

Volume Composition and Evaluation Using Eye-Tracking Data

AIDONG LU

University of North Carolina at Charlotte
and

ROSS MACIEJEWSKI and DAVID S. EBERT

Purdue University

This article presents a method for automating rendering parameter selection to simplify tedious user interaction and improve the usability of visualization systems. Our approach acquires the important/interesting regions of a dataset through simple user interaction with an eye tracker. Based on this importance information, we automatically compute reasonable rendering parameters using a set of heuristic rules, which are adapted from visualization experience and psychophysical experiments. A user study has been conducted to evaluate these rendering parameters, and while the parameter selections for a specific visualization result are subjective, our approach provides good preliminary results for general users while allowing additional control adjustment. Furthermore, our system improves the interactivity of a visualization system by significantly reducing the required amount of parameter selections and providing good initial rendering parameters for newly acquired datasets of similar types.

Categories and Subject Descriptors: I.3.6 [Computer Graphics]: Methodology and Techniques—*interaction techniques*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*color, shading, and texture*

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: Usability and human factors in visualization, eye tracker, interaction, illustrative visualization, volume rendering.

ACM Reference Format:

Lu, A., Maciejewski, R., and Ebert, D. S. 2010. Volume composition and evaluation using eye-tracking data. *ACM Trans. Appl. Percept.* 7, 1, Article 4 (January 2010), 20 pages.

DOI = 10.1145/1658349.1658353 <http://doi.acm.org/10.1145/1658349.1658353>

1. INTRODUCTION

In order to create a meaningful and aesthetically pleasing computer visualization, substantial user effort is involved in manually adjusting the rendering parameters. Generally, each rendering approach has its own special parameters that need to be adjusted, such as color and opacity transfer functions for direct volume rendering. Many volume-rendering approaches also require users to choose the volume placement and viewing direction about the display. While the free selection of these parameters does provide a flexible environment for the user to generate various results, it also requires a significant

This research is supported by DOE DE-FG02-06ER25733, NSF 0081581, 0121288, 0328984, and 0633150.

Authors' addresses: Aidong Lu, University of North Carolina at Charlotte; email: Aidong.lu@uncc.edu

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1544-3558/2010/01-ART4 \$10.00

DOI 10.1145/1658349.1658353 <http://doi.acm.org/10.1145/1658349.1658353>

ACM Transactions on Applied Perception, Vol. 7, No. 1, Article 4, Publication date: January 2010.

amount of work and some knowledge about the applied rendering algorithm in order to obtain a satisfying result. As such, suitable automation of the rendering settings and parameters can simplify some of the tedious user interaction and improve the overall usability of a visualization system. To create such automation, a series of rendering rules must be generated to ensure appropriate placement, view direction, and other key parameter choices.

In building our set of rules, we turn to artistic and scientific illustrations to look at the principles that guide their selection of such parameters when creating images. Such illustrations have already shown their expressiveness in representing various subjects, and they are widely used in science and engineering. Furthermore, illustrators usually follow certain methodologies and procedures to effectively convey important figure aspects to the viewer. A key aspect of the perceived image quality is its composition, which usually emphasizes the focal points and achieves a coherent image structure. While the relative beauty of an image is subjective, some heuristics used by artists to create images are shared by general illustrative works and can be used as guidelines for the selection of rendering parameters.

Although these guidelines are difficult to apply directly in complex environments, it is possible to render volumetric datasets automatically according to several features of volume rendering, such as fixed object shapes and positions. We can treat volume composition as the problem of organizing a set of objects with constant positions and sizes effectively. We then extract features to determine regions of interest within a scene and automate the rendering parameter selection to highlight key features. While some volume features can be extracted through image processing and statistical approaches, determining regions of interest is a subjective issue. For example, given an image of a human hand, a physician might focus on the joints of the hand, while general viewers might be interested in the bone structure. To account for these discrepancies, our system utilizes an eye tracker to determine what areas of the volume attract a subject's general interest. This approach creates a high-quality rendering tuned to a user's regions of interest while still providing users with the flexibility to adjust parameters to their liking.

In order to determine if our automated parameter selection is effective in drawing attention to regions of interest, we have conducted a user study in which subjects are asked to find the important/highlighted regions of previously rendered volumes. Each of these previously rendered volumes has been automatically rendered with one region of interest. Our composition rules give more illustrative emphasis to the important areas, and as such, these areas should be discernible from other structures in the images. Our results show that the parameters generated for different rendering styles are able to highlight importance information for a significant portion of our subjects.

In the following sections, we will first summarize related work on importance-based rendering, composition, and eye-tracker studies. Then, in Section 3, we describe the procedure for gathering importance information from a user with an eye tracker. In Section 4, we present the procedure to process the importance information according to eye movement behaviors. In Section 5, we summarize a set of heuristic composition rules and use them to generate a plausible general-volume visualization automatically. Section 6 describes our user study, and the results are reported in Section 7. Finally, we discuss our results and propose future work in Section 8.

2. RELATED WORK

2.1 Importance-Based Rendering

Illustrative renderings can be more expressive than photographs because of their ability to emphasize important objects and simplify irrelevant details [Gooch and Gooch 2001; Strothotte and Schlechtweg 2002]. The important objects are given different meanings under different contexts, such as calculated salience or user-specified importance information. In computer graphics and visualization, there are several research topics that are closely related to this control of level of detail. First, importance-driven

approaches render objects at different levels according to their importance information. For example, Seligman and Feiner [1991] presented an automated intent-based approach using rule-based methods and evaluators, and Viola et al. [2004] extended the cut-away views to an automatic focus, and context display based on assigned importance values. Similarly, focus+context visualization renders objects in focus with more obvious styles. For example, Helbing et al. [1998] used emphasized rendering to communicate relevance and guide user focus, and Svakhine et al. [2005] presented an interactive medical illustration system with various combinations of rendering enhancements/styles. Third, cut-away views [Diepstraten et al. 2003; McGuffin et al. 2003] can also be included as one approach to emphasize important objects through cutting away or distorting the less important objects in the front. Fourth, stylization and abstraction topics emphasize important features or salience for both images and 3D objects. For example, Hamel and Strothotte [1999] used templates to describe and transfer rendering styles to save tuning the parameters. DeCarlo and Santella [2002] used eye-tracking data to stylize and abstract photographs into a line-drawing style with highlighted interested regions. They further validated their approach through user studies [2004]. Design principles have also been explored for creating effective assembly instructions based on cognitive psychology research [Agrawala et al. 2003]. Stylized focus for 3D models has also been developed to draw viewer's gaze to an emphasized area [Cole et al. 2006]. These approaches utilized techniques from different fields to emphasize important objects during the rendering process, which is one important criterion we consider to develop this composition approach.

2.2 Composition

Both interactive and automatic methods have been developed to reduce user interaction for selecting rendering parameters. Among these methods, rule-based approaches have been developed for several graphics and visualization applications. For example, Mackinlay [1986] developed a framework for automatic graphical presentation creation and evaluation. Beshers and Feiner [1993] designed rule-based visualization principles that take into account the characteristics of data, tasks, and hardware capability. Gooch et al. [2001] presented an overview of compositional principles and an approach for finding a good composition for 3D objects. Strothotte et al. [1994] designed sketch rendering and discussed the function of rendering choices on the perception and understanding of the viewers. Also, suitable user interaction has been introduced to guide composition algorithms. For example, Rist et al. [1994] argued that semiautomation was a reasonable compromise for computer-generated illustrations and it could release users from routine subtasks. Bergman et al. [1995] presented an interactive approach to guide the user's selection of color maps for various visualization tasks. Kowalski et al. [2001] guided the rendering parameter selections based on compositional needs in an interactive animation scene. Similar to these approaches, our method is also based on composition rules and incorporate a small amount of user interaction. The main difference is that we fully automate volume renderings according to regions of interest collected from the user's eye movements.

2.3 Eye-Tracker Studies

To acquire the information of regions of interest, eye tracking has been one popular approach in computer graphics, human-computer interaction, and psychology. We have divided the related work into two groups based on application type. The first group focuses on using the eye-gaze information to improve the understanding of knowledge. Research in psychology has shown that the eye movement patterns during complex scene perception are related to the information of the scene, as well as to the cognitive processing of the scene [Rayner 2004]. For example, eye-tracking data can be used to extract visual features from 2D/3D figures [Chung et al. 2004]. Comparing eye movement strategies can provide further information for professional education training [Law et al. 2004] and 2D/3D display analysis [Tory

et al. 2005]. Eye trackers have also been used to characterize low and high comprehenders [Bednarik et al. 2006].

The second group of approaches focuses on using eye gaze information to improve the performance of various systems. For example, Levoy and Whitaker [1990] let the high-resolution sweet spot follow the user's gaze to achieve a trade-off between the computation resource and the represented details. In human-computer interaction, eye-based interaction has been applied to design a gaze-controlled navigator [Starker and Bolt 1990; Tanriverdi and Jacob 2000] and replace the keyboard and mouse [Majaranta and Raiha 2002]. Eye tracking has also been used to improve the rendering speed in the simulation of a realistic collision [O'Sullivan and Dingliana 2001; O'Sullivan et al. 2003] and natural eye contact between users in a video conference [Jerald and Daily 2002]. In this article, we use eye tracker as a convenient interaction tool, and we use collected information to calculate regions of interest for volumetric datasets.

3. ACQUISITION OF IMPORTANCE INFORMATION

To select regions of interest for composition tasks, a useful technique in visualization is to adjust transfer functions [Bordoloi and Shen 2005; Takahashi et al. 2005]. Although transfer functions have been shown to be powerful visualization tools, they are not necessarily intuitive for general users. Our design principle is to develop an intuitive approach that can be used by both scientific experts and general users. We use the eye tracker as a convenient interaction tool and calculate regions of interest from the patterns of a user's eye movement. A user only needs to look at their interested regions in the volume on the screen and his/her eye movements during that time period will be automatically recorded. In the remainder of this section, we briefly discuss eye movement theory to show the foundation of using an eye tracker for 3D volume composition. Then, we describe our procedure of using an eye tracker to gather importance information.

It is known that eye movements seldom perform wasted motions, and typically focus near the best place to gather desired visual information [Sekuler and Blake 1994]. Therefore, one can determine what regions of an object a user is interested in by analyzing eye movements. Such an analysis has been mainly used for 2D or 2D-oriented applications, such as improving rendering speed [Levoy and Whitaker 1990; O'Sullivan and Dingliana 2001; O'Sullivan et al. 2003] and abstracting data information [DeCarlo and Santella 2002]. There are two types of basic eye movements: saccades and fixations. A saccade is a rapid intermittent eye movement that occurs when the eyes fix on one point after another, or for the purpose of lubrication. A fixation is when eye movements stop and focus on a particular object. It has been shown that fixations correspond to informative locations and the time actually spent fixating on a particular location indicates that processing of the object is taking place [Rayner 2004]. As such, we use the eye fixations to determine the user's regions of interest. We then group the eye data spatially and temporally and calculate the sets of sequential fixations. The lengths of the fixations are related to the interest degrees of the processed information.

To collect the importance information, we have designed a simple procedure for general users. As opposed to most 2D and 2D-oriented research with eye trackers, our objective is to gather the importance information for the 3D voxels of a volume. Since the eye tracker returns 2D point positions on the image plane, we need additional information to reconstruct the 3D points of the focal regions. Our approach is to let users look at a constantly rotating volume while we gather their eye movement data. With the rotation information, we can locate the 3D regions of interest from multiple consecutive eye data if the user keeps looking at the same position. As previously discussed, fixations are the sources for us to locate a user's region of interest and fixations usually take a significantly longer time than saccades; therefore, we can gather the importance information by rotating the observed volumes. The rotation pace is set at 10 seconds per 360 degrees, which is slow enough for a user to observe details and fast enough to




Fig. 1. A sequence of eye movement during the volume rotation and isovalue changing. The red point indicates the position of the eye gaze.

avoid wandering and boredom. The rotation direction can be interactively adjusted by users. During this interaction, their eye data is discarded, as the eye movements may involve other factors due to the interaction. This procedure ends when users feel that they have already explored the volume contents.

To measure the amount of detail from the eye-tracking data, we use the concept of visual acuity, which is a measure of the smallest detail that an observer can resolve under ideal conditions. Reddy [2001] summarized measure models from the vision literature into a single equation for visual acuity, $H(v, e)$. As shown in Equation (1), $H(v, e)$ is computed from the contrast sensitivity $G(v)$ and sensitivity drop-off $M(e)$, when having a velocity of v degrees/second and a peripheral extent e degrees. It shows that the highest sensitivity is located centrally in the fovea and varies across the retina. We use $H(v, e)$ as the weight of observed information for the collected eye data, where v is the volume's rotation speed on the image plane and e is measured from the distance between the user to the screen and the pixel position to the screen center.

$$\begin{aligned}
 G(v) &= \begin{cases} 60.0, & v \leq 0.825 \text{ deg/s} \\ 57.69 - 27.78 \log_{10}(v), & 0.825 \text{ deg/s} < v \leq 118.3 \text{ deg/s} \\ 0.1, & v > 118.3 \text{ deg/s} \end{cases} \\
 M(e) &= \begin{cases} 1.0, & e \leq 5.79 \text{ deg} \\ 7.49/(0.3e + 1)^2, & e > 5.79 \text{ deg} \end{cases} \\
 H(v, e) &= G(v)M(e) \tag{1}
 \end{aligned}$$

Another challenge with 3D volumes is that all the information from a volume cannot be shown to a user at once, as opposed to 2D images. To explore the objects in the volume, we use a standard volume rendering approach: transparent isosurfaces. While the volume is rotating, subjects can change the normalized isosurface values from 0 to 1. This allows the subjects to explore any of their regions of interest in the volume. The window for adjusting the isosurface value is independent from the rendering window. The eye tracker returns 0 when the subjects look outside the rendering window; therefore, the eye movements for adjusting the isosurface value do not affect the gathered importance information. Figure 1 shows a sequence of eye gaze positions that we recorded when a user traversed the isosurfaces of a human hand dataset.

4. PROCESSING OF IMPORTANCE INFORMATION

The data collected during the acquisition phase is a list of 2D eye movement points on the image plane and their corresponding volume rotation sequences. With this information, we reconstruct and cluster an eye gaze volume to generate importance maps that indicate the user's interest. We generate importance maps for both 3D space and data value space and will use them in the automatic composition phase.




Fig. 2. A pipeline overview of eye data process. We first collect 2D eye points when users are looking at a rotating volume. These 2D eye points are used to construct an eye gaze volume recording 3D locations of fixations. We then process and cluster this eye gaze volume to generate a 3D importance volume. Also, from the clustering result of Section 4.2, we construct the fixations on a 2D value space (histogram) and generate a 2D importance space in Section 4.3. Both the 3D importance volume and the 2D importance space are used to decide visualization parameters automatically in Section 5.

4.1 Reconstruction of the Eye Gaze Volume

Eye trackers generally collect eye movements at a fixed rate (such as 60 data-points per second) during an indicated time period. To correctly construct an eye gaze volume, we need to correspond eye movements with a volume rotation sequence. Since our eye-tracking data is acquired on a separate machine, we place a time stamp whenever we rotate the volume. This way, we can clearly link the volume rotations to the eye movements by setting each volume rotation matrix to all the eye movements within the corresponding time range.

As discussed in the previous section, we are only interested in using the fixations of the eye movements to gather the information of regions of interest. Therefore, we remove the eye data that moves faster than a normal fixation speed, which indicates a saccade. As Ohshima et al. [1996] suspended rendering when the eye moves faster than 180 deg/s, we use this value as the eye velocity threshold to distinguish saccades from fixations.

For each 2D point on the image plane, given its corresponding volume rotation matrix, it represents a line passing through the volume. Since the eye observance is located on the image plane, we let all the voxels projected on the same area to share its visual acuity value and use the same importance weight for all these voxels, regardless of their depth. According to the 2D visual acuity model, each eye data point represents that viewers received information from all the voxels that are projected on the image plane within a small area, so each eye point corresponds to a subvolume in the 3D space. For orthogonal views, this subvolume is a cylinder, since depth values do not affect projection positions. For projective views, the radius of this subvolume becomes smaller as the depth increases.

When viewing a 3D volume at each time stamp, users generally focus on a 3D point instead of an entire subvolume. Since a single 2D eye point on the image plane cannot suggest any depth information, we use consecutive eye points to locate the focus point. This can also identify the exact 3D location when multiple voxels are projected onto the same eye point. We cannot simply calculate the intersection of subvolumes, since the constructed shape of a focus point will be affected by its fixation duration. The shape of a focus point may vary from a thin line, when fixation duration is short, to a sphere, when fixation duration is long. To better measure the importance degrees around a 3D focus point, we change the 2D visual acuity model into a 3D model by assigning the same weight to all the voxels with equal distance to the center. The model center is set as the intersection point between two consecutive eye focus points. Instead of providing a visual acuity model for 3D space, we use this 3D model to remove the artifacts of the importance map from collected viewing angles.

The process of reconstruction starts from the second eye data. For each eye point, we test the velocity requirement, calculate the focus point with the previous point, and add the weights from the 3D visual acuity model to the eye gaze volume. After we generate the eye gaze volume, the regions with high

values indicate higher interests or importances to the user. We then normalize the generated eye gaze volume for the clustering procedure.

4.2 Clustering of the Eye Gaze Volume

When observing an object, human eyes often choose several suitable locations to look at the same object. Also, the positions of fixations may indicate areas that are difficult for users to understand, instead of a single region of interest. To combine close fixations that correspond to the same regions of interest, we need a clustering algorithm to group focus points from the eye gaze volume.

To prepare for clustering, we first select voxels that have higher importance values above an assigned threshold in the eye gaze volume. This helps to remove extra regions that are introduced just because they were projected close to some eye points on the image plane. We can use these voxel locations as a list of 3D points and input them to clustering algorithms to group fixation locations. A difficulty for such a clustering task is that we do not know the number of clusters ahead of time. This restricts us from using standard clustering algorithms, such as K-means. The mean shift algorithm [Comaniciu and Meer 2002; Georgescu et al. 2003], based on the gradient direction, has been shown to be a flexible and robust clustering algorithm. It can be used without the knowledge of cluster number and cluster shape. We adopted the mean shift procedure to decide cluster number, cluster centers (modes) and assign each input point to the closest cluster. We continue to merge clusters if their center distance is within the eye-tracker accuracy range. Small clusters with the size of less than 5% of the total point number are discarded, since they are not the main focus of the volume.

We further smooth focus regions by fitting a Gaussian model to each cluster set. All the points assigned to a cluster are used to calculate the parameters of a Gaussian model. The valid window size is located by including at least 90% of the points in this cluster. The cluster results are used to generate a smooth importance map, which will be further used to determine rendering parameters.

4.3 Clustering of Value Ranges

In addition to the 3D importance map, we also generate an importance map for value space. Since data values of focus points may indicate value ranges that users are interested in, this information can be used to detect data features of a volume with eye movements. We use important value ranges to design rendering parameters.

The generation process for importance map of value space is similar to the previously mentioned process of 3D importance map. We use a common 2D histogram (voxel value and gradient magnitude) [Kindlmann and Durkin 1998] as our value space. We first generate an initial 2D importance map by collecting the distribution of 3D focus points in the value space. Then, we cluster the initial 2D importance map with the mean shift algorithm and smooth cluster results with Gaussian models. This procedure is similar to the generation of 3D importance map and helps us produce a 2D importance map on the value space. Each cluster is shown as one Gaussian shape on the 2D histogram and corresponds to one 3D object in the volume.

We use the information of value importance map to finalize the 3D importance map for segmented and unsegmented volumes, respectively. For segmented volumes, we use the value importance map to hit on the objects in the volume. The hit numbers are normalized as object importance values so that we can further emphasize the regions that might not be in the user's focus regions but share some common features with them. For un-segmented volumes, we can use the value clusters to segment the volume automatically by revisiting the volume with the value importance map as a transfer function. Voxels that have data values and gradient magnitudes belong to a 2D importance cluster will be classified as portions of the corresponding 3D object. Therefore, we can treat segmented and unsegmented datasets

as volumes that are composed of several stable objects, and we can process these datasets in the same way as described in the following section.

5. AUTOMATIC RENDERING SETTINGS

Based on the importance maps acquired from previous sections, we can automate the process of volume rendering, which usually requires a user to manually adjust some rendering parameters. The main criterion we try to follow is that a good visualization often captures object features and emphasizes regions of interest in a manner that is not only scientifically accurate but also visually pleasing to the eye. Here, we will concentrate on producing a computer-generated visualization with emphasized important regions automatically. We will explore approaches for several necessary rendering parameters shared among general-volume visualization methods. From the previous two sections, we have prepared the following three data types to use in the automation process:

- $I_v()$: the importance value for each voxel (Section 4.2)
- $ID()$: an object ID for each voxel (Section 4.3)
- $I_o()$: the importance value for each object (Section 4.3)

5.1 View Direction

Psychologists use “canonical view” to refer to the viewpoint that is preferred by most viewers [Blanz et al. 1999]. The study of canonical view searches for the factors that can affect our perception and understanding by observing the consistency across subjects, instead of locating a unique best view for certain types of objects. There are consistent heuristics for choosing view directions from both artists and psychology results, such as to pick an off-axis view from a natural eye height. Since most of them match our criteria of a good visualization image, we can use the common factors of a canonical view to guide our automatic viewpoint selection.

Canonical views have been studied for face recognition [Laeng and Rouw 2001], procedural graphics [Krull et al. 2003], 3D models [Denton et al. 2004], and animation generation [Kawai et al. 1993; He et al. 1996]. Gooch et al. [2001] chose initial viewpoint according to the proportions of projection areas and perturbed viewing parameters guided by heuristic rules for layout. Vázquez et al. [2001] used viewpoint entropy to compute good viewing positions and select good views for scene understanding of polygon models. Recently, the amount of information or viewpoint entropy are also explored for selecting good viewpoints for volume, time-varying, and isosurface renderings [Bordoloi and Shen 2005; Takahashi et al. 2005; Ji and Shen 2006; Viola et al. 2006]. Most work measured the “goodness” of a view direction in a certain way, such as designing objective functions or user experiments. For a volume, since we do not have any specific information about the geometry shapes of the objects in the volume, the viewpoint selection becomes more challenging. The major difference in this article from the previous work is that we treat parameter selections as a more subjective issue, and we use an eye tracker to acquire this information from the user, which is more intuitive than adjusting transfer functions and does not require the knowledge of volume rendering.

Here, we briefly list the factors for viewpoint selections and describe our interpretation. We remove the factors that are related to experiences, since these are impossible to quantize without understanding the context of a dataset.

Saliency: A view that shows the saliency and significance of features is preferred by most observers. For volume data, we use several standard values to represent data features, including the gradient magnitude G , curvature magnitude C , and an edge detection value, E . The saliency of a voxel is represented as a weighted sum of all the feature factors.

Occlusion: Occlusion is used to avoid too many crucial features overlapping with each other. A visualization result is always rendered in a view direction that the information would be sufficient or clear for object recognition except for special purposes. Since volume visualization can show the inside information as well as surfaces, occlusion is a very important factor in evaluating the quality of a view direction. We measure the occlusion by projecting voxels onto the image plane and a good view direction should include fewer overlapping voxels that have high salience or importance values.

Stability: A view from which small rotations will not produce significant changes is preferred to avoid ambiguities. Stability is also related to the occlusion factor, since there would be no stability problem if any two objects/salience were not overlapping on the image plane. We measure the variance of a view direction within a small region as the stability factor.

Familiarity: The views that are encountered most frequently or during initial learning period are generally preferred. Since familiarity is difficult to measure without context information, we match the familiarity by presenting the user-interest objects closer to the eye position as an approximation. For instance, when looking at a volume, an observer is usually interested in the facial portion of a head or the bones of a foot.

Combining these four factors, we calculate a weight $W()$ for each voxel v using salience, importance, and the normalized distance to the viewer location $Distance(v)$, as shown in Equation (2). We use equal weights w_s, w_i, w_d , since all the values have been normalized already.

$$W(v) = w_s * \overbrace{(G + C + E)}^{Salience} * I_o(v) + w_i * I_v(v) + \underbrace{w_d * I_v(v) * (1 - Distance(v))}_{Familiarity} \quad (2)$$

To measure the occlusion degree of a view, we adopt a procedure similar to the splatting algorithm [Zwicker et al. 2001]. We use a temporary buffer, with the same size of the image plane and initialize all the values to be 1. Each voxel in the volume is projected onto the image plane and the corresponding buffer values are modified with the voxel weight calculated in Equation (2). To avoid the handling of voxel sequences, we multiply the weights of voxels and their corresponding buffer values so that the occlusion degrees for all the views can be measured with this procedure.

$$O(view) = \sum_{p \in Image\ plane} \overbrace{\prod_{v \in Volume\ and\ v\ projects\ on\ p} (1.0 + W(v))}^{Occlusion} \quad (3)$$

The evaluation (“badness”) function of a view direction $view$ is calculated as the negative of the sum of all the items in the buffer and the variances.

$$Evaluation(view) = -(O(view) + \overbrace{Variance(O(view))}^{Stability}) \quad (4)$$

To find a minimum value of such an implicit function (caused by the variance portion), we use an optimization process, which only requires the function values instead of the derivatives of the objective function. Gooch et al. [2001] use the downhill simplex method, which is well behaved for the problem with a small computational burden. Since the design of our objective function involves more calculations, we adapt the direction set method for faster processing [Press et al. 1992]. We divide the view space into a set of samples by tessellating a sphere. Then, we choose the initial view direction with the minimum evaluation value from the sample set. Finally, a minimum view is searched within the divided range, which is a much smaller space than the whole view space. This process is guaranteed to find a local

minimum value, which is good enough to find a plausible answer. With an increase in the initial view sampler set, we can find a near-global minimum result.

5.2 Volume Center Position

After determining the view direction, we need to locate the volume center in space. Artistic illustrations often achieve a balanced structure, attract viewer's attention, and avoid equal divisions of an image. Also, most visualization systems have fixed rendering window size and always put the volume at the center of the screen. For a practical visualization system, we keep fixed rendering window size and put the volume on the center so that we do not need to constantly change the volume center position to fit all the objects inside the screen while rotating the volume.

We use the golden ratio (1:1.63) as the ratio of object size to the rendering window, since it is shown to be more appealing than others [Livio 2002]. We calculate the bounding box of important clusters on the image plane. The volume center is located by setting the ratio of the bounding box size and the image plane to the golden ratio.

5.3 Rendering Parameters

The rendering settings of view direction and volume center position are designed for the whole volume, while other rendering parameters are different for each object inside the volume. After a volume is put in a good position to observe with the first two settings, we design the rendering parameters for each object in this section so that they can be distinguished from the surroundings and the important objects can stand out in the results. Since rendering parameters are dependent on the specific rendering algorithms, we concentrate on the parameters that are common to general-volume rendering approaches.

Our parameters are designed according to the relations between each pair of objects that are measured through their positions and sizes. Previously, spatial relationships are often considered using 9-intersection model [Egenhofer and Herring 1990] and the dimensional model [Zlatanova et al. 2002]. We calculate the impact factor of one object to another to approximate objects' spatial relations, which will be used later to decide rendering parameters. For every object, we scan the volume from front to back to measure the overlapping area of every other object in front of it. The overlapping area is divided by the object projection area to overcome the impact of the object size and overlapping size, since a smaller sized object is considered to have less influence than a larger one. For objects A and B , the impact factor $Impact(A, B)$ is considered through the overlapping regions $Overlap(A, B)$ on the image plane and the distance between two objects $Dis(A, B)$.

$$\begin{aligned} Overlap(A, B) &= \frac{Area(A, B)}{Area(B)} \\ Impact(A, B) &= Overlap(A, B) + 1/Dis(A, B) \end{aligned} \quad (5)$$

5.3.1 Rendering Elements. When a volume contains multiple objects, a visualization system usually assigns different rendering elements to each object to distinguish them, such as with different colors or patterns. We use the relation matrix $Impact(A, B)$ and the importance map generated from the previous section to assign rendering elements. For each object, we assign several rendering elements from the aspects of distinguishing and emphasizing, assuming that users do not have any preknowledge of the real objects. To effectively visualize a dataset, we choose different elements for each object and keep large perception distances among them. The most effective elements will be selected for emphasizing important objects, which are acquired from the importance map. We concentrate on color selections in this section and show that the same procedure can be used to assign other elements, such as patterns.

Since colors play an important role in data visualization, approaches have been devoted to color-mapping designs [Rheingans and Tebbs 1990; Rheingans and Landreth 1995; Bergman et al. 1995] and

color representation validations [Healey 1996]. When users interactively select colors for rendering, they generally prefer to use more distinctive colors for objects with higher importance values; therefore, we assign colors according to their importance value distribution. To balance the effects of color luminance, our color map is chosen from the perceptual isoluminance color map [Kindlmann et al. 2002] and the Euclidean distance between different colors is measured in a perceptually balanced color model (CIE LUV) [Fairchild 1998]. We use the following procedure to ensure that our color assignments satisfy their importance value distribution. For a volume with n objects, we equally divide the color map into n sections and randomly select one color near the center of each section. To map the color distance and object relation information, we normalize them to the same value region, respectively. Then, we generate all the possible assignments and the final colors are chosen by minimizing the differences between the object relations and the color distances $ColorDis(i, j)$ for colors i and j .

$$M(i, j) = \sum_{i,j} (ColorDis(i, j) - Impact(i, j))^2 \quad (6)$$

The same procedure can be used to assign other rendering elements with predesigned sensitive values, such as patterns. The rendering patterns can also influence the perception effects for users, as shown in Kim et al. [2004]. We simply assign the sensitive values with the size of the pattern primitives. The objects with higher importance values are rendered with the patterns composed of larger primitives, while the least important object is rendered with small points. Figure 5(c, d) shows the automatically generated results with different patch patterns.

5.3.2 Rendering Degrees. We use rendering degrees to refer to the parameters related to the level of details or opacity values. Intuitively, a more important object will be rendered with more detail and a higher opacity value. If we use the object importance value to assign the rendering degree for each object directly, a problem arises when there are objects overlapping each other. The most important object will occlude all the objects behind it and the less important object may occlude part of the more important objects. Therefore, we use the overlapping relationships and the importance map to assign rendering degree.

Our two basic rules are based on the overlapping relations that can affect our understanding of object shapes and locations [Dowling 1987]. If there is no overlapping relation, the object is rendered at the highest degree. When two objects overlap in the image plane, the object at the back is rendered with a degree as high as possible; while the front object is rendered at a suitable degree to show part of the back object, no matter its importance value.

We use the following procedure to calculate all the rendering degrees. Initially, the rendering degrees $RenderDegree()$ of all the objects are set to 1. Then, we traverse the object with positive importance values in the decreasing order. For each object, we update the degree of the objects in front of it using the overlapping area proportion and the importance values. Assuming object A is in front of object B, we update the degree of A with the overlapping ratio of A on B $Overlap(A, B)$ as a weight:

$$\begin{aligned} RenderDegree(A) &= Overlap(A, B) * f(I_o(A), I_o(B)) \\ &+ (1 - Overlap(A, B)) * RenderDegree(A) \end{aligned} \quad (7)$$

where $f(x, y) = 0, x \leq y; \frac{x-y}{x}, x > y.$

The calculated rendering degree can be used to determine the rendering parameters directly for some algorithms. For example, the degree can be used to decide the opacity values for transfer functions. We set the opacities with $Opacity(x) = 0.1 + RenderDegree(x) * 0.7$ according to our experiences, so that all objects are not totally transparent, and we always show the volume inside. For the silhouette effect of nonphotorealistic rendering (NPR), the silhouette power in $Opacity(x) = (\overrightarrow{Gradient} \cdot \overrightarrow{View})^{Sp(x)}$ is




Fig. 3. The reconstructed eye gaze volumes, segmented clusters, and our composition results for a segmented hand dataset (top) and a foot dataset (bottom). For the hand dataset, since the bones are viewed as more important than the skin, they are less transparent and rendered with less silhouette enhancement and a warmer color. The viewpoint is also selected for better bone observance. For the foot dataset, although the user is more interested in the bones on the first and second toes, all the bones are highlighted because of their value similarities.




Fig. 4. Two pairs of eye gaze volumes and composition results for a segmented feet dataset. The left pair focuses on the bones, and the right pair focuses on the skin portion. These images show that different composite visualizations are generated according to the user interests for best observance of objects of interest.

set as $Sp(x) = 10 - 9 * RenderDegree(x)$ so that the least important object will be rendered only with silhouettes. Figures 3 and 4 show several automatic composition results in which colors are only selected from the yellow to red hue range for more natural results.

It is relatively difficult for other rendering algorithms to summarize the changes of parameters according to their desired rendering degrees. This problem corresponds to the parameter selections of different rendering motifs [Svakhine et al. 2005]. For these approaches, the rendering degree is used to determine the parameters on a high level, which are extracted from the input of experienced users, while the lower-level interface is used to explicitly adjust the rendering parameters by users.

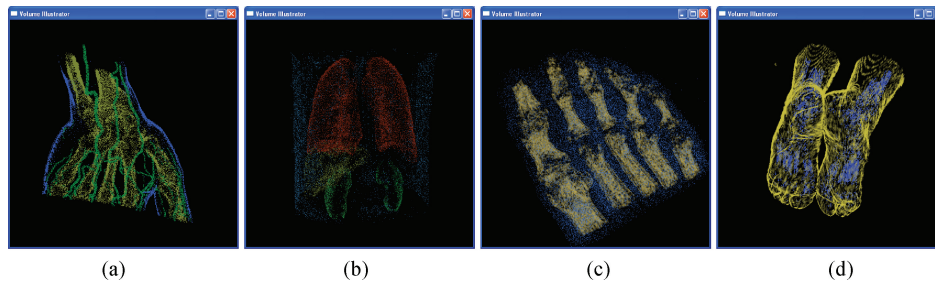


Fig. 5. The parameter selections of two illustrative rendering styles from the collected importance information. (a,b) Composition results of hand and abdomen datasets in a stippling rendering style. The focused regions are the bones of the hand and the lungs in the abdomen. (c,d) Composition results of foot and feet datasets in a patch-based rendering style. The focused regions are the bones of (c) and skins of (d).

Figure 5 (a, b) shows the stipple rendering result with the high-level interfaces that use only several 0 – 1 scalar values, including a gradient slider, sharpness slider, orientation slider, distance slider, and interior slider, to control all the required parameters in the rendering process [Lu et al. 2003].

6. EXPERIMENT DESIGN

We have performed a psychophysical experiment to evaluate the effects of different rendering styles and parameters on user interaction. Since the parameters are designed to reflect important objects in a volume, we expect that these regions of interest in the volume composition results draw more attentions from viewers. Therefore, we hypothesize that subjects shown images rendered with our automatic parameter generation technique will be able to visually locate which area of the image was considered to be a region of interest. This hypothesis leads us to evaluate partial effects of volume composition results and provides some useful results for our future work. We also expect that the effects of parameters vary from different rendering styles and have chosen three rendering styles of both the abdomen and hand datasets in the experiment.

6.1 Apparatus

The system used in our experiment, as shown in Figure 6, consisted of a computer with an Intel Xeon 2.0 GH CPU and an ISCAN ETL-400 eye tracker. In order to acquire accurate results, subjects were required to place their heads in a chin rest to minimize head movement. Rendering was done as a preprocess step, and a set of static images was used for evaluation.

6.2 Subjects

Eight subjects (S1–S8) participated in the experiment. Five subjects were female. All subjects had little to no familiarity with the datasets. All subjects had normal or corrected vision. The subjects ages ranged from 15 to 43 years of age. No testing was performed to assess color vision impairment.

6.3 Stimuli

Two datasets were used in this experiment. They included the hand dataset and the abdomen dataset, see Figure 7. Each dataset was rendered in three different styles and highlighted three different portions of the data based on a predefined region of interest. The rows in Figure 7 represent each rendering style applied to a given dataset. Each column in Figure 7 represent a particular region of interest. For the hand dataset, one rendering highlighted the ulna, another the wrist, and the last, the thumb. For the abdomen dataset, one rendering highlighted the liver, another the kidney opposite the liver, and the



Fig. 6. A user participating in the eye-track study.

last, the heart. Each of the three renderings from the datasets were scaled to 700×700 pixels. These renderings were then rotated and flipped to provide five different viewing angles of each image. This created a set of 90 images.

6.4 Procedures

We tested the subject's ability to visually determine which portion of the dataset was the region of interest by displaying each image for 3 seconds and asking subjects to focus their attention on the portion of the image that they determined to be "highlighted." Prior to each experiment, a subject was presented with a test set of images demonstrating exactly what was meant by highlighting. Subjects were placed 320mm away from the monitor. Images were then presented in random order to each subject. Each image was shown for exactly 3 seconds. This was followed by a black screen with a square placed in one corner of the viewing area for 1 second to redirect the subject's focus. The experiment length was approximately 10 minutes including calibration time. Specifically, subjects were instructed to search the image for what appeared to be a highlighted segment and then focus their gaze on that segment until the the image changed to a black screen.

7. EXPERIMENTAL RESULTS AND ANALYSIS

In order to evaluate the automatic parameter generation, we analyzed the subjects' eye gaze patterns to determine if they were able to accurately identify areas within the image that were highlighted. We did not measure the time which a subject took to find the highlighted area. Instead, we determined the average time it took subjects to come to a consistent gaze location, removing all cases where subjects failed to settle on any location. The average time was approximately 1.5 seconds to locate a region of interest. From there, we focused only on the last 1 second of data collected, as this should be the time in which subjects are most focused on what they consider to be the region of interest in the image.

For each subject, we considered three cases for the hand dataset: (i) the subject looks at the thumb, (ii) the subject looks at the wrist, (iii) the subject looks at the ulna. We also consider three cases for

Table I. Summary of Subjects Able to Distinguish Different Areas in Images

	Stipple	Realistic	Illustrative
Hand	87.5% (7/8)	62.5% (5/8)	62.5% (5/8)
Abdomen	87.5% (7/8)	62.5% (5/8)	25.0% (2/8)

the abdomen dataset: (i) the subject looks at the image with the kidney highlighted, (ii) the subject looks at the image with the heart highlighted, (iii) the subject looks at the highlighted liver opposite of the kidney. For each subject, we looked at each rendering style to see if the subject could accurately determine the highlighted object. To do this, we would compare cases (i), (ii), and (iii) to see if each case was statistically different in each rendering style for each dataset. For each case and rendering style, the subject looked at five images. Each image was rotated or mirrored in each of the five cases such that the subject would not learn the exact location of the highlighted portion for the repeated trials. Each image and the corresponding gaze locations were rotated back to a normal viewing position. The last second for each trial was then aggregated into a single dataset such that for each case and rendering style, we now have a dataset consisting of the last second of data for five trials. These gaze location datasets were then analyzed in a pair-wise comparison test to determine if the subject was able to distinguish different areas in cases (i), (ii), and (iii) of each dataset and rendering style. For example, within the rendering style of stippling, we perform a pair-wise comparison test between the gaze locations from the trials in which the subject looked at images with the thumb highlighted and compared these locations to the gaze locations from the trials in which the subject looked at images with the wrist highlighted, and so on. Under the pair-wise comparison test, we test the null hypothesis that the means of the gaze locations for cases (i), (ii), and (iii) are equal. If the null hypothesis cannot be rejected at the 5% level, then we have shown that the subject was not able to distinguish different areas of the image in cases (i), (ii), and (iii) for the rendering style under question.

We analyze differences within subjects and images across highlighting areas through a Wilcoxon Rank-Sum test to do a pair-wise comparison of cases (i), (ii), and (iii) separately. We use this test, which is the nonparametric analog of the t -test, rather than a t -test because we believe the data are not normally distributed, which is confirmed through visual analysis of histograms and the Shapiro-Wilk test ($p < 0.001$, rejection of normal distribution). Furthermore, under the assumption that subjects have found the area in question during the last second of viewing, it would stand to reason that the data would not be normally distributed during this time. The results of the Wilcoxon Rank-Sum test are summarized in Table I. Note that the location along the x -axis is tested separately from location along the y -axis; however, only one must be significantly different for the location to be confirmed as statistically different. For example, if between cases (i) and (ii) the mean X value was not statistically different, but the mean Y value was, then the 2D gaze point location should be at a different location.

For the stipple hand cases, seven out of eight subjects were able to find each highlighted object ($p < .02$). For the hand images rendered in the realistic style, five of eight subjects were able to find each highlighted object ($p < .04$). For the hand images rendered in the illustrative style, five of eight subjects were able to find each highlighted object ($p < .04$). So, from the eye-tracker results, we can see that for the hand stipple analysis only one out of eight subjects could not identify the highlighted/important region in the images. For the hand images rendered in the realistic style, three out of eight subjects could not identify the highlighted/important region in the images. Particularly, this comes from subjects having difficulty separating the ulna from the wrist. For the hand images rendered in the illustrative style, three out of eight subjects could not identify the highlighted/important region in the images. Here, subjects are unable to separate any of the three objects, indicating this rendering style is the least useful for directing attention.

From the hand images, we can argue that stippling is the most effective for guiding users to a particular focus. This is most likely due to the fact that the detail in the other images drawing attention away from objects. We believe it would be possible to use stippling to guide a user's focus and then change the rendering style to provide more detail.

These same results were confirmed in the abdomen dataset. For the stipple abdomen cases, seven out of the eight subjects were able to find each highlighted object ($p < .02$). For the abdomen images rendered in the realistic style, five out of the eight subjects were able to find each highlighted object ($p < .04$). For the abdomen images rendered in the illustrative style, two out of the eight subjects were able to find each highlighted object ($p < .04$). For the abdomen stipple images, we can see that only one out of eight subjects could not identify the highlighted/important region in the images. This subject was the same subject who had difficulty in the stipple hand images. For the abdomen dataset rendered in the realistic style, three out of eight subjects could not identify the highlighted/important region in the images. Two of the subjects are the same as in the hand images. Here, there is more variance in the data due to the fact that subjects have a larger area to look at in terms of the number of distinct objects visible in the abdomen. For the abdomen dataset rendered in the illustrative style, six out of eight subjects could not identify the highlighted/important region in the image. Here, subjects were unable to separate any of the three objects, again indicating that this rendering style is the worst for directing attention.

From this experiment, we can see that the parameters generated for both stipple and realistic rendering styles are able to adequately highlight importance information for over 50% of our subjects. However, for the illustrative style, we find that our automatically generated parameters do not adequately focus the user's attention to the intended targets. This is most likely due to the inconsistency of texture details and the importance information of the acquired regions of interest, since all the textures are just selected from two medical illustrations. Further studies will be needed for exploring the selections of textures according to the importance information due to their complex factors.

While we have now shown that we can guide subjects to focus on different areas of the image for different highlighting, it is also necessary to see if the focus areas were statistically significantly close to the area of interest that was highlighted. We have plotted the mean and standard deviations of the gaze locations for the subjects to further confirm their gaze locations. In Figure 7, each subject's mean gaze location (over the last second of viewing) is plotted over the image viewed. The radius of the corresponding circle is the standard deviation. Again, we see that in the stippled images (rows 1 and 4) subjects are more likely to find the highlighted portion of the image when compared to the illustrative style (rows 2 and 5) and the realistic style (row 3 and 6).

Since the design of this experiment concentrates on the highlighting aspect of volume composition results, we call it an initial user study. Also, among the three rendering styles, stippling is the easiest to create areas of large contrast, which draw a users gaze, as opposed to other more realistic rendering styles. A more comprehensive study should include more perception and cognition aspects to evaluate rendering parameters and styles.

8. DISCUSSION AND FUTURE WORK

Volume composition aims to improve a practical issue of general-volume visualization systems and reduce the repetitive and tedious user interaction. Although such a rule-based approach cannot compete with the results generated by real users, the user interaction needed for several most common tasks can be significantly reduced. Our interface allows users to concentrate on their specific tasks and are more intuitive for general users, which can be used further to improve the usability of a visualization system.

To produce visualization with different concentrations, we believe that human factors should be included in the visualization design. An eye tracker is a good tool in this case, since it provides input

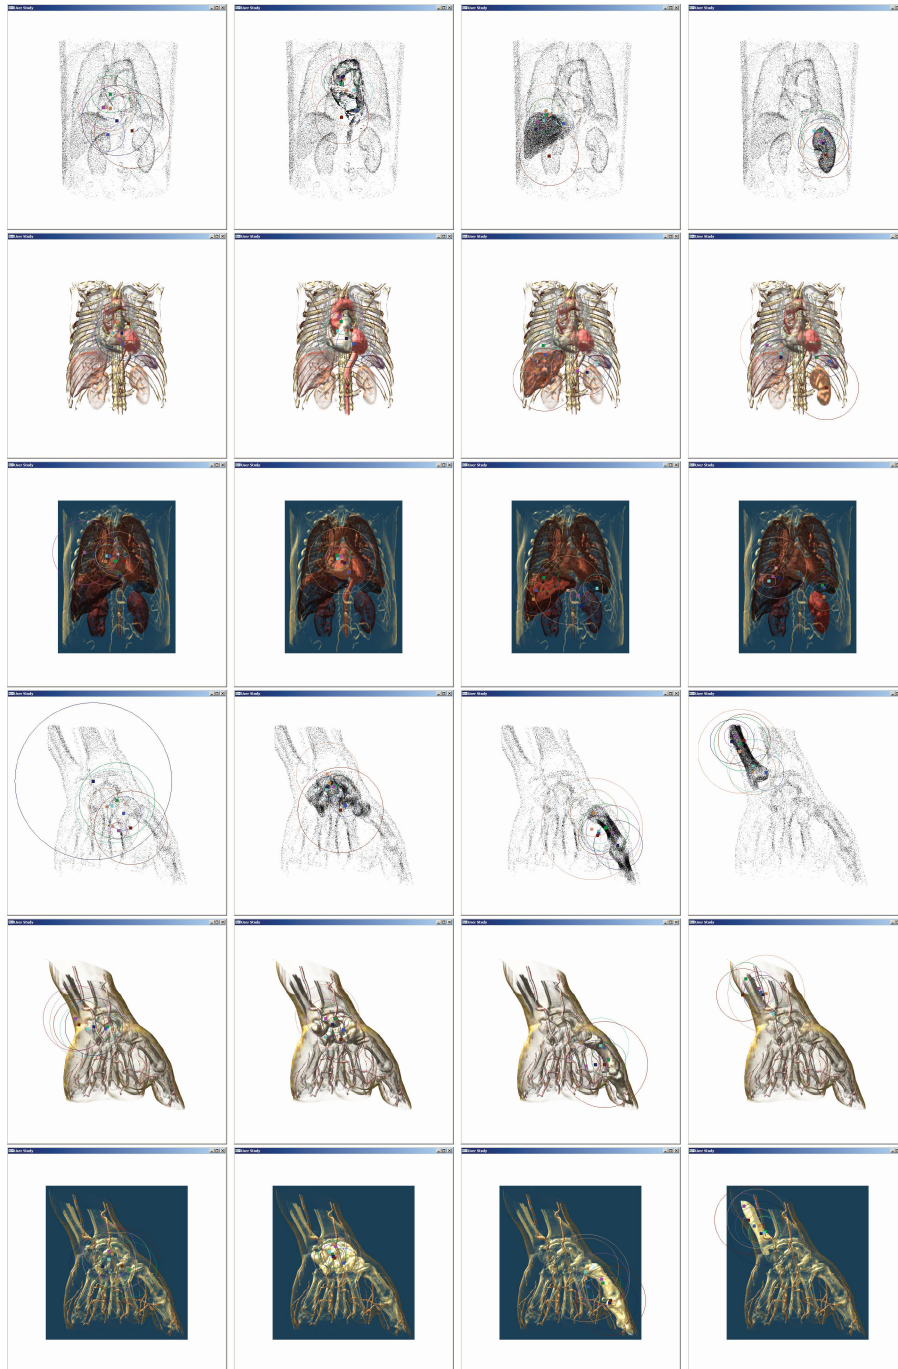


Fig. 7. Example visualization results for the user study. The focuses and variances of individual subjects are represented as circles in the images.

from users and is convenient to use. We show that the importance information acquired using an eye tracker can be used to choose viewpoint, volume center, and multiple rendering parameters. We believe that this importance information can be explored to develop automatic composition approaches for more visualization parameters.

With an eye tracker, we have built a simple interface that can be used by both professional and general users. Without the knowledge of the rendering approach, general users can still explore volume data and achieve satisfying visualization results. Our initial user study demonstrates the potential effectiveness of this automatic parameter selection method. This method is more effective for acquiring immediate instinct reactions from users. We are planning to use an eye tracker as an additional input and evaluation method to further simplify user interaction. There are also limitations of this approach due to the eye movement behaviors. When several small objects are very close to each other, it may be difficult to locate the exact object that users are looking at. We believe that more studies on the correspondence of eye movements and cognition process will help us to improve the accuracy of this approach.

The unoptimized composition algorithm for the results in Figures 3 and 4 takes 10 to 20 minutes, which is mainly spent on the clustering part. Once the algorithm is done, the rendering is interactive and the user can explore the volume based on their interests. We will develop faster algorithms to provide more instant feedback, since our final goal is to use the eye tracker as an interactive input method. We plan to work on more approaches to guide the parameter selections for general and specific rendering methods.

We plan to perform future user studies to assess the effectiveness of visualization results. Beyond our initial experiment, we plan to study the factors due to the rotation of the images, since it is possible that the viewing direction plays a confounding role in the subjects' ability to disseminate information, as these images are naturally occurring scenes. We believe that it is worth a series of experiments to evaluate rendering parameters and styles from more perception and cognition aspects.

ACKNOWLEDGMENTS

We would like to thank the editors and reviewers for their insightful comments.

REFERENCES

- AGRAWALA, M., PHAN, D., HEISER, J., HAYMAKER, J., KLINGNER, J., HANRAHAN, P., AND TVERSKY, B. 2003. Designing effective step-by-step assembly instructions. In *Proceedings of SIGGRAPH*. ACM, New York, 828–837.
- BEDNARIK, R., MYLLER, N., SUTINEN, E., AND TUKIAINEN, M. 2006. Program visualization: Comparing eye-tracking patterns with comprehension summaries and performance. In *Proceedings of the 18th Workshop of the Psychology of Programming Interest Group*. ACM, New York, 68–82.
- BERGMAN, L. D., ROGOWITZ, B. E., AND TREINISH, L. A. 1995. A rule-based tool for assisting color map selection. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA, 118–125.
- BESHES, C. AND FEINER, S. 1993. Auto-visual: Rule-based design of interactive multivariate visualizations. *IEEE Comput. Graph. Appl.* 13, 4, 41–49.
- BLANZ, V., TARR, M. J., AND BULTHOFF, H. H. 1999. What object attributes determine canonical views? *Perception* 28, 5.
- BORDOLOI, U. AND SHEN, H. 2005. View selection for volume rendering. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA, 487–494.
- CHUNG, A. J., DELIGIANNI, F., HU, X.-P., AND YANG, G.-Z. 2004. Visual feature extraction via eye tracking for saliency driven 2D/3D registration. In *Proceedings of the Eye Tracking Research and Applications Symposium*. ACM, New York, 49–54.
- COLE, F., DECARLO, D., FINKELSTEIN, A., KIN, K., MORLEY, K., AND SANTELLA, A. 2006. Directing gaze in 3D models with stylized focus. In *Proceedings of the Euro-Graphics Symposium on Rendering*. ACM, New York, 377–387.
- COMANICIU, D. AND MEER, P. 2002. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 5, 603–619.

- DECARLO, D. AND SANTELLA, A. 2002. Stylization AND abstraction of photographs. In *Proceedings of SIGGRAPH*. ACM, New York, 769–776.
- DENTON, T., ABRAHAMSON, J., AND SHOKOUFANDEH, A. 2004. Approximation of canonical sets AND their applications to 2D view simplification. In *Proceedings of the IEEE Conference on Computer Vision AND Pattern Recognition*. IEEE, Los Alamitos, CA.
- DIEPSTRATEN, J., WEISKOPF, D., AND ERTL, T. 2003. Interactive cutaway illustrations. In *Proceedings of Eurographics*. ACM, New York, 523–532.
- DOWLING, J. E. 1987. *The Retina: An Approachable Part of the Brain*. Belknap Press, Cambridge, MA.
- EGENHOFER, M. J. AND HERRING, J. R. 1990. A mathematical framework for the definition of topological relationships. In *Proceedings of the 4th International Symposium on Spatial Data Handling*. ACM, New York, 803–813.
- FAIRCHILD, M. D. 1998. *Color Appearance Models*. Addison-Wesley, Upper Saddle River, NJ.
- GEORGESCU, B., SHIMSHONI, I., AND MEER, P. 2003. Mean shift based clustering in high dimensions: A texture classification example. In *Proceedings of the 9th International Conference on Computer Vision*. IEEE, Los Alamitos, CA, 456–463.
- GOOCH, B. AND GOOCH, A. 2001. *Nonphotorealistic Rendering*. A. K. Peters, Natick, MA.
- GOOCH, B., REINHARD, E., MOULDING, C., AND SHIRLEY, P. 2001. Artistic composition for image creation. In *Proceedings Eurographics Rendering Workshop*. ACM, New York, 83–88.
- HAMEL, J. AND STROTHOTTE, T. 1999. Capturing AND re-using rendition styles for non-photorealistic rendering. In *Proceedings of Eurographics*. IEEE, Los Alamitos, CA.
- HE, L., COHEN, M. F., AND SALESIN, D. 1996. The virtual cinematographer: A paradigm for automatic real-time camera control AND directing. In *Proceedings of SIGGRAPH*. ACM, New York, 217–224.
- HEALEY, C. G. 1996. Choosing effective colors for data visualization. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA, 263–270.
- HELBING, R., HARTMANN, K., AND STROTHOTTE, T. 1998. Dynamic visual emphasis in interactive technical documentation. In *Proceedings of the European Conference on Artificial Intelligence*. John Wiley and Sons, Hoboken, NJ.
- JERALD, J. AND DAILY, M. 2002. Eye-gaze correction for videoconferencing. In *Proceedings of the Human Factors in Computing Systems Conference*. ACM, New York, 77–81.
- JI, G. AND SHEN, H.-W. 2006. Dynamic view selection for time-varying volumes. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA.
- KAWAI, J. K., PAINTER, J. S., AND COHEN, M. F. 1993. Radioptimization: Goal-based rendering. In *Proceedings of the SIGGRAPH Conference*. ACM, New York.
- KIM, S., HAGH-SHENAS, H., AND INTERRANTE, V. 2004. Conveying shape with texture: Experimental investigations of texture’s effects on shape categorization judgments. *IEEE Trans. Visual Comput. Graph.* 10, 4, 471–483.
- KINDLMANN, G. AND DURKIN, J. W. 1998. Semi-automatic generation of transfer functions for direct volume rendering. In *Proceeding of the IEEE Symposium on Volume Visualization*. IEEE, Los Alamitos, CA, 79–86.
- KINDLMANN, G., REINHARD, E., AND CREEM, S. 2002. Face-based luminance matching for perceptual color map generation. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA, 299–306.
- KOWALSKI, M. A., HUGHES, J. F., RUBIN, C. B., AND OHYA, J. 2001. User-guided composition effects for art-based rendering. In *Proceedings of the Symposium on Interactive 3D*. ACM, New York, 99–102.
- KRULL, R., SHARP, M., AND ROY, D. 2003. Canonical views in procedural graphics. In *Proceeding of the International Professional Communication Conference*. IEEE, Los Alamitos, CA.
- LAENG, B. AND ROUW, R. 2001. Canonical views of faces AND the cerebral hemispheres. *Laterality* 6, 3, 193–224.
- LAW, B., ATKINS, M. S., KIRKPATRICK, A. E., AND LOMAX, A. J. 2004. Eye-gaze patterns differentiate novice AND experts in a virtual laparoscopic surgery training environment. In *Proceedings of the Eye Tracking Research and Applications Symposium*. ACM, New York, 41–48.
- LEVOY, M. AND WHITAKER, R. 1990. Gaze-directed volume rendering. *Comput. Graph.* 24, 2, 217–223.
- LIVIO, M. 2002. *The Golden Ratio: The Story of Phi, the World’s Most Astonishing Number*. Broadway Books, New York.
- LU, A., MORRIS, C., TAYLOR, J., EBERT, D., RHEINGANS, P., HANSEN, C., AND HARTNER, M. 2003. Illustrative interactive stipple rendering. *IEEE Trans. Visual Comput. Graph.* 9, 2, 127–139.
- MACKINLAY, J. D. 1986. Automating the design of graphical presentations of relational information. *ACM Trans. Graph.* 5, 2, 110–141.
- MAJARANTA, P. AND RAIHA, K.-J. 2002. Twenty years of eye typing: Systems and design issues. In *Proceedings of the Eye Tracking Research AND Applications Symposium*. ACM, New York.
- MCGUFFIN, M., TANCAU, L., AND BALAKRISHNAN, R. 2003. Using deformations for browsing volumetric data. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA, 401–408.

- OHSHIMA, T., YAMAMOTO, H., AND TAMURA, H. 1996. Gaze-directed adaptive rendering for interacting with virtual space. In *Proceedings of the IEEE Virtual Reality Annual International Symposium*. IEEE, Los Alamitos, CA, 103–110.
- O’SULLIVAN, C. AND DINGLIANA, J. 2001. Collisions and perception. *ACM Trans. Graph.* 20, 3.
- O’SULLIVAN, C., DINGLIANA, J., AND HOWLETT, S. 2003. Eye-movements and interactive graphics. In *The Mind’s Eyes: Cognitive and Applied Aspects of Eye Movement Research*, J. Hyona, R. Radach, AND H. Deubel, Eds., Elsevier Science, 555–571.
- PRESS, W., FLANNERY, B., TEUKOLSKY, S., AND VETTERLING, W. 1992. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge, UK.
- RAYNER, K. 2004. Eye-movements as reflections of perceptual and cognitive processes. In *Proceedings of the Eye Tracking Research and Applications Symposium*. ACM, New York, 9–10.
- REDDY, M. 2001. Perceptually optimized 3D graphics. *IEEE Comput. Graph. Appl.* 21, 5, 68–75.
- RHEINGANS, P. AND LANDRETH, C. 1995. Perceptual principles for effective visualizations. *Perceptual Issues Visual.* 59–74.
- RHEINGANS, P. AND TEBBS, B. 1990. A tool for dynamic explorations of color mappings. *Comput. Graph.* 24, 2, 145–146.
- RIST, T., KRGER, A., SCHNEIDER, G., AND ZIMMERMANN, D. 1994. Awi: A workbench for semi-automated illustration design. In *Proceedings of the Workshop on Advanced Visual Interfaces*. ACM, New York, 59–68.
- SANTELLA, A. AND DECARLO, D. 2004. Visual interest and npr: An evaluation and manifesto. In *Proceedings of the Conference on Non-Photorealistic Animation and Rendering*. ACM, New York, 71–78.
- SEKULER, R. AND BLAKE, R. 1994. *Perception*. McGraw-Hill, New York.
- SELIGMAN, D. D. AND FEINER, S. K. 1991. Automated generation of intent-based 3D illustrations. In *Proceedings of SIGGRAPH*. ACM, New York.
- STARKER, I. AND BOLT, R. 1990. A gaze-responsive self-disclosing display. In *Proceedings of the Human Factors in Computing Systems Conference*. ACM, New York, 3–9.
- STROTHOTTE, T., PREIM, B., RAAB, A., SCHUMANN, J., AND FORSEY, D. R. 1994. How to render frames and influence people. *Comput. Graph. Forum* 13, 3, 455–466.
- STROTHOTTE, T. AND SCHLECHTWEG, S. 2002. *Non-Photorealistic Computer Graphics: Modeling, Rendering and Animation*. Morgan Kaufmann, San Francisco, CA.
- SVAKHINE, N., EBERT, D. S., AND STREDNEY, D. 2005. Illustration motifs for effective medical volume illustration. *IEEE Comput. Graph. Appl.* 25, 3, 31–39.
- TAKAHASHI, S., FUJISHIRO, I., TAKESHIMA, Y., AND NISHITA, T. 2005. A feature-driven approach to locating optimal viewpoints for volume visualization. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA.
- TANRIVERDI, V. AND JACOB, R. 2000. Interacting with eye-movements in virtual environments. In *Proceedings of the Human Factors in Computing Systems Conference*. ACM, New York, 265–272.
- TORY, M., ATKINS, M. S., KIRKPATRICK, A. E., NICOLAOU, M., AND YANG, G. 2005. Eye-gaze area-of-interest analysis of 2D and 3D combination displays. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA.
- V’AZQUEZ, P., FEIXAS, M., SBERT, M., AND HEIDRICH, W. 2001. Viewpoint selection using viewpoint entropy. In *Proceedings of Vision, Modeling, and Visualization Conference*. 273–280.
- VIOLA, I., FEIXAS, M., SBERT, M., AND GRÖLLER, M. E. 2006. Importance-driven focus of attention. *IEEE Trans. Visual. Comput. Graph.* 12, 5, 933–940.
- VIOLA, I., KANITSAR, A., AND GRÖLLER, M. E. 2004. Importance-driven volume rendering. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA, 139–145.
- ZLATANOVA, S., RAHMAN A.A., AND SHI W. 2002. Topology for 3D spatial objects. In *Proceedings of the International Symposium and Exhibition on Geoinformation*.
- ZWICKER, M., PFISTER, H., VAN BAAR, J., AND GROSS, M. 2001. EWA volume splitting. In *Proceedings of the IEEE Conference on Visualization*. IEEE, Los Alamitos, CA, 29–36.

Received August 2008; revised June 2008, November 2008; accepted November 2008