

LINKÖPING STUDIES IN SCIENCE AND TECHNOLOGY
DISSERTATIONS, NO. 1125

Efficient Medical Volume Visualization

An Approach Based on Domain Knowledge

Claes Lundström



Linköping University
INSTITUTE OF TECHNOLOGY

DEPARTMENT OF SCIENCE AND TECHNOLOGY
LINKÖPING UNIVERSITY, SE-601 74 NORRKÖPING, SWEDEN

NORRKÖPING 2007

**Efficient Medical Volume Visualization
- An Approach Based on Domain Knowledge**

© 2007 Claes Lundström

clalu@cmiv.liu.se

Center for Medical Image Science and Visualization
Linköping University Hospital, SE-581 85 Linköping, Sweden

ISBN 978-91-85831-10-4

ISSN 0345-7524

Online access: <http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-9561>

Printed in Sweden by LiU-Tryck, Linköping 2007.

Abstract

Direct Volume Rendering (DVR) is a visualization technique that has proved to be a very powerful tool in many scientific visualization applications. Diagnostic medical imaging is one such domain where DVR provides unprecedented possibilities for analysis of complex cases and highly efficient workflow. Due to limitations in conventional DVR methods and tools the full potential of DVR in the clinical environment has not been reached.

This thesis presents methods addressing four major challenges for DVR in clinical use. The foundation of all technical methods is the domain knowledge of the medical professional. The first challenge is the increasingly large data sets routinely produced in medical imaging today. To this end a multiresolution DVR pipeline is proposed, which dynamically prioritizes data according to the actual impact on the quality of rendered image to be reviewed. Using this prioritization the system can reduce the data requirements throughout the pipeline and provide both high performance and high visual quality.

Another problem addressed is how to achieve simple yet powerful interactive tissue classification in DVR. The methods presented define additional attributes that effectively capture readily available medical knowledge. The third area covered is tissue detection, which is also important to solve in order to improve efficiency and consistency of diagnostic image review. Histogram-based techniques that exploit spatial relations in the data to achieve accurate and robust tissue detection are presented in this thesis.

The final challenge is uncertainty visualization, which is very pertinent in clinical work for patient safety reasons. An animation method has been developed that automatically conveys feasible alternative renderings. The basis of this method is a probabilistic interpretation of the visualization parameters.

Several clinically relevant evaluations of the developed techniques have been performed demonstrating their usefulness. Although there is a clear focus on DVR and medical imaging, most of the methods provide similar benefits also for other visualization techniques and application domains.

Keywords: Scientific Visualization, Medical Imaging, Computer Graphics, Volume Rendering, Transfer Function, Level-of-detail, Fuzzy Classification, Uncertainty visualization, Virtual Autopsies.

Acknowledgments

Two people deserve extensive credit for making my research adventure such a rewarding journey. My first thanks go to my supervisor Anders Ynnerman, combining extreme competence in research and tutoring with being a great friend. Profound thanks also to another great friend, my enduring collaborator Patric Ljung, who has both provided a technical foundation for my work and been an untiring discussion partner.

Many more people have made significant contributions to my research. Sincere thanks to Anders Persson for neverending support and enthusiasm in our quest to solve clinical visualization problems. The research contributions from my co-supervisor Hans Knutsson have been much appreciated. Likewise, the numerous data sets provided by Petter Quick and Johan Kihlberg and the reviewing and proof-reading done by Matthew Cooper. Thanks to the other academic colleagues at CMIV and NVIS/VITA for providing an inspiring research environment. Thanks also to my other co-authors for the smooth collaborations: Örjan Smedby, Nils Dahlström, Torkel Brismar, Calle Winskog, and Ken Museth.

As an industrial PhD student, I have appreciated the consistent support from my part-time employer Sectra-Imtec and my colleagues there. Special thanks to Torbjörn Kronander for putting it all together in the first place.



It's hard to express my immense gratitude for having my wonderful wife Martina and our adorable children Axel, Hannes and Sixten by my side. Thanks for making my work possible and for reminding me what is truly important in life.

I am also very grateful for the inexhaustible love and support from my mother, father and brother.



This work has primarily been supported by the Swedish Research Council, grant 621-2003-6582. In addition, parts have been supported by the Swedish Research Council, grant 621-2001-2778 and the Swedish Foundation for Strategic Research through the Strategic Research Center MOVIII and grant A3 02:116.

Contents

1	Introduction	1
1.1	Medical visualization	2
1.1.1	Diagnostic workflow	2
1.1.2	Medical imaging data sets	2
1.1.3	Medical volume visualization	5
1.2	Direct Volume Rendering	6
1.2.1	Volumetric data	7
1.2.2	Volume rendering overview	7
1.2.3	Compositing	8
1.2.4	Transfer Functions	9
1.3	Direct Volume Rendering in clinical use	9
1.4	Contributions	11
2	Challenges in Medical Volume Rendering	13
2.1	Large data sets	13
2.1.1	Static data reduction	14
2.1.2	Dynamic data reduction	15
2.1.3	Multiresolution DVR	16
2.2	Interactive tissue classification	17
2.2.1	One-dimensional Transfer Functions	18
2.2.2	Multidimensional Transfer Functions	18
2.3	Tissue detection	19
2.3.1	Unsupervised tissue identification	20
2.3.2	Simplified TF design	20
2.4	Uncertainty visualization	21
2.4.1	Uncertainty types	21
2.4.2	Visual uncertainty representations	21
3	Efficient Medical Volume Visualization	23
3.1	Multiresolution visualization pipeline	23
3.1.1	Pipeline overview	24
3.1.2	Level-of-detail selection	25
3.1.3	Distortion metric	27
3.1.4	Interblock interpolation	28
3.1.5	Virtual autopsy application	29
3.2	Domain knowledge in interactive classification	30
3.2.1	Range weight	31
3.2.2	Tissue separation	32

3.2.3	Sorted histograms	34
3.3	Spatial coherence in histograms	36
3.3.1	α -histogram	36
3.3.2	Partial Range Histogram	38
3.3.3	Evaluations	39
3.4	Probabilistic animation	40
3.4.1	Probabilistic Transfer Functions	40
3.4.2	Probabilistic uncertainty animation	42
3.4.3	Probabilistic Transfer Functions revisited	43
4	Conclusions	47
4.1	Summarized contributions	47
4.2	Beyond the Transfer Function	48
4.3	Future work	49
	Bibliography	51
	Paper I: Transfer Function Based Adaptive Decompression for Volume Rendering of Large Medical Data Sets	57
	Paper II: Extending and Simplifying Transfer Function Design in Medical Volume Rendering Using Local Histograms	67
	Paper III: Standardized Volume Rendering for Magnetic Resonance Angiography Measurements in the Abdominal Aorta	77
	Paper IV: Multiresolution Interblock Interpolation in Direct Volume Rendering	87
	Paper V: The α-histogram: Using Spatial Coherence to Enhance Histograms and Transfer Function Design	97
	Paper VI: Multi-Dimensional Transfer Function Design Using Sorted Histograms	107
	Paper VII: Local histograms for design of Transfer Functions in Direct Volume Rendering	119
	Paper VIII: Full Body Virtual Autopsies Using a State-of-the-art Volume Rendering Pipeline	131
	Paper IX: Uncertainty Visualization in Medical Volume Rendering Using Probabilistic Animation	141

List of Papers

This thesis is based on the following papers.

- I** Patric Ljung, Claes Lundström, Anders Ynnerman and Ken Museth. Transfer Function Based Adaptive Decompression for Volume Rendering of Large Medical Data Sets. In *Proceedings of IEEE/ACM Symposium on Volume Visualization 2004*. Austin, USA. 2004.
- II** Claes Lundström, Patric Ljung and Anders Ynnerman. Extending and Simplifying Transfer Function Design in Medical Volume Rendering Using Local Histograms. In *Proceedings EuroGraphics/IEEE Symposium on Visualization 2005*. Leeds, UK. 2005
- III** Anders Persson, Torkel Brismar, Claes Lundström, Nils Dahlström, Fredrik Othberg, and Örjan Smedby. Standardized Volume Rendering for Magnetic Resonance Angiography Measurements in the Abdominal Aorta. In *Acta Radiologica*, vol. 47, no. 2. 2006.
- IV** Patric Ljung, Claes Lundström and Anders Ynnerman. Multiresolution Interblock Interpolation in Direct Volume Rendering. In *Proceedings of Eurographics/IEEE Symposium on Visualization 2006*. Lisbon, Portugal. 2006.
- V** Claes Lundström, Anders Ynnerman, Patric Ljung, Anders Persson and Hans Knutsson. The α -histogram: Using Spatial Coherence to Enhance Histograms and Transfer Function Design. In *Proceedings Eurographics/IEEE Symposium on Visualization 2006*. Lisbon, Portugal. 2006.
- VI** Claes Lundström, Patric Ljung and Anders Ynnerman. Multi-Dimensional Transfer Function Design Using Sorted Histograms. In *Proceedings Eurographics/IEEE International Workshop on Volume Graphics 2006*. Boston, USA. 2006.
- VII** Claes Lundström, Patric Ljung and Anders Ynnerman. Local histograms for design of Transfer Functions in Direct Volume Rendering. In *IEEE Transactions on Visualization and Computer Graphics*. 2006.
- VIII** Patric Ljung, Calle Winskog, Anders Persson, Claes Lundström and Anders Ynnerman. Full Body Virtual Autopsies using a State-of-the-art Volume Rendering Pipeline. In *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization 2006)*. Baltimore, USA. 2006.
- IX** Claes Lundström, Patric Ljung, Anders Persson and Anders Ynnerman. Uncertainty Visualization in Medical Volume Rendering Using Probabilistic Animation. To appear in *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization 2007)*.

Chapter 1

Introduction

Science is heavily dependent on the analysis of data produced in experiments and measurements. The data sets are of little use, however, unless they are presented in a form perceivable for a human. *Visualization* is defined as the art and science of constructing perceivable stimuli to create insight about data for a human observer [FR94]. Visual impressions are the typical stimuli in question but also audio and touch, as well as combinations of the three are used for the same purpose.

A cornerstone of image-based visualization is the extraordinary capacity of the human visual system to analyze data. Structures and relations are instantly identified even if the data is fuzzy and incomplete. In visualization the human interaction with the presented images is seen as crucial. This emphasis on retaining a *human-in-the-loop* setup in the data analysis separates visualization from other fields, where the ultimate goal can be to replace the human interaction.

The impact of visualization in society is steadily increasing and the health care domain is a prime example. Medical imaging is fundamental for health care since the depiction of the body interior is crucial for the diagnosis of countless diseases and injuries. With this motivation, vast research and industry efforts have been put into the development of imaging devices scanning the patients and producing high-precision measurement data. Capturing the data is only the first step, then visualization is the essential link that presents this data to the physician as the basis for the diagnostic assessment.

Health care is also an area where further substantial benefits can be drawn from technical advances in the visualization field. There are strong demands on providing increasingly advanced patient care at a low cost. In medical imaging this translates to producing high-quality assessments with minimal amount of work, which is exactly what visualization methods aim to provide. In particular, three-dimensional visualizations show great potential for increasing both quality and efficiency of diagnostic work. This potential has not been fully realized, largely due to limitations of the existing techniques when applied in the clinical routine. With the objective to overcome some of these limitations, the methods presented in this thesis embed medical domain knowledge in novel technical solutions.

The overall research topic of this thesis is scientific visualization within the medical domain. The focus is on volumetric medical data sets, visualized with a technique called Direct Volume Rendering (DVR). This first chapter is meant to be an introduction to the domain of the thesis, describing medical visualization in clinical practice as well as the technical essentials of DVR. In chapter 2 a number of central challenges

for DVR in the clinical context are identified and relevant previous research efforts are described. The research contributions of this thesis, addressing these challenges, are then presented in chapter 3. Finally, concluding remarks are given in chapter 4.

1.1 Medical visualization

When Wilhelm Conrad Röntgen discovered x-rays in 1895 [Rön95], imaging of the interior human anatomy quickly became an important part of health care. The radiology department or clinic has for many decades been a central part of the hospitals' organization and the health care workflow. The following sections describe the data sets produced in diagnostic imaging and how volume visualization is currently being performed.

1.1.1 Diagnostic workflow

The workflow for diagnostic imaging at a hospital typically originates at a department dealing with a specific group of diseases, such as oncology or orthopedics, having the main responsibility for the patient. In a vast range of situations, an imaging examination is necessary in order to determine the appropriate treatment. The responsible physician then sends a request to the radiology department, including the diagnostic question to be answered. Based on the question a number of imaging studies are performed. The images are typically produced by a technician/radiographer according to predefined protocols. A radiologist, i.e., a physician specialized in radiology, then reviews the images and writes a report on the findings. Finally, the report is sent to the referring physician who uses it as a basis for the patient's treatment.

The images have traditionally been in the form of plastic films but there has been a strong digitization trend over the last 15 years. Today, virtually all radiology examinations performed in Sweden are digital. Many large hospitals throughout the world are also film-free but the penetration of digitization has not yet been as strong as in Scandinavia. At a film-free hospital there is a digital image management system known as a Picture Archiving and Communication System (PACS). The PACS handles display, storage and distribution of the digital images, replacing light cabinets and film archives. There are many benefits driving the digitization process: unlimited access to images across the hospital, less risk of losing images, no need for developing fluids or space-consuming archives, etc. Reviewing the images by means of computer software also provides unprecedented opportunities to interact with the data.

The diagnostic review is typically performed on a PACS workstation. Routinely used tools to interact with the images include grayscale windowing (brightness and contrast adjustments), zooming, panning, and measurements. Comparisons to prior examinations, if there are any, is another crucial feature to be provided by the PACS workstation.

1.1.2 Medical imaging data sets

There are many types of imaging examinations performed at a hospital, primarily at the radiology department. Many techniques employ x-rays and the resulting measurement values correspond to the x-ray attenuation of the tissues. There are digital 2D imaging methods resembling traditional film-based radiography, such as Computed Radiogra-



Figure 1.1: An example of a medical data set from CT, a slice from a head scan.

phy (CR) and Direct Radiography (DR), where the x-ray “shadow” of the anatomy is registered.

The x-ray technique relevant for this thesis is Computed Tomography (CT), producing volumetric data sets of the patient. The x-ray source and detector are rotated around and moved along the patient, measuring the intensity of x-rays passing through the body as this spiral progresses. The measurement data are then reconstructed into attenuation values on a rectilinear 3D grid using an algorithm based on the Radon transform. In this case, the tissues do not “shadow” each other, instead, each value describes the attenuation at a single point in space as seen in figure 1.1.

CT scanners have developed tremendously over the last decade in terms of higher resolution and decreased acquisition time. The most recent development is dual-energy CT, where two different x-ray energies can be used simultaneously, providing more depiction possibilities. As a result of this progress, there is a strong trend to move many types of examinations to the CT domain. A drawback of all x-ray techniques is the dangers of radiation dose, which is limiting the transition to CT.

Magnetic Resonance (MR) imaging is based on a completely different technique. Here the principle of nuclear magnetic resonance is used. A strong magnetic field is used to align the spin of hydrogen nuclei (protons) in the body. Then a radio-frequency pulse matching the nuclear resonance frequency of protons causes the spins to synchronize. As the pulse is removed, different relaxation times are measured, i.e., times for the spins to go out of sync. The measured value depends on the density and chemical surrounding of the hydrogen atoms. The spatial localization of each value is controlled by small variations in the magnetic field.

An important distinction from x-ray techniques is that there are in general no known harmful effects to the patient. MR is, as CT, a volumetric scanning technique. MR is particularly suitable for imaging of the brain and other soft tissue, where the different tissues cannot be distinguished well in CT. The noise level is typically higher in MR images than in the CT case, which for instance causes tissue boundaries to be less

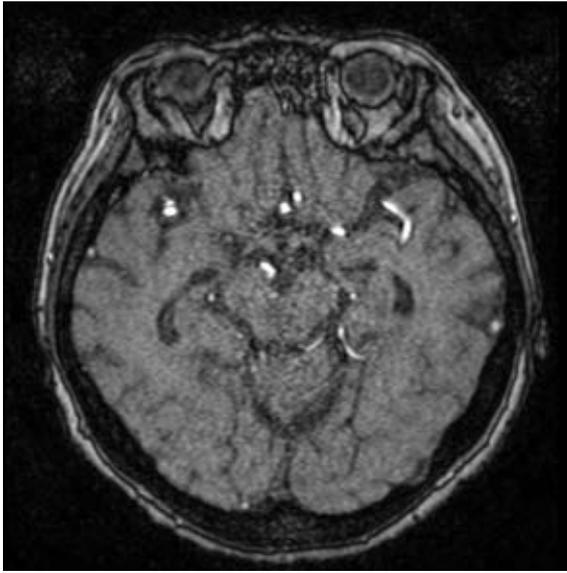


Figure 1.2: An example of a medical data set from MR imaging, a slice from a head scan.

distinct, see figure 1.2. MR methods continue to show tremendous progress and there is a large set of different examination types that can be performed, such as Diffusion Tensor Imaging and Functional MR Imaging.

Ultrasound is another imaging technique without negative side-effects that is widely deployed. The typical use is for 2D imaging, but there are also 3D scanners. As for CT and MR, ultrasound is continuously finding new application areas. Nuclear imaging also constitutes an important branch of medical imaging. In contrast to the typical use of other imaging methods, nuclear imaging shows physiological function rather than anatomy, by measuring emission from radioactive substances administered to the patient. Nuclear imaging data sets are typically of low resolution and 3D techniques are common, such as Positron Emission Tomography (PET). An important hybrid technique is CT-PET, producing multivariate volumetric data sets of both anatomy and physiology.

This thesis studies visualization of volumetric data sets and, as described above, there are many sources for medical data of that type. The emphasis in this work is, however, on CT and MR data sets; there will be no examples from ultrasound or nuclear imaging. Many of the above techniques can also produce time-varying data sets but this is not a focus area in this work. The volumetric data sets are typically formatted as a stack of 2D slices when delivered from the scanning modality. The resolution is usually not isotropic, i.e., the distance between the slices (the z -direction) is not exactly the same as the pixel distance within each slice (the x, y -plane). With modern equipment, there is no technical reason for having lower z resolution, but it is often motivated by reduced radiation dose or decreased examination time.

The scale of the produced values motivates some discussion. In CT the values describe x-ray attenuation that has been calibrated into Hounsfield units (HU), where air corresponds to -1000 HU and water to 0 HU. This means that a given tissue type

will always correspond to a fairly constant HU value. In MR the values can be different types of measurements and they need to be interpreted in the context of the protocol used to capture the data. An important prerequisite for some of the work in this thesis is that MR images do not have any calibrated value range. The value of a specific tissue can vary between patients and between scans of the same patient. This is a major impediment for consistent diagnostic review in situations where high accuracy is needed.

1.1.3 Medical volume visualization

Volumetric data sets are very common in medical imaging and will become even more common as the technologies of CT, MR, ultrasound, and nuclear imaging continue to provide more advanced examination types. The highly dominant visualization method is to show the 2D slices in the format they were delivered from the modality. The volume is reviewed by browsing through this stack of image slices. This approach is sufficient for many examinations but the limitation of being bound to the original slices severely reduces the interaction possibilities.

More interaction is provided by Multiplanar Reconstruction¹ (MPR), where a slice through the volume of arbitrary orientation is displayed. The slicing plane can also be curved. MPR is a routine tool for many radiologists and MPR views of the three main planes are often used as reference views as a complement to other volume visualizations.

There are several techniques that visualize the full volume rather than a slice of it. A commonly used method is Maximum Intensity Projection (MIP). In MIP renderings are constructed from the entire volume or a slab. The volume is projected onto the image plane and each pixel is set to depict the maximum intensity of all data points projected onto it. The viewpoint can be changed freely. MIP is particularly useful for narrow, high-contrast objects such as vessels in angiographies.

Surface rendering (also known as Shaded Surface Display, SSD) is a type of 3D visualization that is less relevant for diagnostic work. In this method a surface is extracted from the data and rendered using a mosaic of connected polygons. Surface rendering is fast and can be useful in some cases, but it is not suitable as a general data exploration tool in clinical use [Rob00]. Medical data sets often have poor contrast between tissues and indistinct boundaries, which makes the extraction of a relevant surface difficult.

In contrast with surface rendering, DVR is considered to be very suitable for diagnostic medical visualization. In DVR, which in the medical community also is known as Volume Rendering Technique (VRT), semi-transparent colors are assigned to the tissues, enabling data points at all depths to contribute to the image. The technical fundamentals of DVR are described in section 1.2. A particularly wide-spread application of volume visualization is Virtual Colonoscopy, which often is based on DVR. This method uses a CT data set to simulate a physical colonoscopy, which is a procedure where an endoscope is inserted into the colon in search for cancer indications.

The motivation for DVR in clinical use is highly relevant for this thesis. Therefore, an analysis of the current and future usage of slice-based viewing and DVR, respectively, is presented in section 1.3.

¹Also known as Multiplanar Reformatting

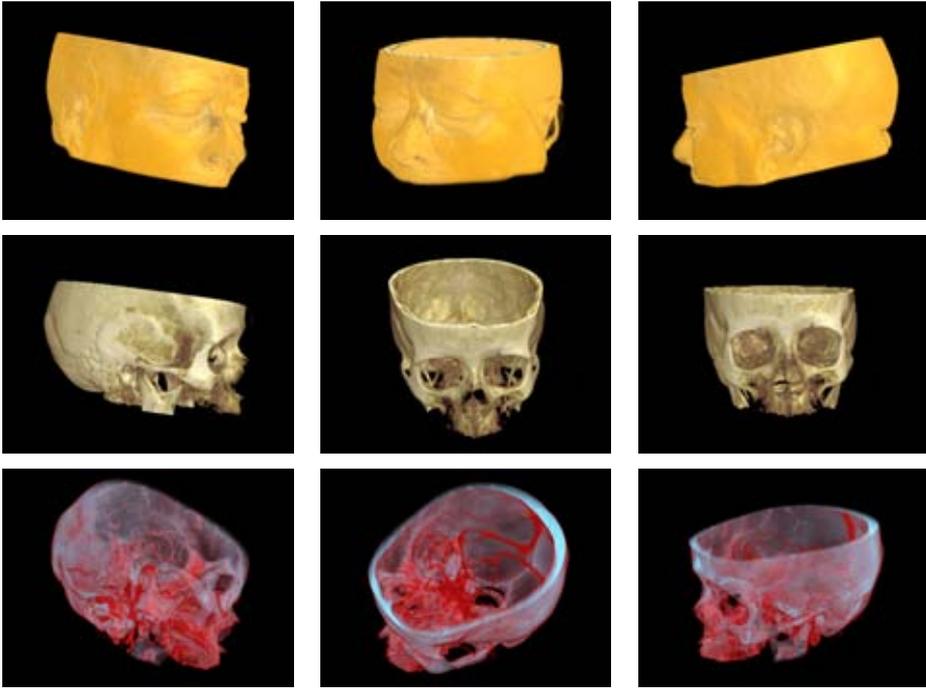


Figure 1.3: A number of medical volume visualizations. The data set is one and the same. The different visualizations are achieved by varying the visual appearance mapping (the Transfer Function) and varying the viewpoint of the virtual camera.

1.2 Direct Volume Rendering

DVR is a visualization technique that aims to convey an entire 3D data set in a 2D image. The key to making this work is to assign semi-transparent colors to the data samples. In this way, objects at all depths in the volume can be seen at once, without obscuring each other. The term “Direct” in DVR stems from the fact that the rendered image is constructed directly from the data, as opposed to techniques that create an intermediate representation, for instance an extracted surface model.

DVR is used in wide range of applications. It is the preferred technique for photo-realistic images of fire, smoke and clouds in computer games and motion picture special effects. It is also a common scientific tool, where visualization of medical data sets is one of the main areas. DVR is often used as an *exploratory* tool where the human user studies an unknown data set. A central component of the exploration is that DVR allows the user to interact with the visualization, to navigate between an immense number of alternative depictions of every single data set. A number of different renderings of a medical data set is shown in figure 1.3.

In an interactive exploration setting, the success of a DVR application is dependent on performance. The rendering must promptly respond to the user’s actions, otherwise the understanding of the visualization will be hampered. Rotation of the volume is a typical operation that needs a perceived real-time performance of 20-30 frames per second (fps). Below 5 fps the response times are usually experienced as very disturbing.

1.2.1 Volumetric data

A volumetric data set, often referred to simply as a *volume*, is usually considered to represent a continuous function in a three-dimensional space. Thus, each point in space corresponds to a function value, formulated mathematically as:

$$f: \mathbb{R}^3 \rightarrow \mathbb{R}$$

Data sets arising from measurements do not have continuous values, they are limited to the points in space where measurements have been collected. A very common case is that the data points constitute a uniform regular grid. Such data points are in the 3D case known as *voxels*, a name stemming from their 2D counterpart *pixels* (picture elements). When values at points in between the original data points are needed, an interpolation between nearby voxels is used as an approximation of the continuous function. The voxels do not need to have the same size in all dimensions.

The values in a volumetric data set can represent many different entities and properties. The interpretation of typical medical examinations is given in section 1.1.2. Examples of properties from measurements and simulations in other domains include temperature, density and electrostatic potential. There can also be more than one value in each voxel, known as multivariate data, which could be flow velocities and diffusion tensors.

1.2.2 Volume rendering overview

The process of constructing an image from a volumetric data set using DVR can be summarized by the following steps, as defined by Engel et al. [EHK*06]:

- **Data traversal.** The positions where samples will be taken from the volume are determined.
- **Sampling.** The data set is sampled at the chosen positions. The sampling points typically do not coincide with the grid points, and so interpolation is needed to reconstruct the sample value.
- **Gradient computation.** The gradient of the data is often needed, in particular as input to the shading component (described below). Gradient computation requires additional sampling.
- **Classification.** The sampled values are mapped to optical properties, typically color and opacity. The classification is used to visually distinguish materials in the volume.
- **Shading and illumination.** Shading and illumination effects can be used to modulate the appearance of the samples. The three-dimensional impression is often enhanced by gradient-based shading.
- **Compositing.** The pixels of the rendered image are computed by compositing the optical properties of the samples according to the volume rendering integral.

In the volume rendering process, two parts are particularly central in the transformation into a visual representation. The compositing step constitutes the optical foundation of the method, and it will be further described in section 1.2.3. The second central part is the classification, representing much of the data exploration in DVR.

Classification is typically user-controlled by means of a Transfer Function as described in section 1.2.4. Further description of the other parts of volume rendering is beyond the scope of this thesis but they are all active research areas within scientific visualization, computer graphics, and/or image processing.

The general pipeline components of DVR can be put together in several different ways. There are two main types of methods, *image-order* and *object-order* methods. Image-order means that the process originates from the pixels in the image to be rendered. Object-order methods approach the process differently – traversing the volume and projecting partial results onto the screen. The most popular image-order method is *ray casting*, where one or more rays are cast through the volume for each pixel in the image. A very common object-order method is *texture slicing*, where the volume is sampled by a number of 2D slices and then the slices are projected onto the image in the compositing step. Both ray casting and texture slicing can be effectively implemented on the processing unit of the graphics board, the GPU.

1.2.3 Compositing

The algorithm for creating a volume rendering image is based on simplified models of the real physical processes occurring when light interacts with matter. These optical models describe how a ray of light is affected when travelling through the volume. The volume is seen as being composed of different materials that cause absorption and emission of light. Illumination models also include scattering effects, but scattering will not be considered in the following for the sake of simplicity. For details beyond the description below, refer to [EHK*06].

In the real world, light rays travel through the material and reach the observer, known as the *camera* in visualization models. When creating an image, each pixel is set to the appearance of a ray ending up in that position. The optical model accounting for emission and absorption results in the *volume rendering integral*, that computes the light reaching the camera:

$$I(b) = I_0 T(a, b) + \int_a^b q(u) T(u, b) du \quad (1.1)$$

The equation is illustrated in figure 1.4. The ray is defined by entry and exit points a and b . The light radiance is given by I , with $I(b)$ being the value at the exit point, i.e., the image pixel value, and I_0 being the light entering from the background. The function $T(u, v)$ is an aggregate of the transparency between points u and v . The function $q(u)$ specifies the emission at a point along the ray. All in all, the first term accounts for the absorption of light as the ray passes through the volume and the second term captures the emission and color contribution from within the volume, which is also affected by absorption.

Typically, numerical methods are used to compute the volume rendering integral (eq. 1.1) in practice. The ray is divided into n small segments, for which the optical properties are assumed to be approximately constant. The emission contribution from a segment i then becomes a single color \mathbf{c}_i . The transparency T_i is usually denoted by the opposite property *opacity*, $\alpha_i = 1 - T_i$. The resulting radiance $I(b) = \mathbf{c}_{\text{tot}}$ is typically computed iteratively from front to back, from b to a :

$$\left. \begin{aligned} \mathbf{c}_{\text{tot}} &\leftarrow \mathbf{c}_{\text{tot}} + (1 - \alpha_{\text{tot}}) \cdot \alpha_i \cdot \mathbf{c}_i \\ \alpha_{\text{tot}} &\leftarrow \alpha_{\text{tot}} + (1 - \alpha_{\text{tot}}) \cdot \alpha_i \end{aligned} \right\} \quad i = n, n-1, \dots, 1 \quad (1.2)$$

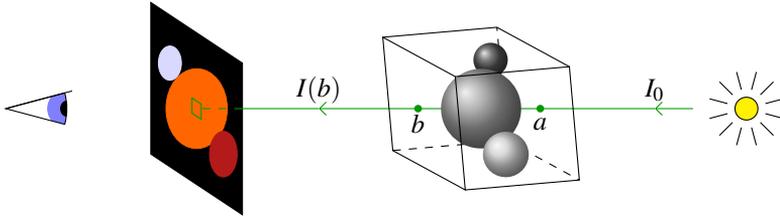


Figure 1.4: Volume rendering light ray model. A pixel value is computed as the radiance of a virtual light ray travelling through the volume. The volume data is given optical properties that modulate the radiance from the initial value I_0 to the outgoing value $I(b)$.

1.2.4 Transfer Functions

The classification part of DVR is achieved through a Transfer Function (TF). The TF constitutes the interface for the user to control what is to be shown in the data set and how it should appear. More precisely, the TF provides the optical properties of a volume sample used in the volume rendering integral. In its basic form, the TF is only dependent on the value of the sample s , i.e., it is a 1D mapping between the sample range and a color vector:

$$\mathbf{c} = \mathcal{T}(s), \quad \mathcal{T}: \mathbb{R} \rightarrow \mathbb{R}^4 \quad (1.3)$$

From here and onwards, unless otherwise stated, a color \mathbf{c} is a four-component vector consisting of red, green, and blue radiance as well as opacity (RGB α format). By changing the TF, a single data set can be given completely different visualizations, as seen in figure 1.3.

A common TF user interface is shown in figure 1.5. Assume that bone is to be studied in a medical data set. A TF definition must then achieve two objectives: defining which data values are to be connected to bone tissue and defining how bone tissue should appear in the rendered image. In this example, the user controls in the TF are trapezoid shapes. The position and width of the rightmost trapezoid (extending beyond the shown scale) define the value range of bone. The height of the trapezoid controls the opacity and its color sets the color of bone in the rendering. Bone is set to be white in this case. Realistic coloring is, however, not a necessary objective. Rather, optimal visual separation of tissues is the typical goal. In manual exploration of an unknown data set, these user interface controls are moved and reshaped in order to best match the feature to be classified. The more distinct features to visualize, the more tuning of visual appearance is needed so that the joint rendering of the entire volume becomes informative. Thus, the user's interaction with the TF constitutes a central part of the exploratory process.

1.3 Direct Volume Rendering in clinical use

The general objective for this thesis is to contribute to increased capabilities and higher efficiency in health care, with the specific aim to facilitate and further empower DVR in clinical use. The main alternative to DVR is slice-by-slice viewing; to browse a stack of 2D images. From a layman's perspective it may seem obvious that the best way

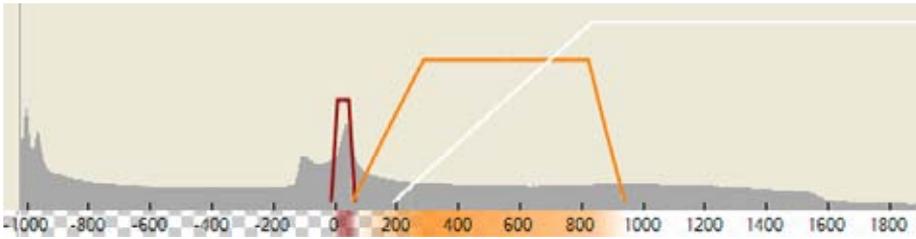


Figure 1.5: Graphical user interface for Transfer Function specification. The gray, filled curve is the histogram of the data. The trapezoids set a visual appearance (color and opacity) for materials in the data set. In the medical case, the materials are tissues that typically have overlapping value ranges, which leads to complex TF definitions.

to depict a volumetric data set is through a volume visualization. This point of view is debatable within diagnostic imaging. As the following discussion shows DVR will not make slice visualizations obsolete but DVR still has a crucial role to fill for routine diagnostic work.

Slice-by-slice viewing is the highly dominant approach for diagnostic review of volumetric data sets. DVR, on the other hand, is not uncommon in the clinical environment but it is far from being an everyday tool for every radiologist. The typical scenario is that if 3D visualization is available, only a fraction of the physicians are using it routinely. This is changing, the deployment of DVR is increasing, but there is still a long way to go before 3D visualization is as natural as slice viewing for volumetric data.

One reason for the slow adoption is limitations in the DVR applications in the clinical context, which is the topic of the next chapter, but there are also other reasons why 3D is not fully deployed. Radiologists are experts at building mental 3D models from stacks of slices, an ability developed through extensive training and experience in slice-by-slice viewing. As this “3D rendering” is already in their mind, spending time constructing it on screen can be unnecessary. In contrast, other physicians who do not have the same experience benefit more from 3D presentations. Moreover, it is likely that the next generation of radiologists will require 3D visualization in their clinical routine since they have become accustomed to these methods during the education.

Another explanation for the relatively low interest in 3D is that the bulk of radiology procedures result in examinations that are not so complex or large that the benefits of 3D apply – this can make DVR seem a peripheral tool. Finally, while the radiology community quickly adopts to technological change, it is more conservative when it comes to changing workflow. The community has spent 100 years streamlining its work for 2D images and this takes time to change.

In spite of its dominant status, slice-by-slice viewing has clear drawbacks compared to 3D visualization in many cases:

- The process of browsing through thousands of slices is very time-consuming and onerous.
- To reduce browsing the stack is often condensed into fewer slices, which discards potentially useful high resolution data.
- It is a difficult and time-consuming process to get an overview of large regions of the anatomy.

- It is difficult to perceive complex structures extending orthogonally to the slice plane, e.g., vessel trees.
- The understanding of non-radiologists reviewing the images is hampered.

The most driving factor for DVR deployment will be the continuous increase in data set sizes. Even the radiologists that prefer traditional slice viewing will need DVR for overview and navigation.

The conclusion to be drawn is that slice visualization suffers from limitations that are increasingly problematic. In light of the additional capabilities provided, DVR is a necessary complement in diagnostic imaging. In order to achieve the required deployment of DVR in clinical use, there a number of challenges to be met. Some of the major obstacles to overcome are presented in the next chapter.

1.4 Contributions

The research contributions of this thesis focus on DVR as a tool for clinical image review. The objective is to enhance the diagnostic capabilities while simplifying the required user interaction. The fundamental means to achieve this goal is to exploit the domain knowledge of the medical professional in the technical solutions. The individual contributions are introduced in the published papers included in this thesis, referred to as papers I – IX throughout the text. In chapter 3 the different components are presented in their relevant context.

Paper I investigates the potential of volume data reduction using a TF-centered multiresolution scheme in combination with wavelet based data compression.

Paper II examines the use of range weights for both detection of characteristic tissue intensities and separation of tissues with overlapping sample value ranges.

Paper III compares different methods for achieving standardized visualizations for uncalibrated MR data sets, where one method is adapted from paper II.

Paper IV presents a technique for direct interpolation of samples over block boundaries of arbitrary resolution differences.

Paper V further investigates spatial coherence to improve histogram presentation and aid in the TF design.

Paper VI presents an extension to traditional histograms in which a sorted, additional attribute is displayed to further improve TF design.

Paper VII extends the techniques from paper II to support additional neighborhood definitions and a spatial refinement of local tissue ranges.

Paper VIII showcases the virtual autopsy application and integrates multiresolution ray casting, TF-based level-of-detail selection, interblock interpolation, and more.

Paper IX presents an animation technique that conveys classification uncertainty in DVR based on a probabilistic formulation of the TF.

Chapter 2

Challenges in Medical Volume Rendering

Direct Volume Rendering is a technique that offers many potential benefits to diagnostic work within medical imaging, as described in the previous chapter. DVR enables analysis of more complex cases than before, while being more efficient and allowing more accurate assessments for certain standard examinations. The work presented in this thesis aims to address some of the specific challenges that DVR needs to meet in the clinical context. These challenges are described in-depth in this chapter along with an overview of previous research efforts in this field.

2.1 Large data sets

There has been a rapid technical development of medical imaging modalities in recent years, which has enabled important benefits for the diagnostic methods. Significantly improved spatial resolution of the data sets has enabled more detailed diagnostic assessments and multivariate measurements lead to unprecedented analysis possibilities. Furthermore, the decreased scan times allow procedures that were previously impossible, for instance high-quality scans of beating hearts. The drawback of this important progress is an enormous increase in data set sizes, even for routine examinations [And03].

Conventional technical solutions have not been sufficient to deal with the continuously growing data sizes. For visualization techniques in general, and DVR in particular, there is an urgent need for improved methods in order to achieve interactive exploration of the data sets. One aspect is the technical limitations in terms of memory capacity and bandwidth that pose serious challenges for the visualization pipeline, making sufficiently high frame rates hard to reach. To achieve the performance needed for DVR in clinical use, methods that can reduce the memory and bandwidth requirements for retrieval, unpacking and rendering of the large data sets must be developed.

There is also a human aspect of the large data set problem. The gigabytes of available data is neither possible nor necessary for the physician to take in. A mere few kilobytes may be enough for the assessment task being addressed, which entails that the task of the visualization is to assist in finding this small subset in an efficient way. The objective of medical visualization is thus transforming from *displaying* all available data to *navigating* within the data.

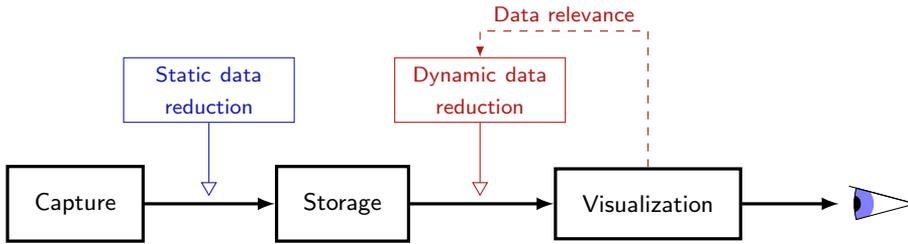


Figure 2.1: Static vs. dynamic data reduction pipelines for DVR. The data flow is from left to right, from capture to human perception. Static data reduction is a one-time process typically applied before the data is stored. Dynamic data reduction continuously adapts to the actual usage of the data in the visualization.

A large number of approaches to handle overly large data sets have been proposed in previous research. In the context of this thesis they can be divided into two main groups: static and dynamic data reduction, illustrated in figure 2.1. The distinction is that a static approach is an isolated preprocessing step that does not take the user interaction with the visualization of the data into account. Dynamic data reduction, on the other hand, employs a feedback loop from the visualization stage to the unpacking stage, where the data is prioritized and the most relevant subset is transferred to the visualization component. A common foundation for data reduction methods is a multiresolution representation of the data, which requires tailored rendering techniques. The rendering aspect is described in section 2.1.3.

2.1.1 Static data reduction

The data compression field of research lays the foundation for static data reduction methods. Compression schemes are either lossless and lossy, the former ensures a perfect reconstruction after decompression, whereas the latter can achieve higher compression ratios by allowing reconstruction errors.

The JPEG2000 image compression standard [Ada01], common in medical imaging, represents a typical approach. The first step is a transformation, which here consists of a wavelet transformation. Wavelets are well suited for image coding since they enable low perceptual error. The second step is an encoding of the coefficients to reduce the size. Methods used are, for example, run-length encoding and/or Huffman entropy encoding. The wavelet coefficients can be effectively compressed at this stage, since they consist of many zeros and values with small absolute value.

The basic setup described so far corresponds to a lossless compression. In lossy schemes, the size can be reduced by modifying the transformed data, the wavelet coefficients. One way is quantization, to decrease the precision of the coefficients, which can effectively reduce data without substantial negative effect on reconstruction quality. Another possibility is to discard certain subsets of the data, which is straightforward in the inherent multiresolution data format.

Many schemes have used wavelet approaches to specifically target volumetric data compression [Mur92, IP99, BIP01a, NS01]. A significant part of these methods is the respective solution for efficient random access to the data, which is important for the total performance.

The Discrete Cosine Transform (DCT) is another transformation used in a similar two-step setup. It is the base of the original JPEG codec [ITU92]. Also DCT has been applied to volumetric data compression [YL95,PW03].

A third class of compression methods is vector quantization [NH92], resulting in lossy compression. This method operates on multidimensional vectors, which typically are constructed from groups of data samples. The vector space is quantized, i.e., reduced to a finite set of model vectors. An arbitrary input vector is approximated by one of the quantized vectors, which allows for efficient subsequent encoding. Vector quantization can also be combined with other approaches, by letting it operate on transformed coefficients [SW03].

An important part of the compression algorithms described above is that they employ blocking in some form. Blocking is a subdivision of the data set into small regions, known as blocks or bricks, which are processed individually. The appropriate choice of block size depends heavily on the characteristics of the hardware components of the computer system in use.

The standard compression methods have been extended for visualization purposes. Bajaj et al. [BIP01b] introduce voxel visualization importance as a weight factor in a wavelet compression approach. Coefficients corresponding to important voxels are prioritized, resulting in higher visual quality in the subsequently rendered image. In a similar approach, Sohn et al. [SBS02] let volumetric features guide the compression, in their case applied to time-varying volumes. A significant drawback with both these approaches is that the important features need to be known at compression time.

A combined compression and rendering scheme for DVR based on vector quantization was proposed by Schneider and Westermann [SW03]. An advantage of this approach is the ability to both decompress and render on the graphics hardware. An example of data reduction for irregular volumetric data is the work of Cignoni et al. [CMPS97]. In this case, compression corresponds to topology-preserving simplification of a tetrahedral mesh.

An important point that sometimes is overseen is that the data reduction should be retained throughout the pipeline, even at the rendering stage. In a traditional pipeline setup, the decompression algorithm restores the data to full resolution even if the quality is lower. This means that the full amount of data must be handled in the rendering, thus disabling the data reduction effect that would be highly desired to increase rendering performance and decrease the need for GPU memory.

Furthermore, it is necessary to bear in mind that maximal compression ratio does not necessarily mean maximal performance. The overall performance of the entire pipeline should be considered and one crucial factor is the decompression speed. The algorithm providing the best compression ratio may be an inappropriate choice if the decompression is slow. Depending on the system characteristics, it may be better to avoid compression altogether if the decrease in transfer time is exceeded by the increase in processing time.

2.1.2 Dynamic data reduction

Dynamic data reduction methods go one step further compared to the standard notion of compression. The idea is that visualization parameters, which in DVR could be the distance to the camera, the Transfer Function (TF), the viewpoint, etc., entails that the demand on precision varies substantially within the data set. Thus, the data can be reduced in these regions, without much loss in the quality of the rendered image. The key feature of a dynamic method is the ability to adapt to changing parameters,

by way of a feed-back loop from the rendering to the unpacking/decompression stage, as illustrated in figure 2.1. In contrast, static data reduction methods are not adaptive since they are defined once and for all at compression time.

A number of volume rendering pipelines employing dynamic data reduction have been developed. Several multiresolution DVR schemes have been proposed that employ an level-of-detail (LOD) selection based on, for example, distance to viewpoint and field-of-view size [LHJ99, WWH*00, BNS01]. A full visualization pipeline was presented by Guthe et al. [GWGS02] where a multiresolution representation is achieved through a blocked wavelet compression scheme. An LOD selection is performed at the decompression stage, prioritizing block resolution according to the distance to viewpoint and the L_2 data error of the resolution level. Gao et al. [GHJA05] presented LOD selection based on TF-transformed data computed using a coarse value histogram. All the above methods are based on an octree structure, i.e., a hierarchical recursive subdivision with increasing resolution. As shown by Ljung [Lju06b], a straightforward flat blocking scheme is a more compact representation for data with abrupt changes in resolution levels.

An important part of a multiresolution scheme is to accurately estimate the impact a lower LOD would have on the rendered image quality. A main challenge is to accurately incorporate the effect of TF transformation without introducing extensive overhead. Furthermore, the method needs to adapt to interactive changes to the TF. Dedicated efforts to estimate LOD impact have been made. One approach is to tabulate the frequency of each actual value distortion and compute the impact after applying the TF [LHJ03], but the size of such tables is feasible only for 8-bit data. Other approaches include a conservative but reasonably fast approximation of actual screen-space error in a multiresolution DVR scheme [GS04] and a computation of block transparency across all viewpoints used for visibility culling [GHSK03].

The research contributions in this thesis include a volume rendering pipeline based on the principles of dynamic data reduction but also allowing static data reduction. The pipeline, presented in section 3.1, combines a TF-based LOD selection with a wavelet compression scheme.

2.1.3 Multiresolution DVR

A common output from both types of data reduction scheme is multiresolution data. This means that the data set is subdivided into small regions, typically blocks, where each region is given an individual resolution. Efficient data reduction is achieved when low-importance regions are given low resolution and vice versa. Having addressed the data reduction issue, there is still a great challenge in developing a rendering method tailored for this situation. First of all, the data reduction should, ideally, be fully exploited in the rendering stage; the frame rate should increase as much as the amount of data decreases. Secondly, high quality images must be rendered. The block-based data reduction schemes will yield clear blocking artifacts in a straightforward DVR implementation, especially in transitions between blocks with very different resolution, as demonstrated in figure 2.2.

A theoretically appealing approach to exploit data reduction in the rendering is to integrate decompression and rendering into a single process. Westermann [Wes94] used this approach to apply a DVR algorithm directly to multiresolution wavelet transformed data. Another example is the visualization pipeline of Schneider and Westermann [SW03] based on vector quantization. For both these schemes, however, the

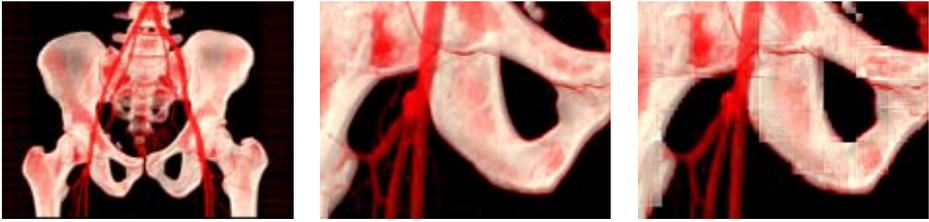


Figure 2.2: Example of blocking artifacts in multiresolution DVR. Left, middle: Original rendering. Right: Multiresolution rendering with high data reduction showing a clear block structure.

additional complexity of the rendering prevents performance benefits from the reduced memory footprint.

The image quality issue in multiresolution DVR has been addressed in GPU-based schemes. LaMar et al. [LHJ99] proposed a multiresolution rendering approach with block-independent processing, where a spherical shell geometry reduces the interblock artifacts. A drawback of the block-independent approach is that it does not provide a continuous transition between blocks. Therefore, data set subdivisions with overlapping blocks have been developed [WWH*00, GWGS02]. Lower resolution samples are replicated at boundaries to higher resolution blocks in order to handle discontinuities. An unwanted side effect is that replication counteracts the data reduction.

The multiresolution DVR pipeline presented in this thesis addresses the issue of blocking artifacts through an interblock interpolation scheme. Without resorting to sample replication, the scheme achieves smooth block transitions, as presented in section 3.1.4.

2.2 Interactive tissue classification

The work of reviewing medical images corresponds, to a great extent, to identifying and delineating different tissues. If this classification is performed correctly, drawing the diagnostic conclusions is often a straightforward task for the trained professional. In the current context, the classification process can be defined as analyzing each voxel with respect to a set of tissues, and for each tissue determining the probability that the voxel belongs to it.

There is an important distinction in what kind of problem a classification scheme attempts to solve. Image processing often deals with precise segmentations where quantitative measurements are a typical objective. The focus in this thesis is the scientific visualization approach, where the classification method is a tool for interactive exploration of the data. In this case, the qualitative aspect of the outcome is more important and an intuitive and direct connection between user and machine is crucial. These characteristics are not common for existing classification and segmentation schemes.

Another relevant grouping of methods is into specialized and general approaches. If the classification is restricted to a narrow domain, good results can be achieved even by fairly automated methods. Examples include the bronchi segmentation of Bartz et al. [BMF*03] and the hip joint segmentation of Zoroofi et al. [ZSS*03]. A different challenge is to create general methods that work for a wide range of image types, which is the objective of the DVR methods presented in this thesis.

2.2.1 One-dimensional Transfer Functions

Medical images have been used for more than a century and for most of that time the diagnostic work flow has been streamlined for the classical x-ray image. One aspect of this is that a scalar value (for example, x-ray attenuation or signal intensity) is by far the most common tissue classification domain. In the DVR context, such a classification corresponds to a 1D TF as described in section 1.2.4.

The 1D TF is, however, not sufficient for the diagnostic tasks of a modern radiologist. Many tissues cannot be separated using only the scalar value; in virtually every examination, the different tissues have fully or partly overlapping value ranges. This is true for nearly every MR data set and for CT data sets in the case of distinguishing different soft tissues. The typical remedy when tissue separation is needed, is to administer contrast agent to the patient before or during the scan. This is very effective when a suitable contrast agent exists, but many diagnostic cases are not yet covered. Furthermore, even data sets with contrast agent can pose problems for DVR. One reason is that new range overlap occurs, as in CT angiographies where blood vessels with contrast agent have the same attenuation as spongy bone. Another common situation is that the differences between the tissues are too subtle to be studied in DVR, for example tumor tissue vs. liver parenchyma in CT examinations. In these cases, informative volume rendering images from 1D TFs are impossible to obtain and the visualization is limited to 2D slices.

In spite of the limitations, the 1D TF is the most important classification interface for the human user. This has been the case ever since the initial DVR method proposed by Drebin et al. [DCH88], where such a one-dimensional mapping is employed. An overview of usability aspects for 1D TF design was presented by König [Kön01]. In many commercial applications for medical DVR, the TF user interface consists of widgets controlling a direct mapping from data value ranges to $rgb\alpha$ vectors. Whenever further tissue separation is needed the user needs to resort to manual sculpting, cutting away disturbing regions of the volume.

2.2.2 Multidimensional Transfer Functions

Many research efforts have been targeted towards overcoming the classification limitations of 1D TFs. A common approach has been to add dimensions to the TF domain, primarily in order to capture boundary characteristics. Two-dimensional TFs using gradient magnitude as the additional dimension were introduced by Levoy [Lev88]. Extensions to this method to further empower classification connected to material boundaries have been proposed over the years [KD98, RBS05, ŠVSG06].

Apart from capturing pure boundary information, measures of local structure computed from second order derivatives have been employed for classification purposes [SWB*00]. For material separation within volumetric surfaces, curvature-based TFs have been shown to add visualization possibilities by conveying shape characteristics [HKG00, KWTM03].

The multidimensional TF schemes above have benefits in the form of enhanced visualization of material boundaries. These solutions are, however, not sufficient in the medical case. One reason is that the boundaries are often far from distinct due to inherent noise [PBSK00]. Even when there are well defined boundaries between tissues, the interior properties of a tissue is often diagnostically important. Moreover, it is not uncommon that separation of unstructured tissues of similar intensity is needed, which is a situation that neither structure nor boundary approaches can handle.

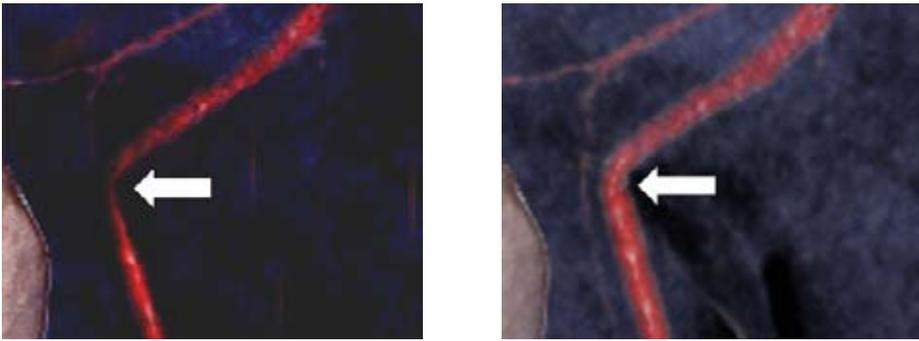


Figure 2.3: Inadequate DVR visualizations can lead to incorrect diagnostic conclusions, even for CT data sets such as this angiography. Left: The rendering using the standard TF indicates a stenosis, pointed out by the arrow. Based on this image, an unnecessary surgical procedure was performed. Right: Exploring other TF settings clearly shows that there is no stenosis.

Usage complexity is an important issue for multidimensional TFs. Even 1D TFs are often overly difficult to master for a physician in the demanding clinical situation. Therefore, adding more dimensions to the TF puts a significant challenge in terms of combining the added capabilities with simplified usage.

Interactive tissue classification methods relevant to clinical work are included in the research contributions of this thesis. In section 3.2.2 simple local histogram analysis is shown to achieve tissue separation beyond the capabilities of boundary-based methods. Furthermore, new types of histogram displays for multidimensional TF design are presented in section 3.2.3.

2.3 Tissue detection

During the research leading to this thesis, much experience has been gathered regarding the clinical usage of DVR. A striking observation is that the radiologists spend a large portion of the work doing tedious manual adjustments of the TF. This work can seldom be characterized as an exploratory navigation within the data, or as an effective visualization session. Rather, it is a time-consuming consequence of overly complex TF construction models. In the clinical context, even the technically straightforward 1D TF model is above the usage simplicity required to enable efficient diagnostic work by the physician. An important step towards wider deployment of medical DVR would be to automatically detect the value ranges of the interesting tissues, since a feasible initial TF then could be provided to the user. Many non-relevant parts of the TF parameter domain would be avoided and the interaction would have a better foundation.

Tissue detection schemes can be needed even if an approximately correct TF is easy to find manually, in cases where consistent fine tuning is needed. A pertinent example is vessel width assessment. A very high accuracy is required, but inadequate visualizations can create seemingly good renderings with large errors in the apparent vessel width [PDE*04], as illustrated in figure 2.3. Visualization consistency over time and between users is important also for the radiologists' gathering of professional experience, since their assessments are continuously "calibrated" through comparison with the previous cases that the physician and his/her colleagues have handled.

2.3.1 Unsupervised tissue identification

There are an immense number of classification approaches that could be used for subsequent TF design. As noted in section 2.2, however, there exists a great challenge in developing generally applicable methods. The most common tissue identification method in DVR is to display the full data set histogram to the user; the idea being that the relevant tissues will stand out as peaks in the histogram. Unfortunately, this is seldom the case, which makes the global histogram a very blunt tool for accurate definition of tissue value ranges.

Some research efforts have been made to find general methods that can be used to predict suitable visualization parameters. Bajaj et al. [BPS97] introduced metrics for identifying appropriate values for isosurface rendering of triangular meshes, the *Contour Spectrum*. A similar method targeting regular cartesian data has also been developed [PWH01].

Tissue detection is particularly important for MR image volumes because of the uncalibrated scale of the voxel values. Previous research efforts have targeted automatic adaptation of visualization parameters for MR data sets [NU99, Oth06]. The adaptations are based on landmarks of the shape of the global histogram. A limitation is that these methods are only applicable for examination types where the global histogram shape is consistent between patients. Rezk-Salama et al. [RSHSG00] use both histograms and boundary measures to adapt TF settings between data sets.

Two different histogram analysis techniques addressing the tissue detection challenge are presented in section 3.3 of this thesis. Both techniques exploit spatial relations in the data to enhance the value of histogram presentations.

2.3.2 Simplified TF design

Many research efforts aim to simplify the manual interaction in TF construction without using automatic tissue detection. In traditional TF specification, the user interaction is directly connected to the volumetric data set, this is known as the *data-driven* approach. Within this class of methods Kindlmann and Durkin [KD98] suggested a simplification of parameters for boundary visualization in DVR where the user defines a mapping between surface distance and opacity. Extensions to this model have been proposed [TLM01, KKH02]. There are also more general data exploration interfaces that are not limited to boundary measures [PBM05].

Another way to simplify traditional TF specification is data probing, i.e., to let the user select representative points or regions in the data set as base for the TF parameters. A data probe is part of the tools presented by Kniss et al. [KKH02]. In the DVR approach of Tzeng et al. [TLM03] probing is used to drive a high-dimensional classification based on an artificial neural network.

With the aim to create a more intuitive 2D TF interface, Rezk-Salama et al. [RSKK06] developed a framework based on a user-oriented semantic abstraction of the parameters. A Principal Component Analysis approach is employed to reduce the TF interaction space.

In contrast with data-driven methods, *image-driven* methods let the user adjust the visualization by selecting from alternative rendered images, changing the TF in an indirect way. The challenge is then to automatically provide the user with a relevant gallery of different settings. He et al. [HHKP96] explore stochastic generation of these alternatives. Further variants and extensions of gallery schemes have been proposed [MAB⁺97, KG01].

2.4 Uncertainty visualization

The challenge of representing errors and uncertainty in visualization applications has been brought forward as one of the main topics for future visualization research [JS03, Joh04]. The benefit of controlling and studying the uncertainty is highly valid within medical DVR. An important aspect is the uncertainty that the user introduces by setting the visualization parameters. A static TF yields a certain appearance in the rendered image, but it is very important that the user can explore the robustness of this setting to make sure that the diagnostic conclusion is not affected by slight mistakes in the TF definition. Figure 2.3 above shows a real-life example of the patient safety risks involved.

In fact, the TF uncertainty is today often assessed by the radiologists through manual adjustments back and forth. A main drawback is that this is very time-consuming. Moreover, the full parameter space relevant to the diagnosis may not be covered by the ad-hoc manual adjustments. This problem grows with the complexity of the TFs used. All in all, there is a substantial risk, especially for physicians with limited DVR experience, that this manual process deviates far from the ideal: an unbiased exploration of all relevant possibilities.

2.4.1 Uncertainty types

A thorough survey of many aspects of uncertainty visualization was presented by Pang et al. [PWL97]. A taxonomy of different methods was proposed, as well as a categorization of uncertainty sources into three groups: acquisition, transformation, and visualization. In the first category statistical variations due to measurement error or simulation simplifications are typical examples. Transformation uncertainty is exemplified by resampling and quantization of data values. Finally, approximations are introduced in the visualization scheme which for instance is manifested by the fact that different DVR schemes do not result in the exact same rendering of a volume.

This thesis focuses on uncertainty arising from the TFs fuzzy classification within DVR, which belongs to the “visualization uncertainty” group. An important distinction is that many previous methods have assumed that the probability values are derived or measured before the visualization stage occurs, whereas the current focus is statistical variation inherent in the visualization process.

2.4.2 Visual uncertainty representations

A number of methods have aimed at representing uncertainty in surface visualizations. The survey of Pang et al. [PWL97] presented several uncertainty representations for surface renderings and proposed additional schemes. One possibility is to connect uncertainty to a proportional spatial displacement of the surface, resulting in a point cloud appearance for low-confidence regions [GR04]. Without changing the spatial extent of the surface, variations in the appearance can be used to convey uncertainty. Hue and texture have been used to visualize isosurface confidence in multiresolution data [RLBS03]. Furthermore, flowline curvature has been employed to represent shape uncertainty of an isosurface [KWTM03].

It is difficult to extend the above surface rendering methods to the volume rendering case. Other researchers have focused on DVR-specific solutions, proposing different ways to incorporate volumetric probabilities. A straightforward solution is to treat

these likelihood values as any other attribute to be visualized, either rendering the likelihood domain itself [RJ99] or applying a multidimensional TF [DKLP01]. In another approach the probability volume is rendered and then used to modulate the pixels of the rendering of the data volume [DKLP01]. A serious limitation of this last approach is that there is no correlation of obscured regions in the two renderings; uncertain regions may affect the final result even if they are not visible in the regular rendering.

The task of visualizing uncertainty from a predefined probabilistic classification was addressed by Kniss et al. [KUS*05], proposing a DVR framework based on statistical risk. The methods include a graph-based data reduction scheme to deal with the challenge of the enlargement of the data sets resulting from the transformation to material classification volumes.

Uncertainty can also be represented by controlled changes in the rendered image. There are examples of such animation schemes in the area of geographical visualization. Gershon [Ger92] used an ordered set of segmentations, which was animated in order to make fuzzy structures stand out. The method of Ehlschlaeger et al. [ESG97] creates a sequence of probabilistically derived rendering realizations to convey spatial uncertainty.

The final research contribution in this thesis is an uncertainty animation technique tailored for clinical use, presented in section 3.4. The foundation is a probabilistic interpretation of the TF, that may become useful also for DVR in general.

Chapter 3

Efficient Medical Volume Visualization

Based on vast research efforts over the past decades both performance and quality of volume visualization have continuously been improved. Medical imaging has been, and is, a prime application area for the developed methods. Despite its success in medical research, volume visualization approaches have not had the same impact in routine clinical situations. To find the reason for this it is important to realize that the needs of a practicing radiologist reading hundreds of examinations per day are not the same as those of a medical researcher, who can spend significant amount of time analyzing individual cases. To reach a more wide-spread use outside of the research field it is thus crucial to work in continuous close collaboration with medical professionals and the medical visualization industry to gain an understanding of the issues involved in the user's clinical workflow.

This chapter describes the research contributions found in the appended papers and puts these contributions in the context of the challenges described in the previous chapter. A main theme of the work is that the clinical usage is put at the core of the research methodology. The common foundation of the developed methods is thus that they are based on the clinical visualization user's perspective and exploit the available medical domain knowledge to address identified pertinent problems. The presentation will show that these methods tailored for the clinical context can lead to increased performance and enhanced quality, thereby increasing the user's ability to fulfill the task at hand.

3.1 Multiresolution visualization pipeline

The first of the identified challenges for DVR in clinical use is the increasingly large data sets, as discussed in section 2.1. With the objective to deal with central parts of this challenge and lay a foundation for future research efforts, a multiresolution DVR pipeline has been developed. The presentation of the pipeline in this section corresponds to papers I, IV and VIII.

The goal is to significantly reduce the amount of data to be processed throughout a Direct Volume Rendering (DVR) pipeline. The path taken in this thesis puts the Transfer Function (TF) at the core, exploiting the user's definition of what to make visible in the data set. When applying a TF large subsets of the volume will give little

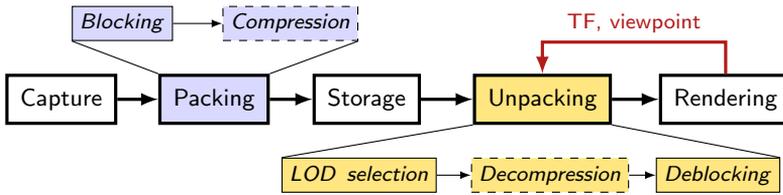


Figure 3.1: Schematic visualization pipeline showing the data flow. A feedback loop provides the LOD selection process with the current TF and viewpoint from the rendering stage.

or no contribution to the rendered image. Thus, these subsets of the original data are of low importance, even if they contain much information from a mathematical point of view. The pipeline is set up to make use of the knowledge encoded in the TF to select an adequate level-of-detail (LOD) in each part of the volume in order to reduce the amount of data to be retrieved, reconstructed and rendered.

3.1.1 Pipeline overview

An overview of the developed pipeline for DVR is given in figure 3.1. When the data has been captured, it is processed in a packing component. The packing primarily includes a blocking step, where the volume is subdivided into small spatial blocks, typically consisting of 16^3 voxels each. The blocking also includes the creation of subsampled versions of each block. The result is a multiresolution representation where the block can be represented by 16^3 , 8^3 , 4^3 , 2^3 , or 1 voxels.

Optionally, the packing stage can include a traditional compression, as long as the blocked multiresolution format is retained. This requirement can, for instance, be fulfilled by the state-of-the-art compression technique for medical imaging – JPEG2000 [Ada01] based on wavelet transforms. In fact, the multiresolution representation is inherent in wavelet-based methods. Paper I employs a wavelet compression approach in this pipeline setup.

The packed data is stored and when it is to be rendered, an unpacking process is applied. The main data reduction is achieved by the LOD selection, described in detail in section 3.1.2. If the data was compressed, a decompression takes place. The final step of the unpacking is to merge the blocks together to a single volume to be rendered. In paper I this deblocking is a separate, explicit step, whereas it is integrated in a GPU-based rendering algorithm in paper IV.

In summary, the multiresolution representation enables a spatially varying resolution level in the volume. This is the foundation for the LOD selection that performs a data reduction to reduce transfer and processing time.

A commonly used alternative to an LOD-based pipeline is to rely on compression alone. There are, however, limitations to such an approach as it is a static reformatting of the data performed without knowledge of how the data is to be used in the visualization. In contrast, the proposed LOD-based pipeline constitutes a dynamic data management approach. The main benefit is that data can be adequately prioritized based on the actual usage in an adaptive manner. Data redundant in the visualization can be discarded, enabling high data reduction even at full quality. Another benefit compared to compression approaches is that the data reduction can be retained throughout the ren-

dering stage. This is not the case for traditional decompression algorithms, where the resolution is brought back to the original, even though the quality may have decreased.

It is important to bear in mind that the dynamic data management does not exclude traditional compression. As shown in paper I, the LOD approach can be added on top of a wavelet-based compression scheme. Thus, the method can at the unpacking/decompression stage enhance an existing compression scheme without counteracting it, which is why it is also referred to as *adaptive decompression*.

3.1.2 Level-of-detail selection

The foundation of the LOD selection algorithm is to effectively estimate the visual importance of each block in the final rendering. The developed method originates from a definition of *TF content*. Let $\mathcal{T} : \mathbb{R} \rightarrow \mathbb{R}^4$ denote a TF, mapping volume sample values to a color and opacity vector. \mathcal{T}_{RGB} denotes the RGB vector and \mathcal{T}_α refers to the alpha component. Consider a block, containing a set of values \mathbf{V}_b . The TF content for the block is defined as the set obtained by applying the TF to each value: $\mathcal{T}(v), v \in \mathbf{V}_b$. The TF content guides the LOD selection for a block as follows:

1. No TF content, $\mathcal{T}_\alpha(v) = 0, \forall v \in \mathbf{V}_b$: The block is completely transparent and can be discarded without introducing distortion.
2. Non-varying TF content, $\mathcal{T}(v) = \mathbf{C}, \forall v \in \mathbf{V}_b$, where \mathbf{C} is a vector constant: The block is completely homogeneous and can be reduced to a single average value without introducing distortion.
3. Varying TF content, $\exists u, v \in \mathbf{V}_b$, such that $\mathcal{T}(u) \neq \mathcal{T}(v)$: Low distortion is achieved by letting the LOD depend on the derivative of the TF in the range of \mathbf{V}_b . A high derivative implies high visual detail, which should be represented by high resolution.

A data reduction is obtained for each block where the selected LOD is less than full resolution. In the cases of no or nonvarying TF content the LOD reduction results in an essentially lossless compression.

A block significance measure, $s_F(\lambda)$, is derived for each subresolution level λ based on TF content. With 16^3 voxels in each block, the measure corresponds to visual distortion when replacing full resolution with resolution level λ , where $\lambda = 8, 4, 2, 1$. The challenge in defining the measure is to make it extremely efficient to compute. The developed approach is based on block specific metadata acquired at compression time, mainly the average value and a simplified histogram. The simplified histogram reduces the full histogram of a block to a number of piece-wise constant segments, see figure 3.2.

Another prerequisite of the significance estimation is a perceptual color difference metric ΔE , described in detail in section 3.1.3. The difference is derived in the CIELUV color space, which requires the use of a transformed TF, defined according to eq. 3.1 where \mathcal{L} represents the RGB to CIELUV conversion.

$$\mathcal{T}^*(v) = \mathcal{L}(\mathcal{T}_\alpha(v) \cdot \mathcal{T}_{\text{RGB}}(v)) \quad (3.1)$$

The significance measure is first derived for $\lambda = 1$, the overall distortion corresponding to replacing full resolution by a single average. The simplified histogram is sampled at small intervals, applying the TF to each sample. The derivation is presented

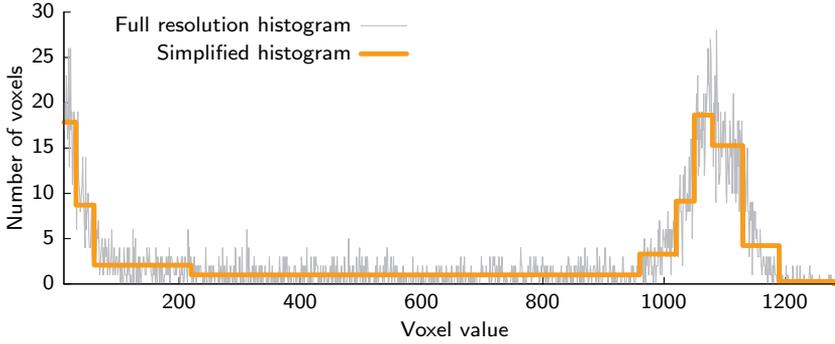


Figure 3.2: The significance measure driving the LOD selection employs metadata for each block in form of a piece-wise constant simplified histogram.

in equation 3.2, where the segments are defined by an index i , a midpoint x_i , a height h_i , and a (common) width w .

$$s_F(1) = \left[\frac{1}{16^3} \sum_i \Delta E (\mathcal{T}^*(x_i), \mathcal{T}^*(\bar{x}))^2 \cdot h_i \cdot w \right]^{1/2} \quad (3.2)$$

The ΔE distortion for each segment is retrieved by comparing transformed mid-points $\mathcal{T}^*(x_i)$ with $\mathcal{T}^*(\bar{x})$, where the block average \bar{x} is used to approximate the TF content at $\lambda = 1$. The distortion is multiplied by the number of voxels in the interval and all interval distortions are added.

The next step is to estimate what portions of the overall distortion originates from each intermediate level. In paper I the distribution of error is done according to the RMS of the wavelet coefficients at each level. This approach is motivated by the fact that large coefficients correspond to more information.

Having derived the significance measures for all levels in all blocks, the actual LOD selection is performed. In an initial step the blocks with no TF content are removed from further processing and the remaining blocks are initialized as the lowest resolution, $\lambda = 1$. A priority queue is then created, containing all possible LOD transitions towards higher resolution. The queue is sorted according to how efficiently the transition reduces distortion: the difference in significance divided by the difference in data size. An LOD selection is performed by traversing this queue and registering the corresponding transitions. A typical setup would be to use a memory budget to define a stopping criterion for the queue traversal – as soon as the LOD selection reaches the budget, the traversal is stopped.

The performance of the data reduction scheme in terms of image quality (rendered distortion, see 3.1.3) at different reduction ratios is shown in figure 3.3. The low-end benchmark is uniform down-sampling (“Uni”). A more relevant comparison comparison is to data reduction based on the reconstruction error (“RecE”), i.e., the distortion due to subsampling without considering how the data is to be visualized. The proposed scheme (“TF”) consistently outperforms the reconstruction error approach. The difference is dependent on the TF in use. The full LOD selection process is performed in a fraction of a second on an off-the-shelf PC.

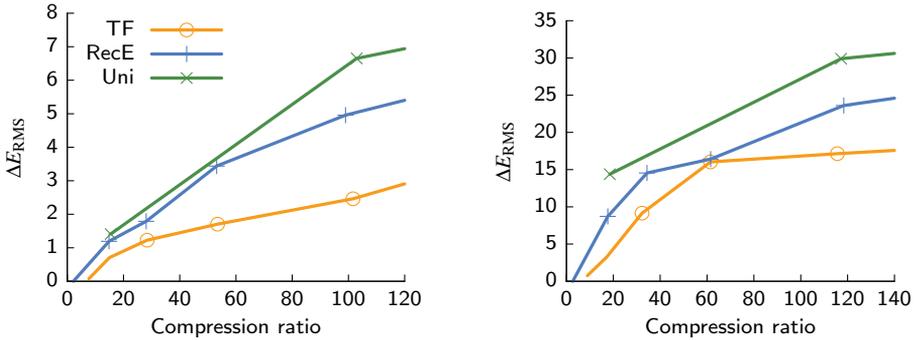


Figure 3.3: Compression versus quality comparison between LOD selection based on TF content (TF), reconstruction error (RecE), and uniform down-sampling (Uni). Distortion is measured compared to a full resolution rendering as ΔE_{RMS} , the root-mean-square of the rendered image quality ΔE . The comparison is made for two data sets rendered with different TFs. Left: low-opacity TF. Right: high-opacity TF.

3.1.3 Distortion metric

Most visualization and computer graphics approaches have high image quality as one objective. The definition of image quality is by no means straightforward and many researchers simply resort to subjective evaluations of artifacts in the rendering. Since the multiresolution rendering aims to maximize image quality at any given data reduction level, the need for an objective distortion metric is insurmountable for tuning and evaluating the methods.

A commonly used measure of distortion in volumetric data reduction schemes is to measure the accuracy of the reconstructed volumetric data compared to the original, here referred to as *reconstruction distortion*. A fundamental concept is the root-mean-square error (RMSE), which is the basis for the often used error measures signal-to-noise ratio (SNR) [SW03] and peak-signal-to-noise ratio (PSNR) [NS01]. LOD selection schemes have been tailored to minimize reconstruction distortion [GWGS02].

Reconstruction distortion does not take the visualization of the data into account. This means that data having no impact on the rendered image may still seem important in terms of such distortion measures. A more sound approach in the current context is to evaluate the error in the rendered image, the *rendered distortion*, which is the error that actually is of relevance to the visualization user. Using such a measure to guide the LOD selection allows large reconstruction distortion while retaining visual fidelity.

The employed distortion metric is a pixel-wise color difference as defined in the CIELUV color space, ΔE . In contrast with RGB, the CIELUV representation constitutes a perceptually adequate metric¹, meaning that a unit difference in ΔE is perceived equally throughout the color space. Two aggregates across the entire image have been found useful, ΔE_{RMS} and ΔE_6 . The former is the root-mean-square average and the latter is the ratio of pixels with $\Delta E > 6.0$, which is a threshold corresponding to perceptually significant errors. The full details on the distortion metric and the RGB to CIELUV transformation can be found in paper I.

¹The original definition has been enhanced to be more perceptually precise, see [LCR01], but the original definition was judged accurate enough for the application in question.

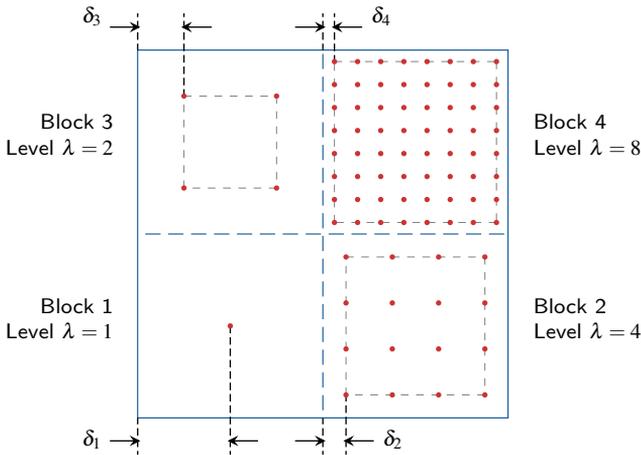


Figure 3.4: A 2D illustration of a four-block neighborhood. Red dots show the location of block samples, the red dashed lines are the sample boundaries. The sample boundary distance, δ_b , is indicated for each block, b .

3.1.4 Interblock interpolation

The developed LOD scheme, as many others, is based on a block subdivision of the volume. An important problem is that the different local resolutions cause discontinuities, block artifacts, in the rendered image. A solution employed by other researchers is to replicate samples between blocks [WWH*00, GWGS02]. This approach can achieve continuity when rendering the blocks independently, which reduces these artifacts. The replication approach has, however, a major drawback in that the redundancy introduced significantly decreases the data reduction achieved by the LOD selection.

A technique has been developed that accomplishes the reduction of blocking artifacts while avoiding sample replication. The method defines interpolation functions in the peripheral parts of the blocks in a way that \mathcal{C}^0 continuity is achieved between arbitrary resolution levels. This interblock interpolation technique was introduced in paper IV, an overview of the main components follow next.

The interblock interpolation is active in the region outside of the sample boundaries of each block, see figure 3.4 for a 2D example. A sample to be taken in such a region in 3D is influenced by eight neighboring blocks. The interpolation algorithm derives a weight w_b for each of the eight blocks. The resulting interpolated value φ is a normalized weighted sum of the block samples φ_b (eq. 3.3). Each block sample is retrieved at the closest point of the block's sample boundary.

$$\varphi = \frac{\sum_{b=1}^8 w_b \varphi_b}{\sum_{b=1}^8 w_b} \quad (3.3)$$

The interpolation should balance the impact of a block sample across the border to a neighboring block, depending on the blocks' resolutions. This is done by deriving the block weights, w_b , from edge weight functions. Each pair of side-facing neighbors are given such a function, defined in between the sample boundaries of the two blocks.

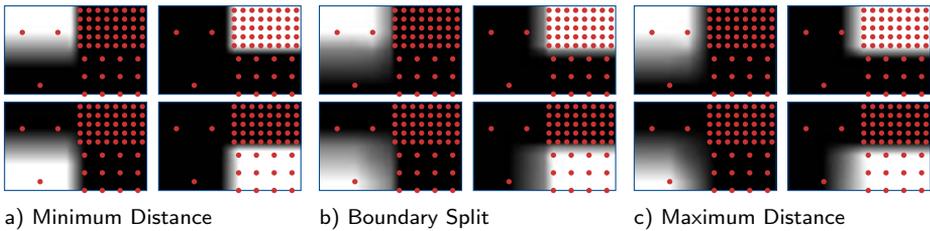


Figure 3.5: Block interpolation weights in 2D for the three interpolation variants. Each data set consists of four blocks at different resolution, where the sample points are denoted as red dots. The block resolutions are 1×1 , 4×4 , 2×2 and 8×8 , from left-bottom to right-top. The interpolation weight for each specific block is illustrated by the four intensity images. The images have been cropped vertically.

An edge weight function is always 1.0 at one sample boundary and 0.0 at the other, being piece-wise linear in between.

Different choices of edge weight functions yield different interpolation characteristics. Three variants have been developed:

Minimum Distance Interpolation: The impact of a value at the sample boundary is limited to the distance it would have if it were fully inside the block. A consequence is that lower resolution samples are extended as constants towards the boundary. See figure 3.5a.

Boundary Split Interpolation: The steepness of the interpolation is not influenced by neighboring blocks. The interpolation is a two-segment linear function that is split at the spatial block boundary. Compared to Minimum Distance, high resolution samples have wider footprints and the constant part of low resolution samples is removed. See figure 3.5b.

Maximum Distance Interpolation: The value is interpolated in one linear segment over the whole distance between neighboring sample boundaries. The derivative of the interpolation is continuous within the interval, which is not the case for the other variants. See figure 3.5c.

All three interpolation variants turn into common trilinear interpolation in between blocks of equal resolution.

The effect of interblock interpolation in terms of increased rendering quality is demonstrated in figure 3.6. The performance drop when adding interblock interpolation heavily depends on which graphics card and base renderer implementation are used. At best, the interblock interpolation doubles the required rendering time.

3.1.5 Virtual autopsy application

The benefits of performing autopsies as part of forensic investigations are undisputed. As shown in paper VIII further benefits can be added by introducing a DVR visualization session to the procedure, known as a *Virtual Autopsy* (VA). The challenge of visualizing very large data sets is particularly relevant within VA, as the data sets are full-body scans captured at maximum resolution – the harm of radiation dose does not need to be taken into consideration. Furthermore, a requirement for an effective VA visualization session is the ability to interactively switch between full-body overviews and close details.

Since conventional visualization applications were unable to meet the demands

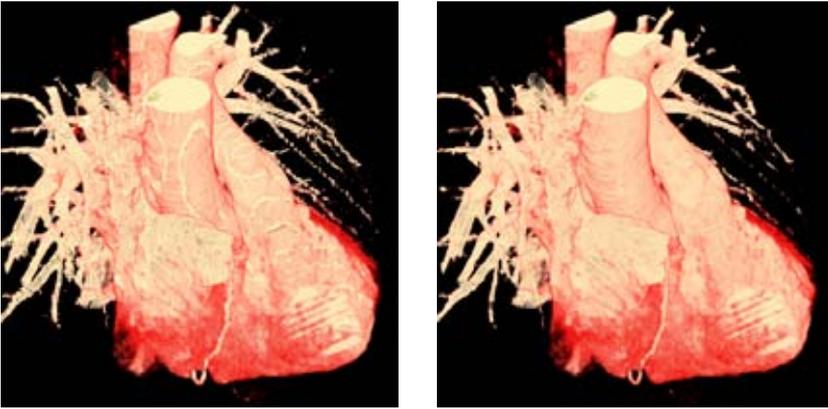


Figure 3.6: Impact of interblock interpolation on highly reduced data. Left: A heart rendering without interblock interpolation, with significant artifacts. Right: Interblock interpolation succeeds in removing the artifacts without adding more data, Maximum Distance variant.

from the VA scenario, a DVR pipeline was developed encompassing the methods presented in sections 3.1.1-3.1.4 and the work by Ljung [Lju06a]. The application, described in paper VIII, enables visualization of these full-body examinations on commodity desktop computers at interactive frame rates, see figure 3.7.

The VA research project, jointly performed by Linköping University and the Swedish National Board of Forensic Medicine, has also defined workflows and developed scanning techniques. The scanning of the cadaver is typically performed through special protocols for high-resolution CT. The images are examined by radiologists and forensic pathologists before the physical autopsy. Findings at this early stage may provide crucial insight in their own right, but also guide crime scene investigators and help the forensic pathologists to plan the physical autopsy. To date over 200 cases have been examined using VA. The significance of volume visualization in the VA procedure has been well established through numerous cases where conclusions could not have been made based on the physical autopsy alone. These substantial added benefits include detection of minute but important features: air pockets, fine fractures, and alien objects such as bullet fragments.

3.2 Domain knowledge in interactive classification

A key reason for the high potential usefulness of DVR in the medical domain is its versatility; DVR can act as a data exploration tool in a wide range of situations. The versatility is represented by the classification component and there exists a significant challenge in providing effective interactive tissue classification, as discussed in section 2.2. This section describes methods and tools that enhance conventional DVR classification by exploiting the knowledge of the visualization user. In many cases, information that is very basic from a user perspective can substantially enhance the effectiveness of the visualization. The presentation corresponds to parts of papers II, VI and VII.

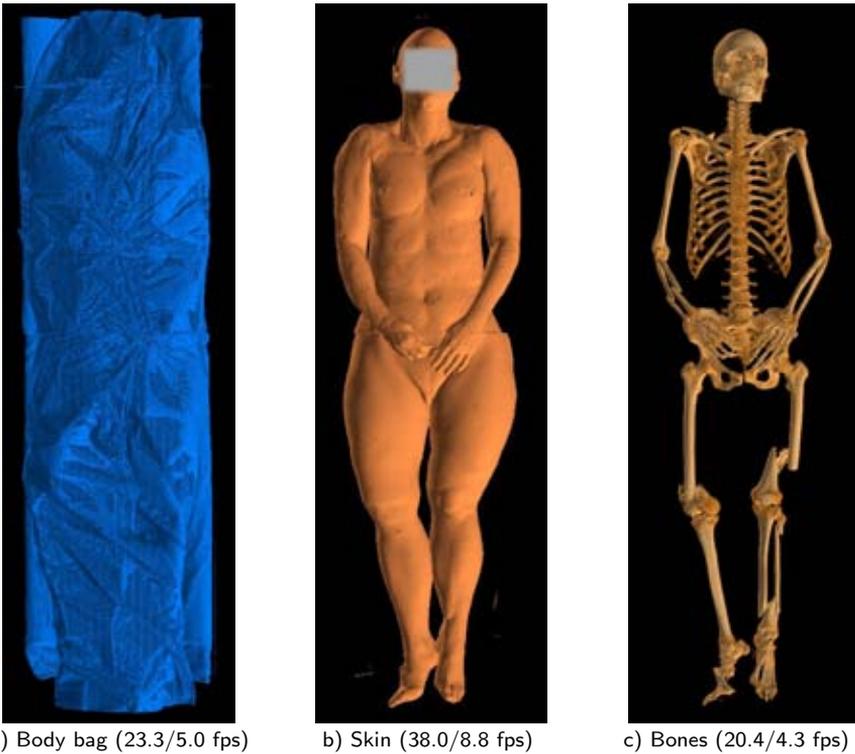


Figure 3.7: Full-body virtual autopsy of traffic accident victim. The resolution is $512 \times 512 \times 3396$ samples, which together with high-precision precomputed gradient data amounts to 5.0 GB. These images are rendered at a data reduction of 30:1 in a 1024×1024 viewport using the adaptive object-space ray casting on the GPU from [Lju06a]. Frame rates are shown without and with interblock interpolation.

3.2.1 Range weight

A first task in visualization based on domain knowledge is to identify information that can readily be input to the pipeline. Intensity and gradient values have been the basis to describe features of interest in the data for most TF models to this date, but much more a priori information can be integrated. A particularly significant type of knowledge is spatial voxel relations. In many domains the features to be studied are spatially connected or concentrated, which is also true for medical imaging where primarily tissues and organs are visualized. This simple fact is a foundation for any segmentation algorithm, but it is not incorporated in traditional DVR approaches, where voxel appearance is not dependent on any wider set of neighbor values.

As first demonstrated in paper II, medical DVR can be effectively enhanced by incorporating combined knowledge on spatial relations and intensity characteristics. This is achieved through simple statistical neighborhood analysis, namely, by studying ranges in local histograms. To this end, the statistical property *range weight* has been introduced.

The range weight measures a neighborhood footprint in a given partial intensity

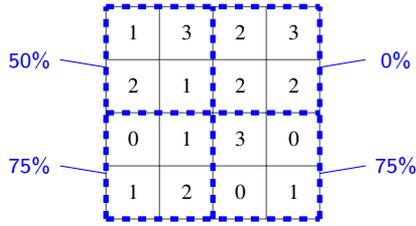


Figure 3.8: Range weight combines data values and spatial neighborhoods. In this simple 2D example, the range weight for range $[0..1]$ is shown for the neighborhoods of 2×2 pixels.

range, see figure 3.8. In other words, when given a partial range, the range weight is a measure of how well a neighborhood is described by that range. The range weight is denoted w_r and the definition is given in eq. 3.4.

$$w_r(\Phi, N) = \frac{|N \cap V_\Phi|}{|N|} \quad (3.4)$$

N is an arbitrary voxel neighborhood, V_Φ is the set of voxels within range Φ . $|V|$ is the cardinality of a set V . The range weight is used both for the tissue classification in section 3.2.2 and some of the tissue detection approaches in section 3.3.

3.2.2 Tissue separation

In medical image data sets, tissues that should be distinguished for diagnostic purposes often have overlapping intensity ranges. A typical example is CT angiographies, where vessels with contrast agent have the same intensity as spongy bone. In MRI, the problem is even more common and it is a main reason why DVR is seldom used for these types of data sets. Traditional DVR approaches based only on intensity and possibly gradients cannot achieve a rendering that separates the different tissues well in these cases.

When the intensity alone cannot separate tissues, a local histogram analysis can often provide differentiating information. In papers II and VII a DVR classification approach based on range weight criteria is presented. As the method is intended for interactive use, an intuitive model for input for the medical end-user is desired. The method allows user input in the form of straightforward statements on tissue characteristics in terms of size, homogeneity, and the existence of neighboring tissues. More precisely, the user would, for each tissue, define levels of these characteristics that correspond to a confident classification. This input can then be automatically translated to range weight criteria.

An example of the resulting tissue separation is given in figure 3.9. The objective is to separate the liver tumor from the diagnostically irrelevant regions. Having defined the intensity range of the tumor as in traditional TF handling, the user can add a homogeneity criterion: “show me homogeneous tissue in this range”. Denoting the tumor range as Φ , the criterion can be formulated as a range weight threshold, e.g.: $w_r(\Phi, N_v) \geq 0.6$. When applying this criterion to the rendering, a neighborhood N_v for each voxel v is analyzed and non-homogeneous obscuring tissue is given full transparency.

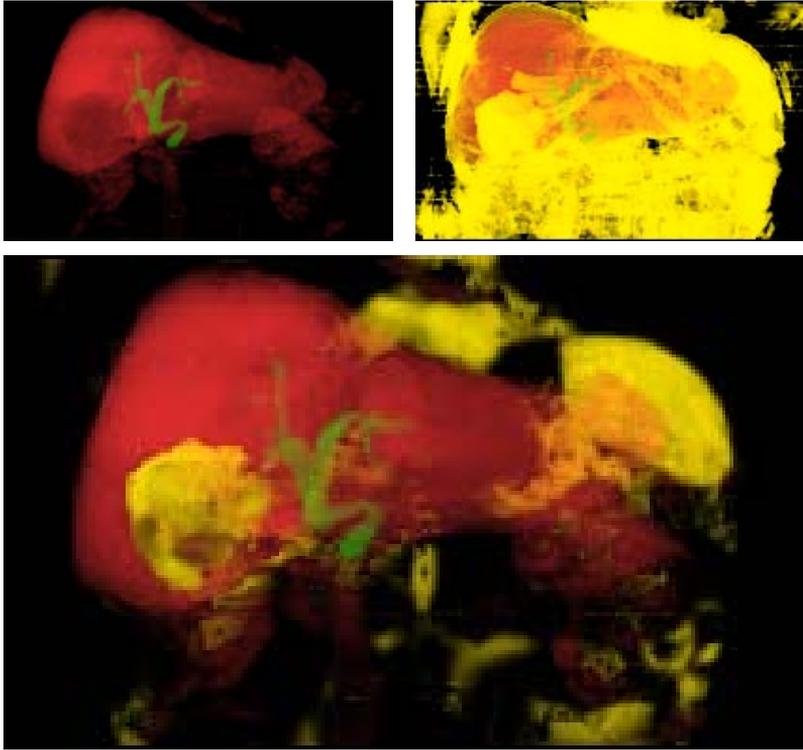


Figure 3.9: Visualization of liver tumor. Top left: 1D TF rendering, the tumor area is slightly darker than the liver. Top right: An attempt to highlight the tumor with a 1D TF. Bottom: A classifying 2D TF clearly highlights the tumor and the spleen.

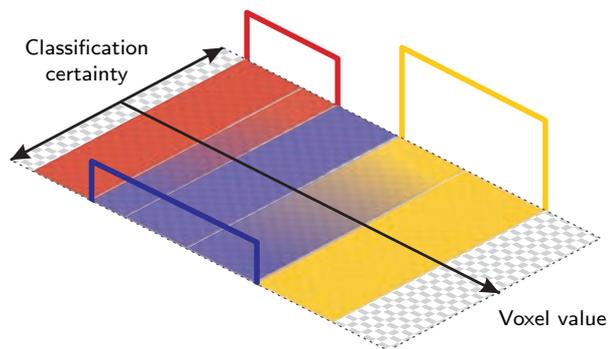


Figure 3.10: Classifying 2D TF. The traditional 1D TF is extended with a dimension corresponding to competitive classification certainty. Trapezoids for overlapping materials are placed at extreme classification values and the 2D TF is interpolated in between them.

The defined set of range weight criteria is combined into a voxel-specific attribute assessing the probability of either tissue, \mathbf{P} , named *competitive classification certainty*. The \mathbf{P} value is integrated in a TF model that constitutes a multidimensional TF in the rendering algorithm, while the user remains in the 1D TF paradigm, as illustrated in figure 3.10. The user defines 1D TFs for the individual tissues. When there are conflicting tissues in an intensity range, classification dimensions are implicitly added to the TF. Allowing at most two overlapping tissue ranges is not very limiting in practice, in that case only one dimension is added. The tissue TFs are placed at either extreme in the classification dimension, so that overlapping tissues' TFs are always on opposite sides, see figure 3.10. In overlapping ranges, the classification dimension is filled by linear interpolation of the 1D TFs at the extremes. Thus, the single classification dimension can harbor several tissue separation attributes, one for each overlapping range.

The range weight criteria have been shown to be effective for tissue separation in the medical volume rendering context. If this method is seen only as a classification algorithm, it must be considered to be quite coarse. It cannot, and should not, compete with advanced segmentation schemes. Its benefits come into play when a traditional 1D TF fails to provide a useful visualization and segmentation algorithms would require substantial manual input. This situation is quite common in clinical work with DVR, where indistinct boundaries constitute a major problem for general-purpose segmentation schemes.

3.2.3 Sorted histograms

An obstacle for further deployment of DVR in routine clinical use is the complexity of TF design. The radiologist must be at least as effective with DVR as with 2D slice viewing in order to adopt the new technique. Today many hospitals organize dedicated 3D labs producing DVR images for the radiologists, but then the radiologist is cut off from exploratory interaction of the data set, which may lead to lower quality in the diagnostic work.

TF design is normally quite complex and time-consuming, but it becomes much more so when a multidimensional TF (MDTF) is to be constructed. It is clear that supporting tools are needed, but the user is typically guided only by a data histogram across the intensity range. An improved histogram display, called *sorted histograms*, was proposed in paper VI along with a tailored workflow that simplifies MDTF design. The key benefit is to enable efficient interactive input of domain knowledge in the visualization process.

A 1D histogram has many benefits due to its simplicity. It provides an assessment of the data set at a glance, thanks to the apparent peak structure. These benefits may be lost when adding more dimensions, for instance the common 2D case of introducing gradient magnitude in addition to intensity. Consider the traditional 2D histogram in figure 3.11, top right. The occurrence of each combination of first and second dimension values is mapped to a color or grayscale value. It is difficult to identify the main characteristics of the data set by comparing voxel clusters in terms of occurrence and value ranges. Moreover, peaks in the intensity scale that are dominating in the 1D case may be smeared across the second dimension and thus lose distinctiveness. Finally, medical data sets typically have a fairly smooth distribution that makes the 2D histogram cluttered.

The *sorted histogram* (fig. 3.11, bottom left), is a 2D histogram display that keeps the 1D histogram shape for the intensity dimension. This approach ensures a direct understanding of the most useful information – the intensity characteristics. The relevant

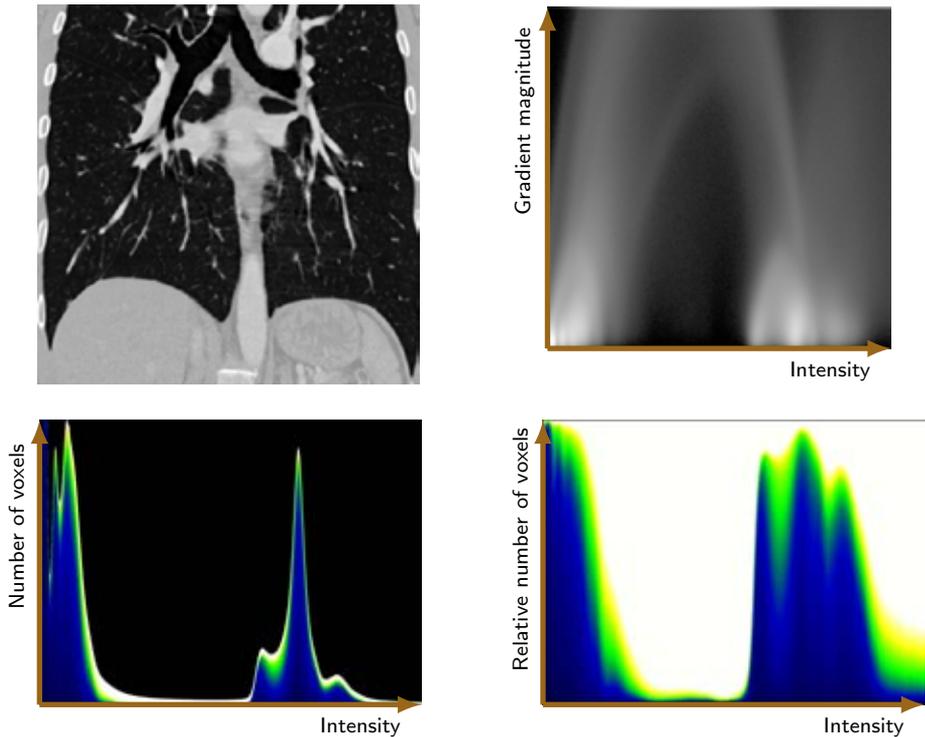


Figure 3.11: 2D histogram displays for intensity and gradient magnitude of a CT angiography representing a typical medical data set. Top left: A data set slice. Top right: The traditional 2D histogram is quite indistinct. The theory suggests that each material boundary should correspond to a distinct arc, but in reality the arcs are fuzzy and overlapping. Bottom left: The proposed *Sorted histogram*, a 1D intensity histogram is sorted and colored (blue-green-yellow-white) according to the second attribute. Bottom right: The proposed *Relative sorted histogram*, a sorted histogram where the relative distribution of the second attribute is shown.

information is the peak silhouette, which means that the area under the curve is free to be used to convey a second attribute. The sorted histogram exploits this fact by sorting the voxels according to the value of their second attribute and then color-coding each histogram column accordingly. The sorted histogram combines the simplicity of the 1D histogram with the additional information of the traditional 2D histogram.

A limitation of the sorted histogram is that the 1D histogram curve often has very different scale across the intensity value range. Therefore, a variant named *relative sorted histogram* has been developed (fig. 3.11, bottom right) as a complement. This relative version conveys changes in the distribution of the second attribute values. The only difference is that each histogram column is extended to cover the full y-axis, resulting in a histogram that always fills a rectangular area.

In addition to the histogram displays, a workflow has been defined that supports an arbitrary number of additional attributes, allowing construction of arbitrary MDTFs. To retain reasonable ease-of-use, the user works with a sequence of 2D histograms,

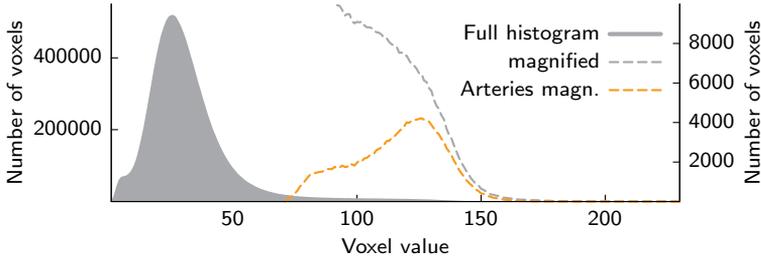


Figure 3.12: A global histogram of an MR data set containing a renal artery angiography, full and magnified scale. The main peak completely dominates the minor peak from the arteries (orange), which is the interesting tissue. In full scale, the artery peak is not visible.

studying one additional attribute at a time. The MDTF design process is then a matter of switching between second attributes while keeping the frame of reference given by the x-axis. The sorted histograms are particularly suitable for this workflow, since their shape does not change when changing second attribute. Materials are defined by direct range selections in the histogram display.

3.3 Spatial coherence in histograms

The next challenge to be addressed is that of tissue detection described in section 2.3. Parts of papers II, III, V, and VII are the basis for the following presentation.

The global histogram is the main guide for manual TF design in DVR. It is also used for tissue and feature detection in MRI and nuclear medicine data sets in order to define visualization parameters for 2D slice images. A fundamental assumption that motivates the use of histograms is that the tissues correspond to visible peaks. This is, however, often not the case and the tissue detection then becomes a time-consuming trial-and-error interaction. An example is seen in figure 3.12, in which there is a single peak consisting of unimportant features that dominates the diagnostically relevant information.

The global histogram is a very attractive tool due to its simplicity, as discussed in section 3.2.3. Its fundamental drawback is that all spatial relations in the data are lost. The idea of the developed histogram analysis techniques is to introduce some of the spatial information while retaining the simple histogram form. Both the α -histogram, described next, and the Partial Range Histogram, described in section 3.3.2, highlight the spatially coherent parts of the data, which typically corresponds to tissues and materials of interest.

3.3.1 α -histogram

The α -histogram, introduced in paper V, is a general enhancement of histograms that amplifies value ranges corresponding to spatially coherent regions. When the tissues of interest do not appear as peaks in the histogram, the α -histogram enhancement can be used to make them emerge. Once the peaks are distinct, the tissues are easily detected regardless of whether manual or automatic TF construction is used.

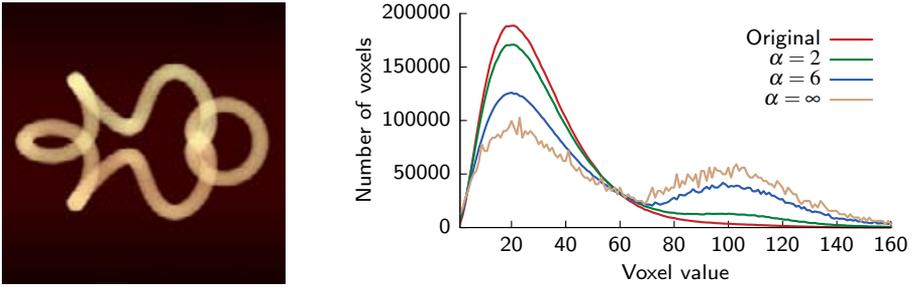


Figure 3.13: Left: Synthetic data set imitating a vessel. The background and the spiral have Gamma and Gaussian distributions, respectively. The “vessel” peak at value 100 is not visible. Right: Different α -histograms. For higher α the amplification of spatially coherent ranges makes the “vessel” peak visible.

The algorithm for creating an α -histogram is summarized as follows. The data set is divided into small neighborhoods for which individual histograms are retrieved. Each of these local histograms are raised to the power of α and they are then summed and normalized into an enhanced global histogram. Since coherent tissue causes local histogram peaks in the corresponding value ranges, these histogram values are proportionally amplified when raised to $\alpha > 1$. This means that value ranges containing spatially concentrated materials will be enlarged in the resulting histogram.

Mathematically formulated, the α -histogram is constructed as a sum of local histogram functions $H_n(N, x)$, counting the instances of the value x in a local neighborhood, see eq. 3.5. N is the set of data values in the neighborhood, D_x is the set of data values equalling x in the data set D . $|S|$ is the cardinality of set S .

$$H_n(N, x) = |N \cap D_x| \quad (3.5)$$

The neighborhood subdivision of the data set creates spatial regions N_1, \dots, N_k , the straightforward example being non-overlapping cubic blocks. The α -histogram $H_\alpha(x)$ is the sum of all the enhanced local histogram functions, see eq. 3.6. Setting $\alpha = \infty$ is equal to applying a maximum operator on the set of local histograms. An initial normalization is made by raising each sum to $1/\alpha$. It is also convenient to scale the α -histogram function to have the same area as the original.

$$H_\alpha(x) = \left(\sum_{i=1}^k H_n(N_i, x)^\alpha \right)^{1/\alpha} \quad (3.6)$$

The ability to amplify spatially coherent material is demonstrated in figure 3.13. The data set contains a synthetic vessel that is designed not to coincide with the cubic block structure. Even though the vessel value range is completely hidden in the original histogram, the α -histogram makes the interesting peak turn up.

In order to use a global histogram for automatic tissue detection, be it an ordinary histogram or an enhanced α -histogram, a peak detection algorithm is needed. In paper V such an algorithm is presented. The key components are a controlled smoothing of the histogram curve and an iterative consolidation of peaks retaining those that emerge the most from the background distribution.

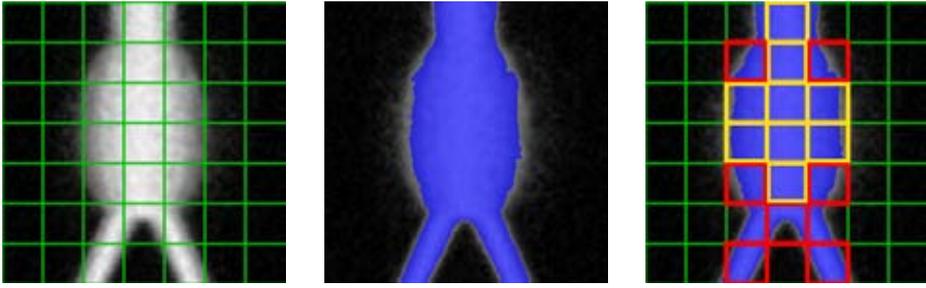


Figure 3.14: Partial Range Histogram (PRH), 2D example. Left: The image is subdivided into block neighborhoods. Middle: Pixels within a given intensity range are shown in blue. Right: A PRH is the histogram for the group of blocks exceeding a range weight threshold. The 80% threshold is shown in yellow, additional blocks for the 50% threshold shown in orange.

There are many possibilities to extend the basic formulation of the α -histogram. A priori information can be exploited through tailored processing of the local histograms. Furthermore, the method is not limited to non-overlapping cubic neighborhoods. For instance, the neighborhoods can be overlapping, anisotropic, and have spatially varying weight.

3.3.2 Partial Range Histogram

The Partial Range Histogram (PRH), first described in paper II, can be used to solve the same problems that the α -histogram targets: identifying tissues of interest when they are not distinct in the global histogram. Another similarity is that the PRH-based method also capitalizes on the spatial coherence of tissues. The methods are, however, quite different in other respects.

The definition of a PRH is: the histogram for a set of neighborhoods that are typical for a given intensity range. More specifically, a PRH groups all blocks that exceed a threshold in terms of range weight, defined in section 3.2.1. A 2D example is given in figure 3.14. As for α -histograms, the first step is to subdivide the data set, but in the PRH case the neighborhoods are limited to a non-overlapping subdivision. Again, the straightforward approach is to use cubic block neighborhoods.

In order to use a PRH for tissue identification, a relevant value range must be defined. One possibility is to let the user manually explore the global histogram by interacting with a range selection. The user can then either study the resulting PRH, or, more effectively, have a TF component that automatically adapts to the PRH shape. This is the idea behind the *adaptive trapezoid* introduced in paper II.

An automatic tissue detection scheme based on the PRH has also been developed. The scheme is unsupervised and exhaustively groups neighborhoods in an iterative approach. The scheme is summarized as follows:

1. Find the highest peak of the main histogram.
2. Create a PRH for the middle part of the peak range.
3. Remove the PRH from the main histogram.
4. Run steps 1-3 until the main histogram is empty.
5. Merge PRHs that are similar.

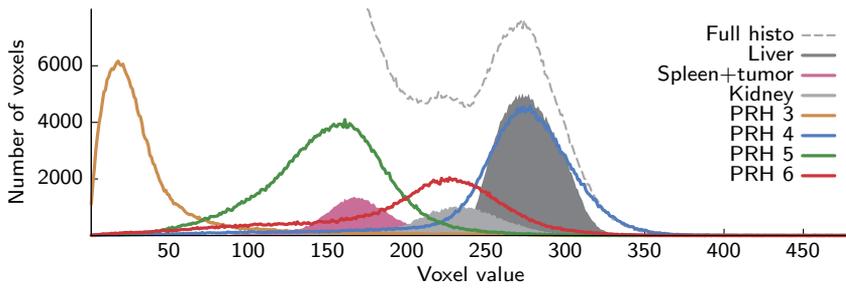


Figure 3.15: Automatic tissue detection using PRHs. In this MR bile duct data set, shown in figure 3.9, the positions of the peaks detected by the PRH-based scheme correspond well to the manually segmented peaks.

An example of the result of the automatic tissue detection for a challenging MR data set is given in figure 3.15. Manual segmentations of the main organs in the volume are used as ground truth. The apex values of the resulting PRHs correspond well to the sought-for tissues. There are significant discrepancies in peak height, but this does not affect the task of detecting the tissue value ranges.

A characteristic of many medical images is that tissues do not have a consistent value range, there is a spatial variation of the typical intensity value for a tissue. This can cause problems for tissue detection schemes, including the PRH-method. In paper VII a tissue detection refinement algorithm is presented. The input to the algorithm is value range definitions from a preliminary tissue detection, for instance in terms of mean value and standard deviation. The value ranges defining the tissues are then iteratively adapted in each spatial location according to the local presence of the tissue. The presence is determined as a range weight measure.

3.3.3 Evaluations

Both the tissue detection schemes, based on α -histogram and PRH, respectively, have been evaluated as tools for TF adaptation. The first evaluation, presented in paper III, was a clinical study of several approaches, where a version of the PRH-based scheme was one. The overall goal was to achieve automatic TF adaptation for MR examinations of abdominal aortic aneurysm (AAA). The diagnostic task in this case is to measure diameters of the aorta. A measurement error of more than 2 mm is usually considered clinically unacceptable, since the consequence could be that inadequate surgical material is chosen which is life-threatening for the patient. Different TF settings can, however, cause variations of the perceived diameter far above the acceptable level [PDE*04].

The detection scheme of section 3.3.2, extended with general a priori knowledge about the examination type, was applied to 20 highly different MR AAA data sets, predicting adequate TF settings. Three radiologists performed measurements using the resulting renderings. It was compared to MIP and DVR from other automatic TF adaptations, with manual measurements from digital subtraction angiography as the gold standard. The PRH method had the lowest average measurement error, 0.9 mm compared to 1.4–1.8 mm for the other methods. Furthermore, the PRH method achieved

results for all data sets, which was not the case for one of the other TF prediction methods based on global histogram shape analysis.

A second evaluation was presented in paper V, where PRH-based and α -histogram-based TF adaptation were compared. The same 20 MR AAA data sets were used, but this time the evaluation was in terms of difference in the apex position of the detected peaks compared to the ground truth peaks. Manual segmentations validated by a radiologist were used as ground truth. The results for the PRH and α -histogram detections were similar. They both showed a high level of accuracy and robustness in the peak detection, a great improvement over peak detection in the original global histogram.

Even though there was no significant difference in performance for this peak detection task, the two methods have different characteristics. When erroneous, the PRH approach can lead to very large errors, since incorrect decisions in early iterations may propagate. In comparison, the α -histogram has never proved to be misleading even though the result may be inconclusive. On the other hand, the PRH effectively isolates individual components of the histogram, enabling analysis of highly overlapping tissue ranges, which is more difficult for the α -histogram.

3.4 Probabilistic animation

There are two major drawbacks of the manual uncertainty assessments employed for conventional DVR in diagnostic use, as discussed in section 2.4. The first drawback is that the procedure is inefficient and the second is that relevant alternatives may be missed.

The main method in this section uses animation to convey diagnostically important uncertainty. Again, medical domain knowledge encoded in the TF is exploited. This time the knowledge used is something that the physicians may be providing unknowingly while designing TFs: a probabilistic tissue classification. The description below starts with an analysis of probabilistic aspects of the TF, which lays the foundation for the animation scheme described in the following section. The presentation in this section mainly corresponds to paper IX.

3.4.1 Probabilistic Transfer Functions

The term “probabilistic TF” can be interpreted in several ways. One line of research in uncertainty visualization aims to create tools to convey probabilistic information that are the result of a separate classification algorithm. An example of such an approach is a partial result in paper VII, where the competitive classification certainty (section 3.2.2) is used to specifically visualize uncertain classifications, see figure 3.16. This effect is achieved by adding a third 1D TF in the middle of the classifying dimension in the 2D TF, applying linear interpolation to the tissue TFs at the extremes.

Another way of defining probabilistic TFs is to focus on the TF as a classification tool, to study the inherent classification within the visualization process. This is the basis for the remainder of this section.

An important categorization of TF models that is often overlooked is whether material probabilities are an *implicit* or *explicit* part of the mapping from sample value to visual appearance. In the explicit case the application of a TF is a two-step approach. First the sample value s is mapped to a set of material probabilities, $p_m(s)$, where m is the index among the M materials. Then the material probabilities are used to combine the individual material colors $\mathbf{c}_m = (r_m, g_m, b_m, \alpha_m)^T$ into the wanted sample color

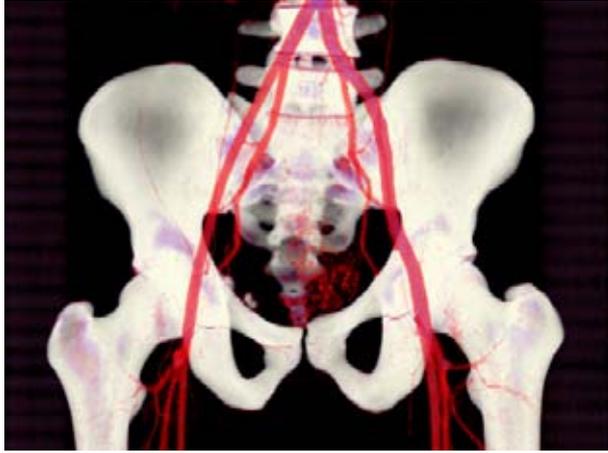


Figure 3.16: Visualizing uncertainty for a CT pelvis examination where vessel and soft bone intensity ranges overlap. Bone is white, vessels are red and uncertain regions are shown in purple.

$\mathbf{c}(s)$. The initial DVR implementation by Drebin et al. [DCH88] employed an explicitly probabilistic approach, according to equations 3.7 and 3.8. Transparent regions are referred to as the *null material* ($m = 0$).

$$\mathbf{c}(s) = \sum_{m=0}^M p_m(s) \cdot \tilde{\mathbf{c}}_m \quad (3.7)$$

$$\tilde{\mathbf{c}}_m = (\alpha_m r_m, \alpha_m g_m, \alpha_m b_m, \alpha_m)^T \quad (3.8)$$

The other view of TF models is the implicitly probabilistic one, where the sample value is mapped directly to sample color. This has for many years been the dominating approach. A common visual representation for uncertain materials is to connect them to low opacity. A typical example is the trapezoid or ramp shapes used for TF components, where the opacity is gradually decreased outside of the main value range of the material.

The uncertainty visualization challenge addressed in this thesis is to explore relevant alternative renderings, for a given TF. Since the explicitly probabilistic representation provides a well-defined uncertainty domain, which is not the case for the implicitly probabilistic approach, the TF model developed is explicitly probabilistic, see figure 3.17. Each material to be visualized is connected to an individual TF component consisting of two parts. The first is the material appearance and the second is the classifying function $\tilde{p}_m(s)$ that maps intensity to material likelihood. The intensity-specific material probabilities are the normalized likelihoods:

$$p_m(s) = \frac{\tilde{p}_m(s)}{\sum_{m'=0}^M \tilde{p}_{m'}(s)} \quad (3.9)$$

It should be noted that the explicitly probabilistic TF is not limited to pure uncertainty visualization applications. There are significant benefits of this approach that makes it viable as the default interaction mode in medical volume rendering. This is further discussed in section 3.4.3.

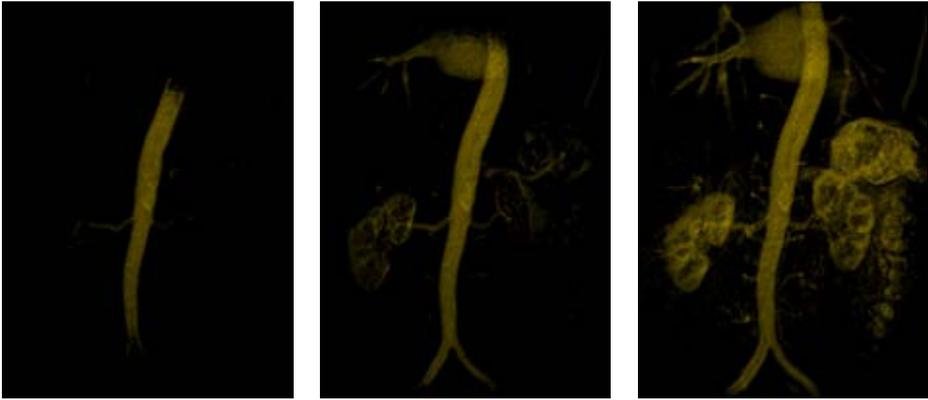


Figure 3.19: Uncertainty animation in material sync mode for a MR renal angiography. Three examples of animation frames connecting visibility to different vessel probability levels.

tries of each material color is not changed, different effects can be achieved by different ordering within the array. Two main types of animation mode have been developed: randomized and synchronized ordering. The latter guarantees that for each material there is a frame where its maximum extent is shown. The resulting animation has a pulsating appearance, see figure 3.19, compared to the unorganized flickering of the randomized mode.

Advanced exploration in the form of uncertainty visualization is seldom needed for the entire volume. The sensitivity analysis is typically wanted for a few essential regions. Based on this observation, a *sensitivity lens* application has been developed. The lens consists of a small user-controlled region-of-interest in which the selected animation technique is applied while the rest of the image is rendered traditionally. The sensitivity lens has great potential to become an important DVR tool for radiologists, since it provides ease-of-use while retaining exploratory abilities.

An evaluation of the benefits of uncertainty animation has been performed, targeting the clinically relevant and challenging area of stenosis assessment in MR angiographies. The form was an experimental study on simulated vessels with twelve experienced physicians as test subjects. The visualizations were set to have fixed and untuned TFs. In this situation, the synchronized animation mode enabled far higher accuracy than both the traditional static rendering and the randomized animation. The results with the synchronized animation are actually fairly close to the “gold standard” of a static rendering with free manual TF adjustments. This is a clear indication that time-consuming TF interaction could, to some extent, be replaced by uncertainty animation.

3.4.3 Probabilistic Transfer Functions revisited

The probabilistic TFs in section 3.4.1 lead to a discussion on how material probabilities are to be handled, with bearing on fundamental aspects of DVR. A first remark is that apart from being a foundation of the developed uncertainty animation, there are significant potential benefits of the TF model itself to be further explored. The separated probability component can be visualized in an arbitrary way. A promising example is

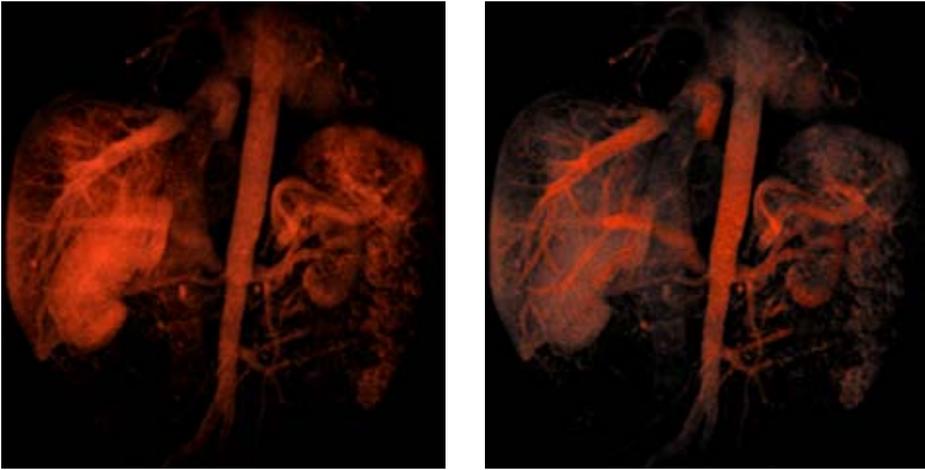


Figure 3.20: Enhanced visualization of an MR renal angiography by modeling uncertainty by desaturation. Left: In the traditional rendering it is hard to distinguish between thin tissue regions and regions with low tissue classification probability. Right: Desaturating uncertain regions using a probabilistic TF model. The TF consists of a single component describing the likelihood of vessel tissue.

to employ (non-animated) renderings where uncertainty is connected to color desaturation, see figure 3.20.

Another advantage of explicitly probabilistic TFs has been identified, which may prove to be an important evolution for clinical DVR. Today, the physicians intuitively look upon TF interaction more as a modification of the rendering appearance than as a classification process. During the evaluation of the uncertainty animation, several radiologists with expertise in DVR expressed that the classification mindset would be an important step towards more effective DVR interaction in diagnostic work. This mindset is promoted by the probabilistic approach.

An in-depth analysis of explicit probabilistic TFs leads to an intriguing discussion about adequate color mixing principles. Assume that there is equal probability, 50%, of materials A and B for a sample value s_0 . Material A is set to be red with high opacity and B blue with low opacity. The established way of finding the TF output for s_0 is to associate colors before mixing them, i.e., to pre-multiply each material color by its opacity and then mix them according to the probability [EHK*06]. This approach will be referred to as *pre-association*. Pre-association is illustrated in the left part of figure 3.21, with s_0 at the midpoint of the data set.

A main reason for this solution is that color-bleeding is avoided. The extreme case of color-bleeding occurs when one of the colors to be mixed has zero opacity. Without pre-association, the zero-opacity color would influence not only the opacity of the resulting mix, but also the pure color components. Thus, pre-associating the pure colors makes intuitive sense.

There are, however, equally valid arguments why the colors should be associated *after* the mixing. In our example above, the mixed color for s_0 will be highly biased towards red, the color of the high-opacity material. This is counter-intuitive: when materials A and B are equally probable, the resulting color should not be much closer to a pure sample of material A. The remedy is to decouple the opacity from the probabilis-

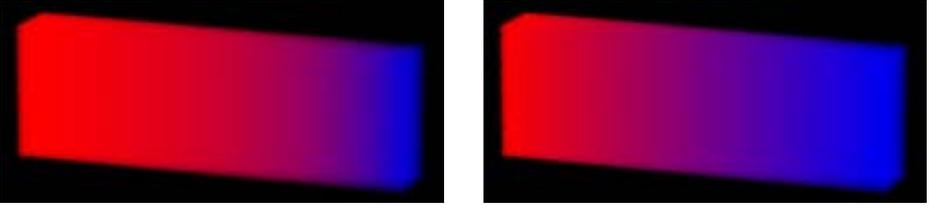


Figure 3.21: Limitations of color mixing with pre-associated colors in a probabilistic context. The data set is defined as a pure material at each end (left end: red, high opacity, right end: blue, low opacity). The material probabilities are linearly mixed in between, which means that there is equal probability of either material at the midpoint. Left: Color-mixing using pre-association. The resulting color is heavily biased towards the high-opacity material (red). Right: Color-mixing using post-association, according to equation 3.10. The visual representation matches the material probabilities without bias.

tic mixing in a *post-association* approach, as shown in eq. 3.10 (notation as in section 3.4, with p_m being material probability). The effect is shown in the right part of figure 3.21.

$$\mathbf{c}(s) = \bar{\alpha} \begin{pmatrix} \sum_{m=1}^M p_m(s) \cdot r_m \\ \sum_{m=1}^M p_m(s) \cdot g_m \\ \sum_{m=1}^M p_m(s) \cdot b_m \\ 1.0 \end{pmatrix}, \quad \bar{\alpha} = \sum_{m=1}^M p_m(s) \cdot \alpha_m \quad (3.10)$$

The different approaches boil down to two ways of interpreting material probability values for a voxel, as illustrated in figure 3.22. On one hand, they can be seen as a portion of the voxel occupied by the material, in other words a partial volume effect. This entails the predominant approach of pre-associated colors, since it is then correct that high-opacity materials dominate the voxel's appearance. On the other hand, there are many uncertainties that are not caused by partial voluming. A clinical example is the value range overlap of vessels and spongy bone in figure 3.16, these two tissues do not coincide spatially but the classification is still uncertain. Thus, material probabilities should, in these cases, be seen from a general statistical standpoint and post-association is a more appropriate color mixing principle.

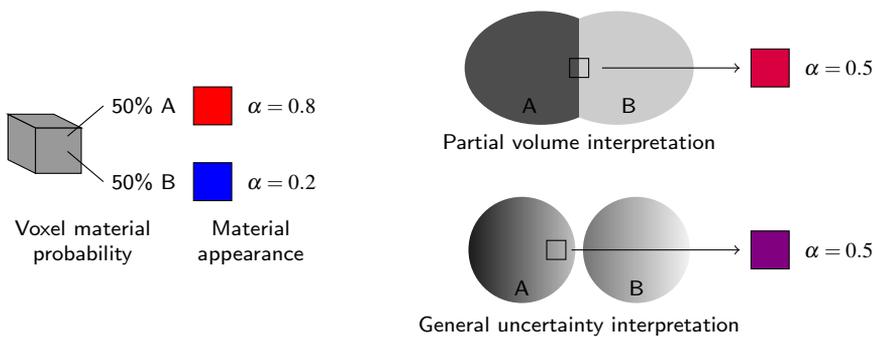


Figure 3.22: Alternative interpretations of material probability within a single voxel. Left: A voxel with equal probability of materials A and B. Material A is set to high-opacity red and B to low-opacity blue. Right, top: Traditional notion where the partial volume effect is considered to cause uncertainty, exemplified by a sample in between the objects corresponding to materials A and B. The resulting color is dominated by the high-opacity material. Right, bottom: Alternative notion that material probabilities are due to a general uncertainty, exemplified by overlapping value ranges for materials A and B and sampling in the uncertain range. The resulting color is neutral to opacity differences of the material appearances.

Chapter 4

Conclusions

In the context of this work, visualization research can be divided into two parts. One track has the objective to work with the technical methods to tune existing approaches and invent new ones. The other track is to build upon the application area of the visualization, to develop tools and find techniques that can make the most of the knowledge and work context of the visualization user. Both tracks are needed for optimal progress, but the purely technical level is, today, so high that the application-centered track seems to have the highest potential for major future advances in the field.

This thesis has described research in the application-centered track, devoted to the development of methods that contribute to the usefulness of DVR. Some concluding remarks will be made in this final chapter. First, the main research contributions of this thesis are summarized. A mindset constituting an adequate context for the proposed methods and future research is described next, followed by an outlook in the form of topics for future work.

4.1 Summarized contributions

Four main obstacles for efficient medical volume visualization in a clinical context have been identified and addressed in this thesis:

- An effective way of handling very large data sets has been presented in the form of a multiresolution DVR pipeline with an LOD selection based on TF content.
- Enhanced capabilities and lowered complexity in interactive tissue classification have been achieved by incorporating basic domain knowledge in the DVR tool set.
- Improved automatic and manual tissue detection for visualization purposes has been demonstrated in several schemes exploiting spatial relations in the data.
- Additional information relevant for data exploration has been achieved by automated uncertainty visualization based on probabilistic TFs.

The methods have been applied in clinical visualization scenarios regarding primarily CT and MR data sets. DVR is commonly applied to CT examinations but is rare for MR data due to the limited distinction between tissues and the noisy boundaries. Many problems concerning visualization of CT data have been addressed. A

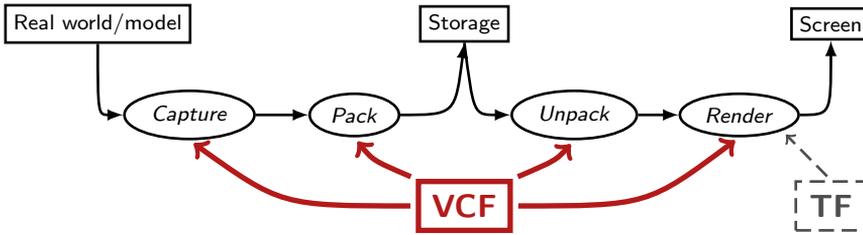


Figure 4.1: A Visualization Control Function (VCF) uses domain knowledge to optimize visualization performance and quality and can potentially affect the entire pipeline. In comparison, a Transfer Function (TF) is limited to the rendering stage.

potentially greater impact of the presented methods is to make DVR feasible for a wide range of MR examinations.

DVR employed in the medical domain is considered to be one of the true success stories in scientific visualization. There are many benefits of this volumetric technique compared to traditional slice-by-slice viewing, especially for exploration and demonstration of complex cases. Even though the clinical use of DVR represents a high impact compared to the deployment in other non-research settings, the usage in absolute terms is still low. A vast majority of the radiologists reviewing digital images do not use DVR or other 3D visualization tools. The methods in this thesis contribute to unleashing more clinical usefulness of volume visualization and comprise a step towards more wide-spread deployment in routine clinical work, ultimately being of benefit for patient care.

4.2 Beyond the Transfer Function

The methods presented in this thesis fit into a wider concept for medical visualization. As the fundamental idea is that the domain knowledge of the visualization user is the key to meeting both performance and quality demands, a new component in the rendering pipeline has been defined: the Visualization Control Function (VCF) [LLY06]. Being an evolution of the TF, the VCF aims to enable the full benefit of domain knowledge. The VCF can collect any type of useful input from the end user and apply it to all parts of the visualization pipeline for optimized efficiency and quality in the rendered image.

A comparison of the VCF and the TF in the context of the full visualization pipeline is shown in figure 4.1. There is much readily available information that is traditionally not included in the TF, but is included in the VCF framework: viewpoint, regions of interest, feature importance, and hardware limitations. A major argument for introducing the VCF concept as a mindset for DVR and other visualizations is that it promotes a holistic approach to visualization pipeline development. Optimal efficiency can only be achieved if the user's needs drive every part of the system. The data reduction method in section 3.1 and the classification method in 3.2 are two examples of method development stemming from this user-centered approach.

The output of the VCF includes the normal TF output, sample color and opacity, for subsequent compositing in the rendering. In addition, there are significance measures

for individual samples or for regions that steer data prioritization and reduction in the other pipeline components.

A final appealing property of the VCF is that it fully embraces further evolution. Any new visualization techniques, processing algorithms or GUI controls can be put into the VCF context and provide wider benefits than as isolated methods.

4.3 Future work

There are a number of interesting areas for further development of the methods and schemes presented in this work. In the multiresolution DVR pipeline, there is room for improvement of the significance measure driving the LOD selection. The histogram metadata is compact and efficient, but its robustness in different rendering situations can be enhanced. The distortion assessment for intermediate levels is another target for future work. Moreover, a pertinent challenge is to exploit the achieved data reduction for rendering performance speed-up. This has, to a large extent, been addressed in the GPU implementation developed by Ljung [Lju06a].

Another natural track for further investigation is the extension of the presented methods to time-varying and multivariate data sets.

The basic operators of the α -histogram would be interesting to employ for other tasks. Segmentation and classification algorithms for noisy data could probably benefit from adding components based on the α -histogram framework, even though no histogram enhancement is involved.

Regarding the probabilistic animation, there are some perceptually-oriented research tracks. A first question is how a color mixture achieved through animation is perceived in comparison with a traditional color interpolation. One hypothesis to be tested is whether the animation with increasing frame rate goes towards an interpolation in a perceptually calibrated color space, such as CIELUV. Another area is the representation of probability in the time domain, how should it be mapped in order to make the certainty correctly perceived.

The discussion of the different interpretations of material probabilities needs further attention. The traditional notion that the partial volume effect is the cause of uncertain material classification conflicts with the more abstract statistical standpoint where no assumptions on uncertainty causes are made. Since each interpretation entails different color-mixing principles, this issue is at the core of DVR. It may be impossible to find a single solution that unites this duality, but future work should aim for pragmatic ways of dealing with it.

Overall, clinical research is a main way forward to achieve the full potential of the methods presented in this work. For each clinical visualization challenge to be addressed, building blocks from the proposed technical foundation can be put together and tuned for the task at hand. The outcome of the clinical validations are likely to lead to refinement of existing methods and ideas for new ones.

Bibliography

- [Ada01] ADAMS M. D.: *The JPEG-2000 Still Image Compression Standard*. ISO/IEC (ITU-T SG8), September 2001. JTC 1/SC 29/WG 1: N 2412.
- [And03] ANDRIOLE K. P.: Addressing the coming radiology crisis: The Society for Computer Applications in Radiology Transforming the Radiological Interpretation Process (TRIP™) initiative. <http://www.siiimweb.org/WorkArea/showcontent.aspx?id=1206>, November 2003. Acquired August 2007.
- [BIP01a] BAJAJ C., IHM I., PARK S.: 3D RGB image compression for interactive applications. *ACM Transactions on Graphics* 20, 1 (January 2001), 10–38.
- [BIP01b] BAJAJ C., IHM I., PARK S.: Visualization-specific compression of large volume data. In *Proceedings Ninth Pacific Conference on Computer Graphics and Applications 2001* (2001), pp. 212–222.
- [BMF*03] BARTZ D., MAYER D., FISCHER J., LEY S., DEL RÌO A., THUST S., HEUSSEL C. P., KAUCZOR H.-U., STRASSER W.: Hybrid segmentation and exploration of the human lungs. In *Proceedings IEEE Visualization 2003* (2003), pp. 177–184.
- [BNS01] BOADA I., NAVAZO I., SCOPIGNO R.: Multiresolution volume visualization with a texture-based octree. *The Visual Computer* 17 (2001), 185–197.
- [BPS97] BAJAJ C., PASCUCCI V., SCHIKORE D. R.: The Contour Spectrum. In *Proceedings IEEE Visualization 1997* (1997), pp. 167–173.
- [CMPS97] CIGNONI P., MONTANI C., PUPPO E., SCOPIGNO R.: Multiresolution representation and visualization of volume data. *IEEE Transactions on Visualization and Computer Graphics* 3, 4 (1997).
- [DCH88] DREBIN R. A., CARPENTER L., HANRAHAN P.: Volume rendering. In *Proceedings of Computer Graphics and Interactive Techniques* (1988), vol. 22, ACM SIGGRAPH, pp. 65–74.
- [DKLP01] DJURCILOV S., KIM K., LERMUSIAUX P. F. J., PANG A.: Volume rendering data with uncertainty information. In *Data Visualization '01* (2001), Ebert D., Favre J., Peikert R., (Eds.), Springer, pp. 243–252, 355–356.
- [EHK*06] ENGEL K., HADWIGER M., KNISS J., REZK-SALAMA C., WEISKOPF D.: *Real-Time Volume Graphics*. A.K. Peters, Ltd, 2006.
- [ESG97] EHLSCHLAEGER C. R., SHORTRIDGE A. M., GOODCHILD M. F.: Visualizing spatial data uncertainty using animation. *Computers & Geosciences* 23 (May 1997), 387–395.
- [FR94] FOLEY J., RIBARSKY B.: Scientific visualization software. In *Scientific Visualization: Advances and Challenges*, Rosenblum L., Earnshaw R., Encarnacao J., Hagen H., Kaufman A., Klimenko S., Nielson G., Post F., Thalmann D., (Eds.). Academic Press, 1994, pp. 103–128.
- [Ger92] GERSHON N. D.: Visualization of fuzzy data using generalized animation. In *Proceedings of IEEE Visualization '92* (1992), pp. 268–273.

- [GHJA05] GAO J., HUANG J., JOHNSON C. R., ATCHLEY S.: Distributed data management for large volume visualization. In *Proceedings IEEE Visualization 2005* (2005), IEEE, pp. 183–189.
- [GHSK03] GAO J., HUANG J., SHEN H.-W., KOHL J. A.: Visibility culling using plenoptic opacity functions for large volume visualization. In *Proceedings IEEE Visualization 2003* (2003), IEEE, pp. 341–348.
- [GR04] GRIGORYAN G., RHEINGANS P.: Point-based probabilistic surfaces to show surface uncertainty. *IEEE Transactions on Visualization and Computer Graphics* 10 (2004), 564–573.
- [GS04] GUTHE S., STRASSER W.: Advanced techniques for high quality multiresolution volume rendering. In *Computers & Graphics* (February 2004), vol. 28, Elsevier Science, pp. 51–58.
- [GWGS02] GUTHE S., WAND M., GONSER J., STRASSER W.: Interactive rendering of large volume data sets. In *Proceedings IEEE Visualization 2002* (2002), pp. 53–60.
- [HHKP96] HE T., HONG L., KAUFMAN A., PFISTER H.: Generation of transfer functions with stochastic search techniques. In *Proceedings IEEE Visualization 1996* (1996), pp. 227–234.
- [HKG00] HLADÚVKA J., KÖNIG A. H., GRÖLLER E. M.: Curvature-based transfer functions for direct volume rendering. In *Proceedings Spring Conference Computer Graphics 2000* (2000), vol. 16, pp. 58–65.
- [IP99] IHM I., PARK S.: Wavelet-based 3D compression scheme for interactive visualization of very large volume data. *Computer Graphics Forum* 18, 1 (1999), 3–15.
- [ITU92] ITU-T/CCITT: *INFORMATION TECHNOLOGY – DIGITAL COMPRESSION AND CODING OF CONTINUOUS-TONE STILL IMAGES – REQUIREMENTS AND GUIDELINES*, September 1992. Recommendation T.81.
- [Joh04] JOHNSON C.: Top scientific visualization research problems. *IEEE Computer Graphics and Applications* 24, 4 (July/August 2004), 13–17.
- [JS03] JOHNSON C. R., SANDERSON A. R.: A next step: Visualizing errors and uncertainty. *IEEE Computer Graphics and Applications* 23, 5 (September/October 2003), 6–10.
- [KD98] KINDLMANN G., DURKIN J. W.: Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings IEEE Symposium on Volume Visualization* (1998), pp. 79–86.
- [KG01] KÖNIG A. H., GRÖLLER E. M.: Mastering transfer function specification by using VolumePro technology. In *Proceedings Spring Conference Computer Graphics 2001* (2001), vol. 17, pp. 279–286.
- [KKH02] KNISS J., KINDLMANN G., HANSEN C.: Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 8 (2002), 270–285.
- [Kön01] KÖNIG A.: *Usability Issues in 3D Medical Visualization*. PhD thesis, Technischen Universität Wien, Austria, 2001. PhD Thesis.
- [KUS*05] KNISS J. M., UITERT R. V., STEVENS A., LI G.-S., TASDIZEN T., HANSEN C.: Statistically quantitative volume visualization. In *Proceedings IEEE Visualization 2005* (2005), pp. 287–294.
- [KWMT03] KINDLMANN G., WHITAKER R., TASDIZEN T., MÖLLER T.: Curvature-based transfer functions for direct volume rendering: Methods and applications. In *Proceedings IEEE Visualization 2003* (2003), pp. 513–520.
- [LCR01] LUO M. R., CUI G., RIGG B.: The development of the CIE 2000 colour-difference formula: CIEDE2000. *Color Research and Application* 26 (2001), 340–350.

- [Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications* 8, 5 (1988), 29–37.
- [LHJ99] LAMAR E. C., HAMANN B., JOY K. I.: Multiresolution techniques for interactive texture-based volume visualization. In *Proceedings IEEE Visualization 1999* (1999), pp. 355–362.
- [LHJ03] LAMAR E. C., HAMANN B., JOY K. I.: Efficient error calculation for multiresolution texture-based volume visualization. In *Hierarchical and Geometrical Methods in Scientific Visualization*, Farin G., Hamann B., Hagen H., (Eds.). Springer-Verlag, Heidelberg, Germany, 2003, pp. 51–62.
- [Lju06a] LJUNG P.: Adaptive sampling in single pass, GPU-based raycasting of multiresolution volumes. In *Proceedings Eurographics/IEEE Workshop on Volume Graphics 2006* (2006), pp. 39–46,134.
- [Lju06b] LJUNG P.: *Efficient Methods for Direct Volume Rendering of Large Data Sets*. PhD thesis, Linköping University, SE-581 83 Linköping, Sweden, October 2006.
- [LLY06] LUNDSTRÖM C., LJUNG P., YNNERMAN A.: *The Visualization Control Function – The natural evolution of the Transfer Function*. Tech. rep., Department of Science and Technology, Linköping University, 2006.
- [MAB*97] MARKS J., ANDALMAN B., BEARDSLEY P., FREEMAN W., GIBSON S., HODGINS J., KANG T., MIRTICH B., PFISTER H., RUML W., RYALL K., SEIMS J., SHIEBER S.: Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proceedings SIGGRAPH 1997* (1997), pp. 389–400.
- [Mur92] MURAKI S.: Approximation and rendering of volume data using wavelet transforms. In *Proceedings IEEE Visualization 1992* (1992), pp. 21–28.
- [NH92] NING P., HESSELINK L.: Vector quantization for volume rendering. In *ACM Symposium on Volume Visualization* (1992), pp. 69–74.
- [NS01] NGUYEN K. G., SAUPE D.: Rapid high quality compression of volume data for visualization. *Computer Graphics Forum* 20, 3 (2001).
- [NU99] NYÚL L. G., UDUPA J. K.: On standardizing the MR image intensity scale. *Magnetic Resonance in Medicine* 42 (1999), 1072–1081.
- [Oth06] OTHBERG F.: *Standardized Volume Rendering Protocols for Magnetic Resonance Imaging using Maximum-Likelihood Modeling*. Master's thesis, Linköping University, SE-581 83 Linköping, Sweden, 2006. LITH-ITN-MT-EX-06/008-SE.
- [PBM05] PRADHAN K., BARTZ D., MUELLER K.: *SignatureSpace: A Multidimensional, Exploratory Approach for the Analysis of Volume Data*. Tech. rep., Department of Computer Science (WSI), University of Tübingen, 2005.
- [PBSK00] PFISTER H., BAJAJ C., SCHROEDER W., KINDLMANN G.: Panel: The transfer function bake-off. In *Proceedings IEEE Visualization 2000* (2000), pp. 523–526.
- [PDE*04] PERSSON A., DAHLSTRÖM N., ENGELLAU L., LARSSON E.-M., BRISMAR T., SMEDBY Ö.: Volume rendering compared with maximum intensity projection for magnetic resonance angiography measurements of the abdominal aorta. *Acta Radiologica* 45 (2004), 453–459.
- [PW03] PATRA A., WANG M.: Volumetric medical image compression and reconstruction for interactive visualization in surgical planning. In *Proceedings Data Compression Conference 2003* (March 2003), p. 442.
- [PWH01] PEKAR V., WIEMKER R., HEMPEL D.: Fast detection of meaningful isosurfaces for volume data visualization. In *Proceedings IEEE Visualization 2001* (2001), pp. 223–230.

- [PWL97] PANG A. T., WITTENBRINK C. M., LODHA S. K.: Approaches to uncertainty visualization. *The Visual Computer 13* (November 1997), 370–390.
- [RBS05] ROETTGER S., BAUER M., STAMMINGER M.: Spatialized transfer functions. In *Proceedings IEEE/EuroGraphics Symposium on Visualization* (2005), pp. 271–278.
- [RJ99] RHEINGANS P., JOSHI S.: Visualization of molecules with positional uncertainty. In *Data Visualization '99*, Gröller E., Löffelmann H., Ribarsky W., (Eds.). Springer-Verlag Wien, 1999, pp. 299–306.
- [RLBS03] RHODES P. J., LARAMEE R. S., BERGERON R. D., SPARR T. M.: Uncertainty visualization methods in isosurface rendering. In *Eurographics 2003, Short Papers* (2003), Chover M., Hagen H., Tost D., (Eds.), pp. 83–88.
- [Rob00] ROBB R. A.: Three-dimensional visualization in medicine and biology. In *Handbook of Medical Imaging – Processing and Analysis*, Bankman I. N., (Ed.). Academic Press, 2000, pp. 685–712.
- [Rön95] RÖNTGEN W.: Über eine neue art von strahlen (vorläufige mittheilung). In *Sitzungsberichte der physikalischmedizinischen Gesellschaft zu Würzburg* (1895), pp. 132–141.
- [RSHSG00] REZK-SALAMA C., HASTREITER P., SCHERER J., GREINER G.: Automatic adjustment of transfer functions for 3D volume visualization. In *Proceedings Vision, Modeling and Visualization '00* (2000), pp. 357–364.
- [RSKK06] REZK-SALAMA C., KELLER M., KOHLMANN P.: High-level user interfaces for transfer function design with semantics. *IEEE Transactions on Visualization and Computer Graphics 12* (2006), 1021–1028.
- [SBS02] SOHN B.-S., BAJAJ C., SIDDAVANAHALLI V.: Feature based volumetric video compression for interactive playback. In *Proceedings IEEE Visualization 2002* (2002).
- [ŠVSG06] ŠEREDA P., VILANOVA BARTROLÍ A., SERLIE I. W. O., GERRITSEN F. A.: Visualization of boundaries in volumetric data sets using LH histograms. *IEEE Transactions on Visualization and Computer Graphics 12*, 2 (2006), 208–218.
- [SW03] SCHNEIDER J., WESTERMANN R.: Compression domain volume rendering. In *Proceedings IEEE Visualization 2003* (2003).
- [SWB*00] SATO Y., WESTIN C.-F., BHALERAO A., NAKAJIMA S., SHIRAGA N., TAMURA S., KIKINIS R.: Tissue classification based on 3D local intensity structures for volume rendering. *IEEE Transactions on Visualization and Computer Graphics 6*, 2 (2000), 160–179.
- [TLM01] TENGINAKAI S., LEE J., MACHIRAJU R.: Salient iso-surface detection with model-independent statistical signatures. In *Proceedings IEEE Visualization 2001* (2001), pp. 231–238.
- [TLM03] TZENG F.-Y., LUM E. B., MA K.-L.: A novel interface for higher-dimensional classification of volume data. In *Proceedings IEEE Visualization 2003* (2003), pp. 505–512.
- [Wes94] WESTERMANN R.: A multiresolution framework for volume rendering. In *1994 Symposium on Volume Visualization* (October 1994).
- [WWH*00] WEILER M., WESTERMANN R., HANSEN C., ZIMMERMAN K., ERTL T.: Level-of-detail volume rendering via 3D textures. In *Proceedings IEEE Volume Visualization and Graphics Symposium 2000* (2000), ACM Press, pp. 7–13.
- [YL95] YEO B.-L., LIU B.: Volume rendering of DCT-based compressed 3D scalar data. *IEEE Transactions on Visualization and Computer Graphics 1* (March 1995), 29–43.

- [ZSS*03] ZOROOFI R. A., SATO Y., SASAMA T., SUGANO N., YONENOBU K., YOSHIKAWA H., OCHI T., TAMURA S.: Automated segmentation of acetabulum and femoral head from 3-D CT images. *IEEE Transactions on Information Technology in Biomedicine* 7, 4 (December 2003), 329–343.