Research issues in . . .

The seven short articles in this special section examine research issues, both ongoing (such as those discussed in "Recent advances" above) and emerging topics. Examples of the latter include volume graphics, multiresolution modeling, visualizing tensor fields, virtual reality interfaces for visualization, automating visualization designs and processing, validation tools, perceptual issues in visualization, and the relation between underlying mathematical models and the visualization process. The increased use of sophisticated mathematics in a trend seen in several of these articles. In most cases proposed research issues are clear-cut, but occasionally they are controversial. This is good, for the resulting discussions will contribute to a clearer vision of future directions.

As Fred Brooks noted in his Visualization 93 keynote address, scientific visualization is not yet a discipline, although it is emerging as one. Too often we still have a collection of ad-hoc techniques and rules of thumb. Perhaps by stepping back and taking a look at where we are going, these articles will assist the field's growth. ❑

Acknowledgments

As workshop chair, I gratefully acknowledge the help of the organizing committee (Rae Earnshaw, José Encarnação, Hans Hagen, Arie Kaufman, Stanislav Klimenko, Wolfgang Krüger, Gregory Nielson, and Frits Post). The workshop was held at the Fraunhofer Institute for Computer Graphics in Darmstadt, Germany. I thank its director, José Encarnação, and Martin Göbel, who handled local arrangements, for the impressive facilities and superb organization. We owe a debt of gratitude to all participants for their willingness to devote so much time and effort to presenting this overview of the research status of visualization today.

Special thanks are due to Andre van Tilborg (ONR Division Director for Mathematical, Computer, and Information Sciences) and Arthur Diness and John Silva (previous and current Technical Directors for ONR Europe) for their sponsorship and active encouragement of this project "from soup to nuts." The Liaison Scientist position has been an immensely enjoyable experience, and I appreciate ONR for providing the opportunity. The assistance and openness of numerous European colleagues, both with the workshop and in my daily tasks, is gratefully acknowledged.

References

1. "Visualization in Scientific Computing," B.H. McCormick, T.A. DeFanti, and M.D. Brown, eds., *Computer Graphics*, Vol. 21, No. 6, Nov. 1987.
2. "Special Issue on Visualization in Scientific Computing," G.M. Nielson, ed., *Computer*, Vol. 22, No. 8., Aug. 1989, pp. 10-101.
3. *Visualization in Scientific Computing*, G.M. Nielson, B. Shriver, and L.J. Rosenblum, eds., IEEE Computer Society Press, Los Alamitos, Calif., Aug. 1990.
4. A. Kaufman, D. Cohen, and R. Yagel, "Volume Graphics," *Computer*, Vol. 26, No. 7, July 1993, pp. 51-64.
5. F.H. Post and J.J. van Wijk, "Visual Representation of Vector Fields," in *Frontiers of Scientific Visualization*, L.J. Rosenblum et al., eds., Academic Press, London (to appear in 1994).
6. J.L. Helman and L. Hesselink, "Visualization of Vector Field Topology," *IEEE CG&A*, Vol. 11, No. 3, May 1991, pp. 36-46.
7. *Frontiers of Scientific Visualization*, L.J. Rosenblum et al., eds., Academic Press, London (to appear in 1994).

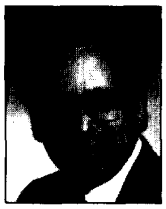

Lawrence J. Rosenblum is currently serving as Liaison Scientist for Computer Science at the US Office of Naval Research European Office in London while on leave from the Information Technology Division of the Naval Research Laboratory in Washington, DC. His research interests involve the application of scientific visualization, machine vision, and virtual reality to scientific data sets. He received his PhD from Ohio State University in mathematics.

Rosenblum is on the editorial board of *IEEE CG&A*, where he created and edits the Visualization Blackboard department. He is a director of the IEEE Computer Society's (IEEE-CS) Technical Committee on Computer Graphics, a former IEEE-CS vice chair for Technical Activities, and a cofounder and steering committee member of the IEEE Visualization conference series. He is a member of the IEEE Computer Society, ACM, Siggraph, the American Geophysical Union, and Sigma Xi.

Readers may contact Rosenblum starting in April at Information Technology Division, Code 5500, Naval Research Lab, Washington, DC 20375.

Research Issues in Volume Visualization

Arie Kaufman, *SUNY at Stony Brook;* Karl Heinz Höhne, *University Hospital Eppendorf, Hamburg;* Wolfgang Krüger, *German National Research Center for Computer Science;* Lawrence Rosenblum, *ONR European Office;* Peter Schröder, *Princeton University*

Volume visualization is a method of extracting meaningful information from volumetric data sets through the use of interactive graphics and imaging. It addresses the representation, manipulation, and rendering of volumetric data sets, providing mechanisms for peering into structures and understanding their complexity and dynamics. Typically, the data set is represented as a 3D regular grid of volume elements (voxels) and stored in a volume buffer (also called a cubic frame buffer), which is a large 3D array of voxels. However, data is often defined at scattered or irregular locations that require using alternative representations and rendering algorithms.

The ONR Workshop on Data Visualization identified eight major research issues in volume visualization.

Volume graphics

Volume graphics is an emerging subfield of computer graphics concerned with the synthesis, manipulation, and rendering of 3D modeled objects, stored as a volume buffer of voxels.[1] Unlike volume visualization, which focuses on sampled and computed data sets, volume graphics primarily addresses modeled geometric scenes, particularly those represented in a regular volume buffer.

Volume graphics has advantages over surface graphics. It is viewpoint independent, insensitive to scene and object complexities, and suitable for representing sampled and simulated data sets and mixtures thereof with geometric objects (see Figure 1). It supports the visualization of internal structures and

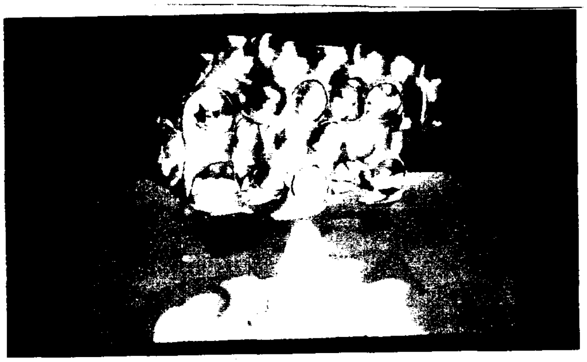

Figure 1. A simulated silicon crystal reflected in a mirror. Data represents charge densities obtained from simulation results on a grid of 32 × 32 × 96. Transfer functions for color and opacity and index of refraction of 1.05 were used to create a translucent isosurface view. The mirror floor is a volume-sampled voxelized polygon. (Data provided by Victor Milman, Oak Ridge National Lab, Oak Ridge, Tenn. Image generated using the VolVis system, State University of New York at Stony Brook.)

Figure 2. Volume rendering via transport theory of a supercomputer solution of the Scrödinger equation for a protein molecule. (Data provided by L. Noodleman and D. Green, Scripps Clinic, La Jolla, Calif.)

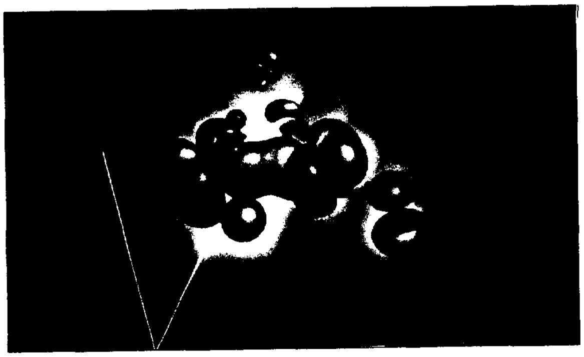

lends itself to block operations, CSG modeling, and hierarchical multiresolution representations. The problems associated with the volume buffer representation—such as discreteness, memory size, processing time, and loss of geometric representation—echo problems encountered when raster graphics emerged as an alternative technology to vector graphics and can be alleviated in similar ways. By offering a comprehensive alternative to traditional surface graphics, volume graphics has the potential to develop into a major trend in computer graphics.

Further research in the field should address modeling and rendering geometric scenes by employing a discrete volumetric representation. The underlying modeling procedure relates to voxelization algorithms that synthesize voxel-based models by converting continuous geometric objects into their discrete voxel-based representation. The voxelization process (also called 3D scan conversion) samples and filters the voxelized objects for conversion. The process must conform to 3D discrete topological considerations to mimic the continuous topological behavior of the continuous objects. It must also consider such issues such as geometric accuracy, minimality, efficiency, and representation and rendering quality (for example, 3D antialiasing).

Other research directions include developing techniques for modeling and sculpturing in discrete space, building CSG models, feature mapping, warping, morphing and changing of the model, and intermixing geometric objects with sampled or simulated data. In another report in this issue, Nielson et al. discuss volume modeling and the possibility of generating renderings without necessarily going through the voxelization step (see pp. 70–73). Rendering issues specifically related to modeled data include global and local illumination parameters, volume and surface features, specification of transfer functions, suppression of artifacts, and the like.

Volume rendering equation

Volume rendering is the direct mapping of the essential content of volumetric scalar data fields onto an intensity field that can be displayed on the screen. The basic model is an adapted transport theory model[2] describing the process of fictitious (light) particles moving through the data field. These particles interact with the data values and preselected features (for example, isosurfaces) and collect information about the light intensity I for screen display.

The general transport theory model for volume rendering is given by the integro-differential equation:

$$(\mathbf{s} \cdot \nabla)\, I(\mathbf{x},\mathbf{s}) = -\sigma_t(\mathbf{x})\, I(\mathbf{x},\mathbf{s}) + q(\mathbf{x}) + \sigma_s(\mathbf{x}) \int ds'\, p(\mathbf{x}, \mathbf{s}' \rightarrow \mathbf{s})\, I(\mathbf{x},\mathbf{s}') \quad (1)$$

This equation describes the gains and losses of the particle beam inside a volume element due to physically based terms such as emission, absorption, and scattering. The extinction coefficient $\sigma_t = \sigma_a + \sigma_s$ accounts for the intensity's attenuation due to absorption and scattering towards other directions. The source term q (for example, point-like, surface-like, or volumetric) intensifies the passing light beam. The last term in the equation describes the amount and angle distribution of the scattering by the scattering cross section σ_s and the scattering phase function p.

The formal solution to this adapted transport equation is given by a Fredholm integral equation of the second kind:

$$I(\mathbf{x},\mathbf{s}) = I_{in}\, e^{-\int_0^R \sigma_t(\mathbf{x}_{in}+R'\mathbf{s})dR'} + \int_0^R dR'\, e^{-\int_{R'}^R \sigma_t(\mathbf{x}_{in}+R''\mathbf{s})dR''}\, Q(\mathbf{x}_{in} + R'\mathbf{s}, \mathbf{s}) \quad (2)$$

where the generalized source Q is given by

$$Q(\mathbf{x},\mathbf{s}) = q(\mathbf{x}) + \sigma_s(\mathbf{x}) \int ds'\, p(\mathbf{x}, \mathbf{s}' \rightarrow \mathbf{s})\, I(\mathbf{x},\mathbf{s}')$$

A Neumann series can evaluate this formal integral solution to account for multiple scattering events. Current volume rendering methods use only a first-order approximation. This approach neglects the scattering term and evaluates the intensity I for each screen point (x, y) for only one viewing direction $\mathbf{s}_0$ for each red, green, blue (RGB) value. We can evaluate the path integral in Equation 2 by discretizing it into an Eulerian sum from front to back, giving an iteration rule

$$I(s+\Delta s) = (1-\alpha_v)\,[I(s) + q(s)\,\Delta s] \quad (3)$$

with opacity

$$\alpha_v = 1 - e^{-\int_{s}^{s+\Delta s} \sigma_t(s')ds'}$$

which a recursion algorithm can easily compute.

Physical considerations underpin the transport equation here, helping it play to the user's physical intuition. However, there is nothing particularly physical about viewing volumetric scalar fields, which do not correspond to the light emission of physical objects. Consequently, in visualization we are usually interested not so much in "physical realism" as in well-defined mappings of the equation parameters onto meaningful visual quantities. These mappings facilitate the reproducible extraction of important data-field features.

Future research should cover topics related to the volume rendering equation, including the following:

- Exploring all degrees of freedom for mapping relevant data structures to the physical parameters of the transport equation. In addition to the currently used simple emission-absorption approaches, surface and volume scattering terms and volume texturing and color shift terms should be incorporated in Equations 1 and 2. For an example, see Figure 2.
- Enhancing the coarse and rigid approximation methods for evaluating Equation 2's path integral, as given by Equation 3, to support methods adapted to the local amount of data detail. The interpolation of the data value and its derivates for the evaluation point from the values of the surrounding points should be extended beyond trilinear interpolation to suppress artifacts.
- Generalizing the transport equation for volume rendering to incorporate a term that generates each pixel color depending on an interesting data feature. Equation 2 can also be transferred to a volume integral equation describing the interaction of the data values in all volume elements with each other. This would support complex transparency and shading models and the use of shadows for depth enhancement as in ray-tracing and radiosity approaches.

The main advantage of this general research approach is the introduction of a well-defined mathematical framework (linear operator theory) for which evaluation algorithms, such as Monte Carlo simulation and series expansions, exist. Such a framework supports the development of advanced evaluation algorithms for interactive, real-time rendering, for example, on massively parallel machines.[3]

Transform coding of volume data

Currently, most volume rendering algorithms work directly on the scalar data field to be visualized. Various acceleration techniques take advantage of local coherence by representing the data in a pyramidal fashion or in octrees. This allows large homogeneous regions to be modeled, traversed, and rendered faster without compromising the quality of the final rendering. Recent compression techniques, such as vector quantization, reduce the storage overhead for volume data and render directly from the compressed data set.

We expect further developments in this area—in particular the use of transform coding, such as representing the data set in a wavelet basis. Wavelets allow efficient multiresolution data representation. Decomposition of the data into a wavelet basis localizes in both frequency and spatial domains. This can reduce the storage overhead when used for a lossy (and possibly with a lossless) compression scheme. It facilitates the analysis of certain features in the data set and the use of asymptotically faster rendering algorithms. A representation in wavelet basis, for example, facilitates the detection of discontinuities, such as edges, in the original data. It has also been used to obtain a continuous multiresolution shape description of volumetric objects.[4] It accelerates rendering because the rendering process evaluates an integral operator defined on the volume.

The integral operator used in rendering has a sparse representation when a finite but arbitrary precision is required. Recent advances in mathematics and numerical analysis support exact error bounds when using fast wavelet approximations. Also, since wavelet decomposition gives an intermediate representation between the frequency and spatial domains, it might be possible to capitalize on the advantages of both to gain speed and quality.

Other coding bases and spaces, such as frequency space, also need further exploration. Using a frequency coding space and the Fourier Projection-Slice Theorem, we can render images directly out of the compressed domain.[5] Especially from a systems perspective, this approach is very promising, since it avoids the overhead of transforming back and forth to the compression domain. Further research should focus on the trade-off between image quality, compression ratio, block size, speed, and system architecture issues.

Scattered data

Research on visualizing scattered data is still in its infancy. Most existing volume rendering algorithms are for data sets of Cartesian (rectilinear) grids with at most nonuniform spacing. Ray-casting methods for curvilinear grids have also been developed. Some of these methods are based on decomposing the cells into tetrahedra; others are based on some special property that the grid may have for certain applications. For example, sorting and interpolation algorithms can be applied to curvilinear grids (as in climate modeling, where the grid is Cartesian in longitude and latitude but with varying altitude spacing). Alternatively, the curvilinear grid data can be interpolated into a Cartesian grid by using some method of scattered data modeling. In principal, this model-based rendering approach applies to data more general than curvilinear data, but difficulties in choosing the proper modeling method require more research before this general approach leads to tools for everyday use.

Scene reconstruction from scattered data is another area where volumetric methods are replacing surface-based techniques (see Nielson et al. in this issue, pp. 70–73). Volumetric algorithms combined with worst-case or probabilistic analysis offer an alternative to classical methods, with many advantages. For example, dense range data from multiple viewpoints taken from optical or laser sensors can create and refine a 3D voxel-based volumetric model. A similar probabilistic approach volumetrically reconstructs 3D scenes in the ocean, where acoustic data from sonars produces far lower resolutions.

Enriching volumes with knowledge

Volume visualization is steadily improving in both rendering speed and image quality. However, current models do not include the semantics of pictures, thus leaving image interpretation generally in the viewer's hands. We can improve image interpretation by enriching the volume with knowledge and by registering, classifying, and segmenting the data.

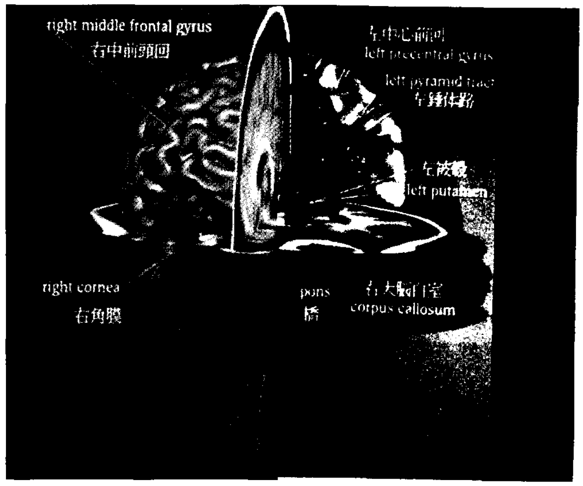

Figure 3. An interactive anatomical atlas example of "intelligent" volume visualization. The picture shows a view of the brain, which has been "sculptured" by the user. Users can get information (even in different languages) from every visible location of the picture. They can also generate pictures by selecting items from the knowledge base.

An association of symbolic and spatial knowledge, however, would benefit many applications, especially in education. An example of such an approach is the "generalized voxel model," which has been extended into the "intelligent volume" approach.[6] Figure 3 shows the application of these approaches in the implementation of an anatomical atlas. A semantic network models the names, relationships, and other symbolic properties of the anatomical objects, while a labeled tomographic volume represents their spatial properties.

Still, many related issues require more detailed research: How can we represent a rather inconsistent (medical) terminology and knowledge? (Modifying it to fit into formal structures is clearly not acceptable.) How can we incorporate different "views" to the same structures (for example, histological and functional)? How we visualize uncertainty? What are optimal data structures for spatial knowledge? What must "visual query languages" look like? To what extent do artificial intelligence and/or database approaches apply?

Segmentation

Measured volume data are typically represented as raw data (for example, densities) with no prior definition of the objects to be visualized. A classification/segmentation step that identifies the regions describing an object must therefore accompany the visualization. Segmentation, a domain related to image processing and pattern recognition, is a field worked on by many researchers over a long time. Nevertheless, automatic segmentation systems have emerged only for very specific tasks and data. At present, there are no methods in sight that could handle, for example, arbitrary medical scenes.

Currently, interactive segmentation is more promising. The user steers the process by viewing the 3D rendered results while changing the segmentation parameters. This combines the unsurpassed capability of humans to recognize patterns and the capability of computers to process very large data sets in near real time.[7] We nevertheless believe that reliable automatic segmentation is a desirable goal. Progress depends heavily on close cooperation with such fields as image processing, computer vision, and computer-aided design.

Real-time rendering and parallelism

Interactivity in scientific visualization is a key to the exploration of data sets. For volume rendering, this implies large computational resources. Currently, these resources are available for reduced resolution primarily on parallel hardware, including multiprocessor workstations.

Since communication costs are high, and will remain so relative to processor speeds, algorithms must be carefully designed to take advantage of the available floating-point resources. Initial work in this area has concentrated on regular volume data and has yielded algorithms with multiple-frames-per-second update rates for low resolution supporting very general rendering options. Additional work is necessary to better map the algorithm to the machine architecture.

Another challenge is the design of rendering algorithms that deliver these speeds for scattered data such as is produced in large finite element computations. Insights gained from implementing such algorithms on parallel platforms will also stimulate the development of custom parallel hardware for volume rendering, which will eventually find its way into workstations.

Real-time visualization, interactive exploration of data content, and steering the data generation and evaluation processes are among the major goals in the development of volume rendering tools. For example, results from large-scale simulations on the new generation of massively parallel computers, or large data sets arriving from measurements in space missions or medical imaging devices, will all benefit from incorporating visualization and data analysis into the data evaluation task itself.

Special-purpose hardware

Both graphics accelerators and special-purpose machines have been used for volume rendering, with reasonable performance gain over software implementations. An obvious way to accelerate volume rendering is to build special-purpose hardware tailored to this task. Several attempts have been launched, but none has gained market acceptance.

Initial attempts employed an octree data structure suitable for relatively uniform and regular objects but ineffective when the application calls for complex objects or many colors or densities, as in sampled data. Consequently, other architectures have stored the volumetric data as a 3D array, while gaining speedup by employing diverse parallel mechanisms (for a survey, see Kaufman[8]). However, Cube[9] uses a unique skewed memory organization that allows parallel read/write of axial beams, while other architectures rely on a modular memory that stores a slice of voxels in each module.

The choice of rendering algorithm is a major decision in hardware design. Many architectures use forward projection, and many use backward projection. Cube, for example, employs backward projection by casting rays parallel to the main axes. An extended architecture of Cube, Cube-3,[10] delivers high-quality 512^3 resolution ray-casting rendering of perspective and arbitrary parallel viewing in real-time. By the turn of the century, we expect add-on accelerators for volume rendering and/or special-purpose machines to be as common as polygon-based accelerators.

True 3D displays, also called direct volume display devices, have been researched for many years.[11] Examples of such devices include computer-generated holography, which has a relatively small angle of view (say, 30 degrees). Another type of device is a varifocal mirror, such as the BBN SpaceGraph, which projects successive 2D slices onto the varifocal mirror. It can display 32,000 points per update of a 26-cm^3 volume at 30 Hz with about a 60 degree angle of view. A third type of device illuminates a 2D rotating screen by the corresponding slice of a cylindrical world. The Texas Instrument OmniView[11] is an advanced device that uses laser beams to illuminate an axially symmetric double-helix display surface. The fourth-generation device is $5 \times 3 \times 3$ feet with a display volume 20 inches in diameter and 10 inches high. Volumetric Imaging's Matrix Imager rotates an active matrix of light-emitting diodes, but has a low voxel addressability ($64 \times 64 \times 48$).

All these devices have proven themselves in prototype form in the laboratory. Once available commercially, they will revolutionize the way we view and interact with volumetric data. However, to be commercially viable, these devices need much higher resolution, more colors, and lower prices.

Conclusions

Volume visualization started out as a research challenge requiring large amounts of time on relatively small data sets. In the past few years it has become a common tool in research domains as well as applications. The rapid development of the underlying hardware has contributed greatly to the wide availability of volume rendering techniques. Algorithmic research has accelerated the manipulation and rendering of volumetric data as well. We are now on the verge of interactive volume rendering on custom workstations for moderately sized data sets. However, the challenges to the research community continue.

There will always be data sets larger than what a given hardware platform can render in real time. Consequently, custom hardware for volume rendering will remain important, especially as the use of volume graphics increases. More advanced algorithms are needed to provide the full palette of rendering options, including global illumination, for volume graphics. Very large data sets force us to deal with compressed or hierarchical data representations.

Once large data sets from many sources are on line, interactively accessible tools to help extract meaning from them become ever more important. We expect major contributions to this area from more rigorous numerical algorithms, interactive segmentation and classification tools, and improved techniques from the database and artificial intelligence communities. The latter two are especially important as we begin to deal with data semantics.

With the diversity of sources of data comes a diversity of data types. Few algorithms efficiently visualize data sets that are not of the regular Cartesian grid type. Development in this area is particularly urgent, not least because advances in numerical techniques, such as adaptive gridding, will likely decrease the use of regular Cartesian gridded data. ❑

Acknowledgments

Special thanks are due Greg Nielson and Roni Yagel for reading and commenting on this article. The work at State University of New York at Stony Brook has been supported in part by the National Science Foundation under grant CCR-9205047 and by a grant from the Department of Energy.

References

1. A. Kaufman, D. Cohen, and R. Yagel, "Volume Graphics," *Computer*, Vol. 26, No. 7, July 1993, pp. 51-64.
2. W. Krüger, "The Application of Transport Theory to Visualization of 3D Scalar Data Fields," *Computers in Physics*, July/Aug. 1991, pp. 397-406.
3. W. Krüger and P. Schröder, "Data Parallel Volume Rendering Algorithms for Interactive Visualization," *The Visual Computer*, Vol. 9, Aug. 1993, pp. 405-416.
4. S. Muraki, "Volume Data and Wavelet Transform," *IEEE CG&A*, Vol. 13, No. 4, July 1993, pp. 50-56.
5. T. Malzbender, "Fourier Volume Rendering," *ACM Trans. Graphics*, Vol. 12, No. 3, July 1993, pp. 233-250.
6. K.H. Höhne et al., "A 3D Anatomical Atlas Based on a Volume Model," IEEE CG&A, Vol. 12, No. 4, July 1992, pp. 72-78.
7. K.H. Höhne and W.A. Hanson, "Interactive 3D Segmentation of MRI and CT Volumes Using Morphological Operations," *J. Computer-Assisted Tomography*, Vol. 16, No. 2, 1992, pp. 285-294.
8. A. Kaufman, *Volume Visualization*, IEEE Computer Society Press, Los Alamitos, Calif., 1991.
9. A. Kaufman, and R. Bakalash, "Memory and Processing Architecture for 3D Voxel-Based Imagery," IEEE CG&A, Vol. 8, No. 6, Nov. 1988, pp. 10-23.
10. H. Pfister and A. Kaufman, "Real-Time Architecture for High-Resolution Volume Visualization," Proc. 8th Eurographics Workshop on Graphics Hardware, Eurographics, Geneva, 1993, pp. 72-80.
11. T.E. Clifton III and F.L. Wefer, "Direct Volume Display Devices," *IEEE CG&A*, Vol. 13, No. 4, July 1993, pp. 57-65.

Research Issues in Perception and User Interfaces

José Encarnação, ***Fraunhofer Institute for Computer Graphics***; **Jim Foley**, ***Georgia Institute of Technology***; **Steve Bryson**, ***NASA Ames Research Center***; **Steven K. Feiner**, ***Columbia University***; **Nahum Gershon**, ***Mitre Corporation***

Current visualization tools are capable but still require too much visualization knowledge on the user's part. This requirement restricts the user in what is possible. Nor do the tools take account of what is known regarding cognition and perception. The User Interface Working Group at ONR's Data Visualization Workshop focused on three things: presentation of information to best match human cognitive and perceptual capabilities, interactive tools and systems to facilitate creation and navigation of visualizations, and software system features to improve visualization tools.

Perception and understanding

The viewer of scientific visualizations must not only perceive but also understand the information presented. Perception and