Volume 0 (1981), Number 0 pp. 1–25 COMPUTER GRAPHICS *forum*

A Survey on Multimodal Medical Data Visualization

K. Lawonn[1,⋆], N.N. Smit[2,⋆], K. Bühler[3] and B. Preim[4]

[1]University of Koblenz · Landau, Germany, [2]University of Bergen, Norway,
[3]VRVis, Vienna, Austria, [4]Otto-von-Guericke-Universität Magdeburg, Germany ⋆ These authors contributed equally to this work

Abstract

Multimodal data of the complex human anatomy contain a wealth of information. To visualize and explore such data, techniques for emphasizing important structures and controlling visibility are essential. Such fused overview visualizations guide physicians to suspicious regions to be analyzed in detail, e.g. with slice-based viewing. We give an overview of state of the art in multimodal medical data visualization techniques. Multimodal medical data consists of multiple scans of the same subject using various acquisition methods, often combining multiple complimentary types of information. Three-dimensional visualization techniques for multimodal medical data can be used in diagnosis, treatment planning, doctor-patient communication as well as interdisciplinary communication. Over the years, multiple techniques have been developed in order to cope with the various associated challenges and present the relevant information from multiple sources in an insightful way. We present an overview of these techniques and analyze the specific challenges that arise in multimodal data visualization and how recent works aimed to solve these, often using smart visibility techniques. We provide a taxonomy of these multimodal visualization applications based on the modalities used and the visualization techniques employed. Additionally, we identify unsolved problems as potential future research directions.

Categories and Subject Descriptors (according to ACM CCS): Medical Imaging [Visualization], Scientific Visualization [Visualization], Volume Visualization [Visualization], Multimodal Medical Data

1. Introduction

Multimodal medical datasets consist of multiple scans of the same subject using various acquisition methods. In this way, it is possible to image different tissue characteristics. Several modalities can also be integrated directly into the hardware of a single scanner. Of these hybrid medical imaging techniques, PET/CT is currently the most widespread. More recently, PET/MR became clinically feasible, and SPECT/CT is also used in clinical practice. The common denominator in the aforementioned hybrid imaging combinations is that the structural modality (CT or MRI) depicts the anatomy of the patient in high spatial resolution, whereas the nuclear medicine modality (PET or SPECT) depicts functional processes, such as metabolism, in a lower resolution. The overall visualization goal in this case is to accurately localize regions featuring abnormal functional values based on their relation to structural anatomical information. Additionally, it is possible to combine structural information from a scan with functional information from the same scanner, such as combining an anatomical MRI acquisition with functional MRI information. In addition to combinations of functional and anatomical imaging, two anatomical imaging techniques can be combined, such as in the work by Beyer et al. [BHWB07]. Since both have a rather high spatial resolution, visualization techniques need to be updated.

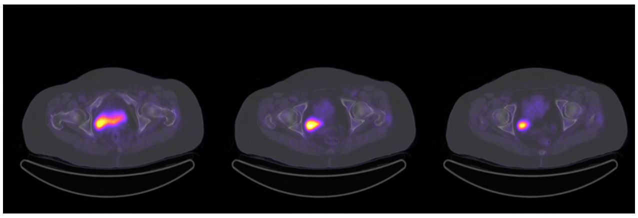

Figure 1: *2D images with superimposed CT and PET data, represented in gray scale and color respectively.*

Typically, radiologists or specialists in nuclear medicine examine data in a slice-based fashion consisting of superimposed 2D images with a combined visualization of CT and PET image data, as is shown in Figure 1. Here, physicians browse through the slices, setting window and level parameters to adjust the brightness and contrast, and examine structures of interest. Unfortunately, presenting these 2D images in such a way can hinder full and quick analysis of the data. The physician needs to check every slice, which may be time-consuming as the number of slices increases. Furthermore, physicians need to mentally fuse all information from these slices to form a correct diagnosis. Therefore, visualization techniques that provide a 3D overview are helpful to see potential abnormalities at a glance. This enables the experts to navigate to these suspicious regions for a more detailed exploration in the slice views.

In recent years, the amount of work focusing on multimodal medical data visualization increased. In 2010, there was an IEEE VIS contest on multimodal visualization for neurosurgical planning [DPL*11], further highlighting the interest in and importance of multimodal visualization. With this STAR, we aim to provide a survey of this literature. A number of medical visualization techniques are described in a survey style in the book by Preim and Botha [PB13]. Multimodal visualization, however, is only slightly touched upon as an add-on to registration, and has not been previously considered in other survey articles. In general, concepts used for visualizing medical data can be extended to other multimodal visualization application domains, but there are some specific challenges in dealing with medical multimodal data. Thus, the presented works are of interest to any visualization researcher dealing with multimodal data.

The contributions of this STAR are the following:

- We provide an overview of multimodal acquisition techniques and relate this to requirements and challenges.
- We propose a taxonomy of multimodal medical data visualization applications.
- We provide an outlook on open problems in multimodal medical visualization and a perspective on future research directions.

STAR scope. There have been several surveys on multi-field data visualizations [STS06, FH09]. In contrast to general multi-field data, measured medical image data is not as 'clean' as simulation data, due to the acquisition process, which results in noise. Living tissue is imaged with scanning parameters that favor the patients' safety over image quality, unlike for instance applications in material testing. Furthermore, several tasks such as searching for metastases, or assessment of infiltration, are unique to the field, and therefore evaluation of multimodal medical visualization needs to consider such or similar relevant tasks. While multimodal medical image data are available at both microscopic and macroscopic level, we focus on the macroscopic level, e.g. radiological image data. We also consider interaction with multimodal data.

Please note that in this survey we mainly focus on combining imaging modes acquired by different scanning techniques, and less on combining imaging modes, e.g. MRI T1 and T2, or visualization of original data combined with derived data from a single scanner. A survey focusing on DTI visualization techniques can be found in the recent work by Isenberg [Ise15]. DTI can additionally be used in multimodal applications when it is combined with fMRI or additional imaging modalities. In contrast to multi-field non-medical data, multimodal medical data is often acquired from separate scanners, which creates a need for software-based registration to align the volumes. Due to the time between the scans and patient pose differences it is not straightforward to register multiple volumes accurately. This in turn generates additional challenges in dealing with uncertainty in the form of processing errors introduced by the registration process. For an overview of medical image registration we refer to the survey by Maintz and Viergever [MV98] and more recent work on mutual-information-based registration methods by Pluim et al. [PMV03]. A multitude of registration techniques have been developed that can handle multimodal medical image registration, for instance methods aimed at multimodal brain image registration [VMN*01]. In essence, there are many registration techniques that can be employed to align multimodal datasets, such as maximizing an information-theoretic measure, e.g. normalized cross-correlation and normalized mutual information. These optimization strategies can be costly in terms of processing time, in which case GPU support is necessary [FVW*11].

While blood flow measurements can be considered functional information, a review of these techniques is out of the scope of this survey (see the recent survey on Cardiac 4D PC-MRI [KBvP*15]).

In our survey we focus on three core applications areas, namely visualizations aimed at research, diagnosis and treatment planning. Within these areas, we distinguish between applications aimed at diagnosis in oncology and cardiology, and applications aimed at treatment planning in neurosurgery and radiotherapy planning.

Paper selection criteria. We searched for papers that were related to multimodal medical visualization using the EG digital library, where we looked for the following conferences and workshops: EuroVIS, VCBM, and VMV in the last ten years. Furthermore, we used the IEEE digital library and the proceedings of the IEEE VIS conference for the last ten years. Additionally, we used Google Scholar for finding papers related to our survey, in order to integrate older papers. Specifically, we looked for the following keywords and combinations of these: combined, concurrent, CT, DTI, dual, fMRI, fused, fusion, hybrid, integrated, intermixing, neurosurgery, MRI, multi-field, multi-variate, multi-volume, multimodal, multiple, PET, planning, simultaneous, SPECT and visualization. We employed "neurosurgery" as the only application-specific search term, since this area is the key application for many multimodal visualization applications so far.

STAR organization. Our survey is structured as follows. In SECTION 2 we provide an overview of (hybrid) imaging acquisition and what the characteristics, advantages and disadvantages of each of these modalities are. In SECTION 3, we describe the typical workflow of a specialist in radiology and nuclear medicine for exploration and analysis of the data. Based on this workflow, we derive requirements for 3D visualizations. In SECTION 4 we provide a brief overview of relevant basic visualization techniques that can be employed in multimodal visualization. Experienced medical visualization researchers readers may safely skip the basic explanations given in Sections 2 and 4. SECTION 5 continues with multimodal rendering and interaction techniques. In SECTION 6, we focus on applications of multimodal visualization techniques to real-world clinical data. SECTION 7 concludes the survey and outlines unsolved problems and challenges for future research. Finally, in SECTION 8 we conclude with a brief summary.

Figure 2: *Volume renderings using dual energy CT scan of an aortic stent. On the left an overview shows the anatomical context, and on the right only the aorta and stent are shown. Springer [SH08], ©Springer Medizin Verlag Heidelberg 2008. With permission of Springer.*

2. Medical Background

In this section, we provide an overview of multimodal and hybrid imaging, as well as their applications in clinical practice. The characteristics of each of the modalities involved are summarized, and their advantages and disadvantages and associated visualization challenges are discussed.

2.1. Computed Tomography (CT)

CT is an X-ray based tomographic imaging technique that creates stacks of 2D cross sectional images. It is especially suitable to distinguish tissues such as bone, water, fat, and the air in the lungs. A contrast agent can be applied to enhance vascular structures. Recently, hybrid scanners such as dual source or dual energy CT scanners became available, delivering a final image which fuses information from high and low voltage image acquisition performed at the same time (see Figure 2) [KSF11]. Depending on the chosen imaging protocol, this technique allows for differentiation of structures like bone and contrast-enhanced blood vessels.

CT data is especially suited for high quality direct volume rendering, due to its high resolution (in general 512×512 in slice resolution and 0.3-2 mm slice thickness), high signal-to-noise ratio and standardized intensity values (Houndsfield Units), allowing the definition of re-usable and task specific transfer functions.

2.2. Magnetic Resonance Imaging (MRI)

In Magnetic Resonance Imaging (MRI), a scan is made using a powerful magnetic field. In contrast to CT, MRI scanners are highly configurable and provide a large variety of imaging protocols, allowing capture of structural as well as functional information. In general, several different MRI sequences, such as T1 and T2-weighted scans, are acquired at the same time, leading to (more or less) co-registered images. Intensities in MR images are not standardized. MRI data often exhibit an inhomogeneous gray level distribution, requiring careful preprocessing of the data, and intensity values vary depending on scanner vendor and clinic. Therefore, MRI data is challenging to visualize. Furthermore, due to the unpredictability of the intensity values, transfer functions are not directly applicable across several datasets without dynamic adaptation [RSHSG00]. Finally, MR images have generally a lower resolution and lower signal-to-noise ratio than CT images.

Besides the standard scanning protocols, there are specific MRI sequences and protocols, such as magnetic resonance spectroscopy imaging (MRSI), DCE-MRI and Diffusion Tensor Imaging (DTI). In MRSI, spatially localized metabolites in body tissues are measured. Dynamic Contrast-Enhanced MRI (DCE-MRI) is a perfusion imaging technique that measures the perfusion of tissues by blood, indicating regions damaged by stroke or infarction as well as characterizing the vascularization of tumors, helping to assess whether they are benign [TBB*99]. DTI is an extension of Diffusion Weighted Imaging (DWI), that detects the direction of white matter tracts in the brain, which represent connectivity between different areas of gray matter. DTI is used in clinical practice to assess the deformation of white matter by tumors, for neurosurgical planning and for (early) diagnosis of brain pathologies such as Alzheimer's disease, schizophrenia and multiple scelerosis [LBMP*01]. DTI data is often visualized as a scalar field of Fractional Anisotropy (FA) values, using glyphs or fiber tracking [HS15].

Functional MRI (fMRI): fMRI records subtle changes in blood flow in response to stimuli or actions and uses this information to visualize cortical activity. The most frequently employed technique is blood oxygenation level dependent (BOLD) fMRI. By having the subject perform tasks categorized into visual, motor, speech or memory tasks, different functional areas of the brain 'light up' and can be associated with the tasks performed. Additionally, fMRI is used in a research context to improve the understanding of neural networks in the brain even when the user has no task, as is the case in resting state fMRI [VDHP10].

2.3. Ultrasound

In medical ultrasound, high-frequency sound waves are employed to characterize tissue. Ultrasound can be used for diagnosis, and to guide interventional therapeutic procedures. Due to the nature of the modality, ultrasound is suitable for imaging soft tissues, such as tendons, vessels and organs, but cannot visualize bone and air, or structures lying underneath these tissue types. Based on the Doppler effect, blood flow in the heart and blood vessels can be detected. The advantages of ultrasound compared to other modalities are that it is cheap, safe, portable and real-time. However, ultrasound is difficult to interpret, due to the low signal-to-noise ratio, artifacts, and the limited field of view. Recent advances in ultrasound technology include 3D ultrasound [VBS*13], elastography and contrast-enhanced ultrasound using microbubbles.

2.4. Modalities from Nuclear Medicine

Positron Emission Tomography (PET) relies on the indirect detection, via gamma rays emitted from inside the patient, of an administered positron-emitting radionuclide (tracer). While CT and MRI scans can provide detailed anatomical data, PET scans are able to reveal functional information, such as metabolism. A common application of PET scans is to search for metastases, for which the radioactive tracer fluorodeoxyglucose (FDG), a glucose analog, is used. The metastases have higher glucose uptake than normal, and specific abnormal metabolic activity can be captured in this way. Besides oncological applications, PET is also used for neurological and cardiological diagnostic purposes. While commonly used, FDG is not the only available tracer for PET, and different tracers may be better suited for specific applications [TCV*13].

The PET data needs to be attenuation-corrected before visualization. Visualizing PET data in 3D is challenging: since normal metabolic information is also contained in PET, the highest activity measures are not always the most interesting.

Single-photon emission computed tomography (SPECT) is a nuclear medicine tomographic imaging technique that uses radioactive tracer material to detect gamma rays. In this way, it is similar to PET, but in contrast to PET, gamma radiation is measured directly from the tracer. SPECT is used for oncological diagnosis, but also for infection, thyroid or bone imaging. Besides these applications, SPECT can also provide localized function within organs for functional cardiac or brain imaging. SPECT suffers from a lower spatial resolution and contrast than PET [BW13]. Similar to PET, SPECT data are not straightforward to render in 3D.

2.5. Hybrid Scanners

When combining acquisitions from multiple modalities from different scanners, registration problems arise. To avoid the registration process, efforts have been made to integrate multiple modalities into hybrid scanners that combine the best of both worlds in structural and functional imaging. Hybrid imaging scanners combining PET or SPECT with CT are already commonplace in clinical practice, while combining PET or SPECT with MRI is more recent development [Che09] for which the first prototypes have been made [PKNS10]. In a preliminary study comparing clinical impact of PET/CT and PET/MRI, PET/MRI imaging outperformed PET/CT and more frequently affected patient management [CRS*13]. Furthermore, integrated whole-body PET/MRI was found feasible in a clinical setting with comparable reliability to PET/CT for this purpose [DSE*12]. Since 2014 the very first devices are legally allowed and used in clinical practice, focused on head and neck imaging initially.

PET/CT: While PET, CT and MRI all provide valuable information by themselves, they can be combined to provide more insight into the exact localization of suspicious metabolic activity. By combining PET and CT/MRI scans, physicians are able to detect (abnormal) metabolic activity using the PET scan and to localize this activity using the CT or MRI scan. This can aid the user in distinguishing which activity is physiologically normal and which is pathological. Since the adaptation of PET/CT is more widespread clinically, we focus on this hybrid imaging techniques in the rest of this section (see Figure 3).

CT data is also used to perform noiseless attenuation correction on the PET data, thus eliminating the need for an additional PET transmission scan for this purpose. This can reduce the total scanning time by up to 40 percent [TCYH04]. A drawback of PET/CT is that the imaging is performed sequentially, which takes more time and eliminates temporal correlation between the modalities [JWN*08].

Figure 3: *A PET/CT scanner developed by Siemens Healthcare. Photo courtesy of the Centre for Nuclear Medicine and PET, Dept. of Radiology, Haukeland University Hospital in Bergen, Norway.*

Clinical oncological applications of the combined PET/CT scanner include diagnosing and staging primary tumors, as well as localization of metastasic disease in almost any region of the body [BTB*00]. Further applications include decision making with respect to surgical operability of the tumor or treatment selection. PET/CT is additionally used to determine if cancer has recurred or to determine the difference between scar tissue and active cancer tissue. Especially when monitoring therapy, PET can detect changes in tumors earlier than CT, because metabolic changes occur sooner than anatomical size changes.

For a more elaborate description of PET/CT acquisition, we refer to Townsend et al. [TCYH04].

SPECT/CT: The benefits of combining PET and CT also extend to combining SPECT and CT. As with PET/CT, CT can provide structural anatomical information for localization, while SPECT can provide metabolic information. Similarly, CT can be used to perform attenuation correction for SPECT. SPECT is less expensive than PET, but suffers from lower contrast and spatial resolution [HH12]. A general description of the technical developments and future directions of SPECT/CT can be found in the review by Buck et al. [BNZ*08]. A more elaborate discussion on the clinical uses of SPECT/CT can be found in Mariani et al. [MBK*10].

In essence, hybrid scanners provide complimentary information that needs to be integrated, and thus fused visualizations are desirable.

3. Clinical Workflow and Requirements

In this section we examine the clinical workflow for the analysis of multimodal medical imaging data. Furthermore, we provide an overview of visualization requirements that should be considered when developing a novel visualization technique for such data.

3.1. Clinical Workflow

Multimodal medical imaging acquisitions are interpreted by radiologists and/or nuclear medicine physicians [ZMA08]. Since each of the two disciplines have their own specific skill-set, ideally, two imaging specialists (one from each discipline) would interpret multimodal images together. However, due to availability and workload this is often not possible [BOB*07]. Fortunately, more and more physicians are specialized in both radiology and nuclear medicine and work in joint departments. While specific protocols may vary per hospital, in general the images are analyzed similarly in radiology workstations of major vendors. The modalities can be separately analyzed in the views that are common for the individual modalities, such as slice-based viewing in the axial, coronal and sagittal direction. For nuclear imaging modalities, Maximum Intensity Projections (MIP) are commonplace in clinical practice. Often these MIP renderings are viewed from the front or the side of the patient. Depending on the clinical indication, more elaborate visualizations can be used, such as polar maps for cardiac SPECT data [LHG*06].

To visualize two modalities simultaneously, the typical approach in for instance PET/CT data, is to examine the CT images in grayscale with the PET images superimposed using a colormap in a 2D fused view (Figure 1). The combination of a colored overlay of functional information with gray-valued anatomical information is quite effective since brightness and color are different perceptual channels, and thus can be processed simultaneously [TG80]. Interestingly, the choice of colors in the functional color map is not standardized and varies per vendor. For instance, while most systems map the highest activity region to red, General Electric (GE) workstations map it to blue. Common clinically used visualization combinations of PET/CT data can be seen in the work by Griffeth [Gri05]. The cases presented here are often shown with a frontal MIP of the PET, axial slices of both the PET and CT separately, and finally a fused superimposed slice of the PET and CT combined. The exact configuration of the available views may vary, but typically there will be a central view with supporting linked views or a display featuring linked equally-sized views of the different modalities and slice directions.

For oncological diagnosis and treatment planning, also the combination of CT and MRI data may be important [BHWB07]. Since these data are both high-resolution data, their combination is more challenging. The simultaneous exploration of CT and MRI, with either PET or SPECT, is currently not feasible with clinically available software.

While more advanced visualization techniques could be employed to visualize multimodal images, these are not yet broadly available in clinical practice. Examples of these techniques can be found in the OsiriX software, which was specifically designed for navigation and visualization of multimodal and multidimensional data [RSPR06]. Visualization options include Multiplanar Reformation (MPR), surface rendering and direct volume rendering for fused multimodal datasets (see Figure 4).

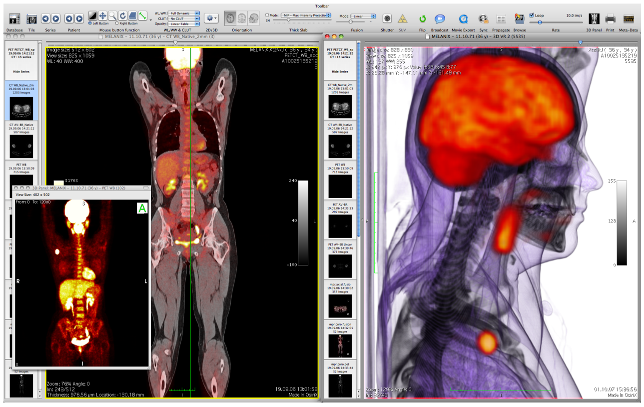

Figure 4: *2D and 3D fused PET/CT images with superimposed CT and PET data rendered in the OsiriX software [RSR04].*

3.2. Requirement Analysis

As opposed to general multi-field data, multimodal data in medicine is acquired from a (living) patient, and scanning results may be influenced by muscle relaxation, breathing, heartbeat, and movement. These local differences present special registration challenges in case of multiple acquisition techniques, and perfect matches are not guaranteed. Even in case of integrated scanners, which should reduce this problem, a perfect match cannot be achieved. The exploration of the original 2D data with the registered data superimposed may help to assess this error. Thus, 2D information are very important. When visualizing multi-field data based on simulation results, e.g., climate or technical simulations, and different parameters, e.g., pressure, temperature, gradient of pressure, that result from the simulation, this registration problem does not occur, because the simulation was performed over the same domain (simulation grid). A further difference is that medical image data result from a regular grid, whereas simulation data are mostly based on irregular grids, e.g., tetrahedron, prism, or hybrid grids. Thus, the resolution is often constant for medical data. For simulation data, on the other hand, coarser and finer grids may occur.

From a medical visualization perspective, the goals of multimodal medical data visualization include:

- reducing complexity and thus cognitive load,
- enabling, improving, or accelerating decision making, and
- providing tailored visualizations for specific applications.

In diagnosis for instance, examining all 2D slices individually to examine anatomy, pathology and metabolic uptake can become time-consuming and cumbersome. Therefore, besides the traditional 2D approaches, 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), which was first developed for use in Nuclear Medicine [WMLK89], can provide such an overview, but suffer from depth perception issues.

3D visualization techniques that are suitable for displaying hybrid multimodal data would allow the users to get a quick overview of areas of interest and localize for instance foci of elevated metabolic activity in an anatomical context. When developing a novel visualization technique or application designed for use with multimodal medical data, there are some general requirements:

1. Visualization parameters should be easily adjustable to fit the needs of the user.
2. In order to be suitable for clinical use, the technique should be fast and interactive, with minimal or no pre-processing required.

When combining functional with structural information for diagnostic purposes, the following additional requirements should be fulfilled [LSPV15]:

1. The technique should show the combination of the two or more modalities in a fused view in which the functional activity of interest is always visible.
2. The technique should relate metabolic activity to nearby anatomical structures for accurate localization.

Such a technique can be used to guide the exploration of the datasets by bringing attention to regions of interest, after which detailed inspection of these regions can be performed in 2D images. Furthermore, besides aiding diagnosis, 3D visualization techniques could be beneficial for research and treatment planning.

For research purposes, the general visualization requirements are similar to the requirements for a diagnostic application. There are differences in the requirements in terms of the time available for users to spend on clinical versus research visualization applications. In a research setting, more time can be allocated for pre-processing and interactive analysis, while in a clinical context, there is far less time available for such tasks. Further specific requirements are application-dependent and should be formulated based on the needs of the domain experts. Surgical treatment planning could benefit from multi-modal 3D patient-specific visualization to support access planning. In oncologic neurosurgery for instance, an access path to the tumor needs to be planned, taking the location of risk structures such as arteries into account. In this neurosurgical context, a visualization application should [BHWB07]:

- Provide high-quality, interactive and flexible 3D visualization
- Offer multimodal visualization for modalities such as CT, MRI, fMRI, PET or DSA.
- Provide interactive manipulation such as simulated surgical procedures, endoscopic views or virtual cutting planes.

In radiotherapy treatment planning, the pathology, i.e., radiation target, should be visualized in the context of the healthy structures at risk to receive undesired radiation damage. The requirements for such an application include [SFNB14]:

- Support for 4D PET and CT data and fusion of these modalities in a 3D view
- Visualization of segmentation data
- Visualization of dose information
- Clipping and/or masking parts of the volume
- No pre-processing and interactive parameter adjustment

These requirements demand advanced visualization techniques to fulfill the presented requirements. In the next section, we present several visualization concepts that can be applied in multimodal medical applications.

4. Visualization Techniques

Due to the nature of multimodal medical data visualization, there are always two or more volumes that need to be visualized. With overlapping extents, each of these volumes may be occluded by the others, and thus no clear view can be given on the features of each of the modalities simultaneously. Visualization techniques to reduce this problem need to incorporate heuristics for *assessing the importance* of information as well as *emphasis techniques* to adjust the visualization to the derived importance. Therefore, smart visibility, focus-and-context techniques and other emphasis techniques are crucial when dealing with multimodal medical data. In this section, we describe basic visualization techniques to cope with challenges such as dealing with occlusion, improving depth perception and presenting relevant information from multiple sources in an insightful way.

4.1. Basic Techniques

For multi-modal visualization, pre-processing is often required in terms of registration or segmentation. When combining datasets from multiple modalities using separate scanners, registration needs to be performed to align the volumes. For example, Wein et al. proposed a fully automatic registration method to align 3D ultrasound with CT scans [WKC*07, W*07]. Segmentation is not performed routinely in all clinical applications, but for instance in radiation treatment planning, the target structure (tumor) and all organs at risk are segmented manually. Here, segmentation may be performed as a basis quantification for further analysis, e.g., volume measurements and selective visualizations.

Both direct volume rendering (DVR) and indirect volume rendering (IVR) techniques can be applied to visualize multimodal medical imaging data. In IVR, the marching cubes algorithm or similar is applied to a segmented subset of the image data to determine the triangulated surface mesh [LC87]. Using the resulting mesh, various shading techniques as well as different visualization methods can be applied. On such surface meshes, it is easier to perform illustrative visualization techniques such as hatching than it would be in DVR, and these techniques have the potential to be applied in multimodal medical visualization. To determine the surface, a segmentation of the medical image data needs to be performed to select and delineate structures of interest. Since discussing alternative methods for image segmentation is out of the scope of this paper, we refer to the books by Bankman et al. [Ban08] and Setarehdan et al. [SS12] for additional information on this topic.

In contrast to IVR, DVR does not require segmentation preprocessing. In this way, there is no need to pre-determine suitable thresholds and no risk of losing critical aspects of the original data based on these choices. However, suitable transfer functions still need to be defined.

4.2. Smart Visibility

Displaying multiple structures from different modalities leads to challenging problems. Mostly, the user is interested in one specific structure, e.g., a tumor, or a class consisting of several important structures. These important structures may be surrounded by other

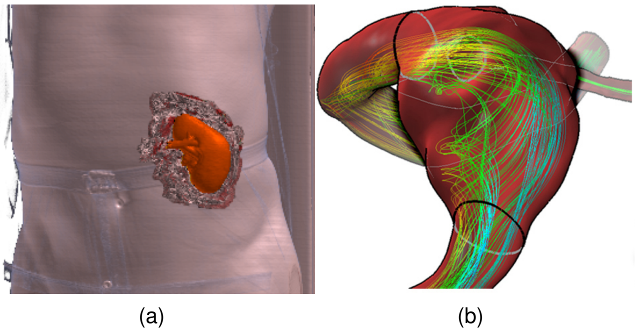

Figure 5: *Examples of smart visibility techniques applied to medical data: a cutaway revealing the kidney inside the abdomen based on CT data [VKG04], and blood flow showed inside an artery using a ghosted view based on a mesh and simulation data [GNKP10].*

objects, resulting in structures that may occlude or overlap each other. In the following, we describe advanced visualization techniques that require the important structures to be segmented.

A primary concern is how to resolve the occlusion problem such that the most important structures are clearly visible, while still indicating surrounding structures. Hiding the surrounding structures is not appropriate, since they provide spatial context. Whenever interesting objects are occluded, smart visibility techniques can be applied. For the smart visibility term, we refer to Viola and Gröller [VG05] who gave the following definition: "Expressive visualization techniques that smartly uncover the most important information in order to maximize the visual information in the resulting images". This visualization technique category comprises several sub-categories and different ways to remedy the occlusion problem. Here, we introduce the most prominent ways to resolve visibility problems comprising of cutaways and ghosted views, and present preliminary guidelines on how they could be applied to multimodal medical data.

4.2.1. Cutaway Views

Cutaway views provide a way to depict a focus object and to separate it from the surrounding context structures. Cutaway views employ, for instance, a conic object placed around the focus object. The conic object is aligned towards the camera in such a way that the camera looks inside the object from the base to the tip. Every non-focus structure that is inside the conic object is cut such that only the focus object inside the cone is visualized. The main advantage is the clear illustration of the main object without occlusion by the surrounding objects. Unfortunately, structures between the focus object and the view point of the camera are lost. In general, cutaways vary from *predefined models* [LHV12] to *focus-oriented objects* [BF08, KTP10].

Predefined models: With predefined models, view-dependent and view-independent approaches can be distinguished. Independently, the predefined model is loaded into the framework and subsequently all structures are tested to see if they are inside or outside the focus object. For instance, a lymph node can be approximated by a cylinder with a suitable size. Here, during the rendering an inside/outside test can be applied to the surrounding (context) objects with the predefined model. In case the context objects overlap with the predefined model, the fragments are discarded; otherwise, the fragments are drawn. Mostly, view-dependent cutaway techniques are used, where the conic objects are oriented along the view-vector. Predefined models can be used if the the focus object has a simple shape. Moreover, it is also only applicable if the model does not vary over time. In case of animated blood flow, predefined models cannot be employed, due to the change of the form and position of the blood flow [LGV*15].

Focus-oriented models: The second class of cutaway techniques uses focus-oriented models. Here, the conic object is constructed based on the focus object and can be arbitrarily shaped. To achieve real time performance, the determination of the conic object and the associated cutaway needs to be efficient. Previous work, such as the work by Viola et al. [VKG04], uses a Chamfer distance transform approximation [Bor86] to the conic object in order to calculate the cutaway (see Figure 5(a)). Thus, a depth image of the conic surface is calculated, which is then used to cut away regions that are closer to the viewer then the conic object itself. Other approaches were developed to speed up this calculation [RT06]. Here, the cutaway surface function C is defined by:

$$C(p) = \max_{q \in R}(q_z - m \cdot \|q_{xy} - p_{xy}\|),$$

where R contains the focus object, m defines the slope of the cone by $m = \tan^{-1}\theta$, and p_{xy} is the pixel's current 2D location. The key idea is to determine the conic object in the view plane with a jump flooding algorithm. Usually, the conic object is generated by applying the standard flood fill, an iterative process in which in every iteration a pixel passes its value to its direct neighbors. Contrarily, the jump flood algorithm uses the same approach, but the neighbors vary for each iteration. Thus, a pixel passes its value to its neighbors, which have a certain distance, and for each iteration the step size is halved. Cutaway generation was used to visualize animated blood flow in such a way that a clear view on the blood flow within a vessel is always guaranteed [LGP14].

Application to multimodal data: Cutaway views can be directly employed in multimodal visualization. For instance, one modality could be displayed outside the cutaway region while another is displayed within the cutaway. Furthermore, cutaways can be used on both modalities to reveal structures of interest hidden within a larger anatomical structure.

For instance, a cutaway can be used to display hybrid imaging PET/CT data, where high PET activity needs to be depicted. By using a cutaway view in these regions, the high PET activity is always visible and it provides a clear view inside the anatomical CT data and can provide a correct localization of the occurrence, whereas a simple maximum intensity projection cannot give reasonable depth cues. Such a view might be useful not only for diagnosis, but also for patient-doctor and interdisciplinary communication, as well as for treatment planning.

4.2.2. Ghosted Views

An easy way to illustrate the focus object and the surrounding objects is to use transparency. The main drawback of employing multiple transparent objects is that they may lead to visual clutter. Furthermore, it is hard to interpret the shape and ordering of overlapping regions. To remedy this, ghosted view techniques attempt to improve the transparency depending on the underlying object. Here, a distinction can be made between smart transparency techniques and interactive approaches. In smart transparency, the strength of transparency depends mostly on curvature or normal vector information. One approach is to set the transparency very high if the surface is nearly planar and to increase the opacity if the curvature varies. That helps to identify regions where the morphology changes strongly. This was successfully applied for instance to vessel visualization with embedded flow [GNKP10, GLH*14], an example of which can be seen in Figure 5(b). The transparency is set by $F = 1 - \langle V, N \rangle^r$, where V represents the view vector, N the surface normal, and r represents the shininess constant. The superiority of this approach over standard semi-transparent visualization related to cerebral blood flow was confirmed in a perceptual study by Baer et al. [BGCP11].

An additional transparency technique was inspired by suggestive contour lines, a line drawing technique that emphasizes characteristics of the surface. This approach was used to develop a shading approach, based on a suggestive contour scalar field [LGP14]. Here, the scalar field is used to determine zero-crossings after which positive and negative values are used to shade the object in orange or cyan, respectively.

Application to multimodal data: Ghosted views can be applied to multimodal data visualization in order to remove emphasis of structures that are less important from one or more modalities. A ghosted view can be employed to show an anatomical context in order to localize functional data without occluding the view. For instance, a brain surface can be rendered as a ghosted view, while functional data is rendered in an opaque style. This yields a visualization where the brain surface can be recognized, without being occluded, whereas the functional brain data is directly depicted. The main challenge is to provide a transparency technique that is not disturbing, but provides as much morphological context as possible, while not occluding underlying structures. Conventionally, a simple transparency setting is used, which does not provide much shape information.

4.3. Focus-and-context visualization

For focus-and-context visualizations, we assume that at least the focus object is segmented. Even though a lot of techniques exist that provide insight into the focus object, the question is how to represent them in the context of the surrounding structures. A straightforward approach is to use the same shading technique on all structures, which may lead to a hindered perception of the focus object by the context objects. Thus, the main goal of the *focus-and-context* visualization is to illustrate the focus object so that it is perceivable, and simultaneously illustrating the surrounding context structures without distracting from the focus. Techniques vary from using unsaturated colors for shading, to different shading methods per object class, to even varying entire rendering concepts.

Line drawings. One of the first line drawings that appeared in combination with focus objects was presented by Interrante et al. [IFP96]. They used a hatching technique for the surrounding objects which are represented by the isosurface of radiation. In this case, anatomical data is combined with a scalar field resulting from the simulation of dose distribution, performed for radiation treatment planning. The lines were aligned along the principle curvatures directions for an improved shape perception, whereas the inside focus object was simply shaded (see Figure 6(a)). Another line drawing technique that appeared in combination with focus objects was presented by Treavett and Chen [TC00]. They developed line drawings that can be applied on volume data sets and furthermore showed how they can be used to illustrate the surrounding context object, e.g., the skin, with the focus object, e.g., the skull. Again in the context of volume rendering, Hauser et al. [HMIBG01] used combined visualization techniques to illustrate different objects. For instance, bones are rendered with direct volume rendering, surface rendering is applied to the vessels, and contours indicate the skin (see Figure 6(b)). Lum and Ma [LM02] presented a way to render a dataset with non-photorealistic techniques efficiently. Several illustrative techniques can be applied with interactive framerates. It was shown that line drawing concepts can successfully be used for focus-and-context visualization [TIP05]. The core idea here was to provide the flexibility to combine surface-, volume-data, and line drawing for a sparse visual representation of context objects. Here, not only line drawings, but illustrative visualization techniques in general were used to visualize context objects. The key idea of illustrative visualization is to provide a simplified and expressive depiction of a scene or problem [Law15]. Here, details are omitted purposefully in order to convey only the important information.

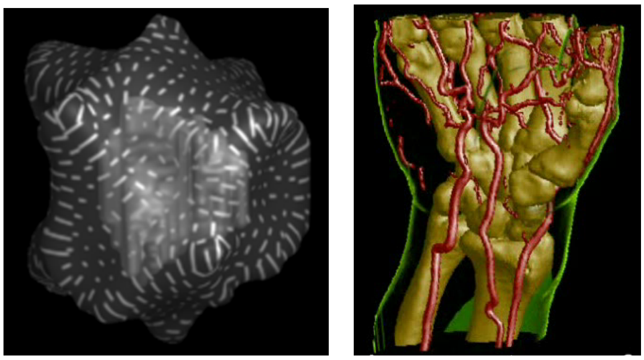

(a) (b)

Figure 6: *One of the first focus-and-context illustrations with hatching patterns used for the isosurface of radiation dose along with a tumor to be destroyed (©[1996] IEEE. Reprinted, with permission, from [IFP96]). On the right, a focus-and-context visualization applied to a hand data set with arteries [HMIBG01].*

Tietjen et al. combined and tested [TIP05] various rendering techniques. While their technique was not shown on multimodal data, it is applicable for feature fusion. Later, Tietjen et al. [TMS*06] used slice-based visualizations in combination with 3D illustration techniques (see Figure 7) for an example of focus-and-context visualization in the medical context. In a similar way, a slice or slab of CT/MRI data may be integrated in a SPECT/PET visualization or vice versa. The slice/slab may be even enhanced with contours from a segmented object. For PET/CT, Kim et al. [KEF07] used different rendering techniques to illustrate the PET enrichment and the anatomical context by using the CT data. Various rendering techniques were also used in the context of simultaneously rendering anatomical and functional brain data [JBB*08]. For an overview of focus-and-context approaches, we refer to Cohen [Coh06] and Bruckner et al. [BGM*10]. Another importance-based direct volume rendering approach was presented by Pinto and Freitas [dMPF10, dMPDSF11]. Furthermore, they combined standard volume rendering approaches with illustrative visualization techniques (see Figure 9). For focus-and-context visualization with PET/CT hybrid data, Bramon et al. [BRB*13] suggested visualization techniques to emphasize the PET activity. Another approach to explore and visualize interesting regions was proposed by Abellán et al. [ATGP13]. They demonstrated their approach on MR/PET data, and used illustrative visualization techniques to convey information.

Lens-based visualization. Focus-and-context visualization can also be applied interactively, e.g., with *magic lenses* that can be used to look through objects and to reveal interesting parts [BSP*93]. The magic lens can be applied in many situations, ranging from whole body that uses a lens to reveal organs, to vessels where a lens depicts special properties of the blood flow [GNBP11]. A special application in medicine is to employ the vertical symmetry of the human body and propagate the movement of the lens in one part of the body to the other part, thus supporting the comparison between a suspicious region and the analogue region on the other part of the human body. For 3D ultrasound, Schulte zu Berge et al. [SzBBKN14] presented a framework where the user can set various predicates such that interesting regions are highlighted. Increasing the slider for the vessels makes it more visible than the surrounding structures (see Figure 8(a)). In the context of vascular models and tumor visualization, Lawonn et al. [LLPH15] presented a way to depict the vessels with line drawings techniques while the tumor is visualized with diffuse shading (see Figure 8(b)). Furthermore, lenses can also be employed to control which source is displayed, e.g., in the lens region PET (or PET and CT combined) and outside only CT. Thus different layers may be shown interactively [KSW06].

The distinction between smart visibility and focus-and-context visualization is not that clear. We argued that focus-and-context visualization techniques are a subcategory within the broader category of smart visibility approaches. The main goal of smart visibility is to provide insights into interesting regions or objects, but this goal can be also be achieved by adding 2D image planes in a 3D visualization representation [HE98, BHW*07]. In this case, there is no focus object involved, although it is still clearly a smart visibility technique. Thus, we consider focus and context visualization a subgroup of smart visibility techniques, which is a broader category including cutaways, ghosted views or exploded views, as well as techniques in which there is no focus object.

Application to multimodal data: Focus-and-context techniques are applicable to multimodal medical data and can be employed in two main ways. First, both modalities may be combined and a region of interest within the datasets can serve as a focus area, while the context is rendered in a different style. Secondly, one modality with a higher importance can be rendered as the focus dataset, while the other is rendered to provide context. The rendering style required strongly depends on the application. For instance, if the focus region consists of a tumor and the physician is interested in the distance to nearby vessels for anatomical context, then the complete morphological vessel structure may be not as important as an indicated position of the nearby vessel structures. In this case, a simple contour that shows only the outline and allows for fast recognition of the spatial extent may be sufficient. But in case of a liver tumor, the spatial impression and the shape of the liver is important for treatment planning. In this case, a more advanced approach needs to be applied to illustrate the tumor in a reasonable way, while simultaneously depicting the liver surface with important spatial and morphological features. Different rendering styles may be used for multimodal visualization, e.g. contours representing boundaries of objects in one modality may be overlaid to the visualization of a second dataset, where no contours are used.

4.4. Summary

There are different channels with which we perceive information, e.g., color, or orientation of elongated objects. Thus, we can perceive information simultaneously if it is encoded in different channels. The major theory here is the Feature Integration Theory, developed by Anne Treisman [TG80]. Multimodal visualization applications may benefit from the use of different visualization techniques that can be easily perceived simultaneously, e.g., color is only used for depicting one dataset, and line drawing techniques for the other. In order to represent multiple modalities in a single rendering, smart visibility techniques are crucial to prevent visual clutter and occlusion problems. Specifically, multimodal medical visualization can benefit from deciding which dataset has priority, i.e., the "focus dataset", and which is the context dataset. In multimodal visualization design, one could consider indicators of importance of a modality, such as the size and resolution, which could for instance be indicative of which dataset is best used for anatomical detail. Then one could think about which techniques are primarily suitable to show the focus dataset, likely revealing more details, and which to show the context dataset. This approach forms a bridge between single data focus-and-context techniques and focus-and-context in the context of multimodal data. There may be also situations where both volumes have the same importance, e.g., CT and MRI, and in these situations, it might be more appropriate to consider a region within both datasets the focus area, while the region outside is rendered as context. In conclusion, visualization techniques that have been developed for single modality visualization can be adapted to multimodal data, but this requires careful consideration of the application area requirements and in some cases extensions of existing techniques.

5. Rendering and Interaction Techniques for Multimodal Data Visualization

In this section, we discuss rendering and interaction techniques for multiple or fused volumes. First, we discuss recent contributions on rendering aimed at fusing multiple volumes together at different points in the rendering pipeline. After this, we discuss interaction

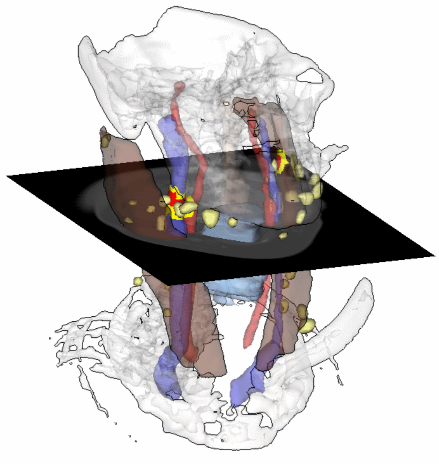

Figure 7: *A focus-and-context visualization applied to a CT scan of the neck revealing lymph nodes and vascular structures with a slice of imaging data [TMS*06].*

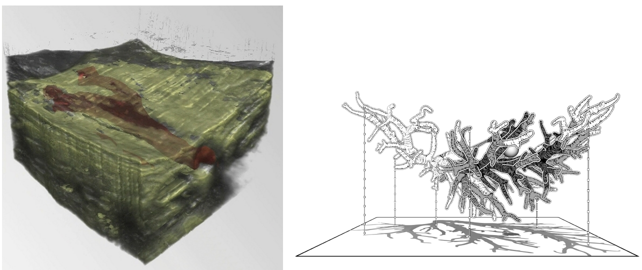

(a) (b)

Figure 8: *Focus-and-context visualization of a 3D ultrasound dataset where the arteries are shown [SzBBKN14] and a visualization where vessels and a tumor, reconstructed from CT data, are illustrated with line drawings [LLPH15].*

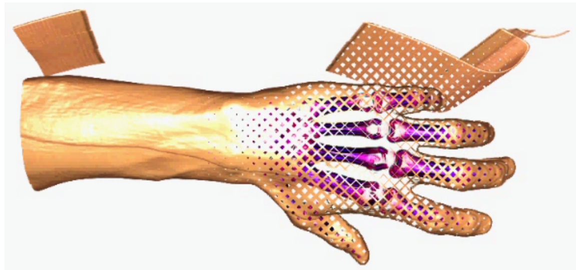

Figure 9: *Focus-and-context visualization based on a CT scan of the hand and wrist. The bones are shown by applying a cutaway technique [dMPF10, dMPDSF11].*

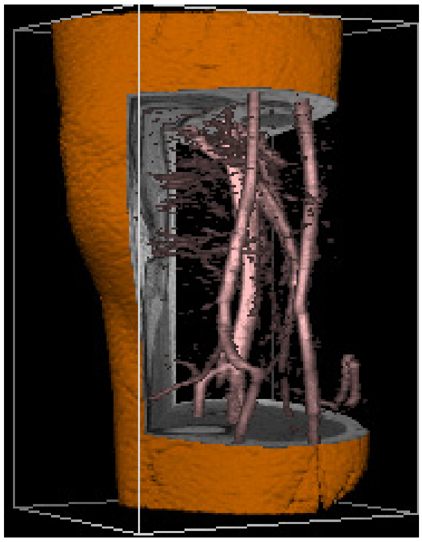

Figure 10: *First multimodal volume rendering technique by Höhne et al. visualizing vascular MRA and anatomical MRI data simultaneously [HBTR88].*

techniques that are designed for use with multimodal medical data, including clipping and specialized transfer functions.

5.1. Rendering Techniques

In this subsection we provide an overview of rendering techniques suitable for rendering two or more volumes simultaneously. Although Drebin [DCH88] and Levoy [Lev88] are often cited as the developers of the first volume rendering approach, actually Höhne [HB86] was the first to render a volume based on CT data. As a follow up, Höhne [HBTR88] was also the first to develop a multimodal visualization technique visualizing MRI combined with MRA (see Figure 10). Cai and Sakas presented a seminal work that allows rendering of multiple volumes [CS99]. They presented a method to achieve a reasonable mix of opacity, color, and illumination. Furthermore, they defined three levels of volume intermixing:

- Illumination model level intermixing
- Accumulation level intermixing
- Image level intermixing

In illumination model level intermixing, the opacity and intensity at each sample point is calculated directly from a multi-volume illumination model, instead of mixing several opacity and intensity values. This type of volume intermixing is the most complex, but also the most realistic. In accumulation level intermixing, the intermixing is performed during ray sample accumulation by mixing opacities and intensities from different volumes. This method requires changes to the rendering pipeline and the time consumption can be high, but provides correct depth cues. The simplest intermixing method is image level intermixing, where the two rendering images are merged on the pixel level. For this, no changes need to be made to the rendering pipeline, but there is a lack of correct depth cueing. This approach was successfully applied in the context of radiotherapy treatment planning, combining CT, dose and segmented object volumes.

Wilson et al. presented a visualization approach to depict different image modalities together at interactive framerates using different rendering styles [WLM02]. To render the result efficiently, they used the GPU and demonstrated their results with two case studies. One of the case studies featured a multimodal mouse data set combining PET and MRI acquisitions. Ferre et al. discussed strategies to visualize multimodal volume datasets using direct multimodal volume rendering (DMVR), in which they determine in which steps of the rendering pipeline data fusion must be performed in order to accomplish the desired visual integration [FPT04]. Furthermore, they stated requirements and proposed five rendering methods that differ in the step of the rendering pipeline at which the fusion is performed: property -, property and gradient -, material -, shading - and color fusion. They used SPECT and MRI data to present a combined visualization and evaluate the suitability of the five methods for this data. Another technique to render multimodal volumes, which additionally solves depth cueing issues with reduced time consumption, was presented by Hong et al. [HBKS05]. They conducted experiments that showed that their method can distinguish the depth of overlapping regions and produces rendering results faster than conventional software using High Level Shader Language (HSLS) for efficient volume fusion. They show the results of their method on PET and CT images in a performance evaluation. Rößler et al. described a GPU-based multi-volume rendering scheme that allows users to visualize an arbitrary number of volumes interactively, focused on functional MRI data [RTF*06]. Subsequently, they introduced dynamic shader generation for multi-volume ray casting, by allowing users to define an abstract render graph, to overcome the fact that typically shaders can only be programmed by GPU experts [RBE08]. In an accumulation-level volume rendering context, their system automatically generates different shaders per volume from the configuration of an interactive abstract render graph. Brecheisen et al. employed a framework to render different volumetric datasets together [BBPtHR08]. They used depth peeling to visualize overlapping volumes that can be intersected with an arbitrary number of geometric shapes. Kainz et al. proposed a framework for rendering multiple volumes on the GPU [KGB*09]. Theirs was the first framework for multi-volume rendering that still provided interactive frame rates while rendering over 50 arbitrarily overlapping volumes. Recently, Sunden et al. introduced a volume illumination technique specifically designed for use with multimodal data [SKR15]. They proposed a new *light-space-based* volume rendering algorithm, that employs illumination-importance metrics to compress and transform multimodal data into an illumination-aware representation. Their method was applied to CT and MRI data, as well as microscopy and simulation data and evaluated based on the quality and performance. Lindholm et al. [LLHY09a] presented a rendering algorithm for hybrid volume-geometry data. They combined volume data and geometry-based data, i.e., a mesh. They applied their approach to proteins, vessels with blood flow, and DTI. Additional related work can be found in the state-of-the-art report on the visualization of multivariate scientific data presented by Fuchs and Hauser [FH09]. Schubert and Scholl provide a performance and perceptual comparison of GPU-based multi-volume ray casting techniques [SS11]. Furthermore, they presented an overview of visualization techniques that use data intermixing approaches and direct volume rendering methods that use ray casting. They add Classification Level Intermixing to the three levels defined by Cai and Sakas [CS99], which mixes volumes by linear combination of the sample values at the beginning of the rendering pipeline. For a general overview of large-scale volume visualization techniques, we refer to the state-of-the-art report by Beyer et al. [BHP15].

In the area of medical image fusion, Gupta et al. [GRB08] presented a measure to evaluate PET/MRI image fusion. They confirmed the usefulness of their approaches in different experiments. For this, they used an entropy measure to combine the interesting parts in PET and MRI, but this approach also was applied to CT combined with MRI. Lindholm et al. introduced fused multi-volume direct volume rendering using a Binary Space Partitioning (BSP), aimed at medical volume rendering of multimodal data [LLHY09b]. Lindholm et al. provided a GPU-based ray casting approach to visualize intersecting volumes with regard to an efficient depth sorting of the resulting fragments. Bramon et al. fused image modalities 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. They evaluated the potential of their technique with medical experts and concluded that it is potentially useful for planning radiotherapy, treatment monitoring, and planning brain surgery. In their evaluation, they proposed several information maps and fused data sets and had the experts vote on the quality. Their method performed well in differentiating between bone and cerebral tissue, as well as between morphological and functional data. More recently, Kim et al. introduced a slab-based intermixing method for fusion rendering of multiple medical objects [KKL*16]. We refer to James and Dasarathy [JD14] and Galande and Patil [GP13] for a more extensive discussion.

5.2. Interaction Techniques

In this subsection we describe several interaction techniques that were successfully used in the context of multimodal medical data. Interaction techniques primarily serve to adjust which portions of the dataset are visible. Visibility may be adjusted with clipping and cutting on a geometrical basis and with transfer functions on the basis of attribute values, e.g., intensity values, gradient magnitude or curvature.

Clipping and cutting. In the work by Hastreiter et al. clipping planes were employed per volume to cut separate modalities in different ways [HE98]. This can be used to look at the slices while simultaneously analyzing the 3D visualization. Furthermore, they combined CT and MR images in a single view. Fairfield et al. proposed a custom clipping solution for co-registered MRI and CT, which they refer to as 'curtaining' [FPJ*14]. This allows the user to define a clipping region in which the other modality will be shown. Manssour et al. propose a method to visualize inner structures in multimodal volume data employing cutting and data intermixing [MFOF02]. They applied it to the skull where cutting planes were used to analyze the brain. More advanced ways of clipping data were developed by Weiskopf et al. via per-fragment operations in texture-based volume visualization [WEE02]. While their method was applied to a single modality, their technique was successfully employed by Rößler et al. to clip the brain based on regions in an atlas of the human brain combined with fMRI data [RTF*06]. More recently, a membrane clipping approach was proposed that avoids cutting through features in the

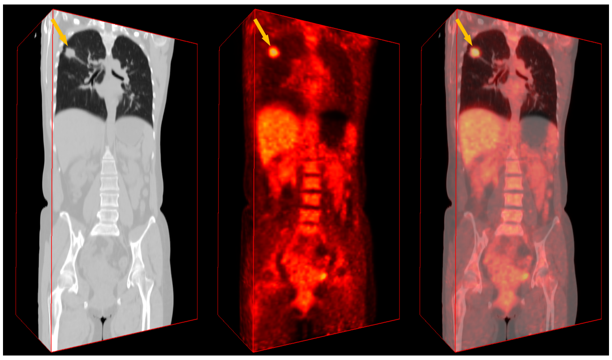

Figure 11: *PET/CT data are mixed to visualize the volume in a single view [KEF07]*

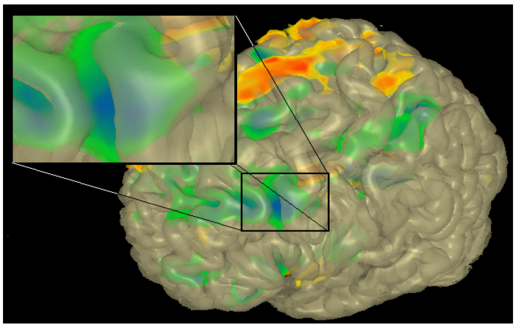

Figure 12: *A non-polygonal isosurface revealing patient-specific functional information in anatomical context of a template brain [RTF*06]. Image courtesy of the authors.*

data [BBBV12]. While it was not directly applied to multimodal datasets, it could be used in such a way. Furthermore, their technique features slab rendering which is a trade-off between slice-based viewing and full 3D rendering.

Transfer functions. To tackle occlusion problems, transfer functions are essential and have been designed for multi-modal data. 2D transfer functions, introduced by Levoy in 1988 [Lev88], are an essential method for multimodal visualization, since they allow the user to emphasize boundaries and details within structures. Kniss et al. applied a multi-dimensional transfer function to multivariate weather data that allows experts to readily identify specific zones for their analysis task [KHGR02]. They presented a case study in which they explore the utility of multidimensional transfer functions for the visualization of multivariate fields. Furthermore, Dobrev et al. proposed to a hierarchical clustering approach for visual analysis of multivariate volume data, which allows the user to operate in cluster space, rather than transfer function space [DVLLL11]. 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 a single view [KEF07]. They provided techniques such that the volumes can be merged properly with different transfer functions, see Figure 11. A general information-based approach for transfer functions was introduced by Haidacher et al. [HBKG08]. They provided a transfer function space to visualize certain tissues of multimodal data. They applied their method successfully to PET/CT data of the brain.

Joshi et al. introduced new interaction techniques to explore and visualize multimodal data [JSV*08]. Their technique allows for a precise control of the shape of the region in the brain that can be used to crop data during exploration and surgery. For their approach they used SPECT/MRI and fMRI data. Additionally, Ropinski et al. proposed interaction techniques to show a close-up of an interesting region in the context of the rest of the data [RVB*09]. They applied their technique to PET/CT data registered with MRI. More recently, Haidacher et al. proposed an approach for volume analysis based on multimodal surface similarity [HBG11]. With this, the similarity space generated can be used for isosurface selection in applications like Dual Energy Computed Tomography (DECT). The similarity map can be used to set different isovalues, e.g., for DECT, such that inner structures can be set opaque whereas outer structures can be set as transparent. Correa and Ma [CM11] employed visibility-driven transfer functions to illustrate important structures in comparisons to the surrounded regions. Their approach provides the user with graphical cues that inform about the contribution of particular scalar values to the final image. Furthermore, they presented a semi-automatic transfer function design that solves an energy minimization problem, such that the visibility of an initial opacity transfer function that provides the desired importance is maximized. A complete survey on transfer functions for volumetric data can be found in the work by Ljung et al. [LKG*16].

6. Applications

In this section, we provide an overview of recent multimodal medical visualization papers that are application-oriented. We further subdivide this section in applications developed in a research-oriented, diagnosis or treatment planning and guidance context. For diagnostic purposes, we distinguish applications developed for cardiology and those aimed at oncology. The treatment planning and guidance applications are further categorized into neurosurgery and radiotherapy planning.

6.1. Medical Research

Multimodal medical data visualization applications were developed in a research context related to vascular pathology development as well as neuroscience. Rößler et al. described a GPU-based multi-volume rendering scheme that allows users to visualize an arbitrary number of volumes interactively (recall Section 5.1.1), and their technique was specifically focused on functional brain images [RTF*06]. Their tool was aimed at supporting cognitive neuroscientists in experimental studies and to communicate results to non-experts and employs a template brain with patient-specific fMRI data (see Figure 12). PET/CT scans are often requested in cases related to clinical oncology for initial cancer staging or a follow-up during or after treatment [BOB*07]. However, Ropinski et al. visualized mouse aorta PET/CT scans in a medical research-oriented application related to the formation of plaque [RHR*09].

They propose a linked multi-view approach, using a specialized

Table 1: *Table of references sorted according to the application type (Type), area (App), and to the acquisition scanners (hybrid scanner, single scanner and multiple scanners). It is stated if any segmentation (Seg) is needed, if the approach was evaluated (qualitatively (Ql), quantitatively (Qt), with domain experts (D) or non-domain experts (N) or the performance (P), and what types of image modalities are used. The visualization techniques that are used are mentioned, subdivided into cutaways (CA), ghosted view (GV), focus-and-context (FC), and illustrative visualization (I-Vis).*

References	Type	App	Scanner	Seg	Eval	Image Modality							Vis technique			
						CT	MRI	fMRI	DTI	PET	SPECT	US	CA	GV	FC	I-Vis
[DHS*13]	Res	Vasc	Hybr	✓	QlD	✓	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗
[RHR*09]	Res	Vasc	Hybr	✗	QlN	✓	✗	✗	✗	✓	✗	✗	✗	✗	✓	✗
[ESM*05]	Res	Neuro	Mult	✓	✗	✗	✓	✗	✓	✗	✗	✗	✓	✓	✓	✗
[vDMvLB12]	Res	Neuro	Sing	✗	QlD	✗	✓	✓	✗	✗	✗	✗	✗	✓	✓	✗
[JNM*09]	Res	Neuro	Sing	✗	QlN	✗	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
[NOE*10]	Res	Neuro	Sing	✗	QtN	✗	✓	✓	✗	✗	✗	✗	✓	✗	✓	✗
[RTF*06]	Res	Neuro	Sing	✗	QtP	✗	✓	✓	✗	✗	✗	✗	✓	✓	✓	✗
[SZV01]	Res	Neuro	Sing	✓	QlD	✗	✓	✓	✗	✓	✓	✗	✗	✗	✗	✗
[SZP*97]	Res	Neuro	Mult	✓	QlD	✗	✓	✗	✗	✗	✓	✗	✗	✗	✗	✗
[JKE*13]	Diag	Onco	Hybr	✓	QtP	✓	✗	✗	✗	✓	✗	✗	✗	✓	✓	✗
[KEF07]	Diag	Onco	Hybr	✗	✗	✓	✗	✗	✗	✓	✗	✗	✗	✓	✓	✗
[LSPV15]	Diag	Onco	Hybr	✗	QlD	✓	✗	✗	✗	✓	✗	✗	✓	✓	✓	✓
[VNØ*08]	Diag	Onco	Mult	❍	QlN	✓	✗	✗	✗	✓	✗	✓	✓	✗	✓	✓
[HSF*08]	Diag	Cardio	Mult	✓	QtD	✗	✓	✗	✗	✗	✗	✗	✗	✓	✗	✗
[OKG*06]	Diag	Cardio	Mult	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗	✓	✗	✗
[TBB*07]	Diag	Cardio	Mult	✓	✗	✗	✓	✗	✗	✗	✗	✗	✗	✓	✗	✗
[KGS*14]	Diag	Cardio	Mult	✓	QtD	✓	✗	✗	✗	✗	✓	✗	✗	✓	✓	✗
[BHWB07]	Trea	Neuro	Mult	❍	QlD	✓	✓	✓	✗	✓	✗	✗	✓	✗	✓	✗
[BBM*07]	Trea	Neuro	Sing	✓	QlD	✗	✓	✓	✓	✗	✗	✗	✗	✓	✓	❍
[BLE*13]	Trea	Neuro	Mult	✓	QlD	✓	✓	✗	✗	✗	✗	✗	✓	✓	✓	✗
[BJH*09]	Trea	Neuro	Sing	✓	✗	✗	✓	✓	✓	✗	✗	✗	✓	✓	✗	✓
[DPL*11]	Trea	Neuro	Sing	✓	QlD	✓	✓	✓	✓	✗	✗	✗	✓	✓	✓	✗
[JBB*08]	Trea	Neuro	Sing	✓	QlD	✗	✓	✓	✗	✗	✗	✗	✓	✓	✗	✓
[JFS*00]	Trea	Neuro	Sing	✓	QlD	✗	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
[JSV*08]	Trea	Neuro	Mult	✓	QlD	✓	✓	✓	✗	✗	✓	✗	✓	✗	✓	✗
[KNS*12]	Trea	Neuro	Mult	✓	QtD	✓	✓	✗	✗	✗	✗	✗	✗	✓	✗	✗
[MSE*06]	Trea	Neuro	Mult	✗	QtP	✗	✓	✗	✓	✗	✗	✗	✓	✗	✓	✗
[NTS*10]	Trea	Neuro	Sing	✓	QlD	✗	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗
[NMW*04]	Trea	Neuro	Mult	✓	QlD	✓	✓	✗	✗	✗	✗	✗	✗	✓	✗	✗
[RRRP08]	Trea	Neuro	Sing	✓	QlD	✗	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗
[RSHP08]	Trea	Neuro	Mult	✗	✗	✗	✓	✗	✗	✗	✗	✗	✓	✗	✓	✗
[SKG*98]	Trea	Neuro	Mult	✓	Ql/QtD	✓	✓	✗	✗	✗	✗	✗	✓	✗	❍	✗
[WRD*11]	Trea	Neuro	Sing	✓	QlD	✗	✓	✗	✗	✗	✗	✗	✓	✗	✓	✗
[NRS*14]	Trea	Radio	Sing	❍	Ql/QtD	✗	✓	✗	✗	✗	✗	✗	✗	✗	❍	✗
[SFNB14]	Trea	Radio	Hybr	✓	✗	✓	✗	✗	✗	✓	✗	✗	✓	✓	✓	✗

straightened multipath curved planar reformation combined with a multimodal vessel flattening technique (see Figure 13). In a follow-up work, Diepenbrock et al. additionally looked into PET/CT visualization of mice arteries [DHS*13]. They used the vessel wall extracted from the CT scan to perform a normalized circular projection which allows the user to judge PET signal distribution in relation to the deformed vessel.

Nguyen et al. used an approach to visualize and interact with real-time fMRI data in a neuroscience research context [NOE*10]. They treat the fMRI signal as light emission and render it in the context of a patient-specific high resolution reference MRI scan. Using this technique, the brain glows and emits light from active functional regions with a 2 second delay of the measured fMRI signal. In the work by van Dixhoorn et al., the focus is on examining whole brain functional network connectivity at a voxel level [vDMvLB12]. They visualize the correlation of the functional activity, where fMRI time-signals at each voxel are correlated with every other voxel in the brain to determine functional connectivity. Therefore they propose an application for the interactive visual analysis of this high resolution brain network data, both in a linked matrix representation as well as in its anatomical context based on MRI in a GPU raycasting framework.

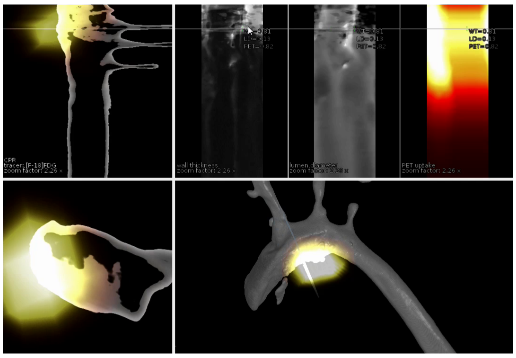

Figure 13: *Multiple linked views allow for comparative visualization of vascular PET/CT in mice [RHR*09].*

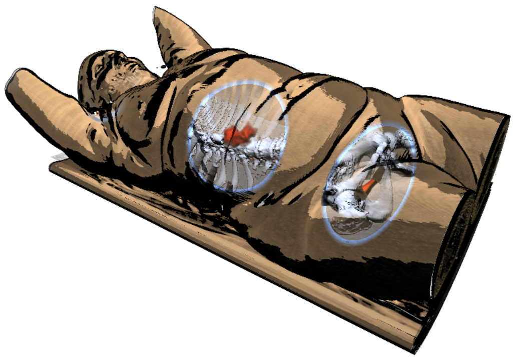

Figure 14: *PET/CT visualization with PET activity presented as a focus area using dynamic cutaways and CT providing anatomical context for accurate localization [LSPV15].*

6.2. Diagnosis

Visualization techniques aimed at improving diagnostic value of multimodal medical data have been primarily developed in the fields of oncology and cardiology.

6.2.1. Oncology

Among the oncological applications, Kim et al. introduced a dual-lookup table specifically designed for use with PET/CT data such that medical experts can set different transfer functions for every volume in a single view [KEF07]. 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. Lawonn et al. recently developed an illustrative technique for focus-and-context rendering of PET/CT data [LSPV15]. In their approach, the functional information from the PET data is used as the focus, while the CT provides anatomical context using dynamic cutaway views (see Figure 14). Viola et al. presented illustrative techniques to visualize liver ultrasound combined with CT information serving as anatomical context [VNØ*08]. They aim to shorten the long learning curve for ultrasound practitioners in training, as well as assist the interpretation of liver examinations.

6.2.2. Cardiology

Due to the severity and frequency of cardiac diseases, this application area is of utmost importance. Essential diagnostic tasks include plaque assessment in the coronary arteries, detailed diagnosis of a heart infarction, e.g. the extent of infarct core, assessment of the heart valves, and assessment of abnormalities, such as congenital heart failures. Morphological information, extracted from CT and MRI, as well as functional information, e.g., wall motion extracted from ultrasound, are important.

An example of a multimodal visualization problem that was considered in visualization research is the use of perfusion data combined with other MR imaging modes of for the diagnosis of the coronary heart disease. Perfusion data of the heart, which indicates the blood perfusion in the myocardium and can be acquired from MRI or SPECT scanners, needs to be combined with anatomical data to provide morphology of the heart muscle and the coronary arteries. Similar to combining PET and CT data, the resolution of the perfusion data is much coarser. Moreover, the acquisition of cardiac perfusion data exhibits gaps, i.e., there are regions in the morphological data where no corresponding perfusion data is available. Simply interpolating the missing information is not a satisfactory solution. Thus, the overlay of both information must convey where slices of the morphological data correspond to morphologic slices and where no such correspondence exists. Termeer et al. proposed a visualization for the diagnosis of Coronary Artery Disease using cardiac perfusion MRI data [TBB*07]. They extend the traditional bull's eye plot to a continuous volumetric version that reveals transmurality of scar tissue and links this to an anatomical view of the heart. Transmurality indicates whether the whole wall of the (left) ventricle is affected by an infarction—information that is essential for prognosis and treatment. Oeltze et al. presented an integrated multi-modal visualization of morphologic and cardiac perfusion data for the analysis of coronary artery disease [OKG*06]. Myocardial perfusion is measured using an MRI scanner and combined with CT Coronary Angiography (CTCA) which depicts the anatomy. In their visualization, colored icons, heightfields and lenses are used to visualize and explore the different parameters measured. Hennemuth et al. developed a comprehensive approach to the analysis of contrast-enhanced cardiac MR images [HSF*08]. They propose a full workflow to align the different datasets, extract surfaces and analyze them in a combined way (see Figure 15). Kirişli et al. combine CT Angiography (CTA) and SPECT myocardial perfusion imaging (MPI) in a visualization aimed at assessing coronary artery disease [KGS*14]. They evaluated the diagnostic value of a software-based image fusion system over conventional side-by-side analysis and found improved diagnostic performance when using their application.

6.3. Treatment Planning

Multimodal medical data visualization applications have been developed mainly in the areas of neurosurgical planning and guidance. Some recent works have focused on radiotherapy treatment planning.

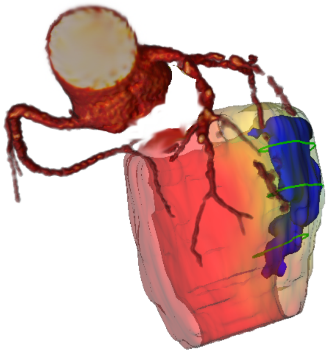

Figure 15: *Combined visualization revealing an infarcted area in dark blue, with hypo-perfused regions represented by contours. The myocardium is shown as a transparent colored surface that encodes the distance to the infarction [HSF*08].*

6.3.1. Neurosurgery

A primary focus in neurosurgical planning applications is identifying surgically relevant structures that are at risk to be damaged, such as the functional areas in the gray matter and the white matter fiber tracts, and understanding how they are related to the lesion that needs to be removed. The surgically relevant structures can be acquired using CT, MRI, fMRI, DTI as well as different MR protocols. Additionally, neurosurgeons need to plan a safe access path to such a lesion, with the least amount of damage to vital functional areas, derived from fMRI, and fiber tracts, derived from DTI. Visualization applications aimed at neurosurgical planning and guidance often focus on brain tumor resection, but also applications aimed at Deep Brain Stimulation (DBS) and Arteriovenous Malformations (AVM) surgery exist.

Jannin et al. fused several modalities and imaging modes from MRI for neurosurgery navigation [JFS*00]. They combined segmented structures and vessels with functional areas from magnetoencephalography (MEG) and fMRI data in a neuronavigation application. As one of the most notable early techniques where surfaces are used for combined information, Stokking et al. presented a visualization method that combines functional input data and a surface extracted from anatomical data [SZV01]. They map fMRI values onto a brain surface using the surface normal and evaluate their technique on both registered (S)PE(C)T/MRI and fMRI/MRI. This technique was an extension from earlier work on combining software-registered SPECT with a surface extracted from an MRI scan, by mapping functional values of the SPECT to the surface of the brain along the normal [SZP*97]. While these techniques were applied to the brain, the idea of using a surface can be transferred to other organs, such as the heart muscle. Blaas et al. proposed an approach that fuses fMRI and DTI data for planning brain tumor resections [BBM*07]. In this fused volume, users can extract fiber bundles that pass through a region around the tumor. These bundles can then be explored by filtering on distance to the tumor, or by selecting a specific functional area using arbitrary convex geometries as selection criteria (see Figure 16).

Rieder et al. employed a combination of fMRI and DTI data for neurosurgical planning [RRRP08]. Their application visualizes pathologies using a distance-based transfer function, as used in the work by Tappenbeck et al. [TPD06], and only shows functional data in close proximity to the lesion. Furthermore, they increase depth perception by including a distance-ring, which visualizes how deep the lesion is situated in the brain from the current viewpoint. For the surgical planning itself, they provide access path visualization and rely on identification of superficial landmarks which can be translated to the per-operative context (see Figure 17). A visualization approach for combined MRI and fMRI brain data was presented by Jainek et al. [JBB*08]. Various rendering styles were used, e.g. ambient occlusion, and illustrative techniques were employed to enhance the visual output. Following up on this work, Born et al. extended it to include DTI tracts, which reveal reconstructed nerve fibers that connect functional areas [BJH*09]. They also enhanced depth and shape perception by applying silhouettes and dithered half-toning (see Figure 18). Janoos et al. proposed a method to visually analyze brain activity from fMRI data, with a special focus on temporal dependencies [JNM*09]. They propose a methodology to analyze the time dimension through volumes-of-interest, of which the selection is guided by a hierarchical clustering algorithm in the wavelet domain. They visualize these volumes-of-interest overlaid onto MRI data of the brain and show the cluster time-series of selected clusters in a separate view.

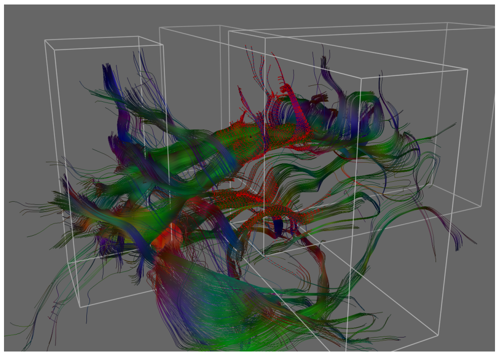

Figure 16: *Fiber bundle filtering using multiple boxes to select bundles with arbitrary logic combinations for selection of fibers to be shown in multimodal visualization [BBM*07].*

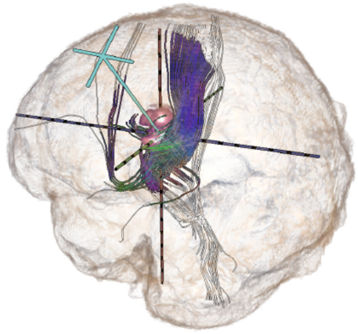

Figure 17: *Path planning application for oncologic neurosurgery based on MRI, fMRI and DTI data [RRRP08].*

Diepenbrock et al. were the winners of the IEEE VIS Visualization contest on multimodal visualization for neurosurgical plan-

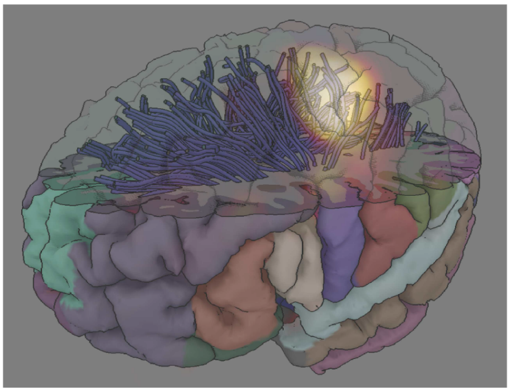

Figure 18: *Visualization of cortical anatomy (MRI), brain activity (fMRI), and nervous pathways (DTI). Springer [BJH*09], ©Springer-Verlag Berlin Heidelberg 2009. With permission of Springer.*

ning [DPL*11]. They provided a 3D view inside the brain featuring an interactive probe containing structures acquired from fMRI, DTI and various additional MR imaging modes. Uncertainty in the fMRI and DTI data is revealed using uncertainty borders and a different rendering style. Furthermore, they provide a cylindrical access path projection representing the distance to the structures at risk along the planned path. A close-up view of the tumor allows a more detailed view on the area that the surgeons are planning to resect.

Serra et al. developed a multimodal pre-operative neurosurgical planning system designed for use with a mirror-based Virtual Reality system workbench [SKG*98]. They render CT, MRI and MRA scans using 3D textures and allow the user to interact with the volumes using a tracked 3D pointer. Neubauer et al. proposed a surgical simulation application for endonasal transsphenoidal pituitary surgery, which is a minimally invasive endoscopic procedure [NMW*04]. They employ CT and MRI to simulate this endoscopic procedure for training and pre-operative planning purposes. They used CT for skull anatomy, MRI for the tumor, optical nerve and pituitary gland, and contrast-enhanced CT and MRI to highlight the internal carotid artery. They provide interactive threshold adjustment to adjust the surface visualizations.

For planning neurosurgical procedures, Beyer et al. presented a seminal paper on multimodal medical visualization that illustrates volumes such as CT, MRI, fMRI, PET and Digital Subtraction Angiography (DSA) [BHWB07]. They propose a skull peeling algorithm that can automatically remove an occluding bony area to reveal the underlying brain. Furthermore, they developed a masked-based approach to show multiple modalities concurrently and a rendering technique to render binary segmented objects with a smoothed appearance when available (see Figure 19). Within masked parts of the volume, different datasets are rendered, each with their own transfer function settings. Joshi et al. presented interaction techniques in a neurosurgical planning application that includes stereotactic navigation [JSV*08]. They visualize MRI, fMRI, DTI and SPECT in a visualization that allows the user to crop volumes using an interactive line widget. Kin et al. developed a neurosurgical planning tool specifically focussing on brainstem malformations in which they fuse MRI, CT, and 3D rotational angiography together [KNS*12]. They mainly rely on combining multiple surface reconstructions of the individual modalities, which demands a significant amount of preprocessing time in manually segmenting important structures.

Bock et al. recently presented a tool for planning and guiding deep brain stimulation (DBS) interventions by fusing multimodal data that includes uncertainty regions from the acquisition process [BLE*13]. In order to guide these procedures, CT and MRI are fused and visualized integrated with results from Microelectrode Recordings (MER), which measure the electric field in the brain intra-operatively. Their tool features a planning, recording and placement phase in which the corresponding steps of the intervention can be performed (see Figure 20).

In the area of neurosurgical planning, Rieder et al. proposed a multimodal visualization of intracerebral pathological tissue [RSHP08]. They use multiple MRI sequences (T1, T1ce (contrast enhanced), T2, FLAIR) as the input for their visualization application and perform clustering to determine pathologic regions. Next, they blend these pathologic regions with the anatomical context information using an automatically calculated transfer function. Furthermore, they propose an automatic cutting tool and brain peeling to reveal hidden structures of interest.

Weiler et al. dealt with neurosurgical planning for treatment of Arteriovenous Malformations (AVMs) in the brain [WRD*11], an example of which can be seen in Figure 21. In these procedures, a precise identification of the arteries and veins is crucial to understand the complex inflow and outflow in these vascular pathologies. They combine several differently weighted MRI scans, such as T1-weighted images with and without contrast agent, arterial time-of-flight (TOF) and MR Venography (MRV), into a single multi-volume visualization to facilitate understanding of the lesion's angio-architecture. To prevent clutter, they propose a focusing technique, based on the distance to a point of interest, that allows the user to attenuate importance using transparency and saturation manipulation for structures outside the region of interest. Navkar et al. developed visualization tools for planning neurosurgical interventions with straight access, such as biopsies, deep brain stimulation and ablation of brain lesions [NTS*10]. For this, they generate various access maps based on vascular structures on the surface of the skin of the patient's head for guidance in selecting a safe entrance point.

The visualization of fiber tracts from DTI scans is an active research area. Enders et al. proposed to visualize white matter tracts reconstructed from DTI with wrapped streamlines [ESM*05]. They generate surfaces that wrap around the convex hull of fiber bundles for a more intuitive representation of tracts and combine this with anatomical information from a T1-weighted scan. Merhof et al. proposed a new visualization technique for white matter tracts using triangle strips and point sprites [MSE*06]. Their novel DTI hybrid rendering approach speeds up the rendering process and can then be combined with DVR of anatomical data to provide additional information for surgical planning. For an overview of these visualizations, we refer to a survey by Tobias Isenberg on illustra-

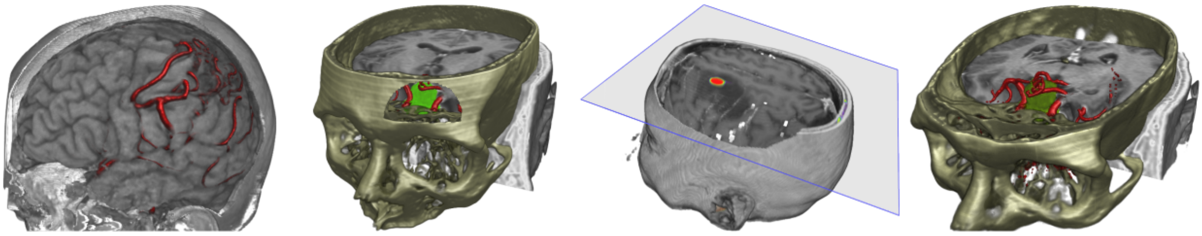

Figure 19: *From left to right: skull peeling, Multi-volume rendering of segmented data (green: tumor (MRI), red: vessels (MRA), brown: skull (CT)), Multi-volume blending (black/white: brain (MR), red: metabolic active part of tumor (PET) and perspective volume rendering for simulating keyhole surgery [BHWB07].*

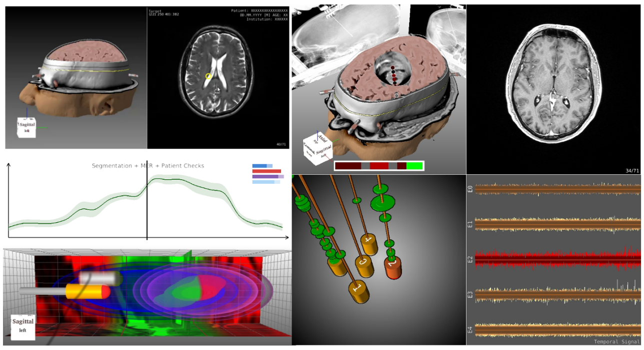

Figure 20: *The DBS planning phase in the top left, recording phase on the right, and the placement phase n the lower left. [BLE*13].*

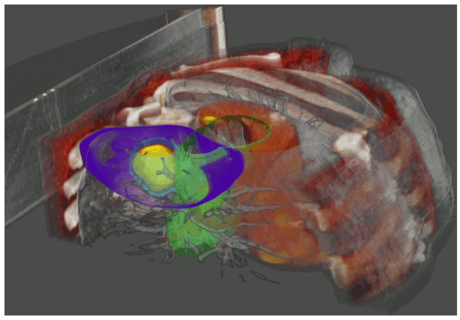

Figure 22: *Fusion visualization of 4D PET/CT with segmentation and dose infomation combined with clipping [SFNB14].*

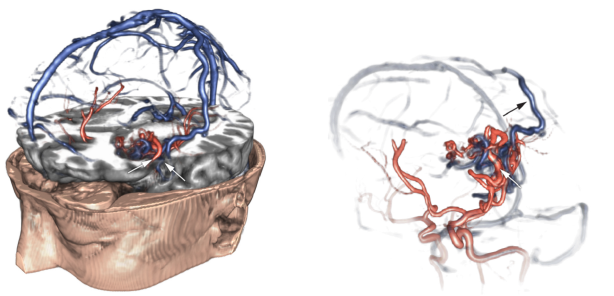

Figure 21: *Visualization of neurovascular anatomy for treatment planning of AVM surgery with context (left) and focused on vascular structures only (right), based on MR-Venographies and T1-weighted MRI [WRD*11].*

tive visualization techniques for diffusion-weighted MRI tractography [Ise15].

6.3.2. Radiotherapy

Radiotherapy planning is generally based on CT, MRI and PET/CT scans. These volumes are used to define target areas for radiation and also organs at risk that should not receive high radiation. Especially for target areas such as chest and upper abdomen, which move strongly due to breathing, the recent possibility to acquire 4D PET/CT data opens the possibility to capture and integrate the movement of tumors into the radiation target volume definition, building the basis for the dose calculation. In contrast to many other diagnostic and treatment planning procedures that occur under severe time pressure, radiation treatment planning is advanced and time-consuming. Thus, advanced multimodal visualization techniques along with appropriate interaction techniques to explore and focus are required to convey this information.

Schlachter et al. [SFNB14] showed that 3D and 4D visualization of images, combined with delineated regions, and the calculated dose (also available as volume dataset), complements the common slice views (see Figure 22). This provides a fast overview over the spatio-temporal configuration of all delineated areas and the related dose distribution resulting finally in a faster quality checks of the radiation plan.

Magnetic Resonance Spectroscopy Imaging (MRSI) data provides metabolic information that quantifies the concentrations of multiple brain metabolites, such as choline and creatine, per voxel. Currently MRSI is only performed in a clinical research context. In the work by Nunes et al., MRSI data is fused with multimodal radiology imaging in an integrated visual analysis system aimed at radiotherapy treatment [NRS*14]. They linked the medical imaging framework MITK and the general purpose data exploration tool ComVis to analyze, relate and visualize MRSI data together with multimodal images.

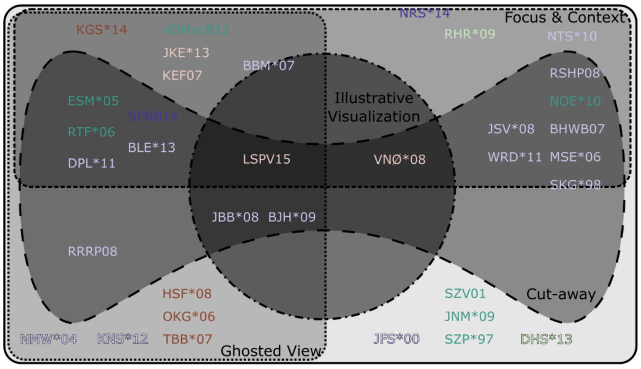

Figure 23: *A Venn-Edwards diagram of the references categorized according to the applied visualization techniques. The color of the references corresponds to the application type colors: vascular research, neuroscience, diagnostic oncology, diagnostic cardiology, neurosurgical treatment planning and radiotherapy planning.*

6.4. Discussion

An overview of all papers according to the application type and application area can be found in Table 1. Here, we stated if any segmentation is needed for processing the date in order to visualize it, which is important to estimate the amount of preprocessing an application requires. We mention which type of evaluation (if any) was performed to verify the utility of the applications. For this, we distinguish between quantitative (Qt) and qualitative (Ql) studies with domain experts (D), non-domain experts (N), or if the performance was tested (P). Furthermore, we mention the image modalities visualized and describe the visualization techniques that are applied for the visualization of the multimodal medical data. Additionally, we have composed a Venn-Edwards diagram of all application papers categorized according to the visualization techniques used and the application type in Figure 23.

From the table we can see that a fair amount of work has been done on neurosurgical treatment planning, as well as work on guidance. Not much work has been done yet on radiotherapy planning, while this field especially uses multimodal data and is in need of suitable visualization methods.

There has only been one visualization application paper combining ultrasound with another modality, although ultrasound has great potential and benefits that can be complimentary to other modalities, such as the ability to provide real-time information. The limited field of view and noise in ultrasound data could be alleviated by combining it with additional modalities, and could be of value also in for instance biopsy guidance. SPECT is involved only in cardiac diagnostic and neuro-science and -surgical planning applications. A popular modality combination for both vascular research and oncologic diagnosis is PET/CT, which is used in all the application papers we listed. Due to the nature of MRI, fMRI and DTI data, these modalities are often involved in both neurosurgical treatment planning and research. For obvious reasons, fMRI is completely neuro-specific, but DTI is also clinically applied so far mainly in neurological contexts.

While the visualization techniques used per application area vary too much to make strong conclusions based on the limited number of samples, it is clear that ghosted views are often employed for research papers intended to support diagnostic purposes (see Figure 23). Furthermore, cutaways are most successfully applied to neurosurgical planning, which makes sense due to the nested anatomical and pathological structures involved and the need for exact access path planning through the skull. There are five papers not employing smart visibility or illustrative techniques, mainly in a research context.

In most works qualitative evaluations with domain experts are presented, but quantitative evaluations as well as more elaborate clinical studies are often still needed. Furthermore, our overview revealed that only a few works applied illustrative visualization techniques [LSPV15, BJH*09, JBB*08, VNØ*08]. However, the application of illustrative visualization technique seems to be promising [TIP05, Law15] and we see a potential to use these techniques for further research on multimodal medical visualization. In multimodal medical visualizations, there are often many overlapping and nested structures, which causes occlusion problems that the abstraction in illustrative techniques can help alleviate. However, the segmentation required often limits the clinical uptake. It might also be helpful to employ novel focus-and-context techniques, by defining one modality as the focus and the other(s) as the context.

7. Conclusion and future challenges

For many techniques described in this survey, it is not clear how they fit in clinical workflows, how much additional and relevant information they provide, and whether this justifies a potentially larger effort, e.g., due to the necessity to segment structures or adjust (2D) transfer functions.

Dealing with medical imaging data involves uncertainty in the form of imaging errors, for instance resulting from noise, field bias, patient motion, or imaging artifacts. Additionally, processing errors may occur due to errors in the segmentation or registration process. When visualizing multi-modal datasets resulting from combining acquisitions of multiple scanners, additional care has to be taken to visualize the uncertainty resulting from these processing errors in the registration process. This is a challenging and so far unsolved research challenge [RPHL14]. When visualizing combined data the reliability of the result needs to be clearly conveyed. One possible extension to medical multimodal visualization solutions would be to incorporate the uncertainty in the visualization, which leads to the question of how to display this uncertainty without distracting the expert. This results in visualization approaches that cope with the problem of showing different types of information combined, e.g., CT, PET and the registration error. This might be an interesting question for further consideration in the future.

Besides the challenge of representing uncertainty in an appropriate way, existing multimodal visualization techniques need to be better evaluated and compared. On the one hand, evaluations that consider effectiveness in specific medical tasks are needed. On the other hand, perception-based evaluations are essential, e.g., to understand how effective certain emphasis techniques actually are, how well users could discriminate values with certain color scales (used as overlays) and how accurate they can locate functional abnormalities. Furthermore, this survey shows that different illustrative visualization techniques can be employed to cope with the challenge of showing different data structures. Also here, more evaluations are needed that analyze if one technique is preferred over other approaches for certain scenarios. This question does not only concern user preferences, but also effectiveness of the visual encoding. For example, one can imagine an evaluation involving a high amount of physicians tasked with identifying and staging a tumor scanned with, e.g., PET/CT. Then, various visualization techniques could be employed and the identification time could be measured.

The existing techniques are spread over many research prototypes and a few commercial solutions. It is thus very difficult to compare them. A framework that integrates at least the most commonly used techniques and flexible parameterization options would be very valuable. A comparison also benefits from benchmark tasks and data. A first step in this direction was the IEEE Vis. 2010 contest on multimodal brain data for neurosurgery planning. A few more such tasks, e.g., in cardiology, would be beneficial.

We have seen a multitude of visualization techniques. Interestingly, these are exclusively techniques that were developed first for displaying single datasets and then adapted and refined, typically for use with two modalities. An open question is whether truly multimodal techniques can be developed, or techniques that go beyond combining two modalities. Multimodal visualization obviously benefits from multiple views. More research is necessary to understand which views are essential, how flexible viewing configurations should be and which synchronization/coordination techniques are needed.

The visualization of single medical datasets may benefit from automatic viewpoint selection, i.e., choosing a good initial view on a dataset as starting point for further exploration [BS05, MNTP07, KBKK07, KBKG08, RCL16]. Measures of visibility, size of objects and viewpoint stability are employed for the viewpoint selection. Automatic viewpoint selection is also promising for multimodal data visualization and requires careful adaptation of measures and application to selected case studies. Showing one viewpoint may be helpful to let the user focus to the specific region, but it might be even more helpful to let the camera follow a path. In this way, the user may get insights into surrounding regions and to acquire a better spatial feeling with regards to, for example, distance to surrounding structures. Also the user interface must be considered to enable a physician to specify what is important as input for an algorithm.

Neurosurgery is among the most advanced user communities in this field, while other disciplines like radiotherapy are not yet exploiting the full potential of visualization techniques available today to support therapy planning and monitoring. We believe there are many applications in neurosurgery because of the relatively easy and reliable registration process, which can be automated, as well as specific application demands. First, CT and MRI contribute a significant amount of relevant and complementary information, while in lung surgery for example most relevant information can be extracted from CT data. Second, neurosurgery has much higher accuracy demands compared to any type of abdominal surgery. A few millimeters in abdominal surgery usually do not matter, while in neurosurgery 2mm may be a lot.

Challenges also appear in case when the input does not consist of surface meshes. For volumetric data, some visualization techniques are more challenging than when applying them to surfaces. Feature lines are mostly defined on surfaces, but some methods can also be applied on volumetric data. Burns et al. [BKR*05] extended the suggestive contours to volumetric data. They also extract the lines as objects, and with this the lines could also be stylized. Kindlmann et al. [KWTM03] visualized ridges and valleys depending on curvature values. With an adapted transfer function, these lines can be highlighted, which is similar to the approach by Lawonn et al. [LSPV15]. In contrast to feature lines, hatching approaches were also applied to volume datasets [DCLK03, CD05, PVW08]. Mostly particles or points are placed in the dataset on a specific isovalue and then the points are traced along the principle curvature directions. Due to the challenging character of feature lines, these techniques were not commonly applied to multimodal data visualization so far. Although its potential was shown to illustrate surfaces [Law15] more work need to be done in this field.

Developing efficient and effective visualization methods with real impact on daily clinical or research routine of potential users requires a highly interdisciplinary effort fusing knowledge from physicians, visual computing specialists, Human Computer Interaction (HCI), imaging and image analysis specialists. This becomes even more important with the increasing availability of new imaging modalities that can be fused with imaging pipelines delivering sets of highly heterogeneous data at different scales. A deep understanding of this data and its application context is necessary to generate useful results; on the other hand, technical skills are getting more and more important, as the amount of data to be processed increases tremendously. As a consequence, multimodal visualization will be even more interdisciplinary than it has been ever before.

8. Summary

We have presented a survey of the current state of the art in medical multimodal visualization. We introduced relevant medical imaging modalities and acquisition techniques as well as the associated visualization challenges. Afterwards, we examined the current clinical workflow for the exploration and analysis of multimodal medical modalities. We summarized the requirements in designing a visualization technique to maximize the insights into relevant details in the depiction of multimodal medical data.

Subsequently, we highlighted the most common visualization techniques that support this visualization problem. These techniques need to incorporate heuristics for assessing the importance of information and emphasis techniques to adapt the importance. For this reason, smart visibility approaches, including focus-and-context techniques, ghosted views and cutaways are highly relevant for visualizing multiple volumes. While these techniques are frequently used in visualization research, they are not part of any commercial solution or available in radiology workstations. This is probably because these techniques require time-consuming preprocessing and are unfamiliar to the physicians.

Rendering multiple volumes is an associated research challenge for which many approaches have been developed. Furthermore, special interaction techniques have been designed for exploring multimodal datasets, which can be as important in providing insight in multimodal data as rendering refinements.

In the main part of our survey, we provided an overview of 35 visualization application papers. We defined the scope of the paper by focusing our survey on multimodal medical visualization applications associated to research, diagnosis and treatment planning or guidance. Within these categories, we found that the main application areas were related to oncology, cardiology, radiotherapy and neurology. We summarized our findings in a table featuring the three aforementioned categories, a further subdivision according to the medical application domain, whether preprocessing is required in terms of segmentation, image modalities used, and visualization techniques employed. We concluded from the table which application areas are active, upcoming or sufficiently researched. While a fair amount of work has been done on neurosurgical planning, radiotherapy planning seems to be a rising field with many opportunities for interesting research, and the same holds for combining ultrasound with other modalities.

Most developed techniques were extended from single modality techniques, while truly multimodal techniques have yet to be developed. With the increase in the amount of data acquired and new modalities being brought into clinical practice, multimodal medical visualization remains a promising research area for future developments.

Acknowledgements

We would like to thank the authors of the cited papers for their contributions in providing figures to illustrate their work. Furthermore, our gratitude goes to the anonymous reviewers for their time and valuable feedback. The contribution of K. BÃijhler has been supported by the project 854174 of the COMET - Competence Centers for Excellent Technologies Program managed by FFG.

References

[ATGP13] ABELLÁN P., TOST D., GRAU S., PUIG A.: Regions-based illustrative visualization of multimodal datasets. *Computerized Medical Imaging and Graphics 37*, 4 (2013), 263–271. Elsevier, Amsterdam. 9

[Ban08] BANKMAN I.: *Handbook of medical image processing and analysis*. 2008. Academic Press, Cambridge, Massachusetts. 6

[BBB*12] BRAMON R., BOADA I., BARDERA A., RODRIGUEZ J., FEIXAS M., PUIG J., SBERT M.: Multimodal data fusion based on mutual information. *Visualization and Computer Graphics, IEEE Transactions on 18*, 9 (2012), 1574–1587. IEEE, New York. 11

[BBBV12] BIRKELAND Å., BRUCKNER S., BRAMBILLA A., VIOLA I.: Illustrative membrane clipping. *Computer Graphics Forum 31*, 3pt1 (2012), 905–914. Wiley Online Library. 12

[BBM*07] BLAAS J., BOTHA C. P., MAJOIE C., NEDERVEEN A., VOS F. M., POST F. H.: Interactive visualization of fused fMRI and DTI for planning brain tumor resections. In *Proc. of SPIE Medical Imaging* (2007), pp. 65091P–65091P. San Diego. 13, 15

[BBPtHR08] BRECHEISEN R., BARTROLI A. V., PLATEL B., TER HAAR ROMENY B. M.: Flexible GPU-based multi-volume ray-casting. In *Proc. of VMV* (2008), pp. 303–312. 11

[BF08] BURNS M., FINKELSTEIN A.: Adaptive cutaways for comprehensible rendering of polygonal scenes. *ACM Transactions on Graphics (TOG) 27*, 5 (2008), 154. 7

[BGCP11] BAER A., GASTEIGER R., CUNNINGHAM D., PREIM B.: Perceptual evaluation of ghosted view techniques for the exploration of vascular structures and embedded flow. *Computer Graphics Forum 30*, 3 (2011), 811–820. Wiley Online Library. 8

[BGM*10] BRUCKNER S., GRÖLLER M. E., MUELLER K., PREIM B., SILVER D.: Illustrative Focus+Context Approaches in Interactive Volume Visualization. In *Scientific Visualization: Advanced Concepts*, Hagen H., (Ed.), vol. 1 of *Dagstuhl Follow-Ups*. Dagstuhl, Germany, 2010, pp. 136–162. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. 9

[BHP15] BEYER J., HADWIGER M., PFISTER H.: State-of-the-art in gpu-based large-scale volume visualization. *Computer Graphics Forum 34*, 8 (2015), 13–37. Wiley Online Library. 11

[BHW*07] BURNS M., HAIDACHER M., WEIN W., VIOLA I., GRÖLLER E.: Feature emphasis and contextual cutaways for multimodal medical visualization. In *Proc. of EuroVis* (2007), vol. 7, pp. 275–282. 9

[BHWB07] BEYER J., HADWIGER M., WOLFSBERGER S., BÜHLER K.: High-quality multimodal volume rendering for preoperative planning of neurosurgical interventions. *Visualization and Computer Graphics, IEEE Transactions on 13*, 6 (2007), 1696–1703. IEEE, New York. 1, 5, 6, 13, 16, 17

[BJH*09] BORN S., JAINEK W., HLAWITSCHKA M., SCHEUERMANN G., TRANTAKIS C., MEIXENSBERGER J., BARTZ D.: Multimodal visualization of DTI and fMRI data using illustrative methods. In *Bildverarbeitung für die Medizin*. Springer, 2009, pp. 6–10. Springer, New York. 13, 15, 16, 18

[BKR*05] BURNS M., KLAWE J., RUSINKIEWICZ S., FINKELSTEIN A., DECARLO D.: Line drawings from volume data. *ACM Transactions on Graphics (Proc. SIGGRAPH) 24*, 3 (Aug. 2005), 512–518. ACM, New York. 19

[BLE*13] BOCK A., LANG N., EVANGELISTA G., LEHRKE R., ROPINSKI T.: Guiding deep brain stimulation interventions by fusing multimodal uncertainty regions. In *Proc. of IEEE Pacific Visualization Symposium (PacificVis)* (2013), pp. 97–104. Sydney. 13, 16, 17

[BNZ*08] BUCK A. K., NEKOLLA S., ZIEGLER S., BEER A., KRAUSE B. J., HERRMANN K., SCHEIDHAUER K., WESTER H.-J., RUMMENY E. J., SCHWAIGER M., DRZEZGA A.: SPECT/CT. *Journal of Nuclear Medicine 49*, 8 (2008), 1305–1319. SNMMI. 4

[BOB*07] BENAMOR M., OLLIVIER L., BRISSE H., MOULIN-ROMSEE G., SERVOIS V., NEUENSCHWANDER S.: PET/CT imaging: what radiologists need to know. *Cancer Imaging 7*, Special issue A (2007), S95. BioMed Central. 5, 12

[Bor86] BORGEFORS G.: Distance transformations in digital images. *Computer vision, graphics, and image processing 34*, 3 (1986), 344–371. Elsevier, Amsterdam. 7

[BRB*13] BRAMON R., RUIZ M., BARDERA A., BOADA I., FEIXAS M., SBERT M.: Information Theory-Based Automatic Multimodal Transfer Function Design. *Biomedical and Health Informatics, IEEE Journal of 17*, 4 (2013), 870–880. IEEE, New York. 9

[BS05] BORDOLOI U. D., SHEN H.-W.: View selection for volume rendering. In *Proc. of IEEE Visualization* (2005), pp. 487–494. Minneapolis. 19

[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. Anaheim. 9

[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. 4

[BW13] BAILEY D. L., WILLOWSON K. P.: An evidence-based review of quantitative spect imaging and potential clinical applications. *Journal of Nuclear Medicine 54*, 1 (2013), 83–89. SNMMI. 4

[CD05] CAI Y., DONG F.: Surface hatching for medical volume data. In *Proc. of International Conference on Computer Graphics, Imaging and Visualization* (2005), pp. 232–238. Beijing. 19

[Che09] CHERRY S. R.: Multimodality imaging: Beyond PET/CT and SPECT/CT. *Seminars in nuclear medicine 39*, 5 (2009), 348–353. Elsevier, Amsterdam. 4

[CM11] CORREA C. D., MA K.-L.: Visibility histograms and visibility-driven transfer functions. *Visualization and Computer Graphics, IEEE Transactions on 17*, 2 (2011), 192–204. IEEE, New York. 12

[Coh06] COHEN M.: *Focus and Context for Volume Visualization*. Phd thesis, University of Leeds, 2006. 9

[CRS*13] CATALANO O. A., ROSEN B. R., SAHANI D. V., HAHN P. F., GUIMARAES A. R., VANGEL M. G., NICOLAI E., SORICELLI A., SALVATORE M.: Clinical impact of PET/MR imaging in patients with cancer undergoing same-day PET/CT: initial experience in 134 patients - a hypothesis-generating exploratory study. *Radiology 269*, 3 (2013), 857–869. RSNA, Chicago. 4

[CS99] CAI W., SAKAS G.: Data Intermixing and Multi-volume Rendering. *Computer Graphics Forum 18*, 3 (1999), 359–368. Wiley Online Library. 10, 11

[DCH88] DREBIN R. A., CARPENTER L., HANRAHAN P.: Volume rendering. In *ACM Siggraph Computer Graphics* (1988), vol. 22, pp. 65–74. ACM, New York. 10

[DCLK03] DONG F., CLAPWORTHY G. J., LIN H., KROKOS M. A.: Nonphotorealistic rendering of medical volume data. *IEEE Comput. Graph. Appl. 23*, 4 (2003), 44–52. IEEE, New York. 19

[DHS*13] DIEPENBROCK S., HERMANN S., SCHÄFERS M., KUHLMANN M., HINRICHS K.: Comparative visualization of tracer uptake in in vivo small animal PET/CT imaging of the carotid arteries. *Computer Graphics Forum 32*, 3pt2 (2013), 241–250. Wiley Online Library. 13

[dMPDSF11] DE MOURA PINTO F., DAL SASSO FREITAS C. M.: Illustrating volume data sets and layered models with importance-aware composition. *The Visual Computer 27*, 10 (2011), 875–886. 9, 10

[dMPF10] DE MOURA PINTO F., FREITAS C.: Importance-aware composition for illustrative volume rendering. In *Proc. of Conference on Graphics, Patterns and Images (SIBGRAPI)* (2010), pp. 134–141. Gramado. 9, 10

[DPL*11] DIEPENBROCK S., PRASSNI J.-S., LINDEMANN F., BOTHE H.-W., ROPINSKI T.: 2010 IEEE visualization contest winner: Interactive planning for brain tumor resections. *IEEE Computer Graphics and Applications*, 5 (2011), 6–13. IEEE, New York. 2, 13, 16

[DSE*12] DRZEZGA A., SOUVATZOGLOU M., EIBER M., BEER A. J., FÜRST S., MARTINEZ-MÖLLER A., NEKOLLA S. G., ZIEGLER S., GANTER C., RUMMENY E. J., M S.: First clinical experience with integrated whole-body PET/MR: comparison to PET/CT in patients with oncologic diagnoses. *Journal of Nuclear Medicine 53*, 6 (2012), 845–855. SNMMI. 4

[DVLLL11] DOBREV P., VAN LONG T., LINSEN L., LINSEN L.: A cluster hierarchy-based volume rendering approach for interactive visual exploration of multi-variate volume data. In *VMV* (2011), pp. 137–144. Berlin. 12

[ESM*05] ENDERS F., SAUBER N., MERHOF D., HASTREITER P., NIMSKY C., STAMMINGER M.: Visualization of white matter tracts with wrapped streamlines. In *Proc. of IEEE Visualization* (2005), pp. 51–58. Minneapolis. 13, 16

[FH09] FUCHS R., HAUSER H.: Visualization of multi-variate scientific data. *Computer Graphics Forum 28*, 6 (2009), 1670–1690. Wiley Online Library. 2, 11

[FPJ*14] FAIRFIELD A. J., PLASENCIA J., JANG Y., THEODORE N., CRAWFORD N. R., FRAKES D. H., MACIEJEWSKI R.: Volume curtaining: a focus+ context effect for multimodal volume visualization. In *Proc. of SPIE Medical Imaging* (2014), pp. 903527–903527. San Diego. 11

[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. Wiley Online Library. 11

[FVW*11] FLUCK O., VETTER C., WEIN W., KAMEN A., PREIM B., WESTERMANN R.: A survey of medical image registration on graphics hardware. *Computer methods and programs in biomedicine 104*, 3 (2011), e45–e57. Elsevier, Amsterdam. 2

[GLH*14] GLASSER S., LAWONN K., HOFFMANN T., SKALEJ M., PREIM B.: Combined visualization of wall thickness and wall shear stress for the evaluation of aneurysms. *Visualization and Computer Graphics, IEEE Transactions on 20*, 12 (2014), 2506–2515. IEEE, New York. 8

[GNBP11] GASTEIGER R., NEUGEBAUER M., BEUING O., PREIM B.: The FLOWLENS: A Focus-and-Context Visualization Approach for Exploration of Blood Flow in Cerebral Aneurysms. *Visualization and Computer Graphics, IEEE Transactions on 17*, 12 (2011), 2183–2192. IEEE, New York. 9

[GNKP10] GASTEIGER R., NEUGEBAUER M., KUBISCH C., PREIM B.: Adapted surface visualization of cerebral aneurysms with embedded blood flow information. In *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2010), pp. 25–32. 7, 8

[GP13] GALANDE A., PATIL R.: The art of medical image fusion: a survey. In *Proc. of Advances in Computing, Communications and Informatics (ICACCI), International Conference on* (2013), pp. 400–405. Mysore. 11

[GRB08] GUPTA S., RAMESH K., BLASCH E.: Mutual information metric evaluation for PET/MRI image fusion. In *Proc. of Aerospace and Electronics Conference* (2008), pp. 305–311. 11

[Gri05] GRIFFETH L. K.: Use of PET/CT scanning in cancer patients: technical and practical considerations. *Proceedings (Baylor University. Medical Center) 18*, 4 (2005), 321. Baylor Health Care System. 5

[HB86] HÖHNE K. H., BERNSTEIN R.: Shading 3d-images from ct using gray-level gradients. *Medical Imaging, IEEE Transactions on 5*, 1 (1986), 45–47. IEEE, New York. 10

[HBG11] HAIDACHER M., BRUCKNER S., GRÖLLER M. E.: Volume analysis using multimodal surface similarity. *Visualization and Computer Graphics, IEEE Transactions on 17*, 12 (2011), 1969–1978. IEEE, New York. 12

[HBKG08] HAIDACHER M., BRUCKNER S., KANITSAR A., GRÖLLER M. E.: Information-based Transfer Functions for Multimodal Visualization. In *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2008), pp. 101–108. Delft. 12

[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. Springer, New York. 11

[HBTR88] HOHNE K.-H., BOMANS M., TIEDE U., RIEMER M.: Display of multiple 3d-objects using the generalized voxel-model. In *Proc. of Medical Imaging II* (1988), pp. 850–854. International Society for Optics and Photonics. 10

[HE98] HASTREITER P., ERTL T.: Integrated registration and visualization of medical image data. In *Proc. of Computer Graphics International* (1998), pp. 78–85. Hannover. 9, 11

[HH12] HICKS R. J., HOFMAN M. S.: Is there still a role for SPECT-CT in oncology in the PET-CT era? *Nature reviews Clinical oncology 9*, 12 (2012), 712–720. Nature Publishing Group. 4

[HMIBG01] HAUSER H., MROZ L., ITALO BISCHI G., GROLLER E.: Two-level volume rendering. *Visualization and Computer Graphics, IEEE Transactions on 7*, 3 (2001), 242–252. IEEE, New York. 8

[HS15] HOTZ I., SCHULTZ T.: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Springer, 2015. Springer, New York. 3

[HSF*08] HENNEMUTH A., SEEGER A., FRIMAN O., MILLER S., KLUMPP B., OELTZE S., PEITGEN H.-O.: A comprehensive approach to the analysis of contrast enhanced cardiac MR images. *Medical Imaging, IEEE Transactions on 27*, 11 (2008), 1592–1610. IEEE, New York. 13, 14, 15

[IFP96] INTERRANTE V., FUCHS H., PIZER S.: Illustrating transparent surfaces with curvature-directed strokes. In *Proc. of IEEE Visualization* (1996), pp. 211–218. San Francisco. 8

[Ise15] ISENBERG T.: A survey of illustrative visualization techniques for diffusion-weighted MRI tractography. In *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*, Hotz I., Schultz T., (Eds.). Springer, 2015, pp. 235–256. Springer, New York. 2, 17

[JBB*08] JAINEK W., BORN S., BARTZ D., STRASSER W., FISCHER J.: Illustrative hybrid visualization and exploration of anatomical and functional brain data. *Computer Graphics Forum 27*, 3 (2008), 855–862. Wiley Online Library. 9, 13, 15, 18

[JD14] JAMES A. P., DASARATHY B. V.: Medical image fusion: a survey of the state of the art. *Information Fusion 19* (2014), 4–19. Elsevier, Amsterdam. 11

[JFS*00] JANNIN P., FLEIG O. J., SEIGNEURET E., GROVA C., MORANDI X., SCARABIN J.-M.: A data fusion environment for multimodal and multi-informational neuronavigation. *Computer Aided Surgery 5*, 1 (2000), 1–10. Wiley Online Library. 13, 15

[JKE*13] JUNG Y., KIM J., EBERL S., FULHAM M., FENG D. D.: Visibility-driven PET-CT visualisation with region of interest (roi) segmentation. *The Visual Computer 29*, 6-8 (2013), 805–815. Springer, New York. 13, 14

[JNM*09] JANOOS F., NOUANESENGSY B., MACHIRAJU R., SHEN H. W., SAMMET S., KNOPP M., MÓROCZ I. Á.: Visual analysis of brain activity from fMRI data. *Computer Graphics Forum 28*, 3 (2009), 903–910. Wiley Online Library. 13, 15

[JSV*08] JOSHI A., SCHEINOST D., VIVES K. P., SPENCER D. D., STAIB L. H., PAPADEMETRIS X.: Novel interaction techniques for neurosurgical planning and stereotactic navigation. *Visualization and Computer Graphics, IEEE Transactions on 14*, 6 (2008), 1587–1594. IEEE, New York. 12, 13, 16

[JWN*08] JUDENHOFER M. S., WEHRL H. F., NEWPORT D. F., CATANA C., SIEGEL S. B., BECKER M., THIELSCHER A., KNEILLING M., LICHY M. P., EICHNER M., KLINGEL K., REISCHL G., WIDMAIER S., RÖKEN M., NUTT R., MACHULLA H., ULUDAG K., CHERRY SR C. C., BJ P.: Simultaneous PET-MRI: a new approach for functional and morphological imaging. *Nature medicine 14*, 4 (2008), 459–465. Nature Publishing Group. 4

[KBKG08] KOHLMANN P., BRUCKNER S., KANITSAR A., GRÖLLER M. E.: Livesync++: Enhancements of an interaction metaphor. In *Proc. of Graphics Interface* (2008), pp. 81–88. Windsor. 19

[KBKK07] KOHLMANN P., BRÜCKNER S., KANITSAR A., KANITSAR A.: Livesync: Deformed viewing spheres for knowledge-based navigation. *Visualization and Computer Graphics, IEEE Transactions on 13*, 6 (2007), 1544–1551. IEEE, New York. 19

[KBvP*15] KÖHLER B., BORN S., VAN PELT R. F., PREIM U., PREIM B.: A survey of cardiac 4D PC-MRI data processing. In *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2015), pp. 139–148. Chester. 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. IEEE, New York. 9, 12, 13, 14

[KGB*09] KAINZ B., GRABNER M., BORNIK A., HAUSWIESNER S., MUEHL J., SCHMALSTIEG D.: Ray casting of multiple volumetric datasets with polyhedral boundaries on manycore gpus. *ACM Transactions on Graphics (TOG) 28*, 5 (2009), 152. ACM, New York. 11

[KGS*14] KIRIŞLI H. A., GUPTA V., SHAHZAD R., AL YOUNIS I., DHARAMPAL A., VAN GEUNS R.-J., SCHOLTE A. J., DE GRAAF M. A., JOEMAI R. M., NIEMAN K., VAN VLIET L., VAN WALSUM T., LELIEVELDT B., NIESSEN W.: Additional diagnostic value of integrated analysis of cardiac CTA and SPECT MPI using the SMARTVis system in patients with suspected coronary artery disease. *Journal of Nuclear Medicine 55*, 1 (2014), 50–57. SNMMI. 13, 14

[KHGR02] KNISS J., HANSEN C., GRENIER M., ROBINSON T.: Volume rendering multivariate data to visualize meteorological simulations: A case study. In *Proc. of the EG Visualization Symposium* (2002), pp. 189–194. Barcelona. 12

[KKL*16] KIM D.-J., KIM B., LEE J., SHIN J., KIM K. W., SHIN Y.-G.: High-quality slab-based intermixing method for fusion rendering of multiple medical objects. *Computer Methods and Programs in Biomedicine 123* (2016), 27–42. Elsevier, Amsterdam. 11

[KNS*12] KIN T., NAKATOMI H., SHOJIMA M., TANAKA M., INO K., MORI H., KUNIMATSU A., OYAMA H., SAITO N.: A new strategic neurosurgical planning tool for brainstem cavernous malformations using interactive computer graphics with multimodal fusion images: Clinical article. *Journal of Neurosurgery 117*, 1 (2012), 78–88. American Association of Neurological Surgeons. 13, 16

[KSF11] KRAUSS B., SCHMIDT B., FLOHR T. G.: Dual source CT. In *Dual energy CT in clinical practice*, Johnson T., Fink C., Schönberg S. O., Reiser M. F., (Eds.). Springer, 2011, pp. 11–20. Springer, New York. 3

[KSW06] KRÜGER J., SCHNEIDER J., WESTERMANN R.: Clearview: An interactive context preserving hotspot visualization technique. *Visualization and Computer Graphics, IEEE Transactions on 12*, 5 (2006), 941–948. IEEE, New York. 9

[KTP10] KUBISCH C., TIETJEN C., PREIM B.: GPU-based Smart Visibility Techniques for Tumor Surgery Planning. *International Journal of Computer Assisted Radiology and Surgery 5* (2010), 667–678. Springer, New York. 7

[KWTM03] KINDLMANN G., WHITAKER R., TASDIZEN T., MOLLER T.: Curvature-based transfer functions for direct volume rendering: Methods and applications. 67–. IEEE, New York. 19

[Law15] LAWONN K.: *Illustrative visualization of medical data sets*. PhD thesis, University of Magdeburg, 2015. 8, 18, 19

[LBMP*01] LE BIHAN D., MANGIN J.-F., POUPON C., CLARK C. A., PAPPATA S., MOLKO N., CHABRIAT H.: Diffusion tensor imaging: concepts and applications. *Journal of magnetic resonance imaging 13*, 4 (2001), 534–546. Elsevier, Amsterdam. 3

[LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3D surface construction algorithm. In *ACM siggraph computer graphics* (1987), vol. 21, ACM, pp. 163–169. Anaheim. 6

[Lev88] LEVOY M.: Display of Surfaces from Volume Data. *IEEE Computer Graphics Applications 8*, 3 (1988), 29–37. IEEE, New York. 10, 12

[LGP14] LAWONN K., GASTEIGER R., PREIM B.: Adaptive surface visualization of vessels with animated blood flow. *Computer Graphics Forum 33*, 8 (2014), 16–27. Wiley Online Library. 7, 8

[LGV*15] LAWONN K., GLASSER S., VILANOVA A., PREIM B., ISENBERG T.: Occlusion-free Blood Flow Animation with Wall Thickness Visualization. *IEEE Transactions on Visualization and Computer Graphics 22 (1)*, 1 (2015), 728–737. IEEE, New York. 7

[LHG*06] LIN G. S., HINES H. H., GRANT G., TAYLOR K., RYALS C.: Automated quantification of myocardial ischemia and wall motion defects by use of cardiac SPECT polar mapping and 4-dimensional surface rendering. *Journal of nuclear medicine technology 34*, 1 (2006), 3–17. SNMMI. 5

[LHV12] LIDAL E. M., HAUSER H., VIOLA I.: Design principles for cutaway visualization of geological models. In *Proc. of the Spring Conference on Computer Graphics* (2012), pp. 47–54. Smolenice. 7

[LKG*16] LJUNG P., KRÜGER J., GROLLER E., HADWIGER M., HANSEN C. D., YNNERMAN A.: State of the art in transfer functions for direct volume rendering. In *Computer Graphics Forum* (2016), vol. 35, pp. 669–691. Wiley Online Library. 12

[LLHY09a] LINDHOLM S., LJUNG P., HADWIGER M., YNNERMAN A.: Fused multi-volume dvr using binary space partitioning. In *Proceedings of the 11th Eurographics / IEEE - VGTC Conference on Visualization* (Chichester, UK, 2009), EuroVis'09, The Eurographs Association; John Wiley; Sons, Ltd., pp. 847–854. 11

[LLHY09b] LINDHOLM S., LJUNG P., HADWIGER M., YNNERMAN A.: Fused multi-volume DVR using binary space partitioning. 847–854. 11

[LLPH15] LAWONN K., LUZ M., PREIM B., HANSEN C.: Illustrative visualization of vascular models for static 2D representations. In *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2015), pp. 399–406. Munich. 9, 10

[LM02] LUM E. B., MA K.-L.: Hardware-accelerated parallel non-photorealistic volume rendering. In *Proc. of Non-photorealistic Animation and Rendering* (2002), pp. 67–74. Annecy. 8

[LSPV15] LAWONN K., SMIT N., PREIM B., VILANOVA A.: Illustrative multi-volume rendering for PET/CT scans. In *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2015), pp. 103–112. Chester. 6, 13, 14, 18, 19

[MBK*10] MARIANI G., BRUSELLI L., KUWERT T., KIM E. E., FLOTATS A., ISRAEL O., DONDI M., WATANABE N.: A review on the clinical uses of SPECT/CT. *European journal of nuclear medicine and molecular imaging 37*, 10 (2010), 1959–1985. Springer, New York. 4

[MFOF02] MANSSOUR I. H., FURUIE S. S., OLABARRIAGA S. D., FREITAS C. M.: Visualizing inner structures in multimodal volume data. In *Proc. of Computer Graphics and Image Processing* (2002), pp. 51–58. Fortaleza. 11

[MNTP07] MÜHLER K., NEUGEBAUER M., TIETJEN C., PREIM B.: Viewpoint selection for intervention planning. In *Proc. of EuroVis* (2007), pp. 267–274. Norrköping. 19

[MSE*06] MERHOF D., SONNTAG M., ENDERS F., NIMSKY C., HASTREITER P., GREINER G.: Hybrid visualization for white matter tracts using triangle strips and point sprites. *Visualization and Computer Graphics, IEEE Transactions on 12*, 5 (2006), 1181–1188. IEEE, New York. 13, 16

[MV98] MAINTZ J. A., VIERGEVER M. A.: A survey of medical image registration. *Medical image analysis 2*, 1 (1998), 1–36. Elsevier, Amsterdam. 2

[NMW*04] NEUBAUER A., MROZ L., WOLFSBERGER S., WEGENKITTL R., FORSTER M. T., BÜHLER K.: Steps-an application for simulation of transsphenoidal endonasal pituitary surgery. In *Proc. of IEEE Visualization* (2004), pp. 513–520. Austin. 13, 16

[NOE*10] NGUYEN T. K., OHLSSON H., EKLUND A., HERNELL F., LJUNG P., FORSELL C., ANDERSSON M., KNUTSSON H., YNNERMAN A.: Concurrent volume visualization of real-time fMRI. In *Proc. of IEEE/EG International Symposium on Volume Graphics* (2010), Eurographics-European Association for Computer Graphics, pp. 53–60. Norrköping. 13

[NRS*14] NUNES M., ROWLAND B., SCHLACHTER M., KEN S., MATKOVIC K., LAPRIE A., BÜHLER K.: An integrated visual analysis system for fusing mr spectroscopy and multi-modal radiology imaging. In *Proc. of Visual Analytics Science and Technology (VAST), IEEE Conference on* (2014), pp. 53–62. Paris. 13, 17

[NTS*10] NAVKAR N. V., TSEKOS N. V., STAFFORD J. R., WEINBERG J. S., DENG Z.: Visualization and planning of neurosurgical interventions with straight access. In *Proc. of Information Processing in Computer-Assisted Interventions* (2010), pp. 1–11. Geneva. 13, 16

[OKG*06] OELTZE S., KUSS A., GROTHUES F., HENNEMUTH A., PREIM B.: Integrated visualization of morphologic and perfusion data for the analysis of coronary artery disease. In *Proc. of EuroVis* (2006), pp. 131–138. Lisbon. 13, 14

[PB13] PREIM B., BOTHA C. P.: *Visual Computing for Medicine: Theory, Algorithms, and Applications*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2013. 2

[PKNS10] PICHLER B. J., KOLB A., NÄGELE T., SCHLEMMER H.-P.: PET/MRI: paving the way for the next generation of clinical multimodality imaging applications. *Journal of Nuclear Medicine 51*, 3 (2010), 333–336. SNMMI. 4

[PMV03] PLUIM J. P., MAINTZ J. A., VIERGEVER M. A.: Mutual-information-based registration of medical images: a survey. *Medical Imaging, IEEE Transactions on 22*, 8 (2003), 986–1004. IEEE, New York. 2

[PVW08] PELT R. V., VILANOVA A., WETERING H. V. D.: GPU-based Particle Systems for Illustrative Volume Rendering. In *Proc. of IEEE/ EG Symposium on Volume and Point-Based Graphics* (2008). Los Angeles. 19

[RBE08] RÖSSLER F., BOTCHEN R. P., ERTL T.: Dynamic shader generation for GPU-based multi-volume ray casting. *Computer Graphics and Applications 28*, 5 (2008), 66–77. IEEE, New York. 11

[RCL16] RANON R., CHRISTIE M., LINO C.: Algorithms and Techniques for Virtual Camera Control. In *EG 2016 - Tutorials* (2016), Sousa A., Bouatouch K., (Eds.), The Eurographics Association. 19

[RHR*09] ROPINSKI T., HERMANN S., REICH R., SCHÄFERS M., HINRICHS K.: Multimodal vessel visualization of mouse aorta PET/CT scans. *Visualization and Computer Graphics, IEEE Transactions on 15*, 6 (2009), 1515–1522. IEEE, New York. 12, 13, 14

[RPHL14] RISTOVSKI G., PREUSSER T., HAHN H. K., LINSEN L.: Uncertainty in medical visualization: Towards a taxonomy. *Computers & Graphics 39* (2014), 60–73. Elsevier, Amsterdam. 18

[RRRP08] RIEDER C., RITTER F., RASPE M., PEITGEN H.-O.: Interactive visualization of multimodal volume data for neurosurgical tumor treatment. *Computer Graphics Forum 27*, 3 (2008), 1055–1062. Wiley Online Library. 13, 15

[RSHP08] RIEDER C., SCHWIER M., HAHN H. K., PEITGEN H.-O.: High-quality multimodal volume visualization of intracerebral pathological tissue. In *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2008), pp. 167–176. The EuroGraphics Association, Geneve. 13, 16

[RSHSG00] REZK-SALAMA C., HASTREITER P., SCHERER J., GREINER G.: Automatic adjustment of transfer functions for 3D volume visualization. In *Proc. of VMV* (2000), pp. 357–364. Stuttgart. 3

[RSPR06] ROSSET A., SPADOLA L., PYSHER L., RATIB O.: Navigating the fifth dimension: Innovative interface for multidimensional multimodality image navigation 1. *Radiographics 26*, 1 (2006), 299–308. RSNA, Chicago. 5

[RSR04] ROSSET A., SPADOLA L., RATIB O.: OsiriX: an open-source software for navigating in multidimensional DICOM images. *Journal of digital imaging 17*, 3 (2004), 205–216. Springer, New York. 5

[RT06] RONG G., TAN T.-S.: Jump flooding in gpu with applications to voronoi diagram and distance transform. In *Proc. of the Symposium on Interactive 3D Graphics and Games* (2006), pp. 109–116. Redwood City. 7

[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. 305–18. Magdeburg. 11, 12, 13

[RVB*09] ROPINSKI T., VIOLA I., BIERMANN M., HAUSER H., HINRICHS K.: Multimodal visualization with interactive closeups. In *Proc. of Theory and Practice of Computer Graphics Conference* (2009), pp. 17–24. Cardiff. 12

[SFNB14] SCHLACHTER M., FECHTER T., NESTLE U., BÜHLER K.: Visualization of 4D-PET/CT, target volumes and dose distribution: Applications in radiotherapy planning. In *Proc. of MICCAI Workshop on Image-Guided Adaptive Radiation Therapy* (2014). 6, 13, 17

[SH08] SEIDENSTICKER P. R., HOFMANN L. K.: *Dual source CT imaging*. 2008. Springer, New York. 3

[SKG*98] SERRA L., KOCKRO R. A., GUAN C. G., HERN N., LEE E. C., LEE Y. H., CHAN C., NOWINSKI W. L.: Multimodal volume-based tumor neurosurgery planning in the virtual workbench. In *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 1998, pp. 1007–1015. Springer, New York. 13, 16

[SKR15] SUNDEN E., KOTTRAVEL S., ROPINSKI T.: Multimodal volume illumination. *Computers & Graphics 50* (2015), 47–60. Elsevier, Amsterdam. 11

[SS11] SCHUBERT N., SCHOLL I.: Comparing GPU-based multi-volume ray casting techniques. *Computer Science-Research and Development 26*, 1-2 (2011), 39–50. Springer, New York. 11

[SS12] SETAREHDAN S. K., SINGH S.: *Advanced algorithmic approaches to medical image segmentation: state-of-the-art applications in cardiology, neurology, mammography and pathology*. 2012. Springer, New York. 6

[STS06] SAUBER N., THEISEL H., SEIDEL H.-P.: Multifield-graphs: An approach to visualizing correlations in multifield scalar data. *Visualization and Computer Graphics, IEEE Transactions on 12*, 5 (2006), 917–924. IEEE, New York. 2

[SzBBKN14] SCHULTE ZU BERGE C., BAUST M., KAPOOR A., NAVAB N.: Predicate-based focus-and-context visualization for 3D ultrasound. *Visualization and Computer Graphics, IEEE Transactions on 20*, 12 (2014), 2379–2387. IEEE, New York. 9, 10

[SZP*97] STOKKING R., ZUIDERVELD K. J., POL H. E. H., VAN RIJK P. P., VIERGEVER M. A.: Normal fusion for three-dimensional integrated visualization of SPECT and magnetic resonance brain images. *The Journal of Nuclear Medicine 38*, 4 (1997), 624. SNMMI. 13, 15

[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. Wiley Online Library. 13, 15

[TBB*99] TOFTS P. S., BRIX G., BUCKLEY D. L., EVELHOCH J. L., HENDERSON E., KNOPP M. V., LARSSON H. B., LEE T.-Y., MAYR N. A., PARKER G. J., RE P., J T., RM W.: Estimating kinetic parameters from dynamic contrast-enhanced t 1-weighted mri of a diffusable tracer: standardized quantities and symbols. *Journal of Magnetic Resonance Imaging 10*, 3 (1999), 223–232. Wiley Online Library. 3

[TBB*07] TERMEER M., BESCÓS J. O., BREEUWER M., VILANOVA A., GERRITSEN F., GRÖLLER M. E.: Covicad: comprehensive visualization of coronary artery disease. *Visualization and Computer Graphics, IEEE Transactions on 13*, 6 (2007), 1632–1639. IEEE, New York. 13, 14

[TC00] TREAVETT S. M. F., CHEN M.: Pen-and-ink rendering in volume visualisation. In *Proc. of IEEE Visualization* (2000), pp. 203–210. Salt Lake City. 8

[TCV*13] TREGLIA G., CASTALDI P., VILLANI M. F., PEROTTI G., FILICE A., AMBROSINI V., CREMONINI N., VERSARI A., FANTI S., GIORDANO A., RUFINI V.: Comparison of different positron emission tomography tracers in patients with recurrent medullary thyroid carcinoma: our experience and a review of the literature. In *Theranostics, Gallium-68, and Other Radionuclides*, Baum R. P., Rösch F., (Eds.). 2013, pp. 385–393. Springer, New York. 4

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

[TG80] TREISMAN A. M., GELADE G.: A feature-integration theory of attention. *Cognitive psychology 12*, 1 (1980), 97–136. Elsevier, Amsterdam. 5, 9

[TIP05] TIETJEN C., ISENBERG T., PREIM B.: Combining silhouettes, surface, and volume rendering for surgery education and planning. In *Proc. of EuroVis* (2005), pp. 303–310. Leeds. 8, 18

[TMS*06] TIETJEN C., MEYER B., SCHLECHTWEG S., PREIM B., HERTEL I., STRAUSS G.: Enhancing slice-based visualizations of medical volume data. In *Proc. of EuroVis* (2006), pp. 123–130. Lisbon. 8, 10

[TPD06] TAPPENBECK A., PREIM B., DICKEN V.: Distance-based transfer function design: Specification methods and applications. In *Proc. of SimVis* (2006), pp. 259–274. Magdeburg. 15

[VBS*13] VIOLA I., BIRKELAND Å., SOLTESZOVA V., HELLJESEN L., HAUSER H., KOTOPOULIS S., NYLUND K., ULVANG D. M., ØYE O. K., HAUSKEN T., GILJA O. H.: High-Quality 3D Visualization of In-Situ Ultrasonography. In *Proc. of Eurographics - Dirk Bartz Prize* (2013), Hege H.-C., Vilanova A., (Eds.). Girona. 3

[VDHP10] VAN DEN HEUVEL M. P., POL H. E. H.: Exploring the brain network: a review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology 20*, 8 (2010), 519–534. Elsevier, Amsterdam. 3

[vDMvLB12] VAN DIXHOORN A. F., MILLES J., VAN LEW B., BOTHA C. P.: Braincove: A tool for voxel-wise fMRI brain connectivity visualization. In *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2012), pp. 99–106. Norrköping. 13

[VG05] VIOLA I., GRÖLLER M. E.: Smart visibility in visualization. In *Proc. of Computational Aesthetics in Graphics, Visualization and Imaging* (2005), pp. 209–216. Girona. 7

[VKG04] VIOLA I., KANITSAR A., GROLLER M. E.: Importance-driven volume rendering. In *Proc. of IEEE Visualization* (2004), IEEE Computer Society, pp. 139–146. 7

[VMN*01] VIERGEVER M. A., MAINTZ J., NIESSEN W., NOORDMANS H., PLUIM J., STOKKING R., VINCKEN K.: Registration, segmentation, and visualization of multimodal brain images. *Computerized Medical Imaging and Graphics 25*, 2 (2001), 147–151. Elsevier, Amsterdam. 2

[VNØ*08] VIOLA I., NYLUND K., ØYE O. K., ULVANG D. M., GILJA O. H., HAUSER H.: Illustrated ultrasound for multimodal data interpretation of liver examinations. In *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2008), pp. 125–133. Eurographics Association, Geneve. 13, 14, 18

[W*07] WEIN W., ET AL.: *Multimodal Integration of Medical Ultrasound for Treatment Planning and Interventions*. PhD thesis, Technische Universität München, 2007. 6

[WEE02] WEISKOPF D., ENGEL K., ERTL T.: Volume clipping via per-fragment operations in texture-based volume visualization. In *Proc. of IEEE Visualization* (2002), pp. 93–100. Boston. 11

[WKC*07] WEIN W., KHAMENE A., CLEVERT D.-A., KUTTER O., NAVAB N.: Simulation and fully automatic multimodal registration of medical ultrasound. In *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2007, pp. 136–143. Brisbane. 6

[WLM02] WILSON B., LUM E. B., MA K.-L.: Interactive Multi-volume Visualization. In *Computational Science*, vol. 2330. 2002, pp. 102–110. Springer, New York. 11

[WMLK89] WALLIS J. W., MILLER T. R., LERNER C. A., KLEERUP E. C.: Three-dimensional display in nuclear medicine. *Medical Imaging, IEEE Transactions on 8*, 4 (1989), 297–230. IEEE, New York. 5

[WRD*11] WEILER F., RIEDER C., DAVID C. A., WALD C., HAHN H. K.: AVM-Explorer: Multi-Volume Visualization of Vascular Structures for Planning of Cerebral AVM Surgery. The EuroGraphics Association, Geneve. 13, 16, 17

[ZMA08] ZAIDI H., MONTANDON M.-L., ALAVI A.: The clinical role of fusion imaging using PET, CT, and MR imaging. *PET clinics 3*, 3 (2008), 275–291. Elsevier, Amsterdam. 5