

INTEGRATING VOLUME VISUALIZATION TECHNIQUES INTO MEDICAL APPLICATIONS

*Stefan Bruckner**, *Peter Kohlmann**, *Armin Kanitsar†*, *M. Eduard Gröller**

*Institute of Computer Graphics and Algorithms
Vienna University of Technology, Austria

†AGFA HealthCare
Vienna, Austria

ABSTRACT

One of the main obstacles in integrating 3D volume visualization in the clinical workflow is the time-consuming process of adjusting parameters such as viewpoint, transfer functions, and clipping planes required to generate a diagnostically relevant image. Current applications therefore make scarce use of volume rendering and instead primarily employ 2D views generated through standard techniques such as multi-planar reconstruction (MPR). However, in many cases 3D renditions can supply additional useful information. This paper discusses ongoing work which aims to improve the integration of 3D visualization into the diagnostic workflow by automatically generating meaningful renditions based on minimal user interaction. A method for automatically generating 3D views for structures in 2D slices based on a single picking interaction is presented.

Index Terms— Volume visualization, viewpoint selection

1. INTRODUCTION

Recent years have seen great improvements in the quality, accuracy, and performance of volume visualization algorithms. However, their use is still limited to relatively few applications. The prime reason that prevents a broad usage of 3D visualization in clinical routine is the considerable effort required to generate a relevant image. Volume rendering algorithms commonly feature many parameters which are often unintuitive even for visualization experts. In order to permit a wider usage of these techniques, the complexities need to be eliminated and volume visualization has to be tightly integrated into the diagnostic workflow.

In this paper, we discuss recent developments and strategies to achieve these goals [1, 2]. We present an approach which attempts to automatically generate a good 3D view for

a structure in a slice triggered by a simple picking interaction. The user just presses a hot-key and moves the mouse pointer over the area of interest and our algorithm derives all required rendering parameters. This procedure requires little additional effort and supplies the clinician with a synchronized 3D view instantly. This concept allows a non-invasive integration of 3D volume rendering into the diagnostic workflow, as no additional parameters have to be specified and no new interaction metaphors have to be learned by the user.

2. METHODOLOGY

A common scenario where medical applications could benefit from volume rendering techniques is the closer inspection of suspicious regions. Vessels, fractures, or tumors can have complex spatial structures and may be difficult to assess using two-dimensional slices. A 3D view might provide additional information leading to a better or faster judgement. The process of finding a good corresponding 3D view for a structure in a slice, however, is not trivial as 3D navigation is a difficult task. To make matters worse, structures of interest may be occluded by other objects which need to be removed while still retaining sufficient anatomical context. Thus, manually specifying a good 3D view for structures detected in cross-sectional images is a prohibitively time-consuming task in many cases. The user has to edit several parameters such as camera position and clipping planes in order to get an expressive visualization. This process has to be repeated for each new structure of interest.

Our approach attempts to simplify this procedure by automatically generating good views for interactively picked structures in 2D slices. This functionality can be activated on demand by pressing a hot-key while pointing the mouse over a structure of interest. Several heuristics are used to estimate viewpoint, zoom factor, clipping planes, and transfer function from this single interaction. The 3D view is instantly updated to present the picked structure to the user. Depending on the quality of the generated result, the user can manually

The work presented in this paper has been funded by AGFA HealthCare in the scope of the DiagVis project.




Fig. 1. The viewing sphere which is generated for the local shape estimation is deformed according to the major volumetric extent of the structure.

fine-tune the 3D view by adjusting the view parameters. This concept allows an efficient and non-intrusive integration of 2D and 3D visualizations in medical workstations.

2.1. Viewpoint Selection

Previous approaches for automatically choosing viewpoints for volumetric data have mostly focussed on finding globally optimal views under a given transfer function, such as the approaches by Bordoloi and Shen [3] and Takahashi et al. [4]. Chan et al. [5] concentrated on viewpoint selection specifically for angiography data sets. We attempt to find a good view for a specific feature identified by a single picking interaction on a slice. A number of different factors influence the choice of a good viewpoint for a particular region of interest. These factors include the shape of the structure, its spatial location with respect to other objects, the viewpoint history, and the preferred orientation for the data set. In order to represent these criteria in a compact manner, a spherical parametrization of the viewing space is chosen: Spheres centered around the picked structure encode the quality of all viewpoints with respect to specific criteria. Deformation of these viewing spheres are employed to represent variations in the viewpoint quality. The following criteria are used in our viewpoint selection process:

Feature Shape. An important input parameter for viewpoint selection is the local shape of the structure of interest. For example, if the picked point is part of a blood vessel, a good viewpoint shows the path of this vessel and does not cut through it. With a fast local segmentation and a principal component analysis the shape information can be derived from the data values surrounding the picked point. We employ a segmentation algorithm based on an approach presented by Huang and Ma [6]. The neighborhood of the seed point is analyzed regarding scalar value and gradient magnitude distributions to initialize the parameters of the cost function. If the seed is located near a boundary the growing captures the boundary of the object while a seed point within a homogeneous area results in a more compact growing process. The growth is influenced by the spatial




Fig. 2. Starting from the picked position, visibility rays are cast to points which are equally distributed on the surface of the viewing sphere. Samples along the rays are analyzed to detect when it exits the structure of interest and at which position the object becomes occluded by other structures.

distribution of the points marked as region members so far. Growing progresses until the object-oriented bounding box of the included voxels reaches a certain limit. During the growing process the bounding box is updated at variable intervals which are computed to estimate how many more voxels may be added until the limit is reached. If just the boundary of a structure is segmented, a lower number of voxels is sufficient to estimate the feature shape than if the growing is performed in a homogeneous area. Thin structures, like blood vessels, for instance, require only a small number of object voxels for a good shape estimation. As soon as the region growing process terminates, a principal component analysis is performed on the member voxels to extract the three feature vectors and the corresponding eigenvalues. A metric of Westin et al. [7] is used to measure the local shape of the segmented feature from the relation between the eigenvalues. This metric allows to classify if a structure has an isotropic, a planar or linear shape. According to the local shape of the object, the viewing sphere has to be deformed as illustrated in Figure 1. If the object has a volumetric extent (blob), then basically all viewpoints are of the same quality (left). For a planar structure (sheet) the viewpoints which are orthogonal to the sheet are favored (middle). If a tubular structure (line) is determined, the preferred viewpoints are aligned along a ring which is orthogonal to this line (right).

Feature Visibility. A further building block for estimating a good viewpoint is the visibility of the structure of interest. To determine whether a certain viewpoint provides good visibility of the selected structure, rays are cast from the picked point. As a local segmentation was performed for the local shape estimation, this information is utilized to determine when a ray exits the tissue of interest. The opacity is accumulated along the ray and as soon as a small opacity threshold is exceeded,




Fig. 3. Automatic viewpoint selection process. Left: Picking position on the slice. Right: Automatically generated 3D view.

the ray is terminated. Even if the structure of interest is occluded by other objects, a higher visibility value is assigned to a viewpoint if there is more space between the structure and the occluding objects as this space allows for a more flexible placement of a clipping plane. This process is illustrated in Figure 2.

Viewpoint History. As a specific view was selected by the system based on estimated demands of the user, the current viewpoint will also be considered for the estimation of the quality of the next viewpoints. In particular, big shifts of the viewpoint for two successive pickings should be avoided. This means that if there is a good viewpoint for the picked structure close to the current one, this viewpoint is preferred to others which are positioned farther away on the viewing sphere.

Patient Orientation. For many types of medical examinations there is a preferred orientation. Our approach also incorporates this information. The head-feet axis serves as a rough estimation to derive the preferred viewpoints. The viewing sphere is deformed in a way that it prefers viewpoints which are orthogonal to this axis.

Each of these components is represented by a deformed viewing sphere. In a final step the individual spheres are combined and weighted. The best viewpoint from this final deformed sphere is then chosen. Figure 3 shows a result generated using this process.

As picking is performed only for a single point, the user is not required to indicate the size of the feature of interest. Instead, a simple interactive strategy is used to control the zoom factor of the generated view: When the user performs the initial picking operation, the zoom factor of the 3D view is derived from the zoom specified in the slice view. If the hotkey is kept pressed, the local region growing proceeds and the zoom is continuously updated to reflect the increased extent of the segmented region's bounding box. Thus, the 3D view depicts a progressive zoom-out on the picked point and the user may stop this process at any time by releasing the hotkey.

2.2. Transfer Function Adjustment

Most volume rendering algorithms rely on a transfer function which maps data values to visual attributes such as colors and opacities. The manual definition of transfer functions, however, is a non-trivial process and frequently relies on a tedious trial-and-error process to generate an acceptable result. While there has been significant progress in the automatic specification of transfer functions [8, 9], the problem remains unsolved for the general case. Approaches like the work of Rezk-Salama et al. [10] attempt to simplify the user-interface for defining a transfer function. However, the procedure is still likely to be too complex for real-world medical applications. Most medical workstations provide color and opacity presets tailored for specific types of examinations. The transfer function is defined by a color look-up table and linear opacity ramp specified by a window/level mechanism. Inspired by the definition of a transfer function based on partial region growing in the work by Huang and Ma [6], knowledge about the distribution of scalar values within the extracted object can be utilized to fine-tune an existing transfer function. The adjustment is based on the mean value and the standard deviation of the structure of interest's scalar values. An example is shown in Figure 4.

2.3. Clipping Plane Specification

As transfer functions have global influence, they are frequently unable to reveal certain features without simultaneously causing occlusion. Thus, medical workstations typically provide view-aligned and object-aligned clipping planes to resolve this conflict. Based on the visibility calculation described in Section 2.1, our approach sets clipping planes to remove objects occluding the picked structure. The near clipping plane is placed along the visibility ray where the structure of interest is not yet occluded by other structures. This preserved the largest possible context while still avoiding occlusion. As view-aligned clipping planes may not be suited for further manual adjustment of the viewpoint, alternatively object-aligned clipping planes along the three main patient axes can be generated by our system. The six available clipping planes can clip from left, right, anterior, posterior, head and feet. To decide which of these planes has to be set to allow an unoccluded view on the structure of interest, the plane which is most perpendicular to the viewing direction has to be identified. Then this plane is placed at the clipping position along the visibility ray. An object-aligned clipping plane sticks to its position when the viewpoint is changed manually and cannot clip away the structure of interest unintendedly.

3. RESULTS

An informal evaluation of the proposed approach was performed. An experienced radiology technician was asked to generate diagnostically relevant volume renderings of




Fig. 4. Transfer function adjustment and clipping. Left: Picking position on the slice. Middle: The structure of interest is not visible with the current setting of the opacity transfer function. Right: Automatically adjusted transfer function makes the vessel visible.

pathologies in different data sets. This task was first performed manually and then using the proposed method. For all data sets the overall expenditure of time ranged from 5 to 20 minutes in the manual case and from 2 to 10 minutes using our method. The conclusion of this initial evaluation is that the proposed approach significantly reduces the effort to localize pathologies in 3D views. In some cases manual adjustments were still required to generate an optimal image. However, the overall burden to benefit from the potential additional information provided by 3D volume visualization was considerably reduced. These initial results form a good basis for future extensions and evaluations of our approach.

4. CONCLUSION

In order to benefit from volume visualization techniques, applications need to integrate them in a non-invasive manner. Keeping existing interaction patterns intact allows the user to incorporate 3D visualization into their workflow. We described a basic approach how this can be accomplished. Initial tests indicate that our method significantly reduces the effort of finding good 3D views for structures of interest in 2D slices. In future work, we plan further evaluations of the benefits of this approach in clinical settings.

5. REFERENCES

- [1] P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller, “LiveSync: Deformed viewing spheres for knowledge-based navigation,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1544–1551, 2007.
- [2] P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller, “The LiveSync interaction metaphor for smart user-intended visualization,” Tech. Rep. TR-186-2-08-01, Institute of Computer Graphics and Algorithms, Vienna University of Technology, 2008.
- [3] U. D. Bordoloi and H.-W. Shen, “View selection for volume rendering,” in *Proceedings of IEEE Visualization 2005*, 2005, pp. 487–494.
- [4] S. Takahashi, I. Fujishiro, Y. Takeshima, and T. Nishita, “A feature-driven approach to locating optimal viewpoints for volume visualization,” in *Proceedings of IEEE Visualization 2005*, 2005, pp. 495–502.
- [5] M.-Y. Chan, H. Qu, Y. Wu, and H. Zhou, “Viewpoint selection for angiographic volume,” in *Proceedings of the International Symposium on Visual Computing 2006*, 2006, pp. 528–537.
- [6] R. Huang and K.-L. Ma, “RGVis: Region growing based techniques for volume visualization,” in *Proceedings of Pacific Graphics 2003*, 2003, pp. 355–363.
- [7] C.-F. Westin, A. Bhalerao, H. Knutsson, and R. Kikinis, “Using local 3D structure for segmentation of bone from computer tomography images,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition 1997*, 1997, pp. 794–800.
- [8] G. Kindlmann and J. W. Durkin, “Semi-automatic generation of transfer functions for direct volume rendering,” in *Proceedings of the IEEE Symposium on Volume Visualization 1998*, 1998, pp. 79–86.
- [9] F.-Y. Tzeng and K.-L. Ma, “A cluster-space visual interface for arbitrary dimensional classification of volume data,” in *Proceedings of VisSym 2004*, 2004, pp. 17–24.
- [10] C. Rezk-Salama, M. Keller, and P. Kohlmann, “High-level user interfaces for transfer function design with semantics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 1021–1028, 2006.