

Toward a Multimodal Diagnostic Exploratory Visualization of Focal Cortical Dysplasia

Wu, Shin-Ting^{*‡}, Raphael Voltoline^{*‡}, Wallace S. Loos^{*‡}, J.A. Iván Rubianes^{*‡}, Lionis S. Watanabe^{*‡§},
Bárbara J. Amorim[§], A. Carolina Coan^{†‡}, Fernando Cendes^{†‡} and Clarissa L. Yasuda^{†‡}

Email: ting@dca.fee.unicamp.br

^{*}School of Electrical and Computer Engineering [†]School of Medical Sciences [‡]BRAINN Research, Innovation and Dissemination Center of the University of Campinas [§]Clinics Hospital
University of Campinas, São Paulo, Brazil

Abstract

Focal cortical dysplasia (FCD) is a malformation of cortical development and a common cause of pharmacoresistant epilepsy. Resective surgery of clear-cut lesions may be curative. However, the localization of the seizure focus and the evaluation of its spatial extent can be challenging in many situations. For concordance assessment, medical studies show the relevance of accurate correlation of multi-source imaging sequences—to improve the sensitivity and specificity of the evaluation. In this paper, we share the process we went through to reach our simple, but effective, solution for integrating multi-volume rendering into an exploratory visualization environment for the diagnosis of FCD. We focus on fetching of multiple data assigned to a sample when they are rendered. Knowing that the major diagnostic role of multiple volumes is to complement information, we demonstrate that appropriate geometric transformations in the texture space are sufficient for accomplishing this task. This allows us to fully implement our proposal in the OpenGL rendering pipeline and to easily integrate it into the existing visual diagnostic application. Both time performance and the visual quality of our proposal were evaluated with a set of clinical data volumes for assessing the potential practical impact of our solution in routine diagnostic use.

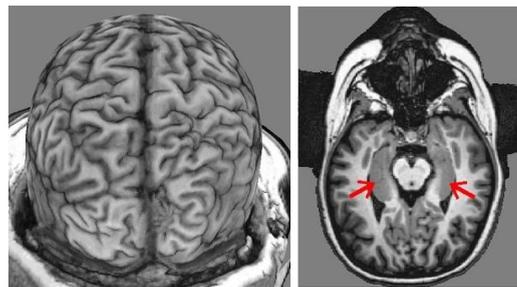
Keywords—J.3 Life and Medical Sciences. D.2.6.d Interactive Environments. I.6.9.h Volume visualization.

I. INTRODUCTION

Focal cortical dysplasia (FCD) is a highly epileptogenic cortical lesion and a major cause of pharmacoresistant epilepsy. It is also the most common histopathological diagnosis in children undergoing surgery for epilepsy. Approximately 50–70% of patients with FCD are seizure-free after removal of well-defined lesions [1]. An accurate delineation of a structural lesion is, therefore, critical for the prognosis of surgical resective treatment.

Although electroencephalogram abnormalities can suggest the presence of a structural abnormality, the location of subtle cortical lesions cannot be accurately inferred from the scalp electrode positions. Advances in modern neuroimaging have considerably increased accuracy in finding suspected epileptogenic lesions. Since cortical gray-white matter blurring is frequently an evidence of FCD in magnetic resonance imaging, a way to enhance

detection and diagnosis of focal lesions is to improve the visualization of those junctions [2]. Given the importance of lesion detection to improve surgical outcomes, the technological team of the School of Electrical and Computer Engineering and the medical team of the School of Medical Sciences at the University of Campinas have jointly engaged (since 2008) in prospecting highly interactive environment that facilitates such a visualization.



(a) CR: cerebral convolutions (b) MPR: hippocampus

Fig. 1. Unique views of brain structures.

Differential interaction requirements in imaging studies after subtle lesions infuse new challenges that neither open-source software applications nor proprietary ones have fully addressed. This led us to implement from scratch a GPU-based interaction architecture proposed by Batagelo and Wu [3] on top of the GPU-based ray-casting architecture developed by Stegmaier et al. [4]. The wxWidgets GUI library was used for creating an interactive interface, and the open source Grassroots DiCoM library for parsing Dicom medical files. We adopted a user-centered design approach and successfully integrated into our proposed architecture a curvilinear (CR) [5] and multiplanar reformatting (MPR) to yield a unique view of complex cerebral cortical convolution patterns (Figure 1(a)) and hippocampus (Figure 1(b)), respectively. Regular meetings and studies of existing medical software applications were essential for this phase of development. Very often, mock interfaces were built for assessing its potential clinical values.

In Section II we show that a single imaging modality can only reveal a limited number of tissue features. Multimodality imaging has rapidly evolved and making viable a concordance-based diagnosis which relies on the degree of agreement among the findings from different studies [6]. The sensitivity (probability of sick people to be diagnosed) and the specificity (probability of healthy people not to be considered as sick) of imaging examination have significantly improved. Multimodality study—became a reality with the design of multimodal medical scanners, along with the development of co-registration algorithms for scans from different sources and the—deployment of frameworks that support multimodal visualization. The latter software environments have been devised with sophisticated interaction tools to assist physicians in improving their findings [7], [8]. Unfortunately, none of existing medical image processing algorithms is reliable enough for automatic detection—of FCD. Major image interpretations still require human expertise.

In Section III we shortly describe the multimodality studies commonly adopted at our university hospital, and the limitations of visualization system physicians should transcend. By observing their routine procedure, we formulated a hypothesis that integrating a multimodal rendering into our monomodal diagnostic environment would meet their expectation.

A review of the existing approaches to multimodal rendering in Section IV indicates a series of ray-casting-based solutions. However, the purpose of multimodal rendering in our specific FCD diagnostic application is to simply integrate complementary tissue features in the existing monomodal OpenGL-based environment. It requires us to look for yet another simpler sustainable rendering solution. From a visualization perspective, the primary contribution of this paper is to handle a multi-volume rendering as a variant of a single-volume one, such that slight changes in the fragment shader suffice. Nevertheless, from a medical perspective, this slight modification can substantially increase the diagnostic value of our previous environment.

Our solution is based on distinguishing five reference systems presented in Section V: DICOM raw data coordinate reference (DCR), physical patient-based coordinate reference (PhCR), normalized physical coordinate reference (NPhCR), texture coordinate reference (TCR), and viewport coordinate reference (VCR). We demonstrate that with these references we replace pre-classification strategies with merely converting the co-registration transformation matrices from the patient space to the texture space.

In Section VI we present an implementation of our proposal in the context of a prototype of bimodal interactive visualization system in which multiplanar and curvilinear reformatting algorithms are integrated. Both time performance and clinical utility of our proposal are assessed with a set of clinical data volumes. The potential and the

limits of our solution, as well as works in progress are discussed in Section VII.

II. HOW FCD IS CURRENTLY VISUALLY DETECTED?

In this section we briefly describe the practical role of diverse imaging modalities and the visual inspection physicians conduct for FCD detection at our university hospital.

The abnormalities associated with FCD can be identified in three different, but usually simultaneously acquired, magnetic resonance imaging (MRI) sequences: T1-weighted, T2-weighted and fluid-attenuated inversion recovery (FLAIR) sequences. Cortex thickening, cortical and subcortical signal changes, blurring of the gray-white matter interface, the “transmantle sign” (an imaging of high intensity signal on T2-weighted images extending from the ventricles to the cortex), and abnormalities in sulcus-gyral pattern are well-known imaging evidences [1]. Figures 2(a), 2(b) and 2(c) illustrate, respectively, diffuse thickened gray signal in a T1-weighted MRI scan, hypersignal in light gray in the both T2-weighted and FLAIR sequences of lesioned areas. However, other information, such as clinical, physiological, functional and metabolic data, is also necessary for confirming the FCD diagnosis. Two other imaging modalities allow for assessing cerebral functional and metabolic activities: single-photon emission computed tomography (SPECT) and 18F-fluorodeoxyglucose-positron emission tomography (FDG-PET).

SPECT is a nuclear medicine tomographic imaging technique that permits the quantification of changes of blood flow (perfusion) in the capillaries of the scanned cerebral regions. Studies have demonstrated that there is a close association between seizure states and blood flow changes. During a seizure (ictal state) the perfusion increases (hyperperfusion) in lesioned regions, while a decrease (hypoperfusion) is usually observed between the seizures (interictal state). These two states are evaluated separately with interictal and ictal studies. Ictal SPECT is performed by injecting radiotracer intravenously during seizures. When the tracer injection occurs at the very beginning of the seizure, the exam is highly sensitive for the detection of the epileptogenic foci [8]. Figure 2(f) illustrates the signals of an ictal SPECT image colored with the rainbow color palette where red indicates hyperperfusion (that correlates with higher cell activity) and blue shows decreased perfusion. Sometimes, the hyperperfusion in an ictal SPECT is not so evident. In these cases, the subtraction of ictal from interictal SPECT co-registered to MRI, widely known as SISCOM, has been proved to be more sensitive than the visual individual analysis [9].

FDG-PET imaging modality analyzes the concentration of a glucose analogue (the 18F-FDG), to assess functional and metabolic activities in the body. Studies have shown that, in the interictal state, the epileptogenic regions present lower glucose uptake, indicating lower metabolic activity. Figure 2(d) illustrates the signals of an FDG-PET study

colored with the rainbow color palette. Although the enhanced areas can be more extensive than the ones revealed by structural and physiological studies, FDG-PET is useful in localization of suspected epileptogenic foci when MRI scans are negative [9].

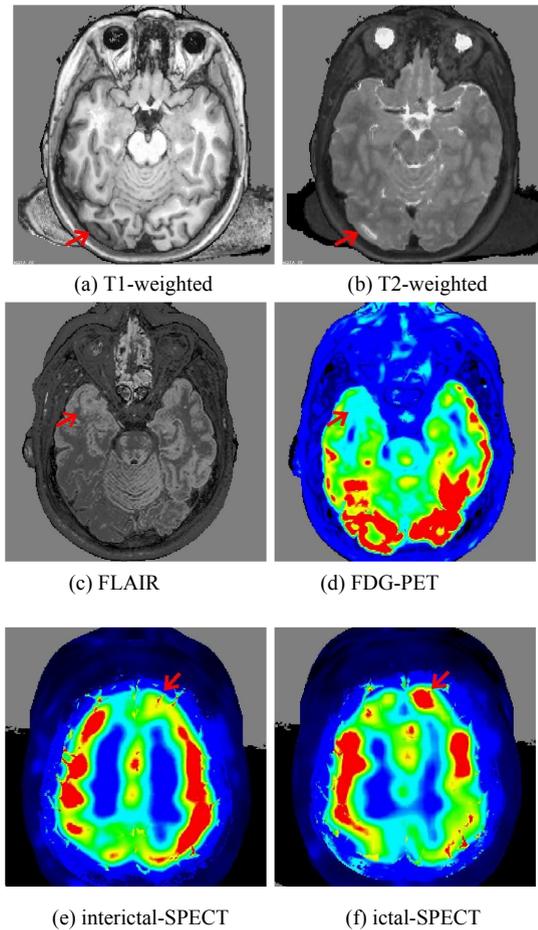


Fig. 2. Relevant imaging features for the diagnosis of focal cortical dysplasia in different patients highlighted with a red arrow: (a) blurring of the gray-white matter, (b,c) signal of higher intensity, (d) low metabolic activity, (e) low perfusion, and (f) hyperperfusion.

These imaging studies are usually inspected in slice-by-slice fashion. Often, the experts do not access the patient's clinical data in order to avoid biased first impression. With the use of a visualization system, they scroll through the slices of a reconstructed volume and adjust the image contrast for better visualization of abnormalities. Three mainly inspected orthogonal direction angles are: axial (head \leftrightarrow feet), sagittal (left \leftrightarrow right), and coronal (back \leftrightarrow front) views. The axial and the coronal views help them with comparison studies between the left and the right brain hemispheres. Sometimes specialists look at MRI sequences from different angles by reslicing the reconstructed volume with an arbitrary user-defined cutting plane. The multiplanar reformatting (MPR) tool delivered together with the MRI scanners and the curvilinear reformatting (CR) tool, such as the one implemented in the BrainSight[®] from the Rogue Research Inc. , are commonly

used. MPR allows better visualization of subtle thickening of the cortex and gray-white blurring that does not lie orthogonal to the direction of the scan, and CR improves the detection of subtle lesions in the depth of complex curvilinear gyral structure. Co-registration of anatomical (MRI) and functional (FDG-PET and SPECT) studies are commonly conducted by nuclear medicine physicians using different types of application software, such as Scenium, a Siemens software program that is part of the Syngo.via Neurology package. Lately, specialists create their mental 3D-models of the region of interest from 2D images and discuss their findings in multidisciplinary meetings to decide on surgical resection according to the degree of agreement of clinical and imaging findings.

III. WHAT IS STILL MISSED?

The major complaint from physicians about systems they used in their routine visual inspection is that such systems do not offer simultaneous visualization of multimodal acquisitions. Apparently,—knowing that an interactive rendering of multiple co-registered neuroimaging data into a single 3D image is already supported by the proprietary software Amira [7] and that an interface with linking multimodal views can be found in the open-source ITK-SNAP [8], it seems that there are no grounds for complaint.

However, careful observation of the major issues discussed in weekly epilepsy meetings brought to light the problem: although physicians can explore each imaging study thoroughly with available tools, they can neither reformat a number of imaging modalities simultaneously nor easily overlay the reformatted slices for concordance assessment and spatial evaluation. This medical desire poses a new challenge. It requires not only appropriate fusion of a sample's data values, but also novel interaction tools for analyzing fused data.

IV. IS THERE AN OFF-THE-SHELF SOLUTION?

Having clarified the goal that we should reach, we undertook research into an existing alternative that reuses interaction tools we developed for a monomodal application [5],[10] in a multimodal environment. First, spatial and timely distributed medical volumes of rectilinear regular grids defined in the PhCR reference system should be in a common reference system. Then, the co-registered multimodal data should be visualized either in the form of multiple coordinated views (Figures 3(a)–3(c)) or in a single multimodal view (Figure 3(d)).

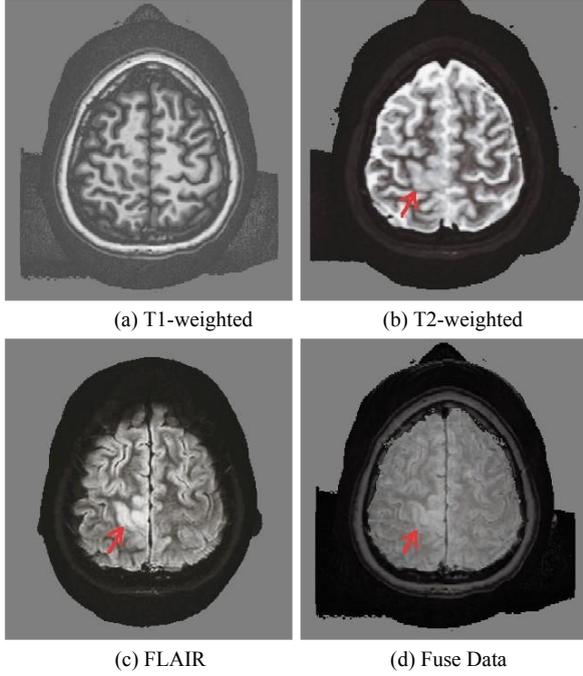


Fig. 3. Coordinated axial views of (a) T1-weighted, (b) T2-weighted, (c) FLAIR scans, and (d) a fused view from three modalities of a patient with FCD in the pre-central gyrus.

In multiple coordinated views, volumes from different modalities of a subject are displayed in multiple windows, usually one modality per window, and the operations on the views are coordinated. Hence, the multi-volume rendering may be reduced to the ray-casting of single volumes. Ikits et al. show in [11] that the ray-casting algorithm for a single volume can be divided into three stages: (1) Initialize, (2) Update, and (3) Draw. In the Initialize stage, a scanned volume in DCR coordinates is loaded, processed and downloaded to the texture memory on the GPU, where it is represented as a 3D texture in TCR coordinates. The transfer functions that map the signal intensities to the optical properties are also downloaded to the GPU as 1D textures. Whenever a view-aligned slicing plane is changed, the bounding volume of the clipped texture, also known as proxy geometry, is computed in the PhCR coordinates. The corresponding 3D texels in TCR coordinates should be assigned to the new vertices of the proxy geometry. Finally, after this custom setup, the rendering is fired in the Draw stage for effectively generating images in the VCR reference. Even if they are well linked, views of separately rendered single volumes still require a mental combination of multimodal attributes displayed in different windows. Very subtle lesions can be easily overlooked.

For concordance assessment, it seems to be more appropriate a single multimodal view that permits visualization of multimodal attributes in a single image. Leu and Chen presented in [12] an implementation of multi-volume raycasting on CPUs. Provided the scanned floating volume datasets $V_i \in \{1, n\}$ and the respective transformation matrices T_i to the scanned reference volume

V_0 in the PhCR reference system, each ray is cast into V_0 for resampling it at discrete positions $Pt = (x, y, z)$ along the ray. The contribution of each sample to the pixel's color and opacity is fetched from a transfer function and accumulated. If the multi-volumes are axis-aligned and have the same resolution, we can simply fetch the data values in each floating volume with the same position Pt (T_i^{-1} is an identity matrix in line 8) and their optical attributes (line 11) are accumulated partially to the current pixel's color (line 12). This gives the perception that the scanned volumes are translucent and are laid one over another, as illustrates Figure 3(d). However, if the reference and the floating volumes are not grid aligned, a non-identity geometric transformation T_i^{-1} is required to get correct correspondence lPt in V_i (line 8). In addition, if the correspondence lies between grid points of V_i , the trilinear interpolation is commonly carried out (line 10).

Algorithm 1 Raycasting multi-volumes

Input: Reference volume V_0 and floating volumes V_i

Output: Multi-modal image

- 1: Determine entry (E) and exit (O) position of V_0 .
- 2: Compute ray direction \vec{d} .
- 3: $\Delta t \leftarrow \frac{\text{distance}(O, E)}{(\text{number of samples} - 1)}$
- 4: $t \leftarrow 0$; $fColor \leftarrow (0, 0, 0, 0)$.
- 5: **while** $t \leq \text{distance}(O, E)$ **do**
- 6: $Pt \leftarrow E + t \cdot \vec{d}$.
- 7: **for** $i = 1 \rightarrow n$ **do**
- 8: $lPt \leftarrow T_i^{-1} \cdot Pt$.
- 9: **if** $lPt \in V_i$ **then**
- 10: Get the data value of V_i at lPt .
- 11: Get (R, G, B, α) attributes.
- 12: Compose and accumulate in $fColor$.
- 13: **end if**
- 14: **end for**
- 15: $t \leftarrow t + \Delta t$.
- 16: **end while**

Because the brain's shape vary very slowly with age, a rigid co-registration algorithm, such as the interactive algorithm described in [13], is sufficient for estimating T_i between V_0 and V_i . In [13], after background noise removal, an initial guess of alignment $T_{i,0}$ is made from the direction cosines and a user-defined correspondence pair of two volumes. Iteratively, the mutual information between the two approximately aligned volumes, V_0 and ($T_{i,j-1}^{-1} V_i$), is calculated from their joint histogram. In sequence, a new geometric transformation $T_{i,j}$ that improves the mutual information is re-estimated with the optimization algorithm NEWUOA. The process is repeated successively until convergence to T_i . The authors also suggested to apply the partial volume interpolation and a multi-resolution strategy for improving the convergence robustness and the optimization efficiency, respectively.

Though quite simple on CPUs, the implementation of Algorithm 1 on GPUs is a challenging issue. It is because that the per-pixel ray processing can only occur either in the texture (TCR) or in the viewport (VCR) space. Several works have been devoted to this issue. Pre-classification of the overlapped volumes along each ray seems to be the preferred paradigm. There are essentially two proposals for this pre-classification: (1) before drawing the volumes with assigned multi-textures (in TCR), the overlapped volumes are divided into disjoint volumes in the PhCR/NPhCR space [14], and (2) before ray shooting, the depth layers of overlapped volumes are ordered from the viewing direction in the VCR space and stored on the GPU [15]. The major drawback of the first approach is the intensive CPU–GPU data transfer, while the scalability of the second approach is very low due to its $O(n^2)$ complexity (n rendering passes of n objects). In this paper, we show that, for diagnostic purpose, a single-pass rendering with just one draw call is sufficient for rendering multimodal scalar properties of the brain tissue.

As in [15], Kainz et al. [16] also did research on ray-casting a large number of volumes intermixed with complex translucent and concave polyhedral objects. For avoiding the multi-pass depth pre-sorting and for gaining access to the GPU shared memory, they fully implemented a single-pass depth-sort-based multi-volume ray-casting algorithm within CUDA. The major problem of this approach is that we should tailor previously developed interaction tools to this new rendering architecture. We show in this work that for specific medical application areas we can give up the generality and develop a clinically effective solution reusing the existing tools.

V. A NOVEL APPROACH

The main issue for implementing Algorithm 1 on GPU is to tailor it to its rendering pipeline in which the scanned volumes are downloaded as 3D textures with their own TCR spaces. Our finding that the data available on CPU and on GPU differ only on the spaces they are represented, led us to reduce multimodal rendering problem to a simple space transformation problem. We present in this section a novel approach—for fetching corresponding data and for simple volume clipping in TCR spaces.

A. Reference Systems

Essentially our proposal relies on the fact that there is a well-defined invertible affine transformation between the patient space PhCR, in which the co-registration procedure is usually estimated, and the texture space TCR, in which the scanned data are fetched for rendering. In order to attain this transformation, we distinguish five reference systems (Figure 4):

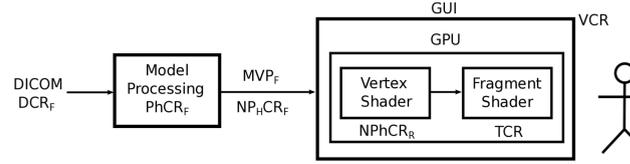


Fig. 4. Data flow for multimodal visualization: the raw data in the DCR space is loaded, transformed to PhCR space for CPU processing, and then transformed to NPhCR (vertex shader) and to TCR (fragment shader) for GPU processing. The MVP viewing transformation matrix is sent to GPU for specifying the desired angle of view on the VCR output window.

- DCR is the reference of the raw data. In the DICOM protocol, they are organized as a regular rectangular array with $Rows \times Columns \times Slices$ voxels addressable by integer coordinates (i, j, k) .
- PhCR is the scanned sample reference where the co-registration is usually performed. We consider that their coordinates (ph_x, ph_y, ph_z) differ from the DCR coordinates by the voxel spacing data (s_x, s_y, s_z) available in the DICOM-compliant file.
- NPhCR is the proxy geometry reference. The coordinates (nph_x, nph_y, nph_z) are, indeed, the normalized coordinates of (ph_x, ph_y, ph_z) with respect to $mS = \max(s_x \cdot Rows, s_y \cdot Columns, s_z \cdot Slices)$.
- TCR is the texture reference of scanned volumes on GPU. Their coordinates (t_x, t_y, t_z) are limited to the range $[0.0, 1.]$ and keep with the PhCR coordinates the relation

$$t_x = \frac{ph_x}{s_x \cdot Rows} \quad t_y = \frac{ph_y}{s_y \cdot Columns} \quad t_z = \frac{ph_z}{s_z \cdot Slices} \quad (1)$$
- VCR is the physical reference of the display device. Each sample in NPhCR coordinates (nph_x, nph_y, nph_z) are transformed to the VCR by premultiplying the matrix $MVP = (Viewport * P * V * M)$, where $Viewport$, P , V and M are the viewport, the projection, the view and the model transformation matrices, respectively.

B. Correspondences

In Algorithm 1 we consider that all volumes are represented in the PhCR reference systems and the transformation matrix T_i^{-1} (line 8 of Algorithm 1) brings a point in the PhCR of a floating volume to the PhCR of the reference one. To get correspondences in a fragment shader where the data are represented in the TCR space, we need a TCR-compliant co-registration matrix $T_{i(TCR)}$. In this section we present a simple way to get $T_{i(TCR)}$ from T_i .

Let $P_{0(PhCR)}$ and $P_{i(PhCR)}$ be a point in the reference volume V_0 and a point in the floating volume V_i , respectively. If they correspond to the same anatomical point, they satisfy

$$P_{i(PhCR)} = T_i^{-1} P_{0(PhCR)}. \quad (2)$$

From Section V-A these points given in the space $PhCR$ can be transformed to the space TCR with Eq. 1:

$$T_{i(PhCR \rightarrow TCR)}^{-1} P_{i(TCR)} = T_i^{-1} T_{0(PhCR \rightarrow TCR)}^{-1} P_{0(TCR)},$$

where $T_{ind(X \rightarrow Y)}$ denotes the transformation T from the space $PhCR$ to the space TC for a volume with the index ind . Isolating the term $P_{i(TCR)}$, we obtain

$$P_{i(TCR)} = T_{i(PhCR \rightarrow TCR)} T_i^{-1} T_{0(PhCR \rightarrow TCR)}^{-1} P_{0(TCR)}. \quad (3)$$

It follows immediately that the co-registration matrix $T_{i(TCR)}$ of a point $P_{0(TCR)}$ in the reference volume with the point $P_{i(TCR)}$ in a floating volume is

$$T_{i(TCR)} = T_{i(PhCR \rightarrow TCR)} T_i^{-1} T_{0(PhCR \rightarrow TCR)}^{-1}. \quad (4)$$

If the transformed point does not fall on the center of the texel of the volume i , we explore the interpolation mechanism provided by the texture mapping unit (TMU) to get missed values.

C. Clipping Planes

Clipping volumes are widely applied in uncovering hidden anatomical structures. Existing practice is either to build a selection volume or to compute the intersection of the clipping plane with respect to the scanned volume on the CPU and to draw this intersection with a single GPU drawing call. However, both procedures are not appropriate for exploratory investigations due to intensive GPU-CPU transfers. With the use of our proposed linear transformation, we can compute visible samples with respect to an implicitly represented clipping geometry directly in a fragment shader.

For example, if we pass to the GPU the clipping plane in TCR coordinates, only an inequality computation suffices for discarding a sample (t_x, t_y, t_z) of the reference volume in texture coordinates:

$$n_x(TCR)t_x + n_y(TCR)t_y + n_z(TCR)t_z - (n_x(TCR)P_x(TCR) + n_y(TCR)P_y(TCR) + n_z(TCR)P_z(TCR)) > 0,$$

where $(n_x(TCR), n_y(TCR), n_z(TCR))$ and $(P_x(TCR), P_y(TCR), P_z(TCR))$ are, respectively, the normal vector and a point on the clipping plane. This allows us to remove on-the-fly samples of the reference volume on the basis of their position in the clipping volumes before accumulating complementary multimodal data in the fragment shader.

VI. DOES OUR SOLUTION MEET REQUIREMENTS?

An evaluation of the multi-data processing strategy given in Section V was performed (1) to evaluate upgrade effort, (2) to assess the appropriateness of our proposal in enhancing the clinical value of a monomodal visualization system, and (3) to demonstrate the suitability of our idea for interactions. MRI volumes were acquired by an MR 3T Philips Intera-Achieva Scanner and the CT/PET volumes in the Siemens Biograph mCT40 TrueV at the hospital of the University of Campinas. Note that all patients enrolled in the present study were informed and signed the consent form approved by the Ethics Committee of the University of Campinas.

A. Software Application Upgrade

The best way to demonstrate the effectiveness of our proposal in upgrading an existing single to a multi-volume one is to upgrade our monomodal application. This should be carried out with minimal software architectural change and with maximal reuse of existing interactive tools.

Considering the volume processed in a monomodal system as a reference volume and other additional volumes as floating ones, our development effort is reduced to: (1) the change of the fragment shader on the GPU, (2) the configuration of texture mapping, from single to multi-texturing, and (3) the redesign of user interface on the CPU. All previously implemented interaction algorithms for a single volume can be straightly applied in the reference volume. The fragment shader is simply an implementation of Algorithm 1 with the transformation given in Eq. 4, while the setup of texture mapping consists essentially in creating more 3D-texture objects. Furthermore, we implemented the clipping procedure presented in Section V-C as a fragment shader and replaced the volume-plane intersection on the CPU with it.



Fig. 5. Bi-modal visualization interface with blending slider and modal tabs highlighted.

The major challenge was undoubtedly the redesign of a user interface. Simply fusing the data as in Figure 3(d) may be perceptually worse than visualizing the corresponding linked data as shown in Figures 3(a)–3(c). From weekly epilepsy meetings we learned that the medical users only expect to have additional support in quickly flipping intermixed co-registered imaging studies such as a flipbook for assessing the degree of agreement of individual findings.

After a series of experiments with prospective users, we decided in favour of a slider that allows them to manually configure the degree of blending of fused data and to scroll through images for building a mental 3D fused model. A new interface has been developed on top of the Qt-GUI cross-platform software due to its portability to mobile devices. Figure 5 illustrates the redesigned interface for a bimodal imaging study. Each study is associated with one

tab. In this figure, the modality T1-weighted MR (in the tab “MR-T1”) and the modality FDG-PET (in the tab “PET”) are loaded. The third tab “Reg”, visible in the figure, contains windows displaying the fused data of loaded volumes after co-registration. There is a slider, highlighted in the left corner of the figure, for configuring the degree of blending of two volumes. We can easily switch among the imaging studies by simply clicking on the corresponding tab. This prototype is available online at [17].

B. Clinical utility

With the purpose of evaluating the clinical value of our proposal, we show in this section three difficult histologically proven FCD cases due to either lesion subtlety or cortical convolution complexity.

In a routine investigation, the typical MRI epilepsy protocol at our university hospital includes T1-weighted, T2-weighted and FLAIR MRI sequences. Many times, these three studies are sufficient for revealing a dysplastic lesion as demonstrates the first selected clinical case (Figure 3). In this case, the three imaging volumes were downloaded to the GPU memory as 3D textures and rendered separately. They are presented in Figures 3(a), 3(b) and 3(c). Then, fixing T1-weighted MR sequence as reference volume, we applied Algorithm 1 to mix three volumes into a single one in the TCR space, as shown in Figure 3(d). Scrolling these coordinated volumes and varying blending factors in the fused one, we found brighter signals pointed by a red arrow on T2-weighted and FLAIR MR images. This suggests the presence of a lesion. Although T1-weighted MR study is negative, our finding agrees with the patient’s medical history and electroencephalogram abnormalities. The medical team decided for surgery and, after surgical removal of that lesion, guided by intraoperative electrocorticography, histologic examination proved that it is FCD.

Retrospectively, we applied multiplanar reformatting with different clipping planes defined in Section V-C to investigate abnormalities in volumes shown in Figure 6. In this second clinical case, the initial MRI evaluation was negative in the T1-weighted MRI sequence (Figure 6(a)) and a very subtle signal change (red arrow) in the floating T2-weighted (Figure 6(b)) and FLAIR MRI sequences (Figure 6(c)). Hence, the functional SPECT and FDG-PET studies were ordered. Hypoperfusion (Figure 2(e)) (red arrow) and hypometabolism (Figure 6(d)) (blue arrow) in the left superior frontal gyrus were observed. The superimposition of FDG-PET on MR images (Figure 6(e)) makes us to perceive the overlooked abnormal cortical thickness and subtle hyperintense FLAIR signal that was confirmed to be a FCD lesion after surgery.

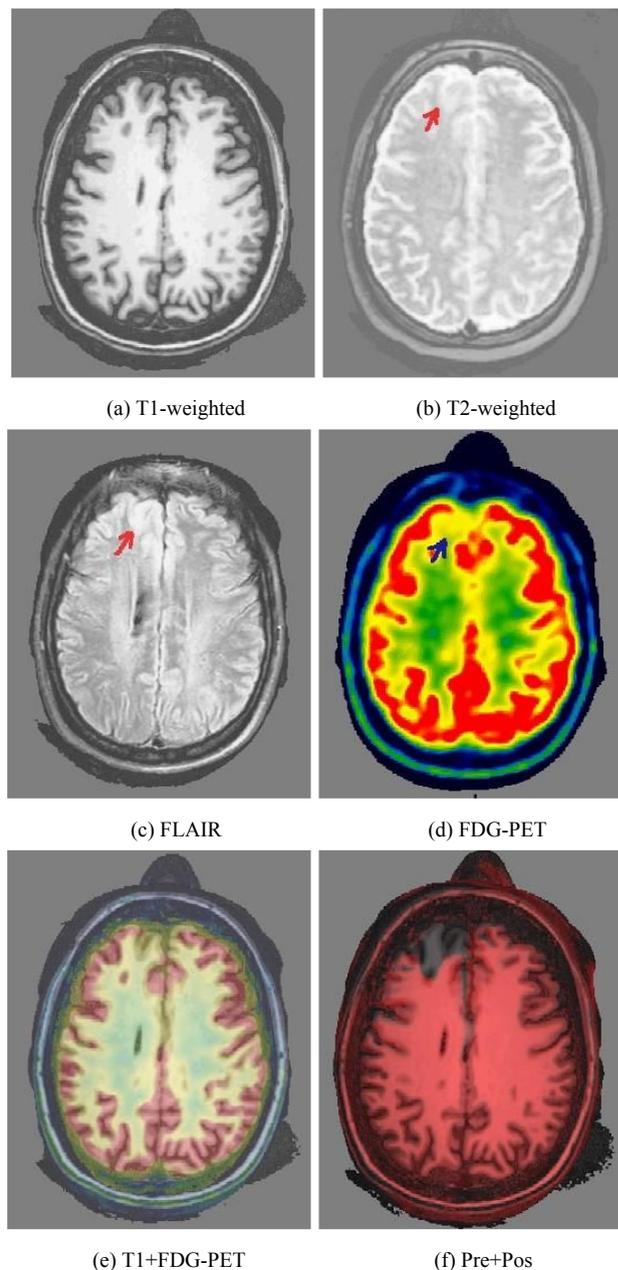


Fig. 6: Coordinated MPR views of (a) T1-weighted, (b) T2-weighted, (c) FLAIR MRI scans, (d) FDG-PET, (e) a fused view of T1-weighted MRI and FDG-PET, and (f) a fused view of pre-(in grayscale) and post-operative (in red) T1-weighted MRI volumes of a patient with FCD.

Just for exploring the potential of multimodal visualization in assessing the extent of resected area, we present additionally in Figure 6(f) a view of blended preoperative T1-weighted MR volume (in grayscale) and postoperative scan (in red). Assessing the exact amount and location of brain tissue that was removed and its relationship to the putative preoperative lesion, is another important clinical application of our tool. It can provide a visual estimation of completeness of surgical resection that is essential for a good prognosis in FCD, as well as in other types of brain pathology

To evaluate the usefulness of our tool in a complex cortical convolution we present in Figure 7(a) the curvilinear version of the volume shown in Figure 6(a). A subtle lesion spreading over the curvilinear cerebral cortex, that is not visible in Figure 6(a), becomes much more visible in Figure 7(a) (red arrow). When superimposed on the FDG-PET volume, the evidence becomes even greater in Figure 7(b). It reveals the blurring as well as lower metabolism (yellowish colored) in the lesioned area. Such dynamic display allows for better visualisation of these subtle lesions and their relationships with the complex brain anatomy by imaging experts and non-experts alike.

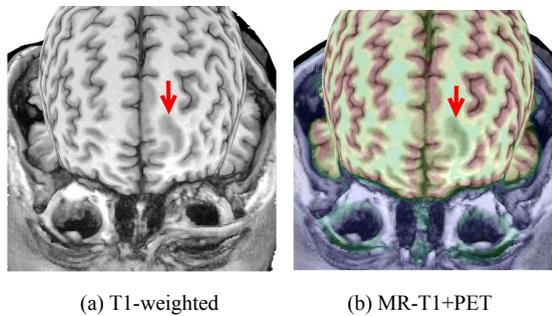


Fig. 7: Coordinated CR views of (a) T1-weighted MRI and (b) its fusion with FDG-PET, corresponding, respectively, to Figures 6(a) and 6(e).

Curvilinear reformatting is also helpful in the evaluation of the spatial extent of a lesioned area, as illustrates the case of histologically confirmed FCD shown in Figures 2(c) and 2(d). In these figures, the evidence of FCD are clearly noticeable: blurring in the T1-weighted MR scan and hypometabolism in the FDG-PET study (red arrows). Nevertheless, it is hard to assess the spatial extent of these lesions in conventional planar images. With curvilinear reformatting we can gain better insight into both the extent and the spatial location of the lesion as depicted in Figure 8(a) and 8(b).

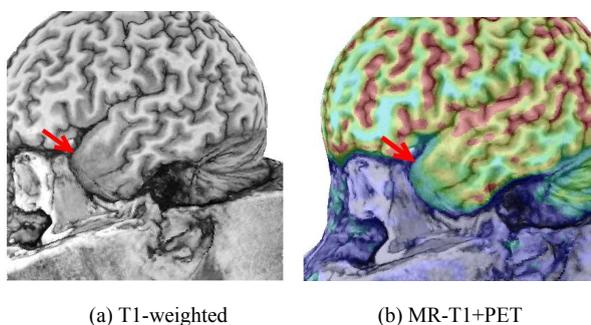


Fig. 8: Coordinated CR views of (a) T1-weighted MRI exam and (b) its fusion with FDG-PET of a patient with FCD.

Through these three cases, we show the effectiveness of our proposal in the localization of subtle epileptogenic lesions, providing a more accurate surgical plan. Also, the curvilinear reformatting allows a more realistic interaction of the neurosurgeon with the MRI anatomical localization of

the lesion for preoperative planning, since it provides a view similar to the neurosurgeons' intraoperative view.

C. Assessment of the Level of Interactivity

In order to evaluate the response time performance, we devised a series of tests for collecting render time with a microbenchmark [18]. The evaluation platform was a Windows 7 desktop with processor Intel® Core(TM) i7-2600 CPU 3.40 GHz with 8 GB RAM and a NVIDIA GeForce GTX 560 with 1 GB VRAM. We profiled our test code and extracted the render time per frame with use of NVIDIA® Nsight™.

The basic structure of our test code is divided into two passes: the first one is for determining the entry and exit points for each ray with respect to the reference volume (PreR); and the second for traversing, with fixed number of steps in the range of the entry and exit points, the volumes along a shooting ray as described in Algorithm 1 (R).

We present in Table I the amounts of time, in milliseconds (ms), each pass takes for rendering on the GPU single and fused volumes of a patient. The fused volumes are T1-weighted, T2-weighted, FLAIR, PET, ictal SPECT (iSPECT) and interictal SPECT (iiSPECT) sequences. We fixed T1-weighted sequence as the reference volume, and the other volumes were processed as floating ones. The number of traversing steps is fixed in 100 and the output image resolution in 600×600. Because only the faces of the proxy geometry are rendered, and the physical sizes of these volumes are close, almost there is no variation in the measures PreR. Moreover, comparing them with the time measures in the column R, its impact on the overall render time is negligible. Much of the time is spent in traversing rays for collecting the color contribution of each intersected sample to pixels. We equally clipped all tested volumes and recorded the amounts of time spent for traversing the clipped rays (Clip-R) to demonstrate that the computation of the inclusion test against a clipping plane does not incur any performance penalty. The response times are even drastically reduced, as the traversing stops at the intersection.

The time required for ray traversing is highly dependent on the number of texture (volume) fetches, the-step size, and the number of pixels in the output image, which are affected by the number of fused volumes, the resolution of the reference volume and the size of the screen display, respectively. From Table I we observe that additional texture fetches cause changes in time performance. Fixing the number of fused volumes in 6 and the output resolution in 600×600, we evaluated the impact of the number of steps in the render time. From Figure 9 we learn that the variation is almost linear and up to 400 steps the response can be considered instantaneous. Since the adult head sizes are commonly less than 270mm and the resolution of our clinical scanners is usually larger than 1mm, traversing with 300 steps is sufficient for our purpose.

TABLE I
RENDER TIMES WITH 100 STEPS PER FRAME OF 600×600 (MS)

Study	Dimension	PreR	R	Clip-R
T1	356×512×512	0,0293	6,60	2,20
T2	327×448×448	0,0309	6,43	2,48
FLAIR	327×448×448	0,0309	6,38	2,52
FDG-PET	400×400×222	0,0388	8,29	2,92
iSPECT	64×64×38	0,0390	4,38	2,78
iiSPECT	128×128×128	0,0388	4,66	3,05
T1+T2		0,0293	9,43	2,26
T1+T2+FLAIR		0,0292	23,78	2,45
T1+T2+FLAIR+PET		0,0292	25,53	2,54
T1+T2+FLAIR+PET+iiSPECT		0,0293	27,76	2,57
T1+T2+FLAIR+PET+iiSPECT+iSPET		0,0294	28,29	2,71

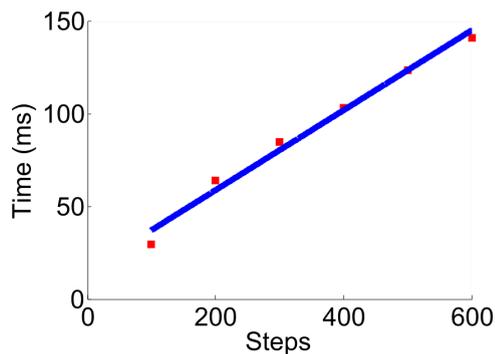


Fig.9. 600×600 image render time of 6 volumes (ms) × number of steps.

Diagnostic display devices play an important role in visual scrutinizing. Typically medical screen display has a resolution of 2018×1536. Since we do not have this sort of device, we decided to move to a computer with 4GB GPU for assessing the render time of 6 fused volumes with 300 steps onto a GPU framebuffer object. Although the new test platform (consisting of a Windows 7 desktop with processor Intel® Core(TM) i74790 CPU 3.60 GHz with 8 GB RAM and a NVIDIA GeForce GT 630 with 4 GB VRAM) delivers worse performance than the previous one, it was enough for demonstrating the performance trend of our proposed multimodal rendering. The plot in Figure 10 shows that, despite cubic behavior, the response times of our proposal are under 1.5s for diagnostic screens. According to Nielsen [19], they are within acceptable response limits for having the user feel interacting smoothly with a software application.

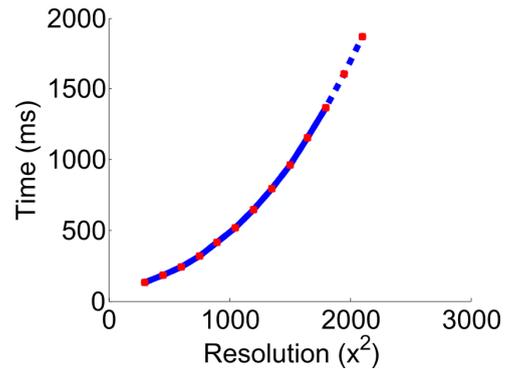


Fig.10. Render time in 300 steps of 6 volumes (ms) × output resolution.

VII. CONCLUDING REMARKS

Our work has focused on the visualization of non-invasive imaging studies to improve the detection and the delineation of a major cause of pharmacoresistant epilepsy – the focal cortical dysplasia (FCD). For reducing implementational effort, we devised a single-pass fragment shader that allows for the reuse of a previously developed interaction architecture. Though our solution is technically simple, we showed that its impact on clinical utility is considerable. Physicians were allowed to expose hidden fused information from an arbitrary angle of view and the system's response time is less than the limit (0.1s) for having the user feel that the system is reacting promptly without delay [19].

Nevertheless, our solution is not flaw free. In our proposal, only the commands for drawing reference volume are actually issued. Hence, the samples that do not overlap the reference volume are not visible. As for diagnostic purpose, the role of floating volumes is to complement the reference's signals, this is less of a problem. Otherwise, a strategy for rendering missed data should be devised. Another issue is that the correspondences of reference volume samples are seldom grid-aligned. Therefore, in theory, reconstruction and resampling on the basis of anatomical and physiological properties are necessary both for the data fusion and clipping. In our algorithm, we simply rely on the TMU that returns the best trilinearly interpolated values. It may introduce artifacts if the sampling resolution is low. Fortunately, the quality of our rendered images is considered acceptable from a medical point-of-view. Finally, the performance of our algorithm can be improved. There are some optimization techniques presented in [20] that we have not integrated in our implementation yet.

As further work, we plan to improve our GPU-based implementation and to integrate other functional and anatomical imaging studies to our diagnostic environment expanding the possibilities for physicians to explore them in coordinated views. This may help us to deepen our investigation about the color usage in data fusion. The major issue we should deal with is the user interface design. Our main goal is, however, the development of an

exploratory visualization environment of multimodal neuroimages that allows neuroscientists to deepen their understanding of the brain signals. This may help them to improve the rates of FCD detection and eventually help to predict surgical outcome.

Working closely and open-minded listening are the best practices for multidisciplinary projects. Without synergy between the medical and technological teams, the present work would not have been possible. The most important lesson we have learned from this project is that often the visualization solution required by a prospective user group is much simpler than what state-of-the-art technologies can offer or what we think it should be. The most challenging from a visualization perspective is to actually understand the user's needs and to accept the idea that the most effective solution could be trivial but highly effective. Besides, we understood that the design of an effective visualization interface requires good knowledge of application domain and visualization tools. Fairly large percentage of workload has been dedicated to the development of a "very friendly" and fast user interface that not only facilitates but stimulates, its application to clinical (very busy) routine.

ACKNOWLEDGMENT

The authors would like to acknowledge Augusto Valente Cavalcante for the co-registration codes and all inspiring works that we did not explicitly cite because of limit on the number of references. The research was supported by a CNPq-Brazil fellowship (305785/20125, 308764/2015-3), a CNPq-Brazil scholarship (165777/20141, 142018/2017-1, 141068/2017-5), the Fapesp Individual Project grant #2011/02351-0, and the Fapesp-Brazil grant #2013/07559-3 to the BRAINN Research, Innovation and Dissemination Center of the University of Campinas.

REFERENCES

- [1] C. L. Yasuda and F. Cendes, "Neuroimaging for the prediction of response to medical and surgical treatment in epilepsy," *Expert Opinion on Medical Diagnostics*, pp. 295–308, Jul 2012.
- [2] A. C. Bastos, R. M. Comeau, F. Andermann, D. Melanson, F. Cendes, F. Dubeau, S. Fontaine, D. Tampieri, and A. Olivier, "Diagnosis of Subtle Focal Dysplastic Lesions: Curvilinear Reformatting from Three-Dimensional Magnetic Resonance Imaging," *Annals of Neurology*, vol. 46, no. 1, pp. 88–94, 1999.
- [3] H. C. Batagelo and S.-T. Wu, "A Framework for GPU-based Application-independent 3D Interactions," *Vis. Comput.*, vol. 24, no. 12, pp. 1003–1012, Oct. 2008. [Online]. Available: <http://dx.doi.org/10.1007/s00371-008-0296-y>
- [4] S. Stegmaier, M. Strengert, T. Klein, and T. Ertl, "A simple and flexible volume rendering framework for graphics-hardware-based raycasting," in *Volume Graphics, 2005. Fourth International Workshop on*. IEEE, 2005, pp. 187–241.
- [5] S.-T. Wu, C. L. Yasuda, and F. Cendes, "Interactive curvilinear reformatting in native space," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 2, pp. 299–308, Feb. 2012. [Online]. Available: <http://dx.doi.org/10.1109/TVCG.2011.40>
- [6] A. P. James and B. V. Dasarathy, "Medical image fusion: A survey of the state of the art," *Information Fusion*, vol. 19, pp. 4–19, Sep. 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.inffus.2013.12.002>
- [7] D. Stalling, M. Westerhoff, and H.-C. Hege, "Amira: A Highly Interactive System for Visual Data Analysis," in *Visualization Handbook*. Elsevier, 2005, pp. 749–767. [Online]. Available: <https://doi.org/10.1016/b978-012387582-2/50040-x>
- [8] P. A. Yushkevich, J. Piven, H. Cody Hazlett, R. Gimpel Smith, S. Ho, J. C. Gee, and G. Gerig, "User-Guided 3D Active Contour Segmentation of Anatomical Structures: Significantly Improved Efficiency and Reliability," *Neuroimage*, vol. 31, no. 3, pp. 1116–1128, 2006. [Online]. Available: <https://doi.org/10.1016/j.neuroimage.2006.01.015>
- [9] W. H. Theodore, "Presurgical Focus Localization in Epilepsy: PET and SPECT," *Seminars in Nuclear Medicine*, vol. 47, no. 1, pp. 44–53, jan 2017. [Online]. Available: <https://doi.org/10.1053/j.semnucmed.2016.09.008>
- [10] S. Wu, J. E. Y. Vidalon, and L. de Souza Watanabe, "Snapping' a cursor on volume data," in *24th SIBGRAPI Conference on Graphics, Patterns and Images, Sibgrapi 2011, Alagoas, Maceió, Brazil, August 28-31, 2011*, 2011, pp. 109–116. [Online]. Available: <http://dx.doi.org/10.1109/SIBGRAPI.2011.32>
- [11] M. Ikits, J. Kniss, A. Lefohn, and C. Hansen, "Volume rendering techniques," in *GPU Gems*, R. Fernando, Ed. Addison-Wesley, 2004, pp. 667–692.
- [12] A. Leu and M. Chen, "Modelling and rendering graphics scenes composed of multiple volumetric datasets," *Computer Graphics Forum*, vol. 18, no. 2, pp. 159–171, 1999. [Online]. Available: <http://dx.doi.org/10.1111/1467-8659.00366>
- [13] S. Wu, A. C. Valente, L. de Souza Watanabe, C. L. Yasuda, A. C. Coan, and F. Cendes, "Pre-alignment for Co-registration in Native Space," in *27th SIBGRAPI Conference on Graphics, Patterns and Images, SIBGRAPI 2014, Rio de Janeiro, Brazil, August 27-30, 2014*, pp. 41–48. [Online]. Available: <http://dx.doi.org/10.1109/SIBGRAPI.2014.40>
- [14] J. Beyer, M. Hadwiger, S. Wolfsberger, and K. Buhler, "High-quality" multimodal volume rendering for preoperative planning of neurosurgical interventions," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, pp. 1696–1703, November 2007.
- [15] R. Brecheisen, A. V. Bartroli, B. Platel, and B. M. ter Haar Romeny, "Flexible GPU-Based Multi-Volume RayCasting," in *VMV*, O. Deussen, D. A. Keim, and D. Saupe, Eds. Aka GmbH, 2008, pp. 303–312. [Online]. Available: <https://goo.gl/7vgFJP>
- [16] B. Kainz, M. Grabner, A. Bornik, S. Hauswiesner, J. Muehl, and D. Schmalstieg, "Ray Casting of Multiple Volumetric Datasets with Polyhedral Boundaries on Manycore GPUs," in *ACM SIGGRAPH Asia 2009 Papers*, ser. SIGGRAPH Asia '09. New York, NY, USA: ACM, 2009, pp. 152:1–152:9. [Online]. Available: <http://doi.acm.org/10.1145/1661412.1618498>
- [17] S.-T. Wu, W. S. Loos, R. Voltoline, and J. A. I. R. Silva, "VMTK," http://www.dca.fee.unicamp.br/projects/mtk/wu_loos_voltoline_rubianes/index.html, 2017, accessed in September 2017.
- [18] S.-T. Wu, W. S. Loos, R. Voltoline, and J. A. I. R. Silva, "Sample Code of VMTK," https://github.com/vmtkUnicamp/vmtk_Multi-volume-Raycasting, 2017, accessed in September 2017.
- [19] J. Nielsen, *Usability Engineering*, 1993, vol. 44, no. 3. [Online]. Available: <http://www.useit.com/jakob/useengbook.html>
- [20] M. Hadwiger, J. M. Kniss, C. Rezk-Salama, D. Weiskopf, and K. Engel, *Real-time Volume Graphics*. Natick, MA, USA: A. K. Peters, Ltd., 2006.