

Volume Composition Using Eye Tracking Data

Aidong Lu^{†1}, Ross Maciejewski², and David S. Ebert^{‡2}

¹University of North Carolina at Charlotte ²Purdue University

Abstract

This paper presents a method to automate rendering parameter selection, simplifying tedious user interaction and improving the usability of visualization systems. Our approach acquires regions-of-interest for a dataset with an eye tracker and simple user interaction. Based on this importance information, we then automatically compute reasonable rendering parameters using a set of heuristic rules adapted from visualization experience and psychophysics experiments. While the parameter selections for a specific visualization task are subjective, our approach provides good starting results that can be refined by the user. Our system improves the interactivity of a visualization system by significantly reducing the necessary parameter selection and providing good initial rendering parameters for newly acquired datasets of similar types.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Interaction Techniques; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

To create a meaningful and aesthetically pleasing computer visualization, substantial user effort is usually involved to manually adjust the rendering parameters. For most volume rendering approaches, the users are required to put a volume in a suitable location in space and choose their preferred viewing directions. Different rendering approaches may have their special parameters, such as color and opacity transfer functions for direct volume renderings, and the enhancement degrees for illustrative visualizations. Although the free selection of these parameters provides a flexible environment for generating various results, it also requires a significant amount of work and some knowledge of the rendering algorithms in order to obtain a satisfying result. Therefore, suitable automation of the rendering settings and parameters can simplify some of the tedious user interaction and improve the usability of a visualization system.

Artistic and scientific illustrations have already shown their expressiveness in representing their subjects and they are widely used in science and engineering. Illustrators usually follow certain methodologies and procedures to effectively convey the important aspects of the subjects. One as-

pect of image quality is composition, which usually emphasizes the focal points and achieves a coherent image structure. While the aesthetic properties of an image are subjective, some heuristics used by artists to create images are shared by general illustrative works and can be used as general guidelines for the selection of rendering parameters.

Because of the subjective nature of judging image quality, these composition rules have limited application in computer graphics. Successful examples include the automatic selection of camera settings for animation generation and the automated parameterization of motion. Complex application environments can also increase the difficulty of using composition rules.

However, several characteristics of volume rendering make automatic volume composition possible, such as fixed object shapes and positions (except for volume deformation). Therefore, we can treat volume composition as a problem with a set of fixed objects with constant positions and sizes and generate automatic settings.

While some volume features can be extracted through image processing and statistical approaches, determining the regions-of-interest is a subjective issue. Some features may be interesting to most people, while the regions-of-interest for other subjects can be very different. For example, given an image of a human hand, a physician might focus on the

[†] alu1@uncc.edu

[‡] rmacieje, ebtert@purdue.edu

joints of the wrist, while general viewers might be interested in the bone structure. To account for these discrepancies, our system utilizes an eye tracker, which uses a camera to focus on one eye and record the eye movements as the user observes the volume. This data is then processed to discern the regions-of-interest for a particular user. The information of these regions are used as input for generating a good composition for volume rendering. Once a good composition is chosen, system parameters are automatically chosen to generate the initial rendering. This approach gives users the flexibility to adjust the parameters to their liking, while quickly providing a reasonable quality image.

In the following sections, we first summarize related work on importance-based rendering, composition, and eye tracker studies. In Section 3, we describe the procedure for gathering importance information from a user with an eye tracker. This importance information is processed (Section 4) according to eye movement behaviors and is used to generate a plausible general volume visualization automatically by summarizing a set of heuristic composition rules (Section 5). Finally, we discuss our results and propose future work.

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. 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 the level of detail. First, importance-driven approaches render objects at different levels according to their importance information, such as an intent-based system using rule-based methods by Seligman and Feiner [SF91]. Similarly, focus+context visualization renders objects in focus with more obvious styles. For example, Helbing et al. [HHS98] used emphasized rendering to communicate relevance and guide user focus. Second, cut-away views can also be included as one approach to emphasize important objects through cutting away or distorting the less important objects on the front. Third, stylization and abstraction topics emphasize important features or salience for both images and 3D objects. Hamel and Strothotte [HS99] used templates to describe and transfer rendering styles. DeCarlo and Santella [DS02] used eye tracking data to stylize and abstract photographs with highlighted interested regions. They further validated their approach through user studies [SD04]. Fourth, design principles have also been explored for creating effective assembly instructions based on cognitive psychology research.

2.2. Composition

While composition is not a popular topic, many methods have been developed to reduce user interaction for selecting rendering parameters. Mackinlay [Mac86] automated graphical presentation creation and evaluation. Beshers and Feiner [BF93] designed rule-based visualization principles from multiple characteristics. Rist et al. [RKSZ94] argued that semi-automation was a reasonable compromise for computer-generated illustrations and it could release users from routine subtasks. Bergman et al. [BRT95] guided the user's selection of colormaps interactively for various visualization tasks. Gooch et al. [GRMS01] presented an overview of compositional principles and an approach to find a good composition for 3D objects. Kowalski et al. [KHRO01] guided the rendering parameter selections based on compositional needs for animation.

2.3. Eye Tracker Studies

To acquire user-interest information, eye tracking has been one popular approach in computer graphics, human-computer interaction, and psychology. According to the application type, we roughly divide the related work into two groups. The first group focuses on using eye gaze information to improve the performance of various systems, such as volume acceleration following the users' gaze by Levoy et al. [LW90]. In human-computer interaction, eye-based interaction has been applied to design gaze-controlled navigator and replace keyboard and mouse. It has also been used to improve the rendering speed in simulating realistic collision [OD01] and natural eye contact between users in video conference [JD02]. The second group of approaches focus on using the eye-gaze information to improve the understanding of knowledge. The research in psychology has shown that eye movement patterns during complex scene perception are related to the information in the scene as well as the cognitive processing of the scene [Ray04]. For example, eye tracking data can be used to extract visual features from 2D/3D figures [CDHY04]. Comparing eye movement strategies can provide further information for professional education training [LAKL04] and 2D/3D display analysis [TAK*05].

3. Acquisition of Importance Information

In visualization, a common technique for acquiring importance information is through the adjustment of transfer functions [BS05, TFTN05]. Although transfer functions have been shown to be powerful visualization tools, they are not intuitive for general users who are not familiar with visualization techniques. We treat "volume composition" as a problem for both scientific experts and general users, and our approach is to use an eye tracker as a convenient tool to determine the regions-of-interest of the user. The user only needs to look at their volumetric regions of interest on the

screen, and we will automatically record their eye movements during this time period. In the remainder of this section, we first briefly discuss eye movement theory to show the foundation of 3D volume composition based on eye tracking devices; then, we describe our designed procedure of gathering importance information using an eye tracker.

It is known that human eyes seldom perform wasted motions and typically eye movements focus near the best place to gather the desired visual information [Red01]. Therefore, one can determine what regions of an object a user is interested in by analyzing their eye movements. The analysis of eye movements has been mainly applied for 2D or 2D-oriented applications, such as improving rendering speed [LW90, OD01, ODH03] and abstracting data information [DS02]. There are two types of basic eye movements: fixation and saccade. A fixation occurs when eye movements stop and focus on a particular object, indicating informative locations and the ongoing information processing by the user [Ray04]. A saccade is a rapid intermittent eye movement that occurs when either the eyes fix on one point after another or for the purpose of lubrication. Therefore, fixations are the main sources to indicate the viewer's regions-of-interest. When the eye data is grouped spatially and temporally, we can calculate a list of sequential fixations and their lengths are related to the degrees of user interest in the processed information. Because of the usage of fixations, we can gather information for a 3D volume based on users' eye movements and ensure that they can be used to acquire the importance information for the volume composition.

To collect the importance information, we designed a simple procedure for general users. Different from most 2D and 2D-oriented research with eye tracker, our objective is to gather the information for the 3D voxels of a volume. Since eye tracking data are 2D points on the image plane, we need additional information to reconstruct the 3D positions of the focal points. 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 discussed in the foundation of eye movements in 3D space, fixations are the sources for us to locate users' region-of-interest and fixations usually take a significantly longer time than saccades; therefore, we can gather the importance information with rotating volumes. The rotation pace is set as 10 seconds per 360 degrees, which is slow enough for a user to observe details and fast enough to avoid wandering and boredom. The rotation direction can be interactively adjusted by users, while their eye data during this period are discarded since the eye movements may involve other factors from the interaction. This procedure ends when users feel that they have already explored the volume contents.

To measure the amount of importance from the eye tracking data, we use the concept of visual acuity, which measures

the smallest detail that an observer can resolve under ideal conditions. Reddy [Red01] summarized the measure models from vision literature into a single equation for visual acuity $H(v, e)$ from contrast sensitivity $G(v)$ and sensitivity drop-off $M(e)$ giving a velocity v deg/s and peripheral extent e deg, as shown in Equation 1. 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/importance for the collected eye data, where v is the volume rotating 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.

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

Another challenge with 3D volumes is that all the information from a volume cannot be shown to a user at one time as with 2D images. To help general users to explore the volume contents and features, we use a standard volume rendering approach: semi-transparent isosurfaces. With the volume rotating, the subjects can simply change the normalized isosurface values from 0 to 1. This allows the subjects to explore most of their interest regions in the volume. Extremely complex situations, such as nested isosurfaces, can be modified by adding an arbitrary clipping plane. The window for adjusting the isosurface value is independent of 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 the recorded eye gaze positions when a user traversed the isosurfaces of a human hand dataset.

4. Processing of Importance Information

The eye tracking data collected during the acquisition phase are a list of 2D eye movement points on the image plane and their corresponding volume rotation matrices. With these two, we can generate importance maps for both volume space and data space. These importance maps indicate the user's interest degrees (importance values) and are used in the automatic composition phase in Section 5. Next, we discuss the reconstruction and clustering processes for generating importance maps from eye tracking data.

4.1. Reconstruction of the visiting volume

We first construct an "eye visiting volume" with the corresponding eye movements and volume rotation matrices. An




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

eye tracker collects 2D eye movements at a fixed rate (such as 60 data points per second) during the indicated time period. Since our eye tracking data is acquired on a separate machine, we place a time stamp at the start of each rendering frame. 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, only the fixations of eye movements will be used to calculate the regions-of-interest. Therefore, we remove the eye data that moves faster than normal fixation speed, which indicates a saccade. As Ohshima et al. [OYT96] suspend rendering when the eye moves faster than 180 deg/s, we use this value as the eye velocity threshold to distinguish between fixations and saccades.

Each 2D eye point on the image plane, given its corresponding volume rotation matrix, represents a line passing through the volume. Since the eye observance is located on the image plane, we let all the voxels with the same projection locations to share the same visual acuity. Therefore, the same importance weight is used for all the voxels regardless of their depth. With the 2D visual acuity model, each eye point represents information received from a sub-volume, which is composed by all the voxels that are projected within the positive weight area on the image plane. Depending on the projection options, this sub-volume is a cylinder for orthogonal views since the depth does not affect the projection positions. For projective views, the radius of the sub-region becomes larger with the depth goes further. Different projection sub-volume shapes are used for the corresponding rendering setting.

When viewing a volume, a user is usually interested in a focal point in the space instead of the whole sub-volume. Since a single eye movement on the projection plane cannot suggest any depth information, we use two consecutive eye points to locate the focal point. One issue arising during this process is that the constructed visiting volume will be affected by the time a user spends on the fixation. Due to this, the shape of a focal point varies from a thin line to a sphere. Since the importance degree of a focal point is indicated by the fixation period instead of the viewing angle, 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 points. Instead of providing a visual acuity model for 3D space, we use this 3D model to remove the artifacts in the importance map from its collected viewing angles.

The process of reconstruction starts from the second eye data point. For each eye point, we test the velocity requirement, calculate the location of the focal point using the previous eye point, and add the weights from the 3D visual acuity model to the visiting volume. After we generate the visiting volume, the regions with larger values indicate higher interest or importance to the user. We then normalize the generated visiting volume for the clustering part.

4.2. Clustering of the visiting volume

Since fixations indicate regions-of-interest or portions that are complex to understand, the position of a fixation does not necessarily correspond to a particular object. Instead, human eyes often choose several suitable locations to observe their focal point. Therefore, to combine the close fixations that correspond to the same region-of-interest, we need a clustering algorithm to group together the focal points from the visiting volume.

Here, a list of 3D points is generated from the normalized visiting volume with a scale function and used as the input of the clustering algorithms. The problem 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 [CM02, GSM03], based on the gradient direction, has been shown to be a flexible and robust clustering algorithm. It can be used without the knowledge of the cluster number and cluster shape. Therefore, we adopted the mean shift procedure to produce the modes (cluster centers) and the cluster index for each input point. Later, we merge the clusters if their center distance is within the eye tracker accuracy range. Small clusters are also discarded since they are not the main focus of the volume.

Next, each point cluster set is used to calculate the correlation matrix based on the statistical 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 the importance map for the volume, which will be further used to determine the rendering parameters.

4.3. Clustering of data ranges

Besides clustering in the volume, we also collect the importance map in the data space. The focal points in the data space indicate the data ranges of interest to the users. This information can be used to detect the data features in a volume through eye movements, and the collected data ranges are then used to design the rendering parameters.

The importance map for the volume is mapped onto the transfer function to generate the initial visiting map. We use a standard scalar value and gradient magnitude transfer function [Lev89, KD98]. Then, in the same sequence of generating 3D focal clusters, the initial visiting map is used to generate 2D points and produces a set of clusters using the mean shift algorithm.

For segmented volumes, we use the importance map to hit on the objects in the volume. The hitting numbers are normalized as the object importance values. We intend to further emphasize the portions which are not in the regions-of-interest, but share some common features with them. For unsegmented volumes, we can use the data clusters to segment the volume automatically by re-visiting the volumes from the transfer functions. Therefore, we can treat segmented and un-segmented datasets as volumes that are composed of several stable objects, and process them in the same way in the following section.

5. Automatic Rendering Settings

Volume rendering approaches usually require a user to manually adjust the rendering parameters. A good visualization often captures object features and emphasizes interest regions in a manner that is not only scientifically accurate but also visually pleasing to the eye. Based on the importance map acquired from the previous sections, we can automatically compose a computer-generated visualization with emphasized important regions. Here we will only concentrate on several necessary rendering parameters for general volume visualization approaches. From the previous two sections, we have prepared the following data to use in the automation process:

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

5.1. View direction

Psychologists and vision researchers use “canonical view” to refer to the viewpoint that is preferred by most viewers [BTB99]. 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 typed 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. Most of them match our criteria of a good visualization image. Therefore, we can use the common factors of a canonical view to guide our automatic viewpoint selection.

Canonical views have been studied for face recognition [LR01], procedural graphics [KSR03], 3D models [DAS04], and animation generation [KPC93, wHCS96]. Gooch et al. [GRMS01] chose initial viewpoint according to the proportions of the projection areas and perturbed viewing parameters guided by heuristic rules for layout. Vázquez et al. [VFSSH01] used viewpoint entropy to compute good viewing positions and select good views for scene understanding for polygon models. Bordolai and Shen [BS05] extended this work to select good viewpoints by measuring transfer function, data distribution, and voxel visibility. Takahashi et al. [TFTN05] also extended viewpoint entropy to select optimal viewpoint for isosurface renderings. 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 don’t have any specific information about the geometry shapes of the objects in the volume, the viewpoint selection becomes more challenging. The major difference of this paper from the related work is that we treat parameter selections as a more subjective issue and we use an eye tracker to acquire this information from a 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 dataset content.

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 the 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 from which the information is 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 to evaluate good view directions. We measure the occlusion by projecting voxels onto the image plane and a good view direction should include fewer overlapping voxels that have high saliency 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/saliencies are not overlapping on the image plane. We measure the vari-

ance of a view direction within a small region as the stability factor.

Familiarity: The views that are encountered most frequently or during initial learning are preferred because of object representations and the strategies applied for recognition. We match the familiarity by presenting the user-interest objects closer to the eye position. 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 the weight $W()$ of each voxel with salience, importance, and the normalized distance to the eye point, as shown in Equation 2. We use equal weights since all the values are normalized.

$$W(x) = w_s * \overbrace{(G + C + E)}^{\text{Salience}} * I_o(x) + w_i * I_v(x) \quad (2)$$

$$+ w_d * \underbrace{I_v(x) * (1 - \text{Distance}(x))}_{\text{Familiarity}}$$

Similar to the splatting algorithm [ZPvBG01], each voxel throws its weight on the image plane. A buffer that has the same size of the image plane is initialized with 1. To avoid the handling of the voxel sequences, we multiply the weight of the voxel with the value in the buffer.

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

The “goodness” function of a view direction is calculated as the negative of the sum of all the items in the buffer and the variances.

$$\text{Good}(v) = -(O(v) + \overbrace{\text{Variance}(O(v))}_{\text{Stability}}) \quad (4)$$

To find a minimum value of such an implicit function, we use an optimization process which only requires the function values instead of the derivatives of the objective function. Gooch et al. [GRMS01] use the downhill simplex method, which is well behaved for the problem with a small computational burden. Since our design of the objective function involves more calculations, we adapt the direction set method for faster processing [PFTV92]. We divide the view space into a set of samples by tessellating a sphere. Then we choose the initial view direction with the minimum “goodness” value from the sample set. Finally, a minimum view is searched within the divided range, that is a much smaller space than the whole view space. This process is guaranteed to find a local minimum value that is good enough for a plausible answer. By increasing 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 in space. Artistic illustrations often achieve a bal-

anced structure, attract viewer’s attention, and avoid equal divisions of an image; while most visualization systems have fixed rendering window size and always put the volume at the center of the screen. For a practical visualization approach, we keep the fixed rendering window size and put the volume at the center, so that we do not need to constantly change the volume center position to fit all the objects inside the screen when rotating the volume.

We use the golden ratio (1:1.63) as the ratio of the object size to the rendering window, since it is shown to be more appealing than others [Liv02]. We calculate the bounding box of the important clusters on the image plane. The volume center is located by setting the ratio of the bounding box size and the image plane as the golden ratio.

5.3. Rendering degrees

We use rendering degree to refer to the parameters related to the level of detail or opacities. Intuitively, a more important object will be rendered with more details and a higher opacity value. If we use the object importance values to assign the rendering degrees directly, one problem arises when there are objects overlapping with 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 determine the rendering degrees.

Our two basic rules are based on the overlapping relations. If there is no overlapping regions, the object will be 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 how important it is [Dow87].

Initially, the rendering degree $D()$ of all the objects is set to be 1. Then we traverse the objects with positive importance values in a decreasing order. For each object, we update the degree of the objects in front of it using the overlapping area proportion and their importance values. Assuming object A is in front of object B, we update the degree of A only if $I_o(A)$ is smaller than $I_o(B)$. The overlapping ratio of A on B, $O_v(A, B)$, is used to decide the weight of the updating function:

$$D(A) = O_v(A, B) * f(I_o(A), I_o(B)) + (1 - O_v(A, B)) * D(A)$$

$$\text{where } f(x, y) = 0, x \leq y; \frac{x-y}{y}, x > y. \quad (5)$$

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

$10 - 9 * D(x)$, so that the least important object will be rendered only with silhouettes. With a set of selected sample colors listed according to their hue values, we use warmer colors for objects with higher rendering degrees. Figure 2 shows several composition results using our eye movements. For other rendering algorithms using more complex parameters, we will need further work to explore more relationship between the parameters or combine sample inputs.

6. Discussion and Future Work

Volume composition aims to improve a practical issue of general volume visualization systems: reducing the repetitive and tedious user interaction. Although such a rule-based approach cannot compete with the intelligent results generated by users, the user interaction needed for many common tasks can be significantly reduced. Our interface allows users to concentrate on their specific tasks and is convenient to use. This convenience can further improve the usability of a visualization system.

Since parameter setting of a good visualization is a subjective issue, we believe that human factors should be included into the visualization design. An eye tracker is a good tool in this case because it can provide input from users in a simple way. We show that the importance information acquired with an eye tracker can be used to choose viewpoint, volume center, and rendering degrees. 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. We are planning to use an eye tracker as an additional input and evaluation method to further simplify the user interaction. This is very effective at acquiring instinct reactions from users. The eye movements of an expert can also be studied for training and education.

The unoptimized composition algorithm for the results in Figure 2 takes 10-20 minutes, which is mainly spent on clustering. 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. The future work also includes designing and performing user studies to assess and validate the effectiveness of this approach. We plan to work on more approaches to guide the parameter selections for general and specific rendering methods.

This project is supported by the National Science Foundation under Grant Nos. 0081581, 0121288, 0222675, 0328984, 9978032.

References

- [BF93] BESHERS C., FEINER S.: Autovisual: Rule-based design of interactive multivariate visualizations. *IEEE Computer Graphics and Applications* 13, 4 (1993), 41–49.
- [BRT95] BERGMAN L. D., ROGOWITZ B. E., TREINISH L. A.: A rule-based tool for assisting colormap selection. In *IEEE Visualization 1995* (1995), pp. 118–125.
- [BS05] BORDOLOI U., SHEN H.: View selection for volume rendering. In *IEEE Visualization* (2005), pp. 487–494.
- [BTB99] BLANZ V., TARR M. J., BULTHOFF H. H.: What object attributes determine canonical views? *Perception* 28, 5 (1999).
- [CDHY04] CHUNG A. J., DELIGIANNI F., HU X.-P., YANG G.-Z.: Visual feature extraction via eye tracking for saliency driven 2d/3d registration. In *Proceedings of the Eye tracking research and applications symposium* (2004), pp. 49–54.
- [CM02] COMANICIU D., MEER P.: Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 5 (2002), 603–619.
- [DAS04] DENTON T., ABRAHAMSON J., SHOKOUFANDEH A.: Approximation of canonical sets and their applications to 2d view simplification. In *IEEE CVPR* (2004).
- [Dow87] DOWLING J. E.: *The Retina: An Approachable Part of the Brain*. Belknap Press, 1987.
- [DS02] DECARLO D., SANTELLA A.: Stylization and abstraction of photographs. In *SIGGRAPH* (2002), pp. 769–776.
- [GRMS01] GOOCH B., REINHARD E., MOULDING C., SHIRLEY P.: Artistic composition for image creation. In *Eurographics Rendering Workshop 2001* (2001), pp. 83–88.
- [GSM03] GEORGESCU B., SHIMSHONI I., MEER P.: Mean shift based clustering in high dimensions: A texture classification example. In *PROC. of 9th International Conference on Computer Vision* (2003), pp. 456–463.
- [HHS98] HELBING R., HARTMANN K., STROTHOTTE T.: Dynamic visual emphasis in interactive technical documentation. In *Proc. of ECAI* (1998).
- [HS99] HAMEL J., STROTHOTTE T.: Capturing and re-using rendition styles for non-photorealistic rendering. In *EUROGRAPHICS* (1999).
- [JD02] JERALD J., DAILY M.: Eye gaze correction for videoconferencing. In *Human Factors in Computing Systems Conference* (2002), pp. 77–81.
- [KD98] KINDLMANN G., DURKIN J. W.: Semi-automatic generation of transfer functions for direct volume rendering. In *IEEE Symposium on Volume Visualization* (1998), pp. 79–86.
- [KHRO01] KOWALSKI M. A., HUGHES J. F., RUBIN C. B., OHYA J.: User-guided composition effects for art-based rendering. In *Proceedings of the 2001 symposium on Interactive 3D* (2001), pp. 99–102.
- [KPC93] KAWAI J. K., PAINTER J. S., COHEN M. F.: Radioptimization: goal based rendering. In *SIGGRAPH* (1993).
- [KSR03] KRULL R., SHARP M., ROY D.: Canonical views in procedural graphics. In *IPCC: The Shape of Knowledge* (2003).
- [LAKL04] LAW B., ATKINS M. S., KIRKPATRICK A. E., LOMAX A. J.: Eye gaze patterns differentiate novice and experts in




Figure 2: (Left) The left six images show the reconstructed visiting volumes, the cluster results, and the composition results for the segmented hand dataset and a foot dataset. For the hand, since the bones are viewed as more important than the skin, they are less transparent, with less silhouette enhancement and a warmer color. The viewpoint is also selected for better bone observance. For the foot, 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. (Right) The right four images show two pairs of visiting volumes and composition results for a segmented feet dataset. The top pair focuses on the bones and the bottom pair focuses on the skin. Different user interests result in different composite visualizations, which are adjusted specifically to observe the objects of interest using our automatic approach.

- a virtual laparoscopic surgery training environment. In *Proceedings of the Eye tracking research and applications symposium* (2004), pp. 41–48.
- [Lev89] LEVOY M.: *Display of Surfaces from Volume Data*. PhD thesis, University of North Carolina at Chapel Hill, 1989.
- [Liv02] LIVIO M.: *The Golden Ratio: The Story of Phi, the World's Most Astonishing Number*. Broadway Books, 2002.
- [LR01] LAENG B., ROUW R.: Tcanonical views of faces and the cerebral hemispheres. *LATERALITY* 6, 3 (2001), 193–224.
- [LW90] LEVOY M., WHITAKER R.: Gaze-directed volume rendering. *Computer Graphics* 24, 2 (March 1990), 217–223.
- [Mac86] MACKINLAY J. D.: Automating the design of graphical presentations of relational information. *ACM Transaction on graphics* 5, 2 (1986), 110–141.
- [OD01] O'SULLIVAN C., DINGLIANA J.: Collisions and perception. *ACM Transactions on Graphics* 20, 3 (2001).
- [ODH03] O'SULLIVAN C., DINGLIANA J., HOWLETT S.: Eye-movements and interactive graphics. In *The Mind's Eyes: Cognitive and Applied Aspects of Eye Movement Research*, Hyona J., Radach R., Deubel H., (Eds.). Elsevier Science, Oxford, April 2003, pp. 555–571.
- [OYT96] OHSHIMA T., YAMAMOTO H., TAMURA H.: Gaze-directed adaptive rendering for interacting with virtual space. In *IEEE VRAIS* (1996), pp. 103–110.
- [PFTV92] PRESS W., FLANNERY B., TEUKOLSKY S., VETTERLING W.: *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.
- [Ray04] RAYNER K.: Eye movements as reflections of perceptual and cognitive processes. In *Proceedings of the Eye tracking research and applications symposium* (2004), pp. 9–10.
- [Red01] REDDY M.: Perceptually optimized 3d graphics. *IEEE CG&A* 21, 5 (2001), 68–75.
- [RKSZ94] RIST T., KRÜGER A., SCHNEIDER G., ZIMMERMANN D.: Awi: a workbench for semi-automated illustration design. In *Proc. of the workshop on Advanced visual interfaces* (1994), pp. 59–68.
- [SD04] SANTELLA A., DECARLO D.: Visual interest and npr: an evaluation and manifesto. In *Proc. of NPAR* (2004), pp. 71–78.
- [SF91] SELIGMAN D. D., FEINER S. K.: Automated generation of intent-based 3d illustrations. In *Proc. of SIGGRAPH* (1991).
- [TAK*05] TORY M., ATKINS M. S., KIRKPATRICK A. E., NICOLAOU M., YANG G.: Eyegaze area-of-interest analysis of 2d and 3d combination displays. In *IEEE Visualization* (2005).
- [TFTN05] TAKAHASHI S., FUJISHIRO I., TAKESHIMA Y., NISHITA T.: A feature-driven approach to locating optimal viewpoints for volume visualization. In *IEEE Visualization* (2005).
- [VFSH01] VÁZQUEZ P., FEIXAS M., SBERT M., HEIDRICH W.: Viewpoint selection using viewpoint entropy. In *Proc. of VMV* (2001), pp. 273–280.
- [wHCS96] WEI HE L., COHEN M. F., SALESIN D.: The virtual cinematographer: A paradigm for automatic real-time camera control and directing. In *SIGGRAPH* (1996), pp. 217–224.
- [ZPvBG01] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: EWA volume splatting. In *IEEE Visualization 2001* (2001).