

Visibility Guided Multimodal Volume Visualization

Lin Zheng, Carlos Correa, Kwan-Liu Ma

Department of Computer Science

University of California, Davis

Email: lzheng@ucdavis.edu, cdcorrea@google.com, ma@cs.ucdavis.edu

Abstract—With the advances in dual medical imaging, the requirements for multimodal and multifield volume visualization begin to emerge. One of the challenges in multimodal visualization is how to simplify the process of generating informative pictures from complementary data. In this paper we present an automatic technique that makes use of dual modality information, such as CT and PET, to produce effective focus+context volume visualization. With volume ray casting, per-ray visibility histograms summarize the contribution of samples along each ray to the final image. By quantifying visibility for the region of interest, indicated by the PET data, occluding tissues can be made just transparent enough to give a clear view of the features in that region while preserving some context. Unlike most previous methods relying on costly-preprocessing and tedious manual tuning, our technique achieves comparable and better results based on on-the-fly processing that still enables interactive visualization. Our work thus offers a powerful visualization technique for examining multimodal volume data. We demonstrate the technique with scenarios for the detection and diagnosis of cancer and other pathologies.

I. INTRODUCTION

In interpreting volume data for medical diagnosis or surgical planning, the information which can be visualized from a single modality, e.g. Computed Tomography (CT), may be insufficient. A number of factors influence this, including limited resolution, sensitivity to tissue properties, noise, etc. For this reason, radiologists often make use of additional modalities that provide complementary or supplementary information. In this way, radiologists are able to extract more clearly the structures of interest and the spatial relationships among them. For example, CT provides the most detailed anatomical information from the human body, usually at high resolution. It helps depict high dense structures such as bone, as well as the shape of internal organs. On the other hand, the acquisition of metabolic activity must rely on a modality like Positron Emission Tomography (PET). In general, metabolic activity is important to detect cancer, since cancer tumors and other malignancies are usually located in regions with high rate of metabolic activity, such as regions with high blood flow. To obtain the best of the two modalities, recent visualization systems attempt at fusing both types of information in a single meaningful image.

With hardware acceleration, volume rendering has become very attractive to many applications. To be more widely adopted, however, its usability remains to be enhanced. In particular, the task of classifying volume data before rendering as well as the task of manipulating potentially a large number of rendering and viewing parameters to achieve desired visualization are often time-consuming and tedious. Recent research results show some good progress on visualizing individual

volume data, but multimodal volume rendering presents additional challenges, from the problems of superimposing dual modality data and highlighting objects of interest, to the desire to suppress occluding materials while maintaining the context and to enhance structural and spatial clarity of the objects.

Until now, the most common way in multimodal visualization is rendering the multiple volumes into one visualization by fusing them based on their weighted intensity values. CT is usually used as the contextual volume to show the anatomical relationships between the pathological tissue and other non-affected regions. However, these visualizations are only effective when the non-pathology tissues have sufficiently different intensities from that of the pathological ones. If the intensity of a target tissue is similar to that of the nearby tissues, then the user has to manually segment out that part and proceed to define a separate transfer function for highlighting the target issue. Such tedious, manual tasks are exactly the reasons 3D visualization is not widely used in medicine. An ideal visualization tool should do most of the work for the users to allow them concentrate on examining the data and the revealed structures.

We introduce a technique that can automatically generate focus+context visualization of multimodal volume data in a view-dependent fashion. In our study, we used PET and CT datasets. The PET data was used to mark the region of interest (ROI) and also modulate the opacity values for rendering the CT data. Rendering the affected organs in the ROI does not require the user to specify a special modulation function or a secondary transfer function for surrounding tissues. Instead, a local visibility histogram is computed for the target tissue and used to adjust the transfer function for rendering the occluding tissues. The adjustment is automatically and iteratively made to reach a balance between ROI visibility and context disambiguation. Immediate visual feedback enables the user to determine if the iteration should stop.

In this paper, we describe how we incorporate visibility information to fuse multimodal data sets for focus+context visualization. Using multiple-pass visibility calculations, we can tradeoff between visibility and spatial clarity. We demonstrate this new technique using four CT and PET data sets and discuss some of the initial feedback we have obtained from our collaborating radiologists and surgeons.

II. RELATED WORK

With the popularity of medical imaging technology in the recent years, it has been possible to extract different images from the human body under multiple modalities, including CT, MRI and PET. Dual modality medical scanners would

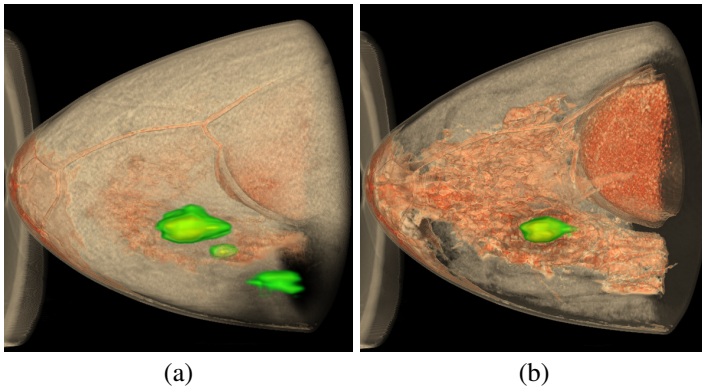


Fig. 1. A breast cancer tumor. As we increase the PET signal threshold (left to right), we can get a more precise view of the tumor in CT. Depending on the threshold we set, we can obtain different types of images

likely become the norm for diagnosing and detecting malignancies. For visualization, the existence of multiple sources for imaging a structure implies handling complementary and supplementary information in a single image. This can be quite challenging.

Previous work in multimodal visualization has addressed this issue as a fusion problem, where the final image is the result of combining different modalities at the data, accumulation, illumination or pixel levels [1]–[3]. Recent approaches combine different visual analysis to fuse information from perfusion data [4]. As GPUs become more widely available, the fusion of multiple volume data sets has become a commodity. Obtaining good images that convey the information in complementary modalities remains a challenge. Although registration and segmentation remain a challenge [5], [6], from the point of view of visualization, there are challenges in transfer function design and focus+context rendering. First, transfer functions are often limited by the particular dimension they classify, such as intensity, gradient magnitude [7], [8], curvature [9] or local structure [10]. Features in a volume data coexist in different regions of these spaces that often overlap, making the separability of structures of interest difficult [11].

On the other hand, the issue can be posed as a focus+context problem, where a data, geometric or optical transformation is applied to a region of interest, called the focus, and a different transformation to the remaining data, called the context. Focus+context techniques include interactive cuts [12], [13], ghosted views [14]–[16], opacity peeling [17], distorted views [18], or deformations [19]. Our approach also fits into the notion of focus+context techniques, especially tailored for multi-modality volumes. Some existing techniques can be extended for multi-modal visualization. Two-level volume rendering, for instance, uses different rendering styles to display segmented structures [20]. Importance-driven techniques can help visualizing hidden structures [16], but require a prior definition of the context regions. This notion was extended for the visualization of multimodal medical images [21]. Instead of requiring segmentation, users simply define importance as an extra dimension of the transfer function. While more interactive, it suffers from the separability issues known for intensity-based transfer functions.

Unlike previous approaches, we present a method that does

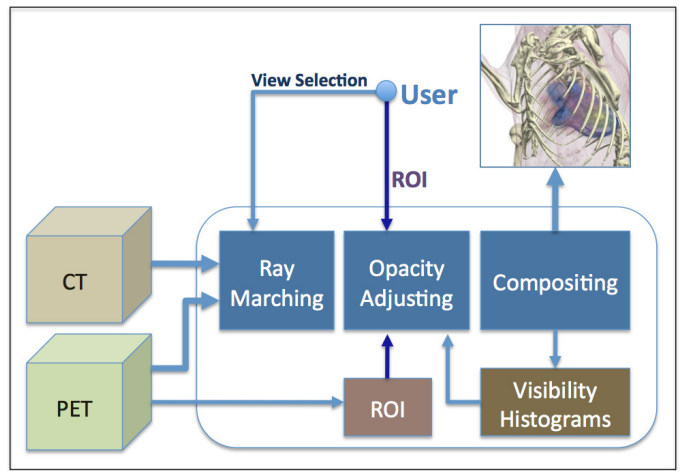


Fig. 2. Process of obtaining multimodal visualization using our visibility based pattern.

not require detail object segmentation, but rather uses the complementary information from another modality, e.g., PET, to provide an image of the other modality, e.g. CT, or their combination, in a way that ensures visibility of the region of interest. To this end, we exploit the notion of visibility histograms [22], which summarize the distribution of visibility of structure of interest from a given viewpoint. Techniques such as context-preserving volume rendering [14] use accumulated opacity to provide focus+context visualizations. We believe that visibility histograms encode more concisely the information required to measure the efficacy of a transfer function. Different from the approach by Correa and Ma [22], [23] which requires to calculate the visibility of all voxels in the volume, we employ local histograms to ensure that features of interest, automatically obtained from other modalities, e.g. PET, are sufficiently visible in the final image.

The idea of guiding the rendering of CT images from PET information was proposed by Kim et al. [24], [25]. These techniques, however, require a fuzzy segmentation of the PET data. Beyer et al. [26] combine segmentation with cuts, while Ropinski et al. [27] take a different approach and use closeups that decouple the multimodality rendering into several informative images.

III. TECHNICAL APPROACH

Figure 2 illustrates our technical approach and the corresponding processing. A ray-casting volume renderer is enhanced to render a multimodal volume data set, in this case including a CT and a PET. The PET data is used to mark the ROI and modulate opacity for the CT data (Section 4). The CT data is used to show anatomical structures and compute the visibility histograms (Section 5.2). Depending on the users requirements, after visibility histograms have been obtained through several passes of computation (Section 5.3), opacity attenuation of the influenced rays is accumulated to get the final rendering result (Section 5.4). Users are able to manipulate the area of ROI, control the visibility of this area, and vary view direction during the entire visualization process. When the view direction changes, the visibility histograms are recalculated and composition is redone to generate a new




Fig. 3. Global visibility histograms. We obtain a single visibility histogram for the volume in front of the ROI. A single histogram has the advantage that we can capture aggregated visibility information and is usually coherent across viewpoints and small TF changes.

image. A quantitative analysis about the cost of our technique is presented in Section 6.

A. PET-guided ROI

Radiologists diagnose a considerable number of patient cases every day. Therefore, the entire process of diagnosis and detection using multimodal data sets should be as simple and intuitive as possible. Currently, analysis of multimodal data sets requires a preprocessing stage that includes registration and segmentation for different data sets. This stage can be quite time-consuming. However, thanks to the increasing availability of multimodal medical scanners, this stage is likely to be done much faster and transparently to the user. Therefore, it is expected that the creation of meaningful images will remain a challenge. One technical challenge lies in ROI definition. One way to solve this problem is to set up the ROI around a tumor manually by identifying some key boundary points. However, if the user wants to set up a complex shape, he/she has to define a large number of boundary points around the ROI, which becomes time-consuming and error-prone. This process becomes more tedious when there are multiple tumors of interest. One way to alleviate this task is to set up spheres covering the ROI. This approach, however, may be difficult to adapt when the situation merits a more complex shape of the pathological tissue, or when the PET signal is noisy and scattered. For this reason, we would like to use the PET data set directly as a guide to define the ROI automatically.

If we only have the CT data set, getting a clear view of a region of interest requires segmentation. In the case of multimodal visualization, we have a functional data set, PET, which produces a 3-dimensional image of the metabolic process in the body. With the PET signal, radiologists can find where important pathological changes occur.

Until now, in most cases, PET serves as a data set that visually defines where is the interesting part. Radiologists visualize these using semi-transparent rendering and try to form an idea of the ROI shape and location. However, PET can be used also to supply us the spatial information. The higher the metabolic activity is in this area, the larger the corresponding intensity is. They provide the possibility to differentiate the interesting parts in the corresponding anatomical data set. The operation to define the PET-guided opacity is quite intuitive. The basic idea is to range the PET data set for obtaining a 3D mask that defines the contribution of the ROI into another modality, such as CT. We do this by evaluating the following equation at each sample s in the volume rendering process:

$$MASK_{value}(s) = g(PET(s)), \quad (1)$$




Fig. 4. Our per-ray visibility histogram dependent method. Top: original rendering without regards of visibility for different transfer functions and views. Without visibility, the PET could be visible but semi-transparent (left), totally (middle) or partially occluded (right). Bottom: our method provides visibility of the ROI without destroying the context. In left and middle images, the contours of ribs are retained automatically to emphasize the spatial relation. This effect is more subtle in the right image where the contour of pelvis bone is also preserved.

where $PET(s)$ is the intensity value at sample point s in PET data set, and $g(\cdot)$ is a windowing function for the local intensity value. This function can be a trapezoid, which causes a blurred boundary, or a box filter, which produces a clear-cut shape.

A breast cancer example is shown in the Figure 1. In this case there are several tumors clustered together. With manipulation of the range of PET data set, we are able to show these tumors separately since their intensities (activeness) in PET are different. In Figure 1(b), we are able to isolate one of the tumors. However, this method suffers from some limitations. It largely depends on the nature of the corresponding PET data set. There may be a scattered collection of automatically generated ROIs, some of which may not carry important information. For this reason, we turn to our main contribution, the visibility-guided fusion of dual modalities. By analyzing the visibility histogram of the ROIs, we can determine the best way to show the surrounding occluding tissue. Small isolated ROIs, possibly due to metabolic activities of healthy tissues, can be discarded with a single threshold based on the visibility histogram, alleviating the issues of automatic ROI definition.

IV. VISIBILITY-GUIDED FUSION

In this paper, we exploit the notion of visibility histograms to automatically find the best way to render multiple modalities, by maximizing the visibility of an ROI, in our case the PET signal, while restricting the visibility of the context, in our case the CT signal.

The visibility T of a point refers to the optical contribution of that point to the final image. For a given point p :

$$T(p) = e^{-\int_0^s \tau(t) dt} \quad (2)$$

where $\tau(t)$ is the attenuation coefficient of a sample and s is the parametrized value of p along a viewing ray. We define attenuation in terms of an opacity mapping $A: R \mapsto R$, such that $A(I(p))$ defines the opacity of p as a function of the intensity I of p .

In this work, we are interested in the visibility of PET structures, as a function of the attenuation introduced by the CT signal. The PET-visibility for a signal strength ϕ of a given pixel in the image is determined by:

$$T_{PET}(q) = e^{-\int_0^{t_{pet}} \tau(t) dt} \quad (3)$$

where $q = \vec{v} t_{pet}$, and t_{pet} is the parameter along the view direction where $PET(t) \geq \phi$. Note that in the above equation, $\tau(t)$ is the attenuation of a point parameterized by t , according to the CT intensity of the point. Therefore, the PET-visibility of a pixel depends on the opacity of CT samples in front of the first hit on the PET signal. To quantify this visibility we turn to the visibility histogram.

Correa and Ma [22] introduced the visibility histogram of a rendered image, which bins the visibility of each rendered voxel according to their intensity values, and can be computed with front-to-back compositing, as follows:

$$\begin{aligned} VH[x] &= VH[x] + T(p)A(x) \\ T(p + \Delta p) &= T(p)A(x) + (1 - T(p)) \end{aligned} \quad (4)$$

where $x = \text{round}(V(p)) \in [1, \dots, N]$ are quantized intensity values in N bins, and VH denotes the visibility histogram. Unlike the original approach, we are not interested in measuring the full visibility of an image, but we restrict ourselves to local visibility histograms around the ROI. We show two ways for doing this: first, as a single visibility histogram for the ROI, and second, as a set of per-ray visibility histograms. While the former is more robust to view changes, the latter can adapt to different types of occluding tissues more easily. In this paper we highlight the benefits of each of these approaches.

A. ROI Visibility Histogram

In this case, we compute the visibility histogram only for the rays in the ROI defined by the PET signal. We restrict the histogram to only 16 bins, which allows us to compute the histogram in a single pass and keep it available in GPU memory at all times using the available color channels. This capability enables real-time analysis of the visibility with very little impact on performance, if perceivable at all. Since the purpose is to adjust the opacity of occluding tissues to get a clear view of the PET signal, it is sufficient to use 16 bins here. Because this method aggregates the visibility into a single histogram, it is not sensitive to small sparse structures, like small vessels. Those structures will appear in the final image, providing context, but the PET signal will still have a clear view. In contrast, highly occluding tissue, such as muscle or fatty tissue, will affect considerably the overall visibility distribution, and therefore are likely candidates to be made transparent. Figure 3 illustrates this process. Once the visibility histogram is known, we adjust the opacity of samples in the ROI with the following formula:

$$A'(x) = A(x)(1 - VH(x))^e \quad (5)$$

where $A(x)$ and $A'(x)$ are the old and new opacity mappings for an intensity x , $VH(x)$ is the visibility associated with that intensity in the visibility histogram, and e is an exponent that defines the strength of such a mapping.

According to this formula, points that dominate the visibility, i.e., $VH(x)$ is high, are likely to be made more transparent, while samples that do not occlude much – $VH(x)$ is low, are kept to provide context. The use of an aggregated mapping for each intensity bin, instead of an individual opacity mapping for each sample point results in a coherent opacity mapping from pixel to pixel and across view changes.

B. Per-Ray Visibility Histograms

One of the issues with an ROI-wide visibility histogram is that we treat all samples under the same intensity, e.g., bone, with the same opacity adjustment. This may be counter-productive to provide contextual information. For example, for bone, the middle solid part may be more occluding than the contours. By keeping the contours and remove the middle solid part, we can still retain the shape without much impact on visibility. Thus, we want to automate this decision using visibility histograms.

The process is essentially the same as described above, except we keep a visibility histogram for each ray hitting the ROI. Because the opacity transformation in eq. 5 is applied per ray, different parts of a tissue with the same intensity value end up with different transparencies, providing the best trade off between visibility and context. An example is shown in Figure III-A. We see a traditional rendering on top, where the visibility of the ROI (green and orange) is affected by occluding tissue, to the point that it becomes totally (middle) or partially occluded (right). At the bottom, we show the automatic adjustment based on PET. We see that the most of the occluding tissue is removed, namely the muscle, fatty tissue and the liver. Tissues that do not affect much visibility, such as the ribs and skin are retained. The per-ray visibility histograms let us automatically emphasize the contours, as seen more clearly in the middle figure. The use of contours is essential for disambiguating depth relationships.

C. Visibility Refinement

One may notice that, once we apply opacity adjustment, the visibility of the PET signal changes. Therefore, we can refine the opacity based on the new visibility histogram in the search for better visibility. We can do this for any number of iterations, but, in the long run, this process converges to the occluding tissue being completely transparent. For this reason, we let the user explore up to 3 iterations of visibility refinement, each of which represents and different combination of the ROI and context. An example is shown in Figure 5. We can see that, for the first iteration, fatty tissue and muscle dominate visibility. On the second iteration, these tissues are removed, and now bone tissue dominates visibility and are made more transparent. Note that the opacity transformation for the bone is not homogeneous: near the contours, where the view ray is tangential, the contribution of bone tissue to visibility is small and does not need to be transformed. Near the center, however, the ray has to go through more samples and contribute more to the visibility histogram, and thus are




Fig. 5. Effect of several iterations in visibility-weighted adjustment. (a) For one iteration, the ribs, which are originally very opaque, are not entirely removed, since they do not affect visibility as much as other tissue. (b) In a second iteration, we see that ribs can be made more transparent, while still retaining the contours. (c) By the time we reach a third iteration, the bones can be even made more transparent, giving a clear view of the ROI. (d) For a different transfer function, which produces an opaque rendering and highly occluding tissue, our approach still provides a good visibility of the tumor while retaining the skin and ribs contours for context




Fig. 6. Figure (a) shows a mouse with tumor in its left lung. Figure (b) shows with only one-pass, the skin and part of the muscle are removed, the bone of the arm and ribs become visible. In Figure (c), with two-passes, only contours of the arm bone and ribs are left. In Figure (d) we can see only faint signal of the ribs.

made more transparent. This explains why the contours are likely to be preserved. On a third and final iteration, almost all occluding tissues are removed. For this particular case, a user may find the second iteration suitable planning a surgical intervention.

Another example is shown in Figure 5(d), where we show the result of applying our approach when the original mapping depicts tissues as very opaque. Our approach is able to remove all unimportant tissues in the third iterations, while still providing context information. Notice the shape of skin and ribs as depicted in their contours.

The main advantage of this approach compared to existing approaches is that we no longer need to specify separate measurements for the focus and context regions. A naive reduction of the opacity of all occluding samples would penalize the regions that do not affect visibility but that otherwise could provide contextual information. Figure 7 compares the result of separate focus+context region (left) with our automatic

approach (right). We notice that our method can provide equivalent or even better results without the time-consuming task of finding suitable transfer function for the context regions.

D. ROI Visibility Guarantee

One problem with the approach is providing visibility guarantees of the ROI. So far, we let the user explore up to 3 iterations and let them decide which visibility configuration is best. Automating this process proves more challenging. One simple approach is to make the occluding tissue completely transparent, but it completely loses all the context information. An alternative, common in existing approaches is to adjust the context region with an independent transfer function. While a manual design of such a transfer function may work, at the expense of time, automatic construction cannot guarantee visibility of the ROI as the mapping is computed per intensity value. What we propose is an approach that automates this process, given an ROI and its desired visibility.




Fig. 7. Comparison between manual generation of separate focus and context transfer functions (left) and our automatic visibility-based opacity adjustment (right), for an abdomen lymphoma. Our automatic approach offers equivalent visual results without the need for user intervention to highlight the focus region.

First, we let the user specify a desired visibility of the ROI. A high value indicates that the ROI should be as visible as possible, maybe at the cost of losing context information. A low value indicates that the ROI should be visible, but can be occluded to some extent. Our approach then iterates over the visibility computation up to a maximum number of times. After each iteration finishes, we accumulate the visibility of the points in the ROI. If the result is larger than the visibility threshold, the process is stopped and the final image is rendered. This mechanism applies both for the ROI- and per ray- visibility histograms. This automated approach alleviates the need for manually tweaking a transfer function and relying on anatomical knowledge. Instead, our approach makes decisions based on spatial relationships alone and can be applied to a variety of cases and visibility configurations. We expect that people with limited training in medical images will benefit from our approach.

V. RESULTS AND EVALUATION

We now discuss the results and evaluation of our approach.

Lymphoma: Figure 5 shows an abdominal lymphoma. Initially, the area contains the lymphoma is surrounded by fat. A transfer function that highlights only the tumor is difficult to obtain, since its intensity is close to that of the fat/muscle tissue boundary and the liver. If we remove the fatty tissue in a transfer function, we lose the CT signal for the lymphoma. Therefore, we isolate the lymphoma using the PET signal and set up the ROI with an ellipsoidal 3D mask. Our visibility-guided method provides views of the tumor under different transfer functions and viewpoints. With a careful study of the isolated signal in the CT image, we see that the PET-signal follows the overall shape of the tumor in the CT signal.

Mouse: Figure 6 shows a mouse cancer. The cancer is surrounded by lung and ribs. This time, it is the pathological part of the lung which shares the same intensity value of surrounding tissue, making it hard to isolate using transfer functions. Again, we turn to the PET signal to visualize the tumor. We show the result of our approach for a number of visibility modulation iterations.

Brain: Figure 8 shows a brain tumor using a combination of CT and multiplanar reformat (MPR) modalities. In the





Fig. 8. Automatic visibility-guided rendering of ROI on a brain dataset. (a) Original visualization of CT and MPR. (b) Multimodal visualization of CT and MPR highlighting a brain tumor. We define the ROI as a sphere, but the opacities are adjusted automatically based on visibility. We clearly visualize the tumor but also retain some of the important vessels for context.

absence of a smooth PET signal, we instead define the ROI as a sphere containing the tumor. Our approach automatically makes the bone tissue fully transparent, so that the interior can be visible, but also retains nearby vessels for context.

Breast: Figure 9 shows a fused PET-CT scan of a female breast. Here, we see two kinds of tumors. One is DCIS (ductal carcinoma in situ) and the other is benign. The ability to isolate ROIs using PET signal becomes crucial, since only the DCIS is captured in the PET signal. Existing visualization systems typically remove the vessels surrounding a tumor due to occlusion and noise. However, for a radiologist it is important to consider the neighboring vessels, since the PET signal indicates a high rate of metabolic activity. In our example, we use a transfer function that can highlight glandular tissue, skin and vessels. Without any visibility provision, these tissues occlude each other and it is not possible to obtain any meaningful information, as shown in Fig.9(a). With our approach, we can now clearly see the tumor and surrounding vessels, as shown in Figure 9(c), which gives a more complete view of the affected area.

A. Informal Evaluation

We validated our results with radiologists who provided us the data sets. Radiologists often need to see the actual CT signal associated with strong PET signal. Our approach makes a better job at showing the CT signal associated with an ROI than current technologies. Radiologists also found attractive the ability to preserve contextual information around the space-occupying lesions. When treating cancer, it is important to distinguish four kinds of space-occupying lesions: hamartoma, inflammation, benign tumor and malignancy. The benign tumor usually shows expansive growth and pushes away the surrounding tissues, while the malignant tumor is always associated with invasive growth and infiltrates the surrounding tissues. Besides the invasive growth, another biological behavior of malignancy is metastasis. This is often manifested in lymphatic and blood vessels, often seen as small, sparse features in the CT signal. Our approach, unlike existing technologies, is able to retain these as part of the CT context. Surgeons also expressed the need for the CT context to be



(a) Original Visualization (single TF) (b) with 2 TFs (c) PET-defined ROI after 1 iteration

Fig. 9. Automatic visibility-guided rendering of ROI on a breast PET-CT data set. (a) With a roughly-defined transfer function that highlights tumor and vessels, it is not possible to see the tumor. Obtaining a better depiction is time-consuming and may require multi-dimensional TFs. (b) 2 transfer functions help visualize the tumor inside. (c) Automatic PET-guided ROI and visibility-guided rendering provides a better view of the tumor and still provides the surrounding vessels for context

able to plan a treatment procedure, such as a virtual biopsy or a puncture. Although the images we provided contain a fair balance between focus and context, we found that interactive exploration was most useful.

VI. IMPLEMENTATION

We have implemented our approach entirely on the GPU through fragment shaders. In general, our approach follows a traditional volume rendering approach, to make it easy to deploy in current visualization systems. We use multiple-passes, in order to obtain the maximum possible number of bins for the visibility histogram and the flexibility of multiple iterations. The first pass records a 2D mask in screen space for the ROI, while subsequent passes compute the visibility histogram. A final pass finally composites the result image using visibility-guided opacity modulation. The shader is essentially a raycasting cycle, where we gather the contribution of visibility of each sample as we traverse them. For speed, we can gather up to 16 bins in a histogram in a single pass shader and gather them in four textures, each using four color channels. For an ROI visibility histogram, we simply gather and sum the information from the four targets. We combine the automatic ROI from the PET as a mask to define the region in the view which require visibility histogram computation.

This implementation is much faster than previous approaches for computing the histogram [22] since it keeps everything on the GPU memory and only requires a small number of bins. In our experiments, we did not notice a particular disadvantage on using a small-sized histogram. In practice, the result with eight bins histogram is also sufficient for our needs. An important consideration of our approach is that we made sure to keep the bins close to the intensities of interest and avoid using bins for the entire intensity range, where some intensities have no contribution at all in the original data set.

A. Performance

Our approach is not asymptotically more expensive than typical volume rendering with multiple modalities, except that now it depends on: (1) the number of visibility modulation

iterations, and (2) the number of passes required to compute a visibility histogram. However, this cost in practice increases much slower than that of multiple rendering passes. First, the number of visibility modulation iterations is usually constant, and in most cases fewer than three iterations are required. Second, the visibility histogram is not computed for the entire viewport, but just for the pixels occupied by the ROI projection onto the view plane. Finally, the number of passes required to compute a visibility histogram is exactly one, since we can write up to 16 bins simultaneously, using four texture targets with four channels each.

We tested our approach on a standard PC with a 2.40 GHz Intel Core(TM)2 Quad CPU with 3GB of memory and a NVIDIA GTX280 graphics card with 1GB of graphics memory. We tested three multimodal datasets and compare the performance for different opacity modulation iterations and sampling rates.

TABLE I. PERFORMANCE COMPARISON FOR LYMPHOMA

Method	frame rate with 0.01 stepsize	frame rate with 0.002 stepsize
normal way	26.7 fps	8.9 fps
1 iteration	17.0 fps	6.1 fps
2 iterations	14.0 fps	5.1 fps
3 iterations	13.0 fps	4.8 fps

TABLE II. PERFORMANCE COMPARISON FOR MOUSE

Method	frame rate with 0.01 stepsize	frame rate with 0.002 stepsize
normal way	25.0 fps	10.3 fps
1 iteration	17.8 fps	7.0 fps
2 iterations	17.4 fps	6.3 fps
3 iterations	16.8 fps	5.8 fps

TABLE III. PERFORMANCE COMPARISON FOR BREAST

Method	frame rate with 0.01 stepsize	frame rate with 0.002 stepsize
normal way	24.3 fps	8.5 fps
1 iteration	18.1 fps	7.6 fps
2 iterations	14.2 fps	5.9 fps
3 iterations	13.3 fps	4.9 fps

The first test data set depicts an abdominal lymphoma, and consists of a $256 \times 256 \times 292 \times 4$ bytes CT data set and a

168 × 168 × 344 × 2 bytes PET data set. The frame size is 400 × 600 pixels while the ROI covers about 120 × 180 pixels. At first, we test the performance with typical volume rendering. Table I summarizes the rendering rate for multiple visibility modulation iterations compared to single-pass volume rendering. Even with the extra cost, our approach runs at 4.8 frames per second at the finest step size, which is acceptable in interactive visualization. The second dataset contains a cancer tumor inside a mouse lung with a 272 × 216 × 480 × 2 bytes CT data set and a 128 × 128 × 133 × 4 bytes PET data set, an a ROI of 70 × 50 pixels. Results are shown in Table II. The third data set is a PET-CT breast data set, composed of a 256 × 256 × 256 × 2 bytes CT data set and a 256 × 256 × 256 × 4 bytes PET dataset. In this case, the frame size is 512 × 512 pixels. The ROI is about 130 × 130 pixels. Results are shown in Table III.

VII. CONCLUSION

We have introduced an effective approach to multimodal volume rendering with a focus on presenting complementary information such that regions of interest, e.g., cancer tumors, can be examined in context. Our results show that visibility histograms are useful aids to quantifying the information derived from structures of interest in one modality as suggested by the other modality. Previous attempts require manual intervention to obtain desirable images, mainly due to the lack of interactive tools for quantifying key properties such as visibility. Our design and implementation enable interactive visualization, and represent a step forward in a direction towards automatic view-dependent classification of complex volume data sets. Future generations of medical visualization systems should incorporate visibility measures, along with others, such as automatic contrast and smoothness control, which are high-level, perceptual-based metrics that alleviate the tedious task of manipulating individual lighting parameters or transfer functions. The initial feedback obtained from the surgeons and radiologists whom we have worked with encourages us to further our study towards more intuitive and efficient multimodal volume data visualization.

ACKNOWLEDGMENT

This research is supported in part by the U.S. National Science Foundation through grants IIS-1320229 and CCF-0811422. Datasets were provided by John Boone, Ramsey Badawi, and Angeliqe Louie at UC Davis.

REFERENCES

- [1] W. Cai and G. Sakas, "Data intermixing and multi-volume rendering," *Computer Graphics Forum*, vol. 18, pp. 359–368, 1999.
- [2] M. Ferre, A. Puig, and D. Tost, "A framework for fusion methods and rendering techniques of multimodal volume data: Research articles," *Computer Animation and Virtual Worlds*, vol. 15, pp. 63–77, 2004.
- [3] A. Ghosh, P. Prabhu, A. E. Kaufman, and K. Mueller, "Hardware assisted multichannel volume rendering," *Computer Graphics International Conference*, vol. 0, pp. 1530–1052, 2003.
- [4] B. Preim, S. Oeltze, M. Mlejnek, E. Gröeller, A. Hennemuth, and S. Behrens, "Survey of the visual exploration and analysis of perfusion data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, pp. 205–220, 2009.
- [5] M. Sonka and M. J. Fitzpatrick, *Handbook of Medical Imaging Volume 2. Medical Image Processing and Analysis*, M. Sonka and M. J. Fitzpatrick, Eds. SPIE Press, 2000.
- [6] J. Maintz and M. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, pp. 1–36, 1998.
- [7] M. Levoy, "Display of surfaces from volume data," *IEEE Computer Graphics and Applications*, vol. 8, pp. 29–37, 1988.
- [8] G. Kindlmann and J. W. Durkin, "Semi-automatic generation of transfer functions for direct volume rendering," in *Proceedings of the 1998 IEEE symposium on Volume visualization*, 1998.
- [9] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Moller, "Curvature-based transfer functions for direct volume rendering: Methods and applications," in *Proceedings of IEEE Visualization*, 2003.
- [10] Y. Sato, C.-F. Westin, A. Bhalerao, S. Nakajima, N. Shiraga, S. Tamura, and R. Kikinis, "Tissue classification based on 3d local intensity structure for volume rendering," *IEEE Transactions on Visualization and Computer Graphics*, vol. 6, pp. 160–180, 2000.
- [11] H. Pfister, B. Lorensen, C. Bajaj, G. Kindlmann, W. Schroeder, L. S. Avila, K. Martin, R. Machiraju, and J. Lee, "The transfer function bake-off," *IEEE Computer Graphics and Applications*, vol. 21, pp. 16–22, 2001.
- [12] J. Diepstraten, D. Weiskopf, and T. Ertl, "Interactive cutaway illustrations," in *Proceedings of Eurographics*, 2003.
- [13] D. Weiskopf, K. Engel, and T. Ertl, "Interactive clipping techniques for texture-based volume visualization and volume shading," *IEEE Transactions on Visualization and Computer Graphics*, vol. 9, pp. 298–312, 2003.
- [14] S. Bruckner, S. Grimm, A. Kanitsar, and M. Gröller, "Illustrative context-preserving volume rendering," in *Proceedings of EuroVis*, 2005.
- [15] J. Kruger, J. Schneider, and R. Westermann, "Clearview: An interactive context preserving hotspot visualization technique," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, pp. 941–948, 2006.
- [16] I. Viola, A. Kanitsar, and M. E. Groller, "Importance-driven feature enhancement in volume visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, pp. 408–418, 2005.
- [17] C. Rezk-Salama and A. Kolb, "Opacity peeling for direct volume rendering," *Computer Graphics Forum*, vol. 25, pp. 597–606, 2006.
- [18] L. Wang, Y. Zhao, K. Mueller, and A. E. Kaufman, "The magic volume lens: An interactive focus-context technique for volume rendering," in *Proceedings of IEEE Visualization*, 2005.
- [19] C. Correa, D. Silver, and M. Chen, "Feature aligned volume manipulation for illustration and visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, pp. 1069–1076, 2006.
- [20] M. M. Hadwiger, C. C. Berger, and H. H. Hauser, "High-quality two-level volume rendering of segmented data sets on consumer graphics hardware," in *Proceedings of IEEE Visualization*, 2003.
- [21] M. Burns, M. Haidacher, W. Wein, I. Viola, and E. Gröller, "Feature emphasis and contextual cutaways for multimodal medical visualization," in *Eurographics / IEEE VGTC Symposium on Visualization*, 2007.
- [22] C. Correa and K. L. Ma, "Visibility driven transfer functions," in *Proceedings of IEEE Pacific Visualization 2009 Symposium*, 2009.
- [23] —, "The occlusion spectrum for volume classification and visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, pp. 1465–1472, 2009.
- [24] J. Kim, W. Cai, and D. Feng, "Dual-modality pet-ct visualization using real-time volume rendering and image fusion with interactive 3d segmentation of anatomical structures," in *Engineering in Medicine and Biology Society*, 2005.
- [25] J. Kim, W. Cai, S. Eberl, and D. D. Feng, "Real-time volume rendering visualization of dual-modality pet/ct images with interactive fuzzy thresholding segmentation," *IEEE Transactions on Information Technology in Biomedicine*, vol. 11, pp. 161–169, 2007.
- [26] J. Beyer, M. Hadwiger, S. Wolfsberger, and K. Bhler, "High-quality multimodal volume rendering for preoperative planning of neurosurgical interventions," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, pp. 1696–1703, 2007.
- [27] T. Ropinski, I. Viola, M. Biermann, H. Hauser, and K. Hinrichs, "Multimodal visualization with interactive closeups," in *EGUK Theory and Practice of Computer Graphics*, 2009.