Visibility-driven PET-CT Visualisation with Region of Interest (ROI) Segmentation

Younhyun Jung[1], Jinman Kim[1], Stefan Eberl[1,2], Michael Fulham[1,2,3], and David Dagan Feng[1,4]

[1] Biomedical and Multimedia Information Technology (BMIT) Research Group, University of Sydney, Australia.

[2] Department of Molecular Imaging, Royal Prince Alfred Hospital, Sydney, Australia.

[3] Sydney Medical School, University of Sydney, Australia.

[4] Med-X Research Institute, Shanghai Jiao Tong University, China.

Abstract

Multi-modality positron emission tomography – computed tomography (PET-CT) visualises biological and physiological functions (from PET) as region of interests (ROIs) within a higher resolution anatomical reference frame (from CT). The need to efficiently assess and assimilate the information from these co-aligned volumes simultaneously has stimulated new visualisation techniques that combine 3D volume rendering with interactive transfer functions to enable efficient manipulation of these volumes. However, in typical multi-modality volume rendering visualisation, the transfer functions for the volumes are manipulated in isolation with the resulting volumes being fused, thus failing to exploit the spatial correlation that exists between the aligned volumes. Such lack of feedback makes multi-modality transfer function manipulation to be complex and time-consuming. Further, transfer function alone is often insufficient to select the ROIs when it comprises of similar voxel properties to those of non-relevant regions.

In this study, we propose a new ROI-based multi-modality visibility-driven transfer function (m^2-vtf) for PET-CT visualisation. We present a novel 'visibility' metrics, a fundamental optical property that represents how much of the ROIs are visible to the users, and use it to measure the visibility of the ROIs in PET in relation to how it is affected by transfer function manipulations to its counterpart CT. To overcome the difficulty in ROI selection, we provide an intuitive ROIs selection tool based on automated PET segmentation. We further present a multi-modality transfer function automation where the visibility metrics from the PET ROIs are used to automate its CT's transfer function. Our GPU implementation achieved an interactive visualisation of multi-modality PET-CT with efficient and intuitive transfer function manipulations.

Keywords: Multi-modality volume rendering and Visibility histogram and Transfer function and PET-CT imaging and Image segmentation

1 Introduction

Multi-modality (MM) biomedical imaging devices such as positron emission tomography and computed tomography (PET-CT) have enabled improvements in cancer staging and assessing the response to treatment [1]. PET-CT allows visualisation of biological and physiological function (PET) within the spatial context of anatomy (CT). Thus a key element in PET-CT visualisation is to

optimally visualise the region of interest (ROI) from PET, such as tumours and other abnormal structures, while preserving as much visibility of the underlying anatomy from CT, without compromising the visibility of the ROIs. The ability to efficiently assess and assimilate the information contained in these two volumes simultaneously, however, has raised new 3D visualisation challenges.

Volume rendering algorithms are an effective approach to 3D visualisation to discover meaningful data in biomedical images [2]. In these renderings, transfer functions, which control opacity and colour parameters, are commonly used to manipulate the volumes [3-8]. Although the transfer functions provide powerful manipulation capabilities, e.g., there are algorithms that can use varying voxel properties including gradients [3], textures [7], or size [8], these algorithms are typically manually-driven and need careful inspection of the changes in the resultant renderings. The complexity of applying transfer functions is compounded when multi-modality volumes such as PET-CT are considered since they require a combined manipulation of a pair of transfer functions. In a typical multi-modality visualisation [9-12], transfer functions are defined independently on the individual volumes with the resulting volumes being fused, e.g., PET transfer function is initially defined which is then followed by manipulating its counterpart CT transfer function to select corresponding anatomical regions while minimising the occlusion of PET. Unfortunately, such individual transfer function manipulations are complicated and time-consuming. Pre-defined transfer functions also do not sufficiently cater for inter-patient variation and variations with different pathologies.

In this study, we propose a new multi-modality visibility-driven transfer function (m^2-vtf) for PET-CT visualisation where we employ **'**visibility**'** metrics to measure the level of occlusion caused by CT structures in front of and/or overlapping the PET ROIs. The use of visual cues such as the visibility has demonstrated assistance in the transfer function manipulation process [13-16]. The visibility histogram proposed by Correa and Ma [13] used the level of visibility of all samples along a view-point as a visual feedback attribute in real-time volume rendering. They showed that the visibility histogram could help discover occlusion patterns and provide a powerful feedback mechanism to guide / automate the transfer function manipulation. We extended the visibility calculations into multi-modality volumes to provide a mechanism for visual feedback of the occlusion, in this case, occlusion from the CT to its counterpart PET, and vice versa. Further, we introduce a visibility-driven transfer function automation to optimise the transfer functions for PET-CT volumes. In PET-CT it is necessary to separate ROIs from non-relevant regions for further analysis. As an example, abnormalities in PET images, including tumours, can occupy the same intensity ranges as normally occurring structures in the body such as with the liver, kidneys, bladder, and brain (see Fig. 1). In such cases, even with the aid of the visibility metrics, the transfer functions manipulations (in 1D or multi-dimensions) are unable to depict all the tumours seen throughout the body. This further limits the ability of the visibility metrics since they are unable to separate the visibility of the ROIs from other non-relevant regions. We, therefore, integrate an automated segmentation algorithm to our rendering framework such that the ROI from PET can be intuitively selected. The visibility of the ROIs can then be used as an optimisation constraint to derive optimal CT transfer functions. We present the capabilities of our method in the visualisation of clinical PET-CT studies with several different anatomy and disease types.

2 Related Work

2.1 Transfer Function Manipulation

Transfer functions, which are used to isolate and manipulate ROIs, play an important role in the visual quality of image volumes generated by direct volume rendering [2]. The most common transfer function is one-dimensional (1D) and it uses voxel intensities to classify the ROIs. To improve the manipulation capabilities, a number of multi-dimensional transfer functions have been proposed. The basic idea is to incorporate varying voxel properties derived from a volume into the transfer function manipulations. In a pioneering study, Kindlmann et al. [3-4] introduced two-dimensional (2D) transfer

functions based on first- or second-order derivatives of voxels within a volume; thus enabling users to emphasise the boundaries of adjacent anatomical regions. Kniss et al. [5] suggested a set of widgets to interactively manipulate these transfer functions. Caban and Rheingans [7] used local statistical texture properties to distinguish ROIs from regions with the same intensity and gradient. Recently, Correa and Ma [6, 8] used spatial voxel properties, i.e., size and occlusion in transfer functions to manipulate ROIs in complex datasets. However, these studies are inherently designed for single-modality images, and their application to multi-modality volumes is often restricted to the manipulation of the transfer function for individual volumes in isolation.

Specific transfer function manipulation algorithms for multi-modality volumes have also developed. Cai and Sakas [9] introduced a data intermixing algorithm to fuse multiple volumes in volume rendering. Bramon et al. [10] suggested an information-theoretic method to automatically select the most informative voxels from two overlapping volumes. Kim et al. [12] used a pair of 1D transfer functions for PET and CT volumes, with the resultant volumes being fused (data intermixing). In another study, Haidacher et al. [11] proposed a 2D transfer function space that was made of a fused volume for multi-modality visualisation and an additional θ value, which described the complementary information contained from the volume pair, was able to separate the different tissue types from the two volumes.

In all these studies, transfer functions were applied to individual volumes, and therefore disregarded how a change in the transfer function of one volume may affect the second volume. Our motivation for this work stemmed from a lack of occlusion feedback from spatially-aligned multiple volumes in the manipulation of multi-modality transfer functions. To equip the user with a feedback mechanism for the transfer function manipulation process, we use the visibility property that represents how much of each voxel in a volume is visible to the user.

2.2 Visibility-Driven Transfer Function

The visibility of a voxel has been used as a fundamental optical property in numerous visualisation applications e.g. to guide the selection of the optimal view-point of a volume [17]. Correa and Ma [13] showed a usage of the visibility to guide the transfer function manipulation for single-modality volumes. They introduced the notion of visibility histogram, which represents the opacity contribution of each intensity bin in the histogram to the resultant rendering, as an interactive aid for generating effective transfer functions. Furthermore, they suggested a semi-automatic transfer function manipulation approach, which can identify optimal opacity parameters to maximise the visibility of the user-defined ROIs in the volume. Ruiz et al. [14] then proposed an automatic transfer function manipulation scheme, where the user assigned the desired visibility and the opacity mapping was then automatically optimised. Wang et al. [15] instead employed a '*feature visibility*' where the visibility was calculated per 'feature' from the volume. Here, features were user-selected visual attributes in multi-dimensional transfer function space and these were used to select and manipulate the volume rendering.

Jung et al. [16] recently considered visibility in multi-modality volumes where the visibility histogram was calculated for the two volumes, and then used to guide the transfer function of one volume, while observing how it affects the visibility of its counterpart volume. Although this approach showed the feasibility of using visibility in multi-modality transfer function, it relied exclusively on manual transfer function manipulations. So in this study, we introduce a new visibility-driven transfer function for multi-modality visualisation with intuitive ROI selection and automated transfer function manipulations.

2.3 Segmentation-based ROI Visualisation

The use of ROIs has received attention due to its ability to prioritise and provide focus in volume rendering visualisations. Viola et al. [18] introduced importance-driven volume rendering where an '*importance value*' which indicates visibility priority, was assigned to the segmented ROIs. Using the ROIs, it generated a cut-away view which retained the visibility of ROIs with high importance by removing ROIs with lower importance. One of its drawbacks, however, was that it removed less

important ROIs and often related information was lost. In Tzeng and Ma [19], individual segmented ROIs were given their own transfer function, with all the ROI renderings fused into a single scene. Kim et al. [20] used segmentation by rendering the segmented ROIs (tumours from PET) together with PET and CT volumes.

In our paper, we have used automated segmentation to select the ROIs, which are then used in visibility-driven transfer function manipulation. Segmentation enables separation of different structures in PET which can then be used for ROI-based analysis and importance-driven visualisation approaches. The visibility histogram, when applied to the image volume, is unable to provide a visibility metric only for the ROIs, since it is unable to separate the visibility of ROIs from non-relevant regions [3].

3 Methods

3.1 Overview

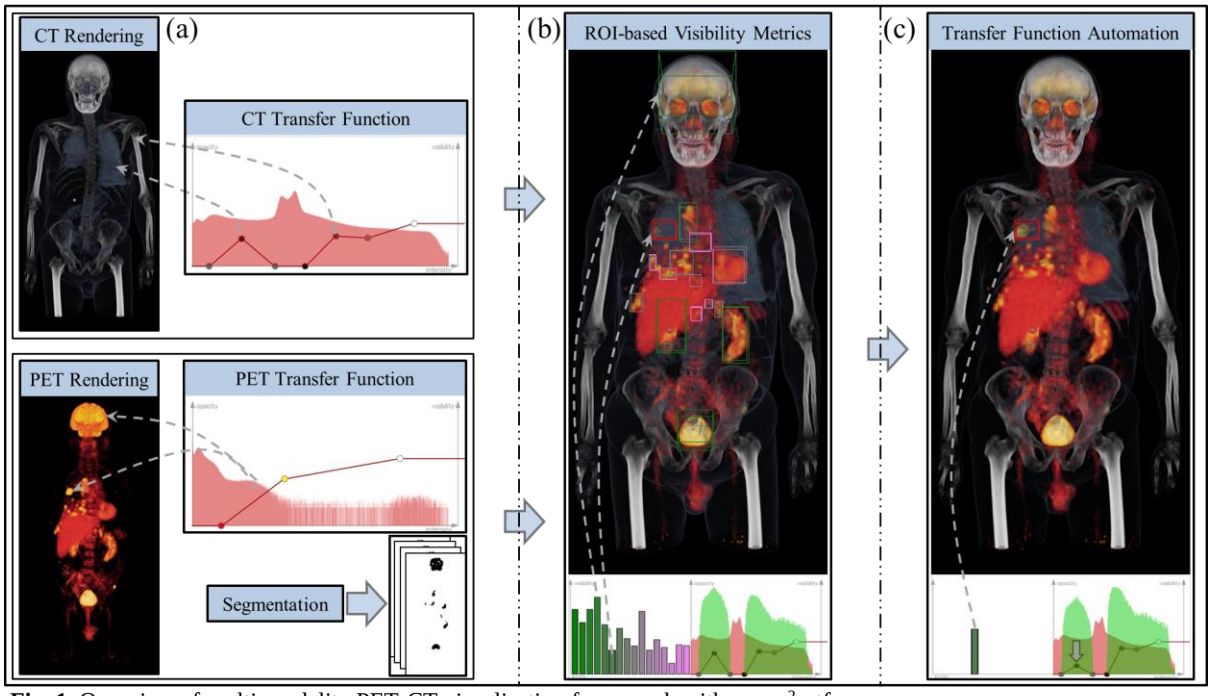

Fig. 1. Overview of multi-modality PET-CT visualisation framework with our m²-vtf.

Our PET-CT visualisation has three major steps as depicted in Fig. 1: (a) individual PET and CT volumes are rendered with their own transfer functions. In these transfer functions, a control line (red) drawn on top of the intensity histogram (pink bars in the x-axis) describes the opacity mapping (y-axis). Here, the CT transfer function control line is assigned to two opacity peaks corresponding to lungs and skin (first peak), and bones (second peak), as indicated by arrows. Similarly for PET, a control line is defined to select high functional activities, which includes the tumour but also other non-relevant structures. Prior to the rendering, the PET image is automatically segmented into ROIs to separate high uptake structures; (b) the two volumes are fused (data intermixing) into a single rendering scene, which then computes the visibility histogram for the PET and CT, by taking into account the overlapping voxels from the two volumes as well as their opacity mappings. From the two overlapping volumes the visibility histograms for PET and CT are calculated: Segmented ROIs in PET are used to calculate the visibility for each ROI to create a ROI-based visibility histogram (green-pink bars); and for CT, the visibility histogram (green bars) is calculated using its intensity range and drawn on top of its intensity histogram. As the user redefines the opacity mapping from one of the transfer functions, the visibility histograms are recalculated and show the user its changes on

the PET and CT visibility histograms; and (c) based on the user ROI selection, our transfer function operates by automatically optimising the transfer function of the CT to minimise its occlusion to the selected ROI. In this example, the selected tumour ROI (red box) in the lung is made more apparent compared to the rendering in (b), with increased visibility of the PET ROI (compared to (b)) and reduction in the CT visibilities belonging to the lung structures (first peak).

3.2 ROIs Segmentation

We have adopted an automated PET thresholding segmentation algorithm [22] to construct our ROIs based on the PET response criterion (PERCIST) [21]. PERCIST is an approach to derive metabolic changes in malignant lesions. It uses a reference region placed on the right lobe of the liver or the descending aorta in the case of abnormality in the liver, with the metabolism in the region used to calculate a threshold value for PET segmentation. This threshold value can then be used to separate malignant tumours from non-tumours. However, because it is a thresholding approach, it also results in the inclusion of other non-relevant regions that are above the threshold, such as the brain, bladder, and kidneys. From the thresholded results, all grouping of voxels (greater than 10 voxels) were defined as an ROI.

3.3 ROIs-based Multi-modality Visibility Histogram

Visibility histograms were introduced to guide transfer function manipulation, which represent the opacity contribution of intensity in the histogram to the rendered image. This metric is computed by measuring the visibility of each voxel in a viewpoint and then adding the visibility to the corresponding bins in the histogram [13]. In this study, we have extended the visibility histogram calculation into two overlapping volumes as illustrated in Fig. 2 where the ray passes through voxel pairs in PET and CT volumes, V_{PET} and V_{CT}. The combined visibility for a coordinate position p, $T(p)$, is the accumulation of the opacity contributions of all the voxel pairs starting from a view-point E to p (front-to-back composition) according to:

$$T(p) = e^{-\int_P^E A(t)dt} \tag{1}$$

$$A(p) = A(p - \Delta p) + \big(1.0 - A(p - \Delta p)\big) * O(p) \tag{2}$$

where $A(p)$ is the composition of opacities from the voxel pair, and $O(p)$ is the voxel's opacity. The visibilities of the voxels, computed as the product of the voxels' opacities, were then added into their corresponding bins of the visibility histogram $VH[x]$ with x as the bin number such that:

$$VH[x] = \sum_{p \in x} O(p) \cdot T(p) \tag{3}$$

For a CT volume, visibility histogram, VH_{CT}, was constructed with N bins representing the entire intensity range. For a PET volume, each ROI occupied a bin in the visibility histogram, VH_{PET}. For improved presentation, we normalise our histograms into a logarithmic scale to reduce the wide-ranging quantities among the intensities.

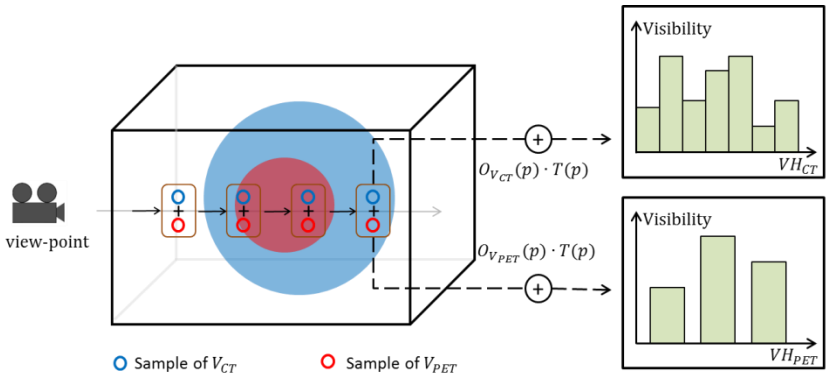

Fig. 2. Calculation of the ROI-based multi-modality PET-CT visibility histograms.

Multi-modality visibility histogram represents the visibilities of the composited voxels from multiple volumes. As such, these visibility histograms introduce '*composite dependency*', which measures how much of the voxel is visible, in addition to the visibility histogram's existing opacity and view dependencies. Fig. 3 demonstrates the application of our proposed multi-modality visibility histogram on two synthetic volumes. In this scene, the spheres contain low to high intensities spread out in Gaussian distributions from the centre to its outer rim. Here, a volume is a small sphere (*A*), and another volume consists of two larger spheres (B_1 and B_2), B_1 of which has the same centre but larger radius compared to *A* and therefore completely encapsulates the sphere. These spheres are assigned to different colours but the same opacities and are controlled by a pair of 1D transfer functions. In (i), only the red sphere, *A*, is visible with the opacities of the blue and green spheres, B_1 and B_2, set to zero. The increase in the opacities of B_1 and B_2 results in their visibility increasing while lowering the visibility of A (maximum visibility decreased from 24.10 to 9.56 (60.33%)), as shown in (ii), affecting both its rendering and its visibility histograms. In (iii), the two volumes are rotated such that the overlapping *A* and B_1 are now closer to the view-point and completely in front of B_2, thus resulting in zero visibility for B_2, even though its opacity is unchanged from (ii), while we see a large increase in the visibilities of the overlapping *A* and B_1 (46.97% increase for *A* and 49.37% for B_1). Finally, (iv) is the opposite of (iii) where B_2 occludes A completely and B_1 partially. Here, the visibilities of the two overlapping *A* and B_1 are at their lowest, due to them being occluded from B_2, while B_2 is at its highest visibility due to it being closer and fully visible in the current view-point. Note the zero visibility in the lower intensity ranges of B_1 in (iv) which represent the voxels closer to the occluded section of the sphere.

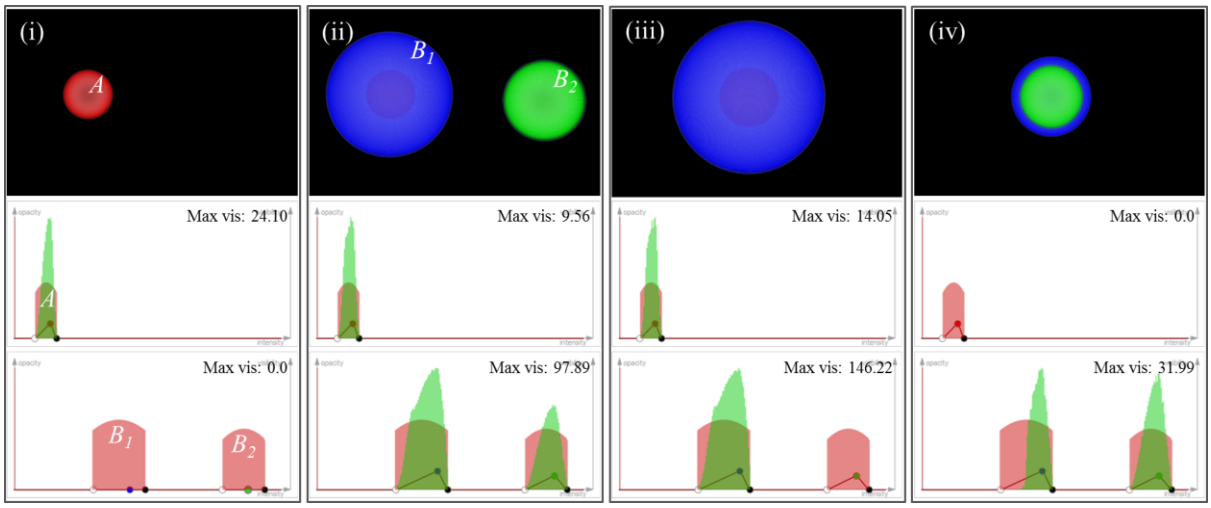

Fig. 3 Illustration of composite and view dependencies of multi-modality visibility histograms on two aligned volumes. A volume has small red sphere (*A*); top visibility histogram, whereas another volume consists of two larger blue and green spheres (B_1 and B_2); bottom visibility histogram.

3.4 Automatic Transfer Function Optimisation

Our m^2-vtf enables optimisation of CT transfer function based on maximising the visibility of an ROI from its counterpart PET. We formulate the manipulation of the CT transfer function as an energy minimization problem. Here, set of opacity parameters, θ, of an initial CT transfer function (preset) is optimised by minimising an energy function, *E*, consisting of: (i) tolerance of PET ROIs visibility defined as E_T and (ii) freedom of transfer function movements for CT defined as E_F. The tolerance of ROIs visibility is used to control how much of it is to be made visible in the resultant renderings. It can be calculated via sum of the squared difference between user-defined visibility, V_D, and resultant visibility, V_R, for the individual ROIs:

$$E_{\mathrm{T}}(\theta) = \sum_{n=1}^{N}(V_{D_n} - V_{R_n})^2 \qquad (4)$$

where N is the number of ROIs. The freedom of the opacity parameters is used to supress the voxels that are occluding the ROIs. The minimum and maximum in p_{th} opacity parameter, θ_p^{min} and θ_p^{max}, respectively, are limited such that:

$$E_F(\theta) = \sum_{p=1}^{P} \left|\theta_p^{min} - \theta_p\right|_+^2 + \left|\theta_p - \theta_p^{max}\right|_+^2 \qquad (5)$$

where P is the number of the parameters and a clamping operator $[x]_+$ = x if x > 0 or 0 otherwise. Finally, the optimal CT transfer function is generated with an iterative manner satisfying:

$$argmin_\theta \, wE_T(\theta) + (1-w)E_F(\theta) \qquad (6)$$

where w is the weight for visibility tolerance. Downhill simplex optimisation [23] was used to solve the minimisation problem for its effectiveness in non-linear optimisation with multiple local optima [24]. Compared to gradient-based approaches, it is less dependent on the initial transfer function [25] which is an important characteristic to ensure convergence. It is computationally more efficient as it only needs an energy function while retaining competitive performance depending on the application, in terms of convergence accuracy and computation time [26].

3.5 User Interactions for ROI Definition

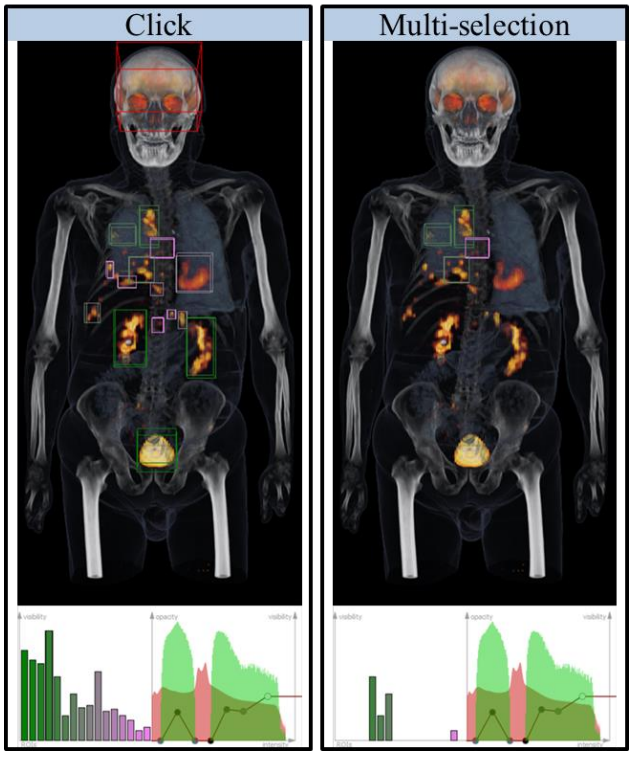

Fig. 4 ROI selection in our m^2-vtf

A key requirement for any ROI-based visualisation is an intuitive and simple selection of the ROIs. In our visualisation, we used PET segmentation to automatically derive ROIs and used it to construct the ROI-based visibility histogram. In this approach, each ROI was represented by a bar and the visibility among the voxels in that ROI was calculated into a single visibility metric. These bars were then sorted according to intensity and colour-coded green and pink (left to right in Fig. 4). The ROIs in the rendering were represented by a bounding box with the same colour code as its corresponding visibility bars. The ROI visibility histogram can be used to select the ROI by clicking on bars (a). Non-relevant regions can be selected and then removed. By removing all non-relevant regions, the resulting ROI visibility histogram is shown in (b) with only four ROIs, extensive glucose avid foci in the right hemithorax.

4 Implementation

4.1 Datasets and Pre-processing

PET-CT studies were acquired from a Siemens Biograph TruePoint PET-CT scanner at the Department of PET and Nuclear Medicine, Royal Prince Alfred (RPA) Hospital, Sydney, Australia. All studies have 326 slices with slice thickness of 3 mm to cover the whole body from the top of the head to the upper thighs. PET image slices were reconstructed to 168 x 168 matrix with pixel resolution of 4.07 mm^2 and CT slices were reconstructed to 512 x 512 matrix at 0.98 mm^2. The hardware co-registered PET was then resampled to the CT dimensions. The CT scans were processed to remove the background and bed/linen via adaptive thresholding and image subtraction from a bed template. The voxel intensity of the CT was in Hounsfield units (HU) with intensity ranges of 0 to 4095 (12 bit). For PET, we applied linear intensity normalisation and matched the intensity range of the CT. Low count intensity ranges (less than 10 voxels) for both volumes were truncated.

4.2 Visualisation Implementation

To be used as a visual feedback mechanism, m^2-vtf needed to be computed in interactive volume rendering. To achieve real-time performance, we used graphical process unit (GPU) optimised Voreen volume rendering library [27]. Voreen is an open source texture-based volume rendering engine that allows interactive visualisation of volumetric data sets with high flexibility for integrating new algorithms and optimisations. In the rendering framework, we used a fragment shader to calculate the visibility of the samples from the two volumes with a single raycast. Parallel processing, together with vertex fragment streams were used to calculate the visibility metrics of several slices simultaneously in a single-pass using frame buffer objects [13].

5 Results

5.1 Whole-body PET-CT studies

Fig. 5 demonstrates our m^2-vtf for whole-body PET-CT visualisation. In Fig. 5 individual PET (panel a) and CT renderings (panel b) with the visibility histograms and transfer functions are depicted. Here, the visibility histograms are individually calculated for each volume. The two volumes can be fused into a single scene as in (c) with the recalculation of the visibility histograms from the overlapping voxels. By using the transfer functions pair, individual PET and CT transfer functions can be manipulated while observing the effect of the visibility from changes in e.g., PET on the CT counterpart. However, a limitation is the inability to classify the tumour ROI in the PET volume, in this case the large mass in the lung, from other non-relevant regions that are also visualised. With the lack of classification ability, the visibility metric becomes irrelevant since all these structures share the same visibility. By using PET segmentation, we can compute the ROI visibility histogram and select only the tumour as in (d). The visibility of the PET is then recalculated to only represent the occlusions to the ROI. In this rendering, the tumour's visibility is not high due to the occlusion from the lung tissues and the rib cage (bones) in front of and/or overlapping the ROI. To reduce such occlusion, we can adjust the CT opacity corresponding to the occluding structures by manually manipulating its transfer function. This results in reduced opacities for the lung structures and also the bone structures in (e) (see grey arrows) while increasing the visibility of the PET ROI in (e) (indicated by red arrow).

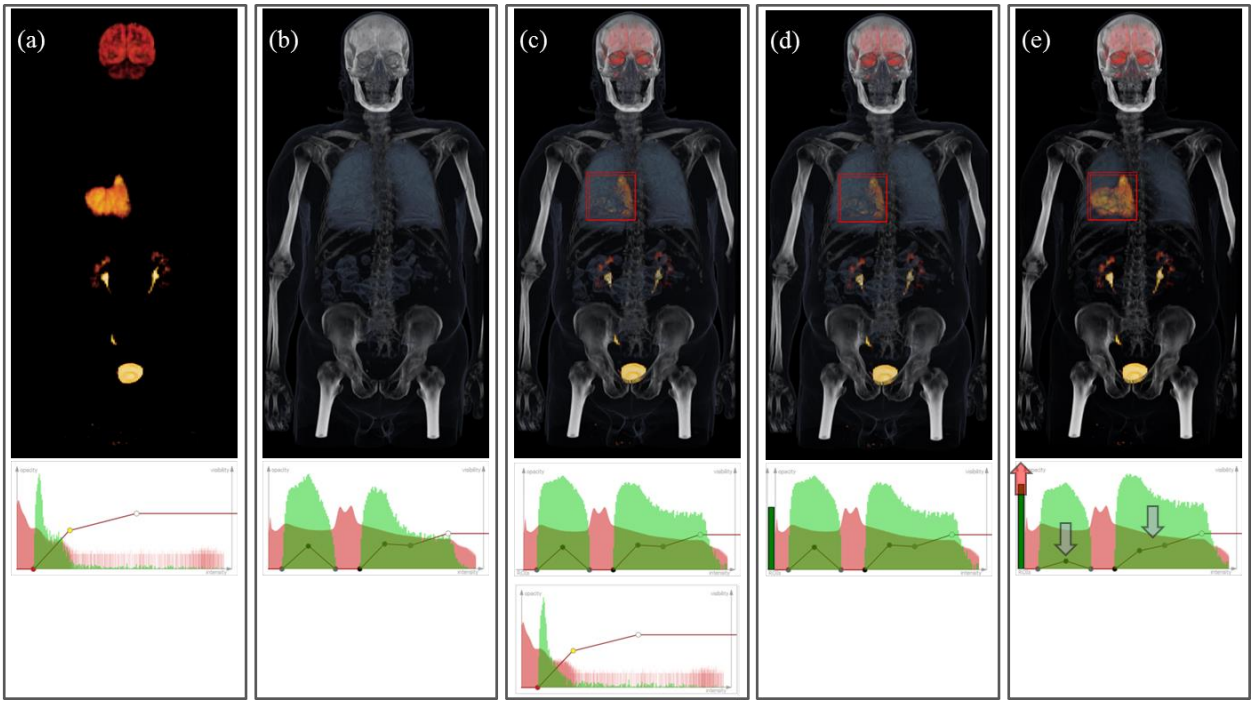

Fig. 5 A comparison of PET-CT visualisations between using a pair of single-modal visibility-driven transfer functions ((a) and (b)) and with m^2-vtf on a PET-CT study ((c) to (e)), with a large high-grade lung tumour patient study.

Our multi-modality transfer function automation is exemplified in Fig. 6. A tumour in the right lung from PET, which is faintly visible in (a), is selected as an ROI (from the ROI visibility histogram). Based on the selection, (b) to (d) represents the visualisations according to varying user-defined constraints for ROI's visibility tolerance. In (b), a low tolerance of 0.2 was set for the visibility of the ROI, and therefore resulted in CT structures predominating in the rendering. When the tolerance was increased to 0.5, and then to 0.8, respectively of (c) and (d), we can clearly notice the reduction of the lung and bone structures in the CT, thus making the PET ROI more visible. In this example, the second constraint of the CT transfer function freedom was set to 0.9 for all variations to give higher degrees of movement.

In Fig. 7, a whole-body PET-CT study of a patient with disseminated disease is visualised. The centre images (d) and (e) show the same PET-CT study in a coronal view (d) and a rotated view (e). The transfer function was set to visualise the bone, lung, and skin from the CT. For PET, it was again set to visualise the segmented ROIs. These ROIs can be selected individually or as a set, and used to automate the CT transfer function. We can see that the bulk of the disease is in the abdomen where there are very large masses on both sides; there are innumerable, markedly glucose avid regions in the bone marrow (flexible tissue found in the interior of bones) e.g., in (a), (c), (f), and (g); there is a soft tissue abnormality in right side of the upper thoracic spine as in (b). These findings are consistent with a high-grade lymphoma. A powerful feature in the visibility histogram is its view dependency, which calculates the visibility according to the user's viewpoint. In this study, we can see the changes in the ROI visibility from the rotation of the volume in (e) for PET, whereas there is only a small difference in the CT.

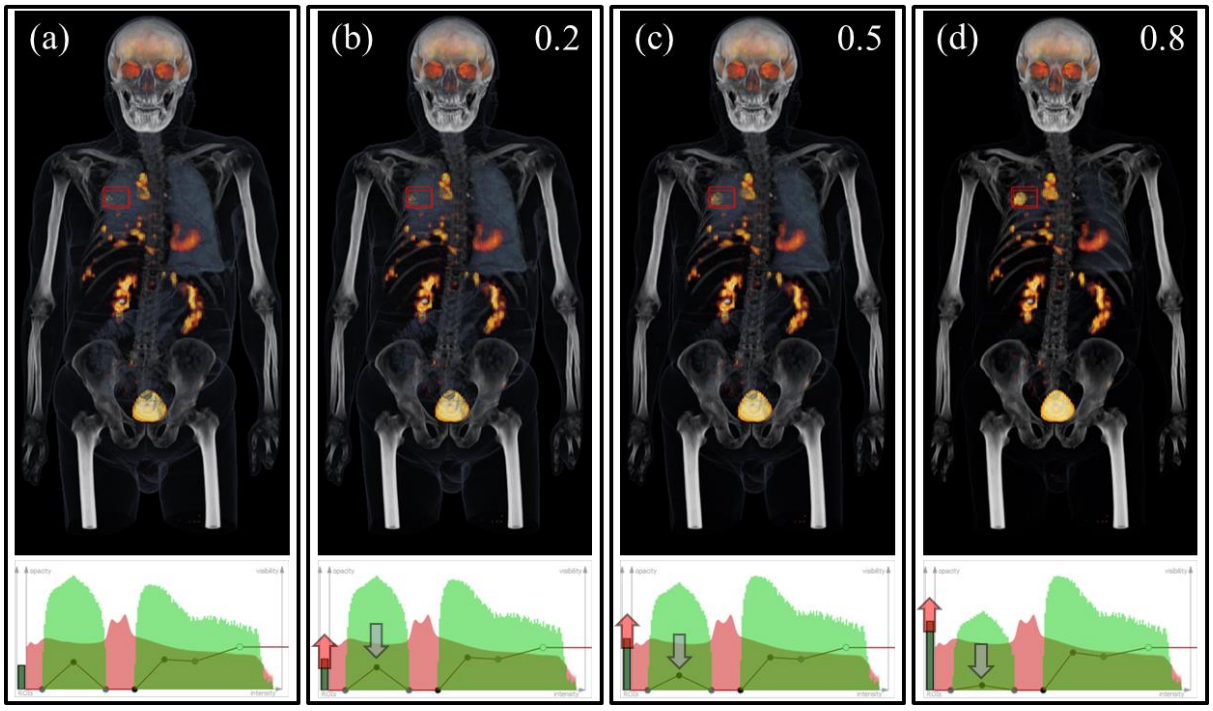

Fig. 6. Automated CT transfer function according to varying optimisation constraint of ROI's visibility tolerance from low (0.2) to high (0.8).

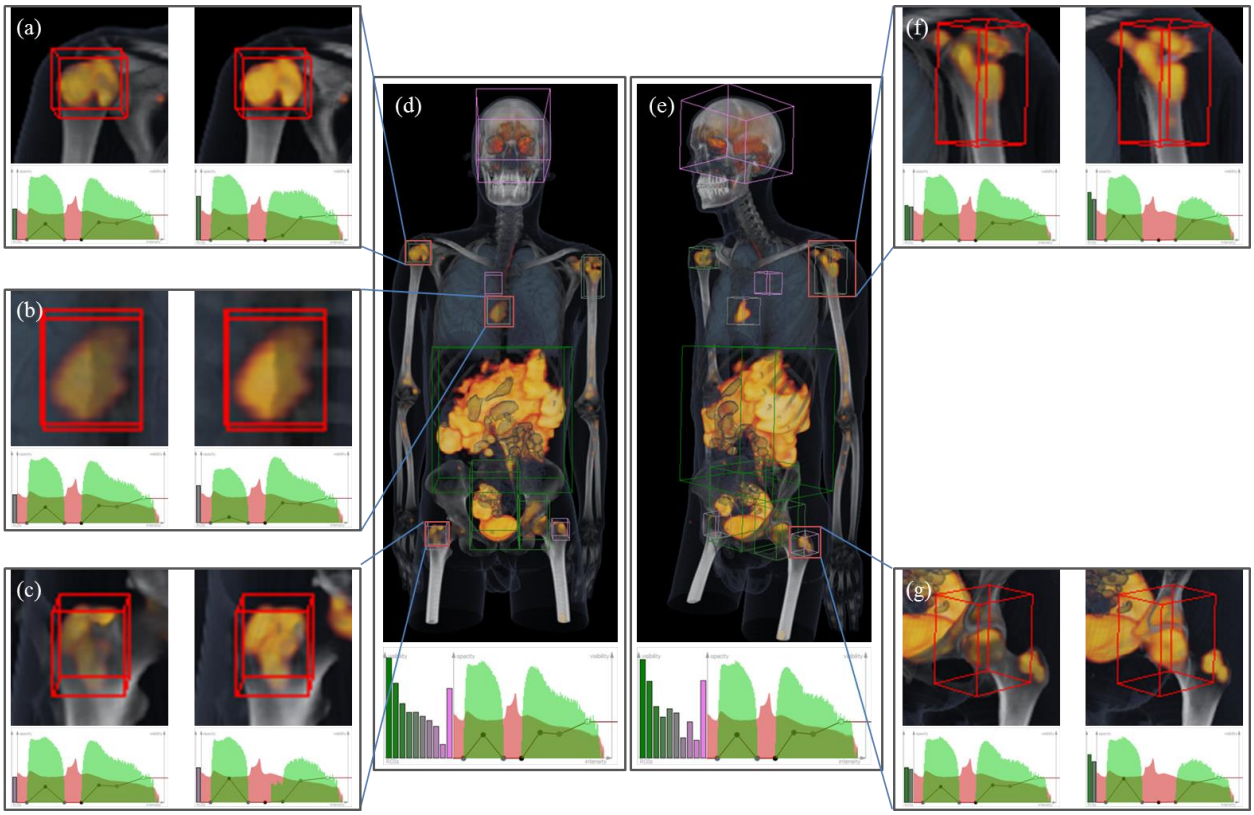

Fig. 7. A PET-CT study of disseminated disease consistent with a high-grade lymphoma, consisting of multiple abnormalities spread throughout the body. Our transfer function automation is used to visualise the different abnormalities based on the user-selected ROIs in (a) to (c) and multiple ROI selections of (f) and (g).

5.2 Computational Performance

We measured the average time, in frames per second (fps), using various volume sizes, for the following interactions: (i) the calculation of visibility histogram from random volume manipulations including view-point rotations, panning, and transfer function changes; (ii) calculations of view-independency visibility histogram based on averaging 6 rotations as in [13]; and (iii) single ROI-based transfer function automation. All computations were performed using a PC with nVIDIA GTX 590 1.5G GPU; Intel i7 CPU @3.20 GHz; running 64-bit Win 7.For PET-CT studies with dimensions of 256 x 256 x 256 x 2 volumes, which is an approximation of a typical study used in this paper, 10.14 fps was measured for various random volume manipulations, which enables our framework to achieve interactive visualisation of multi-modality PET-CT volumes. As expected, the interactivity decreases based on the volume dimensions, but with our algorithm still achieving minimal rates of 3.446 fps for 512^3 x 2 volumes manipulations. In regards to automation, it is interesting to note that regardless of the volume size, automation was relatively similar (from 2.369 to 0.980 fps). In the case of multiple ROIs selection, where we varied between 1 to 16 ROIs, we observed negligible differences (0.23 seconds) in its computation. This is because the visibility histogram construction is only based on the ROIs of PET which occupy far less bins when compared to the full PET intensity histogram. The reduction of the bins enables the computation to be performed in a single rendering pass.

We further computed view-independency as in [13], which is the average of the visibility calculations from six different viewing angles (y-axis rotation), in an attempt to remove the inherent limitations from view-dependency. In our results, we achieved 1.752 fps using 256^3 x 2 volumes and 0.467 fps for 512^3 x 2 volumes. These results show that with our multi-modality ROI-based visualisation, we can easily compute the view-independent visibility histogram after changes to its transfer function.

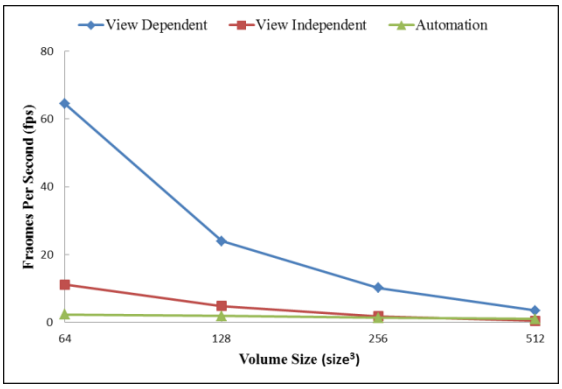

Fig. 8. Measured fps for volume rendering manipulations (both view dependent and independent) and also in calculating our multi-modality transfer function optimisation (automation).

6 Discussions and Future Work

This study presented a PET-CT visualisation with ROI-based visibility-based multi-modality transfer function manipulations and automation. Although we have presented our results designed specifically for PET-CT visualisation, we suggest that our approach is applicable to any other co-aligned multi-modality images (volumes aligned to the same spatial coordinates). This alignment can be from a multi-modality scanner or through software registration algorithms [28]. Our future studies will investigate the potential of our visualisation with different multi-modality volumes such as PET and magnetic resonance imaging (PET-MRI).

In our results, we used a pair of 1D transfer functions for PET and CT volumes. We demonstrated that the use of 1D-based visibility histogram, together with ROI segmentation, was able to provide an efficient approach to visualising PET-CT volumes. Nevertheless, our multi-modality visibility

calculation is not restricted to 1D transfer functions and it can accommodate multi-dimensional transfer function, such as gradient-based approach [13].

In our experiments, we adopted a well-known and robust PERCIST segmentation for the automated delineation of abnormalities as ROIs from PET volumes. PERCIST has the advantage of being a simple yet robust and reliable algorithm on inherently noisy and low signal-to-noise (SNR) PET images [21]. By using a fully automated approach, a user is able to interactively and conveniently select the ROIs to visualise without relying on a complex and manual ROI segmentation process. Although we only used PERCIST segmentation algorithm, our visualisation is not restricted to this algorithm. It can, therefore support algorithms that are disease / modality specific, e.g., using region growing algorithm for brain PET tumour ROIs [29]. Our future work will investigate the integration of segmentation algorithm directly into our rendering framework, by leveraging the general purpose GPU (GPGPU) computation, and thus being able to interactively perform the segmentation in real-time volume rendering.

7 Conclusions

We proposed a new visualisation for multi-modality PET-CT by integrating visibility-driven transfer function together with ROI selection from automated segmentation. Our visualisation enables the selection of an ROI from PET which is then used to calculate the visibility-driven transfer function automation from its counterpart CT. Our results using clinical PET-CT studies with different types of diseases demonstrated the capabilities of our visualisation for multi-modality volume rendering

Acknowledgements

We would like to thank our collaborators at the Royal Prince Alfred (RPA) Hospital. This research was funded by Australian research council (ARC) grants.

References

1. Rosset, A., Spadola, L., Pysher, L., Ratib, O.:, Navigating the fifth dimension: innovative interface for multidimensional multimodality image navigation., Radiograph, 26(1), 299-308 (2006).
2. Preim, B., Bartz, D.,: Visualization in medicine theory, algorithms, and application., Morgan Kaufmann Series in Computer Graphics, (2007).
3. Kindlmann, G., Durkin, J.W.:, Semi-automatic generation of transfer functions for direct volume rendering., In: Proc. IEEE Visualization, pp. 79-86 (1998).
4. Kindlmann, G., Whitaker, R., Tasdizen T., Moller, T.:, Curvature-based transfer functions for direct volume rendering: methods and applications., In: Proc. IEEE Visualization, pp. 513-20 (2003).
5. Kniss, J., Kindlmann, G., Hansen, C.:, Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets., In: Proc. IEEE Visualization, pp. 255-62 (2001).
6. Correa, C.D., Ma, K.:, The occlusion spectrum for volume classification and visualization., IEEE T. Vis. Comput. Gr., 15(6), 1465-72 (2009).
7. Caban, J.J., Rheingans, P.:, Texture-based transfer functions for direct volume rendering., IEEE T. Vis. Comput. Gr., 14(6), 1364-71 (2008).
8. Correa, C.D., Ma, K.:, Size-based transfer functions: a new volume exploration technique., IEEE T. Vis. Comput. Gr., 14(6), 1380-7 (2008).
9. Cai, W., Sakas, W.:, Data intermixing and multi-volume rendering., Comput. Graph. Forum., 18(3), 359-68 (1999).
10. Bramon, R., Bardera, A., Rodriguez, J., Feixas, M., Puig, J., Sbert, M.:, Multimodal data fusion based on mutual information., IEEE T. Vis. Comput. Gr., in press (2012).

11. Haidacher, M., Bruckner, S., Kanitsar, A., Grller, M.E.:, Information-based transfer functions for multimodal visualization., In: Proc. Visual Comput. Biomed., pp. 101-8 (2008).
12. Kim, J., Eberl, S., Feng, D.:, Visualizing dual-modality rendered volumes using a dual-lookup table transfer function., COMPUT SCI ENG, 9(1), 20-5 (2007).
13. Correa, C.D., Ma, K.:, Visibility histograms and visibility-driven transfer functions., IEEE T. Vis. Comput. Gr., 17(2), 192-204 (2011).
14. Ruiz, M., Bardera, A., Boada, I., Viola, I., Feixas, M., Sbert, M.:, Automatic transfer functions based on informational divergence., IEEE T. Vis. Comput. Gr., 17(12), 1932-41 (2011).
15. Wang, Y., Zhang, J., Chen, W., Zhang, H., Chi, Xuebin.:, Efficient opacity specification based on feature visibilities in direct volume rendering., Comput. Graph. Forum., 30(7), 2117-26 (2011).
16. Jung, Y., Kim, J., Feng, D.:, Dual-modal visibility metrics for interactive PET-CT visualization, Proc. IEEE EMBC, 2696-9 (2012).
17. Bordoloi, U., Shen, H.W.:, View selection for volume rendering., In: Proc. IEEE Visualization 487-94 (2005).
18. Viola, I., Kanitsar, A., Eduard, M.:, Importance-driven feature enhancement in volume visualization., IEEE T. Vis. Comput. Gr., 11(4), 408-18 (2005).
19. Tzeng, F.Y., Ma, K.:, A Cluster-space visual interface for arbitrary dimensional classification of volume data., In Eurographics. IEEE TCVG Symposium on Visualization, pp. 17-24 (2004).
20. Kim, J., Cai, W., Eberl, S., Feng, D.:, Real-time volume rendering visualization of dual-modality PET/CT images with interactive fuzzy thresholding segmentation., IEEE T. Info. Tech. Biomed., 11(2), 161-9 (2007).
21. Wahl, R., Jacene, H., Kasamon, Y., Lodge, M.:, From RECIST to PERCIST: evolving considerations for PET response criteria in solid tumors., J. Nucl. Med., 50(suppl 1), 122S-50S (2009).
22. Bi, L., Kim, J., Wen, L., Feng, D.:, Automated and robust PERCIST-based thresholding framework for whole body PET-CT studies., Proc. IEEE EMBC, 5335-8 (2012).
23. Nelder, J.A., Mead, R.:, A simplex method for function minimization., COMPUT J., 7, 303-13 (1965).
24. Barton, R.R., Ivey, J.S.:, Nelder-mead simplex modifications for simulation optimization, INFORMS., 42(7), 954-73 (1996).
25. Lagarias, J.C., Reeds, J.A.:, Convergence properties of the nelder-mead simplex method in low dimensions, SIAM J. OPTIMIZ., 9(1), 112-47 (1998).
26. Manousopoulos, P., Michalopoulos, M.:, Comparison of non-linear optimization algorithms for yield curve estimation, EUR J. OPER. RES., 192(2), 594-602 (2009).
27. Voreen: Volume Rendering Engine, http://www.voreen.org/.
28. Maintz, J.B.A., Viegever, M.A.:, A survey of medical image registration., Med. Imag. Anal., 2(1), 1-36, 1998.
29. Kim, J., Feng, D.D., Cai, T.W., Eberl, S.:, Automatic 3D temporal kinetics segmentation of dynamic emission tomography image using adaptive region growing cluster analysis, IEEE Proc. NSS-MIC., pp 1580-83 (2002).