Eurographics Workshop on Visual Computing for Biology and Medicine (2015)
K. Bühler, L. Linsen, and N. W. John (Editors)

Illustrative Multi-volume Rendering for PET/CT Scans

Kai Lawonn[1,2,3], Noeska Smit[2], Bernhard Preim[1,3] and Anna Vilanova[2]

[1]University of Magdeburg, Germany; [2]TU Delft, The Netherlands; [3]Research Campus STIMULATE

Abstract

In this paper we present illustrative visualization techniques for PET/CT datasets. PET/CT scanners acquire both PET and CT image data in order to combine functional metabolic information with structural anatomical information. Current visualization techniques mainly rely on 2D image fusion techniques to convey this combined information to physicians. We introduce an illustrative 3D visualization technique, specifically designed for use with PET/CT datasets. This allows the user to easily detect foci in the PET data and to localize these regions by providing anatomical contextual information from the CT data. Furthermore, we provide transfer function specifically designed for PET data that facilitates the investigation of interesting regions. Our technique allows users to get a quick overview of regions of interest and can be used in treatment planning, doctor-patient communication and interdisciplinary communication. We conducted a qualitative evaluation with medical experts to validate the utility of our method in clinical practice.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation I.3.8 [Computer Graphics]: —Applications

1. Introduction

Positron emission tomography (PET) is a medical imaging modality that is used to detect metabolic activity and is able to highlight functional pathologies. However, exact localization of these abnormalities is not possible since PET is only able to display metabolic information and not able to distinguish different tissue types. X-ray-based computed tomography (CT) scanners provide structural anatomical information at a much higher spatial resolution, but lack the metabolic information that PET can provide. Therefore, combining the PET information with CT scans can depict metabolic abnormalities enriched with the spatial information of the anatomical imaging provided by CT.

Despite research projects aiming at fused 3D visualizations for clinical practice, radiologists currently explore the stack of two superposed 2D images to analyze abnormalities in a slice-based approach for PET/CT. They are able to browse through the stack of slices and specify different image properties, e.g., window and level to adjust the contrast and examine structures of interest. The major problem with this approach is that physicians have to mentally fuse all information. As the number of slices per CT scan increases, browsing through all of these slices manually becomes more involved and time-consuming. An overview could potentially help to guide the user towards areas that warrant closer examination at a glance. Furthermore, patients and other specialists might be less accustomed to the slice-based representation and might benefit from a more familiar representation.

To help alleviate these problems, we provide a focus-and-context visualization for PET/CT images, see Figure 1. By combining direct volume rendering with our illustrative visualization techniques, we aim at improving shape perception and minimizing occlusion. We visualize the CT scans as an anatomical context using line drawing techniques, while the PET scans are visualized as the focus object, drawing the attention of the viewer to the metabolic activity. Our visualization technique allows for a fast overview visualization that can be easily applied to PET/CT data and extended to other focus-and-context applications. To validate the utility of our visualization technique and its application to PET/CT data, we conducted a qualitative evaluation with medical experts. With this, the contributions of our work are as follows:

- An illustrative focus-and-context visualization for multi-volume rendering applied to PET/CT data.
- An application that demonstrates the utility of our visualization technique.
- An evaluation with domain experts that confirms the advantages of our work.

DOI: 10.2312/vcbm.20151213

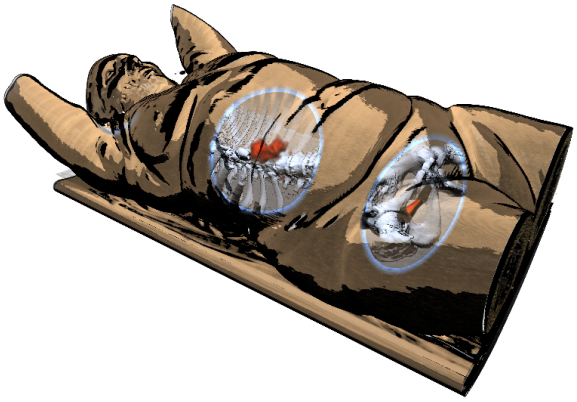

Figure 1: *Our illustrative visualization method encodes PET as a focus object, while the CT allows surrounding bone structures and skin to be visualized as the context.*

2. Related Work

Multi-volume/variate volume rendering: Much work has been published regarding rendering of multiple overlapping volumes as well as rendering multi-variate datasets [CS99, KHGR02, WLM02, KSW*04, FPT04, HBKS05, BVPtR08]. These approaches traverse both volumes with different compositing options and transfer functions or attempt to fuse multiple volumes together into a single volume. A general information-based user-friendly approach for transfer functions was introduced by Haidacher et al. [HBKG08]. Recent work can be found in the state-of-the-art report on the visualization of multivariate data presented by Fuchs and Hauser [FH09]. Several works have developed and used multi-volume rendering techniques specifically to visualize multi-modal medical data. Rößler et al. described a GPU-based multi-volume rendering scheme that allows users to visualize an arbitrary number of volumes interactively, specifically focused on functional brain images [RTF*06]. For planning neurosurgical procedures, Beyer et al. introduced a framework that can visualize multimodal volumes [BHWB07]. In contrast to our work, cutaways are user-indicated instead of automatically generated. Nguyen et al. used an approach to visualize and interact with real-time fMRI data [NEO*10]. We use the previous works on multi-volume rendering as a basis for our technique and further tailored it to fit the specific needs of PET/CT visualization. Since the PET activity is vital for diagnosis, it is important that it is not occluded by other structures as it may occur with existing multi-volume rendering techniques.

Focus-and-Context visualization: Several works have been published on visualization techniques that are also applicable to multi-volume rendering, such as focus-and-context techniques and dynamic cutaways. Viola et al. proposed an importance-driven approach that highlights embedded parts of the volume that are more important than the regions that would occlude it [VKG04]. Krüger et al. presented Clearview, a context-preserving hotspot visualization technique [KSW06]. They use texture-based volume rendering techniques to provide a focus-and-context visualization from a single volume in contrast to our multimodal approach. Burns et al. introduced an importance-driven approach to visualize multiple volume data sets [BHW*07]. They used novel importance specifications combined with cutaway surfaces. A hybrid visualization approach of anatomical and functional brain data was presented by Jainek et al. [JBB*08]. Various rendering styles were used, e.g., ambient occlusion to enhance the visual output. For an overview of general focus-and-context visualization techniques, we refer to Bruckner et al. [BGM*10]. Additionally, the advantages of recent illustrative visualization techniques are described by Lawonn [Law15]. Based on the previous work in this area, we have developed an illustrative focus-and-context rendering technique specifically designed for PET/CT data, where the PET is used as the focus area, and the CT is used as an anatomical context.

PET/CT Visualization: Previous work has also been done specifically to visualize PET/CT data and its combination of metabolic and anatomical information. Stokking et al. presented a visualization method that combines functional input data and a surface extracted from anatomical data [SZV01]. Kim et al. introduced a dual-lookup table for PET/CT data such that medical experts can set different transfer functions for every volume in an single view [KEF07]. Bramon et al. used an approach to fuse different image modalities together and applied their technique to CT, MRI and PET data [BBB*12]. This information-theoretic framework automatically selects the most informative voxels from two volume data sets. Jung et al. employed a novel visualization approach by integrating a visibility-driven transfer function specifically for PET/CT data [JKE*13]. Furthermore, they provided an intuitive region of interest selection tool for further exploration. While these different approaches are suitable for PET/CT applications, in our work we minimize the amount of pre-processing required on the datasets and can directly work with the originally acquired scans without segmentation. Additionally, we designed a simple transfer function setting for PET data that is able to highlight the PET activity of interest using a single value.

3. Medical Background

We provide a short overview on CT, PET and PET/CT acquisition and its applications. For a more elaborate description, we refer to Townsend et al. [TCYH04]. Finally, we present a requirement analysis for visualizing PET/CT data.

Computed Tomography (CT): In CT, cross-sectional (tomographic) images are created by rotating an X-ray source around the patient. In this way, the whole body or parts of it can be scanned. The X-rays are detected by sensors after being attenuated by the various tissue types in the human

body. Tissues such as bone will absorb more of the X-rays, while softer tissues will absorb less and this difference creates contrast between the various tissues. To further enhance the visibility of vessels, a contrast agent that is radio-opaque can be administered intravenously.

Positron Emission Tomography (PET): PET is also a tomographic imaging modality, but unlike CT, PET relies on the detection of gamma rays emitted by the patient after a positron-emitting radionuclide (tracer) is introduced. Using this technique, metabolic processes can be visualized. While CT scans can provide detailed anatomical data, PET scans are able to reveal functional information. This information is typically displayed in a slice-based approach or visualized using a maximum intensity projection (MIP). A common application of PET scans is to search for metastases and for this the radioactive substance fluorodeoxyglucose (FDG) tracer is used, a substance similar to glucose. The metastases have higher glucose uptake than normal and specific abnormal metabolic activity can be captured in this way. Besides oncological applications, PET can also be used for neurological and cardiological applications.

PET/CT: PET scanners exhibit a relatively low spatial resolution, low contrast, higher levels of noise and require long scanning times. As a functional imaging modality, it is not directly suitable to image anatomical structures, but is able to detect metabolic activity. The absence of identifiable anatomical structures complicates localization of foci that have abnormal uptake. CT scanners on the other hand provide good quality high spatial resolution imaging quickly. They are well suited for visualizing anatomical and structural information, but cannot provide metabolic information. By combining PET and CT scans of a single patient, physicians are able to detect (abnormal) metabolic activity using the PET scan and to localize this activity using the CT scan. Clinical oncological applications of the combined PET/CT scanner include diagnosing and staging primary malignancies as well as localization of metastasic disease in almost any region of the body [BTB*00]. Further applications include decision making about the surgical operability of the tumor or best course of treatment. PET/CT can also be used to determine if cancer has recurred or to determine the difference between scar tissue and active cancer tissue. We refer to Kluetz et al. for a case study of several applications of PET/CT scanners in clinical practice [KMV*00].

3.1. Requirement analysis

There are several options for visualizing combined PET/CT images. In the current clinical workflow, the typical approach is to examine the CT images in grayscale with the PET images superimposed using a colormap in 2D. Browsing through all the 2D slices individually to examine metabolic uptake can become time-consuming and cumbersome as the number of slices increases. Besides the 2D techniques, 3D techniques can additionally be used to give an overview of the full datasets at a glance. Existing methods such as maximum intensity projection (MIP), can provide such an overview, but suffer from depth perception issues. In case of PET/CT data, since the intensities are not comparable, a combined volume MIP can suffer from occlusion by the highest intensity regions. Ideally, a combined volume 3D technique would allow the users to get a quick overview of areas of interest and to localize foci of suspicious metabolic activity in an anatomical context. The requirements for such a technique are the following:

Req. 1. : It should show the combination of the two modalities in a single view in which the PET activity is always visible.

Req. 2. : It should relate the metabolic activity to nearby anatomical structures for accurate localization.

Req. 3. : Visualization options should be easily adjustable to fit the needs of the pathology being examined.
Such a technique can be used to guide the exploration of the datasets, bringing the attention to regions of interest, after which detailed inspection of these regions can be performed in 2D images. Furthermore, it could be beneficial for treatment planning, doctor-patient communication and interdisciplinary communication.

4. Method

In this section, we describe the different illustrative visualization techniques designed to depict the combined PET/CT datasets. Based on the requirements specified in the previous section, we divided the visualization tasks into three areas. The PET activity should be the focus area and always visible (Req. 1). The CT scan is used as an anatomical context, and can provide structural information revealing the skeletal structures near the focus (Req. 2). Further away from the focus region, the structural information can be abstracted to show only the skin to aid shape perception and orientation. Our technique then consists of a focus-and-context visualization and can be subdivided in three steps:

1. Illustrating the context, e.g., the skin.
2. Illustrating the focus, e.g., the PET activity.
3. Visualizing the region closely surrounding the focus area.

As a reminder for the reader, in the rest of the paper we will use the following notations. The Houndsfield values of the voxels in the CT scan data are given by the function $f : \mathbb{R}^3 \to \mathbb{R}$ and represent the tissue density. For further calculations, we need the curvature measurements, which are based on the work by Kindlmann et al. [KWTM03]. The gradient ∇f of the CT data is determined using central differences, which gives us the normal vectors: $\mathbf{n} = -\nabla f / \|\nabla f\|$ of the dataset. Afterwards, we calculate the Hessian $\mathbf{H}(f)$ of f by central differences again. Then, we set $\mathbf{P} = \mathbf{I} - \mathbf{n}\mathbf{n}^T$ and determine $\mathbf{G} = -\mathbf{PHP}$. This gives us the curvatures κ_1, κ_2

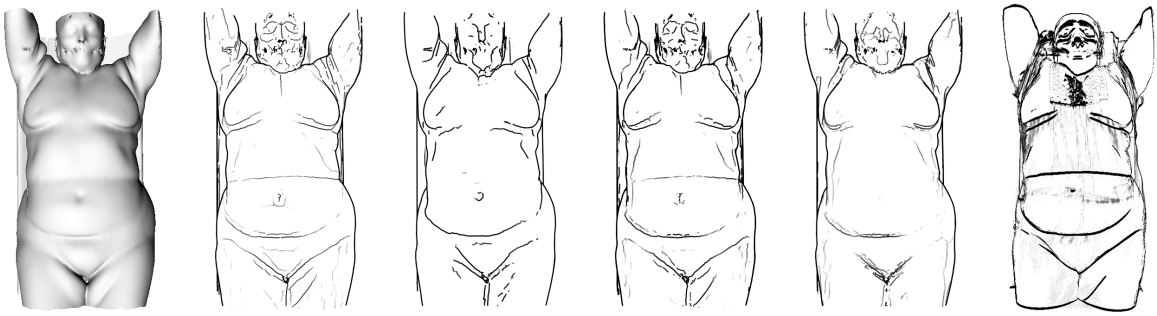

(a) Sh (b) RV (c) SC (d) AR (e) PEL (f) LG

Figure 2: *The shaded model (Sh) with different feature line techniques:* Ridges & Valleys *(RV),* Suggestive Contours *(SC),* Apparent Ridges *(AR),* Photic Extremum Lines *(PEL), and our approach* LG. *Compared to the feature line techniques, our approach enhances salient regions whereas the feature line methods detect lines based on different feature measures.*

by determining the trace: $tr(\mathbf{G}) = \kappa_1 + \kappa_2$ and the Frobenius norm: $|\mathbf{G}|_F = \sqrt{\kappa_1^2 + \kappa_2^2}$. Using the quadratic formula, the curvatures can be computed. Principle curvatures directions are then given by the eigenvectors of $\mathbf{G}$.

4.1. Context Visualization

We decided to use a combination of shading styles for the context object, e.g., the skin given by the CT dataset. For this, we have a given transfer function to determine the structure of interest. Afterwards, we illustrate this structure by using toon shading. The light source is placed top left and it follows the camera, as in medical atlas illustrations normally no headlights are used [Hod03]. To provide an enhanced contextual visualization, we used additional features lines. Most feature line techniques are defined on triangulated surface meshes, where a feature scalar-field is determined. The isolines of this scalar-field then yield the feature lines. As we are working with volumetric datasets, we decided to determine feature regions instead of lines. For this, we calculate the light gradient, which is obtained by determining the shading values per voxel $l = \langle \mathbf{n}, \mathbf{v} \rangle$, where $\mathbf{v}$ is the normalized view vector. Afterwards, we calculate $\|\nabla l\|$ (LG) and whenever this value exceeds a certain user-defined threshold the corresponding regions are illustrated in black. As a remark, sometimes a CT dataset features too much noise to get reasonable curvature and light gradient values. Therefore, we applied an anisotropic diffusion filter on these datasets for the purpose of determining the feature regions only. Figure 2 shows a comparison of the *PETCETIX* [osi] dataset rendered with different feature lines techniques. Here, the shaded model (Sh) is illustrated with the *Suggestive Contours* ([DFRS03]), *Ridges & Valleys* ([IFP95]), *Apparent Ridges* ([JDA07]), *Photic Extremum Lines* ([XHT*07]), and our approach *LG*. For the comparison, we extracted the surface of the volume dataset and applied the feature line techniques to it. Unlike the previous approaches, we only take the light gradient into account. The results of our approach are comparable to the presented feature line techniques. Furthermore, we add valley lines in white and thickness-controlled contours as suggested by Kindlmann et al. [KWTM03]. For the contour, we use the curvature along the view direction κ_v and determine if $|\langle \mathbf{n}, \mathbf{v} \rangle| < \sqrt{T\kappa_v(2 - T\kappa_v)}$ holds, where we set $T = 3$.

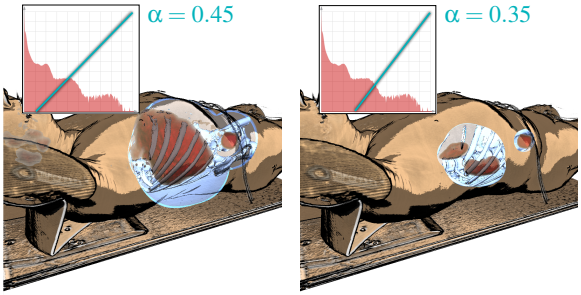

Figure 3: *Our initial TF for the PET dataset is defined using only a single value. After defining the* α *value, the TF can further adjusted if needed.*

4.2. Focus Visualization

Since we want the focus region to draw the attention of the viewer immediately, the PET activity visualization is rendered by using standard Phong shading in red, see Req. 1. The contour region of the PET activity is illustrated in white, similar to Fresnel shading. We used this red color to ensure that the experts' attention is drawn to these regions. Since this hue immediately stands out among the other colors, red is in our case used as a pre-attentive visual feature. For the transfer function (TF) used in our PET visualization, we used a single value $\alpha \in [0,1]$ to set an initial TF. It starts with a vertical line when $\alpha = 0$ and as the alpha increases,

the slope linearly decreases, as shown in Figure 3. This simplified transfer functions setting fulfills Req. 3.

4.3. Contextual Visualization of the Surrounding Focus

To illustrate the surroundings of the PET enhancement, we decided to use the lens concept as it was introduced by Bier et al. [BSP*93]. For a survey on interactive lens visualizations, we refer to Tominski et al. [TGK*14]. Instead of using interactive lenses, we used an automatic approach to set the position of the lens. For this, we used the PET scan and applied Gaussian blurring. This results in a severely blurred version of the PET foci. Afterwards, we can use a similar TF as used for the original PET on the smoothed PET volume. We use the blurred result to create an inside view with a similar effect as cut-away views. During raycasting we detect if the ray hits the surrounding region. If this is the case, we apply a different shading technique as well as a different TF on the structures in this region. For these regions, we employ Phong shading and as a standard TF that illustrates bones as anatomical context. Thus, whenever the ray hits the surroundings a different visualization technique is applied. This gives us the illusion of having a window that allows the user to look inside the body to reveal local context structures. For further emphasis on the lens concept, we also use an additional illustration on the skin inspired by Lerotic et al. [LCMY07]. To create a smooth transition between the context structure and the surrounding focus, we decrease the intensity of the skin relative to the distance to the PET activity. By doing this, the skin slightly disappears for enhancement of the surrounding structures, e.g., the bones. In this way, bones close to the border are more faded than structures closer to PET activity. Furthermore, we use a blueish color to emphasize the active lens area and add a blue border around the lens. However, the feature lines are always drawn on top of the lens to represent the skin surface. Finally, we illustrate the PET data and mix it with the skin illustration. In order to do this, we determine the distance R from the skin along the view direction to the PET data. The distance is then used as a weighting parameter for the transition of the skin with the PET. The closer the PET is to the skin, the brighter it appears. Here, we use the idea of the Rayleigh scattering and mix the image data with $1/R^2$ as a weighted coefficient. In case the PET activity is occluded by the bone structure, we slightly change the brightness of these regions such that the PET regions are still visible.

5. Implementation

The implementation of our algorithm consists of several steps, see Figure 4. For fast rendering, we use the GPU and perform all computations using the OpenGL shader framework in order to be independent of graphics card vendors. In the first step, we apply an anisotropic diffusion smoothing on the CT dataset. We use an anisotropic method to ensure that edges are mostly retained during the smoothing process. The smoothed result is more appropriate for the gradient calculation and for the illustrations of the feature lines. For the PET data, we use a simple Gaussian smoothing to define the surrounding region. In the next steps, we use standard direct volume rendering to composite the intensities along a ray. For the sampling of the ray, we use equally-spaced distances to calculate the intensities. The interpolation is achieved by using trilinear interpolation at the sampling points.

6. Results

To confirm the utility of our method, we applied our illustrative technique on different clinical and research-oriented datasets. We used the DICOM sample image sets that are provided by the publicly available *Osirix* [osi] database. First, we examined the *PROSTATIX* dataset, consisting of a PET/CT study with F18-fluorocholine in a male patient with prostate cancer. The results of our technique on this study can be seen in Figure 3. The enhancement shows high-activity regions in the pelvis and abdomen. Using the surrounding illustration, the location of these regions can be identified. Combined with the structural information from the CT scan, normal metabolic uptake can be differentiated from suspicious regions. Next, we applied our visualization technique to the *COMUNIX* dataset, which consists of a PET/CT study of a patient with a neck tumor, as can be seen in Figure 5. There are two regions of PET activity, one in the neck and one in the brain. By adjusting the transfer function for the CT scan, we can either emphasize the skin or the bones as CT context information to determine the proximity of the PET activity with respect to the skin or bones. As a third dataset, we selected the *MELANIX* dataset, consisting of a PET/CT study of a patient with melanoma (see Figure 9(c)). Here, the PET activity is high in several regions, such as the armpit, heart, abdomen and pelvis. The CT-based context visualization allows us to localize the PET activity in the armpit as right in front of the shoulder joint. Lacking further information about the patient, it could be either activity in the primary tumor or a metastasis to a lymph node. In Figure 6(a) our visualization technique is applied to the PETCETIX dataset once more. Since this study also features contrast-enhanced CT scans, it is also possible to visualize vascular information as the CT context instead of the bones. In this way, it might be possible to see if there is any tumor angiogenesis.

To test the flexibility of our technique in another context, we applied it to the *PANORAMIX* CT dataset that contains contrast-enhanced vascular structures. This dataset contains CT scans of a patient with an abdominal aortic aneurysm (AAA). Instead of using the PET region as the focus object, we segmented the aorta and use this as the main object of interest, using the rest of the CT scan as the context to visualize the surrounding bones and skin (see Figure 6(b)). The location of the aneurysm is visible using our technique, and the context information provided by the bones and skin, can help

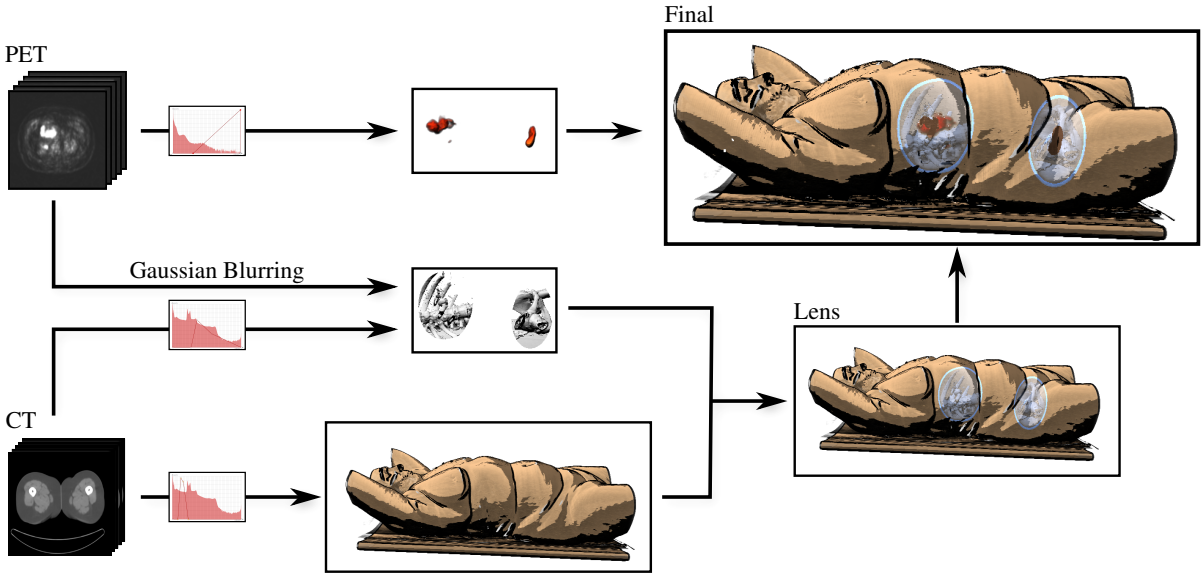

Figure 4: *Illustrative multivolume visualization pipeline: The PET data is used to produce a Phong shaded result. We use the blurred PET combined with the CT data to create a visualization of the surroundings of the PET and we illustrate the skin of the CT data. Afterwards, we combine them and create a lens effect. Finally, we add the PET and create the final image.*

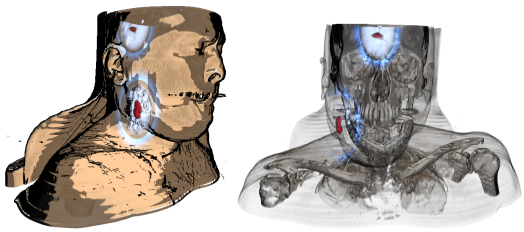

Figure 5: *Different transfer functions applied to the COMUNIX dataset. The first image shows the skin as the context. In the second illustration, we increased the transparency of the skin to reveal the skeletal structures.*

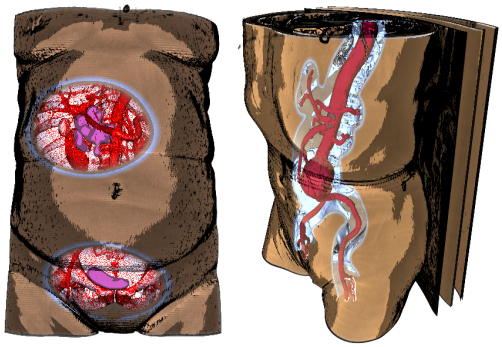

(a) (b)

Figure 6: *(a): The PETCETIX dataset with the arteries as well as the bone structures as a surrounding context object for the PET region. (b): the PANORAMIX CT dataset with the aorta as the focus object, the bones as immediate context and the skin as the context.*

to localize the pathology. This particular dataset shows that our technique can be applied to other cases, such as highlighting atlas information in CT scans.

In addition to clinical applications, our technique can also be used in a research setting. We have applied our technique to the PET/CT data in the publicly available digimouse project [SCS*02, DSCL07], a full digital mouse atlas consisting of labels, CT and PET data. This data was not acquired in a PET/CT scanner, but registered afterwards using markers. Another difference between this and the previously mentioned results is that the CT scan was acquired in a micro-CT scanner. The results of our method applied to this dataset can be seen in Figure 7. Normal metabolic PET activity is shown in several organs and the skin and bones from the CT are used as anatomical context.

7. Evaluation

To assess the clinical applicability of our approach, we conducted two informal evaluations with three physicians in the initial evaluation and eleven physicians in the second phase. The setup for the experiments was as follows:

Preparation: We provided the participants with a short introduction to PET/CT imaging data and the visual encoding

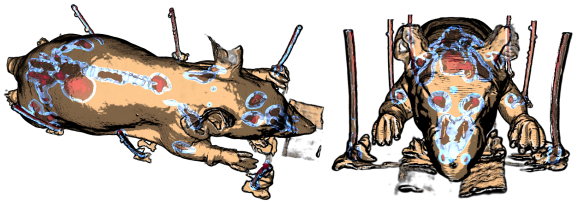

Figure 7: *Our technique applied to the digimouse set.*

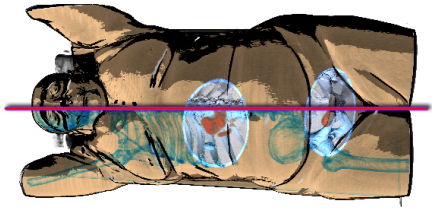

Figure 8: *Optional bone context illustration.*

chosen. First, we introduced the types of data used in the evaluation and what they represent. We showed a 2D slice figure showing the fused PET and CT data and the colormap used for the PET enhancement. Afterwards, we presented a 3D rendering example figure and demonstrated how the PET activity and CT information is shown in these. This helps them to understand and rate the usefulness and the usability of the different rendering techniques presented in the video.

Presentation: In order to gather as many participants as possible and because we are not evaluating interaction, we presented the results using a video. The video features three datasets represented in 2D slices as well as several 3D rendering methods: maximum intensity projection (MIP), standard raycasting with pre-defined transfer functions (Multi-Vol), and our approach (Illustra-Vis). The goal of this phase was to enable participants to assess the quality of our approach and how it compares to existing methods. All datasets were first displayed in the traditional 2D slice-based views, by presenting all the axial images from bottom to top and back. After this, they were shown using three 3D rendering techniques by rotating the volume around the z-axis. First, we decided to show a maximum intensity projection (MIP) as it is currently used in medical practice for analyzing scans. MIP visualizes only the structures in the combined dataset that have the highest intensity along the ray. The second technique we chose is standard multi-volume (Multi-Vol) rendering using a separate transfer function for both modalities. The Multi-Vol approach accumulates the values along the ray with a given transfer function (TF) for both volumes. To make the comparison fair, we used pre-defined transfer functions for the CT scan with a similar appearance to our result. We showed the participants the results of all techniques on all three datasets (see Figure 9).

Questions: To conclude the evaluation, we asked the participants to give us feedback on the presented techniques via a series of multiple choice and open questions. These questions allowed us to get their opinion on preferences regarding 2D slices versus 3D techniques as well as to compare the 3D techniques with each other. Additionally, we asked several questions regarding our technique specifically, to evaluate the benefits and limitations. Finally, we asked the users open questions on potential applications and suggestions.

7.1. First evaluation

For the first evaluation session, we asked three physicians to provide feedback on our illustrative visualization technique. All three women had different specializations in the field of medicine and various levels of experience. The first participant $P1$ (45 years) is a radiologist with 20 years of experience mainly focused on CT. The second participant $P2$ (29 years) has a surgical background with 8 years of experience. The third participant $P3$ (22 years) is a medical student with 4 years of experience. $P1$ stated that to get a good orientation in 3D, the bones are the most important part for her. Therefore, the Multi-Vol visualization was rated less useful, as other organs occluded it as well as the PET enhancement under different viewing angles. Furthermore, she stated that our technique is applicable for patient-doctor communication and education, but for an overview the lens is too small to give a good insight in the full skeletal context. Hence, she preferred the MIP technique as an additional technique to use besides the 2D slices she is accustomed to. $P2$ rated our technique as the best 3D techniques for an overview. She argued that it provides a nice visualization of the skin that facilitates the depth and shape perception. The lens helps to focus on the PET activity and avoids occlusion, which is an issue in Multi-Vol rendering. Furthermore, she stated that our technique could be applicable to surgical planning, if surrounding organs and vessels were made visible. She also saw potential in using the technique for interdisciplinary communication (in so called Tumorboards) and in education for courses on radiology and surgery. $P3$ argued similar to $P1$, she stated that the full skeletal information is important to localize the PET in terms of the whole body. Our technique can do that as well, but then the lens should be enlarged so far that the shape perception from the skin is reduced. Based on the feedback of this initial user study, we adjusted our technique, taking their feedback into account. As all physicians argued that full skeletal visualization is important for the orientation and localization, we decided to optionally include the bones in the context visualization. For this, we computed the contour of the bone structure as well as the bones themselves and added this with a blue color on the skin (see Figure 8). Here, the contours are more visualized more prominently than the complete bone structure. Drawing the full bones on the skin could be either too subtle or too distracting. Therefore, as an optional feature besides

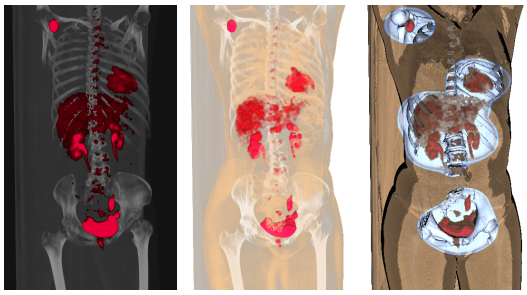

(a) MIP (b) Multi-Vol (c) Our

Figure 9: *Comparison of different volume rendering techniques: MIP, Multi-volume, and our technique.*

only illustrating the skin, we highlight the bones slightly and add the contours as the stronger indication for the bones.

7.2. Final evaluation

After we adjusted our visualization, we conducted a second evaluation session to gather more feedback from the three participants. *P*2 already liked our technique the most, but appreciated the changes made and argued that it adds a better context to the skin. *P*1 and *P*3 argued that our techniques improved significantly and that it combines the advantages of MIP and Multi-Vol. Additionally, we increased the number of users for the second evaluation in order to gather more feedback. We gathered a group of eleven participants (one male, ten female) ranging from ages 22 to 45 (average: 29, median: 28) with diverse backgrounds and years of experience ranging from 1 to 20 years (average: 6.59 years, median: 6 years). Using a questionnaire, we asked several specific questions where the participants had to answer what technique they would choose (2D slices, MIP, Multi-Vol, Illustra-Vis) in four scenarios:

- Showing the PET activity to a patient to explain diagnosis or treatment planning
- Discussing the results as a radiologist with another medical specialist
- Approximately localizing the PET in relation to the anatomy
- Visualizing the structural CT anatomical information

In all scenarios, Illustra-Vis was chosen most often, followed by MIP and Multi-Vol (see Figure 10). Finally, we asked several questions on the practical use of Illustra-Vis. We asked if they would use the Illustra-Vis technique in: diagnosis (Q1), treatment planning (Q2), patient-doctor communication (Q3), interdisciplinary communication (Q4). Most participants would use Illustra-Vis in all cases (see Figure 11). Additionally, we asked whether the skin illustration supports the form/shape impression, here three stated they strongly agree and the other seven agreed. Furthermore, regarding the depth/spatial impression, three stated they strongly agree that the skin illustration is helpful, while six agreed and two neither agreed or disagreed. In the open questions, one user commented that he would use our technique specifically for patient-doctor communication. He mentioned: "Patients with a lack of knowledge probably find it easier to identify themselves when using the Illustra-Vis image to explain diagnosis or treatment. It shows the contours of the body, making it easier to imagine where the area of interest is roughly located." For this purpose, he would use the illustration technique showing just the skin in the context, as the bones are less relevant to patients. He states that MIP reveals the PET activity more clearly, but our technique and specifically the skin illustration supports the depth and shape perception. A participant who chose MIP for patient-doctor communication argued that the other 3D techniques show too much information. Another participant discusses the main advantage of our method: "The main advantage is you have a complete image of the patient's body and you could do some prediction about what is normal or not from the image. For example: when you have an obese patient you could overestimate the shape or size of an organ and make a wrong diagnosis." Additionally, a participant argued: "Illustra-Vis focuses on relevant areas (PET) and is not occluded by tissues as in Multi-Vol. Furthermore, it gives a fast overview of the PET activity as well as a good localization as the surrounding bones are illustrated." Further comments were: "Illustra-Vis is more easily accessible for non-professionals." and "It is an illustrative way to show patients their diagnosis and possible treatment planning."

8. Discussion

In the evaluation, we compared our technique to the standard maximum intensity projection (MIP) approach and the multi-volume rendering (Multi-Vol) approach. Illustrating PET/CT with the MIP approach results in a dominant representation of the PET such that the PET region of high intensities are always drawn on top of the bones. MIP is a good technique to get a fast and simple overview of the data and many medical experts already have previous experience with MIP. Mostly medical experts use this technique to identify areas of interest and afterwards, they use the 2D slice view to analyze the affected regions. Nevertheless, this approach does not provide depth cues and therefore the slice view is needed in addition to accurately locate the suspicious activity correctly. Multi-Vol rendering, unlike MIP, can ensure that the correct drawing order of the structures is shown and has improved depth perception in this way. This technique however, can result in occluded regions such that the important PET activity is not visible anymore under all viewing angles. Our approach illustrates the surrounding regions of the PET activity without causing occlusion, while providing more shape cues than MIP. This can be especially useful when the PET activity of interest has a lower intensity than the surrounding structures. With our technique it is possible

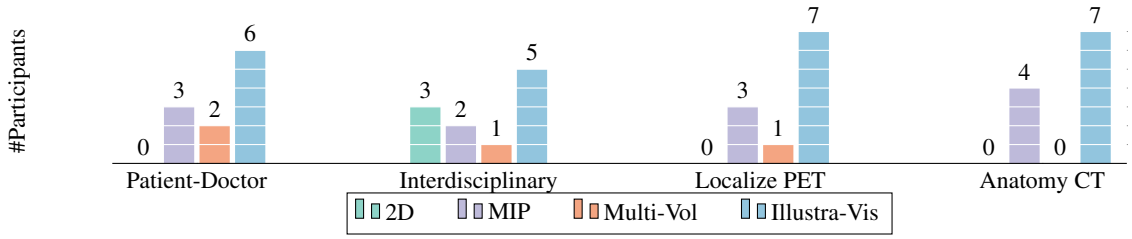

Figure 10: *Answers of the eleven participants regarding which of the four techniques they would prefer in four scenarios.*

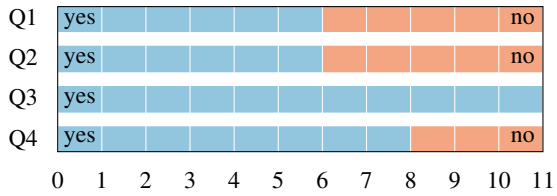

Figure 11: *Would you use the Illustra-Vis technique in: diagnosis (Q1), treatment planning (Q2), patient-doctor communication (Q3), interdisciplinary communication (Q4).*

to have a contextual overview by means of visualizing the skin and bones as well as identifying the PET foci in a single view. Using a different rendering technique around suspicious metabolic activity, allows the users to identify the regions more quickly and precisely. In case the PET activity is occluded by bone, we attenuate the color of the bones such that the PET is always perceivable. Theoretically, our approach combines the advantages of MIP (no occlusion of the PET activity) and Multi-Vol (improved depth perception) in a single comprehensive solution. During the evaluation, while most experts preferred our technique and noted the advantages, several medical experts still preferred the MIP. These experts state that they appreciate the simple view a MIP provides and are accustomed to interpreting this representation, while they are not used to our technique.

9. Conclusion

In this paper, we presented an illustrative visualization technique for PET/CT datasets. We used different illustrative techniques to represent the functional PET activity as a focus area combined with the anatomical CT information as a context object. First, we applied feature line techniques to the CT dataset and demonstrated that the results of our simple approach are comparable to other state-of-the-art feature line techniques. Afterwards, we created a halo around the PET activity by blurring as a simple way to create a region of interest in the dataset. We used this region as the basis for a lens metaphor that allows the user to see the inside of the body surrounding the PET activity. Here, we used different shading techniques to illustrate the surroundings of the PET to give a better indication of the spatial relationships to the anatomy and thus a better way to localize the suspicious PET activity in the data. Furthermore, we applied our illustrative multi-volume rendering to different datasets to show the flexibility of our technique. Besides freely available clinical data, we also applied our technique to a research dataset, acquired using a Micro-CT scanner. Additionally, we applied our technique to a single CT dataset, using contrast-enhanced vascular information as the focus object to demonstrate the flexibility of the method. In the future, we envision extending this technique to other multi-modal visualization applications. A qualitative user evaluation showed the utility of our approach for treatment planning, doctor-patient communication, interdisciplinary communication, and diagnosis (when combined with 2D slice information). The evaluation revealed that our approach gives a fast and clear summary and can be used as a basis for further analysis. In summary, we have developed an illustrative visualization technique designed with PET/CT data in mind, which is potentially applicable for general (multi-)volume focus and context applications.

Acknowledgements

Kai Lawonn was partially funded by the BMBF (STIMULATE-OVGU: 13GW0095A) and Noeska Smit was supported by the Technology Foundation STW.

References

[BBB*12] BRAMON R., BOADA I., BARDERA A., RODRIGUEZ J., FEIXAS M., PUIG J., SBERT M.: Multimodal Data Fusion Based on Mutual Information. *IEEE Trans. Vis. Comput. Graph. 18*, 9 (2012), 1574–1587. 2

[BGM*10] BRUCKNER S., GRÖLLER E., MUELLER K., PREIM B., SILVER D.: Illustrative focus+context approaches in interactive volume visualization. In *Scientific Visualization: Advanced Concepts*. 2010, ch. 10. 2

[BHW*07] BURNS M., HAIDACHER M., WEIN W., VIOLA I., GROELLER E.: Feature emphasis and contextual cutaways for multimodal medical visualization. *IEEE Trans. Vis. Comput. Graph.* (2007), 275–282. 2

[BHWB07] BEYER J., HADWIGER M., WOLFSBERGER S., BÜHLER K.: High-Quality Multimodal Volume Rendering for

Preoperative Planning of Neurosurgical Interventions. *IEEE Trans. Vis. Comput. Graph. 13*, 6 (2007), 1696–1703. 2

[BSP*93] BIER E. A., STONE M. C., PIER K., BUXTON W., DEROSE T. D.: Toolglass and magic lenses: The see-through interface. ACM SIGGRAPH, pp. 73–80. 5

[BTB*00] BEYER T., TOWNSEND D. W., BRUN T., KINAHAN P. E., CHARRON M., RODDY R., JERIN J., YOUNG J., BYARS L., NUTT R.: A combined PET/CT scanner for clinical oncology. *Journal of Nuclear Medicine: Official Publication, Society of Nuclear Medicine 41*, 8 (2000), 1369–1379. 3

[BVPtR08] BRECHEISEN R., VILANOVA A., PLATEL B., TER-HAAR ROMENIJ B.: Flexible GPU-based multi-volume ray-casting. In *Proc. of Vision, Modelling and Visualization* (2008), pp. 1–6. 2

[CS99] CAI W., SAKAS G.: Data Intermixing and Multi-volume Rendering. *Computer Graphics Forum 18*, 3 (1999), 359–368. 2

[DFRS03] DECARLO D., FINKELSTEIN A., RUSINKIEWICZ S., SANTELLA A.: Suggestive contours for conveying shape. *ACM SIGGRAPH* (2003), 848–855. 4

[DSCL07] DOGDAS B., STOUT D., CHATZIIOANNOU A. F., LEAHY R. M.: Digimouse: a 3d whole body mouse atlas from ct and cryosection data. *Physics in medicine and biology 52*, 3 (2007), 577. 6

[FH09] FUCHS R., HAUSER H.: Visualization of Multi-variate Scientific Data. *Computer Graphics Forum 28*, 6 (2009), 1670–1690. 2

[FPT04] FERRE M., PUIG A., TOST D.: A framework for fusion methods and rendering techniques of multimodal volume data. *Computer Animation and Virtual Worlds 15*, 2 (2004), 63–77. 2

[HBKG08] HAIDACHER M., BRUCKNER S., KANITSAR A., GR"OLLER M. E.: Information-based transfer functions for multimodal visualization. In *Proc. of Vis. Comp. for Biology and Medicine* (2008), pp. 101–108. 2

[HBKS05] HONG H., BAE J., KYE H., SHIN Y.: Efficient Multimodality Volume Fusion Using Graphics Hardware. In *Computational Science*, vol. 3516 of *Lecture Notes in Computer Science*. 2005, pp. 842–845. 2

[Hod03] HODGES E. R.: *The guild handbook of scientific illustration*. John Wiley & Sons, 2003. 4

[IFP95] INTERRANTE V., FUCHS H., PIZER S.: Enhancing transparent skin surfaces with ridge and valley lines. In *Proc. of IEEE Visualization* (1995), pp. 52–59. 4

[JBB*08] JAINEK W., BORN S., BARTZ D., , STRASSER W., FISCHER J.: Illustrative Hybrid Visualization and Exploration of Anatomical and Functional Brain Data. In *Computer Graphics Forum* (September 2008), vol. 27, pp. 855–862. 2

[JDA07] JUDD T., DURAND F., ADELSON E.: Apparent ridges for line drawing. In *ACM* (2007), p. 19. 4

[JKE*13] JUNG Y., KIM J., EBERL S., FULHAM M., FENG D. D.: Visibility-driven PET-CT Visualisation with Region of Interest (ROI) Segmentation. *Vis. Comput. 29*, 6-8 (2013), 805–815. 2

[KEF07] KIM J., EBERL S., FENG D. D.: Visualizing dual-modality rendered volumes using a dual-lookup table transfer function. *Computing in Science and Engineering 9*, 1 (2007), 20–25. 2

[KHGR02] KNISS J., HANSEN C., GRENIER M., ROBINSON T.: Volume rendering multivariate data to visualize meteorological simulations: A case study. In *Proceedings of the Visualization Symposium* (2002), Eurographics Association, pp. 189–194. 2

[KMV*00] KLUETZ P. G., MELTZER C. C., VILLEMAGNE V. L., KINAHAN P. E., CHANDER S., MARTINELLI M. A., TOWNSEND D. W.: Combined PET/CT Imaging in Oncology: Impact on Patient Management. *Clinical Positron Imaging 3*, 6 (2000), 223 – 230. 3

[KSW*04] KNISS J., SCHULZE J. P., WÖSSNER U., WINKLER P., LANG U., HANSEN C. D.: Medical applications of multi-field volume rendering and vr techniques. In *Proc. of the Visualization Symposium* (2004), no. 350, Eurographics Association, pp. 249–254. 2

[KSW06] KRUGER J., SCHNEIDER J., WESTERMANN R.: ClearView: An Interactive Context Preserving Hotspot Visualization Technique. *IEEE Trans. Vis. Comput. Graph. 12*, 5 (2006), 941–948. 2

[KWTM03] KINDLMANN G., WHITAKER R., TASDIZEN T., MOLLER T.: Curvature-based transfer functions for direct volume rendering: Methods and applications. In *IEEE TVCG* (2003), pp. 67–. 3, 4

[Law15] LAWONN K.: *Illustrative Visualization of Medical Data Sets*. Thesis, Jan. 2015. 2

[LCMY07] LEROTIC M., CHUNG A., MYLONAS G., YANG G.-Z.: PQ-space based non-photorealistic rendering for augmented reality. In *MICCAI*, vol. 4792 of *Lecture Notes in Computer Science*. 2007, pp. 102–109. 5

[NEO*10] NGUYEN K. T., EKLUND A., OHLSSON H., HERNELL F., LJUNG P., FORSELL C., ANDERSSON M. T., KNUTSSON H., YNNERMAN A.: Concurrent Volume Visualization of Real-Time fMRI. In *Proc. of IEEE/EG Volume Graphics* (2010), pp. 53–60. 2

[osi] OSIRIX, url = http://www.osirix-viewer.com/datasets/, note = (last accessed: 09/06/2015). 4, 5

[RTF*06] RÖSSLER F., TEJADA E., FANGMEIER T., ERTL T., KNAUFF M.: Gpu-based multi-volume rendering for the visualization of functional brain images. In *Proc. of SIMVIS* (2006), pp. 305–318. 2

[SCS*02] STOUT D., CHOW P., SILVERMAN R., LEAHY R. M., LEWIS X., GAMBHIR S., CHATZIIOANNOU A.: Creating a whole body digital mouse atlas with pet, ct and cryosection images. *Mol. Imaging Biol 4*, 4 (2002), S27. 6

[SZV01] STOKKING R., ZUIDERVELD K., VIERGEVER M.: Integrated volume visualization of functional image data and anatomical surfaces using normal fusion. *Human Brain Mapping 12* (2001), 203–218. 2

[TCYH04] TOWNSEND D., CARNEY J., YAP J., HALL N.: PET/CT today and tomorrow. *The Journal of Nuclear Medicine 45*, 1 (Jan 2004), 4S – 14S. 2

[TGK*14] TOMINSKI C., GLADISCH S., KISTER U., DACHSELT R., SCHUMANN H.: A survey on interactive lenses in visualization. In *EuroVis State-of-the-Art Reports* (6 2014). 5

[VKG04] VIOLA I., KANITSAR A., GRÖLLER M. E.: Importance-Driven Volume Rendering. In *IEEE Trans. Vis. Comput. Graph.* (2004), pp. 139–145. 2

[WLM02] WILSON B., LUM E. B., MA K.-L.: Interactive Multi-volume Visualization. In *Computational Science*, vol. 2330. 2002, pp. 102–110. 2

[XHT*07] XIE X., HE Y., TIAN F., SEAH H.-S., GU X., QIN H.: An effective illustrative visualization framework based on photic extremum lines (pels). *IEEE Trans. Vis. Comput. Graph. 13* (2007), 1328–1335. 4