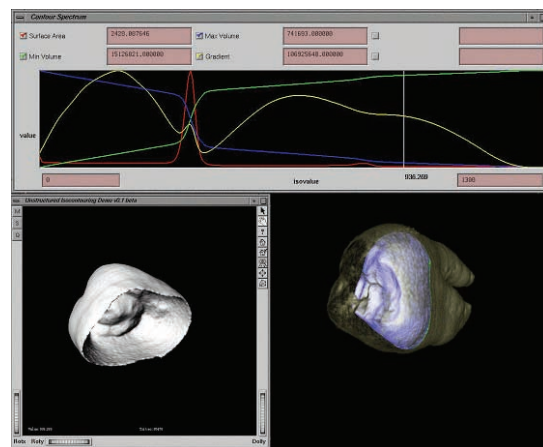
Data-centric, without data model

Chandrajit Bajaj
University of Texas at Austin


In addition to computational and space complexity issues, user interfaces have a tremendous impact on a visualization environment's level of interactivity.

A contour spectrum consists of computed metrics over a scalar field. On the basis of such metrics you can define a set of functions that provide a useful tool to enhance a data set's interactive query. One primary advantage of the contour spectrum interface is that it lets you display—in a 2D image—a global view of the examined scalar field, independent of its dimension. For example, in a 3D isosurface display, one contour component may be hidden inside another. If you associate the isocontour display with the contour tree, it becomes immediately clear that the current isosurface has two components. Hence, you might need a clipping plane to look inside the current isosurface. For time-varying data, we can compute functional properties over time and display it with a 2D interface. This gives users a global overview of the time-varying function and lets them interact with both the isovalue and time step.


All three of the challenge data sets were imaging data sets (static scalar fields over structured rectilinear meshes). The primary characteristic function I used was the gradient integral function curve (shown in yellow in Figures 7 through 10), which automatically separated the various materials in each of the imaging data and generated the appropriate color and opacity map for the final volume rendering. For details of the signature function computations and the contour spectrum, please see the IEEE Visualization Conference 1997 paper¹ or the Web pages where these tools have been applied to various domains <http://www.ticam.utexas.edu/CCV/projects/VisualEyes>.



10 Alternate transfer function selection for the tooth data set highlights the inner tooth surface cap of different material types (and higher density) than the outer surface shown in Figure 9.



11 Renderings of the knee data set using a semiautomatically generated 1D opacity function.



13 Renderings of the heart data set using the transfer functions in Figure 12. The two segments in the opacity function correspond with two boundaries: one between the heart and the background (left) and one for fine structures within the heart tissue (right).


Data-centric, with data model

Gordon Kindlmann
University of Utah


For many medical volume data sets, a good transfer function makes opaque only those values consistently associated with a boundary between materials. The semiautomatic method borrows edge detection concepts from computer vision in order to “locate” boundaries in the 1D space of data value (voxel intensity), since that’s the domain in which transfer functions are specified. The method starts with a preprocess that takes a few minutes and requires minimal user input: creating a 3D histogram of data value versus first and second derivatives and then distilling this into a *distance map* that records the relationship between data value and boundary proximity. Using the distance map, users can interactively experiment with different settings, but the transfer functions are usefully constrained by the boundary information measured in the given data set. (For more details, see the 1998 IEEE Symposium on Volume Visualization paper² or visit <http://www.cs.utah.edu/~gk/MS>.)

Of course, this method has trouble on data sets in which there are noise and coarse boundary sampling, such as the knee MRI scan. As Figure 11 shows, the method detected the boundary between air and skin and rendered it clearly, but the boundaries among the various internal tissues are less clear.


One benefit of the semiautomatic method is the abil-



12 Making a 2D transfer function (from left to right): an automatically generated distance map (white indicates the boundary center), a semiautomatic opacity function, and a manually created color map.




14 Automatically created 2D distance map for the tooth data set. In the cross-section, the colors of the small rectangles (marking the different boundaries) indicate the corresponding boundary colors in the transfer function.

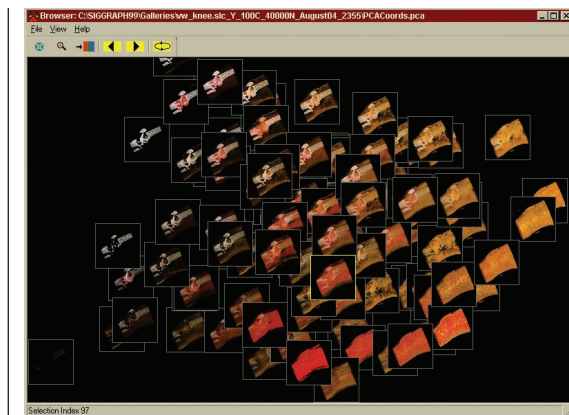


ity to create transfer functions that map not just data value, but also a 2D space of data value and gradient magnitude. Although they can often disambiguate complex material boundaries, 2D transfer functions are even more difficult to specify by hand than traditional 1D functions. The sheep heart MRI data set is a good example. Figure 12 shows the 2D distance map and transfer function used to make the renderings in Figure 13.

You might think that for a clean data set such as the tooth CT scan, transfer functions based on data value alone could accurately convey all the boundaries. However, the tooth cross-section in Figure 14 shows the data set to have four boundaries between four materials (from dark to bright: pulp, background, dentine, and enamel). The presence of four arcs in the distance map



15 Renderings of the tooth data set. Opacity functions are inset in each image; surface colors are determined by the color map in Figure 14.



16 VoIDG user interface.

(Figure 15) show that the semiautomatic method successfully discerned all the boundaries and thereby facilitated the renderings in Figure 15, which emphasize and color the boundaries in sequence. In particular, with an isosurface or 1D transfer function, it would have been impossible to isolate the dentine-enamel boundary shown in cyan.

The current version of the semiautomatic method assumes a specific mathematical boundary model; this may reduce its effectiveness on some data sets. However, the method ultimately derives its utility from combining a common volume-rendering task (“show me the boundaries”) with a characterization of boundaries in terms of easily measured derivatives. A tool such as Design Galleries can answer What’s possible? in the

space of all transfer functions, while this approach intends to answer What’s probable?—that is, What’s most likely to be a good transfer function, assuming the goal of visualizing boundaries? Unconstrained exploration of transfer functions is sometimes needed, but interactivity in a visualization tool proves more valuable when the interface itself embodies information and constraints derived from the data in question.

Image-centric, using organized sampling

Hanspeter Pfister

Mitsubishi Electric Research Laboratories

Raghu Machiraju and Jinho Lee

The Ohio State University

Along the lines of the Design Gallery approach,³ my colleagues and I developed VoIDG—Design Galleries for Volume Graphics as a viable alternative to facilitate transfer function selection. The image-centric transfer function design of VoIDG focuses on what matters most to the user: the image. VoIDG evaluates transfer functions on the basis of the images they produce, not in terms of data set properties. Instead of asking the computer What’s best? we ask the computer What’s possible? The computer picks a set of input-parameter vectors that span the space of output values as much as possible; the user simply selects from among the presented possibilities.

As Figure 16 shows, VoIDG interfaces present the user with the broadest selection—automatically generated and organized—of perceptually different images that can be produced by varying transfer functions.


The VoIDG approach’s principal technical challenges are dispersion (finding a set of input-parameter vectors that optimally generates dissimilar output values) and arrangement (arranging the resulting designs for easy browsing). For dispersion, we use a form of evolutionary computation. For arrangement, we use multidimensional scaling. The dispersion process can require rendering hundreds or thousands of candidate images and therefore benefits greatly from hardware acceleration by Mitsubishi’s VolumePro board. In addition, expedient rendering aids the interrogative process between user and computer.

We built our current system on top of the popular Visualization Toolkit (vtk) to using the VolumePro board. The real-time volume-rendering speed of the VolumePro board lets large galleries be generated in minutes. VoIDG is freely available at <http://www.merl.com/projects/dg/>.


Even though VoIDG can manipulate both color and opacity transfer functions, we believe that generating gray-scale images leads to the most insight into unknown data sets. Figure 17 shows two representative images of the knee and sheep heart data set. Note that VoIDG automatically detected the interior structure in both images. This is remarkable because MRI data is notoriously difficult to deal with.

Figure 18 shows the transfer function for one MRI knee image. Note that the function is piecewise linear. More sophisticated curve parameterizations, such as B-splines or wavelets, could improve the results of VoIDG.


Not surprisingly, the tooth CT scan shows better interior structure and detail (see Figure 19). Each gallery



17 VoIDG-generated images of the knee (top) and sheep heart (bottom).



18 Transfer function generated by VoIDG.



19 VoIDG-generated images of the tooth.

with 200 images and 20,000 iterations took approximately 7 hours to generate. Absolutely no user intervention was necessary. A smaller gallery with 50 images and 1,000 iterations only takes about 3 to 4 minutes. Plus, users can interactively browse the gallery.

Discussion

Bill Lorensen

GE Corporate Research and Development

Since its introduction in the late 1980s, volume visualization has had limited success in cost-effective applications. Advances in image quality and feature sets continue to outpace the technology's acceptance in commercial products.

Several factors contribute to the slow adoption of volume visualization:

- A lack of proven application areas. Routine use of 3D in medicine is still, for the most part, limited to research and teaching hospitals.
- There's no agreement on software application programming interfaces (APIs) for volume visualization. This limitation translates to risk for commercial products that adopt one vendor's API over another's.
- Volume visualization is slow, requiring expensive

workstations with large amounts of memory and special graphics hardware extensions.

- The volume techniques are difficult to use by all but an experienced engineer or scientist.

Fortunately, various universities and companies are addressing these limitations:

- New scanners are presenting much more information than a radiologist can possibly review on a slice-by-slice basis. Three-dimensional visualization could be the key to increasing productivity.
- APIs are emerging that fit within current graphics and visualization systems.
- Low-cost, special-purpose hardware is now available for personal computers. And, the general-purpose processor speeds continue to improve. Texture-mapping hardware is available on cheap graphics cards. In addition, 3D texture mapping could also get cheap if the gamers find a use for it.

However, ease of use is still an issue. Volume visualization has the potential to significantly reduce the amount of time to segment medical data. We need fast, robust techniques to create color and opacity transfer

functions before volume rendering can move from the lab to the hospital. It doesn't matter whether the techniques are automatic, semiautomatic, or manual. They just need to be fast and simple.

Summary and conclusions

Bill Lorensen

GE Corporate Research and Development

The four groups on the panel weren't given a specific task to perform. Their sole goal was to produce images that would impress the panel judge (me). All the groups gave a short presentation of their results. Each team performed admirably, presenting a variety of renderings for all the data sets.

The Kitware panelist produced "artistic" renderings of each data set by manually choosing transfer functions. The MERL approach presented dozens of alternatives for each data set, requiring the user to choose an appropriate rendering. The Texas algorithm automatically created transfer functions based on metrics derived from the data. The Utah panelist also presented an automatic technique that followed a more traditional feature-extraction approach, designed to find boundaries in the data.

As the judge, I was biased against techniques that required too much or too little human interaction. This bias eliminated the manual Kitware approach and the automatic MERL technique. I had a difficult time deciding the winner between the two remaining datacentric approaches. The Texas reliance on observable metrics in the data seems to be more intuitive than the Utah approach. However, in my opinion, the Utah algorithm shows the most promise and is most likely to stimulate future research in the area of automatic transfer function synthesis.

Rectilinear, static grids are the simplest volumetric data, yet it's obvious that many problems exist in devel-

oping effective renderings. Other types of volumetric data are even more challenging, such as time-varying data, time-varying grids, irregular grids, scattered data, or nonscalar fields. We hope that this panel encourages further research in transfer function design, particularly for more complex, difficult-to-visualize volumetric data sets. ■

Acknowledgments

We'd like to thank the National Library of Medicine, National Institutes of Health, and Visible Human Project for hosting the data distribution Web site (<http://visual.nlm.nih.gov/>).

Also, we'd like to acknowledge that Vamsidhar Juvvignunta, Joe Marks, and Kathy Ryall were the main collaborators on VolDG.

References

1. C. Bajaj, V. Pascucci, and D. Schikore, "The Contour Spectrum," *Proc. 1997 IEEE Visualization Conf.*, IEEE CS Press, Los Alamitos, Calif., Oct. 1997, pp. 167-173.
2. G. Kindlman and J. Durkin, "Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering," *Proc. 1998 IEEE Symp. Volume Visualization*, IEEE CS Press, Los Alamitos, Calif., Oct. 1998, pp. 79-86.
3. J. Marks et al., "Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation," *Proc. Siggraph 97*, ACM Press, New York, Aug. 1997, pp. 389-400.

Readers may contact Pfister at MERL, 201 Broadway, Cambridge, MA 02139, email pfister@merl.com.

Contact department editors Rhyne and Treinish by email at rhyne@siggraph.org and lloyd@us.ibm.com.