



DISSERTATION

LiveSync: Smart Linking of 2D and 3D Views in Medical Applications

ausgeführt zum Zwecke der Erlangung des akademischen Grades eines
Doktors der technischen Wissenschaften

unter Anleitung von
Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller
Institut für Computergraphik und Algorithmen
der Technischen Universität Wien

eingereicht an der Technischen Universität Wien,
Fakultät für Informatik,

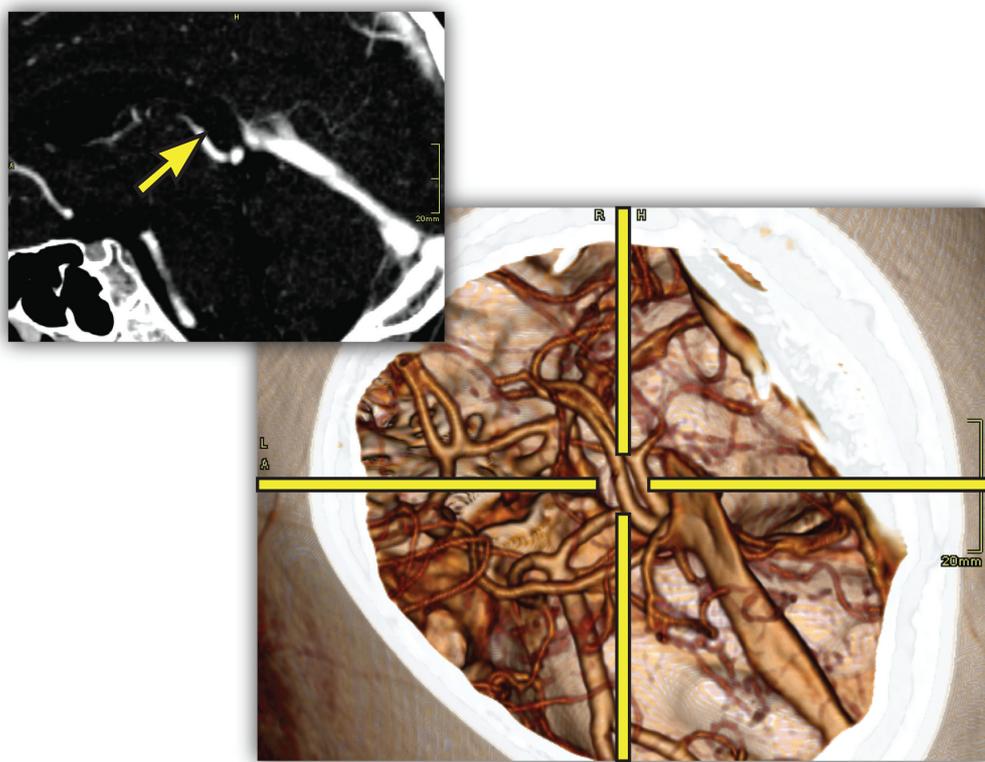
von

Dipl.-Inform. Peter Kohlmann
Matrikelnummer: 0827227
Spengergasse 31/5
1050 Wien, Österreich
geboren am 15.02.1979

Wien, im Dezember 2008

LiveSync: Smart Linking of 2D and 3D Views in Medical Applications

Dissertation



Peter Kohlmann

Institute of Computer Graphics and Algorithms

Vienna University of Technology, Austria

kohlmann@cg.tuwien.ac.at

Abstract

In this thesis two techniques for the smart linking of 2D and 3D views in medical applications are presented. Although real-time interactive 3D volume visualization is available even for very large data sets, it is used quite rarely in the clinical practice. A major obstacle for a better integration in the clinical workflow is the time-consuming process to adjust the parameters to generate diagnostically relevant images. The clinician has to take care of the appropriate viewpoint, zooming, transfer function setup, clipping planes, and other parameters. Because of this, current applications primarily employ 2D views generated through standard techniques such as multi-planar reformatting (MPR).

The LiveSync interaction metaphor is a new concept to synchronize 2D slice views and 3D volumetric views of medical data sets. Through intuitive picking actions on the slice, the users define the anatomical structures they are interested in. The 3D volumetric view is updated automatically with the goal that the users are provided with diagnostically relevant images. To achieve this live synchronization a minimal set of derived information, without the need for segmented data sets or data-specific precomputations, is used. The presented system provides the physician with synchronized views which help to gain deeper insight into the medical data with minimal user interaction.

Contextual picking is a novel method for the interactive identification of contextual interest points within volumetric data by picking on a direct volume rendered image. In clinical diagnostics the points of interest are often located in the center of anatomical structures. In order to derive the volumetric position, which allows a convenient examination of the intended structure, the system automatically extracts contextual meta information from the DICOM (Digital Imaging and Communications in Medicine) images and the setup of the medical workstation. Along a viewing ray for a volumetric picking, the ray profile is analyzed to detect structures which are similar to predefined templates from a knowledge base. It is demonstrated that the obtained position in 3D can be utilized to highlight a structure in 2D slice views, to interactively calculate approximate centerlines of tubular objects, or to place labels at contextually-defined 3D positions.

Kurzfassung

In dieser Dissertation werden zwei Techniken vorgestellt, welche 2D und 3D Ansichten in medizinischen Anwendungen geschickt miteinander verknüpfen. Obwohl interaktive 3D Volumenvisualisierung selbst für sehr große Datensätze verfügbar ist, wird diese recht wenig in der klinischen Praxis verwendet. Der größte Hinderungsgrund für eine bessere Integration in den klinischen Arbeitsablauf ist der hohe Zeitaufwand, um die Parameter für diagnostisch relevante Bilder einzustellen. Der Arzt muss sich um das Einstellen eines geeigneten Blickpunktes, des Zooms, einer Transferfunktion, von Schnittebenen und anderer Parameter kümmern. Deshalb werden in aktuellen Anwendungen hauptsächlich 2D Ansichten verwendet, welche durch Standardverfahren wie Multiplanare Reformation (MPR) erzeugt werden.

Das *LiveSync Interaktionsmetapher* ist ein neuartiges Konzept zur Synchronisierung von 2D Schichtbildern und 3D Volumenansichten auf medizinische Datensätze. Die relevanten anatomischen Strukturen werden vom Benutzer durch intuitives Auswählen auf dem Schichtbild definiert. Die 3D Volumenansicht wird automatisch aktualisiert, um dem Benutzer ein diagnostisch relevantes Bild anzubieten. Um diese direkte Synchronisierung zu erreichen, wird eine minimale Menge abgeleiteter Information verwendet. Hierbei werden keine vorsegmentierten Datensätze oder datenspezifische Vorberechnungen benötigt. Das vorgestellte System liefert dem Arzt synchronisierte Ansichten, welche dabei helfen können mit minimaler Benutzerinteraktion einen besseren Einblick in die medizinischen Daten zu bekommen.

Contextual Picking ist eine neuartige Methode, um relevante Positionen in volumetrischen Daten abhängig von ihrem Kontext interaktiv zu bestimmen. Erreicht wird dies durch das Auswählen eines Punktes in einem Bild, welches mittels direktem Volumenrendering erzeugt wurde. In der klinischen Diagnostik befinden sich die relevanten Positionen häufig im Zentrum anatomischer Strukturen. Um diese 3D Positionen, welche eine komfortable Untersuchung der gewünschten Struktur ermöglichen, abzuleiten, extrahiert das System kontextabhängige Metainformation aus den DICOM (Digital Imaging and Communications in Medicine) Bildern und der Konfiguration der medizinischen Arbeitsstation. Entlang eines Sichtstrahls für eine volumetrische Auswahl wird das Strahlenprofil analysiert, um Strukturen zu ermitteln, welche Ähnlichkeiten zu vordefinierten Vorlagen in einer Wissensdatenbank aufweisen. Es wird demonstriert, dass eine zurückgelieferte 3D Position dazu verwendet werden kann, eine Struktur in 2D Ansichten hervorzuheben. Desweiteren können angenäherte Zentrallinien röhrenförmiger Objekte interaktiv berechnet oder Beschriftungen kontextabhängigen 3D Positionen zugewiesen werden.

Acknowledgments

This work would not have been possible without the help of a lot of people. First of all, I would like to express my gratitude to Meister Eduard Gröller for his excellent supervision. He made the time I spent at the institute fruitful but also very enjoyable. Further, I am very grateful to Armin Kanitsar and Stefan Bruckner. Their ideas and the discussions with them were a great source of inspiration. Moreover, I want to thank all past and present members of the vis-group. They all contributed to the enjoyable working environment and gave great feedback whenever I presented ideas or results to them. Very special thanks to my family for their enduring emotional support throughout the years and my girlfriend Nora for enjoying life together with me.

The work has been funded by AGFA HealthCare in the scope of the `DiagVis` project (<http://www.cg.tuwien.ac.at/research/vis/diagvis>). I want to thank Rainer Wegenkittl, Lukas Mroz, and Matej Mlejnek (AGFA HealthCare) for their collaboration and for providing various CT data sets. Additional data sets are courtesy of OsiriX's DICOM sample image sets website [39].

Contents

1	Introduction	1
1.1	Medical Imaging in Clinical Practice	1
1.1.1	Data Storage	1
1.1.2	Workflow in Digital Radiology	4
1.1.3	Clinical Image Viewing	5
1.1.4	Computer-Aided Detection	11
1.2	Linked Views in Medical Applications	12
1.2.1	Virtual Colonoscopy	13
1.2.2	Slice-Based Visualizations	14
1.2.3	Textual Descriptions and 3D Models	15
1.2.4	Electronic Health Records and 3D Models	16
1.3	Scope of this Thesis	17
2	The LiveSync Interaction Metaphor	19
2.1	Introduction	19
2.2	Related Work	20
2.3	The LiveSync Workflow	22
2.4	Viewing Sphere	24
2.4.1	Sphere Parameterization	24
2.4.2	Sphere Map	25
2.4.3	Sphere Deformation	25
2.5	Viewing-Sphere Manipulators	26
2.5.1	Patient-Orientation Viewing-Sphere	26
2.5.2	Viewpoint-History Viewing-Sphere	27
2.5.3	Local Shape-Estimation Viewing-Sphere	28
2.5.4	Visibility Viewing-Sphere	30
2.6	Viewing-Sphere Operators	32
2.6.1	Weighing of Viewing Spheres	32
2.6.2	Combination of Viewing Spheres	33
2.7	Derived Viewport Parameters	34
2.8	Results and Evaluation	35
2.9	Conclusion	39
3	LiveSync++: Enhancements of an Interaction Metaphor	41
3.1	Introduction	41
3.2	Related Work	43
3.3	LiveSync Workflow	44
3.4	Sphere Parameterization	45

3.4.1	Visibility Calculations	47
3.4.2	Sphere Partitioning	47
3.5	Feature-Driven Transfer Function Tuning	49
3.5.1	Feature Extraction	50
3.5.2	Transfer Function Tuning	51
3.6	Interaction Modes	53
3.6.1	LiveSync Mode	54
3.6.2	LiveSync++ Growing Mode	55
3.7	Feature-Driven Clipping	56
3.7.1	LiveSync Clipping Strategies	56
3.7.2	LiveSync++ Smooth Importance-Driven Clipping	58
3.8	Performance and Qualitative Results	59
3.9	Conclusion	61
4	Contextual Picking of Volumetric Structures	63
4.1	Introduction	63
4.2	Related Work	65
4.3	Contextual Picking Overview	66
4.4	Knowledge Base	68
4.4.1	Ray-Profile Library	68
4.4.2	Contextual Profiles	69
4.5	Initialization	71
4.6	Contextual Picking	73
4.6.1	Profile Matching	74
4.6.2	Contextual Picking Action	76
4.7	Performance and Results	77
4.8	Conclusion	79
5	Summary and Conclusions	83
A	Curriculum Vitae	93

Stay committed to your decisions, but stay flexible in your approach.

Tom Robbins

1

Introduction

Medical workstations are designed to assist radiologists in making a diagnosis. Modern imaging modalities produce a huge amount of images and the viewing software has to be capable of handling the image stacks. Today's computer hardware allows the display of powerful 3D visualizations of the medical data sets at interactive frame rates. Algorithms for the computer-aided detection of certain pathologies are integrated to provide a highly productive viewing environment.

This chapter first gives an overview of medical imaging in a clinical environment with a focus on the visualization techniques. Then, some current efforts to link different views in medical applications are presented.

1.1 Medical Imaging in Clinical Practice

Medical volume data can be acquired by various imaging devices, such as computed tomography (CT), magnetic resonance imaging (MRI), or positron emission tomography (PET). There are several manufacturers of acquisition hardware and the recorded data is displayed on quite heterogeneous output devices. It is important to ensure that the appearance of the data is equivalent on the appropriate devices and that data exchange is possible between the devices [42]. Because of these requirements the following sections first describe the standardized format for medical datasets and the integration of the imaging software into the workflow of a radiology department. Afterwards, some key capabilities and concepts of the clinical-image-viewing software are presented.

1.1.1 Data Storage

In 1983 a working group was established by the American College of Radiology (ACR) and the National Electrical Manufacturers Association (NEMA) with the goal to develop a standard for medical images. Their efforts led to the presentation of the DICOM (Digital Imaging and Communications in Medicine)

standard [53] in 1993. At the same time publicly available software tools were released which supported the DICOM format. DICOM is an open, vendor-independent platform for the exchange of medical images and related data between computers and medical devices. This characteristic is the main reason for the widespread acceptance of the format and its key role in the establishment of Picture Archiving and Communication Systems (PACS). In 2008 the base standard consists of 16 parts and 139 supplements. If, e.g., a CT or MRI examination is performed, the output of the scanner is a series of DICOM files where each of them represents a single slice [53, 42].

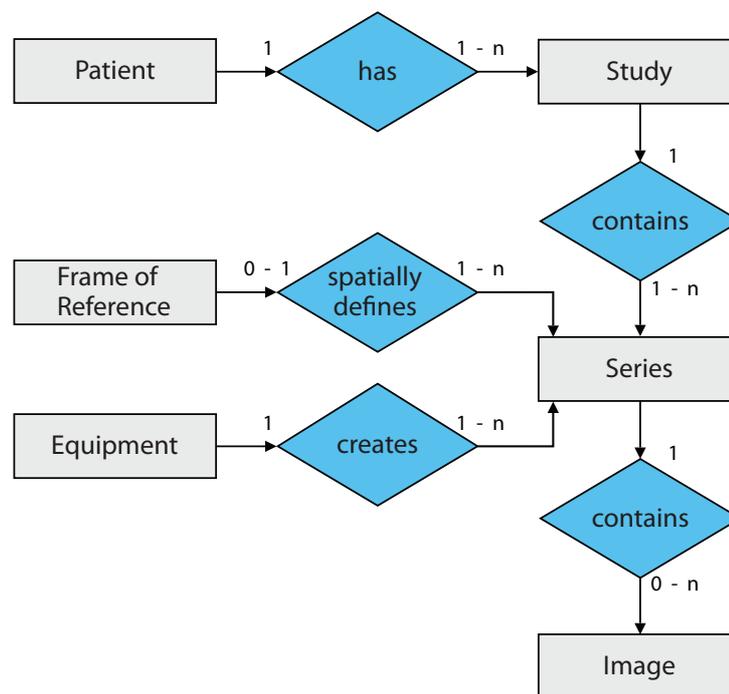


Figure 1.1: DICOM's object-oriented information model to describe a real-world medical scenario (adapted from Gld et al. [17]).

The DICOM objects which are relevant to a certain imaging process are specified by an information object definition (IOD) like, e.g., the CT image IOD. Figure 1.1 depicts the relationships between real-world objects in the DICOM standard. Each patient is subject of one or more medical studies. The patient information entity (IE) and the study IE are modality independent. Attributes of the patient IE are the patient's name, date of birth, and sex, as well as a unique patient ID. Further attributes which are optional are, e.g., weight and height of the patient. The study IE contains data about the examination, like the name of the patient's referring physician, a description of the study, or date and time of the study. A single study may contain one or more series which are created by a particular

1.1 Medical Imaging in Clinical Practice

imaging device. A description of the imaging device is provided by the equipment IE. It specifies, e.g., the manufacturer of the equipment, its software version, or the resolution of the acquisition equipment. The optional frame of reference IE defines the spatial relationship of images within a series. Some important attributes of the series IE are the type of the modality that was used to create the images, the patient's position relative to the equipment, and the part of the body which was subject of the examination. Finally, there is a set of two-dimensional slice images for each series. Each image IE contains various attributes like, e.g., the slice thickness and the pixel spacing, the number of columns and rows in the image, or details about the image compression [53].

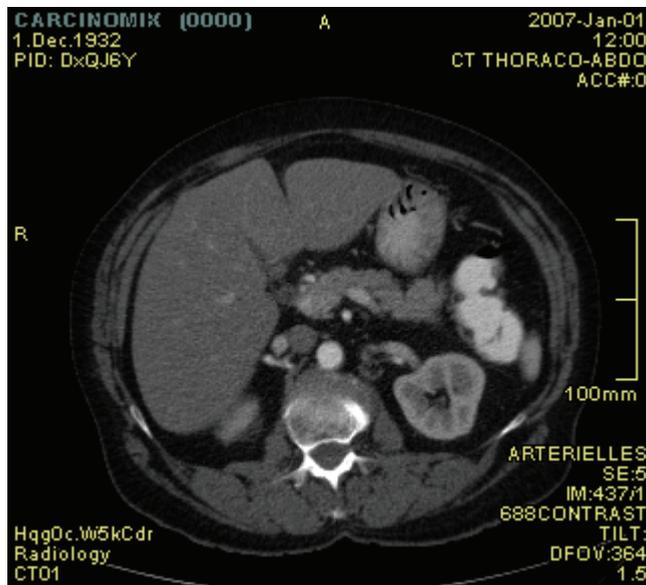


Figure 1.2: Additional information is included in the display of a CT slice.

Some of the described attributes are typically displayed as a legend on the screen of a radiological workstation when the data is examined. Figure 1.2 shows a single slice together with an exemplary legend. For this particular data set the information about the patient and the institution were made anonymous. The capital letters at the top (A) and at the left (R) represent the anatomical directions *anterior* and *right*. Further denotations for anatomical directions are *posterior*, *left*, *head*, and *foot*. The legend at the right edge gives information about the actual extent of the displayed structures. In the upper left corner, patient-specific information is displayed and the upper right shows information about the study. The lower right shows information which describes the series and the image. Finally, in the lower left the institution's name, department name, and station name are displayed.

1.1.2 Workflow in Digital Radiology

There are four key components of a hospital communication system: The Hospital Information System (HIS), the Radiology Information System (RIS), and the Picture Archiving and Communication System. The HIS is designed to manage financial, administrative, and clinical aspects of the hospital. RIS and PACS are the systems which support the workflow in the radiology department. Whereas the RIS is dealing with the information exchange with the HIS and with patient registration and scheduling, the image data is stored by a PACS server and is displayed by the PACS clients. Usually the PACS also supports diagnosis and reporting, as well as archiving the reports. Both, the RIS and the PACS are connected to the image acquisition modalities. Finally, there is the Electronic Health Record (EHR) which allows the storage of all medical data about a patient in electronic form. The EHR consists of the images and all reports which are gathered during the treatment of a patient with the goal that one system integrates all patient information. In reality, typically the health records are stored in different systems of the hospital, such as the HIS and the RIS. A detailed survey about the key capabilities of an EHR system can be found in a report issued by the Institute of Medicine's Committee on Data Standards for Patient Safety [22].

Advantages of Filmless Radiology

In the 1980s the implementation of PACS systems in hospitals started the replacement of the roles of conventional radiological film. More and more radiologists began to view medical images on computer monitors. There are several advantages of a PACS system starting with the fact that once an image is acquired it cannot be misfiled, lost, or stolen. The possibility of distributed viewing is also very beneficial. Whereas a conventional hard copy can only exist in one place at a specific time, a digital image can be viewed simultaneously at any location. The database of a PACS ensures an automatic chronological ordering of the images. Also the grouping of the images into the correct examination is handled by the PACS. As all studies of a patient are immediately available, more often current studies are compared to prior studies and images acquired by different modalities are taken into account. Digital images allow the efficient use of computer tools to annotate, manipulate, or post-process the images. A crucial benefit of a filmless radiology can be seen in economic savings. No film, film processing chemicals, or darkroom technicians are needed anymore. Also the former storage space for the films can be redeployed. The cost aspect caused by salary savings is put into perspective by the need of qualified information technology managers and computer personnel. Overall it is realistic that a PACS is at least cost neutral compared to conventional radiology [49].

Radiology Workflow

The following description of the radiology workflow is based on the work of Preim and Bartz [42] and on an online knowledge base for AGFA HealthCare's current PACS client [1]. Initially, the patient has to be registered at the HIS. Data which is maintained by the HIS include patient demographics, patient visits to the hospital, transfers of the patient, and patient location. Besides the patient registration component the HIS has an order/scheduling component. Typically the HIS sends high-level diagnostic requests to a department information system such as the RIS. Within the radiology department the care for a registered patient is started with an order of the referring medical specialist. Now, the RIS is used to manage the patient's visits to radiology, to schedule diagnostic procedures, and to store the diagnostic results. The referring physician is automatically notified about all carried out examinations. If there is prior image data for the patient available, which is relevant to the current study, it can be provided by the PACS. The RIS transfers a worklist to the corresponding modality. A radiology technician uses this worklist to perform the requested imaging procedure on the patient. The images which are captured by the modality include the patient demographics and are sent to the PACS systems for storage. The PACS is not exclusively a storage device but is also responsible for the transmission, displaying, and printing of medical images. Medical specialists use the display component of the PACS to view the images and to generate a report on studies. This report is finally transferred to the referring physician.

1.1.3 Clinical Image Viewing

The diagnosis starts with the selection of the relevant image series for a particular patient. Different series can be acquired by different modalities. Further they can be part of prior studies.

Slice Viewing

Figure 1.3 shows the simultaneous display of nine successive slices of a head CT data set. Browsing through the slices is one of the most important interactions with the slices. This allows the user to step forward and backward to examine the data set. If there are n successive slices $slice_1$ to $slice_n$ shown simultaneously, one step forward can either lead to the display of $slice_2$ to $slice_{n+1}$ or to the display of $slice_{n+1}$ to $slice_{n+n}$.

The fact that the stack of slices is available in digital form allows reformatting of the data to generate and display three orthogonal views (axial, coronal, and sagittal) as shown in Figure 1.4. A fourth section of the display area is used to show

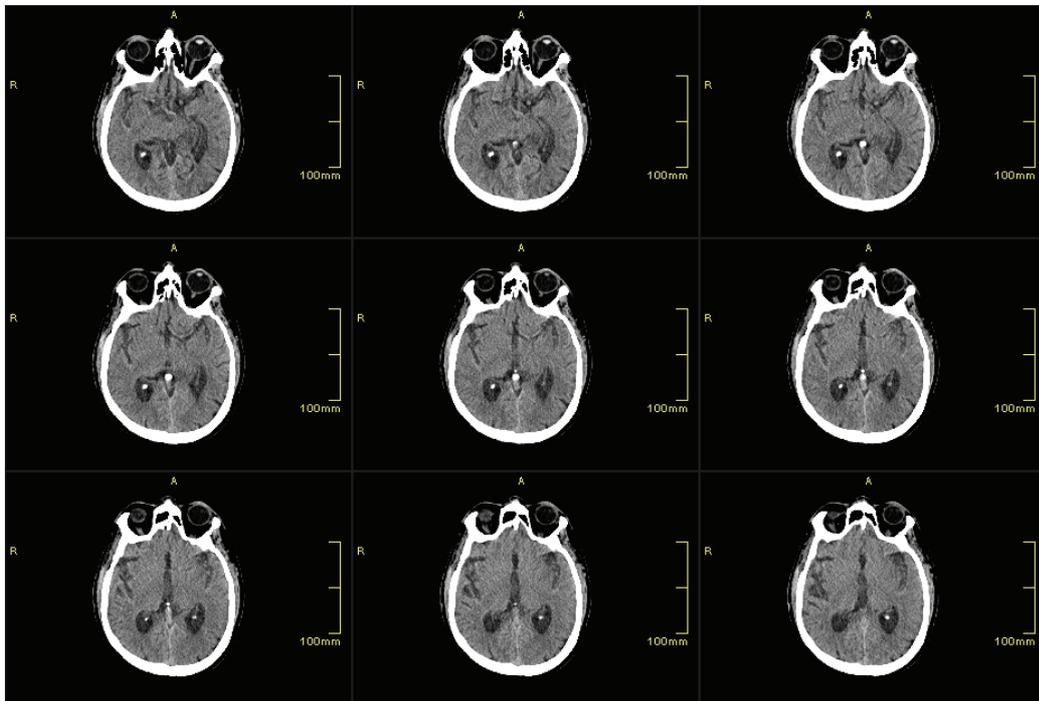


Figure 1.3: Nine successive slice images are displayed simultaneously.

a 3D volume rendering of the data. The exact display layout for a class of images is defined by a digital hanging protocol. A hanging protocol describes, e.g., the order and orientation of images on the screen or the mapping of data values to colors. Further, the synchronization between the different views can be defined in the hanging protocol. Synchronized views can be used, e.g. to link the zooming which means that when the zoom factor is changed in one view, the zooming is automatically adjusted in other views. Whenever a data set is loaded into the workstation the potentially applicable hanging protocols are retrieved from an archive. A scoring function evaluates the DICOM tags of the data set to select the most appropriate hanging protocol automatically. An important feature is the option to customize a hanging protocol to match the viewing preferences of the radiologist.

Windowing

The information content of digital images can be enhanced by image manipulation techniques. A simple form of image manipulation is *windowing* which is used to map the data values to gray values. Imaging modalities like MRI or Ultrasound currently produce images with a resolution of up to 16 Bit. The output

1.1 Medical Imaging in Clinical Practice

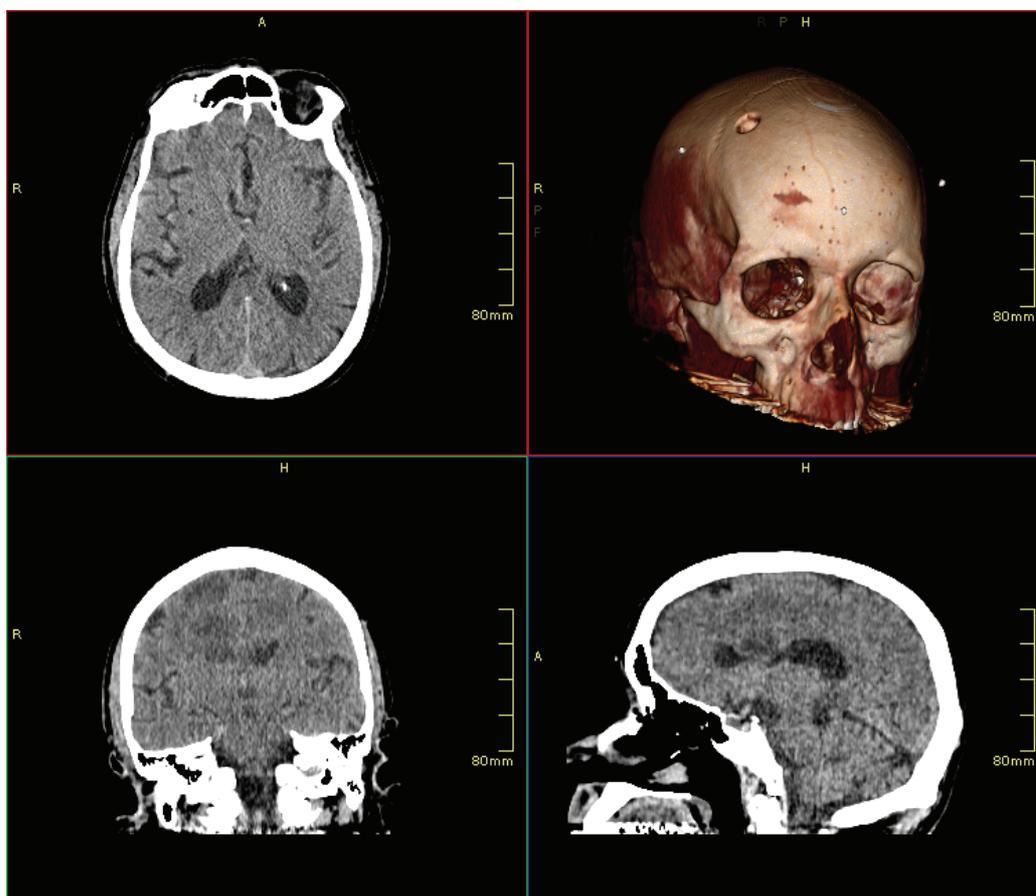


Figure 1.4: Reformatting the image stack allows the display of axial, coronal, and sagittal slices. The upper right section of the display shows a 3D volume rendering of the data.

of CT scanners is transformed to Hounsfield Units (HU) which are in the range from -1000 to +1000. On the Hounsfield scale water has the HU zero, air is represented by the HU -1000, and high density tissue like, e.g., bones can be close to HU +1000. Some modern CT scanners provide a range of 4096 values. It is not possible to display up to 4096 shades of gray on a conventional screen which offers 256 gray levels. In fact, such high granularity is not even necessary because the maximum number of gray values which can be distinguished accurately by the human eye is between 60 and 80. Windowing is a simple mapping technique which requires only little user interaction. The available range of displayable gray values is assigned to a certain window which defines the intensity interval of interest. Figure 1.5 illustrates this mapping. The *window center* defines a central intensity value and the *window width* defines a range of intensity values. Intensity values which are covered by the window width are mapped to the corresponding

gray values. All intensity values outside this interval are mapped either to black or white [24, 23].

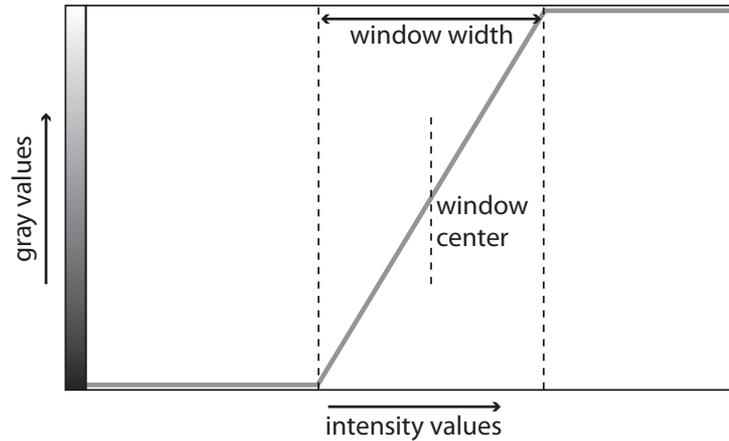


Figure 1.5: Windowing maps intensity values to gray values.

Figure 1.6 demonstrates the impact of windowing on a slice of an abdominal CT data set. In Figure 1.6(a) the density range from -1000 to +1000 HU is mapped to the gray values (window center: 0, window width: 2000). With a reduced window width (window width: 500) as shown in Figure 1.6(b) there is an enhanced contrast resolution, e.g., between muscle and fat tissue. Figure 1.6(c) shows the result of a pulmonary windowing (window center: -600, window width: 1500).

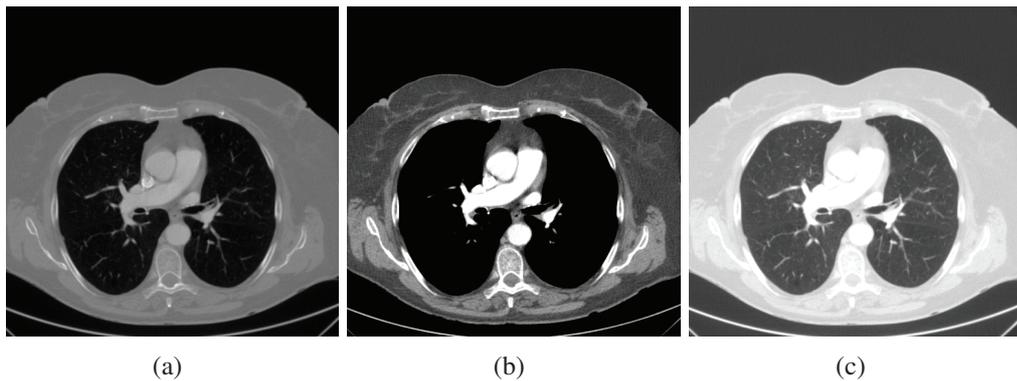


Figure 1.6: Different windowing. (a) Window center: 0, window width: 2000. (b) Window center: 0, window width: 500. (c) Window center: -600, window width: 1500.

3D Visualization

Besides the display of the 2D slices there are several widespread 3D visualization techniques used for diagnostic readings. Most similar to the slices is the 3D visualization technique called multi-planar reformatting (MPR). In contrast to conventional slices which are parallel to the faces of the volume, MPR calculates oblique cross sections through the volume data. An exemplary result of this visualization technique is shown in Figure 1.7(a). MPR enables the generation of slices which are oriented along relevant structures. Often anatomical structures like, e.g., blood vessels, are non-planar and thus difficult to follow in MPR views. Curved planar reformation (CPR) is a visualization technique to overcome this limitation. High-level information like the centerline of a vessel is utilized for the resampling and visualization of the volumetric data. CPR allows the display of entire tubular structures in a single image [25, 26].

For the visualization of surfaces of anatomical structures, often surface shaded display (SSD) is the visualization technique of choice. The key idea is to extract an intermediate surface description to approximate the surface of an object from the volume data. The surface which has to be extracted is defined by a specific threshold value. A polygonal mesh is generated which connects neighboring voxels close to the chosen threshold value. The Marching Cubes [31] algorithm is the most popular technique for surface extraction. SSD is especially suited to extract the surface of bones from CT data as shown in Figure 1.7(b) where lighting is applied for an enhanced depth perception [62].

With volume rendering it is possible to generate an image directly from the volume data without producing an intermediate representation. Volume rendering is a computer-based technique to mimic the physics of light reflection. For each pixel in the output image the reflected amount of light from a virtual light source is calculated as a summation of all contributions along a ray through the volume. The contribution of a voxel is determined by the opacity which is assigned to its intensity value. An opacity transfer function is used to define the mapping from intensity values to opacity values. A second transfer function is used to map the intensity values to gray values or colors which is similar to windowing. Often transfer function presets are provided by diagnostic workstations, e.g., for abdominal, vascular, skull, or lung examinations. Figure 1.7(c) shows the volume rendering of an abdominal CT data set [2].

Maximum intensity projection (MIP) is similar to volume rendering and has proven to be quite useful for the visualization of blood vessels from MR imaging and CT angiography (CTA) where contrast dye is injected into the blood stream prior to the image acquisition. The data values of vascular structures in MRI and CTA data sets are higher than the data values of the surrounding tissue. In contrast to volume rendering not all sampled data values along a ray are taken into

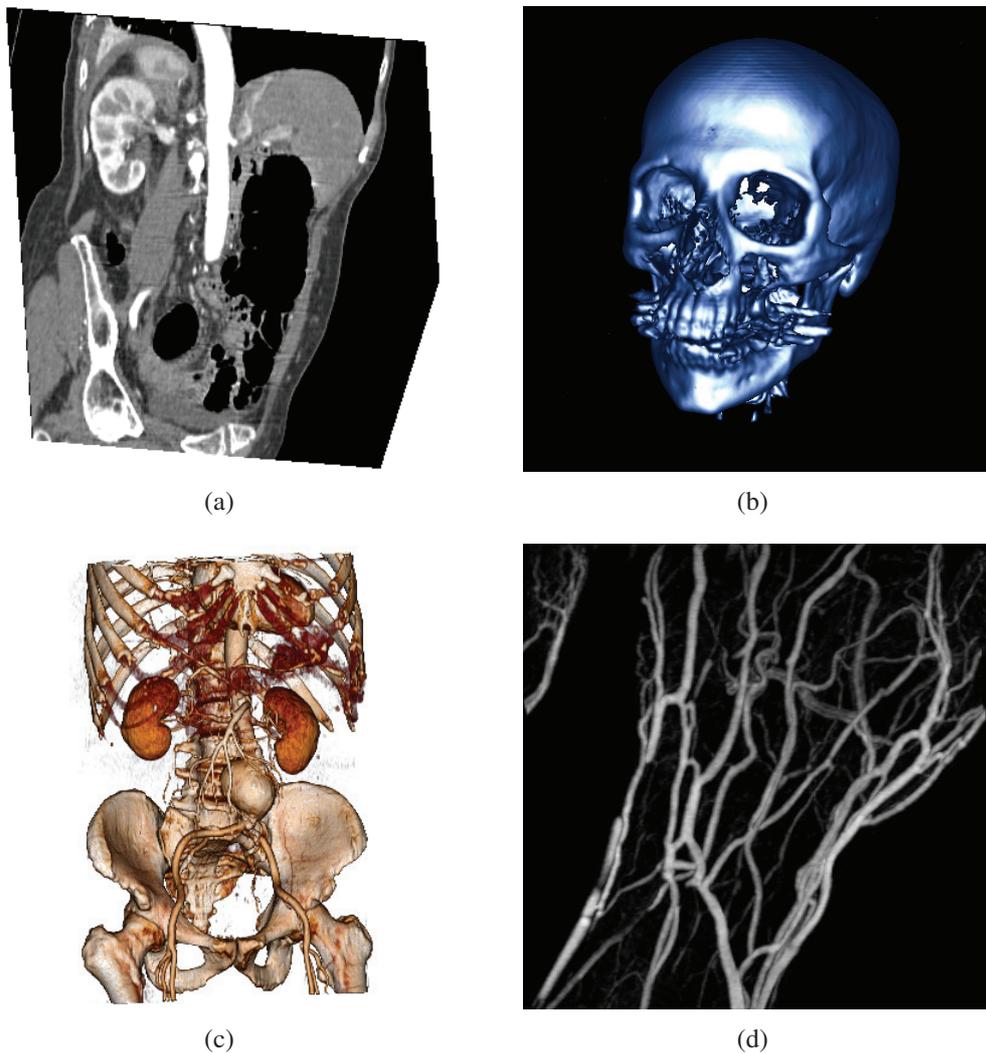


Figure 1.7: Medical volume data displayed with different rendering techniques. (a) Multi-planar reformatting of an abdomen. (b) Surface shaded display of a skull (generated with Klaus Engel's pre-integrated volume renderer [12]). (c) Volume rendering of an abdomen. (d) Maximum intensity projection of a hand (generated with Lukas Mroz's interactive high-quality MIP [37]).

account but only the voxel with the highest intensity contributes to the final image. A drawback of MIP is that the image usually contains no shading information. Thus, valuable depth and occlusion information is missing. An interactive change of the viewpoint helps to ease the interpretation of such images [37]. In CTA, often the bones have to be removed from the data set by segmentation algorithms

1.1 Medical Imaging in Clinical Practice

prior to the MIP calculation, because their intensity is similar or even higher than the intensity of the contrast-enhanced vascular structures. Figure 1.7(d) shows a MIP image of a hand data set. Variations of MIP are closest vessel projection (CVP) [45] and local maximum intensity projection (LMIP) [46]. For CVP the local maximum values which are closest to the viewer are taken as contribution to the image whereas for LMIP the first local maximum values above a certain threshold are taken [62].

Virtual endoscopy has proven its clinical value especially for a non-invasive examination of the colon and the bronchi. In virtual colonoscopy the detection of polyps which can cause cancer is the primary goal. The camera is positioned inside the CT data set and allows a flythrough through the center of an organ to explore the areas of concern. A non-trivial task is the calculation of an optimal path for the flythrough. The path planing can be automated, e.g., by methods which allow automatic centerline extraction as presented by Wan et al. [66]. Advanced visualization techniques facilitate the screening of colon foldings for polyps. Vilanova et al. [21] presented nonlinear colon unfolding with the goal to offer the physician as much information as possible about the inner surface of organs in one image. This method generates images as shown in Figure 1.8(a). Volumetric curved planar reformation presented by Williams et al. [70] is a recent extension to CPR which enhances the examination of the inside of tubular structures. This method combines conventional CPR and volume rendering techniques to produce images as shown in Figure 1.8(b).

1.1.4 Computer-Aided Detection

Features of a modern workstation which gain more and more importance are intelligent tools for the computer-aided detection (CAD) of suspicious regions of interest. In the ideal case this process is automated to a certain degree to speed up the diagnosis. Some main application areas of CAD are the detection of tumor lesions in mammography, polyp detection in virtual colonoscopy, and the detection of lung nodules. Other application areas include the inspection of skin spots for melanoma, the analysis for leaking blood vessels as they are an early indicator of diabetes, the detection of aorta plaques, and the detection and quantification of calcified coronaries [2].

The CAD system supports the decision making process in pinpointing to abnormalities but cannot replace a human observer. The final diagnosis is always made by the radiologist. Figure 1.9 shows the pipeline of a typical CAD system. After the image acquisition is finished a segmentation or detection algorithm is performed on the images to narrow down the search space for anatomical abnormalities. In the feature extraction phase, the general idea is to identify structures which are very similar to objects in a target category and very different to objects

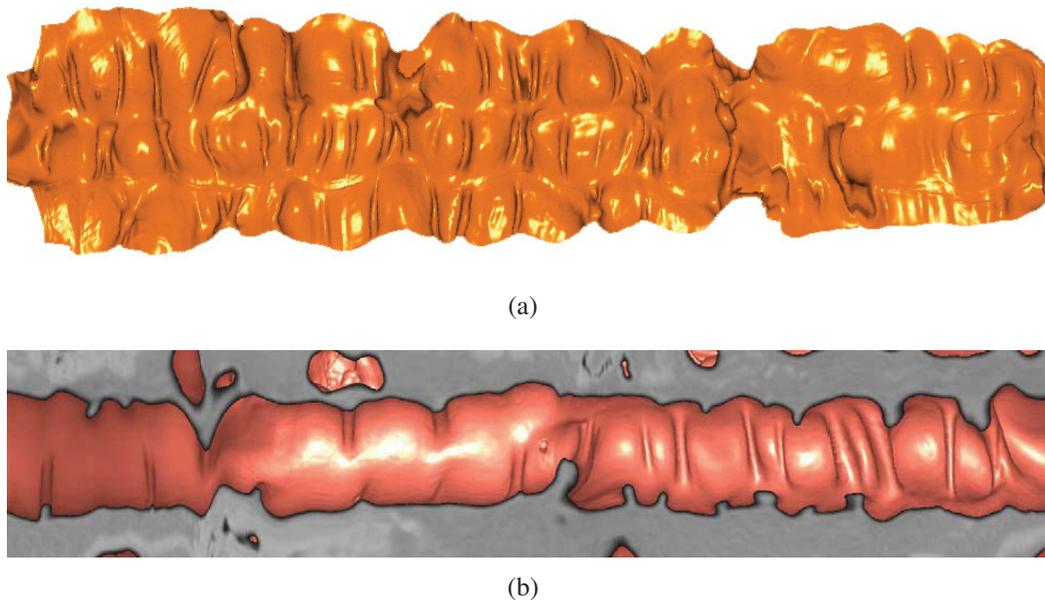


Figure 1.8: Advanced visualization techniques for virtual colonoscopy. (a) Nonlinear colon unfolding [20]. (b) Volumetric curved planar reformation of the colon [68].

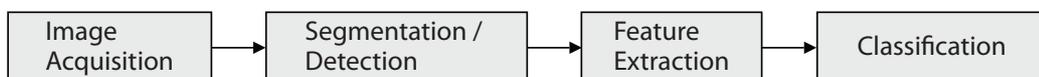


Figure 1.9: General CAD pipeline (adapted from Masala [35]).

in other categories. Usually a supervised classification process is supported by statistical methods and in many commercial systems neural networks are applied for the classification. The performance criteria of a diagnostic system are *sensitivity* and *specificity*. High sensitivity means that the system performs good in detecting actual abnormalities. High specificity means that a minimal number of healthy structures are classified as suspicious structures [35].

1.2 Linked Views in Medical Applications

In medical applications frequently multiple views are provided for the medical volume data. Moreover, textual content like, e.g., the description of anatomical structures or the patient's electronic health record, is relevant for the current examination. This section presents a selection of current approaches to link these different views.

1.2.1 Virtual Colonoscopy

Virtual colonoscopy is an application area where multiple views are displayed simultaneously. Figure 1.10 shows a possible layout for a virtual colonoscopy system. The system provides an overview (top-left) where the current position inside the colon is marked, a 3D volume-rendered flythrough view (top-right), a MPR view (center-left), and two Volumetric CPR views (bottom) facing in opposite directions. All of these views have to be synchronized to ensure that meaningful information can be conveyed. The synchronization is not a trivial task because of the different camera models which are employed for the views. The following camera models are used [68]:

MPR and Flythrough: Camera position and orientation are necessary to display these visualizations. A single 3D point defines the position and the orientation is described by a local coordinate system. The direction of the flythrough (*forward* vector) is the first component of this coordinate system. The other two vectors are an *up* and a *right* vector to specify the rotation of the flythrough. Further, up and right vector define the sampling directions for the MPR image.

Overview: The camera is placed outside the volume to give an overview of the patient's colon. A marker indicates the current flythrough position of the camera. To avoid ambiguous situations in some cases the overview is rotated.

Volumetric CPR: Again the flythrough camera position is required. An additional parameter is an angle of rotation around the position vector. A zoom factor should be provided to define the mapping of centerline points to scanlines.

The synchronization process has to be performed automatically to allow a seamless integration of the multiple views into the virtual colonoscopy system. A precalculated centerline through the data is utilized to provide an automatic flythrough and the synchronized Volumetric CPR views. There are two options to specify the rotation for the Volumetric CPRs. First, the user could freely adjust the rotation which implies considerable additional interaction effort. The second option allows the selection from predefined rotations. This option is less intrusive and thus implemented in the presented system. If the user clicks on a point on the Volumetric CPR, the flythrough camera jumps to the centerline position closest to the clicked point and faces towards it [68].

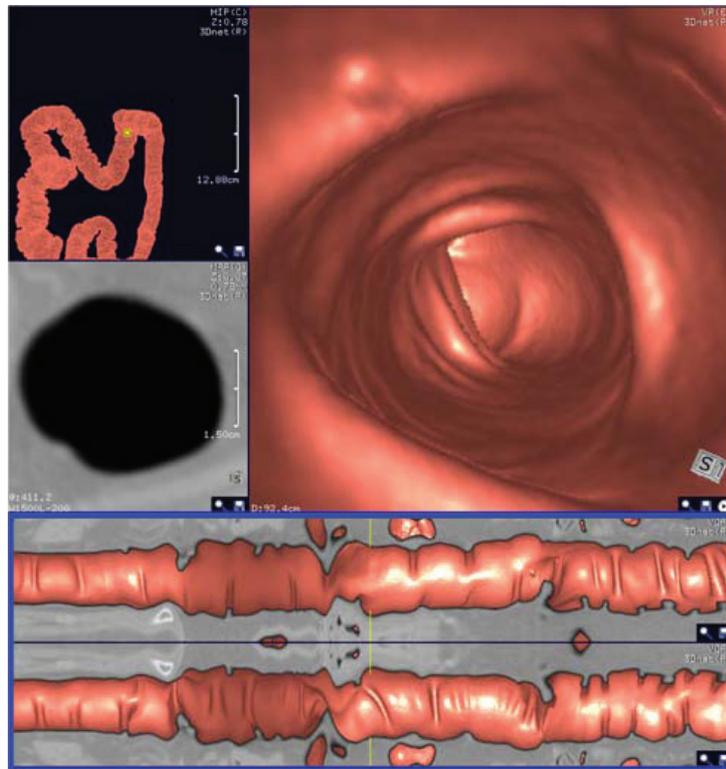


Figure 1.10: Possible layout for a virtual colonoscopy system [69].

1.2.2 Slice-Based Visualizations

Tietjen et al. [54] presented LIFTCHARTs to enhance the slice-based visualization of medical volume data if segmentation information is available. Their system provides a graphical overview for stacks of slices which allows smart navigation to move through the slices. Figure 1.11 illustrates the concept of a LIFTCHART widget. A narrow frame which represents the overall extent of slices in the volume data is displayed next to a slice image. Inside this frame colored bars are arranged for each segmented structure. Each bar represents the extent of the corresponding structure in z-direction. A horizontal line through the frame shows the position of the currently displayed slice within the volume data. LIFTCHARTs provide an overview which indicates in which slices certain structures are located. The horizontal slice indicator can be moved to navigate through the slices and tool tips with specific informations about the underlying anatomical structures are shown. If a particular bar is selected all the slices which are within its range may be displayed in a cine mode which is an animated movement through these slices [54]. The LIFTCHART concept is integrated into the Medical Exploration

1.2 Linked Views in Medical Applications

Toolkit (METK) [36] presented by Tietjen et al. [55] which bundles various concepts for loading, visualizing, and exploring segmented medical data sets.

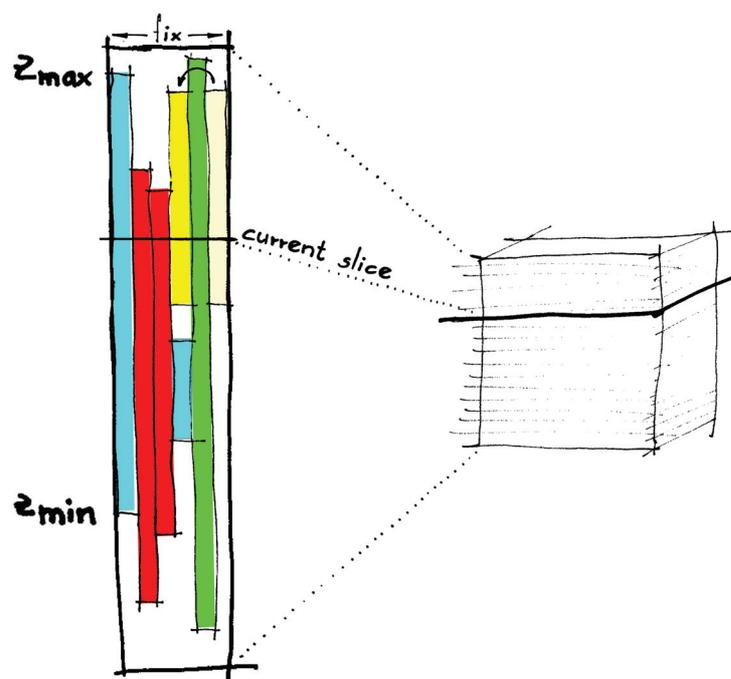


Figure 1.11: Conceptual illustration of the LIFTCHART widget [54].

1.2.3 Textual Descriptions and 3D Models

Götzelmann et al. [16] presented an educational system where learners can interactively explore textual descriptions and 3D visualizations. Their tutoring system utilizes queries which are processed by an information retrieval system to allow a two-directional synchronization of side-by-side views of textual descriptions, e.g., from an anatomical textbook and annotated 3D models. The *text* \mapsto *3D model* synchronization allows medical students to select segments in a textual description for which they want to get additional background information. Figure 1.12 shows how the selection of a text segment (left) provides an annotated 3D model in the linked view (right). The system suggests further appropriate views on other 3D models in small overview windows (center). A colored sphere (top-center) indicates the quality of the current view on the 3D model with respect to the selected text segment. The inverse *3D model* \mapsto *text* synchronization facilitates the interactive exploration of the 3D anatomical model. For the current view on

the 3D model a *visual query* is constructed to determine appropriate text segments which describe the displayed structures [16].

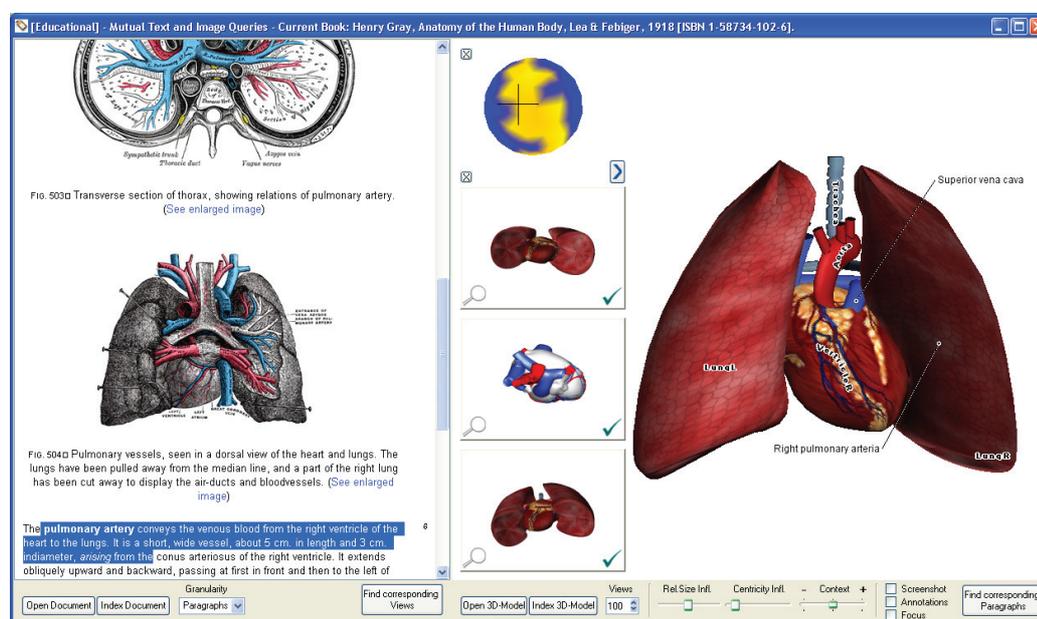


Figure 1.12: Linking of textual descriptions and 3D models. The selection of a text segment leads to the display of an appropriate 3D model [16].

1.2.4 Electronic Health Records and 3D Models

A research team at IBM currently develops a system called the Anatomic and Symbolic Mapper Engine (ASME) [10] to link the patients' electronic health records with 3D anatomical models of the human body. A typical use-case scenario is that the doctor clicks on a particular part of the body in the 3D model. This triggers a search process in the EHR which corresponds to a patient. The retrieved information which corresponds to the selected body part is then displayed. This includes all relevant text entries, lab results, and data from imaging modalities like CT or MRI. The doctor has several possibilities to influence the search parameters, like changing the zoom factor of the model view or by setting a time frame for the results. To further assist the doctor, the anatomical model is linked to 300.000 medical terms. These terms are defined by a hierarchical classification system called Systematized Nomenclature of Medicine (SNOMED). Elisseeff who leads the research team describes the system as *Google Earth* for the human body. Besides its clinical usage the ASME has high potential as a communication tool. It can be used for the communication with patients or other health care profession-

1.3 Scope of this Thesis

als. Further, the ASME might help to improve the teaching of anatomical concepts in medical education [10].

Figure 1.13 shows the search for pain in selected parts of the body. All corresponding medical entries are extracted from the EHR of the patient and the relevant body parts are marked in the anatomical model. A selection of a body part in the model view highlights the corresponding text entries.

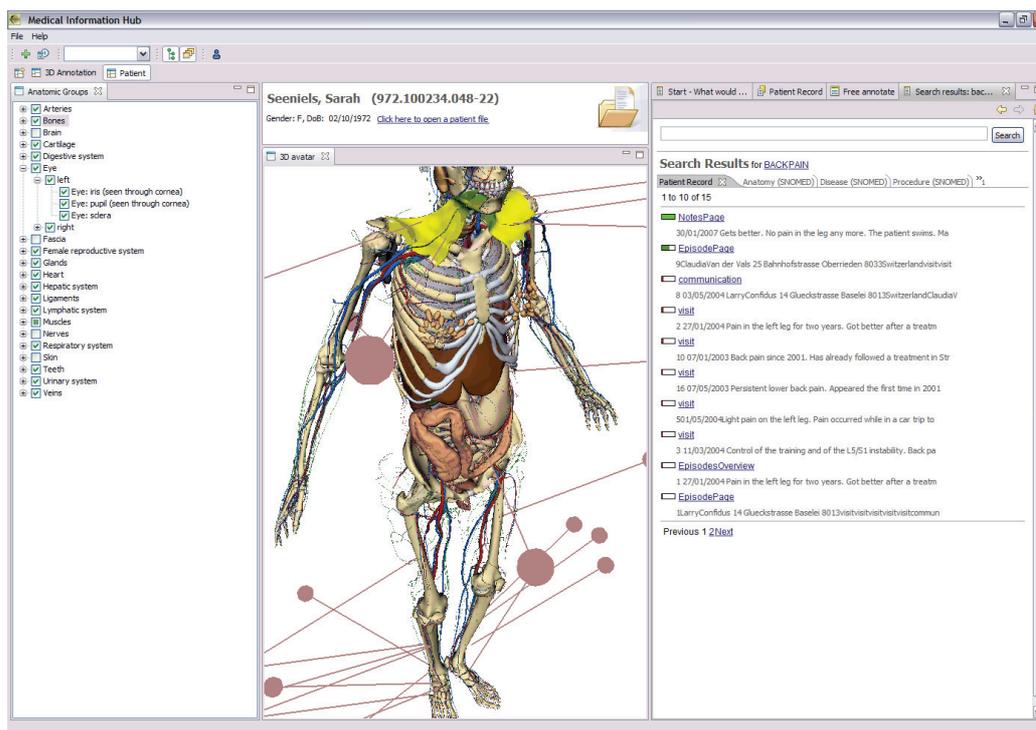


Figure 1.13: The Anatomic and Symbolic Mapper Engine (courtesy of IBM Zurich Research Laboratory).

1.3 Scope of this Thesis

In this work two approaches are presented for a better integration of 3D visualizations in the clinical workflow. The key is to provide methods which require very little user interaction to lead to results which improve the efficiency in diagnosis.

Chapter 2 introduces the LiveSync interaction metaphor. This concept targets the problem that it is quite time consuming to set up the parameters for a diagnostically relevant 3D volume visualization. Due to the necessary effort, most often primarily 2D views such as MPRs are employed in clinical readings. Because 3D renditions can provide additional useful information, LiveSync allows

a seamless integration of 3D visualizations. The presented approach attempts to automatically generate a meaningful 3D view for a structure in a 2D slice image. A simple picking interaction on the slice adjusts the parameters to set up the 3D view automatically. The live synchronization is achieved by utilizing a minimal set of derived information. Considered components are the picked point, slice-view zoom-factor, patient orientation, viewpoint history, local object shape and visibility.

Chapter 3 presents substantial enhancements to the LiveSync interaction metaphor to improve its usability. First, an efficient parametrization for the derived parameters is presented, which allows hierarchical refinement of the search space for good views. Second, the extraction of the feature of interest is performed in a way, which is adapting to the volumetric extent of the feature. The properties of the extracted features are utilized to adjust a predefined transfer function in a feature-enhancing manner. Third, a new interaction mode is presented, which allows the integration of more knowledge about the user-intended visualization, without increasing the interaction effort. Finally, a new clipping technique is integrated.

Chapter 4 addresses the identification of contextual interest points within volumetric data. A novel method allows an interactive identification of these points by picking on a direct volume rendered image. Along a viewing ray for a volumetric picking, the ray profile is analyzed to detect structures which are similar to predefined templates from a knowledge base. The technique can be utilized to highlight a structure in 2D slice views, to interactively calculate approximate centerlines of tubular objects, or to place labels at contextually-defined volumetric positions. Finally, Chapter 5 summarizes the research and the achieved results.

*Things should be made as simple
as possible, but not any simpler.*

Albert Einstein

2

The LiveSync Interaction Metaphor

This chapter is based on the following publication:

P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. LiveSync: Deformed viewing spheres for knowledge-based navigation. In *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1544-1551, 2007.

2.1 Introduction

Modern modalities for medical imaging (e.g., computed tomography) provide large quantities of data at an unprecedented resolution. Presenting this enormous amount of information is a challenging task for today's radiology workstations. Volumetric rendering is the current method of choice for providing a good survey of the data. Combining the information provided by two-dimensional cross sections and three-dimensional visualization can improve the diagnosis process. Linking the different representations of the data has the potential benefit to provide significant enhancements in efficiency. Usually the volumetric display acts as an overview display in this context. The cross-sectional images contain diagnostically relevant information.

Pinpointing a pathological area in the volumetric display selects the corresponding cross-sectional images to be displayed in the two-dimensional display area. From a technical point of view this process is relatively easy to implement. The 3D position of the interesting point can be deduced from the given viewport specification (i.e., transfer function and viewing direction). Very important to note is the reduced degree of freedom in highlighting the position on the corresponding cross-sectional image. The reverse operation is, however, not that straightforward. Picking a 2D position on a cross-sectional slice should result in an expressive unobstructed 3D view. Even though the interesting position is well defined by select-

ing a point in the cross-sectional image, the appropriate highlighting of the area of interest in 3D is challenging. The general motivation for emphasizing a structure selected in 2D in its three-dimensional setting is to get the contextual information. A short example illustrates the situation: A frequently occurring request during reading cross-sectional images of computed tomography angiography is to determine to which anatomical structure a specific partially visible vessel belongs. In this case a volumetric rendering of the depicted vessel and its spatial vicinity would be desired. For optimal results the selected structure should be visible to a large extent and must not be occluded by structures of lower importance.

In the following, a novel concept for interactive viewpoint estimation based on a limited set of input parameters, preserving the generality of the approach, is presented. The only necessary user interaction to derive all the input parameters is given by the picking of a position on a slice. A viewing sphere surrounds the center of a scanned data set and describes all possible camera positions with respect to this object. The input parameters are utilized to encode viewpoint quality in deformed viewing spheres whenever a picking action is performed. After combining the deformed spheres for the different parameters, the estimated quality for all possible viewpoints on the picked structure of interest can be determined from the resulting sphere.

This chapter is structured as follows: Section 2.2 provides an overview of the relevant previous work. In Section 2.3, the workflow and the input parameters are described. Section 2.4 introduces the concept of the viewing sphere. In Section 2.5, it is shown how the input parameters are utilized for the deformation of the viewing sphere. The weighing of the spheres for the different parameters and the combination operators are described in Section 2.6. Section 2.7 explains how the viewing setup can be arranged with the derived viewport parameters. In Section 2.8, LiveSync-generated images are presented for different data sets and the results of an informal evaluation are summarized. Finally, Section 2.9 concludes the chapter.

2.2 Related Work

Viewpoint selection is a well investigated research area for polygonal scenes but relatively few research has been done in the scope of volumetric data. Moreover, the combination of optimal viewpoint estimation and synchronized views has received little attention within the community.

Fleishman et al. [13] presented an approach for an automatic placement of the camera for image-based models with known geometry. A quality measure is applied for the visibility and the occlusion of surfaces. Methods like canonical views are investigated by Blanz et al. [5] for aesthetic aspects of a viewpoint. In

2.2 Related Work

their experimental setup users assign goodness ratings to viewpoints for three-dimensional object models. Based on the feedback a set of criteria for good viewpoints is defined. To determine the viewpoint quality for virtual scenes Sbert et al. [48] applied a measure based on the Kullback-Leibler distance of the projected area of the polygons in the scene. The mesh saliency approach introduced by Lee et al. [29] measures a regional importance for meshes. Besides for mesh simplification this can be employed for viewpoint selection as well. Vázquez et al. [60, 61] worked on the problem that in computer graphics there is no consensus about what defines a good view. Viewpoint entropy based on information theory is introduced to compute good viewing positions automatically. Polonsky et al. [41] aimed for the computation of the best view of an object. They define a set of view descriptors to measure the viewpoint quality. Mühler et al. [38] presented an approach for viewpoint selection in medical surface visualizations. Their work aims at the generation of animations for collaborative intervention planning and surgical education.

Inspired by the research for polygonal data there is some recent work on viewpoint selection for volumetric data. Bordoloi and Shen [6] presented an entropy-based approach to determine a minimal set of representative views for a given scene. The data distribution, the transfer function, and the visibility of voxels are taken into account for their viewpoint selection process. A feature-driven approach to select a good viewpoint is proposed by Takahashi et al. [50]. They identified feature components in the volume for the detection of locally optimal viewpoints. These viewpoints are utilized to extract an optimal global viewpoint. Viola et al. [63] introduced an importance-driven approach to focus on structures within volumetric data. The focus object is defined by the user and their system automatically selects a characteristic viewpoint which provides an expressive view on the object of interest. A framework which facilitates viewpoint selection for angiographic volumes is presented by Chan et al. [9]. View descriptors for visibility, coverage, and self-occlusion of the vessels are considered to determine a globally optimal view. This view is selected by a search process in a solution space for the viewpoints.

Besides techniques for viewpoint selection there are numerous approaches to define a region of interest (ROI) in volumetric data. In the scope of volumes this region is also called volume of interest (VOI). Tory and Swindells [56] presented ExoVis for detail and context direct volume rendering. The VOI can be defined by placing a box within the volume. A translation extracts this part from the volume and this 3D cutout can be displayed with different rendering styles or transfer functions. Owada et al. [40] presented volume catcher as a technique to specify a ROI within unsegmented volume data. The user defines this region by drawing a 2D stroke along the contour of the interesting structure and their system performs a constrained segmentation based on statistical region merging.

Zhou et al. [73] proposed focal region-guided feature-based volume rendering to emphasize the VOI. In their approach a geometric shape like a sphere is used to divide the volume into a focal and a context region.

Regarding tissue classification interesting research has been done by Sato et al. [47]. They have taken 3D local intensity structures into account to identify local features like edges, sheets, lines, and blobs which typically correspond to types of tissues in medical volume data. Their local structure filters use gradient vectors along with the Hessian matrix of the volume intensity combined with Gaussian blurring.

2.3 The LiveSync Workflow

The overall goal of this work is to offer the physician an optimal setup of the viewing parameters for the volumetric view with the least possible effort. If slice views and the volumetric view are not linked the navigation has to be done separately. To enable a 2D/3D synchronization, the functionality of LiveSync can be activated on demand by pressing a hot key while pointing with the mouse on the structure of interest on the slice. Based on this picking process, knowledge-based techniques are applied to estimate good viewpoints for the volumetric view, to calculate an appropriate placement of a view-aligned clipping plane, and to adjust the zoom factor. Depending on the user's preferences, the system allows a smoothly animated rotation or an instant switch between two successive viewpoints. In the case the user is not entirely satisfied with a provided view, it can be refined by manually changing the viewpoint, replacing the clipping-plane, or adjusting the proposed zooming to get a better view of the ROI. If LiveSync is not activated the navigation with the slices is done in a traditional manner and does not lead to an update of the volumetric view. The following factors are considered to achieve the live synchronization:

Picked point: The volumetric position of the depicted structure is determined by the position which the user has picked on a slice.

Slice view zoom: The zoom of the slice view serves as an indicator for the size of the interesting anatomical structure. To set up all viewport parameters automatically this zoom factor is considered to adjust the zoom of the volumetric view.

Patient orientation: Scanned medical data contain information about the patient's position and orientation. Taking into account knowledge about the performed procedure, a rough estimate of the preferred viewing directions is possible.

2.3 The LiveSync Workflow

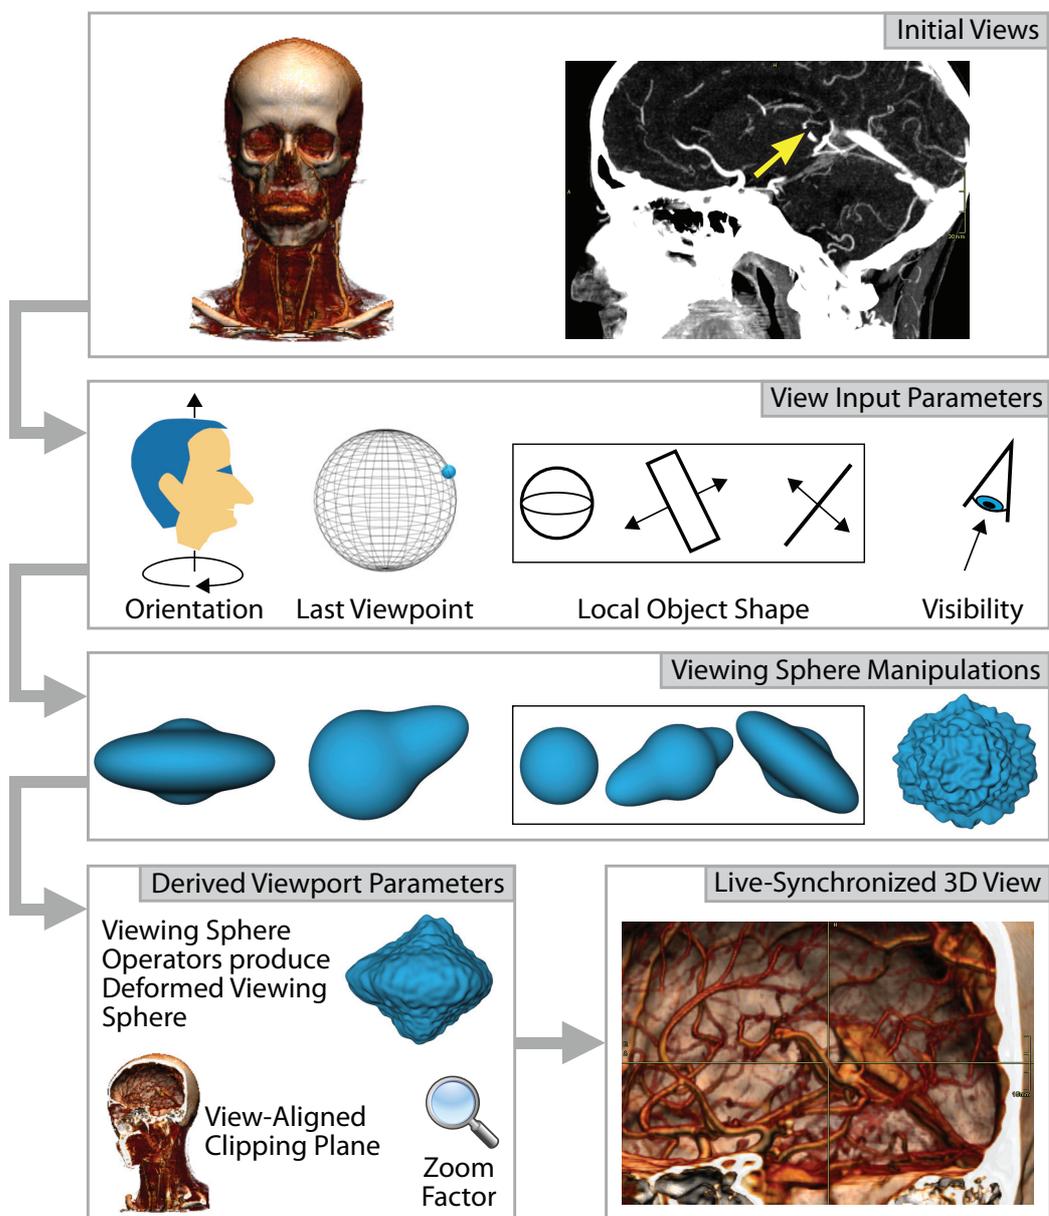


Figure 2.1: LiveSync workflow: Initially there is a volumetric view and a 2D slice image. A picking action on the slice initiates the deformation of viewing spheres for the input parameters: patient orientation, viewpoint history, local shape estimation, and visibility. The combination of these parameters leads to a single deformed viewing sphere which encodes the quality of the viewpoints. In addition, a view-aligned clipping plane is positioned and the zoom is adjusted. These steps generate a live-synchronized volumetric view which provides a good view on the picked structure.

Viewpoint history: The last viewpoint is used as a parameter for the selection of the next viewpoint. This means that the system tries to find a good viewpoint close to the last one if this does not counteract the other parameters.

Local shape estimation: The local shape of the picked structure is estimated based on local segmentation. Three major shapes - lines, sheets and blobs - are assigned to structures to be utilized as parameters for viewpoint selection.

Visibility: Another parameter is the visibility of the picked structure. To compute visibility, rays are cast from the picked point to a certain number of possible viewpoints and analyzed regarding occluding structures.

The parameters patient orientation, viewpoint history, local shape estimation, and visibility are encoded directly in the viewing spheres. If the particular parameter indicates a good viewpoint at a certain position, a unit sphere is deformed in a way that the distance of this point to the sphere's center is increased. Figure 2.1 gives an overview of the LiveSync workflow. Initially there is a volumetric view which is shown from a default viewpoint and a 2D slice view. For each picking action on the slice, the input parameters are used to estimate good viewpoints and to deform the viewing spheres accordingly. This is achieved without any data-specific a priori information or precomputations. The input parameters have to be weighed and combined to get a resulting deformed sphere which encodes the combined quality of the viewpoints. In addition, the zoom factor is adjusted and a view-aligned clipping plane is defined which allows a flexible removal of occluding structures to generate a meaningful visualization.

2.4 Viewing Sphere

The viewing sphere and the camera analogies are well-known concepts for setting up a viewpoint and a viewing direction in computer graphics. Basically, a virtual camera can be placed at any point on the surface of a sphere which encapsulates the scene. To move the camera on this sphere typically rotation operations are performed. In addition, the viewing direction of the camera defines on which location in the scene the camera is focusing. Zooming can be achieved by moving the camera along the surface normal of its position on the sphere.

2.4.1 Sphere Parameterization

As the input parameters have to be encoded directly into the sphere's shape there is need for an intuitive way to parameterize the viewing sphere. In addition this

2.4 Viewing Sphere

parameterization has to be stored efficiently with taking into consideration that operators for the combination of the individual spheres have to be applicable. A convenient parameterization of spheres can be achieved with polar coordinates. In this system each point of a sphere can be characterized by θ and φ , which represent the polar and the azimuthal angle, and its radial distance r . The polar angle starts from the positive z -axis and ranges from 0 to 180° and the azimuthal angle in the xy -plane starts from the positive x -axis with a range from 0 to 360° . With this parameterization several conversions and calculations can be computed very efficiently [59, 71].

2.4.2 Sphere Map

A well-known challenge in computer graphics is the problem of applying a texture map to a sphere. The naive approach performs a direct latitude-longitude mapping onto a sphere by using a single rectangular texture in which the width is twice the height. With uv -mapping u spans the equator and v covers the pole-to-pole range. This is a straightforward mapping with the disadvantage that the sampling becomes higher towards the pole regions. Alternatives for spherical textures are cube, omnitect, icosahedral and octahedral mappings [65].

The inverse problem has to be handled to map a sphere to a structure which facilitates the operations that are performed in the presented concept. Because of memory efficiency and intuitive indexing the direct latitude-longitude mapping was the technique of choice. The rectilinear texture is stored as a two-dimensional array with 360×180 entries. Explicit storing in memory is necessary to facilitate an efficient combination of differently sampled data. In the current implementation information about patient orientation, viewpoint history, and local shape estimation is analytically described, whereas visibility information is sampled in a discrete manner. As the angular position can be calculated from the array indices it is sufficient to write the radial distance values to this array.

2.4.3 Sphere Deformation

The general idea to indicate the quality of viewpoints is the direct deformation of the viewing sphere. Positions on the sphere's surface with a high radial distance represent good viewpoints. To achieve an appropriate deformation of the sphere, the Phong illumination model serves as an analogy. In this model a hemisphere represents the diffuse reflection intensity with a bump which indicates the specular reflection intensity. Phong's model of the specular highlight is adapted for the calculation of the radius r at a certain point on the sphere's surface with the equation

$$r = a \cdot (\mathbf{n} \cdot \mathbf{v})^{m_w}, \quad (2.1)$$

where a is a constant which controls the height of the bump, \mathbf{n} is the surface normal at a specific point on the sphere, \mathbf{v} is the surface normal at a good viewpoint and m_w controls the width of the bump. With slight variations of this formula the deformed spheres for most input parameters used for viewpoint selection can be generated.

2.5 Viewing-Sphere Manipulators

A challenging part in the selection process of a good viewpoint is the identification of the relevant parameters. For a generic solution which works for different types of medical volume data, the definition of objective parameters is important. The patient's orientation, the viewpoint history, the local shape of the structure and its visibility are considered to be of relevance for viewpoint selection. Viewing spheres are deformed to encode the viewpoint quality for each of these components.

2.5.1 Patient-Orientation Viewing-Sphere

The first utilized parameter to construct a deformed viewing sphere is the patient's orientation. According to the type of an examination there exist general preferred viewing directions. In this case the head-feet axis serves as a rough estimation to derive the preferred viewpoints. Figure 2.2 (left) shows the rotation axis which corresponds to the patient's orientation. The viewing sphere is deformed in a way that it prefers viewpoints which are orthogonal to this axis. This deformation is achieved by applying Equation 2.1 as it is described in Algorithm 1 where the z-axis is the main rotation axis.

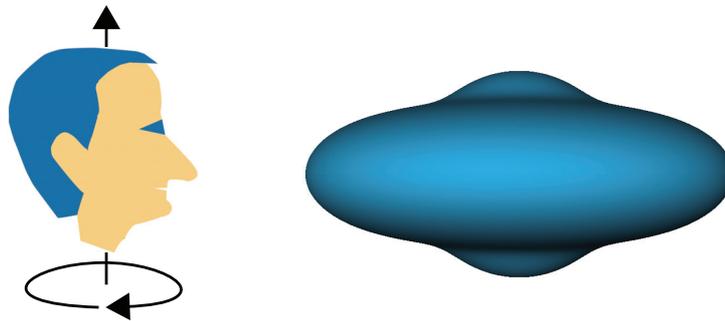


Figure 2.2: The head-feet axis is assumed to be the main rotation axis (left). To encode this information in the viewing sphere it is enlarged around the equator (right).

2.5 Viewing-Sphere Manipulators

Algorithm 1 Generation of the patient-orientation viewing-sphere

```
for each line of longitude  $lon$  do  
  set  $v$  to the surface normal at a latitude of  $90^\circ$   
  for each parametrized point  $p$  of  $lon$  do  
    set  $n$  to the surface normal at  $p$   
    compute radius at this point with Equation 2.1  
  end for  
end for
```

2.5.2 Viewpoint-History Viewing-Sphere

The selection of a good viewpoint is based on different input parameters to provide the user with an intended view. As a specific view was selected by the system based on estimated demands of the user, the current viewpoint will also be considered for the estimation of the quality of the next viewpoints. Especially, big shifts of the viewpoint for two successive pickings should be avoided if possible. This means that if there is a good viewpoint for the picked structure close to the current one this viewpoint is preferred to others which are positioned farther away on the viewing sphere.

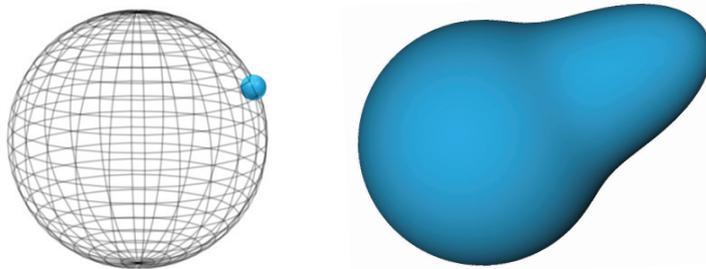


Figure 2.3: *The marked position on the viewing sphere indicates the last viewpoint (left). The deformed sphere has a single bump which indicates the quality of the viewpoints (right).*

Figure 2.3 shows how the deformed sphere for this criterion should look like. The position of the last viewpoint is marked on the viewing sphere. After deformation the resulting viewing sphere should have a bump with a maximum at this position which also encodes the quality of surrounding viewpoints. The according deformation can be generated with Algorithm 2.

Algorithm 2 Generation of the viewpoint-history viewing-sphere

```

set  $v$  to the surface normal of the last viewpoint
for each point  $p$  of the parameterized sphere do
  set  $n$  to the surface normal at  $p$ 
  if  $\text{dot}(v,n) > 0$  then
    compute radius at  $p$  with Equation 2.1
  else
    set radius to 1
  end if
end for

```

2.5.3 Local Shape-Estimation Viewing-Sphere

Another important input parameter for viewpoint selection is the local shape of the structure of interest. If the picked point is, e.g., part of a blood vessel, a good viewpoint shows the course of this vessel and does not cut through it. With a fast local segmentation and a principal component analysis (PCA) the shape information can be derived locally from the data values. Region growing is performed on a $32 \times 32 \times 32$ neighborhood of the picked data point which serves as seed point. The lower and upper threshold for the region growing are calculated by analyzing the distribution of the scalar values at the picked point and its neighborhood. The result of this local segmentation is a connected 3D point cloud. PCA is performed on this point cloud to extract the three feature vectors and the corresponding eigenvalues which are utilized to determine the local feature shape according to a metric of Westin et al. [67]. Figure 2.4 shows how the vector of the first principal component is oriented when picking is performed at three different positions on blood vessels in the head. The local orientation of the vessels is indicated by these vectors quite well. In combination with the orthogonal second and third principal components and the corresponding eigenvalues this information is used to create the deformed spheres for the local shape estimation.

According to the local shape of the object, the viewing sphere has to be deformed as illustrated in Figure 2.5. If the object has a volumetric extent (blob), then basically all viewpoints are of the same quality (left). For a planar structure (sheet) the viewpoints which are orthogonal to the sheet are favored (middle). If a tubular structure (line) is determined, the preferred viewpoints are aligned along a ring which is orthogonal to this line (right). For the planar object the deformation of the sphere is calculated analogous to the deformed sphere for the viewpoint history. To get two bumps on the opposite sides of the sphere Equation 2.1 is adjusted slightly to

2.5 Viewing-Sphere Manipulators

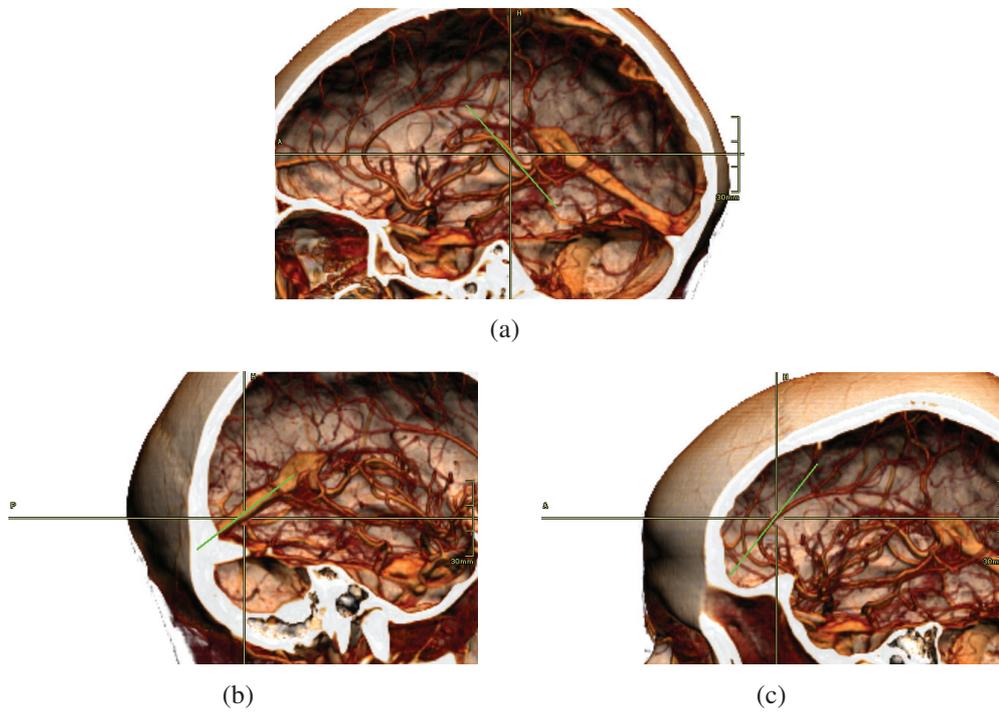


Figure 2.4: The green line displays the orientation of the most important feature vector determined by a PCA for three different positions on blood vessels in the head. These vectors are strongly aligned with the local orientation of the vessels.

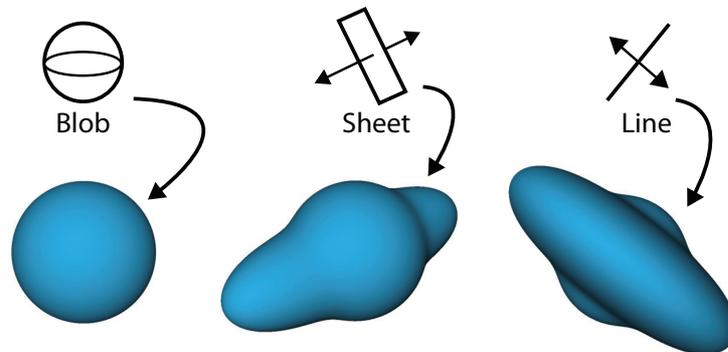


Figure 2.5: The viewing sphere which is generated for the local shape estimation is deformed according to the major volumetric extent of the structure.

$$r = a \cdot \text{abs}((\mathbf{n} \cdot \mathbf{v})^{m_w}). \quad (2.2)$$

If the structure is tubular the deformation process is a bit more complex. It is a generalization of the deformation process of the patient-orientation viewing-

sphere because the tube can be oriented arbitrarily within the volume. Geometrically the good viewpoints are located around a great circle of the viewing sphere, defined by the two points where the vectors of the second and the third principal components intersect the sphere's surface. A great circle is always uniquely defined by two points on the surface of the sphere and its center is the same as the center of the sphere. For each position p on the sphere's surface the vector from the origin to the closest point on the great circle has to be calculated. This can be achieved by projecting the vector from the origin to p onto the plane of the great circle. The procedure to generate the deformed sphere is presented in Algorithm 3.

Algorithm 3 Generation of the local-shape viewing-sphere

```
if shape == blob then
    radius of each point of the parameterized sphere is 2
else if shape == sheet then
    set  $v$  to the vector of the third principal component
    for each point  $p$  of the parameterized sphere do
        set  $n$  to the surface normal at  $p$ 
        compute radius at  $p$  with Equation 2.2
    end for
else if shape == line then
    calculate great circle  $c$  for the two points where the 2nd and the 3rd principal
    component intersect the surface of the unit sphere
    for each point  $p$  of the parameterized sphere do
        set  $n$  to vector from the origin to  $p$ 
        set  $v$  to the projection of  $n$  onto the plane of  $c$ 
        normalize  $v$ 
        compute radius at  $p$  with Equation 2.1
    end for
end if
```

2.5.4 Visibility Viewing-Sphere

A further building block for estimating a good viewpoint is defined by the visibility information. Starting from the picked point visibility rays are cast to determine occluding objects. As stated in Section 2.4.1 the parameterized points of the sphere are not distributed uniformly. It is neither efficient nor necessary to cast visibility rays to all 360×180 positions. Nevertheless it is highly preferable that the positions which are tested are distributed uniformly on the sphere. Bourke [7] provides source code (written by Lettvin) for this purpose. Based on the standard

2.5 Viewing-Sphere Manipulators

physics formula for charge repulsion an arbitrary number of points is distributed over the surface of a sphere. In our experiments we determined that a subset of 36×18 rays provides a good trade-off between performance and quality. The calculation of the uniformly distributed points is performed only once and the result is stored in a look-up table.

To determine whether a certain viewpoint provides good visibility of the selected structure, rays are cast from the picked point. As a local segmentation was performed for the local shape estimation, this information is utilized to determine when a ray exits the tissue of interest. When this has happened the opacity information of the transfer function is considered. The opacity is accumulated along the ray and as soon as a small opacity threshold is surpassed the calculation is terminated for the specific ray. A high visibility value is assigned to a viewpoint if there is much space from the picked point in the direction of this viewpoint until it gets occluded by other structures. Such a situation provides more flexibility for positioning the clipping plane. This allows to position the clipping plane orthogonal to the viewing direction far away from the picked point, so that an unobstructed view of the picked point is possible while the helpful context information is not unnecessarily reduced.

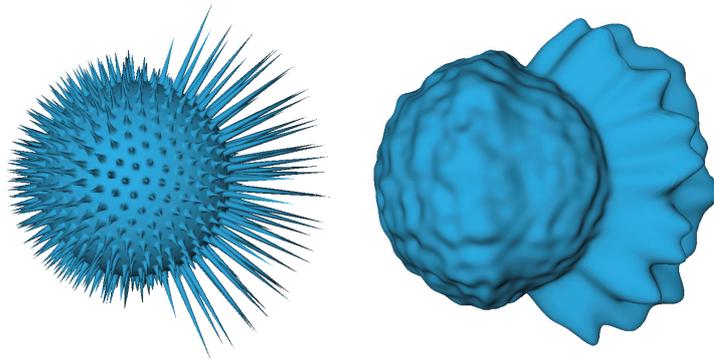


Figure 2.6: *The lengths of the spikes encode the viewpoint quality at a uniformly distributed set of sample positions (left). After reconstructing at all positions a smooth sphere is generated (right).*

The deformed sphere is depicted in Figure 2.6. One important criterion for the viewpoint entropy of Bordoloi and Shen [6] is *view stability* which describes the maximum change in a certain view caused by small camera shifts. The view is defined to be stable if a small camera change implies also only small changes in the view. Transferred to the visibility viewing-sphere there is the possibility to encode view stability derived from the visibility values at the discrete uniformly distributed points. It is heuristically assumed that a viewpoint in between several good ones is also rather good. Such a point offers high view stability, because

small changes of the viewpoint will also lead to good viewpoints. To encode this information into the viewing sphere, for all the parameterized sphere positions which are not explicitly tested for visibility a weighing with the surrounding tested points is performed. With this weighing a smoothly deformed sphere is obtained as shown in Figure 2.6 (right). The pseudocode to generate the deformed sphere for the visibility criterion is presented in Algorithm 4.

Algorithm 4 Generation of the visibility viewing-sphere

```
for each of the uniformly distributed points  $p$  do
    calculate the visibility
    set the radius at  $p$  to the visibility value
end for
for each point  $s$  of the parameterized sphere do
    get all  $ps$  within a certain distance  $d$  to  $s$ 
    for each  $p$  in range  $d$  do
        set  $n$  to the surface normal at  $p$ 
        set  $v$  to the surface normal at  $s$ 
        compute  $r$  with Equation 2.1
        add  $r$  to the current radius at  $s$ 
    end for
    normalize the radius at  $s$ 
end for
```

2.6 Viewing-Sphere Operators

After the generation of the deformed viewing spheres for the various view input parameters the question arises how they have to be weighed and how they can be combined to simultaneously accommodate for all the effects.

2.6.1 Weighing of Viewing Spheres

Equation 2.1 offers different options to weigh the extent of deformation of a sphere. Basically, a controls the height of the bump and m_w its width. To facilitate the combination operators the values of a for the individual sphere deformations are chosen so that their radii vary from 1 to 2 after deformation. For all the input parameters it makes sense that an estimated good viewpoint also influences the quality of viewpoints in a certain neighborhood. For each viewpoint criterion the radius can vary by a factor of two around a good viewpoint at a certain position.

2.6 Viewing-Sphere Operators

The sphere generation for the viewpoint history contains a built-in weighing control. A big shift of the viewpoint is quite disturbing for picking actions within a small spatial area of the data but it is acceptable for two picked points which are located far apart from each other. This just means that the user is switching to a totally different inspection region whereby the viewpoint coherency is less critical. A distance factor d is calculated as the ratio of the spatial distance between two successively picked points to the diagonal extent of the volume. To influence the weighing for the viewpoint-history viewing-sphere Equation 2.1 is modified to

$$r = (1 - d) \cdot a \cdot (\mathbf{n} \bullet \mathbf{v})^{m_w}. \quad (2.3)$$

2.6.2 Combination of Viewing Spheres

As the deformed spheres were calculated for the input parameters individually, they have to be combined into a single sphere which encodes the overall quality. Currently three operators are implemented for this combination - summation, multiplication, and thresholding. Each of these operators emphasizes certain viewpoint characteristics. Figure 2.7 shows the effects of the three operators on the resulting sphere. For this example the visibility viewing-sphere and the local shape-estimation viewing-sphere are chosen as input spheres. The application of these operators and the development of additional operators is easy to achieve because each deformed sphere is parameterized as a two-dimensional array.

As operands for the operators, the offset of the radius which is higher than the radius of the unit sphere is taken. At each position the radius of a deformed sphere has a value between 1 and 2 so that the operations are performed on values between 0 and 1. The implementation and the characteristics of the realized operators are as follows:

Summation: A loop over all entries of the sphere arrays is performed to sum up the corresponding radii. This intuitive approach leads to very good results. Good viewpoints will be detected at positions where at least some of the input spheres indicate a good one. Summation is not as sensitive to outliers as multiplication or thresholding.

Multiplication: To emphasize certain characteristics more strongly an operator is implemented which computes the multiplication of the input spheres. This operator emphasizes positions where good viewpoints are indicated by several source spheres and deemphasizes positions where at least one source sphere indicates a bad viewpoint. Low values have an increased impact on the result. Even if the value of only one input sphere is low the corresponding viewpoint will be rated as a bad one.

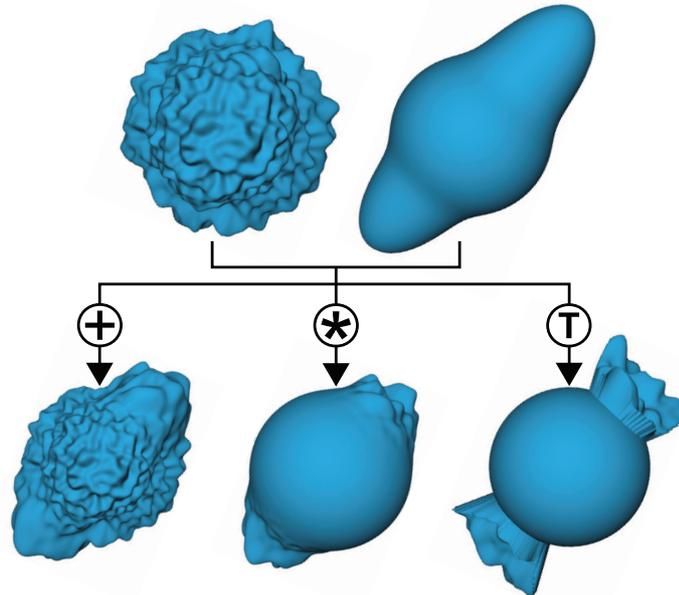


Figure 2.7: The currently implemented operators to combine the deformed spheres are summation, multiplication and thresholding.

Thresholding: For the thresholding operation one specific sphere is taken as the initial one. In a loop over all parameterized points of this sphere the value at this position is only considered if the values of the other spheres at the same position are above a certain threshold. If this is not the case the radius on the specific position is set to 1. This operator filters out the values where the corresponding values on the other spheres are indicating a bad viewpoint. With the thresholding it is possible to define knock-out criteria. Assuming the patient-orientation viewing-sphere is the initial thresholding sphere a window for a certain preferred viewing direction can be defined. By thresholding over the other deformed viewing spheres a good viewpoint within this frame will be estimated.

2.7 Derived Viewport Parameters

After describing the view input parameters, the viewing-sphere manipulators and the viewing-sphere operators, the parameters to set up the volumetric view can be derived. These are the parameters for a good viewpoint, the placement of the view-aligned clipping plane, the zoom factor, and the viewing direction. The

2.8 Results and Evaluation

application of the viewing-sphere operators to the individual deformed viewing spheres produces the combined viewpoint quality at 360×180 positions on the joint viewing sphere. A good viewpoint can be easily determined by the highest entry in the sphere-map array which holds the radial distances of all points. The system can then display the data according to the best estimated viewpoint or suggest a small number of preferred views (e.g., displayed as thumbnails).

With the information obtained by the visibility calculation described in Section 2.5.4, the exact position where the picked point is occluded along each tested visibility ray is known. This information is used for setting up a view-aligned clipping plane to clip away the occluding structures. To position the clipping plane, a location along the ray starting at the picked point where the accumulated opacity is still below a small threshold is selected. This allows an unobstructed view of the picked object while preserving as much context information as possible. The viewing direction is directly defined by the picked point and this point is shown in the center of the volumetric view window. Finally, the zoom factor for this view can be derived from the current settings of the slice view. The zoom of the slice view gives a rough estimation about the size of the interesting anatomical structure. In the current implementation this zoom factor directly determines the zoom of the volumetric view.

2.8 Results and Evaluation

For a convenient evaluation of the results of the LiveSync concept the implementation is integrated into a real-world medical workstation which is under development by our collaborating company partner. All the computations for the LiveSync viewpoint selection can be performed interactively and are not influenced significantly by the size of the data set. The performance was measured on a PC configured with an *AMD Athlon 64 Dual Core Processor 4400+* and 2 GB of main memory. In the unoptimized implementation, the LiveSync-related computations took 70 ms to 150 ms per picking, depending on the number of segmented voxels at the local segmentation step and on the estimated local feature shape. The users get an instant update of the volumetric view whenever they are picking a certain structure on a 2D slice. In the remainder of this section, LiveSync-generated images for three different application scenarios and the results of an informal evaluation will be provided to demonstrate the usefulness of the interactively synchronized views.

For the first scenario an estimated good viewpoint and a rather bad viewpoint will be presented for comparison reasons. Figure 2.8 shows the results for a picking action on a slice as depicted in Figure 2.8(a). The picking is performed with the aim to get information on the course and the spatial vicinity of the partly vis-

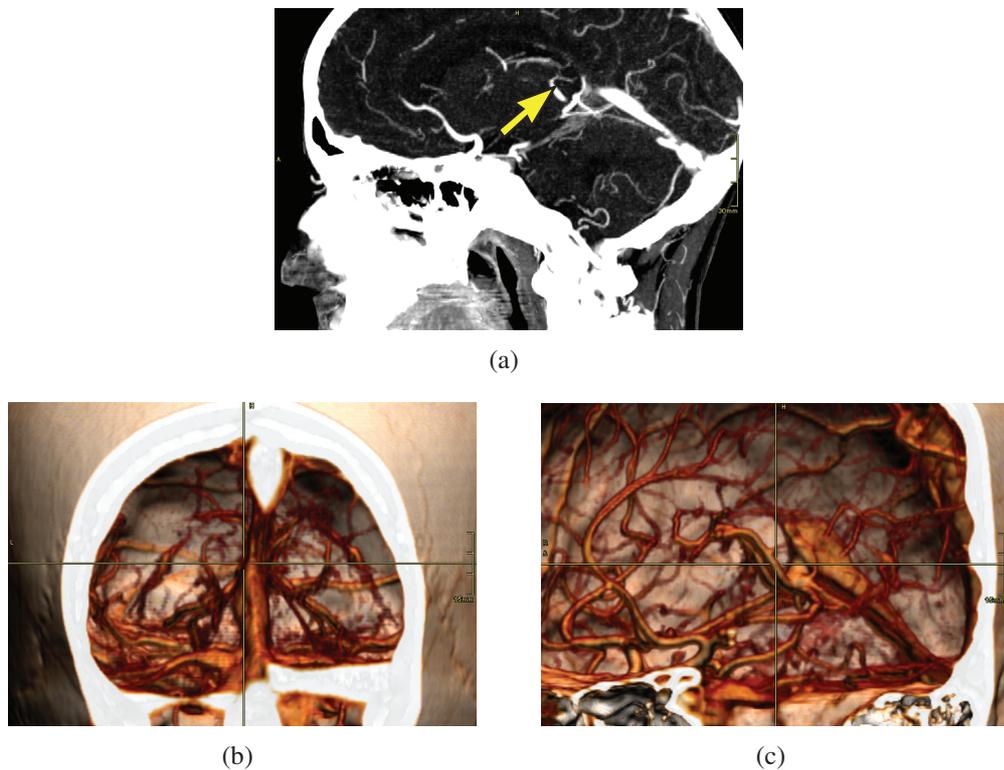


Figure 2.8: (a) The picking on the partly visible vessel in a 2D slice image. (b) A viewpoint which is estimated as being rather bad. (c) A good viewpoint provides useful context information.

ible vessel. A viewpoint which is rated as rather bad by the viewpoint selection process leads to a volumetric view as shown in Figure 2.8(b). Important parts of the vessel are occluded, its course remains unclear and the connectivity to other vessels can hardly be revealed by this viewpoint. In Figure 2.8(c), where a good viewpoint defined by the deformed viewing sphere is considered, the information about the vessel's course and its spatial vicinity can be extracted easily. The other two application scenarios demonstrate that LiveSync is a generic tool for various kinds of clinical examinations. In their typical workflow radiologists search for specific structures in medical data. Although there exist highly sophisticated and specialized methods, e.g., for the detection of polyps in the colon or lung nodules, LiveSync can help to quickly explore these pathological cases. The examination of the colon is a very difficult task with 2D slices only because there it is very hard to see differences between foldings of the colon and polyps. Figure 2.9 shows the LiveSync result for the picking on a suspicious structure in the colon. With the provided volumetric view it can be clearly seen that the picked structure is not

2.8 Results and Evaluation

a colon folding but a polyp. Another challenging task is the detection of lung nodules. In the 2D slices they often look very similar to bronchia or vessels. In Figure 2.10, a structure which is assumed to be a nodule is picked on the slice and LiveSync presents the corresponding volumetric view automatically. This view can clearly help to classify the picked structure as a lung nodule.

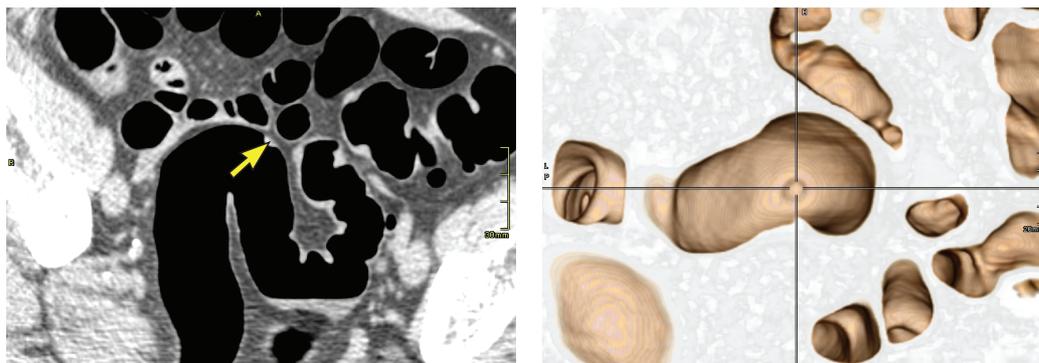


Figure 2.9: The picking on a suspicious structure in the 2D slice of the colon (left) leads to an automatically generated volumetric view (right) in which a polyp can be identified.

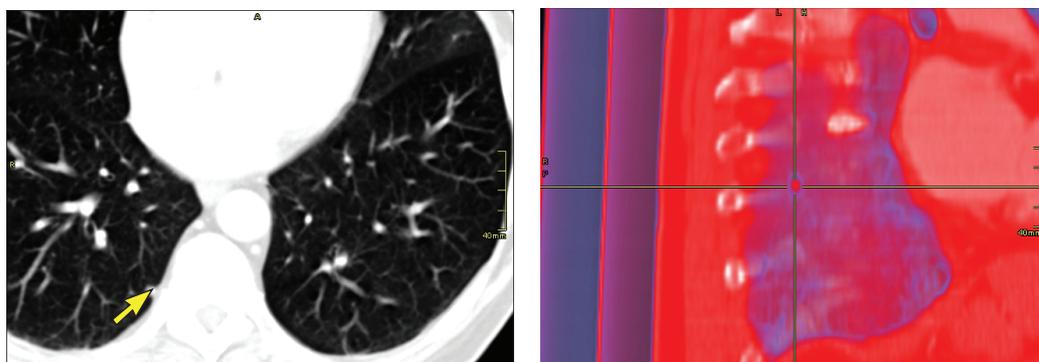


Figure 2.10: Lung nodules can be hardly distinguished from vessels or bronchia in the slice view (left). In the volumetric view (right) the picked structure can be clearly classified as a lung nodule.

In an informal evaluation LiveSync was used by an experienced radiology technician. The used data sets were a head CT angiogram (CTA), a peripheral CTA, and a CT of the chest. The goal was to generate 3-4 diagnostically relevant volume renderings of pathologies in the different data sets. This task had to be done at first manually and then with LiveSync support. In both cases a predefined transfer function was given and the radiology technician had to adjust the other parameters. For all data sets the overall expenditure of time ranged from 5

to 20 min in the manual case and from 2 to 10 min with LiveSync support. Figure 2.11 presents an exemplary result of the evaluation with the head CTA. To get a diagnostically meaningful view on the aneurysm, the parameters were at first adjusted manually. The resulting image is shown in Figure 2.11(a). Figure 2.11(b) shows the instant result of just picking a point on a slice. With very little user interaction a clipping plane was adjusted to generate an almost identical view to the manually adjusted one as shown in Figure 2.11(c). Similar results were achieved for the examination of the ureter. Figure 2.12(a) shows the manually adjusted image. A single picking on the ureter on the 2D slice was sufficient to generate the image shown in Figure 2.12(b). Viewpoint and zooming are very similar in both images. The only clearly visible difference is again caused by the manual clipping of posterior parts in Figure 2.12(a).

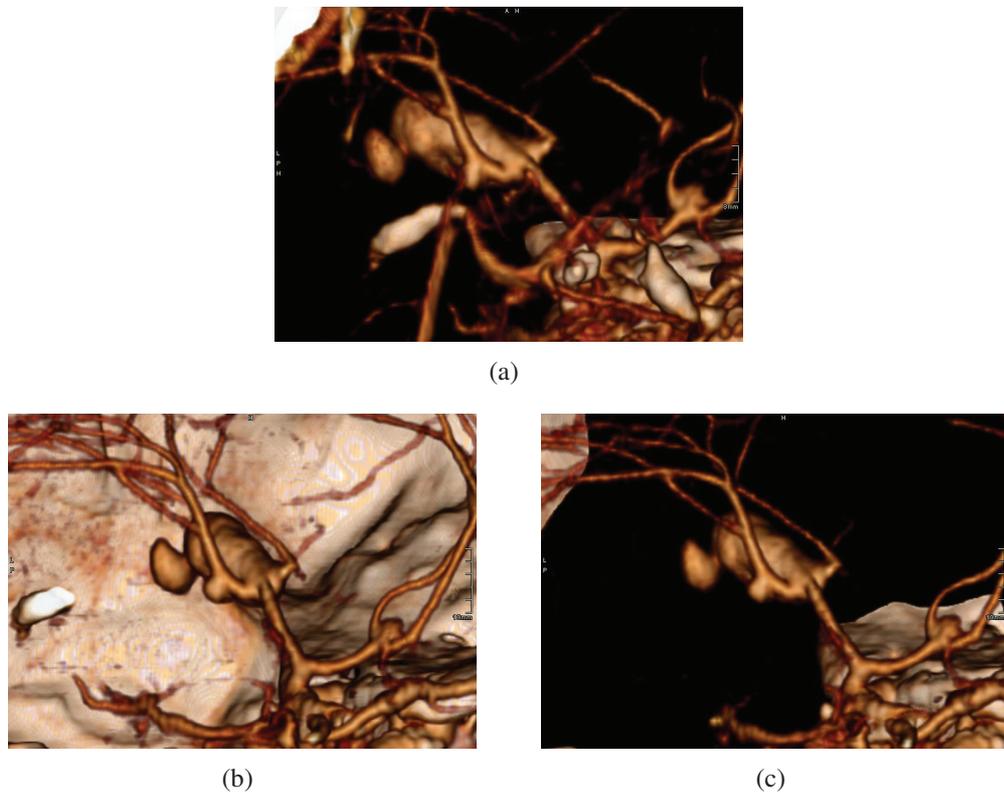


Figure 2.11: Manually-adjusted vs. LiveSync-generated image. (a) The manually adjusted image to get a good view on the aneurysm. (b) The LiveSync-generated image looks very similar. (c) After manual clipping the LiveSync-generated image looks almost identical to the manually-adjusted image.

2.9 Conclusion

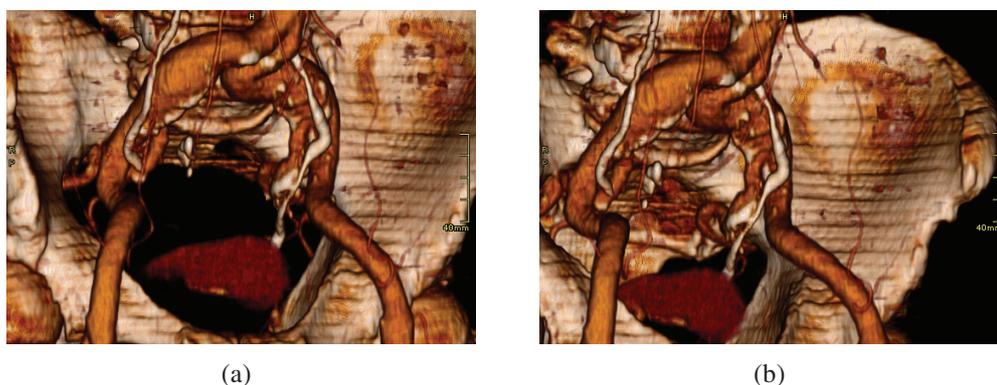


Figure 2.12: Manually-adjusted vs. LiveSync-generated image. (a) The manually adjusted image to get a good view for the examination of the ureter. (b) The LiveSync-generated image looks very similar to the manually-adjusted image.

In this informal evaluation the main conclusion was that the effort to localize pathologies diminished considerably when LiveSync was used. The initial views which were suggested with LiveSync support were already very good, but some additional time was needed to adjust them to get screenshots which are perfectly suitable for diagnosis. In the majority of the LiveSync results only small user interaction was necessary to produce screenshots with the same diagnostic value as fully manually adjusted 3D views. These results are especially satisfying as the radiology technician had no prior experience with the new LiveSync interaction metaphor. The overall impression of the LiveSync feature during the evaluation was that it provides an excellent additional functionality. This opinion is also supported by radiologists getting demonstrations of the LiveSync functionality.

2.9 Conclusion

In this chapter, we presented a novel concept for the live synchronization of 2D slice views and a volumetric view within a medical workstation. Different input parameters are identified to generate deformed viewing spheres which indicate the quality of the viewpoints for the specific criteria. After the combination of individual spheres the combined quality of the viewpoints is encoded for 360×180 parameterized positions on a viewing sphere. Our system provides a good viewport estimation considering patient orientation, viewpoint history, local object shape, and visibility. Further it performs an automated placement of a view-aligned clipping plane and zooming. Picking on an interesting structure in the 2D slice image is the only necessary user interaction to get a meaningful

volumetric view. The performed evaluation indicated that the presented approach might considerably improve the efficiency of diagnosis in clinical routine.

An interesting direction for future research might be the integration of more information gathered by the user interaction history. With this information it should be possible to generate templates of deformed viewing spheres for different clinical application scenarios. These spheres can provide the physicians with optimal views and optimal view paths for their examinations.

A good thing never ends.

Mick Jagger

3

LiveSync++: Enhancements of an Interaction Metaphor

This chapter is based on the following publication:

P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. LiveSync++: Enhancements of an interaction metaphor. In *Proceedings of Graphics Interface 2008*, pages 81–88, 2008.

3.1 Introduction

In the clinical routine the available time to investigate large quantities of patient data, recorded by modern medical imaging modalities (e.g., computed tomography), is very limited. As today's medical datasets may contain several hundreds of slices, it is already very time-consuming to scroll through them. As soon as a potentially pathological area is detected on a slice it can be very helpful to visualize its three-dimensional context. The prime reason which prevents a broad usage of 3D visualizations in the clinical routine is the tedious work to set up all the needed parameters. For example, the rotation in a three-dimensional space to select a good viewpoint is often not very intuitive for people, who are not dealing with computer graphics. Also the definition of a transfer function, which shows the feature of interest and its anatomical context in an expressive way often is a time-consuming trial-and-error process. Another issue which appears by moving from 2D to 3D visualization is occlusion. Clipping planes can be defined to remove structures which occlude the feature of interest. A typical set of clipping planes allows an object-aligned removal along the x-, y-, and z-axis, as well as a near and far clipping along the view direction. Finally the zoom factor for the 3D view has to be adjusted manually.

The goal of the LiveSync interaction metaphor is to derive all the parameters which are needed to present a meaningful 3D view with minimal interaction ef-

fort. As the radiologist is usually examining the slices with a keyboard and mouse as input devices, pointing with the mouse on a structure of interest combined with pressing a hot-key is not very intrusive. Chapter 2 introduced the basic concept of LiveSync. To encode the quality of a viewpoint for different input parameters, viewing spheres are deformed. A viewing sphere surrounds the entire object and all possible viewpoints are located on the surface of the sphere. The viewing direction points to the center of the sphere. This concept presented the basic building blocks to achieve the live synchronization of a 2D slice view and a volumetric view. The following improvements and extensions of the LiveSync interaction metaphor are presented in this chapter:

Sphere parameterization: An efficient sphere parameterization is used to encode the viewpoint quality for the input parameters. With a parameterization in polar coordinates, the distances between neighboring points vary considerably, depending on their distance to the poles. A multi-resolution approach which samples the sphere with uniformly distributed points is used. This allows hierarchical refinement of the sampling for efficient calculations of the overall viewpoint quality. Besides its efficiency regarding memory consumption and performance, it can be employed to find a good viewpoint that is surrounded by other viewpoints which are estimated as good ones.

Feature extraction: The extraction of a feature of interest is a crucial part for finding good viewpoints. Our region growing based segmentation automatically determines the size of the region which is necessary to specify the shape of the object of interest.

Transfer function: The manual setup of a transfer function is rather time-consuming and often not very intuitive. With increasing dimension of the transfer function space the definition becomes more and more complicated. For alleviation, the user can choose from a predefined set of transfer functions which are designed for typical examination procedures. With the knowledge about the distribution of scalar values within the feature of interest, it is possible to fine-tune a predefined transfer function. This is especially important when a structure is picked on the 2D slice which is not visible in the 3D view with the currently defined transfer function.

Interaction modes: An informal evaluation indicated that a single click on the slice often results in a very good view on the interesting object in the 3D view. However, the size of the area of interest is hard to depict without further user interaction. As it is preferred to keep the user effort as low as possible, the time of keeping a hot-key pressed is taken to interactively

3.2 Related Work

increase the area of interest. The view is updated in short time intervals and the user can decide when an intended view is reached. In this mode a mask with the current segmentation result can be displayed to inspect the region growing process. An automatically generated history helps the user to restore and to keep previously generated views. This is especially helpful to review the intermediate states which are generated in this new interaction mode.

Feature-driven clipping: In the initial implementation presented in Chapter 2 only the setup of a view-aligned clipping plane was supported. The performed evaluation showed that often object-aligned clipping planes are preferred, especially if there is further interaction necessary to manually fine-tune the view. In the presented approach the user can choose if the view-aligned or the object-aligned clipping planes are set automatically to remove occluding structures, while preserving as much contextual information as possible. To increase the degree of preservation, importance-driven clipping techniques are integrated.

This chapter is structured as follows: In Section 3.2, the relevant previous work is discussed. A brief summary of the LiveSync workflow and functionalities is given in Section 3.3. Section 3.4 describes the improved parameterization of the spheres. In Section 3.5, it is first shown how the extent of the region growing-based segmentation and the other growing parameters are derived automatically. In the following, an approach is presented to fine-tune a predefined transfer function by taking the distribution of scalar values within the feature of interest into account. Section 3.6 introduces the different interaction modes and Section 3.7 describes the feature-driven clipping. Section 3.8 gives numbers about the performance of the live synchronization and presents qualitative feedback from users. The chapter is concluded in Section 3.9.

3.2 Related Work

A major task of LiveSync is the computation of viewpoint quality. An overview of the related work which is targeting the selection of a good viewpoint for polygonal scenes and for volumetric data is given in Section 2.2.

In the context of volume visualization there are several approaches which attempt to perform either part of or the entire setup of a transfer function automatically. The RGVis techniques introduced by Huang and Ma [18] present a 3D region growing approach to assist the user in locating and defining features of interest in volume data. They perform partial region growing to generate a transfer function, which reveals the full feature of interest. Another work on

semi-automatic generation of transfer functions was presented by Kindlmann and Durkin [27]. They assume that the boundary regions between relatively homogeneous materials are the areas of interest in the scalar volume. After generating a three-dimensional histogram using the scalar values, the first and the second directional derivatives along the gradient direction they can generate the opacity transfer function. A cluster-space approach for the classification of volume data was presented by Tzeng and Ma [58]. In a preprocessing step they transform volumetric data into a cluster space representation. This provides an intuitive user-interface which allows the user to operate in this cluster space. Rezk-Salama et al. [43] presented an approach where high-level transfer function models designed by visualization experts can be controlled by a user interface which provides semantic information. The non-expert user only has to adjust sliders for a goal-oriented setup of a suitable transfer function.

There are various approaches for the rendering and accentuation of identified features. Marchesin et al. [34] introduced locally adaptive volume rendering to enhance features. They modify the traditional rendering equation to improve the visibility of selected features independently of the defined transfer function. Huang et al. [19] presented an automatic approach to generate accurate representations of a feature of interest from segmented volume data. They construct a mesh for the boundary of the volumetric feature to enable high-quality volume rendering.

Regarding sphere parameterization, there is various research on how to distribute points evenly on the surface of a sphere. Bourke [7] presented an approach which distributes an arbitrary number of points over the surface of a sphere based on the standard physics formula for charge repulsion. Leopardi [30] worked on the partition of the unit sphere into regions of equal area and small diameter. Górski et al. [15] presented HEALPix which is a framework for high-resolution discretization and fast analysis of data distributed on the sphere. A hierarchical equal area iso-latitude pixelization produces a subdivision of the sphere where each pixel covers the same surface area.

3.3 LiveSync Workflow

LiveSync aims to provide an optimal setup of the view parameters for the volumetric view with little additional user interaction. Manually specifying a good 3D view for structures detected in cross-sectional images is a time-consuming task. The user has to edit many parameters, such as camera position and clipping planes in order to get an expressive visualization. This process has to be repeated for each new structure of interest. LiveSync makes the process of diagnosis more efficient by automatically generating good 3D views for interactively picked struc-

3.4 Sphere Parameterization

tures. This functionality can be activated on demand by pressing a hot-key while pointing the mouse on a structure of interest on the 2D slice. Depending on the quality of the instantly generated result, the user can manually refine the view by adjusting the view parameters. This concept allows an efficient and non-intrusive integration of 2D and 3D visualizations in medical workstations.

Based on a picking action on a slice, a set of relevant view input parameters is extracted. In the original system (see Chapter 2) these parameters are *patient orientation*, *viewpoint history*, *local shape estimation* and *visibility*. To get a unified representation of the parameters regarding viewpoint quality, a viewing sphere is deformed for each of them. Each of these deformed spheres contains information about the viewpoint quality. They are combined to encode an estimation about the overall view goodness. The combined sphere now contains information on how to set up the viewpoint for the best estimated view on the structure of interest. After the automatic setup of the clipping plane and the zoom factor, the live-synchronized view is provided. Figure 3.1 gives an overview of the LiveSync workflow. The deformation of the spheres works as shown in Figure 3.2. For viewpoints which are indicated to be good ones based on one of the view input parameters, the radial distance of the surface point is increased. In Figure 3.2, the fully opaque eye indicates a good, the half-transparent one an average-rated, and the crossed out one a rather bad viewpoint.

The conceptual design of LiveSync is not limited to a certain number of input parameters. It can be easily extended by defining more parameters which influence the viewpoint quality. This is facilitated by the unified representation of the encoding of the viewpoint quality. Different weights can be assigned to the view input parameters to control their influence on the combined sphere.

3.4 Sphere Parameterization

As the viewpoint quality of all view input parameters is encoded in a sphere parameterization, a smart way for accessing points on the surface of the sphere is important regarding performance and memory efficiency. In the original implementation there was the problem that the spheres for the view input parameters were sampled differently. Information about patient orientation, viewpoint history, and local shape information was analytically described. Due to expensive calculations, the visibility computation had to be performed in a discrete manner for a precomputed set of points on the surface of the sphere. This set of points was much smaller than for the other parameters. The computed radial distances for all spheres were stored in two-dimensional arrays with 360×180 elements via direct latitude-longitude mapping. This mapping results in a very uneven distribution of

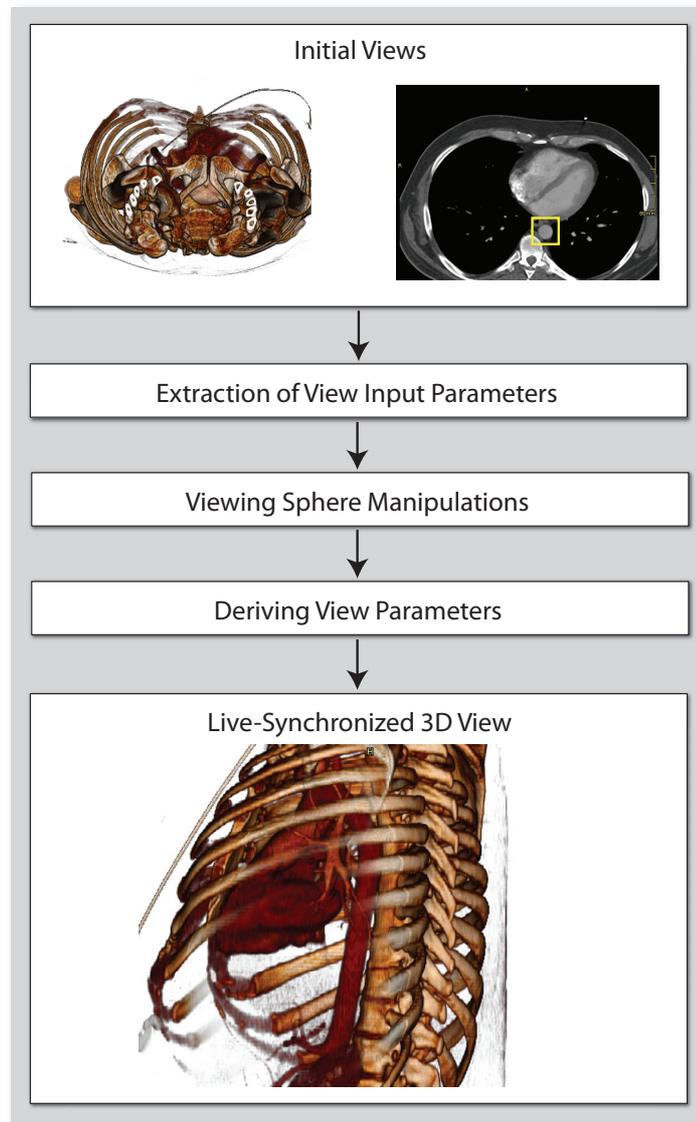


Figure 3.1: LiveSync workflow: The initial views are a 2D slice image and a volumetric view. A picking action on the structure of interest (surrounded by the rectangle) in the slice, starts the deformation of viewing spheres according to automatically derived view input parameters. These parameters are combined to derive the view parameters for setting up the volumetric view. As result an expressive live-synchronized 3D view is generated.

points on the surface of the sphere with much higher sampling close to the poles. Another issue has been the combination of differently sampled spheres.

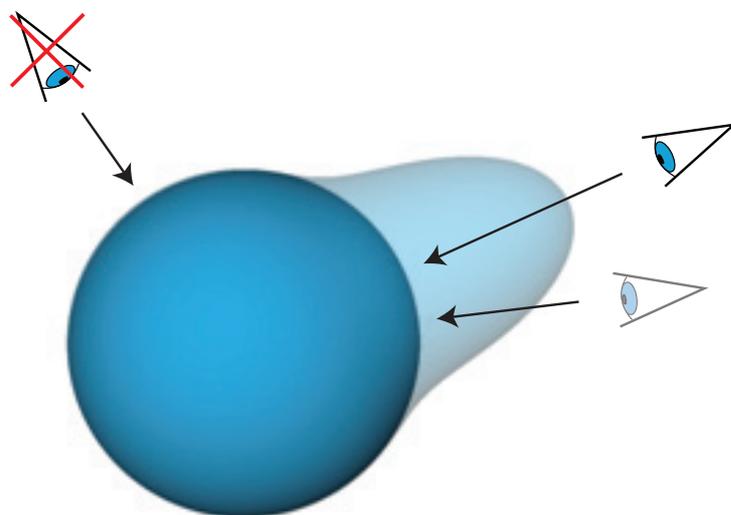


Figure 3.2: The quality of a viewpoint is encoded into the radial distances from a unit sphere. Good viewpoints are located at positions on the unit sphere which refer to high radial distances on the deformed sphere.

3.4.1 Visibility Calculations

For a better understanding of the need for evenly distributed points, this section will provide a brief review of the performed calculations to generate the visibility viewing sphere. Good visibility is given if a ray from the picked point to the viewpoint exits the object of interest within few steps and if the distance until the structure of interest is occluded by other structures is high. In this case, there is high flexibility for positioning a clipping plane to remove occluding structures while preserving important anatomical context around the picked object.

Figure 3.3 illustrates how rays are cast from the picked point on the structure of interest to uniformly distributed points on the sphere. It has to be detected at which distance the ray exits the structure of interest. This is done by analyzing the distribution of scalar values and gradient magnitudes to decide if a voxel belongs to the structure. After a ray exits the structure of interest, opacities depending on the underlying transfer function are accumulated to detect occluding structures. As soon as a small opacity threshold is reached, the computation for a ray is terminated.

3.4.2 Sphere Partitioning

To achieve interactive performance the visibility calculations cannot be performed for 360×180 points. However, it is inaccurate to sample only a small number of points and to combine the sparsely sampled sphere with other, higher-sampled

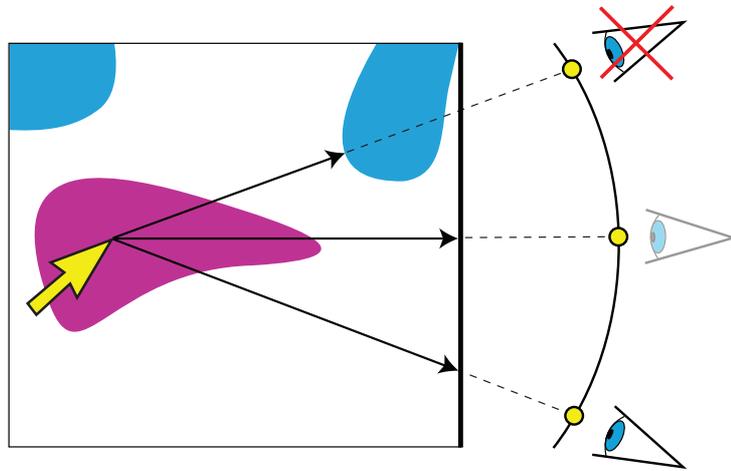


Figure 3.3: Starting from the picked position, visibility rays are cast to points which are equally distributed on the surface of the viewing sphere. Samples along the rays are analyzed to detect (a) when it exits the structure of interest and (b) at which position the object gets occluded by other anatomical structures.

spheres. Moreover, taking the point with the highest radial distance after the sphere combination and before filtering, does not always lead to satisfying results. The structure of interest may be only visible through a small keyhole, and its larger part is hidden by occluding structures. In the ideal case it is very helpful for further interactions with the 3D view if the provided viewpoint offers a good *view stability*. This criterion defined by Bordoloi and Shen [6] describes the maximum change in a certain view caused by small shifts of the camera position.

The capabilities of the HEALPix package [15] offer good properties for the existing needs if the techniques are applied in an elaborate way. Figure 3.4 (top left) shows the HEALPix base partitioning of the surface of a sphere into 12 equally sized areas with dots indicating their center positions. In each subdivision step a partition is divided into four new partitions. A nested indexing scheme can be utilized for the fast access of an arbitrary surface point at different partitioning resolutions.

Applied to the requirements of the LiveSync viewpoint-quality encoding-strategy, the whole sphere is sampled at an initial resolution for all view input parameters. In the following, a partition of the sphere, holding the samples which indicate the best viewpoint quality, is identified. This is done by summing up the radial distances of the sample positions for each partition. To achieve higher resolution, this peak partition is sampled more densely. The process of identifying a peak sub-partition with a successive hierarchical refinement can be repeated if higher resolution is required. To provide good view stability, a final filtering of the points in the target area is performed. The filtered peak is provided as the

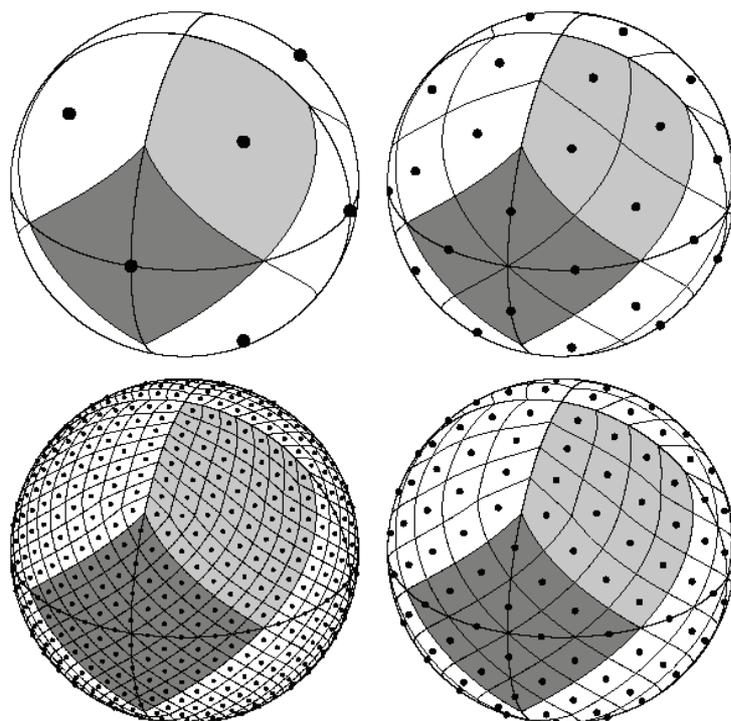


Figure 3.4: In its base partitioning a HEALPix sphere is divided into 12 equally sized areas. In each subdivision step an area is further divided into 4 areas [15].

estimated best viewpoint regarding the view input parameters.

In the presented approach an initial resolution of 3076 points over the whole sphere is chosen, which corresponds to a 3.66 degrees angular distance between neighboring points. After hierarchical refinement, a peak area with 1024 sample points (0.55 degrees angular distance) is identified and low-pass filtering is performed. This approach allows high flexibility concerning the sampling of the sphere, to detect a partition which contains a collection of many good viewpoints.

3.5 Feature-Driven Transfer Function Tuning

A critical point in the LiveSync workflow is the definition or the extraction of the feature of interest. The shape of the feature is important to determine a good viewpoint. For example, if a blood vessel is picked, the user should be provided with a view which shows the course of the vessel and does not cut through it. This section first describes how the parameters for the region growing are controlled, before focusing on how the transfer function can be tuned with knowledge about the extracted feature.

3.5.1 Feature Extraction

A natural choice to segment the object around the picked point is region growing. The picked point is transformed to a voxel in 3D as seed position. Huang and Ma [18] presented with RGVis a cost function which determines if a visited voxel during the region growing process is a member of the region or not. In this approach the neighborhood of the seed point is analyzed regarding scalar value and gradient magnitude distributions to initiate the parameters of the cost function. If the seed is located close to a boundary, the growing captures the boundary of the object, whereas a seed point within a homogeneous area results in a more compact growing process.

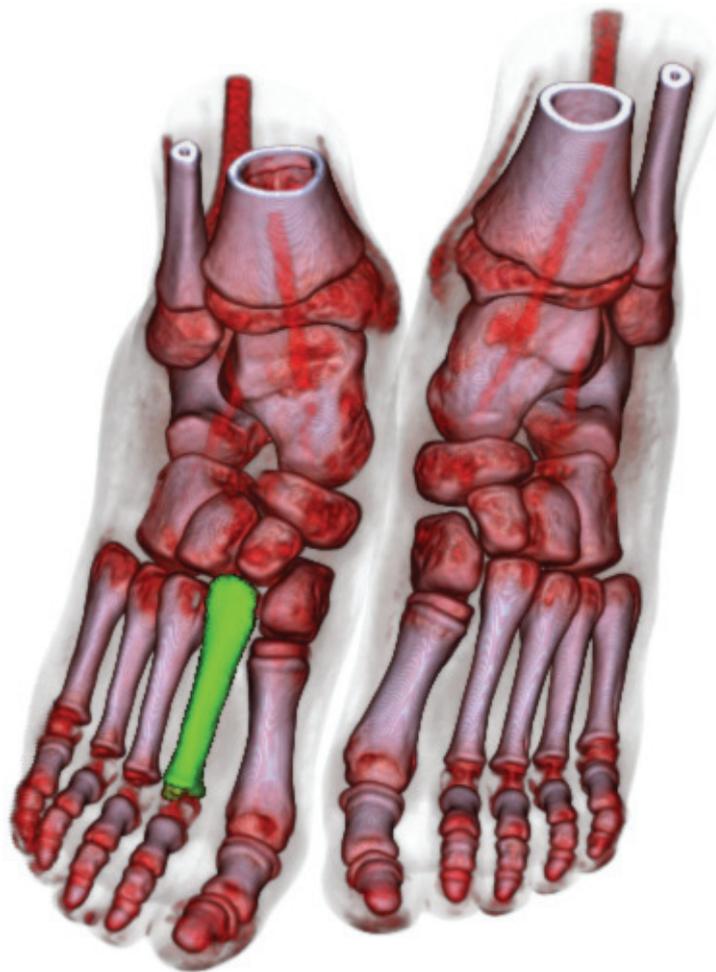


Figure 3.5: The optional rendering of the segmentation mask enhances the structure of interest.

3.5 Feature-Driven Transfer Function Tuning

In the initial LiveSync implementation (see Chapter 2) the growing process was limited to a fixed region of $32 \times 32 \times 32$ voxels. In the presented approach in this chapter the expansion is influenced by the spatial distribution of the points marked as region members so far. Growing progresses until the object-oriented bounding box (OOBB) of the included voxels reaches a certain limit. During the growing process the OOBB is updated at variable intervals to estimate how many more voxels may be added until the limit is reached. This strategy is superior to controlling the growing by the number of detected member voxels. The behavior of the growing process regarding the spreading of the points depends on various aspects. If just the boundary of a structure is segmented, a much lower number of voxels is sufficient to estimate the feature shape, than if the growing is performed in a very homogeneous area. Also with thin structures like blood vessels a small number of object voxels can already indicate the shape, whereas more voxels are needed for more compact areas.

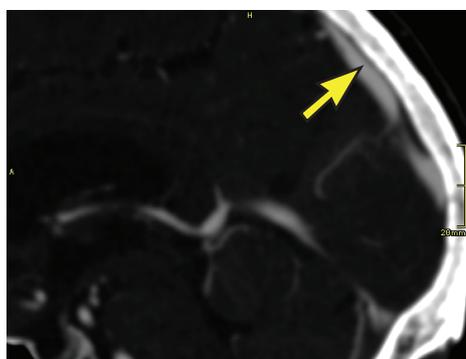
As soon as the region growing stops, a principal component analysis is performed on the member voxels to extract the three feature vectors and the corresponding eigenvalues. A metric of Westin et al. [67] is used to measure the local shape of the segmented feature from the relation between the eigenvalues. This metric allows to classify if a structure has an isotropic, a planar, or a linear shape. By labeling the segmented voxels with a segmentation mask, this information can be displayed in the live-synchronized volumetric view to highlight the structure of interest for a clear visual separation from its context. Figure 3.5 shows this enhanced visualization as result of a picking on the metatarsal in the slice view.

3.5.2 Transfer Function Tuning

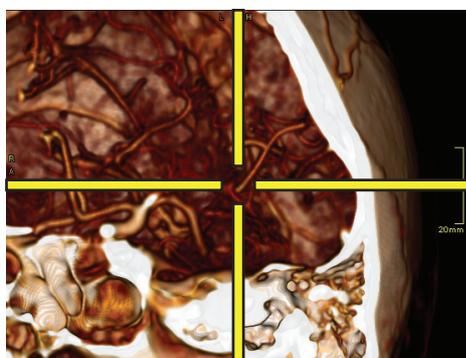
It is frequently stated, that the definition of a transfer function to assign color and opacity to the data values is a time-consuming and not very intuitive task. Even for the developer of a transfer-function editor it often ends in a tedious trial-and-error process to generate the desired result. Usually, the necessary effort increases with the dimensions of the transfer function.

LiveSync aims to generate the 3D visualization without any additional interaction apart from picking. Even the control of sliders for the transfer-function design as presented by Rezk-Salama et al. [43] might lead to an unwanted distraction in the diagnosis process. The presented system allows the user to choose from a set of predefined transfer function templates, which are very well-tailored for different types of examinations. The transfer function is defined by a color look-up table and a simple ramp which assigns zero opacity to scalar values from zero to the start of the slope, increasing opacity along the slope, and full opacity from the peak of the slope to the end of the scalar range.

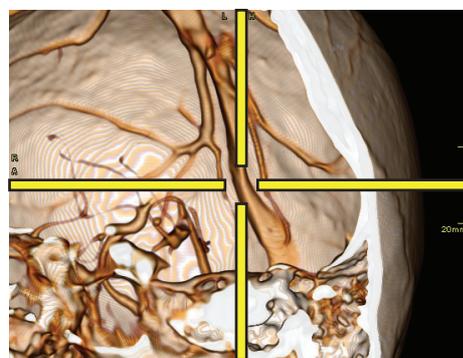
Inspired by the definition of a transfer function based on partial region growing in the work by Huang and Ma [18], knowledge about the distribution of scalar values within the extracted object is utilized to fine-tune an existing transfer function. This feature can be activated on demand and is especially helpful if the structure of interest is occluded to a large extent by densely surrounding objects. It is also helpful if a structure selected on the slice is hardly visible or not visibility at all applying the current opacity transfer function. The adjustment is based on the mean value and the standard deviation of the feature's scalar values which were collected by the region growing. The center of the slope is set to the mean value and the width is adjusted to three times the standard deviation.



(a)



(b)



(c)

Figure 3.6: Transfer function tuning. (a) The position of the picking on the slice. (b) The structure of interest is hidden to a large extent by blood vessels. (c) Allowing LiveSync to fine-tune the transfer function automatically leads to an unoccluded view of the sinus vein.

Figure 3.6 shows how the automatic setup of the ramp adjusts the transfer function. The structure of interest is visible to a much larger extent by making the

3.6 Interaction Modes

surrounding blood vessels transparent. The opposite case is shown in Figure 3.7 as the current transfer function setup does not assign enough opacity to the smaller blood vessels. By moving the center and changing the width of the ramp which defines the opacity transfer function, the object of interest is clearly visible in the volumetric view. Figure 3.8 shows the result of automatic transfer-function adjustment to provide a good view on the metatarsal.

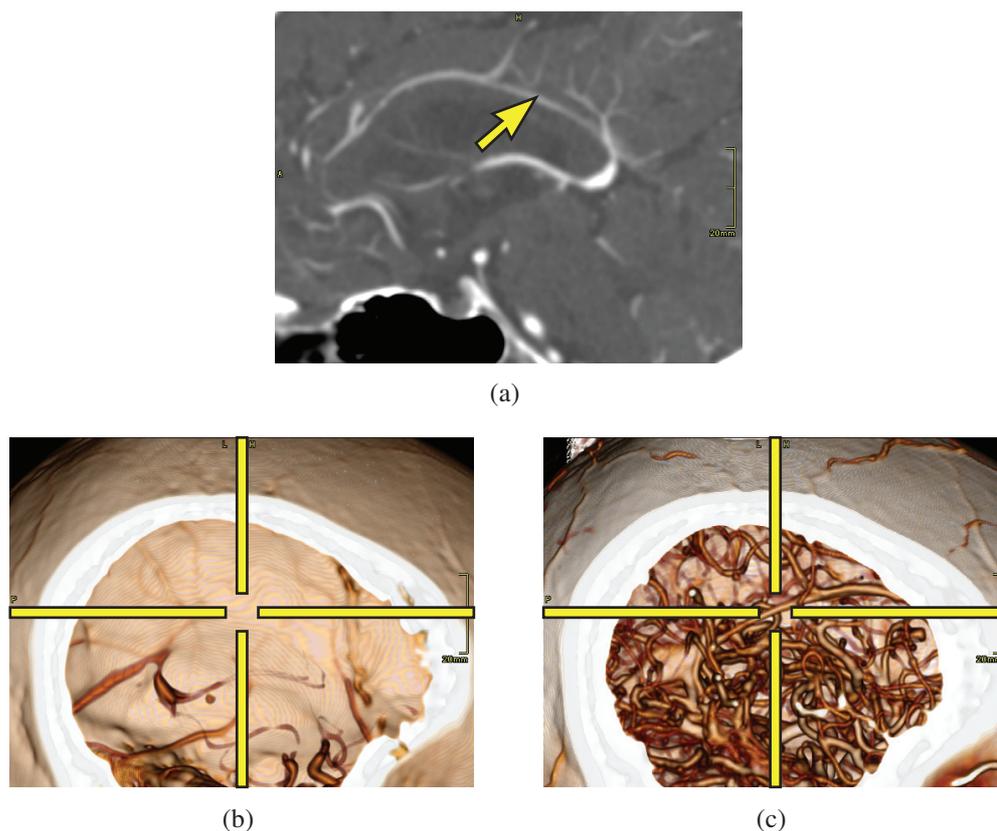


Figure 3.7: Transfer function tuning. (a) The position of the picking on the slice. (b) The structure of interest is not visible with the current setting of the opacity transfer function. (c) Allowing LiveSync to fine-tune the transfer function assigns opacity to the small vessels to make them visible.

3.6 Interaction Modes

The primary conceptual goal for the design of LiveSync was to interactively offer a synchronized 3D view. This shall be achieved without the need to manually

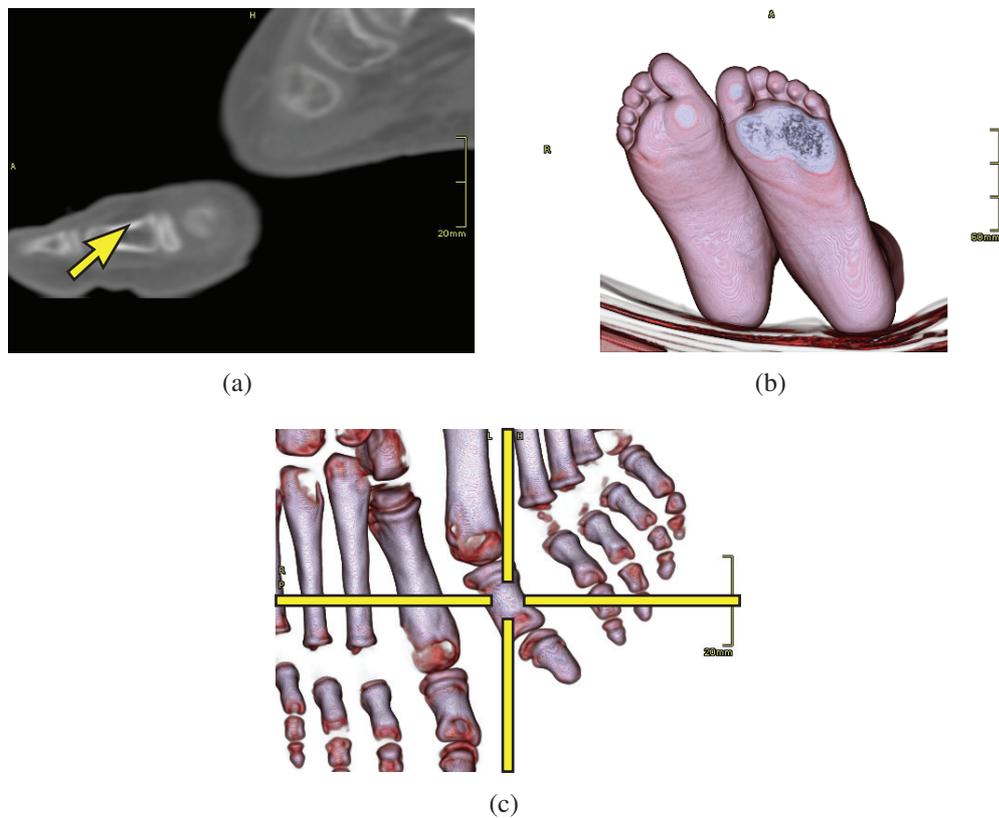


Figure 3.8: Transfer function tuning. (a) The position of the picking on the slice. (b) Initial view with a default transfer function. (c) LiveSync-adjusted transfer function which enables a good view on the metatarsal.

adjust all the parameters to generate an expressive view. Therefore a non-intrusive user interaction technique has to be implemented which lets the user focus on his primary analysis goal. In this section, the value and the limitation of the LiveSync picking interaction will be described and extended to integrate more knowledge about the intentions of the user.

3.6.1 LiveSync Mode

The evaluation with a radiology technician showed that the intuitive picking was sufficient to generate good results (see Section 2.8). One comment was that the handling and the mouse-over/hot-key interaction is very intuitive. However, it is rather difficult to derive in which part of an anatomical structure the user is interested. If the object is very small like a lung nodule or a polyp in the colon,

3.6 Interaction Modes

the user will be provided with a view which shows the whole feature within its three-dimensional context. If the volumetric extent of the feature is rather large, like e.g., the sinus vein, then the user might either be interested in getting a close-up view of the structure around the picked point, or a view which shows the whole vein is preferable. The current zoom factor of the slice view controls the zoom factor of the volumetric view, as it indicates the size of the anatomical structure the user is interested in. To integrate more knowledge about the size of the structure, the goal was to extend the functionality of LiveSync without introducing new interaction methods to the user.

3.6.2 LiveSync++ Growing Mode

Taking into consideration that region growing is performed to segment the structure of interest, the natural choice was to let the user control the size of the area of interest by keeping the hot-key pressed. This strategy allows to control the volumetric extent of the region growing. As mentioned in Section 3.5.1, OOBBs are computed during the growing process to estimate the extent of the initial growing. This is used to derive a reliable estimation of the local feature shape. By keeping the hot-key pressed, the region keeps on growing and the 3D view is updated in fixed time intervals. Within each interval, the growing process is continued until the diagonal length of the OOBB of the so far segmented set of voxels reaches the next predefined target size.

Integrating the knowledge about the user-indicated size of the area of interest, this can be utilized to improve the live-synchronized volumetric view regarding the demands of the user. After each growing step, the viewpoint is optimized, the clipping planes are adjusted and the zoom for the 3D view is changed. Figure 3.9 and Figure 3.10 show intermediate steps of the growing process. These views are provided to the users while they keep the hot-key pressed. Whenever a satisfying result is reached a release of the key stops the update of the volumetric view. There is a smooth zoom out provided for each step as the user wants to examine a larger area of interest. The zoom factor is computed by projecting the OOBB of the currently segmented set of voxels on the screen, and allowing it to cover a certain amount of the available display area, e.g., 50%.

The growing mode generates live-synchronized views at short update intervals. The user releases the hot-key or continues with the mouse movement to stop the update process. To prevent that the users miss and lose a view they like, a history mode is implemented. All live-synchronized views, including the intermediate views, are stored automatically in a view history. This is achieved by storing the adjusted parameters for the current view. The user can select a previous view from a list and store the best ones permanently.

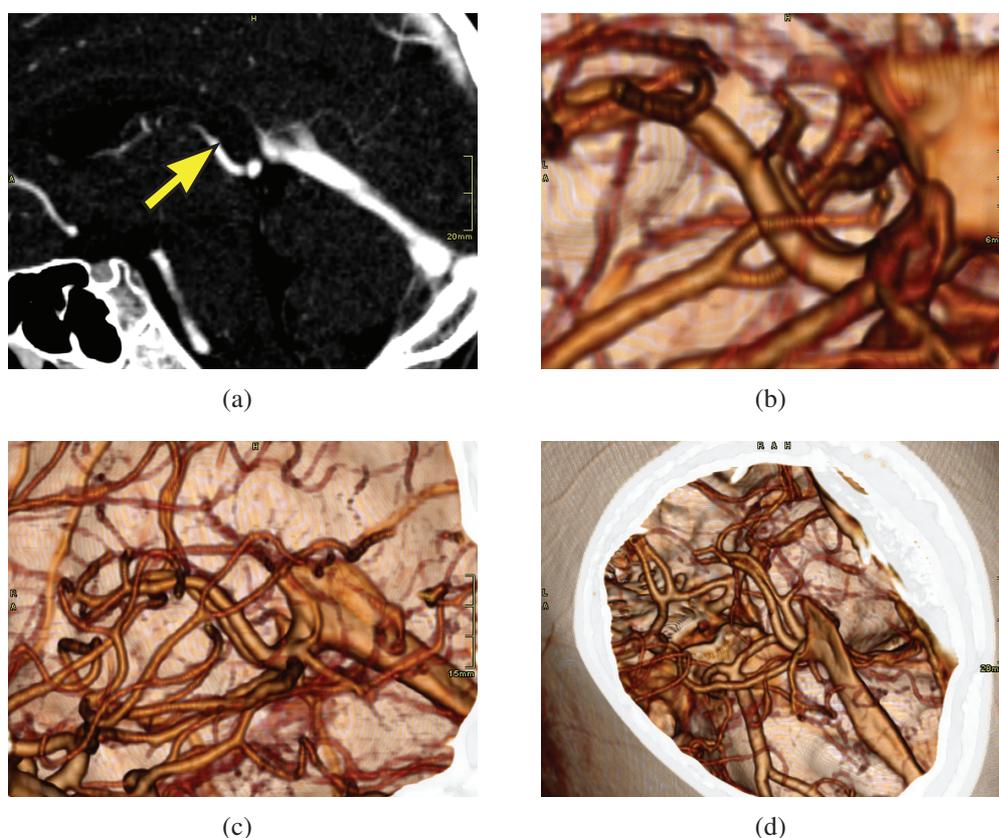


Figure 3.9: The LiveSync++ growing mode. (a) The position of the picking of a vessel on the slice. Initiated by this picking the 3D view is live synchronized at equal time intervals. (b)-(d) Images are provided for increasing OOB diagonal lengths.

3.7 Feature-Driven Clipping

Viewpoint selection is critical for generating the live-synchronized 3D view. In most cases interesting structures are surrounded by other tissues. These can be removed by the definition of an adjusted transfer function or, more straightforwardly, by the definition of a clip geometry. In this section, first the automatic setup of conventional clipping planes is explained before focusing on clipping planes which enable a higher degree of context preservation.

3.7.1 LiveSync Clipping Strategies

Typically, a medical workstation offers view-aligned and object-aligned clipping planes to interact with. As LiveSync aims to minimize the user's effort, the clip-

3.7 Feature-Driven Clipping

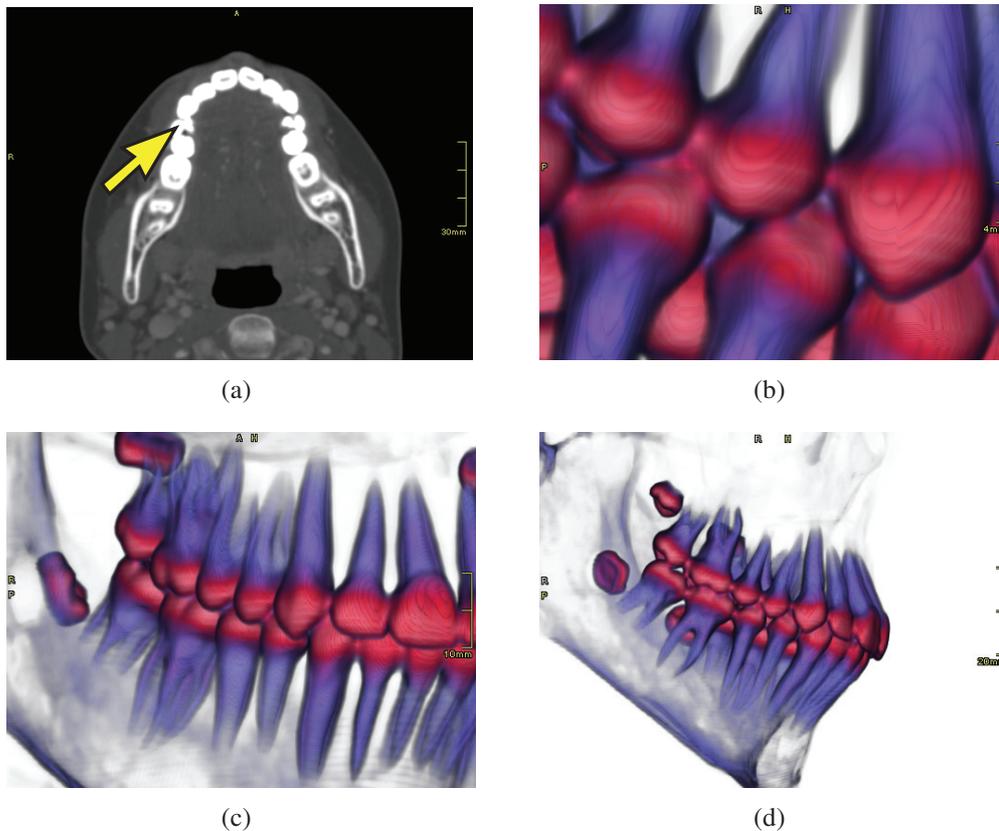


Figure 3.10: The LiveSync++ growing mode. (a) The position of the picking of a tooth on the slice. Initiated by this picking the 3D view is live synchronized at equal time intervals. (b)-(d) Images are provided for increasing OOB diagonal lengths.

ping is performed automatically to provide an unoccluded view on the structure of interest. Based on the visibility calculations described in Section 3.4.1, a position in the volume to place the clipping planes is computed. The plane is placed along the visibility ray where the structure of interest is not yet occluded by other structures. Initially only view-aligned clipping planes were supported. A plane which is orthogonal to the viewing direction is moved to the clipping position and the raycasting process starts where this plane intersects the volume. This strategy provides a good view on the structure of interest. Problems might occur if the user wants to perform further interactions with the 3D view after the live-synchronized view is generated. In the worst case, the structure of interest is clipped away by manual changes of the viewpoint.

Object-aligned clipping planes are moved along the three main patient axes. The six available clipping planes can clip from left, right, anterior, posterior, head

and feet. To decide which of these planes has to be set to allow an unoccluded view on the structure of interest, the plane which is most perpendicular to the viewing direction has to be identified. Then this plane is placed at the appropriate clipping position along the visibility ray. An object-aligned clipping plane remains at its position when the viewpoint is changed manually and cannot clip away the structure of interest unintentionally.

3.7.2 LiveSync++ Smooth Importance-Driven Clipping

The drawback of conventional clipping strategies is that they often remove anatomical context which does not occlude the object of interest. Viola et al. [64] introduced importance-driven volume rendering to remedy this problem. Their strategy uses segmentation information to generate a clipping region which follows the shape of the focus object. This approach relies on a complete segmentation of the structure in question. As our method relies on incomplete information in order to permit parameter-less interaction, the method presented by Viola et al. can lead to misleading results as a full segmentation is not available. To remedy this problem, while still providing more context information for the picked structure, we propose an extension of the context-preserving rendering model of Bruckner et al. [8]. The advantage of this approach is that it still follows the semantics of clipping planes, but retains salient features which are important for orientation.

The opacity α at each sample point along a viewing ray is computed in the following way:

$$\alpha = \begin{cases} \alpha_{tf} (|g| (1 - |n \cdot v|))^{(1-(d-c))^\kappa} & \text{if outside picked region} \\ \alpha_{tf} & \text{otherwise,} \end{cases}$$

where α_{tf} is the opacity as specified in the transfer function, $|g|$ is the gradient magnitude, n is the normalized gradient vector ($g/|g|$), v is the view vector, d is the distance along the viewing ray, c is the location of the view-aligned clipping plane, and κ is a parameter which controls the transition between clipped and unclipped regions. The range of d and c is $[0..1]$, where zero corresponds to the sample position closest to the eye point and one corresponds to the sample position farthest from the eye point. Opacity is only modulated in regions outside the focus structure identified by our local segmentation procedure. The adjustment is based on the idea of smooth clipping: Instead of having a sharp cut, it smoothly fades out structures that would be clipped otherwise. As the gradient magnitude indicates the *surfacedness* of a sample, its inclusion causes homogeneous regions to be suppressed. Additionally, since contours give a good indication of the shape

3.8 Performance and Qualitative Results

of a three-dimensional structure, while being visually sparse (i.e., they result in less occlusion), the gradient magnitude is multiplied by $(1 - |n \cdot v|)$. The parameter κ controls the transition between clipped and unclipped regions. In all our experiments, a value of 16 has proven to be most effective. Thus, this parameter does not require user adjustment and therefore does not introduce additional complexity.

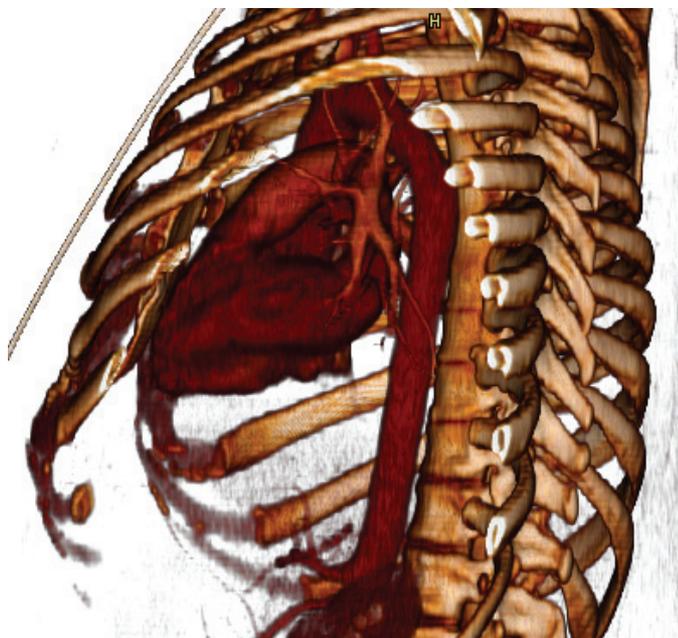
While we already achieve satisfactory results using this context-preserving clipping model, pathological cases might still cause an occlusion of the picked structure. Therefore, we additionally reduce the previously accumulated color and opacity along a ray on its first intersection with the picked region by a factor of $0.5(1 - \alpha_{tf})$. This lets the picked region always shine through – similar to ghosted views commonly found in medical illustrations – even if it is occluded by other structures. An example of this clipping technique is shown in Figure 3.11. The picking was performed on the aorta in the slice view. In Figure 3.11(a) there is a hard cut through the ribs, whereas the smooth clipping in Figure 3.11(b) shows the ribs which are not occluding the interesting part of the aorta with decreased opacity.

3.8 Performance and Qualitative Results

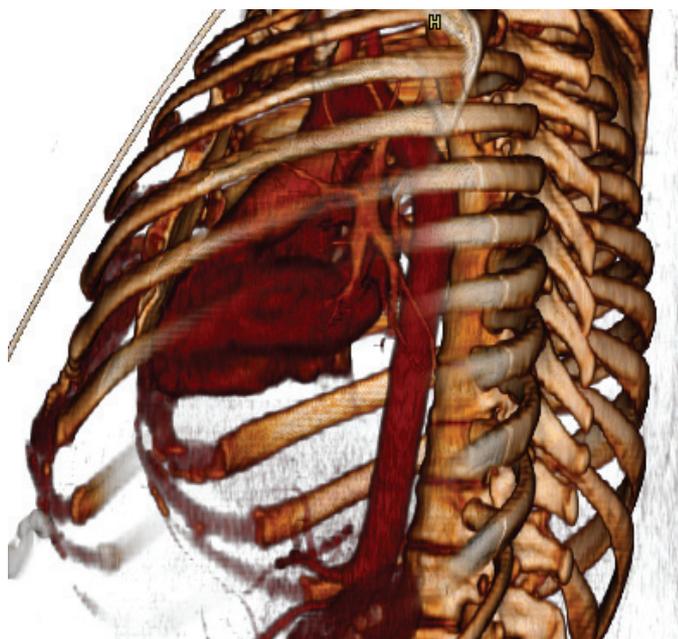
All computations which are done to derive the parameters to set up a live-synchronized 3D view can be performed interactively. The performance was measured on a PC configured with an *AMD Athlon 64 Dual Core Processor 4400+* and 2 GB of main memory. Compared to the approach presented in Chapter 2 where the region growing was fixed to a limit of $32 \times 32 \times 32$ voxels, the performance is more strongly influenced by the growing process. The size of the growing region is calculated adaptively. In the vast majority of the cases it is possible to stop the growing process much earlier than with a fixed growing region and the estimation of the local feature shape is more reliable.

Besides offering better view stability, the efficient and flexible sphere parameterization helps to eliminate the need for precomputing uniformly distributed points on the sphere, to reduce memory consumption, and to increase the overall performance. The LiveSync++-related computations are averaging between 50 ms and 100 ms by measuring with different data sets and picking on various structures. In growing mode as described in Section 3.6.2, the time to compute the intermediate views is in the same range as for the instant updates.

For the initial implementation of LiveSync an informal evaluation with an experienced radiology assistant has been performed (see Section 2.8). The feedback about the handling of the picking interaction and the live-synchronized views was very positive. The extensions presented in this chapter are designed based on



(a)



(b)

Figure 3.11: Clipping strategies comparison. (a) Conventional clipping leads to a hard cut through the ribs. (b) Smooth clipping preserves anatomical context by reducing the opacity of the ribs which are not occluding important parts of the aorta.

3.9 Conclusion

comments for possible improvements. As the LiveSync feature is integrated into a real-world medical workstation which is under development by our collaborating company partner, its functionality is frequently presented in demonstrations to radiologists. Their feedback is very positive as well, and they confirm the high practical value of this feature.

3.9 Conclusion

Live synchronization of the volumetric view from non-intrusive interaction with a slice view, is a powerful concept with the potential to improve the efficiency of diagnosis in clinical routine. In this chapter we presented several enhancements of the existing concept. Our goal was to integrate more knowledge about the anatomical structure the user is interested in to provide an expressive 3D visualization. As the key of the LiveSync concept is to keep the user interaction as small as possible, we did not want to increase the interaction effort or to introduce new interaction techniques.

A more efficient parameterization for the derived view input parameters was integrated to allow a hierarchical refinement of the search space for the best estimated viewpoint. This parameterization is utilized to provide a view with a high degree of view stability. An adaptive region growing was presented that allows to extract the part of the feature which is needed to estimate its local shape in a reliable way. Information about the scalar value distribution within the feature of interest is used to fine-tune the existing opacity transfer function. This helps to enhance the visibility of the object the user is interested in. The growing interaction mode integrates knowledge about the volumetric extent of the area of interest without increasing the user interaction. Especially an appropriate zoom factor for the 3D view can be derived from this information. A new clipping technique was integrated to increase the preservation of valuable anatomical context while guaranteeing an unoccluded view on the structure of interest.

The most successful people are those who are good at Plan B.

James Yorke

4

Contextual Picking of Volumetric Structures

This chapter is based on the following publication:

P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. Contextual picking of volumetric structures. In *Proceedings of the IEEE Pacific Visualization Symposium 2009 (to appear)*, 2009.

4.1 Introduction

Current computer hardware allows to display volume data with different rendering techniques simultaneously and in real time. For a certain medical diagnostic task in the clinical routine, a digital hanging protocol defines how the data is reformatted and arranged on the screen. For some examinations, e.g., in mammography, the hanging protocol is highly standardized, whereas, e.g., in vascular examinations, more often customized hanging protocols are preferred. Frequently, multi-planar reformatting (MPR) is the technique of choice to provide sectional renderings. With a curved planar reformation (CPR) the whole extent of a tubular structure is displayed within a single image. These 2D renderings are often accompanied by a direct volume rendering technique like raycasting. The examination of a structure in its three-dimensional setting often provides better insights into contextual information.

A typical hanging protocol arranges a side-by-side presentation of different views of the volumetric data. The physician performs different interactions during the examination of the data. Examples of frequently recurring interactions are scrolling through slices (2D images), zooming, panning, labeling, windowing (2D/3D images), or viewpoint selection and clipping (3D images). The synchronization of the different views is quite challenging because it is not trivial to determine if an interaction in one view leads to changes within another view. In

Chapter 2 and Chapter 3 the LiveSync interaction metaphor was presented as a solution for the live synchronization of a 2D slice view and a 3D volumetric view. The viewing parameters for the 3D view are derived automatically from a picking of the anatomical structure of interest on the 2D slice. In that work the viewing parameters are viewpoint, zoom factor, clipping planes, and transfer function setup.

This chapter presents a new approach to handle the picking of a structure which is directly performed on the 3D volumetric view in a context-sensitive way. Therefore, the first step is the identification of the 3D position of interest within the volumetric data. For each point on a 2D slice its exact 3D position can be calculated easily, whereas the picking of a point in the 3D volumetric view is not well defined. This is due to the fact, that for each pixel on the screen ray casting is performed, where opacity and color is accumulated along a ray from the eye point through the volume. Potentially each position along this ray might be the desired volumetric position. A simple solution to this problem is the definition of a *first-hit position* as the location along the ray where a certain opacity threshold is exceeded. For some cases this might be sufficient, but often a different, contextually-defined position is of interest. We propose ray-profile templates which are designed to represent anatomical structures like, e.g., a vessel, the aorta, the airway, or a vertebra. The ray profile which is calculated for each picking on the 3D volumetric view is then scanned and analyzed to find similarities to the defined ray-profile templates in a knowledge base.

To narrow down the number of anatomical structures which are of relevance for a certain examination, the presented method extracts contextual information automatically from the DICOM data. Details about the DICOM format are given in Section 1.1.1. Further context is provided by the setup of the medical workstation and the selection of clinical tools. For the examination of anatomical structures, frequently center points are of special interest. The *first-hit* solution cannot provide these positions because of self-occlusion or occlusion by other structures. Our proposed method can either return the first hit of the determined structure of interest or its center along the analyzed ray.

When the 3D position of interest is located, the next important step is the presentation of the result to the physician. This work depicts the highlighting of the results by synchronized 2D slice views as a straight-forward solution. Also the placement of labels at the appropriate 3D positions is demonstrated. For instance, the whole spine can be labeled by picking each vertebra directly in the 3D volumetric view. Further it is shown that approximate centerlines can be calculated interactively by tracing along tubular structures in the 3D view. These centerlines could be utilized to display CPR renderings of the structure or to guide a segmentation process.

4.2 Related Work

In medical visualization some techniques have been developed to ease the interaction with multiple views of a certain data set. The LiveSync interaction metaphor introduced in Chapter 2 is a first attempt to combine optimal viewpoint estimation and synchronized views for the visualization of medical volume data. An overview of selected current approaches to link different views on medical data is given in Section 1.2.

A lot of research has concentrated on the extraction of certain anatomical structures. Especially for the detection of curvilinear structures multi-scale filtering approaches are very popular. Vessel enhancement filters based on eigenvalue analysis of the Hessian matrix have been proposed, e.g., by Lorenz et al. [32], Sato et al. [47], and Frangi et al. [14]. Tek et al. [51] presented an approach which focuses on the segmentation of vessel cross sections. A single click inside the vessel on a slice initiates mean shift-based ray propagation to detect the boundary of the vessel.

Other anatomical tubular structures are, e.g., the airway and the aorta. Tschirren et al. [57] presented an airway segmentation algorithm based on fuzzy connectivity. Their method uses small adaptive regions of interest which follow the airway branches during the segmentation process. Kovács et al. [28] developed a system for automatic segmentation of the entire aorta without any user interaction for treatment planning of aortic dissections. The segmentation is based on a Hough transformation to detect the approximate circular shape of the aorta. To fit this shape more closely to the actual contour of the aortic lumen an elastic mass-spring deformable model is utilized. An interesting concept for the detection of tubular objects in medical imaging is the *Gradient Vector Flow (GVF) snake* introduced by Xu and Prince [72]. This method first calculates a field of forces (GVF forces) over the image domain. These forces drive the snake to fit to the boundaries of an object. Bauer and Bischof [3, 4] utilize the properties of the GVF for the detection of tubular objects and the extraction of curve skeletons, e.g., for virtual endoscopy. They argue that conventional tube detection or line filters which use local derivatives at multiple scales have problems with undesired diffusion of nearby objects. Their GVF-based method allows an edge-preserving diffusion of gradient information. Malik et al. [33] presented a rendering algorithm called feature peeling. They analyze peaks and valleys of intensity ray profiles for a given viewpoint to detect features inside the volume data. By classifying a number of feature layers it is possible to scroll through the layers to inspect various structures.

Sketch-based techniques are employed to classify and segment volume data by painting directly on the volume rendering. Owada et al. [40] presented a sketching interface which allows an intuitive segmentation of volumetric regions. The user

has to draw 2D strokes along the contour of the 3D target region to initiate a constraint segmentation process. Related to this approach, Chen et al. [11] applied sketching to seeded region growing for volume segmentation. Initially, the user specifies a region of interest by placing a closed free-form sketch on the volume rendering. Then a region of interest is extruded to facilitate the definition of seed points for the region growing. Ropinski et al. [44] proposed an interface for the design of 1D transfer functions which is based on the drawing of strokes on the volume rendered image. Features of interest can be identified by strokes which are close to their silhouettes. Based on a histogram analysis, component transfer functions are automatically generated for the identified features. The user can then decide which of the selected features should be integrated into the final transfer function.

4.3 Contextual Picking Overview

The overview of related works indicates that often highly specialized methods are used to detect various anatomical structures within the volumetric data. This chapter presents a generalized system which is not limited to a certain type of structure. An example illustrates one potential application area for the presented system: A frequently needed task during an orthopedic examination of the spine is the labeling of individual vertebrae. The vertebral column consists of 7 cervical, 12 thoracic, 5 lumbar, 5 sacral, and 3 to 5 coccygeal vertebrae. If the labeling is performed in 2D slice views only, quite some scrolling through the slices is involved to navigate to a meaningful 3D position where a single label is placed. A good position for the label is the center of the vertebral body. By picking on the vertebrae in the 3D view, contextual picking allows a convenient labeling. If the label for the first picked vertebra and the labeling direction is given, then a single contextual picking on the following vertebrae is sufficient to add the appropriate labels at a central position within the vertebral body.

Figure 4.1 shows the building blocks to achieve contextual picking. A knowledge base consists of a ray-profile library and contextual profiles. The ray-profile library holds ray-profile samples (intensities and gradient magnitudes) of various anatomical structures. A contextual profile for a certain structure bundles the needed information to react to a contextual picking operation. In the XML format it describes the following components: The type of the structure, a list of keywords, minimal and maximal extent of a structure, a representative mean ray profile built from the samples in the ray-profile library, and the default reaction to a picking operation. An initialization step is performed whenever a new data set is loaded into the workstation. The DICOM header as well as the workstation envi-

4.3 Contextual Picking Overview

Environment is analyzed to extract the relevant meta data for selecting the applicable contextual profiles.

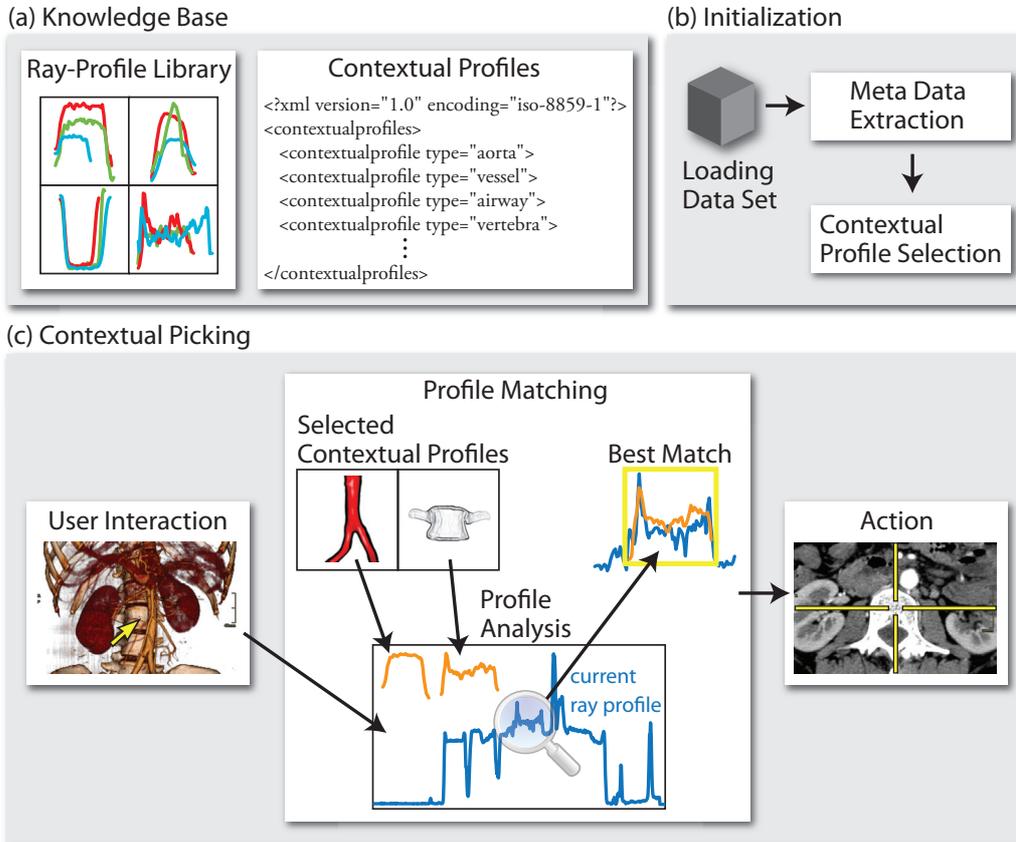


Figure 4.1: Contextual picking overview. (a) A knowledge base provides a library of ray profiles of anatomical structures together with an XML description of available contextual profiles. (b) An initialization step is performed during data set loading. Meta information is extracted to automatically select contextual profiles from the knowledge base. (c) For a picking on the volume rendered image, the current ray profile is analyzed to detect anatomical structures which are represented by the selected contextual profiles. If a good match is given, then this can be utilized for instance to highlight an interest point (e.g., the center of a vertebra) in the slice view.

Contextual picking is initiated by positioning the mouse cursor on the 3D view and pressing a hot-key. Whenever the physician picks on a structure of interest in the 3D view, the following steps are performed: First, information from the current picking, which includes the intensity and gradient magnitude values along the viewing ray, as well as accumulated opacities and positions of clipping planes, is collected. Second, the representative mean ray profiles of the selected contextual profiles are compared to the current ray profile. Finally if a good match is de-

tected, this result is utilized to highlight the anatomical structure of interest in an appropriate way, e.g., in MPR views.

4.4 Knowledge Base

Most approaches to automatically detect features in medical volume data need a considerable amount of user interaction to set up the needed parameters. Besides they are often very specialized on a certain type of anatomical structure as well as on a specific extent of the feature. Two important preconditions for the presented system are that it has to be as generic as possible and easy to extend. For these reasons a ray-profile library was set up which is supported by an easy-to-use interface for the generation of new ray-profile samples. Together with the contextual profiles this knowledge base provides all information to react to a contextual picking.

4.4.1 Ray-Profile Library

The ray-profile library is implemented as an XML file which stores ray-profile samples for various anatomical structures. A ray-profile sample for a certain structure consists of a sample id, a textual description, the spacing along the ray in mm, the extent of the structure in mm, and a list of intensity and gradient magnitude values. The system provides a convenient user interface to add new ray-profile samples to the ray-profile library. Figure 4.2 illustrates the generation of a ray-profile sample of a contrast enhanced blood vessel. After picking on a vessel in the 3D view, a ray profile is generated and displayed. A plot of the ray profile shows the intensity values (blue) and the values of the gradient magnitude (green) along the ray. To ease the selection of a ray-profile sample, the color transfer function for the corresponding scalar values is displayed in a small horizontal bar below the profile plot.

By dragging the mouse with the left button pressed a semi-transparent window is painted over the ray profile which represents a selection along the ray. By confirming this selection, a sample for a new or an already existing structure is written to an XML file. For a certain anatomical structure it is recommended to generate several representative samples. The intensity values of an anatomical structure can vary slightly, e.g., because of the patient's age or sex. Further the size of the structures varies because of the mentioned factors. A set of multiple samples in a contextual profile for example enables to detect vessels of a certain diameter range. The generation of ray-profile samples has to be done only once for each anatomical structure. They are added to the library by a domain expert. This step is not visible to the physician who just uses the contextual picking.

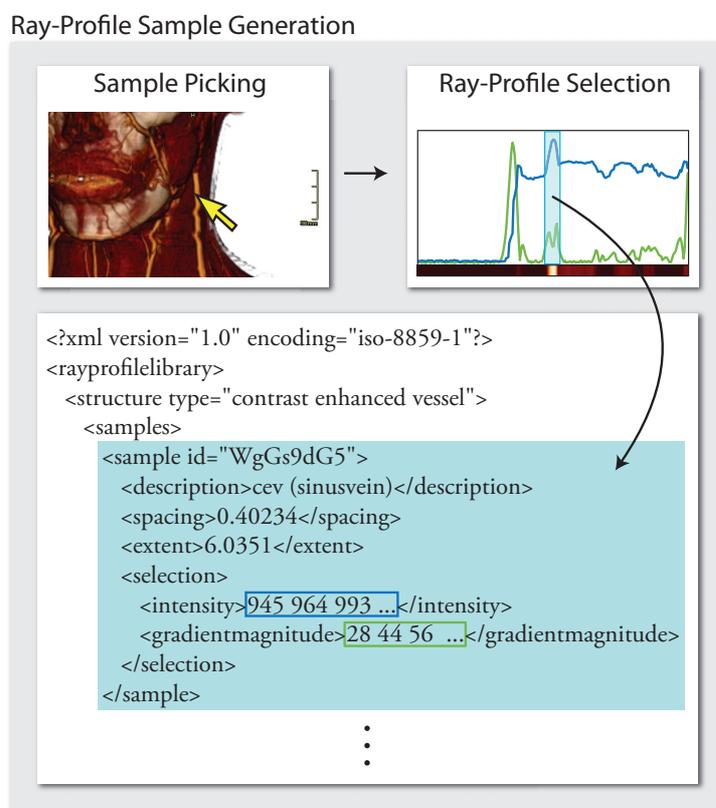


Figure 4.2: The generation of a ray-profile sample for a contrast enhanced vessel is initiated by picking on the structure in the 3D view. A ray profile is displayed which shows the intensity and the gradient magnitude values along the viewing ray. From a selected subset of this ray profile, which represents the picked vessel, a sample is generated and stored in the ray-profile library.

4.4.2 Contextual Profiles

The contextual profiles are stored in an XML file which contains all needed instructions to react to a contextual picking. Listing 4.1 shows the XML skeleton of the contextual profile for an anatomical structure. First of all it has a *type* entry which has to match with the corresponding *structure types* in the ray-profile library. A list of keywords describes the type of data sets and/or the setup of the medical workstation in which the defined structure is of interest. The *extent* defines a range of the typical extent of the anatomical structure. In the *mean-rayprofile* entry, a mean ray profile is stored which is generated from the available ray-profile samples in the ray-profile library. Finally, *return* defines which position will be returned (e.g., the center of the structure) and which default action shall be performed (e.g., highlighting of the obtained position in MPR views).

Listing 4.1: XML skeleton of a contextual profile

```
<contextualprofile type="">
  <keywords>...</keywords>
  <extent>...</extent>
  <meanrayprofile>
    <spacing>...</spacing>
    <intensity>...</intensity>
    <gradientmagnitude>
      ...
    </gradientmagnitude>
  </meanrayprofile>
  <return>
    <position>...</position>
    <reaction>...</reaction>
  </return>
</contextualprofile>
```

Mean Ray-Profile Generation

The generation of a mean ray profile is motivated by the variation of intensities and extents of anatomical structures. To obtain a good representation of a structure, all samples from the ray-profile library which correspond to the type of the contextual profile are collected. Figure 4.3 (left) shows three intensity ray-profile samples of the aorta which were all captured using different data sets. The similarities of the three samples are clearly visible. The intensity values are in a range of Hounsfield Units between about 880 and 1440 (shown on the y-axis of the plots) and the extent of the samples differs in a range between about 17.5 and 28.5 mm (shown on the x-axes of the plots). In general the samples start with a steep intensity ascent followed by a plateau and a steep intensity descent. The generated mean sample (right) shows the same characteristics with the advantage that outliers are filtered out. The algorithm for the generation of the mean ray profile first calculates the mean extent of the available samples. Then the ray-profile samples are scaled horizontally. The scaling factor is determined by dividing the mean extent by the extent of the current sample. Afterwards the mean of the corresponding intensity values is calculated at uniformly distributed positions along the x-axis. Analogous, a mean ray profile is also generated for the gradient magnitudes. Taking the mean extent of the ray-profile samples seems to be an appropriate approach because we assume that there is an approximate Gaussian distribution of the extents for multiple samples of a single structure.

4.5 Initialization

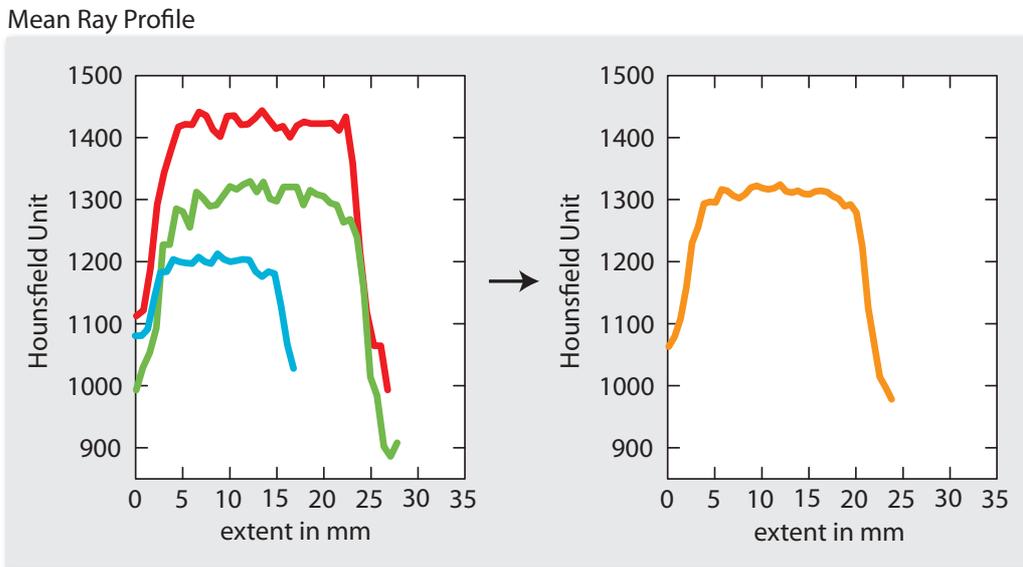


Figure 4.3: Three samples of the aorta (left) with varying extent and intensity range are collected from the ray-profile library. These samples are utilized to construct a representative mean ray profile (right).

The described algorithm for the generation of the mean ray profiles is well suited to preserve slopes which are characteristic for a certain structure. This is due to the fact that the steepness of slopes in shorter samples is decreased and the steepness of slopes in longer samples is increased by the horizontal scaling. Taking the mean of the intensity values results in a mean ray profile which represents the anatomical structure accordingly. Mean calculations have to be performed only when a new sample is added to the ray-profile library.

4.5 Initialization

The major challenge for a system which allows contextual picking of anatomical structures in the 3D view is to gain as much information as possible about the structure of interest. Typically for a certain type of examination only a small number of structures is relevant for the diagnosis. For instance in vascular examinations veins and arteries are of special interest whereas in an orthopedic examination the spine and bones are more important structures. There are two reasons to narrow down the number of target structures in a certain examination. First, our approach is based on a matching between ray-profile samples from a knowledge base and the ray profile extracted for the current picking. In volume data, structures are occluding each other. Thus the analysis of the current ray profile

with respect to all structures given in the contextual profiles, might lead to ambiguous results. For example a vessel in front of a bone is usually not relevant for an orthopedic examination. In such cases the contextual picking has to respond appropriately. Second, the analysis of the current ray profile through matching with a lot of structures of varying extents can lead to high computational costs.

As this work aims to minimize the overhead of the user interaction, the presented system extracts valuable meta information directly from the DICOM headers and the current setup of the medical workstation. The DICOM format is widely used in hospitals as a standard for handling, storing, printing and transmitting information in medical imaging. DICOM files contain the image pixel data of the scanned slices combined with header information. A *Patient* entity contains relevant data about the patient like name, age, and sex. Data concerning the examination like the name of the physician or the description of the examination is stored in a *Study* entity. A *Series* entity represents a logical unit of information about the imaging modality (*Equipment*) and information about the spatial relations of images within a series (*Frame of Reference*). Finally, the image itself and meta data about the image are stored in an *Image* entity. Grouping the image and the meta data prohibits that the image data gets separated from the meta information [53]. To narrow down the number of structures the physician might be interested in, we identified a suitable set of entries in the DICOM header:

- (0018, 0015) - Body Part Examined
- (0008, 1030) - Study Description
- (0008, 103E) - Series Description
- (0040, 0254) - Performed Procedure Step Description
- (0018, 1030) - Protocol Name

After a new data set is loaded, our system extracts the textual description which is stored for these entries. Entries which contain details about the patient like sex, age, and weight are further candidates which could be utilized to gather information about the size and the intensity range of a specific structure. The system also considers if the medical workstation is only used for certain examinations. For instance often a medical workstation with reduced functionality is available, e.g., as a vascular, orthopedic, cardiac, or mammography workstation.

The automatically extracted information is used to select the suitable contextual profiles from the knowledge base. Each contextual profile contains a list of classified keywords to decide if the represented structure is relevant for the currently loaded data set with the current setup of the workstation. Listing 4.2 shows exemplary keywords in a contextual vertebra profile.

4.6 Contextual Picking

Listing 4.2: Keywords of the contextual vertebra profile

```
<keywords>
  <strong>
    Workstation=Orthopedic
    BodyPartExamined=*SPINE ...
  </strong>
  <medium>
    BodyPartExamined=ABDOMEN ...
  </medium>
  <weak>
    BodyPartExamined=HIP ...
  </weak>
  <kickout>
    Workstation=Cardiac
    Workstation=Vascular ...
  </kickout>
</keywords>
```

A list of keywords is categorized into the classes *strong*, *medium*, *weak*, and *kickout*. The information which is extracted during loading the data set is compared with these entries to decide if a certain anatomical structure is relevant for the current examination. In the presented example a vertebra is strongly relevant when using an orthopedic workstation and/or when the examined body part is the spine. Within a cardiac or a vascular workstation a vertebra is typically not a structure of interest. With this approach it is possible to select from a ranked list the contextual profiles which are suitable in the given environment. The ranking is given by a comparison of the extracted information with the classified keywords.

4.6 Contextual Picking

After one or more contextual profiles are automatically selected, the physician can perform the picking directly on the 3D view and the system provides an immediate feedback. For each contextual picking, the current ray profile is analyzed to find close similarities in the selected contextual profiles. This analysis is done by a profile-matching algorithm which evaluates a cost function to measure the degree of similarity. Based on the outcome of this matching the respective action is taken. Three actions have been implemented so far. The default action is the highlighting of the center of the picked anatomical structure in MPR views. Further, contextual

picking is integrated into a spine labeling system to demonstrate its potential to place labels at meaningful 3D positions. Finally the system allows the calculation of approximate centerlines of picked tubular structures.

4.6.1 Profile Matching

The profile matching to detect the anatomical structure of interest is repeated for each of the selected contextual profiles. A contextual profile provides a mean ray profile which represents a structure together with the information about a minimal and a maximal extent of the structure. To allow the search for structures within this range along the current ray profile, a non-uniform scaling of the mean ray profiles is performed. In Figure 4.4 the minimal (left) and the maximal scaling (right) is shown for the mean ray profile (middle) of the aorta to cover a structure extent from about 16 to 32 mm. The scaling algorithm only performs a scaling on positions where the gradient magnitude is low. This avoids an undesirable change of the steepness of the major slopes.

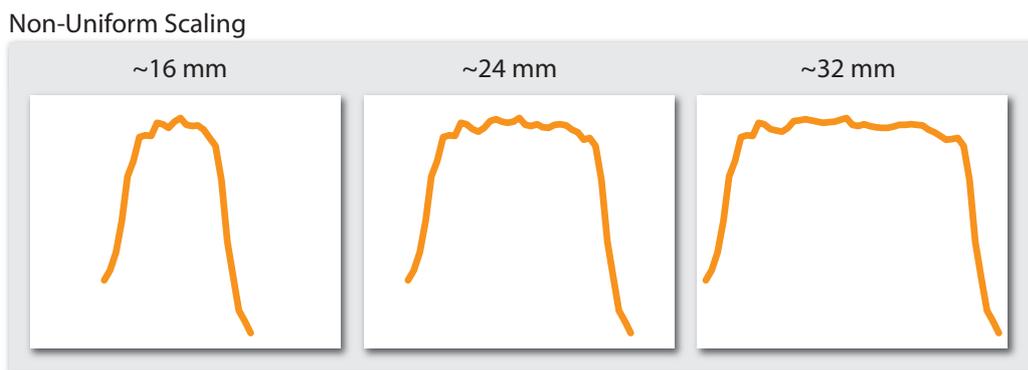


Figure 4.4: The mean ray profile of the aorta (middle) is non-uniformly scaled between a minimal (left) and a maximal range (right) to match the corresponding structure in the current ray profile.

The implementation of the profile matching is described in Algorithm 5. In this algorithm the mean ray profile and the current ray profile are represented by intensities as well as gradient magnitudes. Thus, the summed Euclidean distances are calculated for the corresponding intensities and gradient magnitudes. The matching algorithm detects the section along the current ray profile which is most similar to the mean ray profile at different scales. A length normalization of the fitting costs allows the comparison between different scales, as well as the comparison of responses between different contextual profiles. This is especially important if multiple contextual profiles are selected and thus, the ray profile is scanned for different structures. For instance, if an abdominal data set is loaded

4.6 Contextual Picking

Algorithm 5 Structure detection by profile matching

```
1: set  $pos = 0, width = 0, minCost = MAX\_VALUE$ 
2: for each scaling of the mean ray profile  $m$  do
3:   for each step  $s$  along the current ray profile  $c$  do
4:     if  $s + m.length < c.length$  then
5:       sum up squared Euclidean distances (SEDs) at the corresponding positions along  $m$  and  $c$  ( $s$  to  $s + m.length$ )
6:     end if
7:   end for
8:   set profile matching  $cost$  to  $sum\_of\_SEDs/m.length$ 
9:   if  $cost < minCost$  then
10:    set  $pos = s, width = m.length, minCost = cost$ 
11:   end if
12: end for
```

the contextual aorta-profile and the contextual vertebra-profile might be selected. To decide which structure is detected, a trade off between the cost function and the suitability of the contextual profile for the current environment has to be calculated. In the following, low costs of the profile matching are equivalent to a good or high response of a contextual profile.

Optimizations

To decrease ambiguities of the contextual profile response and to increase the performance of the system several optimizations were applied to the algorithm. First of all, the positions of the clipping planes are considered before the profile matching is performed. Only the part of the volume which is not clipped away is analyzed for matching regions. On the one hand this lowers the computational cost and on the other hand it reduces the chances of ambiguities. For example in the case of the contextual vessel profile, a vessel which is located in front of another vessel along the ray but clipped away could lead to unintended results. A second optimization which is applied for similar reasons takes the opacity transfer function into account. The analysis of the ray profile starts only at a sample position where a small opacity threshold is exceeded. From this position on the remaining part of the ray (if not clipped away) is considered. Third, the cost function for the profile matching is implemented so that the costs slightly increase along the ray. With this adjustment, structures which are closer to the viewer, e.g., vessels which are in front of other vessels, return lower costs. To apply this modified cost function line 8 in Algorithm 5 is replaced by the equation

$$cost = (sum_of_SEDs/m.length) \times (1 + 0.5 \times s/c.length), \quad (4.1)$$

where sum_of_SEDs are the summed-up squared Euclidean distances, $m.length$ is the sample size of the current matching profile, s is the current sample position along the ray profile, and $c.length$ is the total sample size of the ray profile. Our experiments have shown that the multiplication of the costs with a penalty factor of up to 1.5 for distant structures leads to good results for features which appear multiple times along the viewing ray. Despite of the improvements achieved by Equation 4.1 two overlapping target structures which are very close to each other might still be problematic. An example is if the aorta due to the current viewpoint of the volumetric view is right in front of some vertebrae and contextual profiles for both structures are active.

If the user is continuously tracing along the aorta it would be quite disturbing if the contextual vertebra profile and the contextual aorta profile alternate in generating the better response. The proposed solution to this problem is to deactivate all but one contextual profile as long as a continuous tracing goes on. As soon as the hot-key is released the other contextual profiles are re-activated. Finally, a default contextual profile is implemented which returns the *first-hit position*. The opacity is accumulated along the ray until a certain opacity threshold is reached. This contextual profile becomes active if the cost function of the profile matching returns too high values, which means that no contextual profiles are detected along the current ray profile.

4.6.2 Contextual Picking Action

The implemented default action to react to a contextual picking is the highlighting of the detected interest point in the MPR views. For each picking on the volumetric view, a three-dimensional position within the data set is computed. This position can be the center of the target structure along the current viewing ray or the first-hit position if no target structure is detected. The structure's center can be calculated easily as the start and the extent of the structure along the viewing ray is determined by the profile matching algorithm. To show the obtained position in the MPR views, the axial, coronal, and sagittal cross sections for this volumetric position are displayed. The centering of this position on the slices, as well as the overlay of crosshairs is used to highlight the target structure.

Another proposed action following a contextual picking is the labeling of anatomical structures. Often the labeling is performed in the slice views alone although the volumetric view can be very well suited for this task. For instance by utilizing the contextual vertebra profile, the physician gets more flexibility in the

4.7 Performance and Results

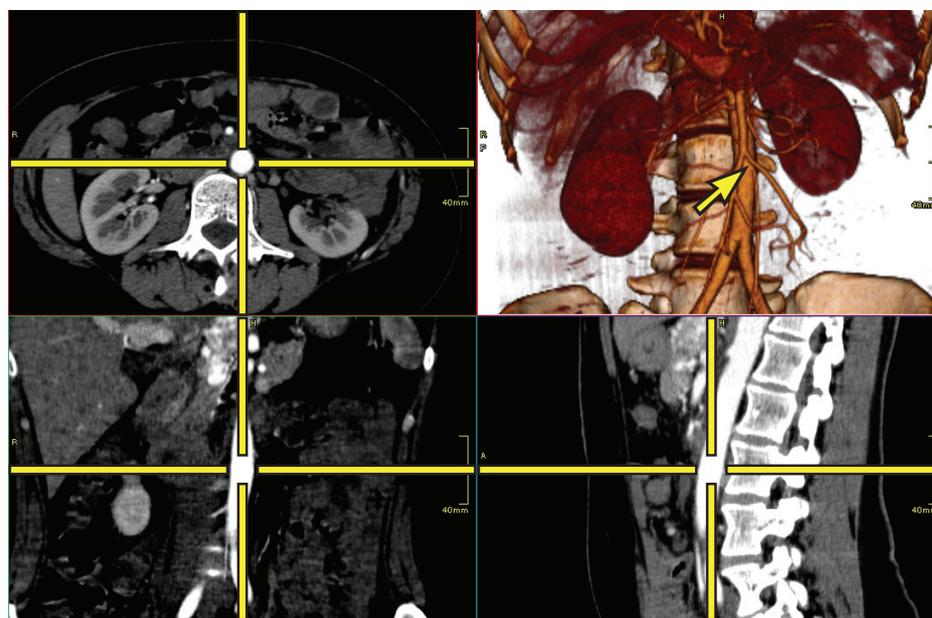
spine labeling process. The whole spine can be easily labeled in the 3D view. A single contextual picking on each vertebra determines the exact three-dimensional position of each label. Finally, the estimation of feature center points during the continuous tracing along a structure can be utilized to calculate approximate centerlines of tubular structures like vessels. If the obtained approximation is not sufficient it can be a helpful initial input for more accurate centerline detection algorithms.

4.7 Performance and Results

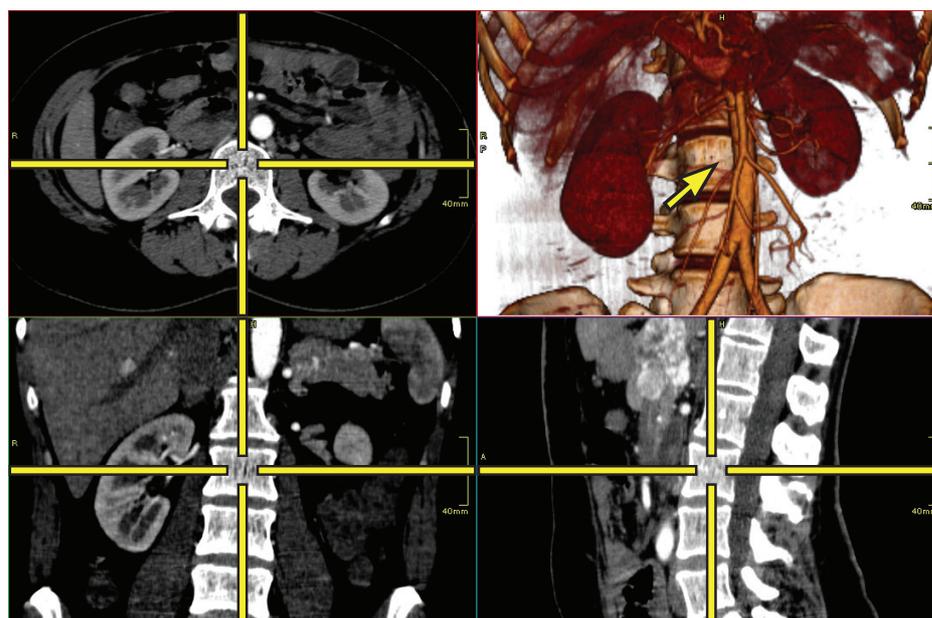
The contextual picking is implemented in Java and the contextual profiles as well as the ray-profile library are stored in the XML format. For parsing and manipulation of the XML files within the Java classes the JDOM API [52] is used. The contextual picking is integrated into a real-world medical workstation which is under development by our collaborating company partner. All contextual picking-related computations are performed interactively. For the generation of the ray-profile samples in our ray-profile library eight different CT data sets were used. Three samples were taken for each anatomical structure from suitable data sets. The data sets shown in the result images of this section are different from the data sets which were used to establish the ray-profile library.

In Figure 4.5 the contextual picking is illustrated for a thoracic-abdominal CT data set. After the data set is loaded into the workstation, meta information is extracted according to the description in Section 4.5. Based on this information, the contextual vertebra profile and the contextual aorta profile are selected automatically. When the aorta is picked as in Figure 4.5(a), the contextual aorta profile gives the best response. The detected center position of the aorta along the viewing ray is utilized to set up the axial, coronal, and sagittal MPR views. Four different views on the picked anatomical structure are provided to the physician. The picking of a close-by vertebra in Figure 4.5(b) leads to analogous results.

Figure 4.6(a) shows the path (indicated by the yellow line) of a continuous tracing along part of the aorta. If the contextual aorta profile and the contextual vertebra profile are active, the contextual vertebra profile has the better response at some positions along the trace path although the user is interested in the examination of the aorta. Figure 4.6(b) shows the resulting sagittal slice views when only the best response is taken into account. The vertebra is captured as the prominent structure at the positions 4 and 5. Whenever a continuous tracing is performed, the assumption can be made that the user currently examines a single anatomical structure. Thus, just a single contextual profile is active during the continuous tracing and all the others are deactivated temporarily. This leads to the results shown in Figure 4.6(c). Along the tracing path, the aorta is always captured as



(a)



(b)

Figure 4.5: Contextual picking on a thoracic-abdominal CT data set. The 3D position which is returned by the contextual profile with the best response is used to provide meaningful MPR views of the picked aorta (a) and the picked vertebra (b).

4.8 Conclusion

the prominent structure and jerky leaps in the MPR views between the aorta and a vertebra are avoided. The tracing along a tubular structure allows the computation of its approximate centerline.

Figure 4.7 depicts the result for the contextual picking of the airway in the 3D view (left) using a head CT data set. The contextual airway profile gives a better response than the contextual vessel profile which is also active. A highlighting of the corresponding position is performed in a 2D slice view (right). Occluding structures do not impede the detection of a central point within the airway.

Figure 4.8 shows some results when the contextual picking is integrated into a spine labeling tool. With this tool the user has to specify the label for the first picked vertebra and a labeling direction (head-to-feet or feet-to-head). Then a single picking on each vertebra leads to the results shown in Figure 4.8(a). Figure 4.8(b) shows the labeling from another viewpoint if just the first-hit position is taken for the 3D placement of the labels. If the placement is done by taking the positions determined by the contextual picking, the labels are in the center of the vertebral body as shown for the same viewpoint in Figure 4.8(c). The exact positions of the labels are depicted on the axial, coronal and sagittal slices for the first-hit approach in Figure 4.8(d) and for the contextual picking approach in Figure 4.8(e).

4.8 Conclusion

In this chapter we presented a novel method for the interactive identification of contextual interest points within volumetric data by picking on a direct volume rendered image. We built a knowledge base which holds characteristic ray-profile samples for different anatomical structures. New ray-profile samples can be added with an easy-to-use interface by domain experts. A contextual profile bundles information like for which kind of data sets it should be selected, the extent of the target structure, or the action which has to be performed if there is a high response to the contextual profile. Based on this knowledge base, the contextual profiles which are applicable in the current environment are selected automatically. We implemented a profile matching algorithm which analyzes the viewing ray for each contextual picking on the 3D volumetric view. It returns a position within the volume with the best response to one of the selected contextual profiles. Based on this result, certain actions can be performed interactively. In the simplest case, the obtained position is the center of the selected structure and is highlighted by crosshairs in multi-planar reformatted views. Because of the interactivity of the underlying computations, contextual picking is well suited to continuously trace the mouse pointer along volumetric structures. This allows to simultaneously examine the selected structure in MPR views. We have also demonstrated that

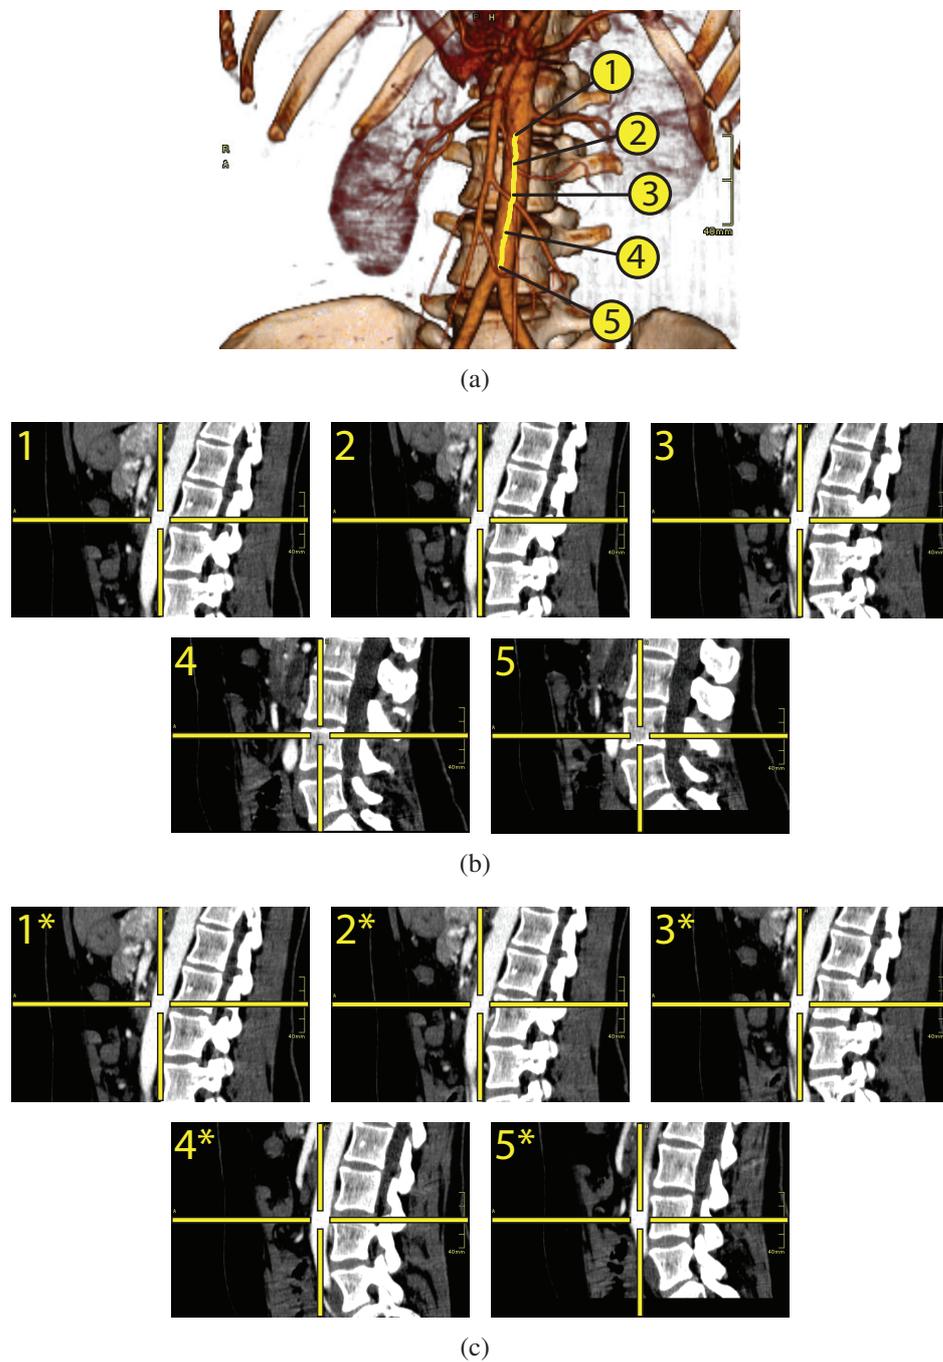


Figure 4.6: (a) Continuous tracing along the aorta. (b) This might lead to unintended responses for the vertebra if the contextual profiles of the aorta and the vertebra are selected. (c) The automatic temporary deactivation of the contextual vertebra profile during the tracing leads to a continuous capturing of the aorta.

4.8 Conclusion

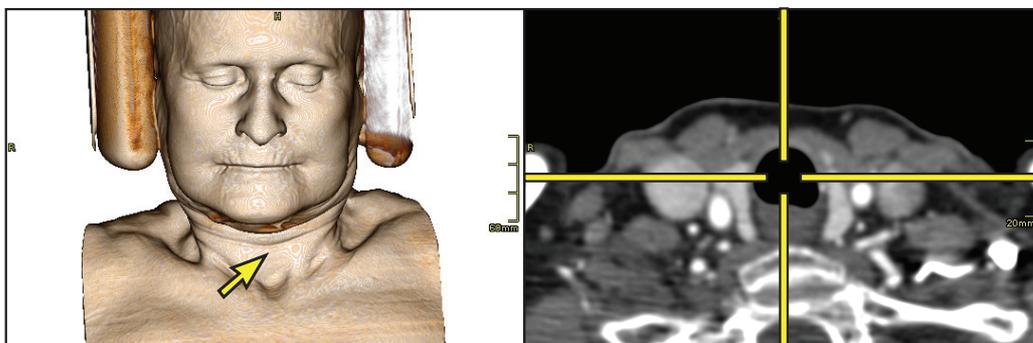


Figure 4.7: Contextual picking of the airway (left). The identified 3D position is used to provide a meaningful MPR view (right).

the contextual picking can be easily integrated into a conventional spine labeling framework to increase its flexibility.

Classifying structures only along ray profiles is not a limitation of the contextual picking framework. Other local classification schemes could be added as well. More research is necessary to investigate if this could help to further improve the detection of a structure's interest point. Until now, our method is to a certain degree dependent on the chosen viewpoint of the 3D view. While this could be a problem in some cases, there are often default viewpoints for 3D diagnostic examination procedures. An integration of shape-based methods would probably help to improve on this issue. The challenge thereby will be to ensure interactivity. Alternatively, multiple contextual profiles could be provided for different viewpoints on a structure. The integration of techniques to display uncertainty information about the currently detected structure is another interesting direction for further research.

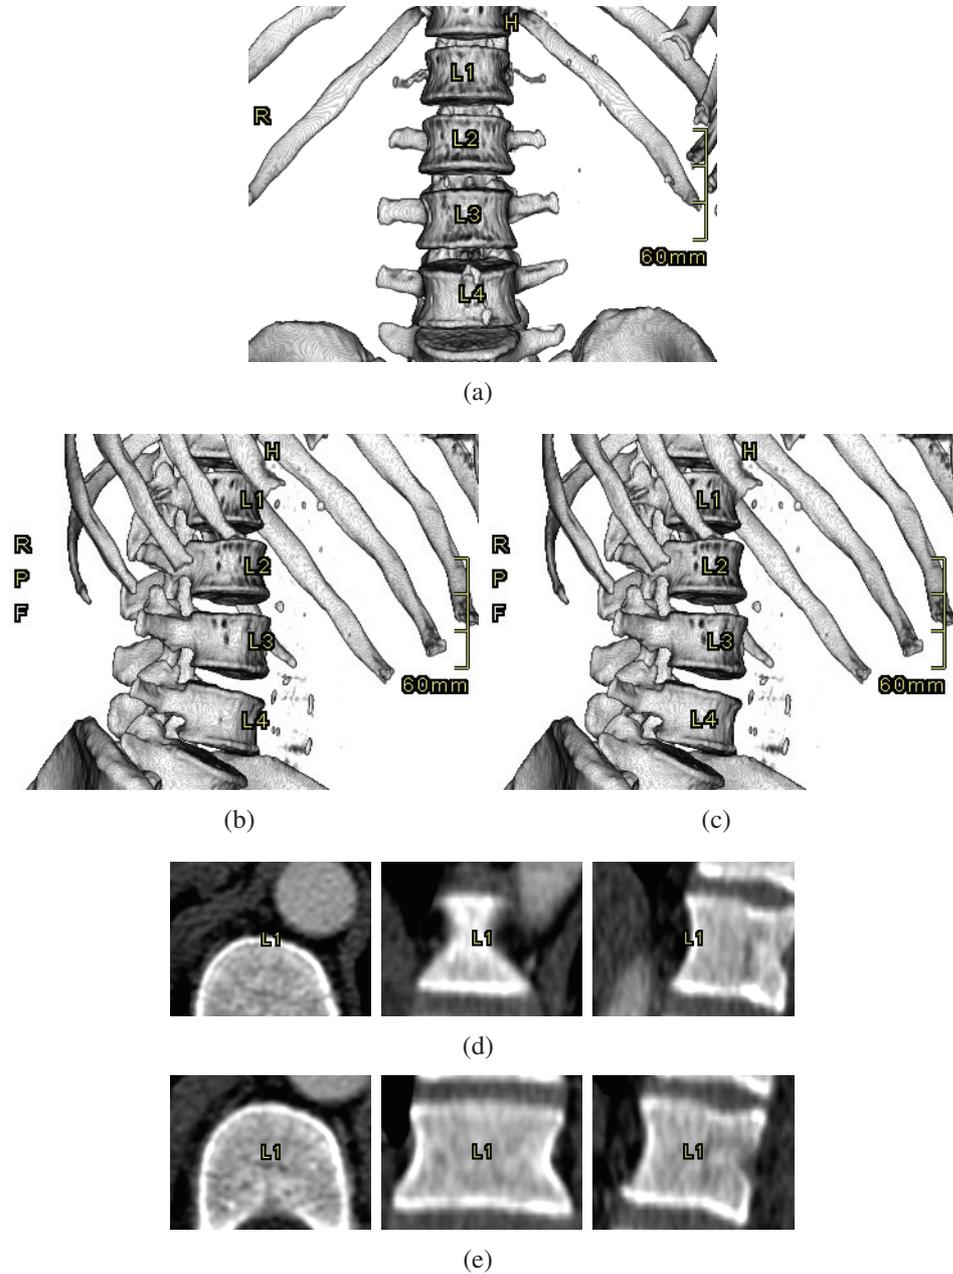


Figure 4.8: Labeling of lumbar vertebrae. (a) A single picking on each of the four lumbar vertebrae is performed for their labeling. (b) The result from another viewpoint when the first-hit position is taken for label placement. (c) The result for the same viewpoint when the contextual picking result is taken for label placement. (d) The exact labeling positions of L1 with just the first-hit approach. (e) The exact labeling positions of L1 with the contextual picking approach.

So long, and thanks for all the fish.

Douglas Adams

5

Summary and Conclusions

In current clinical applications for medical diagnosis the investigation of 2D slice images is still the dominant way of looking at the acquired data sets. The 3D renditions which can supply additional useful information are primarily used to get an overview. This thesis presents two techniques to incorporate 3D visualization into the diagnostic workflow.

The LiveSync interaction metaphor is a method for the automatic generation of meaningful 3D views based on a single picking on structures of interest in 2D slices. A synchronized 3D view helps to gain deeper insight into the medical data with minimal user interaction. From the picking on the slice several input parameters are derived automatically. The parameters which are considered for the estimation of a good viewpoint are patient orientation, viewpoint history, local object shape, and visibility. An automatic placement of the clipping plane is performed in a way which preserves valuable context while occlusion of the anatomical structure of interest is avoided. Further the zoom factor of the 3D view is adjusted to ensure that the structure covers a reasonable fraction of the screen space. An informal evaluation of this approach was performed where an experienced radiology technician was asked to generate diagnostically relevant 3D visualizations of pathologies in various data sets. For each case, this was first done manually and then with the LiveSync functionality. The conclusion of this evaluation was that the effort to localize the pathologies in 3D views was reduced significantly. In some of the cases small adjustments to the initial results were necessary to generate images with high diagnostic value.

In further work the parameterization of the view input parameters was optimized and new features were added to improve the usability of the system. The new parameterization scheme allowed a hierarchical refinement of the search space for the best estimated viewpoint and thus helped to improve the view stability of the proposed viewpoint. In addition, a technique was integrated which allows an automatic fine-tuning of the currently applied opacity transfer function based on the scalar value distribution within the structure of interest. Further added features are an interaction mode in which multiple views with different

zoom factors are provided within short time intervals as well as improved clipping strategies.

This work presented contextual picking as a second technique to improve the integration of 3D visualizations into the diagnostic workflow. Contextual picking targets the interactive identification of contextual interest points within volumetric data. The identification process is triggered by picking on a direct volume rendered image. In the following, the current ray profile is analyzed to detect structures which are similar to predefined templates in a knowledge base. Contextual picking can be utilized to highlight a structure in 2D slice views, to interactively calculate centerlines of tubular objects, or to place labels at contextually-defined volumetric positions.

Bibliography

- [1] AGFA HealthCare. IMPAX 6.3 client knowledge base. Available online at <http://osibsx1.opensystemimaging.com>, November 2008.
- [2] B. M. ter Haar Romeny. Image processing on diagnostic workstations. In E. Neri, D. Caramella, and C. Bartolozzi, editors, *Image Processing in Radiology: Current Applications*, pages 123–134. Springer, Berlin / Heidelberg, 2008.
- [3] C. Bauer and H. Bischof. Extracting curve skeletons from gray value images for virtual endoscopy. In *Proceedings of the 4th International Workshop on Medical Imaging and Augmented Reality 2008*, pages 393–402, 2008.
- [4] C. Bauer and H. Bischof. A novel approach for detection of tubular objects and its application to medical image analysis. In *Proceedings of the 30th DAGM Symposium on Pattern Recognition 2008*, pages 163–172, 2008.
- [5] V. Blanz, M. J. Tarr, and H. H. Bülthoff. What object attributes determine canonical views? *Perception*, 28(5):575–599, 1999.
- [6] U. D. Bordoloi and H.-W. Shen. View selection for volume rendering. In *Proceedings of IEEE Visualization 2005*, pages 487–494, 2005.
- [7] P. Bourke. Distributing points on a sphere. Available online at <http://local.wasp.uwa.edu.au/~pbourke/geometry/spherepoints/>, November 2008.
- [8] S. Bruckner, S. Grimm, A. Kanitsar, and M. E. Gröller. Illustrative context-preserving exploration of volume data. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1559–1569, 2006.
- [9] M.-Y. Chan, H. Qu, Y. Wu, and H. Zhou. Viewpoint selection for angiographic volume. In *Proceedings of the Second International Symposium on Visual Computing 2006*, pages 528–537, 2006.
- [10] R. N. Charette. Visualizing electronic health records with "Google-Earth for the body". *IEEE Spectrum Online*, January 2008. Available online at <http://www.spectrum.ieee.org/jan08/5854/>, November 2008.
- [11] H.-L. J. Chen, F. F. Samavati, M. C. Sousa, and J. R. Mitchell. Sketch-based volumetric seeded region growing. In *Proceedings of the Eurographics Workshop on Sketch-Based Interfaces and Modeling 2006*, pages 123–129, 2006.

- [12] K. Engel, M. Kraus, and T. Ertl. High-quality pre-integrated volume rendering using hardware-accelerated pixel shading. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Workshop on Graphics Hardware 2001*, pages 9–16, 2001.
- [13] S. Fleishman, D. Cohen-Or, and D. Lischinski. Automatic camera placement for image-based modeling. *Computer Graphics Forum*, 19(2):101–110, 2000.
- [14] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention – MICCAI 1998*, pages 130–137, 1998.
- [15] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622:759–771, 2005.
- [16] T. Götzelmann, P.-P. Vázquez, K. Hartmann, T. Germer, A. Nürnberger, and T. Strothotte. Mutual text-image queries. In *Proceedings of Spring Conference on Computer Graphics 2007*, pages 181–188, 2007.
- [17] M. O. Güld, M. Kohnen, D. Keysers, H. Schubert, B. Wein, J. Bredno, and T. M. Lehmann. Quality of DICOM header information for image categorization. In *Proceedings of the SPIE International Symposium on Medical Imaging 2002*, volume 4685, pages 280–287, 2002.
- [18] R. Huang and K.-L. Ma. RGVis: Region growing based techniques for volume visualization. In *Proceedings of the 11th Pacific Conference on Computer Graphics and Applications 2003*, pages 355–363, 2003.
- [19] R. Huang, H. Yu, K.-L. Ma, and O. Staadt. Automatic feature modeling techniques for volume segmentation applications. In *Proceedings of the IEEE/Eurographics International Symposium on Volume Graphics 2007*, pages 99–106, 2007.
- [20] A. Vilanova i Bartolí. *Visualization Techniques for Virtual Endoscopy*. PhD thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, 2001.
- [21] A. Vilanova i Bartolí, R. Wegenkittl, A. König, and M. E. Gröller. Nonlinear virtual colon unfolding. In *Proceedings of IEEE Visualization 2001*, pages 91–98, 2001.

BIBLIOGRAPHY

- [22] Institute of Medicine. *Key Capabilities of an Electronic Health Record System: Letter Report*. The National Academies Press, Washington, DC, 2003.
- [23] S. Jackson and R. Thomas. *Cross-Sectional Imaging Made Easy*. Churchill Livingstone, Edinburgh, 2004.
- [24] W. A. Kalender. *Computed Tomography*. Publicis MCD Verlag, München, 2000.
- [25] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and M. E. Gröller. CPR - curved planar reformation. In *Proceedings of IEEE Visualization 2002*, pages 37–44, 2002.
- [26] A. Kanitsar, R. Wegenkittl, D. Fleischmann, and M. E. Gröller. Advanced curved planar reformation: Flattening of vascular structures. In *Proceedings of IEEE Visualization 2003*, pages 43–50, 2003.
- [27] G. Kindlmann and J. W. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of the IEEE Symposium on Volume Visualization 1998*, pages 79–86, 1998.
- [28] T. Kovács, P. C. Cattin, H. Alkadhi, S. Wildermuth, and G. Székely. Automatic segmentation of the vessel lumen from 3D CTA images of aortic dissection. In *Proceedings of Bildverarbeitung für die Medizin 2006*, pages 161–165, 2006.
- [29] C. H. Lee, A. Varshney, and D. W. Jacobs. Mesh saliency. In *Proceedings of ACM SIGGRAPH 2005*, pages 659–666, 2005.
- [30] P. Leopardi. A partition of the unit sphere into regions of equal area and small diameter. *Electronic Transactions on Numerical Analysis*, 25:309–327, 2006.
- [31] W. E. Lorensen and H. E. Cline. Marching Cubes: A high resolution 3D surface construction algorithm. In *Proceedings of ACM SIGGRAPH 1987*, pages 163–169, 1987.
- [32] C. Lorenz, I.-C. Carlsen, T. M. Buzug, C. Fassnacht, and J. Weese. Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2D and 3D medical images. In *Proceedings of the First Joint Conference on Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery 1997*, pages 233–242. Springer, Berlin / Heidelberg, 1997.

- [33] M. M. Malik, T. Möller, and M. E. Gröller. Feature peeling. In *Proceedings of Graphics Interface 2007*, pages 273–280, 2007.
- [34] S. Marchesin, J.-M. Dischler, and C. Mongenet. Feature enhancement using locally adaptive volume rendering. In *Proceedings of the IEEE/EG International Symposium on Volume Graphics 2007*, pages 41–48, 2007.
- [35] G. L. Masala. Computer aided detection on mammography. In *Proceedings of World Academy of Science, Engineering and Technology 2006*, volume 15, pages 1–6, 2006.
- [36] Medical Exploration Toolkit. Available online at <http://www.metk.net>, November 2008.
- [37] L. Mroz, H. Hauser, and E. Gröller. Interactive high-quality maximum intensity projection. *Computer Graphics Forum*, 19(3):341–350, 2000.
- [38] K. Mühler, M. Neugebauer, C. Tietjen, and B. Preim. Viewpoint selection for intervention planning. In *Proceedings of the Eurographics/IEEE VGTC Symposium on Visualization 2007*, pages 267–274, 2007.
- [39] OsiriX. DICOM sample image sets website. Available online at <http://pubimage.hcuge.ch:8080/>, November 2008.
- [40] S. Owada, F. Nielsen, and T. Igarashi. Volume catcher. In *Proceedings of the ACM Symposium on Interactive 3D Graphics and Games 2005*, pages 111–116, 2005.
- [41] O. Polonsky, G. Patané, S. Biasotti, C. Gotsman, and M. Spagnuolo. What’s in an image: Towards the computation of the "best" view of an object. *The Visual Computer*, 21(8–10):840–847, 2005.
- [42] B. Preim and D. Bartz. *Visualization in Medicine: Theory, Algorithms, and Applications*. Morgan Kaufmann Publishers Inc., San Francisco, CA, 2007.
- [43] C. Rezk-Salama, M. Keller, and P. Kohlmann. High-level user interfaces for transfer function design with semantics. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1021–1028, 2006.
- [44] T. Ropinski, J. Praßni, F. Steinicke, and K. Hinrichs. Stroke-based transfer function design. In *Proceedings of the IEEE/EG Symposium on Volume and Point-Based Graphics 2008*, pages 41–48, 2008.

BIBLIOGRAPHY

- [45] G. Sakas, M. Grimm, and A. Savopoulos. Optimized maximum intensity projection. In *Proceedings of the 5th EUROGRAPHICS Workshop on Rendering Techniques 1995*, pages 55–63, 1995.
- [46] Y. Sato, N. Shiraga, S. Nakajima, S. Tamura, and R. Kikinis. Local maximum intensity projection (LMIP): A new rendering method for vascular visualization. *Journal of Computer Assisted Tomography*, 22(6):912–917, 1998.
- [47] Y. Sato, C.-F. Westin, A. Bhalerao, S. Nakajima, N. Shiraga, S. Tamura, and R. Kikinis. Tissue classification based on 3D local intensity structures for volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 6(2):160–180, 2000.
- [48] M. Sbert, D. Plemenos, M. Feixas, and F. González. Viewpoint quality: Measures and applications. In *Proceedings of the Eurographics Workshop on Computational Aesthetics in Graphics, Visualization and Imaging 2005*, pages 185–192, 2005.
- [49] N. H. Strickland. PACS (picture archiving and communication systems): filmless radiology. *Archives of Disease in Childhood*, 83(1):82–86, 2000.
- [50] S. Takahashi, I. Fujishiro, Y. Takeshima, and T. Nishita. A feature-driven approach to locating optimal viewpoints for volume visualization. In *Proceedings of IEEE Visualization 2005*, pages 495–502, 2005.
- [51] H. Tek, D. Comaniciu, and J. P. Williams. Vessel detection by mean-shift based ray propagation. In *Proceedings of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis 2001*, pages 228–235, 2001.
- [52] The JDOM API Project Website. Available online at <http://www.jdom.org>, November 2008.
- [53] The National Electrical Manufacturers Association (NEMA). The DICOM standard. Available online at <http://medical.nema.org/>, November 2008.
- [54] C. Tietjen, B. Meyer, S. Schlechtweg, B. Preim, I. Hertel, and G. Strauß. Enhancing slice-based visualizations of medical volume data. In *Proceedings of the Eurographics/IEEE VGTC Symposium on Visualization 2006*, pages 123–130, 2006.
- [55] C. Tietjen, K. Mühler, F. Ritter, O. Konrad, M. Hindennach, and B. Preim. METK - The Medical Exploration Toolkit. In *Proceedings of Bildverarbeitung für die Medizin 2008*, pages 407–411, 2008.

- [56] M. Tory and C. Swindells. Comparing ExoVis, orientation icon, and in-place 3D visualization techniques. In *Proceedings of Graphics Interface 2003*, pages 57–64, 2003.
- [57] J. Tschirren, E. A. Hoffman, G. McLennan, and M. Sonka. Intrathoracic airway trees: Segmentation and airway morphology analysis from low-dose CT scans. *IEEE Transactions on Medical Imaging*, 24(12):1529–1539, 2005.
- [58] F.-Y. Tzeng and K.-L. Ma. A cluster-space visual interface for arbitrary dimensional classification of volume data. In *Proceedings of the Joint Eurographics - IEEE TCVG Symposium on Visualization 2004*, pages 17–24, 2004.
- [59] J. M. van Verth and L. M. Bishop. *Essential Mathematics for Games and Interactive Applications: A Programmer's Guide*. Morgan Kaufmann Publishers Inc., San Francisco, CA, 2004.
- [60] P.-P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich. Viewpoint selection using viewpoint entropy. In *Proceedings of Vision, Modeling, and Visualization 2001*, pages 273–280, 2001.
- [61] P.-P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich. Automatic view selection using viewpoint entropy and its application to image-based modelling. *Computer Graphics Forum*, 22(4):689–700, 2003.
- [62] F. P. Vidal, F. Bello, K. W. Brodlie, N. W. John, D. Gould, R. Phillips, and N. J. Avis. Principles and applications of computer graphics in medicine. *Computer Graphics Forum*, 25(1):113–137, 2006.
- [63] I. Viola, M. Feixas, M. Sbert, and M. E. Gröller. Importance-driven focus of attention. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):933–940, 2006.
- [64] I. Viola, A. Kanitsar, and M. E. Gröller. Importance-driven feature enhancement in volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):408–418, 2005.
- [65] Virtual Terrain Project. Spherical textures. Available online at <http://www.vterrain.org/Textures/spherical.html>, November 2008.
- [66] M. Wan, Z. Liang, Q. Ke, L. Hong, I. Bitter, and A. Kaufman. Automatic centerline extraction for virtual colonoscopy. *IEEE Transactions on Medical Imaging*, 21(12):1450–1460, 2002.

BIBLIOGRAPHY

- [67] C.-F. Westin, A. Bhalerao, H. Knutsson, and R. Kikinis. Using local 3D structure for segmentation of bone from computer tomography images. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition 1997*, pages 794–800, 1997.
- [68] D. Williams. *Visualisation of Curved Tubular Structures in Medical Datasets: An Application to Virtual Colonoscopy*. PhD thesis, Centre for Health Informatics, City University, London, 2007.
- [69] D. Williams, S. Grimm, E. Coto, H. Hatzakis, and A. Roudsari. Volumetric CPR as an enhancement to virtual colonoscopy systems. *International Journal of Computer Aided Radiology and Surgery*, 2 (Suppl. 1):4–6, 2007.
- [70] D. Williams, S. Grimm, E. Coto, A. Roudsari, and H. Hatzakis. Volumetric curved planar reformation for virtual endoscopy. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):109–119, 2007.
- [71] E. Williams. Aviation formulary v1.44. Available online at <http://williams.best.vwh.net/avform.html>, November 2008.
- [72] C. Xu and J. L. Prince. Gradient vector flow: A new external force for snakes. In *Proceedings of the Conference on Computer Vision and Pattern Recognition 1997*, pages 66–71, 1997.
- [73] J. Zhou, M. Hinz, and K. D. Tönnies. Focal region-guided feature-based volume rendering. In *Proceedings of the First International Symposium on 3D Data Processing, Visualization, and Transmission 2002*, pages 87–90, 2002.



Curriculum Vitae

Contact Information

Name	Peter Kohlmann
Address	Spengergasse 31/5, 1050 Wien, Austria
Phone	+43 680 2019067
E-Mail	kohlmann@cg.tuwien.ac.at

Personal Details

Date of Birth	February 15 th , 1979
Place of Birth	Neumarkt i.d.OPf., Germany
Nationality	German
Gender	Male
Marital status	Single

Education

since 03/2006	Vienna University of Technology, Austria Doctoral studies (computer science). Working on the dissertation <i>LiveSync: Smart Linking of 2D and 3D Views in Medical Applications</i> at the Institute of Computer Graphics and Algorithms. Advisor: Prof. Eduard Gröller.
---------------	--

- 10/2000 - 02/2006 University of Siegen, Germany
Master's studies (applied computer science and media science). Successfully completed the master's thesis *Transfer Functions Design Based on Statistical Shape Models* at the Computer Graphics and Multimedia Systems Group. Advisors: Prof. Andreas Kolb and Dr. Christof Rezk Salama.
- 10/1999 - 09/2000 University of Regensburg, Germany
Studies of pedagogy.
- 09/1989 - 06/1998 Willibald-Gluck-Gymnasium, Neumarkt, Germany
(secondary school, diploma after 13th grade).
- 09/1985 - 07/1989 Grundschule Bräugasse, Neumarkt, Germany
(elementary school).

Employment History

- since 03/2006 Vienna University of Technology, Austria
Research assistant. Research and teaching at the Institute of Computer Graphics and Algorithms. (Project funding: AGFA HealthCare)
- 02/2005 - 10/2005 Siemens Corporate Research, Princeton, NJ, USA
Intern. Development of software for medical imaging and doing research for master's thesis at the Imaging Architectures Group.
- Summer 2004 Fraunhofer Gesellschaft, Sankt Augustin, Germany
Intern. Development of applications within the range of virtual environments for numerical and computer graphics simulations at the Institute for Media Communications, Competence Center Virtual Environments. (8 weeks)
- 10/2003 - 02/2004 University of Siegen, Germany
Teaching assistant. Tutor for the lecture *Introduction to Computer Science*.
- 10/2002 - 09/2003 University of Siegen, Germany
Student assistant. Design and implementation of interactive e-learning elements for teaching computer networks and distributed systems at the Didactics of Informatics and E-Learning Group.

Summer 2002	Systemhaus Schmaderer, Regensburg, Germany Intern. Working with content management systems and on preparations for ISO certification. (8 weeks)
Summer 2000	Evangelisches Bildungswerk, Regensburg, Germany Intern. Interviews and situation analysis for application and job interview, writing of application papers, and support of public relations. (6 weeks)
09/1998 - 09/1999	Regens-Wagner-Stiftungen, Neumarkt, Germany Alternative to military service. Caring for mentally and physically handicapped people.

Professional Activities

Referee for the International Journal of Computer Assisted Radiology and Surgery (IJCARS)

Referee for international conferences in the area of visualization and computer graphics:

Simulation and Visualization (SimVis), Vision, Modeling and Visualization (VMV), International Conference in Central Europe on Computer Graphics and Visualization (WSCG), Topology-Based Methods in Visualization (TopoInVis)

Publications

P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. Contextual picking of volumetric structures. In *Proceedings of the IEEE Pacific Visualization Symposium 2009 (to appear)*, 2009.

S. Bruckner, **P. Kohlmann**, A. Kanitsar, and M. E. Gröller. Integrating volume visualization techniques into medical applications. In *Proceedings of the IEEE International Symposium on Biomedical Imaging 2008*, pages 820–823, 2008.

P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. LiveSync++: Enhancements of an interaction metaphor. In *Proceedings of Graphics Interface 2008*, pages 81–88, 2008.

P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. LiveSync: Deformed viewing spheres for knowledge-based navigation. In *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1544–1551, 2007.

P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. Evaluation of a bricked volume layout for a medical workstation based on Java. In *Journal of WSCG*, 15(1-3):83–90, 2007.

C. Rezk-Salama, M. Keller, and **P. Kohlmann**. High-level user interfaces for transfer function design with semantics. In *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1021–1028, 2006.

C. Bastuck, T. Hambürger, T. Hof, M. Keller, **P. Kohlmann**, J. Mehnert, S. Nowak, C. Rezk-Salama, and A. Kolb. An open and extensible framework for visualization. In *Tagungsband der Informatiktage der Gesellschaft für Informatik 2005*, pages 151–154, 2005.

Patent Applications

Lead-inventor of patent applications in the field of medical visualization:

European *Method and apparatus for determining a position in an image, in particular a medical image*, EP08170289.6, Nov. 2008.

European/US/World *Method and apparatus for volume rendering of medical data sets*, EP07118075.6, Oct. 2007 – US12/134615, Jun. 2008 – PCT/EP2008/063526, Oct. 2008.

Skills

Languages German (native), English (fluent), French (basic knowledge)

Programming Java, C/C++, Open Inventor/Coin3D, OpenGL, QT, JavaScript, SQL