Eurographics/ IEEE-VGTC Symposium on Visualization (2007)
Ken Museth, Torsten Möller, and Anders Ynnerman (Editors)

Feature Emphasis and Contextual Cutaways for Multimodal Medical Visualization

Michael Burns[1,2], Martin Haidacher[1,3], Wolfgang Wein[1], Ivan Viola[4] and M. Eduard Gröller[3,4]

[1]Imaging & Visualization Department, Siemens Corporate Research, Princeton, NJ USA
[2]Computer Science Department, Princeton University, USA
[3]Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria
[4]Department of Informatics, University of Bergen, Norway

Abstract

Dense clinical data like 3D Computed Tomography (CT) scans can be visualized together with real-time imaging for a number of medical intervention applications. However, it is difficult to provide a fused visualization that allows sufficient spatial perception of the anatomy of interest, as derived from the rich pre-operative scan, while not occluding the real-time image displayed embedded within the volume.
We propose an importance-driven approach that presents the embedded data such that it is clearly visible along with its spatial relation to the surrounding volumetric material. To support this, we present and integrate novel techniques for importance specification, feature emphasis, and contextual cutaway generation.
We show results in a clinical context where a pre-operative CT scan is visualized alongside a tracked ultrasound image, such that the important vasculature is depicted between the viewpoint and the ultrasound image, while a more opaque representation of the anatomy is exposed in the surrounding area.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation - Display algorithms

1. Introduction

While already an active area of research for many years, three-dimensional visualization of medical imaging data is still emerging in clinical settings. Medical imaging devices continue to improve and become more widely available while digital displays emerge from the radiology reading room into interventional settings and even mobile devices. One well established technique is 3D volume rendering of CT and CTA (Computed Tomography Angiography) scans. While many physicians still prefer cross-sectional slice images for diagnosis and interpretation of the data, volume rendering can provide a better global spatial impression of the anatomy. This becomes especially important for fusion of CT/CTA with interventional imaging modalities.

In medical procedures such as biopsy and radio frequency ablation, the insertion path of a needle must be carefully planned on the CT scan, in order to avoid critical structures in the patient. During a procedure, ultrasound is often used as an interventional imaging modality to assist in navigating needles to their correct locations. Medical ultrasound mainly depicts borders between various tissue types in a 2D plane oriented from the transducer into the patient's body. If a tracking system is used to locate the ultrasound probe in 3D, and the pre-operative CT scan is aligned (i.e. registered) correctly with the patient coordinate system, volume rendering of the CT data can be merged with a real-time view of the ultrasound plane. Dense 3D information from the CT data can help the physician relate both the needle and ultrasound plane to the critical anatomical structures. Furthermore, planning information like ablation target volumes and margins, optimal needle path, etc., can be visualized.

The inherent problem here is that it is difficult to show enough data from the dense pre-operative scan without occluding the ultrasound image plane. Motivated by this problem, we present a new approach for visualizing an object of interest embedded within volumetric data. This approach can be broken down into four components, as described in

Section 3. At the core is the specification of importance and implicit ranking of materials within the volume by augmentation of the conventional transfer function (3.1). To facilitate visualization steered by this importance specification, we emphasize rendered materials by their assigned importance, performing importance-driven shading (3.2). We then define a flexible, view-dependent cutaway structure (3.3) and use this definition to cut away occluding material based on its importance (3.4). As such, we are able to clearly display the object of interest embedded in the volume while maintaining view of materials deemed especially important and progressively trimming away material of lesser importance. We conclude by presenting results which show how this approach can be used in our motivating application, visualization of multimodal medical data in an interventional setting.

2. Related Work

Our work aims at one of the main challenges in visualization research—to satisfy needs of a specific diagnostic and treatment environment. The challenge is to provide a clear understanding of complex data and guide the user to the most relevant information. The relevance of features in the data can be assigned in many different ways, i.e., from user steered segmentation to a fully automatic process. Relevance then serves as the controlling parameter for assigning visual representations among those features. Combination of dense (visually prominent) representation of the most relevant features with sparse (visually suppressed) information about other features is often denoted as focus+context visualization [Hau05]. Focus, i.e. the most relevant feature, is represented very densely and context is presented sparsely to indicate overview information.

Several publications have addressed the issue of visually emphasizing features in volumetric renderings. In visualization of volumetric scalar data, two-level volume rendering uses segmentation information to render objects in the data with different compositing and rendering techniques [HMBG01]. In visualization of 3D flow data, a user-specified degree of interest function (DOI) affecting optical properties is shown to help visualize important flow features [HM03]. Like our shading-driven feature emphasis, these techniques apply globally to the visualization. However, they do not use a scalar importance value derived from a transfer function to modulate shading as our technique proposes.

Other works incorporate viewpoint information to modify optical properties of a context region occluding a focus region, and are often inspired by traditional illustration techniques [VG05]. Context-preserving volume rendering [BGKG06] and opacity peeling [RSK06] generate visualizations by ghosting and fading outer layers of material to reveal inner structures. Clearview [KSW06] uses similar techniques while incorporating user specification of a focus region. Special peel-away [CSC06] and exploded visualizations [BG06] have been proposed to expose segmented or user-defined regions of a volume. Cutaways for polygonal rendering [DWE03] and hidden iso-surface display [FBS05] also address these issues.

In general, however, these works do not incorporate an importance assignment mechanism so object emphasis and suppression is harder to specify. In that respect, our work is most similar to importance-driven visualization, which uses data segmentation and relevance to automatically generate cutaway visualizations [VKG05]. However, we do not use explicit segmentation, as the importance classification is done in the transfer function space. Furthermore, we extend and formalize the existing cutaway algorithms to provide more options for display of contextual material.

Much visualization research has been done regarding medical imaging visualization, where cut-away and smart visibility techniques have been applied to the visualization of peripheral arteries in lower extremities [SvC*04] and neck dissection planning for enlarged lymph node removal [KTH*05]. However, our work is the first to develop and apply advanced 3D volumetric rendering techniques to the problem of visualizing live ultrasound data within the context of a 3D volume.

There are two major challenges associated with multimodal medical data. The first challenge is the data registration. We register the CT (or MRI) data with the ultrasound data as described in previous work [WRN05]. The second challenge is how to combine multiple modalities in a single visualization. Magic mirrors have been applied for multimodal visualization of the human brain [KDG99]. The combination of multiple modalities or properties through smooth brushing and linking of scatter plots has also been presented [DGH03]. The product of this multimodal selection is mapped to a one-dimensional DOI classification that is directly mapped to a visual representation. In our work we use multiple modalities for selection of a *feature*. According to the needs of the interventional setup, we use importance parameterization as a steering mechanism for the combination of data from different modalities, meaning that in our visualizations multiple modalities are and need to be explicitly present.

3. Feature Emphasis and Contextual Cutaways

This section describes the pipeline for feature emphasis and contextual cutaway visualization (Figure 1). The top three nodes with the sharp edges are the input sources and the subsequent nodes are processing steps described in the following subsections. The main inputs are the volumetric data and the object of interest, which occupies a region within the space of the volumetric data set. In terms of our motivating application, the CTA data set and the ultrasound plane fill these roles, respectively. The main goal of our approach is to have the object of interest, which may be occluded by ordinarily rendered volumetric data, visible yet shown in the context of the volumetric data.

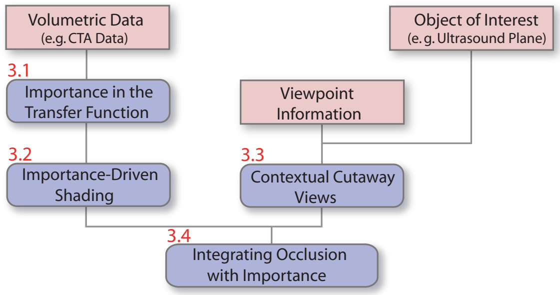

Figure 1: *Overview of the pipeline for feature emphasis and contextual cutaway visualization.*

In the left pipeline, the standard color and opacity transfer function is augmented with an importance value which specifies the relative importance of each regional component of the transfer function (3.1). We then introduce a shading modification that is used to emphasize features based on their importance value (3.2).

In the right pipeline, the geometry of the object of interest and viewpoint information are used to create a cutaway structure, which divides space into different contextual regions and defines a function for determining occlusion (3.3).

Finally the importance of sample points and their occlusion values are used to change their visibility. In section 3.4 we describe the visual effect of both components on the final appearance of a rendered sample.

3.1. Defining Importance via Transfer Function

As previously stated, we wish to create a visualization which displays materials differently based on their relative importance. In order to do so, we must establish a method for assigning importance values to materials in the volumetric data. Many previous approaches have done this by using an auxiliary segmentation volume or a geometric segmentation representation to explicitly assign importance values to 3D space. In our approach, we create an augmented transfer function to classify materials and assign to them an importance value, much as they are assigned a color and opacity. This allows us to avoid pre-computation time and reduce user interaction, since a transfer function that adequately classifies materials can be created once for a particular visualization goal and imaging system.

A conventional transfer function allows the determination of a color and opacity for samples of a volume based on local properties of the sample, such as data intensity value, gradient magnitude, or curvature. We denote a feature vector of these properties as $\vec{x}$. Such a transfer function can be composed of several transfer function components, possibly overlapping, each of which defines a color and opacity ramp over a specified region in the transfer function space (spanning $\vec{x}$). If each component corresponds to a particular material in the data and is assigned a distinct color, then materials can be visually distinguished in the volume rendering.

We augment the conventional transfer function by adding an importance value to each transfer function component. As transfer function components may overlap, we use the importance value to blend components for a particular feature vector as follows, where N_{comp} is the number of components and $\vec{c}_i$, α_i, and I_i are the RGB color, opacity, and importance value of the i'th component, all evaluated at $\vec{x}$.

$$g(\vec{x}) = \left(\frac{\sum_{i=1}^{N_{comp}} \vec{c}_i * \alpha_i * I_i}{\sum_{i=1}^{N_{comp}} \alpha_i * I_i}, \max_{i=1..N_{comp}} \alpha_i, \max_{i=1..N_{comp}} I_i \right)$$

This blending is done once before the rendering and results in a lookup-table which assigns a quintuple consisting of RGB color values, an opacity value, and an importance value, to each point in the transfer function space.

3.2. Feature Emphasis with Importance-Driven Shading

Our goal of emphasis is to guide the viewer quickly to the most important parts of an image, which we can achieve by rendering more important materials with more detail. The importance value of a sample point from the volume, as determined by $g(\vec{x})$, can be used to modify its rendered optical properties, with the goal of emphasizing samples corresponding to important materials. In general, emphasis can be made by color, opacity, and shading style. Since the transfer function already allows us to specify color and opacity directly, there is no need to derive those from importance. However, by modulating shading by importance, we can de-emphasize unimportant materials.

We do this by using the importance value to limit the contribution of a given shading model to the output color, so that materials with low importance are shaded less than those of high importance. We denote the strength of this effect, *emphasis*, by the global parameter E. We then calculate the output color as follows, where $\vec{c}_{shaded}$ is the color value calculated with a given shading method, $\vec{c}$ is the raw color value from the transfer function, and I is the importance value derived for the sample from the transfer function.

$$\vec{c}_{sample} = \vec{c}_{shaded} * (1 - E * (1 - I)) + \vec{c} * E * (1 - I)$$

With no emphasis ($E = 0$), all materials are shaded as normal. With maximum emphasis ($E = 1$), the equation is a linear interpolation between shaded and unshaded colors based on importance, such that important materials are fully shaded and unimportant materials are unshaded. The choice of the emphasis value E depends on the visualization scenario and how much one can afford to discard shading information for low-importance materials. The images in Figures 2 and 3 show how importance-driven shading can emphasize important features.

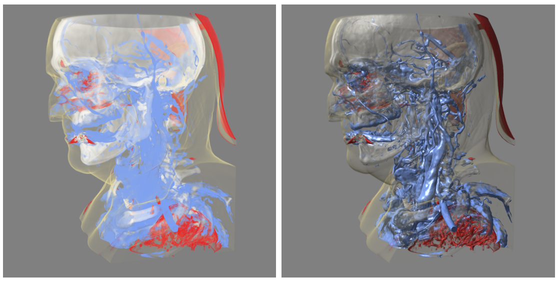

(a) No shading (b) Phong shading, $E = 0$

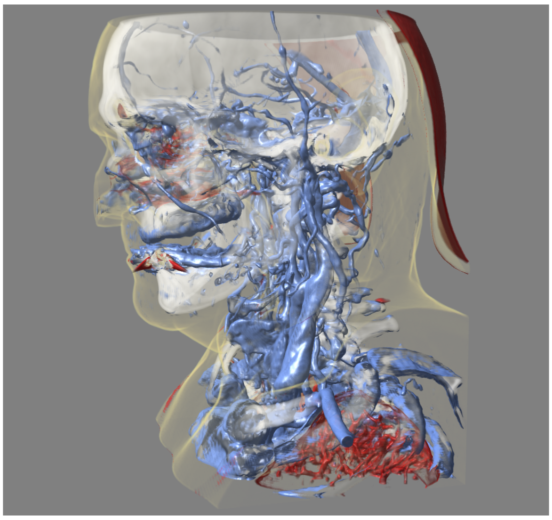

(c) Phong shading, $E = 1$

Figure 2: *No shading (a) results in low contrast and spatial distinction. With no emphasis (b), vessels in the skull are obscured by shading details in the bone. With full emphasis (c), shaded and unshaded samples are selectively combined to accentuate important materials (vessels) and suppress shading detail of non-important materials (skin, bone).*

3.3. Contextual Cutaway Views

As stated earlier, our goal is to present the object of interest in such a way that the material around it is cut away based on the importance of the material. In this section we present a flexible cutaway structure that ultimately allows us to trim away materials of different importance at different levels of the cutaway structure.

The simple cutaway view definition partitions space into two distinct regions: the area inside the cutaway, which we denote the *clear* region, and everything else, which we denote the *base*. We formalize this by defining an *occlusion function*, denoted Ω. This occlusion function will represent the degree to which a point in space occludes the object of interest. In the simple cutaway scenario, at a given point, $\Omega = 1$ if the point is inside the *clear* region and 0 if it is inside the *base* region. For a given ray through the volume in eye-space, let $\xi(\theta)$ be the depth (z component in eye-space) where the ray intersects a cutaway suface with angle θ. We can then define Ω for a given point p along the ray as follows, where p_z is the z component of p and $step(a,x) = 0$ if $x < a$ and 1 if $x \geq a$.

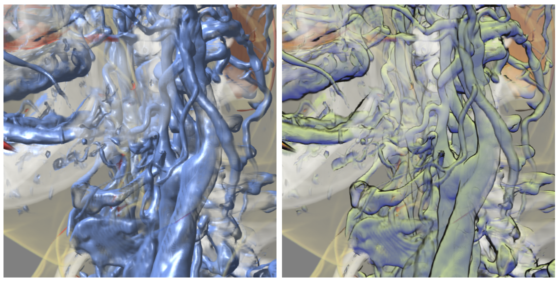

(a) Phong shading, $E = 1$ (b) Gooch shading, $E = 1$

Figure 3: *With full emphasis, importance-driven Phong shading (a) and Gooch cool-to-warm shading [GGSC98] with silhouette enhancement (b) help provide visual distinction between vessels and bone.*

$$\Omega = step(\xi(\theta), p_z)$$

This binary definition suggests rendering can have only two modes: sparse (for the *clear* region) and dense (for the *base* region). In order to give more control over the rendering of materials with multiple importance values, we propose a new cutaway definition where occlusion values vary smoothly between 0 and 1 over 3D space.

We begin by modifying the simple cutaway definition to include a second cutaway surface defined by a wider angle. This new region, which we denote the *transition* region, can have an occlusion function that varies between the two cutaway surfaces. We define this as shown in Figure 4, allowing us to determine the cutaway angle of points located in the *transition* region, relative to the two bounding angles θ_1 and θ_2. This will ultimately allow variation of visibility in the image outside the projected object-of-interest silhouette by letting us cut or fade away materials at varying angles.

To control visibility of the data in front of the object-of-interest, we add another region, the *overlay* region. This region is bounded by the cutaway surface of θ_1, offset a thickness d toward the camera, as shown in Figure 4.

Considering these four regions, we define the occlusion function for a given point in eye space as follows, where θ_1 and θ_2 are the cutaway angles, d is the thickness of the *overlay* region, and $ramp(a,b,x) = 0$ if $x < a$, 1 if $x \geq b$, and $(x-a)/(b-a)$ for $a \leq x < b$.

$$\Omega = \frac{ramp(\xi(\theta_2), \xi(\theta_1), p_z) + ramp(\xi(\theta_1), \xi(\theta_1) + d, p_z)}{2}$$

This definition results in $\Omega = 0$ for points in the *base* area,

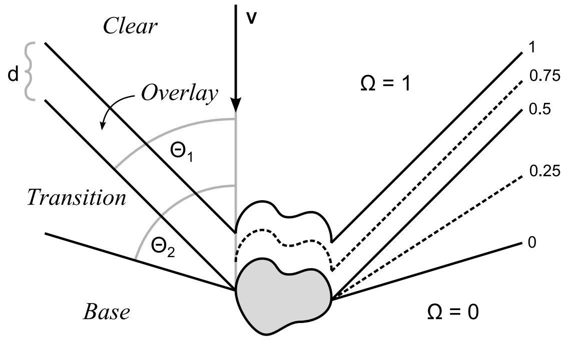

Figure 4: *Given an object of interest (shaded) and a viewing direction, we can construct a cutaway with distinct* base*,* transition*,* overlay*, and* clear *regions, as shown in this cross section diagram. An occlusion function* Ω *can then affect visibility of materials in different areas of the cutaway structure. In the traditional simple cutaway,* $\theta_1 = \theta_2$ *and* $d = 0$*.*

$\Omega = 0.5$ for points on the *transition-overlay* boundary, and $\Omega = 1$ for points in the *clear* area, with the appropriate linear blends for points in between the boundaries, as shown in Figure 4. In the next section, we use this definition of the occlusion function to give control over how materials of varying importance are rendered in the cutaway structure.

3.4. Integrating Occlusion with Importance

Now that we have a flexible cutaway structure, we can use the occlusion function and importance values to render various components of the transfer function in different regions of the cutaway structure. Our goal is to have only materials with the very highest importance rendered in the *clear* region, and all materials rendered in the *base* region. We then wish to fade out materials based on their importance such that materials of low importance are cut away at a wider angle than materials of higher importance in the *transition* region. We also wish to allow some materials of moderately high importance to be present in the *overlay* region, where they will be faded out before the *clear* region.

In order to accomplish this, we modify the opacity of a sample based on the occlusion value at that point. In order to control fading, one needs to establish two occlusion value thresholds, τ_l and τ_u, between which points will be faded. Given these thresholds, we can modify opacity by the following formula, where α is the alpha component of a sample computed from the transfer function for a given point, α' is the alpha component to be used in rendering, and Ω is the occlusion value at the point.

$$\alpha' = \alpha * (1 - ramp(\tau_l, \tau_u, \Omega))$$

This allows us to control how rendered samples are faded in the cutaway structure. Note that by choosing a constant set of τ such as $(0,1)$, materials would be completely transparent in the *clear* region, and opacity would be unchanged in the *base* region, with a uniform fade in the *overlay* and *transition* regions based directly on the occlusion value. In this scenario, the importance value of a sample would not have any effect on the fade.

In order to incorporate importance into this formula, we derive thresholds for each sample based on its importance value I from the transfer function as follows.

$$\tau_u = I$$
$$\tau_l = max(2 * I - 1, 0)$$

Such calculation results in several desirable effects. First, materials of highest importance ($I = 1$) are visible everywhere. Materials of lesser but high importance have sharp fades in the *overlay* area so their cutaway boundaries are easily visible. Additionally, materials of moderate importance are faded gradually within the *transition* area. Finally, materials of low importance are faded near the *base* region.

The effects of the importance values and the cutaway parameters are explored in Figure 5. Here, we have a CT scan of an abdomen and our object of interest is a plane textured with the volume data and embedded in the volume. The transfer function is defined such that the importance values for skin and flesh (brown, red) are lowest (0.1, 0.01), vasculature and bone (blue, white) are highest (0.99), and other organs (yellow) are in between (0.5). In Figure 5(a), we can see the effect of the simple cutaway scenario, where the *transition* and *overlay* regions are nonexistent and a hard boundary separates the *base* and *clear* regions. This result makes no use of the importance values, and our most important materials are not visible.

This is remedied in Figure 5(b), where we see the effect of having a large *overlay* region, with *d* sufficiently large so that none of the vasculature (blue) and bone is faded away. This allows us to view the important materials but maintains the hard cutaway boundary, not allowing for any difference in visibility between the organs (yellow) and the flesh (red).

By setting θ_2 larger than θ_1, this distinction can be made, as we see the skin and flesh cut away at a wider angle and the organs fading into the *transition* region in Figure 5(c). We see this effect most strongly in the bottom of the image, as the complex yellow material is no longer obscured by the flesh and skin. This is easily achieved by having a cutaway with multiple angles, as established by the *transition* region.

If our goal is to view the textured plane with less obstruction, then we can set the thickness of the *overlay* region to a much smaller value, so that vasculature (blue) is only visible just in front of the plane (Figure 5(d)). This allows us to see the important material in context with the object of interest without displaying so much that the object of interest is overly obscured.

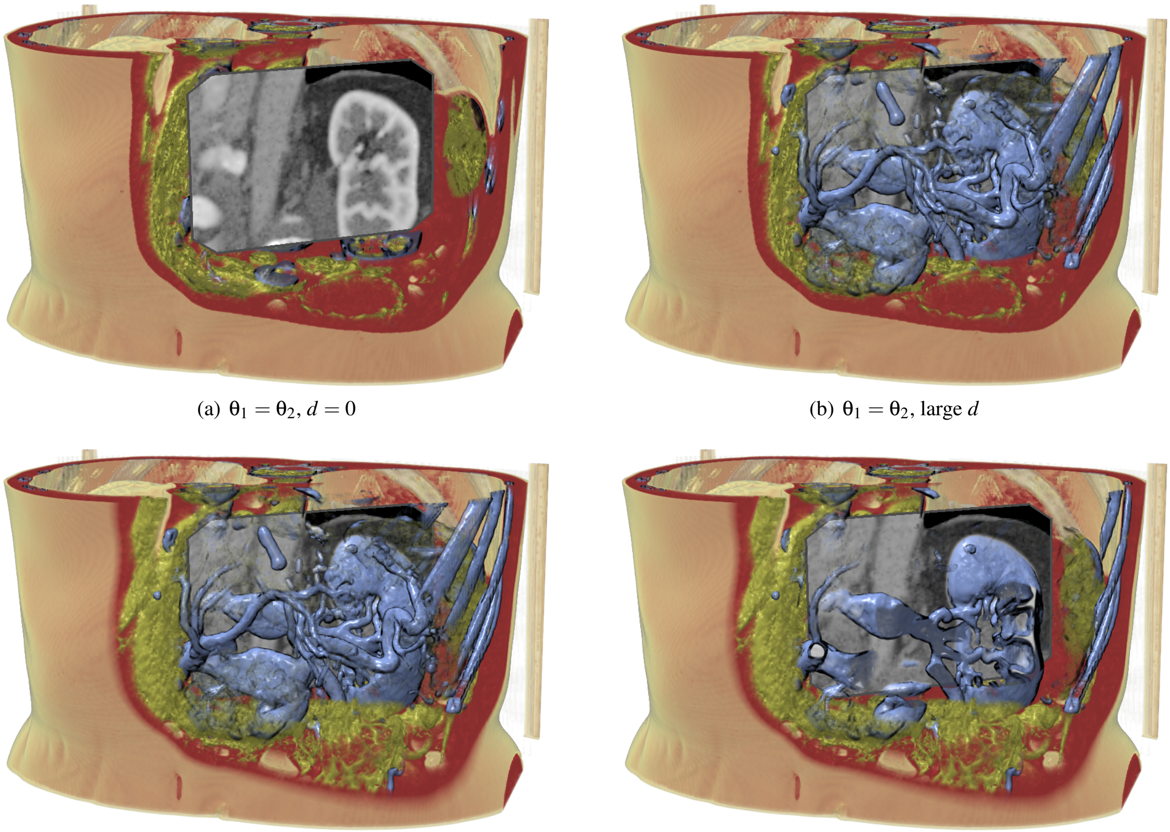

(a) $\theta_1 = \theta_2$, $d = 0$

(b) $\theta_1 = \theta_2$, large d

(c) $\theta_1 < \theta_2$, large d

(d) $\theta_1 < \theta_2$, small d

Figure 5: *The effects of various parameters in the cutaway structure are shown in the above images. Note that (a) and (b) are traditional cutaway views but (c) and (d) expose more contextual information by cutting away less important materials at wider angles. A smaller overlay thickness in (d) causes the embedded object to be less obscured.*

4. Application to Multimodal Medical Data

We use ultrasound exams recorded on patients and volunteers, where the spatial location and orientation of the transducer has been tracked using a magnetic position sensing system. The ultrasound sweeps have been manually registered (i.e. spatially aligned) to the tomographic 3D scans, aided by automatic techniques [WRN05]. To show interventional ultrasound in the context of organ vasculature, we use an early arterial phase CTA scan from a patient's liver. The injected contrast agent causes the vasculature to show up with high intensities in the scan, easily distinguishable from surrounding tissue. Figure 6(a) depicts our proposed rendering for a longitudinal ultrasound image of the liver and the corresponding CTA data. Note that in front of the ultrasound plane, only the vessels are rendered due to their high importance value $I = 1$, allowing a clear view of the ultrasound information as well. Figure 6(b) shows a similar rendering on a transversal ultrasound plane, where a sharper transition of the cutaway cone is used. In figure 7 a similar visualization is applied using a regular CT scan without contrast agent. In this case, bone is given a high importance and is allowed to obscure the ultrasound plane.

Magnetic Resonance Imaging (MRI) is generally less favorable for volume rendering because intensities are not proportional to tissue density (making it difficult to define opacity based on it) and bone is not distinguishable. Using the proposed techniques, we still achieve results that we believe are suited well for interventional visualization. Figure 8 depicts the fused renderings of a liver MRI scan together with a registered ultrasound image.

In all data sets, the ultrasound images are made clearly visible and their spatial context within the 3D volume is well defined due to the view dependent cutaway structure. In addition to the global 3D relation, critical anatomical structures such as liver vasculature that must not be punctured in an interventional scenario can be visualized as well. This is important for image-guided needle procedures and can increase the acceptance of multimodal visualization for such proce-

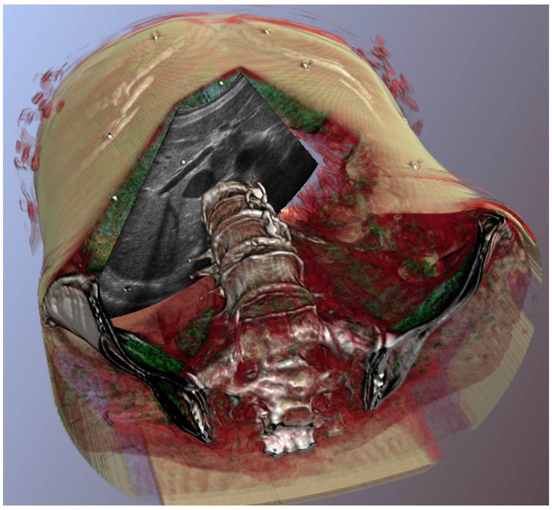

Figure 7: *A transversal ultrasound image of a patient's liver is shown aligned in the context of a (non-contrasted) CT scan. Bone is allowed to obscure the ultrasound plane and as such helps facilitate spatial perception of the plane.*

dures, as the physician always has a concise view of all the needed information, regardless of the ultrasound probe position. It is straightforward to integrate planning information, such as ablation target volumes and needle path, using the visualization tools that we have proposed. We note, however, that our visualizations are meant to aide the doctor in exploring the CT data and relating the position of the ultrasound scan with the CT data and would not be used on their own for clinical purposes. In practice such images would be presented alongside views of the 2D ultrasound data.

Finally, the presented techniques have been designed for and tested with a hardware-based GPU raycaster. As such, the system can operate at interactive frame rates, which is necessary for use in a clinical scenario.

5. Conclusion

We have presented a suite of new techniques for visualizing objects of interest embedded within volumetric datasets. These techniques allow for emphasis of important volumetric features as defined by transfer functions as well as visualization of contextual information relative to the object of interest by means of a new, flexible cutaway structure. We have applied our techniques to the area of image-guided needle procedures, and show fused visualizations of ultrasound with CTA, CT and MRI volumes.

Contemporary radiology software provides a built-in set of transfer functions adapted to various imaging modalities, which doctors use to visualize patient data and distinguish various tissues and organs. As our algorithm classifies materials in the transfer function space, it would be reasonable to expect that these transfer functions could be augmented with importance information for various medical scenarios. As such, we believe that the presented techniques can have a positive impact on the medical visualization community.

Acknowledgments

We would like to thank Gianluca Paladini of Siemens Corporate Research for sponsoring this project and initiating our collaboration.

References

[BG06] Bruckner S., Gröller M. E.: Exploded views for volume data. *IEEE Trans. on Visualization and Computer Graphics 12*, 5 (2006), 1077–1084.

[BGKG06] Bruckner S., Grimm S., Kanitsar A., Gröller M. E.: Illustrative context-preserving exploration of volume data. *IEEE Trans. on Visualization and Computer Graphics 12*, 6 (2006), 1559–1569.

[CSC06] Correa C., Silver D., Chen M.: Feature aligned volume manipulation for illustration and visualization. *IEEE Trans. on Visualization and Computer Graphics 12*, 5 (2006), 1069–1076.

[DGH03] Doleisch H., Gasser M., Hauser H.: Interactive feature specification for focus+context visualization of complex simulation. In *Proc. of VisSym'03* (2003), pp. 239–248.

[DWE03] Diepstraten J., Weiskopf D., Ertl T.: Interactive cutaway illustrations. In *Proc. of Eurographics '03* (2003), pp. 523–532.

[FBS05] Fischer J., Bartz D., Strasser W.: Illustrative display of hidden iso-surface structures. In *Proc. of IEEE Visualization '05* (October 2005), pp. 663–670.

[GGSC98] Gooch A., Gooch B., Shirley P., Cohen E.: A non-photorealistic lighting model for automatic technical illustration. In *Proc. of SIGGRAPH'98* (1998), pp. 447–452.

[Hau05] Hauser H.: *Scientific Visualization: The Visual Extraction of Knowledge from Data*. Springer, 2005, ch. Generalizing Focus+Context Visualization, pp. 305–327.

[HM03] Hauser H., Mlejnek M.: Interactive volume visualization of complex flow semantics. In *Proceedings of VMV'03* (2003), pp. 191–198.

[HMBG01] Hauser H., Mroz L., Bischi G. I., Gröller M. E.: Two-level volume rendering. *IEEE Trans. on Visualization and Computer Graphics 7*, 3 (2001), 242–252.

[KDG99] König A., Doleisch H., Gröller M. E.: Multiple views and magic mirrors - fMRI visualization of the human brain. In *Proc. of SCCG'99* (1999), pp. 130–139.

[KSW06] Krüger J., Schneider J., Westermann R.: Clearview: An interactive context preserving hotspot visualization technique. *IEEE Trans. on Visualization and Computer Graphics 12*, 5 (2006), 941–948.

[KTH*05] Krüger A., Tietjen C., Hintze J., Preim B., Hertel I., Strauss G.: Interactive visualization for neck dissection planning. In *Proc. of EuroVis'05* (2005), pp. 295–302.

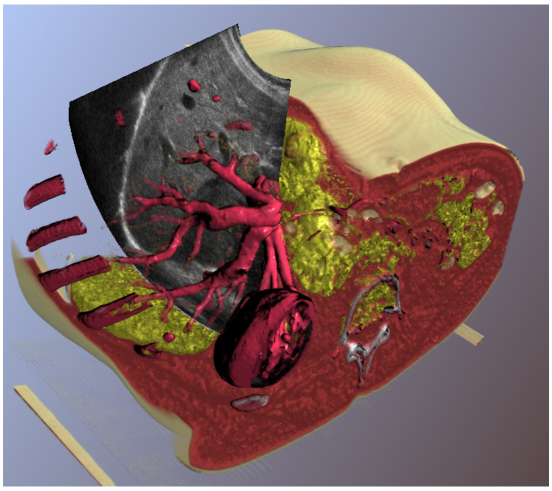

(a) Longitudinal ultrasound image, $E = 0.4$

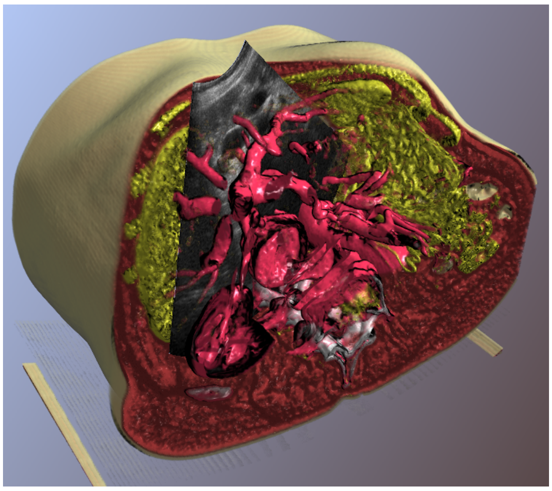

(b) Transversal ultrasound image, $E = 0$

Figure 6: *Liver ultrasound images are shown within a CTA scan. Vessels (pink) with an importance value of 1 are allowed to obscure the ultrasound plane. The effect of shading emphasis can be seen in (a) as shading of skin and flesh is subdued.*

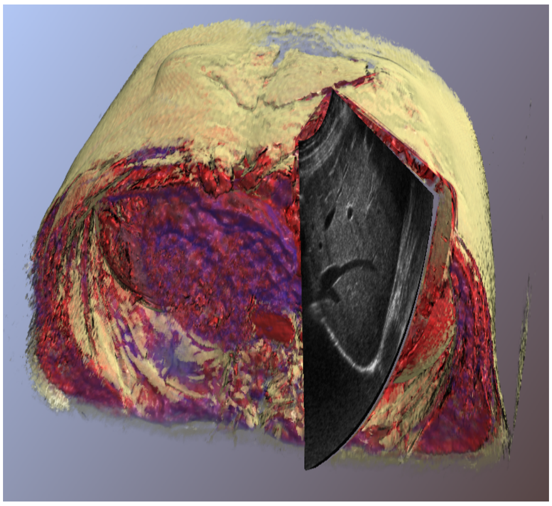

(a) Vessels have $I < 1$, $d = 0$, $\theta_1 = \theta_2$, $E = 0.1$

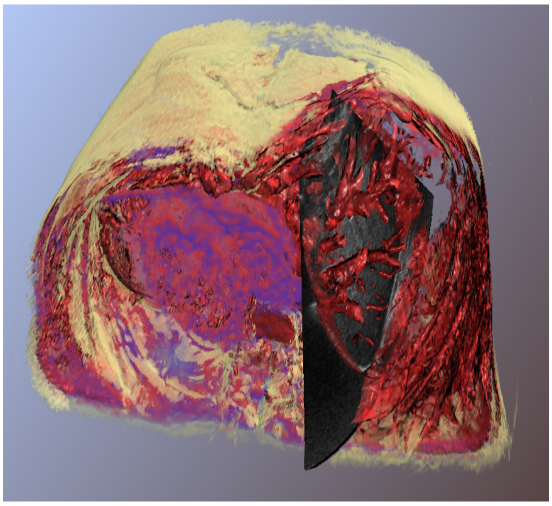

(b) Vessels have $I = 1$, small d, $\theta_1 < \theta_2$, $E = 1$

Figure 8: *A longitudinal ultrasound image of the liver is visualized within an MRI scan of the same volunteer. A simple cutaway scenario with no* transition *or* overlay *regions is shown in (a). The addition of* transition *and* overlay *regions in (b) allows one to see the vasculature emerge from regions with corresponding structure in the ultrasound plane.*

[RSK06] REZK-SALAMA C., KOLB A.: Opacity peeling for direct volume rendering. *Computer Graphics Forum 25*, 3 (2006), 597–606.

[SvC*04] STRAKA M., ČERVEŇANSKÝ M., CRUZ A. L., KÖCHL A., ŠRÁMEK M., GRÖLLER M. E., FLEISCHMANN D.: The VesselGlyph: Focus & context visualization in CT-angiography. In *Proc. IEEE Visualization '04* (2004), pp. 385–392.

[VG05] VIOLA I., GRÖLLER M. E.: Smart visibility in visualization. In *Proc. of Computational Aesthetics in Graphics, Visualization and Imaging* (2005), pp. 209–216.

[VKG05] VIOLA I., KANITSAR A., GRÖLLER M. E.: Importance-driven feature enhancement in volume visualization. *IEEE Trans. on Visualization and Computer Graphics 11*, 4 (2005), 408–418.

[WRN05] WEIN W., RÖPER B., NAVAB N.: Automatic registration and fusion of ultrasound with CT for radiotherapy. In *Proc. of MICCAI'05* (2005), pp. 303–311.