

Dissertation

Real-time Processing for Advanced Ultrasound Visualization

Christian Ulrich Schulte zu Berge





Technische Universität München

Fakultät für Informatik

Lehrstuhl für Informatikanwendungen in der Medizin

Real-time Processing for Advanced Ultrasound Visualization

Christian Ulrich Schulte zu Berge

Vollständiger Abdruck der von der Fakultät für Informatik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender: Prof. Dr.-Ing. Darius Burschka

Prüfer der Dissertation:

1. Prof. Nassir Navab, Ph.D.
2. Prof. Dr.-Ing. Bernhard Preim
Otto-von-Guericke-Universität Magdeburg

Die Dissertation wurde am 15.12.2015 bei der Technischen Universität München eingereicht und durch die Fakultät für Informatik am 23.04.2016 angenommen.

Christian Ulrich Schulte zu Berge

Real-time Processing for Advanced Ultrasound Visualization

Dissertation, angenommen am 23.04.2016

Technische Universität München

Fakultät für Informatik

Lehrstuhl für Informatikanwendungen in der Medizin

Boltzmannstraße 3

85748 and Garching bei München

Abstract

The continuous advances in the field of medical imaging constantly provide clinicians with new technology promising to facilitate diagnosis, interventions and therapies and improve their outcome. Inherent with this progress is that the amount of available data, its resolution and quality is significantly increasing. However, such improvements are only of actual use if the data can be adequately presented to the clinician. Thus, visual computing in medicine is an important field of research. This thesis focuses on the development of such techniques specifically designed for medical ultrasound. One of the key strengths of ultrasound is its real-time capability allowing the clinician to interactively examine the patient's anatomy given the visual feedback while manipulating the ultrasound transducer. Therefore, the methods for advanced ultrasound image processing and visualization presented in this thesis are designed to ensure maintaining this interactivity.

At first, techniques for improved 2D B-mode ultrasound visualization are presented, where the observer is provided with real-time feedback on the uncertainty present in the image. Therefore, in a first step, the necessary information is generated through the computation of ultrasound Confidence Maps estimating per-pixel uncertainty information with respect to the ultrasound signal attenuation. While their original formulation has been proven to improve tasks such as image registration and segmentation, their valuable information has never been directly exposed to the clinician. Therefore, a real-time capable extension in form of an incremental solver scheme is presented and shown to yield precise results for visualization applications. The uncertainty information is eventually fused with the original B-mode images using three carefully chosen visualization schemes. Their benefit has been shown for both educational and clinical applications.

The second part of this thesis' contributions targets improved 3D ultrasound visualization. In order to generate high-quality 3D volumes from tracked ultrasound sweeps, an incremental compounding scheme is presented, which combines orientation-driven correlation terms and the aforementioned Confidence Maps in an information fusion approach. The subsequent visualization of such ultrasound volumes is a particularly challenging task due to the special nature of B-mode intensities. Therefore, a novel classification concept in the form of point predicates is presented, which seamlessly integrates into the standard direct volume rendering pipeline. In conjunction with the proposed predicate histogram as intuitive user interface, this technique allows for much more meaningful visualization and facilitates the understanding of the data for the observer.

Zusammenfassung

Dank des kontinuierlichen Fortschritts auf dem Gebiet der medizinischen Bildgebung haben Ärzte heutzutage eine große Bandbreite an verschiedenen Technologien zur Auswahl, die Diagnosen, Eingriffe und Therapien vereinfachen und deren Ausgang verbessern können. Damit einher geht auch eine deutliche Erhöhung der zur Verfügung stehenden Daten, sowie deren Auflösung und Qualität. Solche Verbesserungen bieten allerdings nur dann einen Nutzen, wenn die Daten auch adäquat visualisiert werden können. Dies macht Visual Computing in der Medizin zu einem wichtigen Forschungsfeld. Diese Dissertation behandelt die Entwicklung solcher Techniken speziell für die Anwendung mit medizinischem Ultraschall. Eine der größten Stärken von Ultraschall ist die Echtzeitfähigkeit, die es dem Arzt ermöglicht Untersuchungen interaktiv auf Basis der visuellen Rückmeldung durchzuführen. Die in dieser Arbeit vorgestellten Techniken für fortgeschrittene Ultraschall Bildverarbeitung und Visualisierung wurden speziell in der Hinsicht entwickelt diese Interaktivität zu erhalten.

Im ersten Teil dieser Dissertation werden Verfahren zur Verbesserung von 2D B-mode Ultraschall Visualisierung vorgestellt. Um dem Anwender Echtzeit-Rückmeldung über die im Bild vorhandene Unsicherheit zu geben, wird diese Information in Form von Ultraschall Confidence Maps berechnet. Diese beschreiben für jeden Bild-Pixel die Abschwächung des Ultraschallsignales. Bis heute wurde jedoch solch eine Information nie dem Arzt als direkte Rückmeldung zur Verfügung gestellt. Daher wird in dieser Arbeit zunächst ein echtzeitfähiges, präzises Löser-Schema für die zugrunde liegenden Gleichungssysteme entwickelt, sowie drei neue Visualisierungsschemata vorgestellt, die die Unsicherheitsinformation mit den originalen B-mode Ultraschallbildern fusionieren. Die ausführliche Evaluation der Methoden zeigt den Nutzen dieser Technik sowohl für didaktische als auch für klinische Anwendungen.

Der zweite Teil dieser Dissertation behandelt Methoden zur verbesserten 3D Ultraschall Visualisierung. Um qualitativ hochwertige 3D-Volumen aus tracked-Ultraschall Aufnahmen zu erzeugen, wird zunächst ein inkrementelles Compounding-Schema entwickelt, welches ultraschallspezifische Korrelationsgrößen mit den zuvor eingeführten Confidence Maps kombiniert. Aufgrund der besonderen Eigenschaften von B-mode Ultraschallintensitäten, ist die anschließende Visualisierung solcher Ultraschall Volumen jedoch besonders herausfordernd. Daher wird ein neues Klassifikationskonzept auf Basis von Prädikaten vorgestellt, welches sich nahtlos in Standard Direct Volume Rendering Algorithmen integrieren lässt. In Verbindung mit dem ebenfalls entwickelten Predicate Histogramm als intuitive Benutzerschnittstelle, ermöglicht diese Technik deutlich aussagekräftigere Visualisierungen und eine einfachere Interaktion mit den Daten.

Acknowledgments

I would like to thank Prof. Navab for his constant trust in me and my work. In particular retrospectively, I appreciate the incredible amount of freedom he provided to me regarding my work and his support in pursuing my own ideas. Second, I would like to express my deep gratitude to my sister Marion who, recently successfully pursuing her own PhD, acted as mentor for me during the last four years. The value of our numerous talks in terms of ideas and mental support is hard to put in words.

Furthermore, I would like to thank Athanasios Karamalis, Christoph Hennersperger and Maximilian Baust for their highly valuable support as senior colleagues. Even if it sometimes might only have been an open ear or a small hint, I feel that your constant encouragement to pursue my own ideas contributed significantly to my success. At this place I would also like to mention Timo Ropinski whose excellent lectures and lab course back in Münster University actually raised my interest in medical visualization and thus made this work possible in the first place.

Finally, I want to express my gratitude to my family and friends for supporting me in everything I do, for providing help and advice in both small and big things, and especially for ensuring that my life is a happy one.

Contents

I	Introduction and Background on Medical Visualization	1
1	Introduction	3
1.1	Motivation	3
1.2	Organization	4
2	Visualization	7
2.1	Disambiguation	8
2.2	Goals and History of Visualization	9
2.3	The Visualization Pipeline	12
2.4	Data Characteristics	13
2.5	Evaluation of Visualization Methods	14
3	Medical Visualization	19
3.1	Imaging Modalities	19
3.1.1	Radiography	20
3.1.2	X-Ray Computed Tomography	20
3.1.3	Magnetic Resonance Imaging	21
3.1.4	Ultrasound	22
3.1.5	Nuclear Imaging	23
3.2	Basic Visualization Techniques for Medical Data	23
3.2.1	Visualization of Scalar Data	23
3.2.2	Visualization of High-dimensional Medical Data	28
3.2.3	Multi-modal Visualization	31
3.3	Ultrasound Visualization	34
3.3.1	Acquisition and Processing	34
3.3.2	Mapping and Rendering	36
4	Perception	39
4.1	Visual Attention	39
4.2	Perceptual Organization	41
4.3	Perceiving Depth and Size	42
4.4	Color Perception	45
II	Methodology	47
5	Real-time Ultrasound Uncertainty Estimation	49
5.1	The B-mode Ultrasound Image Formation Process	50
5.2	Confidence Maps for B-mode Ultrasound	51
5.3	Traditional Implementations	53

5.4	Real-Time Incremental Solving	53
5.5	Evaluation	55
5.6	Conclusion	57
6	Uncertainty Visualization for 2D Ultrasound	61
6.1	Uncertainty in Health Care and Medical Visualization	61
6.2	Clinical Significance and Application	63
6.3	Selection of Visual Variables	64
6.4	Visualization Schemes	65
6.4.1	Uncertainty as Color Overlay	66
6.4.2	Uncertainty Mapping to Chroma	67
6.4.3	Uncertainty Mapping to Fuzziness	68
6.5	Results and Evaluation	69
6.5.1	User Study with Ultrasound Novices	69
6.5.2	User Study with Ultrasound Experts	71
6.6	Conclusion	72
7	Advanced Ultrasound Compounding	75
7.1	Review of Ultrasound Compounding Techniques	76
7.2	Correlation Terms for Tracked Ultrasound Sweeps	78
7.3	Orientation-driven Inter-frame Registration of Ultrasound Sweeps	79
7.4	An Information-fusion Approach to Ultrasound Compounding	80
7.4.1	Clustering of Ultrasound Sweeps by Direction	81
7.4.2	Compounding of each Cluster	82
7.4.3	Uncertainty-based Fusion of the Compounded Clusters	82
7.5	Incremental Compounding System	82
7.5.1	Incremental Compounding Pipeline	83
7.5.2	Real-time Implementation	84
7.6	Results and Evaluation	85
7.6.1	Parameter Evaluation	85
7.6.2	Reconstruction Accuracy	85
7.6.3	Reconstruction Quality	86
7.6.4	System Performance	88
7.7	Conclusion	88
8	Predicate-based 3D Ultrasound Visualization	91
8.1	Background and Related Work	92
8.2	Point Predicates as Generic Concept for Classification	94
8.2.1	Predicate Library	95
8.2.2	Predicate Combination	96
8.3	The Predicate Histogram	97
8.4	Predicate-based 3D Ultrasound Visualization	98
8.4.1	Integration into the Standard DVR Pipeline	100
8.5	Scribble-based Predicate Histogram Setup	101
8.6	Results and Evaluation	102
8.7	Conclusion	105

III	Conclusion	107
9	Discussion and Conclusion	109
9.1	Discussion and Future Work	109
IV	Appendix	115
A	CAMPVis – A Game Engine-inspired Framework for Medical Imaging and Visualization	117
A.1	Motivation	117
A.2	Related Work	118
A.2.1	Entity Component System Architecture	119
A.3	CAMPVis Software Architecture	120
A.3.1	Build System	120
A.3.2	Package Architecture	121
A.3.3	The Entity Component System for CAMPVis	122
A.3.4	Properties	124
A.3.5	Module Architecture	124
A.4	Framework Features	125
A.4.1	Signal Manager	125
A.4.2	Factory Registration	127
A.4.3	Scripting Layer	127
A.4.4	Network Communication	128
A.5	Clinical Application	130
A.5.1	Real-time Uncertainty Visualization for 2D B-mode Ultrasound	130
A.5.2	Multi-modal Image-guided Prostate Biopsy	131
A.6	Conclusion	132
B	List of Authored and Co-authored Publications	135
C	Abstracts of Publications not Discussed in this Thesis	137
	Bibliography	141
	List of Figures	155
	List of Tables	159

Part I

Introduction and Background on Medical
Visualization

Introduction

Ever since the discovery of X-Ray imaging by Wilhelm Conrad Röntgen in the late 19th century, scientists have striven to find new ways of gaining in-vivo insight into the human anatomy and further improving existing imaging techniques. Nowadays, clinicians have a plethora of imaging modalities available for all aspects of their work, including diagnosis, treatment planning and interventional usage. Anatomical imaging modalities such as X-Ray Computed Tomography (CT), Magnetic Resonance Imaging (MRI) and Ultrasound (US) allow for a detailed examination of the the tissue structure while functional imaging modalities such as nuclear imaging or contrast agent-enhanced imaging provide insight into the metabolism and other (patho-) physiological processes within the body.

Ultrasound can be used for wide range of clinical applications including abdominal, cardiac, obstetric and vascular imaging. It exhibits a high spatial and temporal resolution, good soft tissue contrast and is free of ionizing radiation. Furthermore, compared to other imaging modalities, ultrasound devices are rather cheap and small so that sonography is often the primary choice for many examinations. However, ultrasound images are quite hard to interpret due to their noisy appearance and the presence of artifacts and shadowing. It requires a significant amount of experience by the clinician to perform a reliable diagnosis based on ultrasound. Thus, state-of-the-art techniques are needed for the entire processing and visualization pipeline in order to yield meaningful visualizations for the sonographer. This thesis presents novel approaches for visual computing in medical ultrasound covering several aspects of the processing pipeline and both 2D and 3D visualization.

1.1 Motivation

The progress in ultrasound imaging is driven by a close interdisciplinary collaboration between technical and medical experts who are constantly looking for innovative tools for clinicians. However, advanced techniques for the acquisition and processing of ultrasound data alone are only of limited usefulness. It requires visualization to transform such abstract data into a visual representation and to enable the observer to actually read and understand the data. In order to allow clinicians to gain maximum insight, advanced visualization techniques are carefully designed to extract the necessary information for a given task and present it in an intuitive way, taking human perception into account. Thus, visual computing and visualization are very important aspects of medical imaging.

One of the key strengths of ultrasound imaging is its real-time capability and interactivity. This enables the clinician not only to directly see the ultrasound image in the very moment of the acquisition but more importantly also to interact with the image and examine how the tissue is moving when changing acquisition parameters such as probe pressure or patient

positioning. In musculoskeletal applications, for example, it can be beneficial for the clinician to see how tendons and the surrounding tissue interact with each other during certain movements of the extremities. However, when introducing novel image processing or visualization techniques to a real-time modality such as ultrasound, it is crucial to maintain its interactivity. Thus, every single step of the imaging pipeline has to be real-time capable. In the course of this thesis, different techniques for visual computing in ultrasound are presented, each specifically designed to fit into the interactive ultrasound imaging pipeline.

From the perspective of scientific communities, this work is mainly related to the field of scientific visualization since the ultimate goal is an improved visualization. However, the presented methods consider the entire ultrasound imaging pipeline so that techniques from the community of medical imaging and computer aided interventions play an important role as well. Thus, this work is bridging the gap between these two fields and can be placed at their intersection with a considerable bias toward the community of scientific visualization.

1.2 Organization

This thesis is structured into three parts. As previously motivated, this work is mainly related to scientific visualization. Therefore, the first part introduces this research field and provides the reader with a thorough background and review of state-of-the-art techniques in medical visualization:

- *Chapter 2* will introduce the general scientific field of visualization, discuss its various definitions within the community and distinguish it from related fields such as computer graphics and image processing. It will further present general visualization concepts such as the visualization pipeline and data characteristics, as well as discuss how visualization techniques can be evaluated scientifically.
- *Chapter 3* will draw the focus onto medical visualization. Therefore, it provides an overview of the main medical imaging modalities available today and presents existing state-of-the-art methods for their visualization. Since this thesis is focused on visual computing for ultrasound, there is a separate section on ultrasound visualization.
- *Chapter 4* will discuss the human visual system and provide an overview on human perception theory. These insights are important as they can be used to create better visualizations.

The second part of this thesis will present novel methods for advanced ultrasound processing and visualization, first in the context of 2D B-mode ultrasound and later also for volumetric ultrasound:

- Since the availability of uncertainty information is a key aspect of many of the presented algorithms in this work, *Chapter 5* will introduce an incremental computation scheme for ultrasound Confidence Maps allowing for a per-pixel estimation of signal attenuation and uncertainty in the image. In contrast to existing implementations, our technique allows to generate this information in real-time and can therefore be integrated into the interactive ultrasound processing pipeline.

- *Chapter 6* builds on that and proposes different perceptual visualization techniques to expose the previously generated uncertainty information to the clinician. This allows ultrasound novices to understand the complex ultrasound image formation process more easily and helps expert sonographers with an interactive feedback on the image quality while they are manipulating the probe.
- *Chapter 7* discusses the transition from tracked 2D B-mode ultrasound sequences to 3D ultrasound volumes and presents an advanced orientation-driven approach to ultrasound compounding. Our method uses an uncertainty-based information fusion approach to allow for an incremental compounding during the acquisition, so that it can be integrated into the real-time ultrasound processing pipeline.
- Since the subsequent visualization of volumetric ultrasound is particularly challenging, *Chapter 8* introduces a novel generic formulation for classification in the form of predicates in combination with an intuitive user interface. On the technical side, this allows for the evaluation of both global and local information. Practically, the semantic nature of predicates also facilitates the interaction with the visualization for the user.

Finally, the third part will conclude this thesis with *Chapter 9* discussing the presented work in the context of the current state-of-the-art in medical visualization and presenting directions of future work.

All work presented in this thesis has been implemented in the CAMPVis software framework for medical visualization, which has been developed during the course of this thesis. Appendix A will present the initial design process, its software architecture, available features as well as clinical applications.

Even though this work is a PhD thesis and thus all work presented here is my own work if not explicitly declared otherwise, I will use the *we*-form throughout this work. On the one hand, this has stylistic reasons and, on the other hand, I would like to highlight that in my opinion research is not performed by oneself alone but it is always the product of many. Without the valuable input from friends, colleagues and the community, this work would not have been possible as such.

Visualization

” vi·su·al·ize, *verb*

1. *to form a picture of somebody/something in your mind*
2. *to make something visible to the eye*

— Oxford Advanced Learner’s Dictionary

Forming an all-encompassing definition of visualization is not a trivial task. In 1987, when computer-aided visualization started to emerge, McCormick et al. formulated [107]:

“Visualization is a method of computing. It transforms the symbolic into the geometric, enabling researchers to observe their simulations and computations. Visualization offers a method for seeing the unseen.”

In this statement, the authors assess two important aspects: First, they identify the core aspect of visualization, which is the mapping of abstract information into a graphical representation in order to make the information observable. Second, they nicely phrase that effective visualization has the power to show information and relationships, which would not be observable from the raw data itself.

Since both the amount and the density of data is continuously increasing, effective visualization becomes even more important as it can increase the throughput and speed at which the human brain can process information. Friedhoff and Kiley therefore write [53]:

“The standard argument to promote scientific visualization is that today’s researchers must consume ever higher volumes of numbers that gush, as if from a fire hose, out of supercomputer simulations or high-powered scientific instruments. If researchers try to read the data, usually presented as vast numeric matrices, they will take in the information at snail’s pace. If the information is rendered graphically, however, they can assimilate it at a much faster rate.”

One main reason for this inflation of information is the technical progress leading to constantly increasing amounts of instrument data (e.g. resolution) and furthermore shaping the development of novel technologies and modalities. At the same time, today’s complete digital interconnection makes information available even if it is physically located elsewhere. However, this sheer amount of available data is opposed to the limits of human cognition. Therefore, visualization provides effective tools to analyze, aggregate, reshape and present abstract information in a form aimed at facilitating its comprehension, and in the best case

even allowing the observer to discover otherwise hidden features or relationships. Good visualization exploits the power of human perception in order to reduce the observer's cognitive load.

With today's computational capabilities, the interaction aspect of visualization becomes more and more important, in particular considering the large amounts of data to process. In order to explore the available information, understand it correctly and eventually draw the right conclusions, users need to be able to interact with the visualization and adapt it to their needs. With this requirement in mind, Ben Shneiderman formulated his famous *Information Seeking Mantra* [169]:

“There are many visual design guidelines but the basic principle might be summarized as the Visual Information Seeking Mantra:

*Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand
Overview first, zoom and filter, then details-on-demand”*

This mantra iconically summarizes the elemental tasks he identified for visualization users. As a first step, the users have to get an overview of the entire data available. Based on the specific task, they then focus on the items of interest and filter out unneeded information. Since showing all details at once would lead to visual cluttering, it is important to provide details only when needed based on context information such as selections. Furthermore, a good visualization requires the user to be able to relate items of information to each other, to keep a history of actions to support undo, replay and progressive refinement, and finally to extract the gained information and its corresponding query parameters.

2.1 Disambiguation

The field of visualization is related to many other research fields, such as computer graphics, image processing and computer vision:

Computer Graphics Both visualization and computer graphics deal with image synthesis by transforming abstract data into graphical representations. However, while computer graphics deals with the technical and implementation aspects on how one can generate such pictures, visualization focuses on the data abstraction itself and finding the best visual representation for given data and tasks. Thus, computer graphics targets efficient algorithms while visualization targets effectiveness to use.

- Computer Vision** Visualization maps abstract information to visual representations. In computer vision one performs the inverse task and tries to interpret images, usually acquired from the real world by cameras, in order to extract abstract information from them.
- Image Processing** The field of image processing is closely related to visualization and the two are often performed in combination. In fact, as we will discuss in Section 2.3, image processing is one step in the visualization pipeline. However, the input and output data of image processing algorithms remains in the same domain, i.e. image processing maps data to data or images to images. In contrast, visualization techniques map data to images.
- Perception** The science of perception provides important insight into how the human visual system works. We will provide an overview of the most important aspects of perception in Chapter 4, which can be used to create better visualization.
- Augmented Reality** Augmented reality is the interactive enrichment of real-world scenes with computer-generated elements. Both visualization and AR use rendering techniques from Computer Graphics in order to generate their output. However, while visualization focuses on how to visualize abstract data by itself, the augmented reality community investigates how such visualizations can be integrated with the real world as well as possible.

The visualization community distinguishes *information visualization* from *scientific visualization*. Though it is certainly not a less scientific field of research, the former one deals with the representation of abstract, often multi-modal, n -dimensional data, which usually has no spatial domain by default. In contrast, scientific visualization focuses on the representation of physical data, which has an inherent spatial and/or temporal reference and generally features a rather small dimensionality (usually two to four dimensions). Though mixed forms exist, these two classifications of visualization use mostly distinct methods of visual representation. Since this thesis is on medical visualization, the remainder of this chapter will mostly focus on methods for scientific visualization.

2.2 Goals and History of Visualization

Though visualization as a scientific research field has only emerged in the last decades, visualization techniques have always been used to store and transport information. The extensive list on “Milestones in the history of thematic cartography, statistical graphics, and data visualization” provides a detailed overview of the developments [55].

One of the main goals of visualization is *to provide an overview* and *to abstract information*. Today, we use this in our everyday life for instance in the form of street maps and subway plans. The oldest known map of a town is the drawing found at the Çatalhöyük site in Turkey, dating back to 6200 B.C. Figure 2.1 shows a reconstruction. The Peutinger map, named after a 16th century German collector, is considered the first visualization showing a route map, to





Fig. 2.1. Redrawing of the image found at the Çatalhöyük site depicting a map of a town, probably Çatalhöyük itself, from 6200 B.C. The original drawing is considered as the oldest known map produced by mankind. Image from [1].

some extent comparable to today's subway maps (cf. Figure 2.2). The 34 centimeters by 7 meters large parchment depicts the whole Roman world of that time from Vienna, through Italy to Carthage. Most probably intended to be used as a travel map, it does not show places and distances in scale but instead mentions the connections between places and how one has to travel. The exact date of origin can not be determined but the original cartographer is estimated to have lived in the fourth century [54, 6].


Another important goal of visualization is to *record findings* in order to keep/archive the information and make it available to others. Probably the most impressive example of this are the historic notebooks from Leonardo da Vinci in which he kept the findings of his scientific studies on a wide range of physical phenomena. His observational approach to science, through which he tried to understand the human body by describing and depicting each sample in extensive detail, lead to an incredible amount of scientific visualizations. His



Fig. 2.2. First sheet of the Peutinger map showing the eastern part of Britain, Holland, Belgium, a part of France and the west of Morocco. This map is considered as the first map showing an abstract representation of the travel paths instead of exact distances at scale. Image from [6].



(a) The mechanisms of the hand.



(b) Composite view of the internal organs of a woman's body.

Fig. 2.3. Leonardo da Vinci's recordings on his observations of the human anatomy. Images from [30].

work on the human anatomy “*De figura umana*” alone contains 120 chapters recording his anatomical findings such as the ones shown in Figure 2.3 [30].

The power of abstraction allows visualization to *tell a story* from complex data. The map of Napoleon’s Russian campaign of 1812 (cf. Figure 2.4), drawn in 1869 by Charles Minard, is considered as one of the most powerful visualizations in this regard. Showing the number of soldiers over time and covered distance in relation to geographical features (e.g. rivers) and temperature, it succeeds in elegantly mapping this very heterogeneous and multi-dimensional data into a single image [35, 188].




Fig. 2.4. Charles Minard’s visualization of Napoleon’s Russian campaign is considered as one of the most powerful examples of visual storytelling. The strength of the French army in terms of numbers of men is depicted by the widths of the colored lines. The brown line represents troops moving toward Russia and the black lines represent those on retreat.

Furthermore, carefully designed visualization allows to *discover* new facts and otherwise hidden relationships. One such example is John Snow's dot map of the London cholera epidemic of 1854, where he placed a dot on a London street map for each death. The resulting image (cf. Figure 2.5) shows a clear clustering of dots around the Broad Street water pump. This brought the English physician to the conclusion that the water of this particular well had to be contaminated. The subsequent disabling of the water pump then marked the end of the epidemic spread [35, 189].

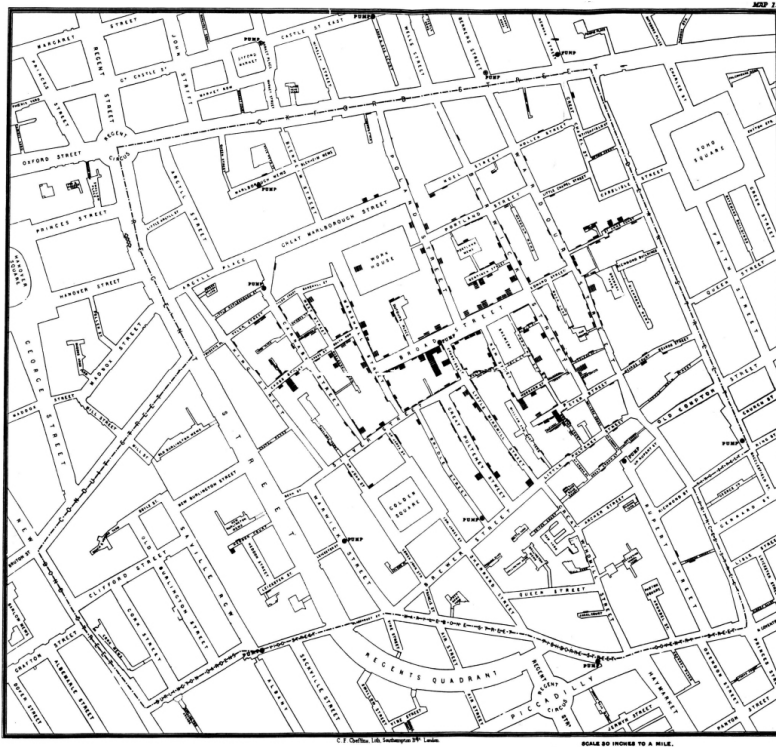


Fig. 2.5. Charles Cheffins' lithography of John Snow's cholera epidemic map showing clusters around the Broad Street water pump. Each bar stands for one death during the London cholera epidemic of 1854.

2.3 The Visualization Pipeline

Though there is no official definition of the term *visualization pipeline*, most authors in the visualization community agree that the workflow of transforming raw input data into graphical visualizations consists of several recurring steps and use this term to refer to them. This series of stages is illustrated in Figure 2.6).

1. As an initial step, the information that should be visualized has to be gathered, which can be done in different ways. Traditionally, raw data is measured using physical processes, such as X-ray attenuation, magnetic resonance or ultrasound in the case of medical imaging. Alternative data sources can be simulation or modeling.
2. In the second step, the acquired raw data has to be processed, e.g. by applying transformations, filtering, interpolation or deriving data otherwise. This step usually falls into the field of signal- or image processing.

3. The third step of the visualization pipeline is the actual visualization part of mapping the processed data to graphical representations. This mapping relates to generating renderable data in terms of graphical primitives and assigning optical properties.
4. In the final step, computer graphics tools are used to render the visualization onto the screen.

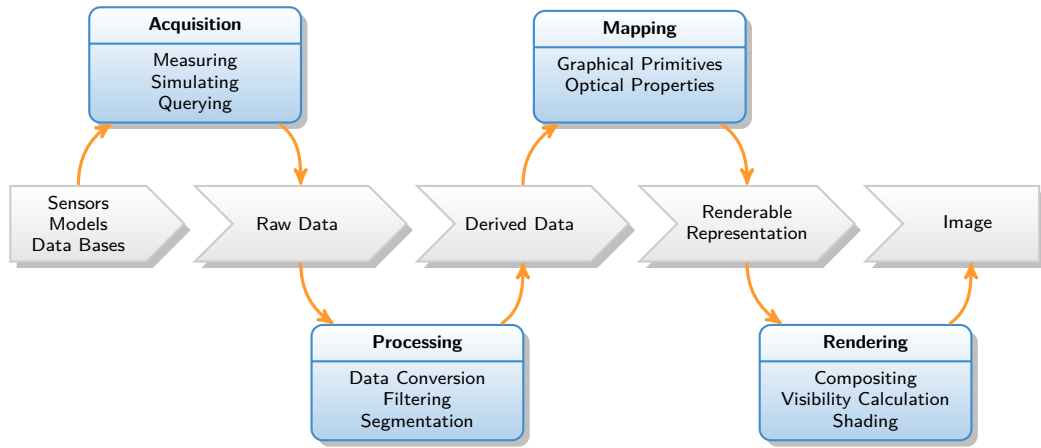


Fig. 2.6. Illustration of the Visualization Pipeline as it is often referred to in the visualization community. It describes the four steps of acquisition, processing, mapping and rendering to get from raw data to the final visualization.

2.4 Data Characteristics

Effective visualization always needs to be tailored to fit the individual characteristics of the input data since the performance of a particular technique to yield an insightful visualization depends on the type of the data. Generally, four main data characteristics can be distinguished [31, 183]:

Source One main characteristic is the origin of the data. As already mentioned in Section 2.3, one can differentiate between three different data sources. The input data can either be an image of the real world (acquired through measuring), originate from theory (acquired through simulation) or a form of design when originating from modeling. Each of the three possible data sources exhibits individual features and characteristics such as data credibility, distribution of uncertainty, presence of noise, etc.

Domain The second data characteristic is the data domain, for which one usually distinguishes three or four classes (cf. Table 2.1). *Nominal* data, sometimes also referred to as categorical data, can be mathematically described as a permutation group [176]. Since nominal data has no assigned ordering, the single comparison operation that can be applied is the test for equality.

In cases where an additional rank-ordering operation can be defined on the data domain, one speaks of *ordinal* data. The ordering measure can determine greater or less but does not support mathematical calculations.

For instance, for two ordinal elements a, b one can define that element a is greater than b but not that a has twice the size of b .

Finally, there is *metric* data, which allows measuring the distance between two elements. This last class is often further split into continuous metric (real values) and discrete metric (distinct values) or into interval scale and ratio scale. The data domain significantly influences the design decisions when creating visualizations. For instance, when representing values through color, the employed color map has to be chosen carefully as discussed in Section 4.4.

Dimensionality

The third data characteristic is the dimensionality of each data element itself. Most data is formed of 1D scalar elements, such as Hounsfield units in CT volumes. However, other element dimensionalities appear as well, such as 3D vector data in flow or strain data, tensor data in Diffusion Tensor Imaging or even higher dimensions in multi-variate data. The data element dimensionality is one of the main constraints defining which visualization techniques are actually applicable.

Structure

Finally, the structure of the data forms an important data characteristic as well. One mainly distinguishes scattered data (e.g. freehand SPECT events over time) from data aligned in a regular grid (e.g. reconstructed CT volume). The data structure determines which algorithms can be applied within the visualization pipeline. For instance, scattered data does not exhibit implicit neighborhood information and therefore needs extra lookup data structures to efficiently support ray casting algorithms as discussed in Section 3.2.

Scale	Comparison Operations	Mathematical Structure	Permissible Statistics
Nomial	Equality	Permutation group	Number of cases, mode, contingency correlation
Ordinal	Greater/less	Isotonic group	Median, percentiles
Interval	Equality of intervals or differences	General linear group	Mean, standard deviation
Ratio	Equality of ratios	Similarity group	Coefficient of variation

Tab. 2.1. List of scales (data domains) and their corresponding properties as described in [176].

2.5 Evaluation of Visualization Methods

Scientific evaluation strategies can be divided into two groups: qualitative and quantitative evaluation. While quantitative methods focus on the systematic empirical investigation and target tangible statistical, mathematical or computational results in form of numbers, qualitative methods focus more on gathering knowledge on how and why something works in order to better understand the results.

It is difficult to quantify how good a novel visualization technique is in comparison to the state-of-the-art due to perception being a very subjective matter and therefore subject to high observer-interrelatedness. Thus, research on visualization is rather hard to evaluate quantitatively. Instead, qualitative evaluation methods such as expert feedback or detailed analysis of case studies play an important role. Nevertheless, quantitative evaluation methods are applicable if they are executed correctly. This section discusses the different approaches to the evaluation of visualization methods.

Researchers can make use of a plethora of different evaluation methods. The following list is inspired from [121] and provides an overview of the possibilities.

- Implementation performance, algorithm complexity analysis
- Quantitative metrics on domain-specific tasks, image quality, etc.
- Qualitative discussion of result pictures
- Case studies, user anecdotes (insights found)
- User studies on usability, intuitiveness, appreciation, etc.
- Design justification from task analysis
- Visual encoding justification from theoretical principles

However, deciding on how and when to apply these methods is not straight-forward as each of them targets specific aspects of visualization design.

Since many visualization techniques are rather specific to their application, task-dependent evaluation of design choices is a frequently used approach [92]. This is particularly the case for medical visualization. Tasks are an essential concept in order to provide a motivation and goal for medical visualization techniques. Such techniques are targeted to support clinicians during their work and in their decision-making and therefore are often specifically tailored to certain activities. However, the term *task* is often used ambiguously within the visualization community, as for instance tasks can be formulated very open-ended (e.g. “find anomalies in the anatomy”) or specific (e.g. “determine the diameter of the aneurysm”). Rind et al. therefore propose a three-dimensional conceptual space of user tasks in visualization. They classify them by abstraction (concrete/abstract), composition (high-level/low-level) and perspective (why/how), which would make task-based evaluation easier to compare [145].

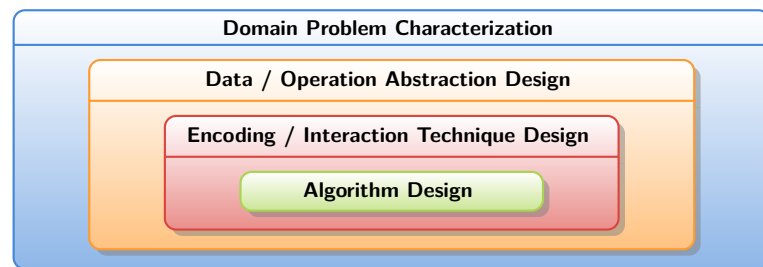


Fig. 2.7. The four-level hierarchical model for the development of visualization as proposed in [120].

There have been different approaches toward generalizing the process of visualization evaluation, providing a wide range of methods and guidelines on how to apply them [32, 133]. Tamara Munzner proposes a four-level hierarchical model (cf. Figure 2.7) of visualization

creation, providing a framework to assess which methodology is appropriate for each of its levels. Based on this model, she identifies possible threats to each individual level and derives applicable methods of validation (cf. Figure 2.8) [120]. Though her work is focused on information visualization, her findings are also applicable to the field of medical visualization.

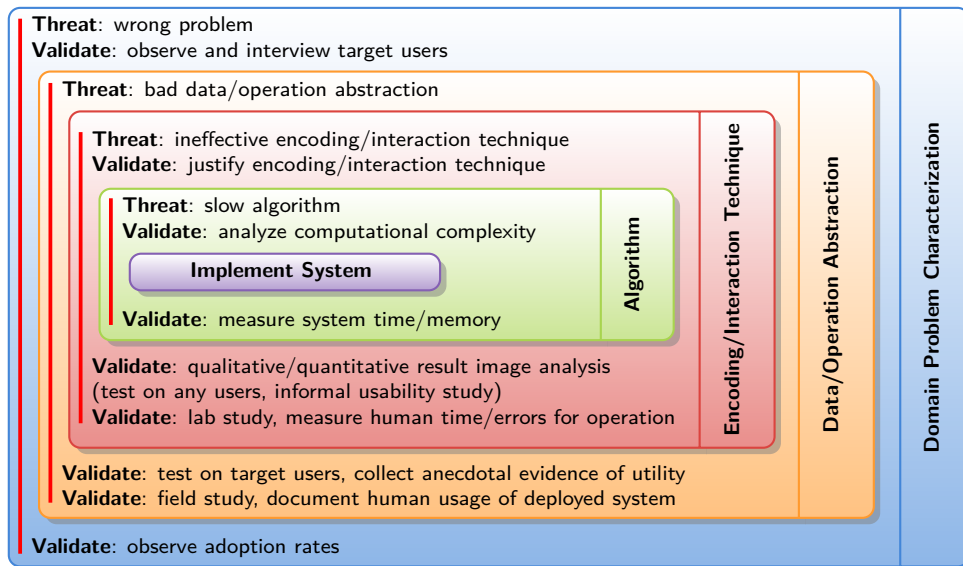


Fig. 2.8. Threats and validation in Munzner's nested model yield a framework for the structured evaluation of visualization techniques as shown in [120].

Another approach to provide an all-encompassing evaluation model is presented by John Stasko who proposes a value-driven evaluation of visualizations. He suggests that the value of a visualization technique is composed by four components and defines it as the sum of [173]:

- The visualization's ability to minimize the total time needed to answer domain specific questions.
- The visualization's ability to spur questions and insight into the data.
- The visualization's ability to provide an overall essence of the data.
- The visualization's ability to generate confidence, knowledge, and trust about the data.

However, the initial question of how to quantify the effectiveness of visualization remains mostly unanswered also for these individual components. Thus, for the time being, the evaluation of novel visualization techniques will remain a highly application- and task-specific process where quantitative measures are defined mostly on a case-by-case basis. A universal, unified framework has yet to be found.

In the recent past, there have been some attempts to introduce new evaluation methods by adopting techniques from other fields. Product reaction cards were originally proposed by Benedek and Miner [7] and provide a way to collect feedback on user experience orthogonally to traditional Likert scales, where the user rates the system on a pre-defined set of questions on a fixed scale. In their approach the users chooses from a set of reaction cards, each containing an adjective, those that best reflect their experience. Since the set contains both positive and

negative cards, the selection allows for an elaborate analysis of the feedback [3, 109]. Blaha et al. propose to build a recommender system out of user studies by conducting them in a structured manner. This allows to apply machine learning techniques to the study results, which then can lead to new insight into the relationships between the experimental data and metrics on the visualized data itself [11]. Perhaps at some day, such machine learning techniques might be capable of attaining the ultimate goal with respect to visualization evaluation: A universal, formal perceptual model that allows to quantify the value of a visualization technique.

Medical Visualization



med·i·cal, *adjective* [usually before noun]

1. *connected with illness and injury and their treatment*
2. *connected with ways of treating illness that do not involve cutting the body*

— Oxford Advanced Learner's Dictionary

Due to the inherent spatial reference of most medical data, medical visualization is a part of scientific visualization. Even though McCormick's report from 1987 [106] is seen as the birth of scientific visualization as a field of research, first examples of medical visualization have been published before. Radiotherapy planning systems were the first to exploit the power of computer-aided visualization. Already in the late 1960's, early systems acquired patient contours with a mechanical digitizer, computed isodose distributions derived from multiple external beams, and displayed both pieces of information on an oscilloscope display [40, 67].

Since then, medical visualization developed considerably and offers benefits in many diagnostic and interventional applications. Since the body of literature on medical visualization is by far too large to be completely covered within this work, the review in the thesis is limited to introducing its general concepts. More detailed discussions are provided only on works that are related to the scope of this thesis such as ultrasound processing and visualization. For a complete and elaborate overview of the entire field of medical visualization, the reader is referred to the book of Preim and Botha [138] and the more recent survey of perceptually motivated 3D visualization of medical image data [139].

3.1 Imaging Modalities

Today's clinicians can make use of a broad range of imaging modalities. Each modality has its own advantages and disadvantages and is therefore commonly used for a certain set of medical application. Generally, two types of imaging modalities may be distinguished. Anatomical imaging shows the underlying anatomy and structure and functional imaging allows for investigating the body's metabolism and other physiological activities.

3.1.1 Radiography

The oldest medical imaging modality is radiography and goes back to the discovery of X-Rays by Wilhelm Conrad Röntgen in the late 19th century, which he named after the unknown variable x [148]. Today, the general setup of a radiography device consists of a X-Ray source, usually formed by a tube emitting X-Ray photons of a specific intensity, and a X-Ray detector. The two are placed at opposing sides of the target anatomy, so that the rays can travel through the tissue. The resulting image represents the originally emitted energy minus the attenuation based on absorption and scattering within the tissue.




Fig. 3.1. Analog X-Ray image showing the forearm and wrist. Today, modern radiography devices are working completely digitally and support much higher resolutions.

Radiography is frequently used in clinical practice due to being cost-effective and real-time capable. Bones and other high-attenuation structures such as surgical tools show good image contrast, while soft tissue differences are often difficult to clearly identify. The major drawback of radiography is the usage of ionizing radiation, which has the potential to impair the genetic material of both the patient and clinical staff. Since the 2D projection shows the attenuation integrated along the entire ray, it is hard to determine positional relationships, particularly in depth. From the visualization perspective, X-Ray images are thus characterized by 1D metric data (intensities), arranged in a regular, 2D rectilinear grid. Traditionally, these intensities are displayed as gray-scale images.

3.1.2 X-Ray Computed Tomography

In order to overcome the limitation of 2D X-Ray projections, clinicians often take multiple X-Rays from different viewing directions. X-Ray Computed Tomography (CT), originally proposed by Godfrey Hounsfield [69], performs this approach in a systematic manner by rotating source and detector axially around the object by at least 180° . Tomographic reconstruction algorithms, such as filtered backprojection, then reconstruct a volumetric representation of the scanned anatomy, where each element describes the local attenuation coefficient [65].

While the drawback of using ionizing radiation remains, CT offers a significantly increased scanning resolution compared to radiography and further allows showing arbitrary cross-sectional images of the anatomy, so that pathologies can be exactly located (cf. Figure 3.2). Modern multi-slice CT scanners even allow the acquisition of volumetric images in real-time in order to study dynamic processes over time. Gating techniques thereby reduce the amount of artifacts when scanning periodically moving anatomy such as the heart. The domain and

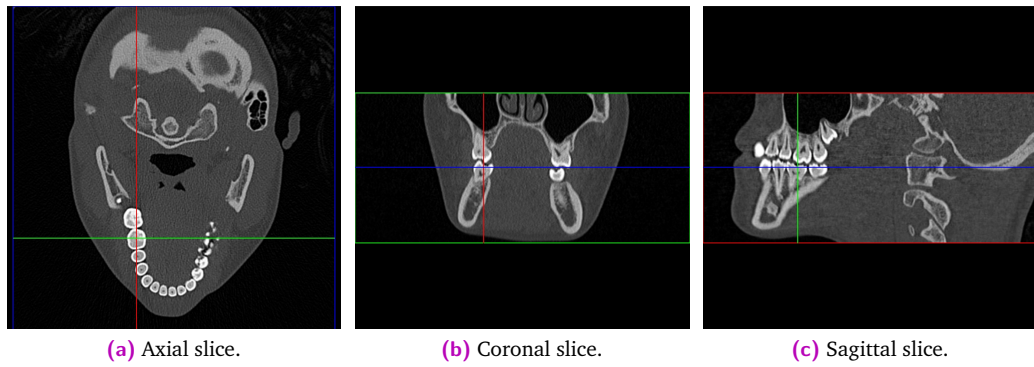


Fig. 3.2. Illustration of the three axis-aligned cross-sectional slices through a CT dental scan, which enable an exact location of potential pathologies. The colored lines within a slice view depict the relative position of the corresponding orthogonal slices.

dimensionality of CT data are still 1D intensities but they are structured in a 3D grid, possibly with an additional time domain. In order to increase the comparability of CT scans, their intensities are normalized into Hounsfield Units (HU), which map water to zero and air to -1000 . Hence, given an attenuation coefficient μ , the corresponding Hounsfield Unit is defined by

$$HU = \frac{\mu - \mu_{\text{Water}}}{\mu_{\text{Water}}} \cdot 1000, \quad (3.1)$$

where μ_{Water} is the attenuation coefficient of water.

3.1.3 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is based on the physical properties of tissue in presence of a magnetic field. A strong, uniform field is applied so that the protons of the hydrogen atoms in the anatomy are aligned while spinning arbitrarily around the axis of the field. An additional pulsating magnetic field causes all protons precessing at the given frequency to resonate and precess in the plane orthogonal to the magnetic field. When the pulse is switched off, the protons return to their equilibrium orientation while emitting electromagnetic signals, which are detected by the scanner. Since the relaxation time is correlated to the concentration of protons, which itself is tissue-specific, it allows for drawing conclusions on tissue type and the reconstruction of a volumetric representation of the anatomy [22, 85].

So far, a plethora of MRI protocols have been developed (cf. Figure 3.3), which allow for both anatomical and functional imaging. For instance, the traditional T_1 and T_2 weighted images show 1D intensities with excellent soft-tissue contrast. The more complex protocol of diffusion-weighted imaging allows to measure the amount of diffusion along certain directions. In Diffusion Tensor Imaging this information is reconstructed into 3D volumes of second-rank tensors modeling the distribution of diffusion present in the anatomy, so that the dimensionality of each data element is six (cf. Section 3.2.2). Compared to CT, MRI has the huge advantage of being radiation-free and showing better soft-tissue contrast while providing a comparable resolution and signal-to-noise ratio. However, acquisition times are much longer and the general costs are considered higher. Furthermore, the inherent strong magnetic field significantly complicates the deployment for applications in the operating room.




Fig. 3.3. Illustration of different MRI protocols applied to brain imaging showing complementary information. Both T_1 weighted images (a) and T_2 weighted images (b) show good contrast between fluid and more solid tissue. In (a), the cerebrospinal fluid of the corpus callosum exhibits low intensities while the same region appears bright in (b). The Fluid-attenuated Inversion Recovery (FLAIR) protocol (c) suppresses fluid structures and highlights certain lesions instead, in this case a brain tumor.

3.1.4 Ultrasound

Ultrasound imaging (US) uses high frequency sound waves, which are emitted from a transducer into the anatomy. In order to improve the transmission, a special ultrasound gel is used as coupling in between. The ultrasound pulse travels as a longitudinal wave through the patient's body and gets partially scattered and reflected at tissue interfaces of different acoustic impedance. When the reflected waves get back to the transducer, their time delay is recorded as scan line data. A post-processing pipeline converts the scan line data into B-mode images showing intensities in a regular 2D grid depicting the echo in the corresponding image plane. A more detailed discussion of the B-mode ultrasound image formation process can be found in Section 5.1.




Fig. 3.4. Example ultrasound images acquired with commercially available systems showing different clinical applications: (a) liver, (b) fetus, (c) Doppler, (d) cardiac 3D ultrasound. Images from [9].

Different imaging modes allow for measuring additional information such as flow velocities (Doppler US) or elasticity (Elastography US) [48, 128]. Special 2D matrix transducers support the recording of 3D B-mode volumes in real-time. Furthermore, when equipping the ultrasound transducer with 6 degrees of freedom (DOF) tracking hardware, one can acquire tracked ultrasound sweeps, a series of ultrasound frames, where each image is accompanied with a pose matrix describing the frame's orientation in space. Ultrasound compounding techniques allow to reconstruct this scattered, unstructured data back into a regular 3D volumetric representation (cf. Chapter 7).

Ultrasound is a widely used imaging modality mainly due to being very cost-effective, real-time capable and radiation-free. Furthermore, ultrasound machines have a very small footprint and are thus mobile and can be specialized for many different clinical applications [123]. However, interpreting ultrasound images correctly requires expertise, since their formation process is very complex and the resulting images often exhibit low contrast, a low signal-to-noise ratio and are prone to artifacts and shadowing. A particular challenge for visualization is the highly non-uniform meaning of B-mode intensities. In contrast to CT Hounsfield units or most MRI intensities, ultrasound intensities do not form an interval scale (cf. Section 2.4) and can therefore not be arithmetically compared within the same image. A more thorough discussion of Ultrasound imaging and visualization can be found in Section 3.3.

3.1.5 Nuclear Imaging

Nuclear imaging allows to show functional processes such as the metabolism within the human body. Therefore, a special radioactive tracer substance is injected into the body, which accumulates in the target anatomy where it decays over time. Positron Emission Tomography (PET) detects the gamma photons emitted by the annihilated positrons using a dense circular detector setup. Single Photon Emission Computed Tomography (SPECT) uses 2D gamma cameras, which are rotated around the target anatomy similar to the X-Ray CT setup [201]. Both modalities reconstruct 3D volumetric images but with significantly lower resolution than CT or MRI. Since they exhibit only functional information, they are often combined with anatomical modalities in order to provide a spatial reference. As of today, there exist fully integrated PET/CT, PET/MRI, and SPECT/CT scanners.

In order to avoid the large footprint of conventional SPECT scanners, recent works introduced freehand SPECT acquisitions using tracked Gamma cameras [104, 200]. From the visualization perspective, this leads to a change of data structure, which is no longer formed by a regular grid but a series of 2D projections arbitrarily oriented in 3D space.

3.2 Basic Visualization Techniques for Medical Data

As discussed in the previous section, most imaging modalities store their acquired data in a structured representation, usually a regular, rectilinear 2D or 3D grid. This simplifies the subsequent visualization techniques since the spatial information (i.e. the position of each sample) is inherently encoded in the data structure. However, depending on the other data characteristics of domain and dimensionality there are different visualization techniques applicable.

3.2.1 Visualization of Scalar Data

Most medical data is formed of metric scalar data stored in a 2D or 3D regular grid. Thus, there is a wide range of visualization techniques available.

Slice Rendering

The traditional technique for displaying scalar medical data is slice rendering. In case of 2D data, such as X-Ray radiographs or 2D B-mode ultrasound images, the transformation of the image on the 2D screen is straight-forward. For volumetric data, however, different methods exist to extract a 2D slice. The classic approach is to display a slice, which is aligned orthogonally to one of the three main axes, since this can be done without the need of interpolation. Actually, early CT scanners did not reconstruct complete 3D volumes but stacks of axial 2D slices and, even today, many clinicians still just cycle through the axial slices and build the 3D model in their head as they have been trained to do so.

However, often the target anatomy is not aligned well with the three main axes so that important information can not be displayed in a single slice. Therefore, multi-planar reformations (MPR) allow to define an arbitrary plane through the 3D volume, which will then be resampled

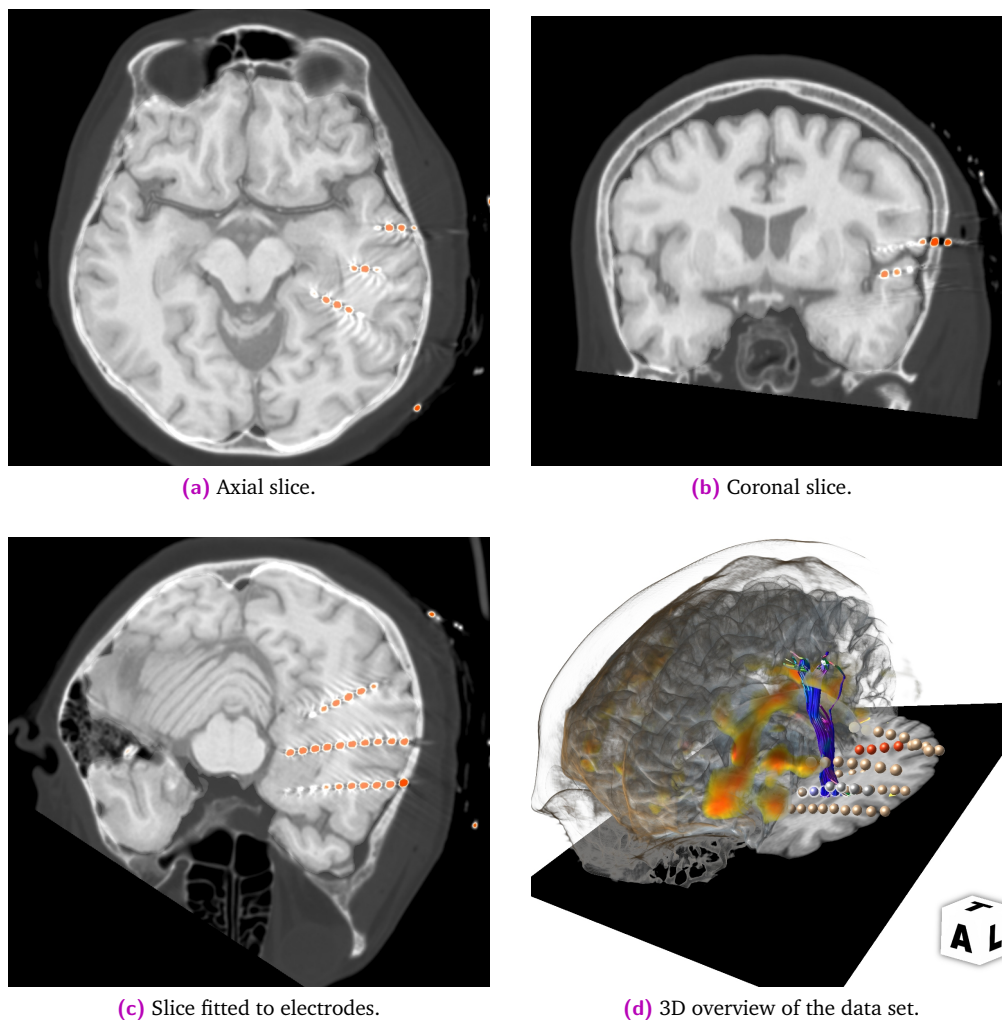




Fig. 3.5. Illustration of multi-planar reformations (MPR) used for neuro surgery planning where the clinician needs to relate implanted depth electrodes (displayed in orange) to their surrounding anatomy. Both axial (a) and coronal (b) slices can not expose all electrodes due to their alignment. A MPR fitted to the electrodes (c) shows almost all electrodes with respect to the anatomy. The volume visualization (d) provides an overview of the data set and electrode configuration.

at the intersection (cf. Figure 3.5). Displaying such 2D views is natively supported by today's GPUs.

Curved planar reformations (CPR) go even further and project an arbitrarily curved connected 2D subspace into a planar 2D image. Kanitsar et al. propose using this technique for vessel visualization since it allows for easier recognition of vessel wall abnormalities (cf. Figure 3.6a) [73, 74]. Williams et al. adapt this technique so that it can also be used with larger tubular structures such as the trachea and the colon [205]. Kretschmer et al. generalize this approach and introduce anatomy-driven reformation. Their straightening and unfolding of bone structures such as the rib cage allows clinicians to easier spot small bone lesions (cf. Figure 3.6b) [87].



(a) Example of a color coded curved planar reformation showing the aorta and surrounding vasculature. Image from [73].



(b) Example of Anatomy Driven Reformation of the rib cage, which allows for inspecting the entire anatomy for lesions in a single 2D image. Image from [87].

Fig. 3.6. Illustration of advanced reformation techniques embedding non-flat surfaces into a 2D image.

Volume Rendering

Even advanced 2D reformations are limited in the amount of information that they can show as they display only a 2D subspace of the data. Thus, they require the clinician to build the complex model of the three-dimensional anatomy in their mind. Volume rendering techniques allow to display such 3D information in a natural way as illustrated in Figure 3.5d. From a technical perspective, one distinguishes *indirect* volume rendering from *direct* volume rendering. The former techniques first extract a geometry representation from the volumetric image data, which can then be rendered using classic rendering techniques from computer graphics. The intermediate geometry can be either defined through the results of segmentation algorithms or as isosurfaces of a certain intensity. The classic algorithm for isosurface extraction is still the Marching Cubes algorithm originally proposed by Lorensen and Cline [100]. However, many more algorithms exist, which are often specialized for individual applications [99].

With the increasing computational power of graphics processing hardware, it has become feasible to implement direct volume rendering techniques in real-time. Instead of extracting an intermediate geometry representation, those algorithms operate directly on the volume data and project it onto the 2D viewport by simulating the physics of light transport through ray casting [47]. Many of today's implementations are based on the image order approach as proposed by Krüger and Westermann, which can be easily implemented in modern GPUs [89].

A set of virtual rays is casted from the camera through the viewport plane into the volume, which are then traversed at a selected sampling interval. For each sample the local image intensity is looked up and mapped to optical properties, which are eventually integrated along each ray using a given compositing scheme as illustrated in Figure 3.7.

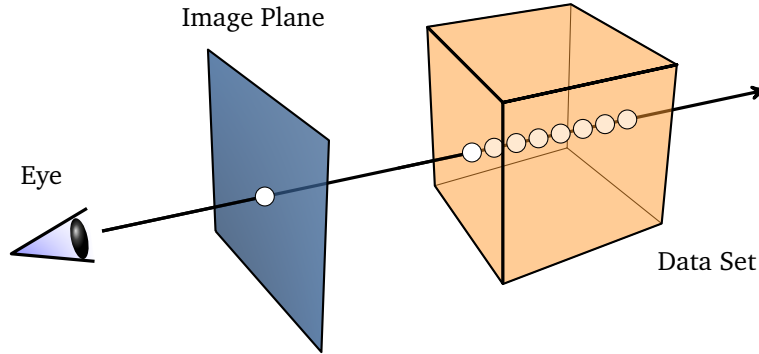


Fig. 3.7. Illustration of image order based ray casting for direct volume rendering. A set of rays is casted from the camera through the image plane toward the image data, where the data is sampled at a given frequency. The image intensity of each sample is then transformed to the optical properties of emission and absorption. Finally, the optical properties are integrated along the ray to yield the final pixel intensity.

For traditional photo-realistic rendering, one often approximates the physics of light transport with an emission-absorption model [105]. This optical model leads to the Volume Rendering Integral where the intensity I at sample s is given by

$$I(s) = I(s_0) \cdot e^{-\tau(s_0,s)} + \int_{s_0}^s q(\tilde{s}) \cdot e^{-\tau(\tilde{s},s)} d\tilde{s}, \quad (3.2)$$

$$\tau(s_1, s_2) = \int_{s_1}^{s_2} \kappa(s) ds,$$

with $q(s)$ being the emission coefficient and $\kappa(s)$ being the absorption coefficient for sample s . This integral can be efficiently computed in an incremental scheme running either front-to-back or back-to-front.

Simpler compositing schemes, such as Maximum Intensity Projection (MIP) or Maximum Intensity Difference Accumulation (MIDA) [25] are often used for visualizing vasculature or other high-contrast anatomy. Digitally Reconstructed Radiograph (DRR) compositing allows to simulate radiography images from CT [110, 114]. In recent years, direct volume rendering techniques have been continuously improved in order to yield a more photo-realistic appearance through sophisticated global illumination models [72, 98], improve depth perception [79, 178], and support even extensively large data sets through out-of-core rendering [41, 42].

Classification

One important step during the visualization of scalar data is the mapping of the scalar values to optical properties, such as color and transparency. This step of defining the look of the data is referred to as classification and usually performed through transfer functions. According to the introduced emission-absorption model (cf. Equation 3.2), a standard 1D transfer function defines an emission coefficient in terms of color and an absorption coefficient in terms of

opacity for each value of the input domain (scalar intensities). Traditional transfer function editors allow to define such transfer functions by defining key/support points in a 2D graph mapping intensity to the horizontal axis and opacity to the vertical axis. To provide a reference to the classified data set, a histogram of the intensity distribution is plotted in the background. Examples of such a transfer function editor are shown in Figure 3.8).

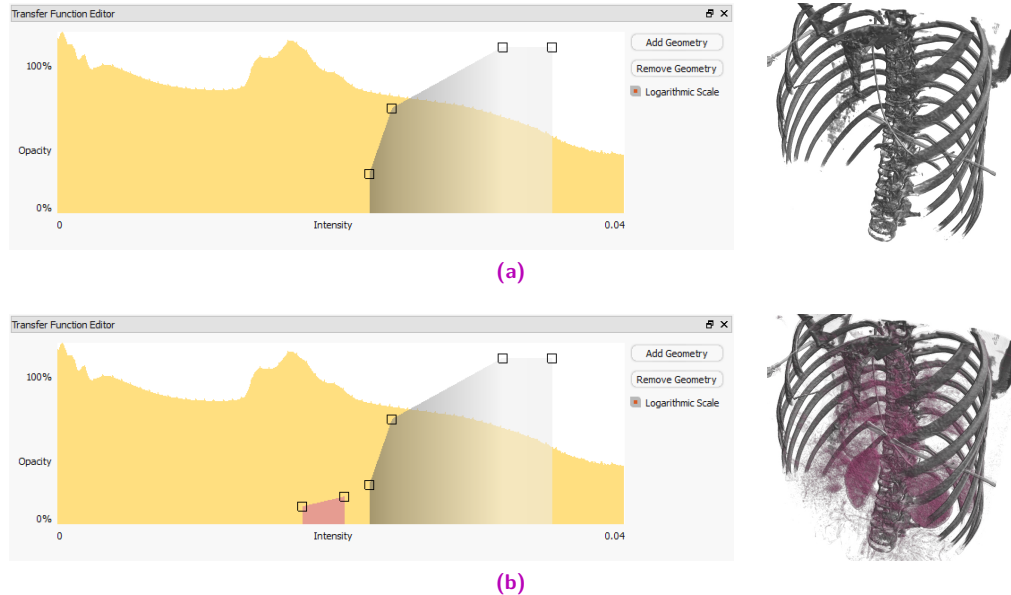


Fig. 3.8. Example of a simple 1D intensity transfer function editor showing a histogram of the intensity distribution in the data as reference. Using geometric primitives to define opacity and color, the user can the optical properties of opacity and color. In (a) a single geometry is used to visualize the bone structure of a thorax CT data set. An additional geometry in (b) yields the purple display of the kidneys.

A standard CT data set usually allows for distinguishing different tissue types solely based on the Hounsfield units. Thus, a 1D transfer function is sufficient to yield high quality visual results. However, in radiocontrast CT, some MRI sequences or ultrasound images, the scalar intensities are ambiguous with respect to the tissue type so that advanced classification methods are required. One solution is to increase the input domain and use multi-dimensional transfer functions as originally introduced by Levoy [97]. Often, gradient magnitude is used as a second input dimension for the transfer function, which allows for a more precise extraction of certain tissue types. Kniss et al. propose a set of direct manipulation widgets to define three-dimensional transfer functions that are based scalar intensity, gradient magnitude, and a second derivative measure [83]. Sereda et al. introduce LH histograms to identify tissue interfaces. Plotting the LH space into a joint histogram reveals clusters for the different tissue types present in the data and allows for an easy definition of the transfer function through 2D geometric primitives [165]. Another approach is to extract local shape information from the data and use this for advanced classification of the data [136].

Though multi-dimensional transfer functions may yield more distinctly segmented visualization, their setup is often highly unintuitive. Especially non-expert users struggle with mapping the complex parameter domain to semantic features for visualization, which makes such techniques unsuitable for clinical practice [142]. Different approaches have been proposed to make the process of transfer function setup more intuitive. Stroke-based techniques let the user delineate directly in the images structures that should be shown and structures




Fig. 3.9. Dynamic galleries as proposed by Jönsson et al. allow novice users to explore large volumetric data sets through intuitive touch interaction. Image from [71].

that should be hidden. These annotations in image space are then mapped back to transfer function space. Thereby, the frequent focus switch between the transfer function editor and the image data can be avoided [150, 190]. Another approach are data-driven techniques to (semi-)automatically generate presets based on clustering or statistical analysis [129, 143]. Introducing a semantic layer into the definition of transfer functions can make their setup more intuitive. Therefore, Rezk-Salama et al. propose to use principal component analysis to map a small set of semantic parameters to the potentially large transfer function parameter space and claim that this can be learned from clinicians [142]. A different approach is the work of Rautek et al. who implement a fuzzy logic evaluation of the semantic descriptions on the GPU and combine it with interaction-dependent rendering [140]. A very recently proposed technique combining the idea of design galleries [103] with modern touch interaction are dynamic galleries presented by Jönsson et al., which allow also novice users to intuitively explore volumetric data sets (cf. Figure 3.9) [71].

3.2.2 Visualization of High-dimensional Medical Data

Not all medical data is formed solely by a scalar field. Special imaging protocols allow to generate high-dimensional images, such as vector- and tensor fields. Furthermore, there is often an additional time domain present.

Vector Data

In vector fields, each element is formed by a (usually three-dimensional) vector. Such data mainly occurs in flow applications such as the visualization of motion of fluids (e.g. blood), geometric boundary conditions, or deformation-/strain fields. Typical visualization strategies include the display of arrow geometry, characteristic lines and dynamic particle simulations.

Direct techniques, such as mapping the vector field to arrow glyphs, are the most straightforward way of displaying such data and are often combined with color coding additional measures such as the gradient magnitude. While they are rather simple to implement and can show direction, magnitude and vorticity of the vector field, they are hard to read since the dense geometry leads to occlusion and visual cluttering. Dynamic particle visualizations simulate the behavior of particles in the flow field over time. Since the user does not need to track single particles but rather observes general flow patterns they can provide a good overview of the properties of the flow field. When combined with an interactive probing technique, this approach allows also for detailed examination of specific regions of interest [131].

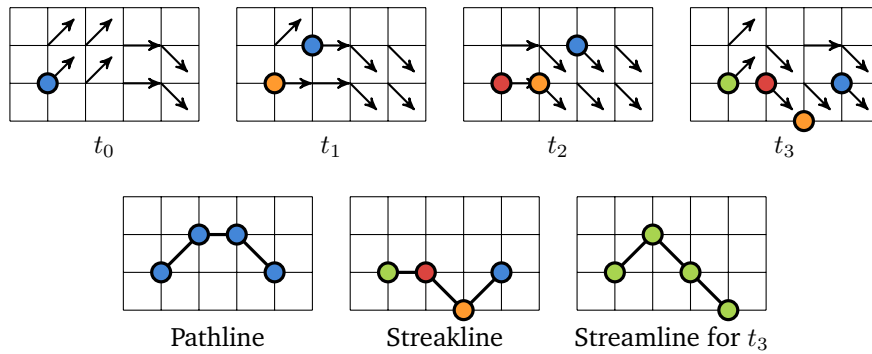


Fig. 3.10. Comparison of the construction of streamlines, pathlines, and streaklines in an unsteady flow field with four time steps. In a steady flow field, all three are identical.

Characteristic lines, such as streamlines, pathlines and streaklines, allow to observe the behavior of the particle flow over time in a single static image. Streamlines only consider a static vector field and propagate a massless particle along the flow. Thus, the resulting line is always tangential to the vector field. Pathlines extend this concept to dynamic vector fields. Streaklines depict the trace of dye that is constantly released into the flow at a fixed position. Thus, they connect all particles that passed through a certain position. Figure 3.10 illustrates the difference between these three types. Once the characteristic lines have been reconstructed from the vector field, they can be rendered with standard geometry-based techniques.

Recent works on medical flow visualization combine these techniques with illustrative focus-and-context rendering techniques of the surrounding anatomy in order to provide a highly integrated result that allows the clinician to assess many different aspects in the data (cf. Figure 3.11) [58, 93, 130].

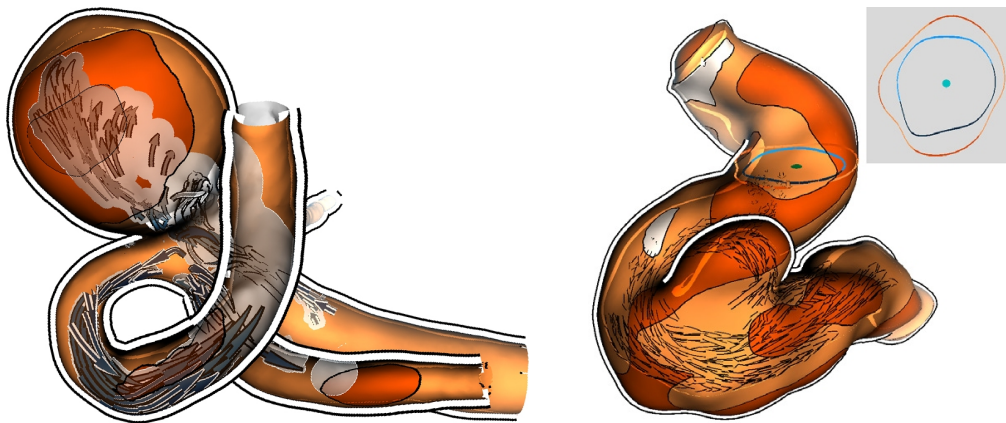



Fig. 3.11. Examples of state-of-the-art flow visualization for vessels combined with a wall thickness estimation as presented by Lawonn et al. Images from [93].

Tensor Data

The mathematical tensor, a generalization of the concepts of scalars and vectors, is a geometric object describing linear relations. In the context of medical visualization, the most common form is the diffusion tensor used in Diffusion Tensor Imaging (DTI)

$$D = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix}. \quad (3.3)$$

It is used to model the amount of water diffusion in tissue, which can be measured by diffusion-weighted MRI [33, 203]. In this diffusion tensor D_{xx} , D_{yy} and D_{zz} represent the amount of diffusion along the three main axes. The other six off-diagonal values represent the correlation between the corresponding two perpendicular axes.



(a) Slice rendering of the fractional anisotropy of a DTI scan of a human brain. The cross-shaped red region in the center represents the corpus callosum, where a high density of fiber bundles is present.

(b) Streamtube rendering of the pyramidal tract reconstructed from DTI. Fiber visualization can depict global connectivity within the anatomy.

Fig. 3.12. Examples of standard diffusion tensor visualization techniques. Both images show a DTI scan of a human brain.

Traditional approaches to visualize this high-dimensional data use dimensionality reduction. Various anisotropy measures (e.g. trace, fractional anisotropy, etc.) can be used to describe certain features of the diffusion tensor in scalar values [5, 203]. The resulting scalar field can then be rendered using standard slice-based and volume rendering techniques as shown in Figure 3.12a. One inherent feature of the diffusion tensor is it being symmetric and positive definite. Thus, its three eigenvalues are real and positive and the corresponding eigenvectors perpendicular to each other. Since the eigenvector corresponding to the largest eigenvalue always points into the main diffusion direction, one can use vector visualization techniques such as streamline renderings in order to show white matter connectivity information of the human brain (cf. Figure 3.12b) [117]. Furthermore, different glyphs have been proposed to represent the diffusion tensor through geometric objects as illustrated in Figure 3.13 [80, 202].

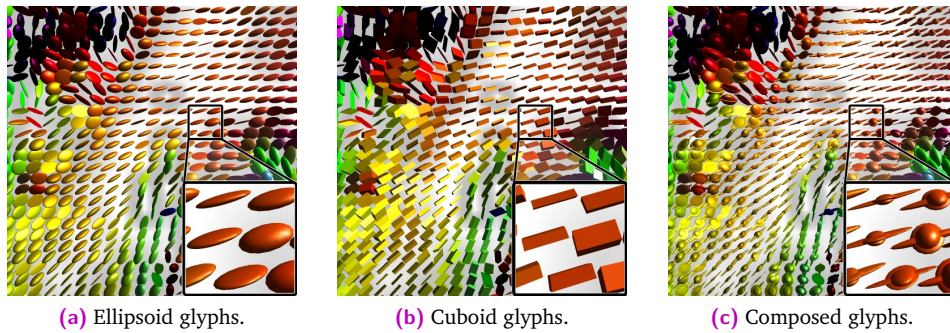


Fig. 3.13. Examples of different glyph visualization techniques for DTI data. Each image shows an axial slice through the corpus callosum of a human brain. Besides the local information, the arrangement of the glyphs can also show some global features of the tensor field.

3.2.3 Multi-modal Visualization

Combining multiple imaging modalities is a powerful tool for clinicians since every modality exhibits certain strengths and weaknesses. Functional imaging modalities such as PET often have a very small resolution. Putting them in context with an anatomical imaging modality such as CT or MR provides the clinician with a spatial reference so that functional hot spots can be associated with the corresponding anatomies. Traditional techniques for multi-modal visualization are juxtaposition, image fusion (superimposition) and glyph techniques. Fuchs and Hauser provide an overview of the state-of-the-art techniques for visualization of multivariate scientific data in [56].

The juxtaposition of multiple images (i.e. showing them side by side) is the most straightforward approach and requires the least amount of preprocessing. However, it also requires the highest amount of mental mapping by the observer, since they have to locate the correspondences between the images themselves. In case the data sets are registered, linking




Fig. 3.14. Illustration of slice-based multi-modal visualization for brain PET-MR imaging. The juxtaposition of the T_1 weighted MR image (a) showing the anatomy and the PET image (b) highlighting functional brain activity requires the clinician to mentally fuse the information. The fused visualization (c) enables the observer to directly relate the PET information to the brain anatomy (i.e. gyri and sulci). All images are registered to each other and show the same cross-sectional slice through the brain.

techniques can facilitate this process – for instance by always showing a slice of the same position and orientation throughout all modalities as illustrated in Figure 3.14.

Image fusion describes the embedding of multiple images into the same reference frame. Unless the images are acquired from the same device in the same configuration, an initial registration process is required to align the images to each other. In the subsequent visualization process the images are either first rendered individually and then blended together into the final image (suitable approach to 2D visualizations or geometry-based rendering techniques) or directly rendered together in a single pass (required for ray-casting volume rendering). Special care has to be taken to adjust the transfer function of each image to suit the entire scene. Failure to do so may result in an unreadable image due to occlusion or visual clutter. Different approaches have been proposed to make this process more user-friendly. Haidacher et al. suggest to use a generic information-based technique to identify complementary information in two data sets and thereby reduce the parameter space [61]. Focus-and-context techniques such as cutaways can further improve the understanding of the complex data and allow deep structures to be shown. Beyer et al. present a fully integrated rendering system for neuro surgical planning supporting such cutaways and an arbitrary number of fused volumes and geometries [8]. A recent approach also using illustrative visualization techniques is the work presented by Lawonn et al. focusing on fusing PET and CT scans [94].

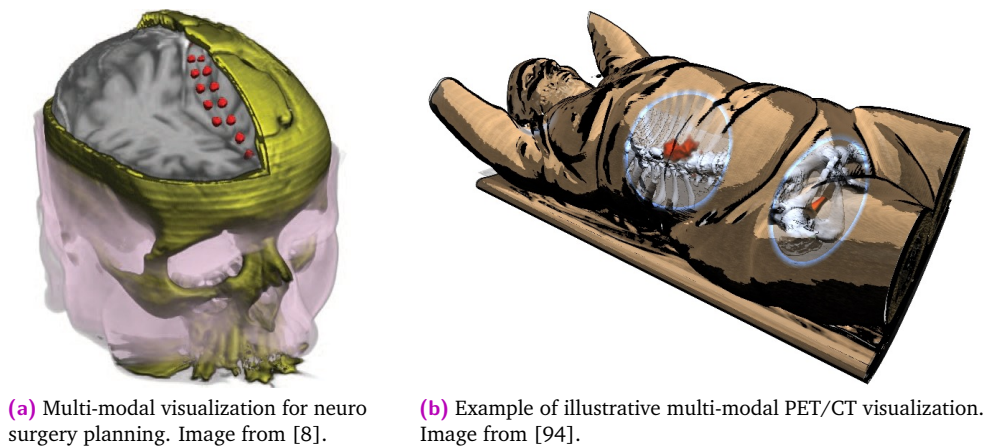


Fig. 3.15. Examples for state-of-the-art multi-modal volume rendering using illustrative techniques such as cutaways and lenses.

Glyph techniques are a very flexible way of visualizing multivariate data, where there are multiple output variables for each data element. Their strength is the support of quantitative analysis of specific features in the data, in particular when used in hybrid visualizations that provide a spatial context. Glyphs are simple polygonal objects with a fixed set of optical properties such as color, size, etc. Besides classic glyphs such as arrows, ellipsoids or cuboids, which have already been mentioned in the previous section on high-dimensional medical data visualization, more advanced glyphs exist that allow to depict more measures through additional optical properties. The supersphere, seamlessly interpolating between cube and sphere, allows to map one additional parameter to roundness. Similarly, the supertorus allows to map two additional parameters to roundness and thickness as illustrated in Figure 3.16.

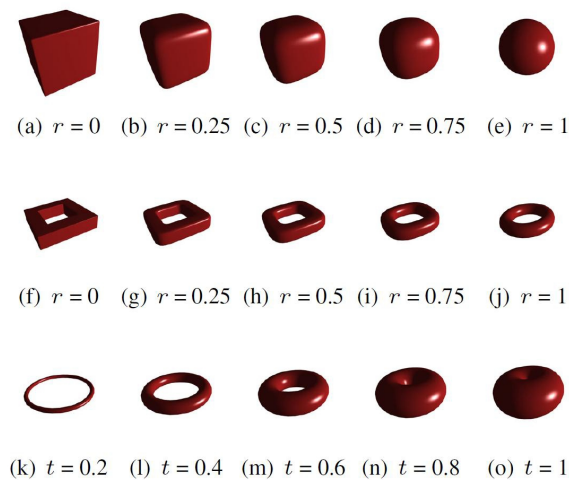


Fig. 3.16. Illustration of the additional parameter space of the supersphere and supertorus glyphs. Both allow to map one parameter to the roundness. The supertorus has the additional property of thickness. Image from [152].

While glyphs offer an effective means of visualizing highly multivariate data, their application and usage, particularly for 3D visualizations, is not straight-forward. In [151], Ropinski and Preim define a taxonomy for glyphs based on the kind of visual information processing (e.g. pre-attentive vs. attentive processing; cf. Section 4.1) and derive usage guidelines for glyph-based medical visualization. They argue that glyph placement is a crucial aspect, since glyphs are usually placed onto the surface of the anatomy of interest in order to display additional information. In such a setup, visual cluttering should be avoided and it must be ensured, that neither do the glyphs occlude each other nor do the glyphs occlude large parts of the anatomy. Therefore, different glyph placement strategies have been proposed [196]. In general, a good




Fig. 3.17. Comparison of different glyph placement strategies. From left to right: regular grid, random distribution, random distribution with relaxation. Image from [111].

glyph-based visualization ensures that the shapes are unambiguously perceivable and that unwanted glyph aggregations in image space due to bad placement is avoided. The parameter mapping should visually emphasize important variables and thereby guide the user's focus of attention. Furthermore, it should incorporate the semantics of the underlying data in order to facilitate the mental reconstruction of the information by the observer. Good demonstrations of the power of such visualizations can be found in the works of Meyer-Spradow et al. [111] and Oeltze et al. [126].

3.3 Ultrasound Visualization

Since 2D images can be directly embedded into the rendering viewport, the display of 2D B-mode ultrasound is very straight-forward as the intensities are directly translated into brightness. Additional information, such as Doppler or elastography ultrasound, is usually displayed through color overlays. However, the special nature of ultrasound images compared to traditional tomographic modalities such as CT or MRI poses special challenges when moving toward 3D visualization. Here, traditional visualization techniques often yield unsatisfactory results and therefore require additional processing. Birkeland et al. provide an in-depth overview of the ultrasound visualization pipeline in [9].

3.3.1 Acquisition and Processing

Consistent with the visualization pipeline (cf. Section 2.3), ultrasound visualization begins with the acquisition of the ultrasound image, which is based on the physical principles of sound pulses interacting with tissue. A detailed overview of the B-mode ultrasound image formation process can be found in Section 5.1.

Usually, further processing has to be applied before the data can be displayed to the user. Depending on the application such processing step can be formed by the execution of algorithms for filtering, segmentation or deriving additional information from the original data. Chapter 5 will present a novel method for generating real-time uncertainty information from 2D ultrasound sequences.

When targeting the visualization of volumetric ultrasound, one approach to generate such data is to equip a traditional 2D probe with 6 degrees of freedom tracking hardware and acquire a sequence of images covering the target anatomy. This data is then compounded into a 3D volume. Section 7.1 will provide an in-depth review of the available techniques for ultrasound compounding. Furthermore, we will develop a novel, orientation-driven compounding scheme in Chapter 7, which allows for an interactive reconstruction of the volume during the acquisition.

Segmentation

It is an inherent property of sonography that its images do not only show the anatomy of interest but the entire tissue in the field of view of the transducer. However, in particular for volumetric ultrasound, unimportant anatomy such as the skin may actually occlude the target anatomy and thereby hide important details. Semi-automatic clipping algorithms enable the user to manually define regions within the image that should be removed [154]. Figure 3.18 illustrates the MagiCut technique [17], which is implemented in the GE Volusion ultrasound machines.

Due to the special nature of B-mode ultrasound intensities, generic segmentation techniques solely based on the local data values are not applicable. Instead, such algorithms are very application-specific and require a certain type of data or make other prior assumptions.




Fig. 3.18. Illustration of the MagiCut technique [17] that is implemented in the GE Volusion ultrasound machines. In a 2D cross-sectional slice, the user manually delineates the skin region, which is subsequently removed from the 3D rendering in order to provide a clearer view onto the target anatomy. Image from [9].

Noble and Boukerroui provide an extensive survey on the recent developments in ultrasound segmentation and additionally highlight 10 particularly influential works [124].

Filtering

While the characteristic speckle pattern provides important cues to clinicians when looking at 2D ultrasound images, the inhomogeneous intensity distribution leads to a very noisy appearance and unwanted occlusion in 3D visualizations. Thus, advanced structure-preserving filtering is essential in order to get satisfactory results. Traditional implementations employ median filtering or other non-linear methods based on wavelets [60], total variation or anisotropic diffusion [88]. A comparison of their performance on ultrasound despeckling can be found in the survey paper of Michailovich and Tannenbaum [113]. In a recent work, Soltészová et al. propose a streamline integration method along the direction of lowest variance [172]. Since it can be assumed that the anatomy structures are oriented in this direction, averaging the samples along the streamlines yields superior results. Exemplary results are shown in Figure 3.19.




Fig. 3.19. Exemplary results of the lowest-variance-based filtering method of Soltészová et al. The filtered images (I.B through IV.B) exhibit a significantly lower amount of speckle compared to the original images (I.A through IV.A) while all important anatomical details are preserved. Image from [172].

Since it is an essential feature that ultrasound imaging is real-time capable, the acceptable complexity of the filtering method is limited by the target frame rate. Though the computational power of today's hardware is constantly increasing and many filtering methods can be parallelized on GPUs, not all methods allow for an interactive processing of the incoming images. To reduce the computational burden, Soltészová et al. propose a visibility-driven approach to the processing of streamed ultrasound volumes [171]. After estimating the set of potentially visible voxels with respect to the current view point, they reduce the problem size by restricting the filtering process to those voxels and can thereby increase the performance.

3.3.2 Mapping and Rendering

Assigning optical properties such as color and opacity to ultrasound intensities for volume rendering poses various challenges as well. Traditional classification methods are difficult to apply due to the low dynamic range and the significant amount of noise. Most importantly, B-mode ultrasound exhibits different intensities for the same tissue. This renders standard intensity-based 1D transfer function approaches mostly useless since the resulting images show low contrast and a high degree of occlusion. As an alternative, Steen and Olstad evaluate different volume projection techniques that do not require an additional classification step, such as maximum/average intensity projection and gradient maximum projection [174].

Classification

One early method, specifically tackling these challenges of visualization, is the work of Sakas et al. [154]. Fattal and Lischinski propose a variational approach based on the local value and gradient that performs filtering and opacity classification in a single unified process. This allows them to extract smooth surfaces from 3D fetal ultrasound volumes [49]. Also for the application with 3D fetal ultrasound, Wang et al. describe a surface rendering technique, which combines a variant of anisotropic diffusion filtering with a novel light absorbing function to remove artifacts. Finally, they map textures acquired from photos of baby faces to the extracted surfaces [193].



(a) Variational classification approach of Fattal and Lischinski. Image from [49].



(b) Adaptive global opacity transfer function setup by Hönigmann et al. Image from [68].

Fig. 3.20. Illustration of rendering results when applying different state-of-the-art classification results to 3D fetal ultrasound.

Another challenge of ultrasound classification is the high variability of intensities between different images, which requires opacity transfer functions to be carefully optimized for each data set. Often, opacity transfer functions for 3D ultrasound distinguish three classes: invisible, partially transparent and opaque. Such a piecewise linear transfer function can be parameterized by a threshold defining which intensity range is invisible and the a slope parameter controlling the increase of opacity for intensities above the threshold. This reduction to two scalar parameters allows for a basic but straight-forward setup of the rendering. Hönigmann et al. propose a method to automatically adjust the global opacity transfer function based on the intensity distribution along selected scanlines [68].

Many implementations for 3D ultrasound visualization additionally apply a color coding of depth information similar to the chroma-depth technique proposed by Ropinski et al. [153]. Coloring close regions in warm tones such as orange and farther regions in colder tones such as blue significantly improves the spatial perception by the observer (cf. Section 4.4).

Chapter 8 will introduce a novel approach to classification based on predicate functions specifically targeted for volumetric ultrasound. Our technique integrates the evaluation of both local and global information (e.g. application-specific segmentation algorithms) into a unified framework for focus-and-context visualization.

Systems

By combining state-of-the-art algorithms throughout all steps of the volumetric ultrasound visualization pipeline, Andrej Varchola developed a fully integrated system allowing for high-quality renderings of prenatal ultrasound in real-time [191]. Advanced structure-preserving



(a) Real image acquired through a fetoscope.



(b) Fetoscopic rendering of a ultrasound volume.

Fig. 3.21. Tight integration of state-of-the-art algorithms allows for high-quality real-time fetal ultrasound imaging mimicking the appearance of fetoscopic imaging. Images from [191].

filtering, dynamic global illumination and realistic rendering of the human skin yield results of great quality mimicking the appearance of fetoscopy images (cf. Figure 3.21). While the visual quality of the fetoscopic rendering is quite impressive, one has to keep in mind that prenatal ultrasound has a much better acoustic configuration compared to other anatomies,

since the fetus is surrounded by amniotic fluid exhibiting excellent transmission properties. Thus, the results are not necessarily transferable to other ultrasound applications.

An orthogonal approach, mainly targeted for applications aside from obstetric ultrasound, is multi-modal visualization, where interventional ultrasound is combined with previously acquired CT or MRI data for spatial reference. Burns et al. present a system for relevance-based visualization of multi-modal data supporting dynamic cutaways. Their visualization interactively shows the 2D ultrasound image in the context of a volume rendering of a CT data set showing the surrounding anatomy. As illustrated in Figure 3.22, this provides the user with an excellent spatial comprehension of the data and relative position of the ultrasound image [27].




Fig. 3.22. Relevance-based multi-modal ultrasound visualization as presented by Burns et al. A volume rendering of a pre-operatively acquired CT data set provides spatial and anatomical context for the interventional 2D ultrasound image. Images from [27].

Perception

” perception, *noun*

1. *the way you notice things, especially with the senses*
2. *the ability to understand the true nature of something*
3. *an idea, a belief or an image you have as a result of how you see or understand something*

— Oxford Advanced Learner’s Dictionary

The human visual system of perception and cognition is highly complex. The eye serves as broadband channel into the mind since for the average human the majority of the environment is perceived visually and roughly half of the human brain deals with processing this visual input. Developing an understanding of the underlying processes helps in designing better visualization since one can explain why a certain output is suitable for clinical practice and another is not.

The human visual system perceives only a limited bandwidth of the entire light spectrum. The rods and cones in the human retina respond to the wavelengths between 400 and 700 nm and generate neuronal signals. They pass through the optic nerve and optic tract to the visual cortex where they are initially processed in the primary visual receiving area (V1). From there they are forwarded along the dorsal pathway for object location and motion, as well as along the ventral pathway for object recognition and form representation [24].

However, what people actually see is not simply a translation of retinal stimuli but actually the result of the complex system of visual processing, which motivates the perception theory. In this chapter, we will discuss different aspects of human perception and how this insight can be used to create better visualizations.

4.1 Visual Attention

According to the feature integration theory, the cognitive process of visual processing can be divided into two stages [186]. It starts with the parallel processing of single features such as form, orientation, color and motion (*preattentive stage*), and then continues sequentially with the combination of the features into 2D patterns, contours and object identification (*focused attention stage*). Preattentive processing applies to various features underlying a pop-out mechanism, which occurs prior to the conscious attention. This salience is mostly independent

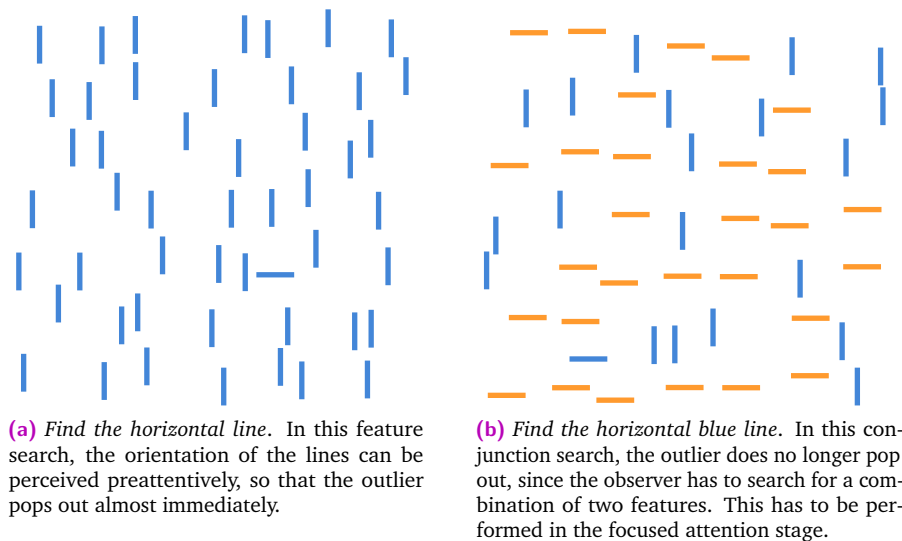


Fig. 4.1. Demonstration of visual processing in the preattentive stage in comparison to visual processing the focused attention stage.

of the number of distractors in feature search tasks, such as the one depicted in Figure 4.1a. Conjunction searches, however, where a combination of multiple stimuli is targeted (cf. Figure 4.1b), can no longer be performed preattentively since the features have to be combined in the focused attention stage. This visual search has to be performed by sequentially looking at each target and thus requires significantly more time (cf. Figure 4.2) [197].

This trait of human perception can be exploited during visualization design. It does not only explain why often less is more but can also be used to guide the observer's attention. In order to highlight a specific detail in the data, one should design an implicit feature search and map the corresponding data domain to a single feature that is subject to the preattentive pop-out effect.

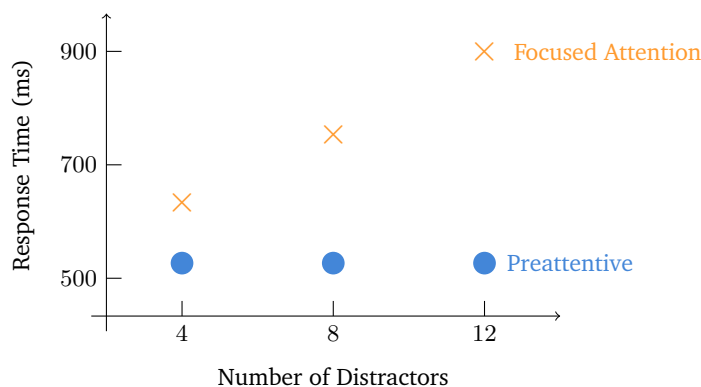


Fig. 4.2. Typical response times for outlier search tasks. The preattentive response times are mostly independent of the number of distractors while the response time for focused attention-based searches increases with the number of elements [197].

4.2 Perceptual Organization

One important aspect of the previously introduced focused attention stage is perceptual organization, in which the visual stimuli and features become perceptually grouped into coherent units (objects). This process is based on two components: Single features are assembled into objects by *grouping*, while *segregation* allows us to discriminate multiple objects from each other. These effects are illustrated in Figure 4.3. Being aware of the perceptual organization mechanism allows us to create better visualizations, for instance, avoid visual cluttering and make it easier for the observer to perceive certain objects as such.



Grouping: All windows belong to one object (building).

Segregation: The two buildings are separated from each other; the left one is in front of the right one.

Fig. 4.3. Illustration of grouping and segregation during the perceptual organization stage. These two components allow us to identify the buildings as single objects and separate them from each other.

Gestalt theory

In the early 1900s, the Gestalt psychologists around Kurt Koffka tried to find an answer to the question of how the human brain is capable of acquiring meaningful perceptions in an apparently chaotic world. Their resulting *Gestalt theory* can be seen as basis for the aforementioned process of grouping and segregation. According to this holistic approach, perceptual organization is based on the four key aspects of *emergence* (objects are perceived as a whole and not as sum of its parts), *reification* (objects are perceived even if they are incomplete), *multistability* (perceived alternative is chosen randomly in case of ambiguous stimuli) and *invariance* (objects are perceived independent of transformation and deformation) [24, 95].

Furthermore the Gestalt psychologists defined a set of organizing principles describing how elements are grouped together:

- Continuity** Points that can be connected in a smooth, continuous fashion are perceived as a group and lines are seen to follow the smoothest path.
- Closure** Objects are perceived as a whole even if they are not complete.
- Similarity** Objects of similar appearance are grouped together.




Fig. 4.4. Illustration of three of the four key gestalt principles. *Left:* One perceives the triangle even though it is only partly drawn (reification). *Middle:* One perceives either two faces or a vase (multistability). *Right:* One perceives all elements as one instance of the same object (invariance). Images from [95].

Proximity	Objects are grouped together based on their proximity.
Symmetry	Objects are perceived as being symmetric and formed around a center point.
Common Fate	Objects that have the same trend of motion are grouped together.
Good Gestalt	Elements are grouped into objects if they form a simple and regular shape.

These principles of grouping are further extended by the principles of figure and ground, which are responsible for perceptual segregation.

Prior Knowledge

Further mechanisms for perceptual organization are given by prior knowledge about the regularities in the environment. For instance, there are certain regularities in the physical properties of the environment, which we have learned to take for granted over time and thus perceive instinctively. One very strong example is the light-from-above assumption as described in [81]. In our daily environment, most light comes from above. Thus, in Figure 4.5, we perceive the left six circles as mounds and the right six circles as indentations solely because of their shading. This is the case even though this is a completely artificial scene and the circles on the right-hand side are simply vertically mirrored versions of the circles on the left-hand side. Further examples of physical regularities are that common objects in our environment are formed by mostly homogeneous colors, which are different from the colors of surrounding objects. Also, vertical and horizontal orientations have a stronger presence than oblique (angled) orientations [39]. Finally, prior knowledge on regularities of semantic nature allow us to identify objects through assigning them a certain meaning. The mechanism allows us to form a semantic context for the observed scene and thereby influence the processing of ambiguous stimuli as objects [127].

4.3 Perceiving Depth and Size

Understanding depth perception is highly important when designing visualization. This holds particularly true for medical applications where one wants to show specific details of the three-dimensional anatomy as well as spatial relationships between multiple features to the clinician. However, even though virtual reality systems such as head-mounted displays and




Fig. 4.5. Illustration of the light-from-above assumption. We perceive the left six circles as mounds and the right six circles as indentations solely because of their shading. This is the case even though this is a completely artificial scene and the circles on the right-hand side are simply vertically mirrored versions of the circles on the left-hand side.

stereoscopic monitors are making their way into clinical practice, the vast majority of display devices are normal 2D screens. Thus, a three-dimensional scene needs to be visualized in a two-dimensional view, which requires the correct usage of depth cues in order to allow the observer to perceive the information correctly.

Generally, perceptual depth cues can be categorized into three groups:

1. *Oculomotor* cues are based on the ability to sense the position of our eyes and the tension on our eye muscles.
2. *Monocular* cues that work with one eye.
3. *Binocular* cues that depend on two eyes.

Oculomotor Cues

The oculomotor depth cues are based on two mechanisms of our eyes: When we are looking at nearby objects, our eyes move inwards synchronously (convergence). Furthermore, when focusing objects at different distances, the shape of our lenses change (accommodation). We are able to feel these two mechanisms through muscle tension and thereby loosely estimate the depth at which we are looking up to a distance of about arm's length [115, 187]. However, since in the context of standard 2D visualization the observer is constantly focusing the display device being fixed at one position, these cues cannot be exploited to improve depth perception in 2D images.

Monocular Cues

Monocular cues work with only one eye and can be grouped into pictorial cues, which can be depicted in a single picture, and motion-produced cues. Therefore, they are the most important ones for visualization purposes. Some of the cues in the following list from [24] are also depicted in Figure 4.6.

Occlusion

Partial occlusion of objects provides a relative depth cue. The occluded object is farther away than the occluding object, but one can not perceive by how much.



Fig. 4.6. Various monocular depth cues allow us to perceive depth in this static 2D image. The perspective convergence of the road and the power line on the right is probably the most powerful one. Furthermore, occlusion shows that the road sign is in front and the shadows indicate that the motorcycles are on the ground and not hovering above it. Relative height and size allow further depth estimation of the vehicles and clouds. Note that the relative height depth cue works inversely for objects above and below the horizon.

- Relative Height** Objects having their bases closer to the horizon are perceived as being farther away. Thus, objects on the ground appear the more distant the higher they are, while in the sky this cue works inversely.
- Relative Size** Objects will become smaller the more distant they are. Thus, if two instances of the same object are present, the smaller one will be perceived as farther away.
- Perspective Convergence** Perspective convergence is the fact that in projective vision parallel lines are projected to intersect in a vanishing point. This provides us with a very strong depth cue, for instance in pictures where train tracks or roads converge in the distance.
- Atmospheric Perspective** Since the air is full of particles (dust, water droplets, airborne pollution), distant objects are perceived as less sharp and often have a blueish tint. While in the real world this cue usually only appears at rather long distances, it can be also used synthetically in visualization for arbitrary depth relations by introducing a layer of fog to depict farther away objects.
- Texture Gradient** Regular textures such as equally spaced features become more closely packed as distance increases.
- Shadows** Shadows occur when the incoming light is blocked by an object and the surroundings appear partially darker. This is a particularly powerful cue to perceive relative positions of objects and can for instance indicate whether an object is floating or touching the ground.

In addition to these stationary monocular depth cues, there are monocular cues that are induced through motion. They are perceived as soon as the observer starts moving, thus, in visualization as soon as it becomes interactive.

Motion Parallax Motion parallax is a phenomenon that is widely used to create an impression of movement and depth in 2D cartoons and video games. It is characterized by the differences in velocity of objects passing by at different depths where closer objects pass by more quickly than farther objects. Many animals use this depth cue as main source for depth information [24].

Deletion and Accretion The combination of the occlusion depth cue with motion parallax leads to deletion and accretion. When we move, some objects become covered while other become uncovered depending on their relative depth and position.

Binocular Cues

Although the above monocular cues are sufficient to perceive complete depth information in a scene, binocular or stereoscopic vision allows for a both quantitatively and qualitatively better perception of depth. When perceiving a scene with two eyes, our visual system is capable of identifying the minimal differences in the two images of the left and right eye that exist due to them being at slightly different positions. Their slight horizontal separation leads to an absolute and relative disparity, where the images of objects at different depths fall on non-corresponding points in the two retinas. Though this disparity is rather a geometrical effect and the perceptual efforts are limited on detecting this, the angle of disparity allows to estimate the relative depths of multiple objects to each other. Similar as with preattentive processing in visual attention (cf. Section 4.1) objects just pop out in depth if their binocular disparity is large enough [24].

An impressive and vivid description of this effect can be found in the story of Susan Barry [4]. Originally being cross-eyed and suffering strabismus, her visual system constantly suppressed the vision in one of her eyes to avoid a double vision. She was seeing her environment with only one eye at a time and was therefore solely relying on monocular depth cues. She reports that she quickly learned to compensate for this and was able to perceive depth very well and could participate in everyday activity such as sports and driving a car without any problems. At the age of 49, she consulted a developmental optometrist in order to learn to align her eyes and one day the steering wheel of her car suddenly popped out from the dashboard. After having lived in a flat world for fifty years, she describes this leap into three-dimensionality as a revelation.

4.4 Color Perception

The perception of color is a psychophysical process. While the physical properties of light, such as wavelengths, spectral distribution and intensity are given by the environment, the psychological and perceptual process in the human visual system transforms these sensations to what we call color. Color can be generalized into the three perceptual dimensions of hue, saturation and intensity, having a non-linear relation to the above-mentioned physical properties [24]. While the hue (i.e. the color of the rainbow such as blue, green, red, yellow) corresponds to the predominant wavelengths, its saturation (i.e. chromacity/paleness)

corresponds to the spectral distribution of the incoming light. Finally, the color intensity (i.e. lightness/darkness) is related to the amount of light entering the retina.

When designing a visualization, one should choose the applied colors carefully, as they have individual perceptual effects. For instance, short wavelengths are usually perceived as being farther away than long wavelengths because the lens of the eye is refracting them at different angles [175]. This phenomenon was exploited by the pseudo-chromadepth technique developed by Ropinski et al. and is illustrated in Figure 4.7. It provides an improved depth perception in monocular angiography images by using a gradient color map running from red (long wavelength, close) to blue (short wavelength, far) [153].




Fig. 4.7. Illustration of chromadepth and pseudo-chromadepth coding of depth in angiography images as shown in [153]. Compared to the (a) and (b), image (c) provides the observer with a much better perception of relative depth just through the use of color mapping.

One very controversial topic in the visualization community is the usage of the rainbow color map for depicting numeric data. According to the perception theory it has two major flaws: The main hues of the rainbow exhibit no intuitive ordering and the non-linear correlation between hue and perceived brightness leads to misleading banding effects [14]. Nevertheless, the rainbow color map is still used quite often in scientific visualizations. This may be due to it being the default setting in many visualization tools, a lack of alternatives or scientists simply being used to reading such visualizations.

However, general guidelines exist for choosing good color tables with respect to the given data type (cf. Section 2.4) [14, 185]. For nominal data, one should choose a selection of distinct colors. High frequency ordinal data can be represented well with a luminance contrast color map such as the heated body scale. When applying color mapping to surfaces, one should use an isoluminant scheme to preserve the perceived shape. Finally, diverging color maps [116] provide an excellent means for showing interval and ratio data on two-ended scales.




Fig. 4.8. Examples of different color maps and when they are used best.

Part II

Methodology

Real-time Ultrasound Uncertainty Estimation

” re·al time, *noun* [uncountable] (computing)

1. *the fact that there is only a very short time between a computer system receiving information and dealing with it*

— Oxford Advanced Learner’s Dictionary

As with every kind of information, medical images comprise uncertainty, which should always be taken into account in order to interpret them correctly. The fact that uncertainty has many different facets and originates from various sources makes it a difficult aspect to understand. Thomson et al. categorize uncertainty into the nine different types of accuracy/error, precision, completeness, consistency, lineage, currency, credibility, subjectivity and inter-relatedness [181]. Some of these facets are directly based on the physics of the image acquisition and can thus be estimated through mathematical models. Others are specific to the individual user, patient anatomy and further features, which can not be estimated directly.

With respect to B-mode ultrasound, the physics of the image formation process poses a significant source of uncertainty. Ultrasound images exhibit a wide range of image artifacts and require a large amount of training in order to be interpreted correctly [2, 158]. Thus, being able to estimate the amount of uncertainty present in an image is an important step to improve ultrasound imaging. One major source of uncertainty is the presence of signal attenuation in terms of shadowing and dropout artifacts. Different approaches have been developed to estimate the amount of signal loss to improve the results of tasks such as segmentation, tissue characterization and image registration [122, 123]. So far, however, ultrasound uncertainty estimation has always been an offline task that is not easily integrated with one of the key benefits of ultrasound imaging, its interactive character. In this chapter, we discuss and present a real-time capable method to estimate ultrasound signal attenuation from the B-mode images themselves (cf. Figure 5.3). We extend the original formulation of ultrasound Confidence Maps [77] by an incremental solver scheme and show that the accuracy of results is sufficient for real-time visual computing and visualization applications.

Parts of this work have been published in [163].

5.1 The B-mode Ultrasound Image Formation Process

While the general term *ultrasound* describes the entire frequency range of acoustic waves that are not perceivable by the human ear (20 KHz to 1 GHz), ultrasound in the context of medical imaging is used as a synonym for the imaging modality using frequencies between 1 and 60 MHz. Historically, the first works on understanding the fundamentals of ultrasound date back to the 18th century to the findings of Euler, Lagrange, Rayleigh and many others [21, 38]. From the technical point of view, a major step toward today's ultrasound machines was the discovery of the Piezo-electric effect [43], allowing to transform mechanical stress into electric potentials and vice versa. While early developments were only considering 1D transmission over time (Amplitude or A-mode ultrasound), the key work toward 2D cross-sectional ultrasound imaging was the system of Wild and Reid [204] as it introduced showing the amount of reflection as brightness and thereby shaping today's B-mode imaging.




Fig. 5.1. Illustration of a longitudinal wave exhibiting compression and rarefaction. λ describes the length of the wave.

The ultrasound image formation process starts with the transducer sending out a short ultrasound pulse sequence into the body. This pulse propagates as longitudinal wave (cf. Figure 5.1) through the tissues and gets partly absorbed and reflected at different depths until these echoes eventually reach the transducer again where they are recorded. After several processing steps and assuming a constant speed of sound of $1540 \frac{m}{s}$, the intensity and runtime of the ultrasound echoes, representing reflectance and depth, can be used to form a 2D image of the underlying anatomy.

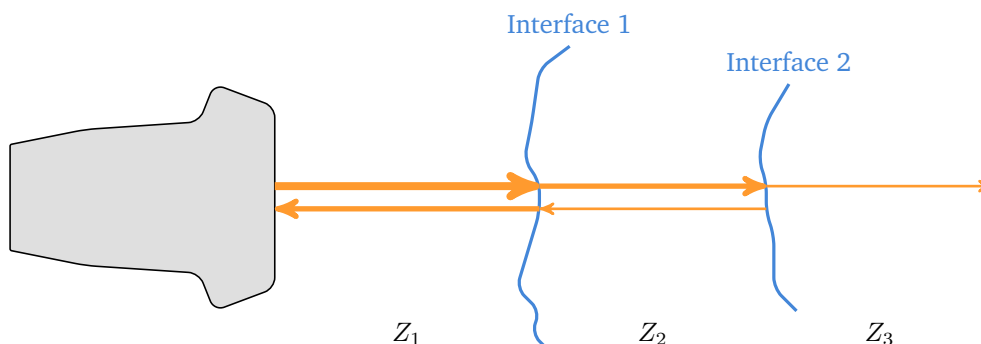


Fig. 5.2. Simplified illustration of reflection and transmission at tissue interfaces. When an ultrasound wave crosses an interface between two tissues of different acoustic impedance, parts of the signal get reflected while the remainder continues traveling through the anatomy (cf. Equation 5.1).

More precisely, the longitudinal ultrasound wave propagates as areas of compression and rarefaction (cf. Figure 5.1) through the anatomy inducing several different interactions with

the tissue. Within homogeneous media, absorption is the predominant effect where energy is dissipated over depth leading to a mostly uniform attenuation of the signal. Furthermore, wave interference and diffraction, as well as further non-linear effects occur. For diagnostic imaging however, the wave interactions with inhomogeneous media, namely reflection, transmission, refraction and scattering, are much more important. Interfaces between two tissues of different acoustic impedance lead to a reflection of parts of the signal, while the remainder of the signal is further transmitted (cf. Figure 5.2), yielding a reflection coefficient R_i and a transmission coefficient T_i :

$$R_i = \left(\frac{Z_2 - Z_1}{Z_2 + Z_1} \right)^2, \quad T_i = \frac{4Z_1Z_2}{(Z_2 + Z_1)^2}. \quad (5.1)$$

According to Snell's law and the Fresnel equations, this reflection and transmission is angle-dependent so that the longitudinal wave is refracted and its direction changed if an interface is not perpendicular to the wave front. Since a tissue is never completely homogeneous, this process happens continuously and yields the characteristic ultrasound speckle patterns where the inhomogeneities are smaller than the wave length λ , often modeled as random point scattering.

This image formation process gives B-mode ultrasound its characteristic gradient-like appearance, where the intensities show the changes in physical properties rather than the physical properties themselves. In addition, these intensities are highly direction-dependent since the amount of observed reflection intensity decreases with increasing angle between sound direction and interface normal. Thus, estimating signal attenuation and general uncertainty from B-mode images is not a trivial task, for which different approaches have been presented [122, 123].

5.2 Confidence Maps for B-mode Ultrasound

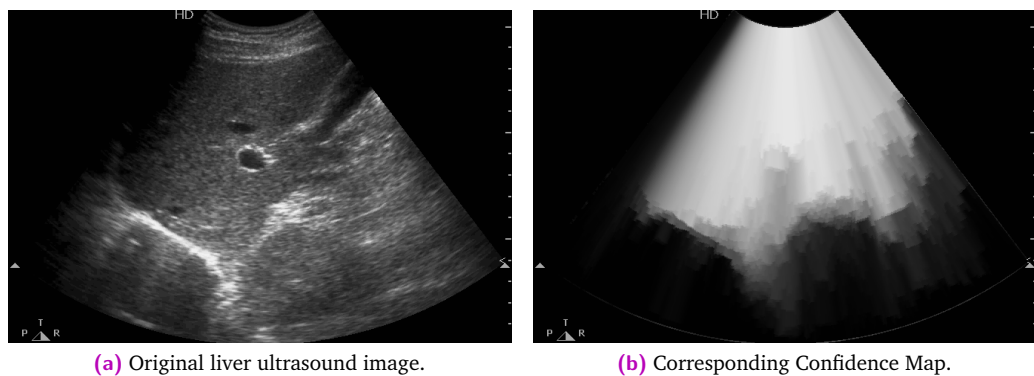


Fig. 5.3. Illustration of Confidence Maps on liver ultrasound: (a) shows the original ultrasound image, (b) shows the corresponding Confidence Map, bright regions depicting regions of high confidence, dark regions depicting regions of low confidence.

A generalization of different concepts of ultrasound uncertainty estimation are ultrasound Confidence Maps, originally proposed by Karamalis et al. [77] and later refined by Hennemersperger et al. [64]. They describe a per-pixel estimation of the amount of uncertainty present in ultrasound images due to signal attenuation and shadowing, for which they employ a random

walks formulation similar to the one by Grady to solve the k -label image segmentation problem [59]. Taking ultrasound physics into account, Confidence Maps describe the probability of a random walker starting from a given pixel to reach the virtual transducer elements.




Fig. 5.4. Illustration of the diffusion graph network for computing Confidence Maps. Each node is connected to its 8-neighborhood where the edge weights are defined as a function of the edge direction and gradients in the B-mode image.

As illustrated in Figure 5.4, the image is represented as an undirected graph $G = (V, E)$, where the nodes $v \in V$ represent the image pixels, each being connected to its 8 surrounding nodes through the edges $e \in E$. Each edge $e_{ij} \in E$, connecting nodes v_i and v_j is assigned a weight $w_{ij} > 0$ describing the likelihood of the random walker crossing that edge. Furthermore, a set of virtual transducer elements is placed at the start and end of each ultrasound scanline to define the signal attenuation boundary conditions. The edge weights are given as

$$w_{ij} = \begin{cases} \exp(-\beta(|c_i - c_j| + \gamma)) & \text{if } i, j \text{ adjacent and } e_{ij} \text{ horizontal edge} \\ \exp(-\beta(|c_i - c_j|)) & \text{if } i, j \text{ adjacent and } e_{ij} \text{ vertical edge} \\ \exp(-\beta(|c_i - c_j| + \sqrt{2}\gamma)) & \text{if } i, j \text{ adjacent and } e_{ij} \text{ diagonal edge} \\ 0 & \text{otherwise,} \end{cases} \quad (5.2)$$

$$c_i = g_i \exp(-\alpha l_i),$$

where g_i is the image intensity at node i , and l_i describes the normalized closest distance to the virtual transducer elements. The three free parameters α, β and γ define the resulting Confidence Maps. While the α parameter describes the likelihood of random walks along vertical edges and thus scales the estimated attenuation with increasing depth, β affects the robustness and accuracy of the result. Furthermore, γ penalizes random walks along the horizontal and diagonal edges and thereby influences the amount of horizontal discontinuities. A more thorough discussion of the edge weights and the free parameters can be found in

the original paper [77]. In their evaluation, Karamalis et al. demonstrate that ultrasound Confidence Maps describe an accurate per-pixel estimate of the signal attenuation, which can be translated to the amount of uncertainty present in the image. Typical applications are ultrasound shadow detection, 3D freehand ultrasound compounding and multi-modal registration.

Due to the formulation as diffusion problem with Dirichlet boundary conditions, where the scanline sources have full confidence and the scanline sinks have zero confidence, Confidence Maps always exhibit the whole value range of $[0, 1]$. Thus, the computed values describe only relative information with respect to the current image. As a consequence, one has to pay special attention when comparing Confidence Maps of different images. While the values of images of the same sequence may be comparable, confidence values between images of different anatomies are not directly numerically comparable.

5.3 Traditional Implementations

Being a Dirichlet diffusion problem, the computation of ultrasound Confidence Maps can be efficiently encoded in a linear system of the form

$$Lx = \mathbf{b}, \tag{5.3}$$

where L describes the Laplacian matrix of the graph, x is the Confidence Map in vectorized form and \mathbf{b} encodes the Dirichlet boundary conditions. Since the graph is undirected and describes an 8-neighborhood, L is symmetric, sparse, and positive definite. Hence, Equation 5.3 is traditionally solved with direct methods such as LU decomposition [77].

Unfortunately, direct solvers do not accommodate the interactive characteristics of ultrasound imaging well. Even with the relatively moderate resolution of ultrasound images, real-time computation of Confidence Maps is not feasible with today's hardware. Karamalis et al. report computation times of just above two seconds [77], which was reproduced by our experiments. Despite this slow computation time, Chatelain et al. use Confidence Maps for servoing of robotic ultrasound [34]. However, in order to yield interactive frame rates they have to significantly downsample the ultrasound images and it is unclear how much this impairs the quality of the results. Thus, other methods are needed in order to reliably solve Equation 5.3 in real-time.

5.4 Real-Time Incremental Solving

In order to yield real-time Confidence Maps for our system, we leverage the dynamic nature of ultrasound acquisitions and the temporal coherency of its images based on the high acquisition rates of today's systems. Consecutive images in ultrasound sequences usually differ very little and therefore the corresponding Confidence Maps will likely be very similar. We exploit this fact by using an iterative Jacobi-preconditioned Conjugate Gradient (PCG) solver instead of directly solving Equation 5.3 through matrix decompositions. This allows us to execute an

incremental computation scheme, where we directly use the resulting Confidence Maps as initialization for the subsequent frame (cf. Figure 5.5). This greatly reduces the number of needed iterations per frame as we have a better convergence to the true solution in the limited time budget available per frame.




Fig. 5.5. Illustration of our proposed incremental computation scheme for Confidence Maps. Due to the temporal coherency of ultrasound imaging, consecutive images and thus also their Confidence Maps are usually very similar. Therefore, we use the computed Confidence Map as initialization for the iterative solving of the subsequent frame.

We construct the equation system matrix L explicitly on each frame. Due to the given graph structure, the matrix has only 9 diagonals with non-zero entries and can thus be efficiently stored using the sparse DIA matrix storage format [6]. For the very first frame, we initialize x with a linear gradient as very rough approximation of the Confidence Map. Then, for every subsequent frame, we use the computed result as initialization. For easier integration with visualization techniques and to allow for an interactive experience in general, we implemented our technique on the GPU using CUDA. The preconditioned conjugate gradient solver has a specific time budget depending on the target frame rate, in which it performs as many iterations as possible. Even though this time budget might be too small to yield exact solutions from the beginning, our later evaluation will show, that the temporal coherency of ultrasound sequences is sufficient to have the solver converge within very few frames.




Fig. 5.6. Processing pipeline of our reference implementation. We receive the original ultrasound B-mode image from the acquisition stream through OpenIGTLink and perform a resampling in order to improve the PCG convergence performance. The Confidence Map solver then incrementally computes the Confidence Map by using the previous image's Confidence Map as initialization. The resulting Confidence Map will be used as initialization for the subsequent frame.

As already discussed in Section 5.2, Confidence Maps exhibit only a relative measure of confidence. Their strength is not the exact confidence value at a single pixel, but rather the distribution within the image. Thus, due to noise in the original ultrasound, the Confidence Maps of multiple consecutive frames may exhibit a flickering behavior when being watched in

a sequence. To introduce a better temporal coherency, we apply an additional alpha beta filter [20]. Being a variant of the Kalman filter, it recursively operates on the stream of computed Confidence Maps and produces a smoothed version by averaging the current image with a prediction based on the previous images. We empirically selected a configuration of $\alpha = 0.36$ and $\beta = 0.005$ providing good results with both damping of flickering and preservation of the original confidence distribution and temporal responsiveness of the estimates.

5.5 Evaluation

For the evaluation of our methods, we had a professional sonographer acquire several sequences of patient abdominal ultrasound. Our system was run on a Linux notebook with an Intel i7 processor and a nVidia GTX 750M GPU. Instead of acquiring live images from the ultrasound machine, we streamed the pre-recorded sequences from the hard disk via OpenIGTLink link, each containing 392 images of 512x512 pixels resolution. To examine the influence of the system performance with respect to the different parameters, we ran our real-time solver in various configurations.

Resample Scale vs. Number of Iterations The analytical solving of Equation 5.3 using standard Cholesky decomposition requires an average of more than 2.2 seconds per frame and is thus far from real-time. In contrast, our proposed incremental solver scheme allows to stop the iteration process at any time and use the best possible solution that can be computed within a give time budget. We set this to 30ms in order to ensure real-time applications with a target frame of 30 fps. As shown in Figure 5.7, the image size has a significant impact on




Fig. 5.7. Solver runtime per frame for the kidney data set with different configurations of PCG iterations and resample scales targeting 30ms per frame. Smaller images yield a smaller problem size and thus allow for a higher number of iterations.

the runtime performance, since the problem size can be significantly reduced by resampling the ultrasound images to a smaller resolution, compute the corresponding Confidence Maps and then upsample the result back to the original resolution. While our test platform allows

for only 30 PCG iterations in full resolution, 50 iterations are feasible with a resample scale of 0.75 in each direction and over 100 iterations are possible with a resample scale of 0.5 in each direction. Though the frame rates are rather constant, a jitter in the computation times is present. Therefore, we do not limit our incremental solver by the number of iterations but by a given time budget.

Resample Scale vs. Image Quality While it is no surprise that smaller images allow for a higher number of PCG iterations, the interesting question is what impact the downsampling has on the image quality of the computed Confidence Maps. We use the Structured Similarity Index Measure (SSIM) with a window size of 9×9 pixels for a quantitative comparison of our results with the exact reference solution. By comparing the patch-wise covariance between the two images with the images' individual variances, the SSIM computes the degradation of structural information [195]. This models that the human visual perception is more sensitive to changes of object structures than to absolute errors of intensities. Since we intend to use our real-time uncertainty estimation framework mainly for visualization purposes, the SSIM is thus an excellent quantitative measure to assess the image quality.

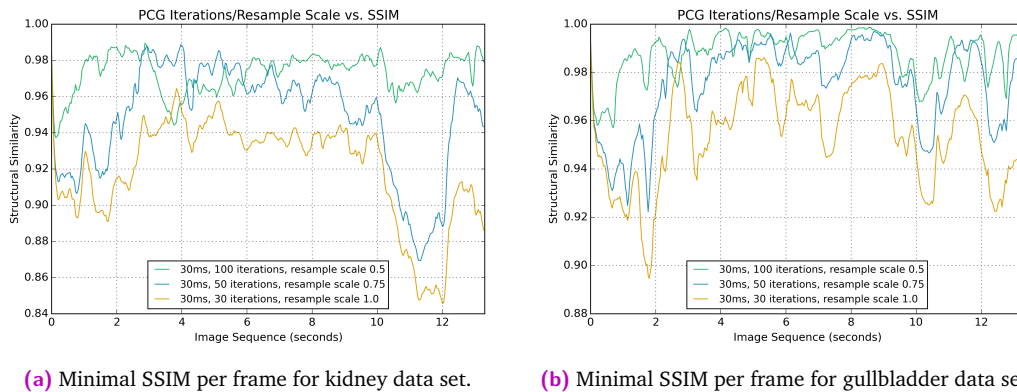



Fig. 5.8. Evaluation of the image quality with respect to the resample scale in terms of structural similarity. When maintaining a target time budget of 30ms, a lower resample scale yields a better image quality.

Smaller images exhibit both a smaller computational burden and a higher convergence rate of the solver, so that the resampling may even be beneficial for the resulting image quality. However, at the same time, the lower resolution of the images may remove small details, which are important for the correct estimation of the signal attenuation. When looking at Figure 5.8, one can see that in general a slight downsampling of the images is beneficial for the final image quality. For a couple of frames however, in particular with the kidney data set, details get lost yielding a worse structural similarity applying less PCG iterations on the full scale images. With these results, we decided that using a resample scale of 0.5 defines a good trade-off between computation time and image quality (cf. Figure 5.8).

PCG Iterations vs. Image Quality As a final quantitative evaluation, we ran our proposed solver scheme with a fixed resample scale of 0.5 but different number of iterations. As illustrated in Figure 5.9, the SSIM never falls below 0.94 when performing 110 PCG iterations, which is an excellent result. For comparison, the dashed lines show the error progression when the solver is reinitialized at each iteration, which yields mostly worse results than 20



(a) Minimal SSIM per frame for kidney data set. (b) Minimal SSIM per frame for gullbladder data set.

Fig. 5.9. Error in terms of structural similarity over time for different PCG iteration counts. The resample scale was fixed to 0.5.


incremental iterations. This demonstrates the importance of leveraging the dynamic nature of ultrasound and the effectiveness of our technique using the solution of the previous frame as initialization. Table 5.1 further summarizes the results and shows the average runtime, average SSIM and minimum SSIM over the entire kidney data set. One can see that the 110 incremental PCG iterations of our reference implementation are real-time capable and yielding almost exact results with an average SSIM of 0.999 and a minimum SSIM of only 0.941. Figure 5.10 presents qualitative results on a representative frame of the same ultrasound sequence.

# Iterations	Avg. Runtime (ms)	Avg. Quality (Structural Similarity)	Min. Quality
Direct solver	2214.68	1.0	1.0
20 (incremental)	10.84	0.991	0.837
50 (incremental)	18.07	0.997	0.899
80 (incremental)	25.37	0.998	0.929
110 (incremental)	32.46	0.999	0.941
110 (discarding)	32.59	0.984	0.856


Tab. 5.1. Aggregated performance results in terms of runtime and error with different configurations for the number of iterations. The resample scale was fixed to 0.5. The data set was a patient kidney ultrasound sequence containing a total of 392 frames.

5.6 Conclusion


We presented a method to bring the estimation of B-mode ultrasound signal attenuation into real-time applications. Building on ultrasound Confidence Maps, originally proposed by Karamalis et al. [77], we propose an incremental solver scheme, which leverages the temporal coherency of ultrasound imaging. In an extensive evaluation using clinical patient data, we assessed the impact of our different optimizations on the quality of the computed Confidence Maps and showed that our methods yield accurate approximations within a 30ms




(a) Original ultrasound frame from kidney sequence.




(b) Corresponding Confidence Map.




(c) Computed Confidence Map with 30 iterations in full resolution (30ms).



(d) Structured Similarity between (b) and (c); average: 0.984, minimum: 0.889.



(e) Computed Confidence Map with 100 iterations in half resolution (30ms).



(f) Structured Similarity between (b) and (e); average: 0.998, minimum: 0.973.

Fig. 5.10. Qualitative results of the presented incremental solver scheme. By comparing (b) with (c) and (e), one can see the impact of the limited amount of iterations on the resulting Confidence Maps. The SSIM in (d) and (f) quantifies the perceptual difference. For better visualization of differences, the color map in (d) and (f) was normalized to the range of $[0.85, 1.0]$ and does not show the full SSIM range of $[-1, 1]$.

time budget. In the remainder of this thesis, we will use this work to generate real-time uncertainty information for different ultrasound processing and visualization techniques.

Uncertainty Visualization for 2D Ultrasound

” un·cer·tain·ty, *noun*

1. *the quality or state of being uncertain*
2. *something that is doubtful or unknown : something that is uncertain*

— Merriam-Webster’s English Dictionary

Though there are many advantages of ultrasound imaging, such as being rather low-cost and real-time capable, the correct interpretation of B-mode images is a challenging task that requires a large amount of experience and training. Even ultrasound experts sometimes struggle in performing an ultrasound-based diagnosis due to the presence of many different kinds of artifacts and its non-homogeneous distribution of uncertainty. In this chapter, we present novel visualization techniques that augment the B-mode images with uncertainty information in real-time in order to support both ultrasound novices and expert users in gaining a better understanding of ultrasound. Therefore, we build upon the previously introduced method of real-time confidence estimation for B-mode ultrasound images (cf. Chapter 5) and interpret these Confidence Maps as per-pixel uncertainty information, which we expose to the user using perceptual visualization techniques. After a review on general uncertainty visualization techniques for medical applications, we discuss the clinical significance of our proposed system. In total, we present three different visualization schemes for both educational and clinical applications and motivate the selection of visual variables to depict uncertainty in B-mode ultrasound images. An extensive evaluation conducted with both ultrasound novices and expert clinicians demonstrates the usefulness of our techniques.

Parts of this work have been published in [163].

6.1 Uncertainty in Health Care and Medical Visualization

A comprehensive look at health care suggests that the concept of uncertainty is ubiquitous throughout all levels. However, since clinicians have been trained to make decisions and provide their patients with answers, many of them struggle to admit that their work is based on a plethora of input factors, each inducing an individual level of uncertainty. Thus their eventual findings represent the most likely interpretation rather than the actual truth.

Han et al. investigate the different concepts and appearances of uncertainty in general health care using cancer treatment as example and derive a conceptual taxonomy in [62]. They point out that a coherent concept of uncertainty in health care is missing and that multiple meanings and interpretations of that term are present, which are not necessarily distinguished from each other. They conclude that uncertainty is actually a threefold concept which can originate either through probability, ambiguity or complexity with issues ranging from disease-centered to patient-centered ones.

While the paper of Han et al. analyzes the problem from the clinical perspective, Ristovski et al. focus on the various technical aspects and investigate the presence of uncertainty in medical visualization [146]. As introduced in Section 2.3, the visualization pipeline ranges from the initial acquisition of data, over several data processing steps to the final rendering output. Though each of these steps is subject to an individual level of uncertainty, the final rendering often shows the data as if it were the only possible truth. However, it requires appropriate visualization to make the amount of uncertainty assessable to the clinician and thereby propagate the information back correctly to the clinical perspective as it was previously presented by Han et al. For instance, Lundström et al. point out that a suboptimal setup of a transfer function during rendering can lead to a false classification of stenosis [101].




Fig. 6.1. Uncertainty visualization of radio-frequency ablation zones as proposed by Rieder et al. [144]. The red isolines around the applicator and the tumor are depicting the estimated ablation zone incorporating the heat sink effect. The blue isolines around the vessels illustrated their cooling factor. Image from [144].

The special setup in which medical visualization is used and its resulting requirements make it difficult to apply techniques from other domains, such as [132, 135], even though they use very similar data. This may be one of the reasons why the body of literature on medical uncertainty visualization is rather shallow and the individual works are very application- and task-specific. One such example is the uncertainty visualization technique of Rieder et al. proposed in the context of percutaneous radio-frequency ablation. For this application, precise planning of the applicator placement is crucial in order to ensure that the malignant tissue is destroyed completely (cf. Figure 6.1). Their work is particularly noteworthy as they use their visualization for the optimization of a seven-dimensional optimization problem (five degrees of freedom for the placement of each needle plus two free simulation parameters) in three-dimensional space [144]. A similar approach is the work of Brecheisen et al. exploring

the parameter sensitivity in fiber tracking algorithms for DTI data [18]. Both techniques rely on interactive exploration by the user in order to cover the available parameter space.




Fig. 6.2. Uncertainty visualization of segmentation algorithm results as proposed by Prassni et al. [137]. The yellow isolines in the 2D depict the uncertainty of the segmentation result. The 3D visualization additionally uses opacity modulation for depicting uncertainty. Image from [137].

One example of providing feedback on the uncertainty present in medical imaging algorithms is the work of Praßni et al. They propose a guided probabilistic volume segmentation technique in which they visualize the resulting uncertainty distribution using iso-lines and opacity modulation (cf. Figure 6.2). As mentioned previously, it is crucial to expose information on the quality of the result to the clinician, since ultimately they have to decide on diagnosis and treatment. The recent progress on ensemble visualizations [63, 125, 141] could be influential for further progress in this direction. In the medical domain such techniques can for instance be used to visualize the results of large random studies on medical imaging algorithms and thereby make them more accessible.

6.2 Clinical Significance and Application

As discussed before, acquiring a good image (e.g. in terms of high diagnostic value) is not a trivial task due to the highly complex ultrasound image formation process (cf. Section 5.1). It is influenced by various physical imaging parameters such as frequency, focus, and depth, as well as by external factors such as probe positioning, probe pressure, patient positioning and patient breathing cycle [2], and can yield a wide range of image artifacts [158]. Furthermore, some target anatomies can not be directly reached but need to be scanned by circumventing strong reflectors such as bones, which prevent the acquisition of images underneath. A classic example of such an anatomy are the kidneys, which can not be scanned from the back as they are then in the acoustic shadow of the spine reflecting almost all ultrasound waves. Instead, sonographers perform kidney ultrasound through the abdomen, usually using the liver as an acoustic window since it has the best transmission properties of the surrounding anatomy.

Sonographers need to be aware of all these caveats of ultrasound imaging in order to correctly understand the image. In particular medical trainees and ultrasound novices have difficulties in getting the right image needed for their clinical objectives since traditional ultrasound imaging does not provide a direct qualitative feedback on the image quality. Already a slight repositioning of the transducer can yield a significantly better acoustic window and thus improve the image for the target anatomy (cf. Figure 6.3). Thus, training is an important aspect for medical students when learning ultrasound [28]. In particular in trauma applications,




Fig. 6.3. Kidney ultrasound using the liver as acoustic window with applied uncertainty visualization using chroma as visual variables (cf. Section 6.4). A slight repositioning of the transducer results in a considerable increase of confidence in image (b) compared to (a).

where time is critical, the surgeon has to determine possible fractures and lesions as quickly as possible and has a minimal margin for error [84]. With this motivation in mind, we developed our work with the aim to support both medical students in learning sonography as well as expert clinicians for a more quick and intuitive interpretation of ultrasound images by providing an interactive feedback on the image quality and uncertainty distribution.

6.3 Selection of Visual Variables

Essentially, we target two different applications with our work: On the one hand, we are convinced that exposing uncertainty information to ultrasound novices and trainees helps them to better understand the complex ultrasound image formation process. At the same time, we also expect expert sonographers to benefit from uncertainty visualization as the additional information may improve the diagnostic value of the image. Regarding these two target applications, we yield the following requirements for ultrasound uncertainty visualization schemes.

- The uncertainty information should be fused directly into the original B-mode sequence. Thus, both the spatial and the temporal domain are fixed.
- The primary information in the uncertainty maps is the distribution of the uncertainty with respect to the anatomy rather than the exact per-pixel uncertainty values (cf. Section 5.2).
- For educational applications, the uncertainty distribution in the image should be easily and intuitively perceivable. Even small changes in the uncertainty distribution should be clearly observable when repositioning the ultrasound probe in order to maximize the learning effect.
- For clinical applications, the diagnostic information in the B-mode image must not be impaired. Thus, the original image intensities should be preserved as good as possible and no image regions should be occluded.

Given these requirements for our uncertainty visualization, the number of viable visual variables for depicting the uncertainty information is limited.

One traditional technique is the usage of glyphs for depicting uncertainty, with error bars in 1D visualization being the classic example. Glyphs have the advantage of offering a large number of visual variables that can be used to alter their appearance. For instance, MacEachren et al. evaluate 11 different mappings of uncertainty to point glyphs [102]. Glyphs excel in depicting uncertainty when used in sparse layouts allowing the observer to individually focus on single glyphs in order to read their information. However, this does not work well for our application where the goal is to visualize the distribution of a dense 2D scalar field. Although dense glyph fields have also been successfully used for depicting global information [13, 80], we do not consider them for our work since early experiments did not show promising results. The work of Sanyal et al. supports this fact as their experiments showed that in many uncertainty related tasks on 2D data sets the different glyph mappings perform significantly worse than surface color mapping [157]. Furthermore, adding glyphs to the B-mode image would occlude the original ultrasound image, which is undesirable for clinical applications.

One intriguing approach is to extend the 2D data to the third dimension and map uncertainty to the Z axis in a 3D rendering [23, 75]. While this may be a valid method for applications such as geospatial visualization, it can not be applied to our use case since, due to the 2D projection of a 3D scene, it requires the user to interact with the camera to get the full information. Furthermore, as mentioned above, the spatial domain of our visualization is fixed as clinicians expect a 2D image when performing 2D B-mode ultrasound.

Another approach is to exploit the spatial domain for depicting uncertainty, for instance through animations, animated jittering or probabilistic animation [23, 101]. However, since we are working with real-time ultrasound sequences, the temporal domain is fixed and such approaches are not applicable.

6.4 Visualization Schemes

Given these considerations and the particular requirements of our intended uncertainty visualization, we selected the visual variables of color and texture. They both do not affect the spatial or temporal image domain, and are very powerful and intuitive for expressing general uncertainty [102]. In total, we propose three different uncertainty mapping techniques for the two applications, which we will discuss in detail in the following sections. Since both the B-mode image and the uncertainty map are in the same image domain, no coordinate transformation is necessary and the mapping techniques are focused on the optical properties.

As illustrated in Figure 6.4, all proposed mapping schemes start with the original B-mode ultrasound image I , for which we compute the corresponding Confidence Map (cf. Chapter 5). We assume it to be inversely related to the amount of uncertainty in the image, more precisely to its facets of accuracy, precision and credibility. Thus, we obtain the per-pixel uncertainty information U by applying a direct inverse linear mapping

$$U(x) := 1 - CM(x), \quad (6.1)$$

where $CM(x) \in [0, 1]$ is the Confidence Map value at pixel x in the image domain.

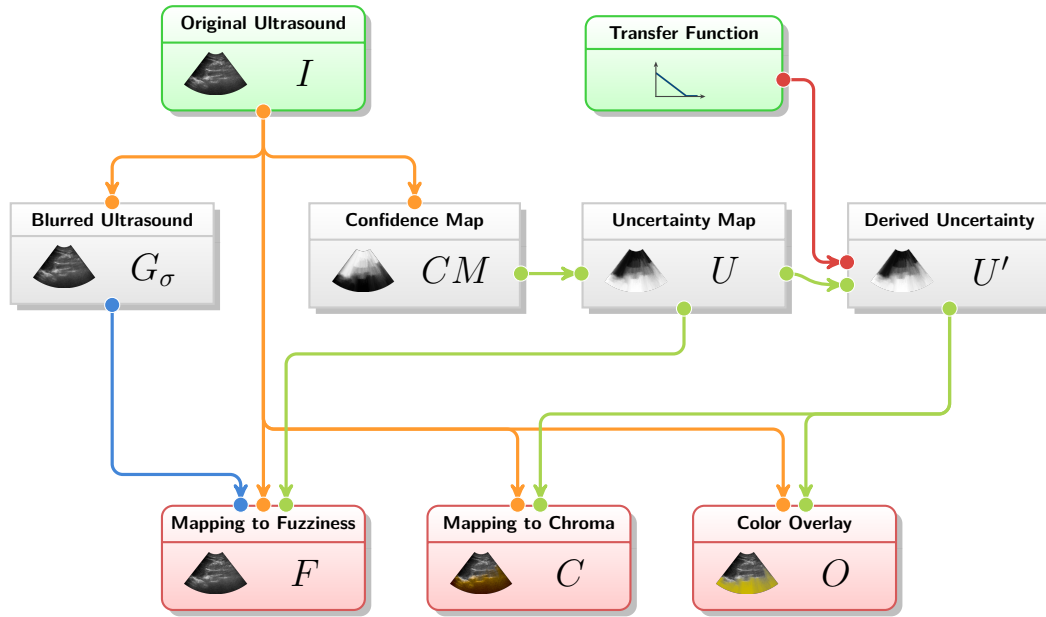


Fig. 6.4. Schematic diagram of the different proposed uncertainty visualization schemes. Given the original B-mode ultrasound image I , we compute its Confidence Map CM and a Gaussian blurred version G_σ . The uncertainty map U is derived from CM by inverse linear mapping (cf. Equation 6.1). For the color overlay and the uncertainty mapping to chroma, we compute a derived uncertainty measure U' using a transfer function. Finally, the uncertainty information is fused into the original ultrasound image using one of the three visualization schemes.

6.4.1 Uncertainty as Color Overlay

For educational applications, the focus of the visualization should be on the uncertainty information and even small changes in the distribution should be clearly distinguishable by the observer. At the same time, the corresponding ultrasound B-mode image should be shown as anatomical reference in order to allow for an understanding of the connection between image features and their effects on the uncertainty. Therefore, we combine the visual variables of hue and value in our proposed color overlay scheme. Compared to the other presented mapping schemes, the combination of the two visual variables makes even subtle changes of uncertainty visible to the observer. We deem this an important feature to teach ultrasound novices the caveats of the ultrasound image formation process.

As previously introduced, for each pixel x , let $U(x)$ be its uncertainty value and $I(x)$ be its original B-mode intensity. Since the color overlay is a very obtrusive mapping scheme, we use a transfer function to apply a thresholding to the uncertainty measure and define the derived uncertainty as $U'(x) := \max(0, 2U(x) - 1)$. Using this derived measure instead of the original $U(x)$ avoids overlaying regions of negligible uncertainty. We first generate the color overlay in HSV color space as

$$C(x)_{HSV} := (H, U'(x), V), \quad (6.2)$$

with constant hue $H \in [0, 1]$ and constant value $V \in [0, 1]$. We chose a bright orange color with $H = 0.15$ and $V = 0.8$ to avoid lowering the contrast to Doppler ultrasound utilizing the

colors blue and red. In a second step, we linearly mix the color overlay transformed to RGB color space with the original B-mode to yield the final pixel color

$$O(x) := U'(x) \cdot C_{RGB}(x) + (1 - U'(x)) \cdot I_{RGB}(x). \quad (6.3)$$

Figure 6.5 shows the uncertainty color overlay applied to two different abdominal ultrasound images.




Fig. 6.5. Illustration of the color overlay mapping scheme applied to kidney ultrasound. The high contrast yellow overlay allows for a clear analysis of the signal loss effects in ultrasound.

Since the original B-mode intensities are altered in this mapping scheme, we propose to use it only for educational purposes to give ultrasound novices a better understanding of the image formation process, but do not consider it for clinical usage. Furthermore, the color overlay also partially occludes the original ultrasound intensities. Although image regions of low confidence may not be reliable enough for diagnosis, they still contain structural information, which can help the sonographer in navigating toward the correct anatomy and optimizing the acoustic window. Thus, hiding these parts completely is disadvantageous for clinical usage, which was later also confirmed by some candidates during our evaluation (cf. Section 6.5).

6.4.2 Uncertainty Mapping to Chroma

For clinical applications, we propose additional uncertainty visualization schemes that maintain the structural information in the ultrasound B-mode image also in unreliable regions. Similar to the color overlay, uncertainty mapping to chroma also uses color to depict uncertainty but uses a different visual variable.

In order to preserve the diagnostic value of the original ultrasound image, we need to ensure that the perceived intensity remains the same when augmenting the image with uncertainty information. Therefore, we perform the chroma modification in the perceptually uniform CIE L*a*b* color space, the polar coordinate derivative of the CIE L*a*b* color model [134].

As illustrated in Figure 6.4, we compute the derived uncertainty from the original Confidence Map as $U'(x) := \max\left(0, \frac{3U(x)}{2} - \frac{1}{2}\right)$ to again avoid coloring regions with negligible uncertainty. The final pixel color in L*a*b* space is given by

$$C(x)_{L^*C^*h^*} := (I(x)_{L^*}, U'(x), H), \quad (6.4)$$

where $I(x)_{L^*}$ is B-Mode intensity transformed to L^* space and $H = 0.23$ is a bright orange. Again, we chose bright orange as hue for depicting uncertainty to avoid lowering the contrast to Doppler ultrasound. Sample images are shown in Figure 6.6.




Fig. 6.6. Illustration of the chroma mapping scheme applied to kidney ultrasound. Unreliable regions in the image are depicted with an orange overlay. Since the color transformation is applied in the perceptually uniform CIE $L^*C^*h^*$ color space, the perceived pixel intensity is maintained.

6.4.3 Uncertainty Mapping to Fuzziness

Many clinicians prefer to look at gray scale ultrasound images as they have been trained to do so. Since texture is a very effective visual variable that keeps the spatial, temporal and color domain of the original image, we selected it as third uncertainty visualization scheme for real-time B-mode ultrasound. More precisely, we use the visual cue of fuzziness to map the uncertainty information. Therefore, we fuse the ultrasound image with its uncertainty map such that regions of low uncertainty appear sharp and regions of high uncertainty appear fuzzy. As a matter of fact, according to MacEachren et al., this visual variable is also the most intuitive to represent uncertainty [102].

Our proposed uncertainty mapping to fuzziness combines a slight Gaussian blur of the original ultrasound image with its unsharp mask (subtraction of the blurred image from the original image). As illustrated in Figure 6.4, we compute the Confidence Map of the original image and apply Equation 6.1 to obtain the uncertainty value $U(x)$ for each pixel x . Since this is a diverging mapping scheme, where regions with high uncertainty will be blurred, and regions of low uncertainty will be sharpened, we can directly use $U(x)$ and do not need to compute a derived uncertainty measure. We compute a Gaussian filtered version of the original image and combine the two to yield the final pixel value $F(x)$ as

$$F(x) = U(x)G_\sigma(x) + ((1 - U(x)) \cdot (2I(x) - G_\sigma(x))), \quad (6.5)$$

where $I(x)$ is the original ultrasound intensity and $G_\sigma(x)$ the corresponding intensity in the Gaussian with parameter σ . We choose $\sigma = 2.5$ to limit the blurring to a tolerable amount but still yield the effect of perceived differences in fuzziness. An exemplary result can be seen in Figure 6.7.

It should be noted that this mapping scheme certainly alters the original B-mode image in a way that may reduce the amount of original information in regions of high uncertainty. However, discussions with clinicians showed that they nevertheless like uncertainty mapping




Fig. 6.7. Illustration of the fuzziness mapping scheme applied to kidney ultrasound. Unreliable regions in the image are blurred while reliable regions are sharpened through unsharp masking. Especially in video sequences this provides a very strong perceptual cue for uncertainty while at the same time maintaining the ultrasound image in its original gray scale color domain.

to fuzziness and appreciate the intuitiveness of the visual variable. Our evaluation results in Section 6.5 underline this fact.

6.5 Results and Evaluation

We performed the evaluation of the proposed methods independently for the two target applications. Therefore, we implemented a fully working reference system using an Ultrasonix RP (Analogic Corporation, Peabody, MA, USA) ultrasound device for the acquisition of images. Using OpenIGTLink [182], we stream the ultrasound frames and imaging parameters to a standard workstation where our computation framework performs the necessary processing steps before eventually routing the fused image back to the device of the ultrasound device. The processing pipeline was implemented in the CAMPVis framework (cf. Appendix A) and is executed entirely on the GPU using both CUDA (solving of Equation 5.3) and OpenGL/GLSL (all other processing steps) in order to achieve optimal performance. After acquiring the B-mode image, we perform a Gaussian filtering as well as a resampling. Gaussian filtering is required in order to remove high-frequency noise as well as to allow for our uncertainty mapping to fuzziness. We perform the downsampling to speed up the computation of the Confidence Maps and achieve better convergence (cf. Chapter 5). For our experiment setup, we used a 0.5 scaling factor and a smoothing factor σ of 2.5. Finally, one of the discussed uncertainty visualization schemes is applied and the rectilinear B-mode image in polar coordinates is scan converted to Cartesian coordinates using the known probe geometry and dynamically queried imaging parameters.

6.5.1 User Study with Ultrasound Novices

In a first user study we evaluated the educational value of our proposed system. We equipped our ultrasound machine with an Ultrasonix C5-2/60 convex abdominal transducer and asked ultrasound novices to locate different structurally deep anatomies in a CIRS abdominal phantom while using our uncertainty visualization techniques. In total we interviewed 13 medical students, which all had very limited experience with ultrasound (median of 3 performed ultrasound examinations, minimum 0, maximum 20).

	Left Kidney	Right Kidney	Portal Vein
	Average Time (seconds)		
Original B-Mode	6.73 ± 4.3	4.96 ± 1.6	4.78 ± 1.2
Color Overlay	4.74 ± 3.8	4.26 ± 1.5	3.30 ± 1.7
Chroma Mapping	4.19 ± 2.3	3.42 ± 1.8	3.44 ± 1.5
Fuzziness Mapping	3.49 ± 3.9	3.15 ± 0.7	2.67 ± 1.1

Tab. 6.1. Quantitative evaluation results of our user study with ultrasound novices. The table shows the average time (*seconds*) required to optimize the view on target anatomies (aggregated results from 9 of the 13 users, since not all acquisitions were complete). With enabled uncertainty visualization, the users managed to decide faster when they had a good view on the target anatomy than with the plain B-mode image.

For a quantitative evaluation, we recorded the acquisitions and measured the time the users required to optimize the view on anatomical targets. After giving the participants some time to familiarize themselves with the phantom anatomy, we asked them to find an optimal view onto the vessel targets in the left and right kidney, as well as onto the portal vein. We then measured the required time to optimize the view for each visualization scheme by counting the number of frames between the first frame where the target anatomy was in the field of view until the frame where the students defined the view as optimal in their personal opinion. To avoid biasing the results, we shuffled the order of the visualizations for each user. While the results in Table 6.1 show no significant differences between color overlay and chroma mapping, the time needed with fuzziness mapping is consistently lower (in average 0.79 seconds) than with the one of other two mapping schemes. Furthermore, the students performed significantly worse (in average 1.86 seconds longer) with only the original B-mode image compared to all of our proposed visualization schemes, which supports our idea that visualizing uncertainty helps the user in interactively assessing the quality of the acquired ultrasound image.




Fig. 6.8. User study results on the educational value of our technique to ultrasound novices. Almost all test subjects appreciated the added uncertainty visualization and confirmed that it helps them in understanding the ultrasound image formation process and getting a better understanding of uncertainty in B-mode imaging.

After the experiment, the test subjects additionally answered a short questionnaire on different aspects of the presented visualization schemes. In the first set of questions we asked the

students whether they generally appreciate the additional information presented to them, whether it helps them to better understand the ultrasound image formation process and whether it helps them with understanding the concept of uncertainty in B-mode images. The results (cf. Figure 6.8) indicate a general level of appreciation for our presented techniques, since especially the color overlay method got very positive results. In the second set of questions, the students were asked to assess the intuitiveness of the presented visualizations and whether our technique helps them in finding anatomical structures faster. As shown in Figure 6.9, ultrasound novices find the colorful visualizations much more intuitive than the fuzziness mapping. However, only a few students found the presented visualizations helpful to find target anatomies faster. Interestingly, fuzziness mapping yielded an overall worse response in the questionnaire, as many students considered this scheme as not helpful for diagnosis nor intuitive to read. This result is particularly interesting as it contradicts the quantitative results of Table 6.1.

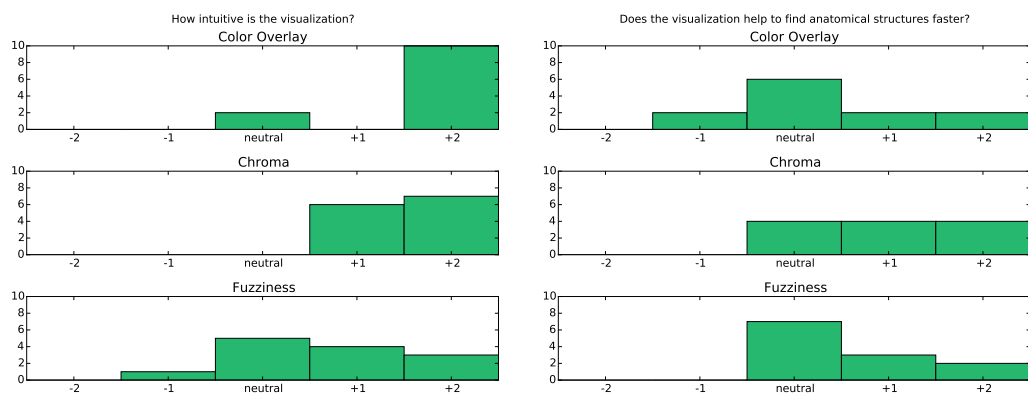


Fig. 6.9. User study results on the perception and clinical value for ultrasound novices on a 5-point Likert scale. While there is no clear favorite visualization scheme, a bias toward the colorful mappings is present. Furthermore, some students said that the added information helps them with finding target anatomical structures faster.

6.5.2 User Study with Ultrasound Experts

In a second user study, we presented our system to expert sonographers in order to evaluate the clinical significance of real-time ultrasound uncertainty visualization. Therefore we applied our visualization schemes to clinical data of abdominal ultrasound and presented the results to 7 experts (5 clinicians, 2 senior researchers). We presented them two patient abdominal ultrasound sequences of three different anatomies (liver, kidney, spleen) and asked them about perception, clinical value as well as whether our techniques assist in finding the optimal acoustic window. Since, for clinical usage, we do not want to hide information in the B-mode image, we evaluated only mapping to chroma and mapping to fuzziness.

As shown in Figure 6.10, the expert sonographers clearly prefer mapping to fuzziness over mapping to chroma. The results on general diagnostic value are quite mixed. While there is a slight positive response for the gray scale fuzziness mapping, the chroma mapping performs rather bad in this regard. This discrepancy compared to ultrasound novices is probably due to the fact that experts got used to the monochrome appearance of B-mode ultrasound over the




Fig. 6.10. Questionnaire results on expert sonographers' uncertainty perception. Each question was answered independently for the two different sequences. Therefore, there are 14 answers in total.

years. Thus, they prefer a visualization scheme that keeps the image in its original gray scale domain.

In addition to asking about the general diagnostic value, we also asked more specific questions regarding the clinical significance. Here, the study shows significant improvements compared to the default visualization in today's ultrasound devices (Table 6.2). The clinicians reported that seeing the amount of uncertainty dynamically adapting to the ultrasound view provides them with a very strong feedback on the image quality and its credibility. Here, 6 out of 7 stated that our uncertainty visualization helps with the correct interpretation of the images and also 6 out of 7 participants confirmed that the proposed real-time visualization schemes assist in optimizing the acoustic window. One clinician found that the computed Confidence Map was unexpected in one of the sequences and therefore confused him with the correct interpretation of the image. Nevertheless, he stated that the technique was helpful for optimizing the acoustic window. One other clinician confirmed the uncertainty visualization to be helpful for the correct interpretation but had doubts that it helps with optimizing the acoustic window.

	Yes	No
Helps with Correct Interpretation	86%	14%
Helps Optimizing Acoustic Window	86%	14%

Tab. 6.2. Questionnaire results on expert sonographers evaluating the clinical value of our uncertainty visualization technique.

6.6 Conclusion

In this chapter, we introduced a novel approach to 2D ultrasound visualization. Instead of only showing the plain B-mode image, we augmented it with additional uncertainty information based on the estimated per-pixel signal attenuation. Though different works have introduced the concept of uncertainty to ultrasound image processing, such information has never been exposed to the user before.

To do so, we build upon the previously introduced technique of real-time uncertainty estimation (cf. Chapter 5) and fuse Confidence Maps with their original B-mode image. Targeting both educational and clinical applications, we presented three individually designed visualization schemes. After implementing a fully working system, we ran two user studies with both ultrasound novices and experts sonographers. The results clearly show the benefit of our technique for educational purposes as the added feedback on the signal attenuation helps students interactively learn how the ultrasound image formation process works. Also most clinicians value the additional information since it can help them with optimizing the acoustic window on target anatomies.

Advanced Ultrasound Compounding



com·pound, *transitive verb*

1. to put together (parts) so as to form a whole
2. to form by combining parts

— Merriam-Webster's English Dictionary

So far in this thesis, we covered only two-dimensional ultrasound, which is also its most widely used form. However, being restricted to a 2D image plane imposes considerable limitations for some applications such as musculoskeletal (MSK) imaging, where ultrasound is one of the traditional imaging modalities amongst X-Ray and MRI. The 2D imaging plane of ultrasound is mostly perpendicular to the skin surface so that some anatomies, such as the tendons, can only be imaged in cross sections and not in their full extent. This complicates the identification of possible lesions and may impair diagnostic results [123].

Thus, extending ultrasound imaging to the third dimension provides an effective means to improve the clinician's workflow. Therefore, two different technical approaches exist. 3D matrix transducers use a 2D grid of elements in order to acquire a full three-dimensional volume at once. However, their field of view and image resolution is limited by the speed of sound and the used frame rate. An alternative approach is to equip a traditional 2D transducer with tracking hardware and acquire a sweep covering the entire anatomy. In a subsequent processing step, called compounding, a 3D volume is reconstructed from the 2D images and their corresponding poses.

In order to yield high quality compounding results, traditional techniques usually require constant probe pressure or linear sweep trajectories. For applications such as MSK ultrasound however, these constraints do not hold since the anatomy is prone to deformation and exhibits high curvature surfaces. In this chapter, we introduce a novel orientation-driven framework for compounding 3D volumes from tracked 2D B-mode ultrasound sweeps. Using a set of orientation-driven correlation terms, we perform a complementary pressure compensation and then cluster the ultrasound frames into groups of homogeneous orientation. These groups of frames are then compounded individually and eventually fused into the final volume using an information fusion approach based on uncertainty. Thereby, our technique yields accurate 3D reconstructions also with challenging acquisition setups where the ultrasound sweeps exhibit pressure changes or curved trajectories – situations where traditional compounding methods often fail. Finally, we propose a novel incremental compounding scheme, which allows for interactively updating the reconstructed volume during the acquisition and can thus be used in real-time applications.

Parts of this work have originally been published in [161].

7.1 Review of Ultrasound Compounding Techniques

Following Solberg et al.'s survey paper [170], ultrasound compounding algorithms can be classified into three different approaches:

Early approaches were the *pixel-based* or *forward warping methods*, which traverse the pixels in each 2D ultrasound frame, project each pixel's location into the coordinate system of the initially empty target volume and write the pixel's intensity information into the corresponding voxel (cf. Figure 7.1a). Pixel-based methods vary in the used averaging of multiple pixels contributing to a single voxel and in the employed hole filling of regions where the grid resolution is higher than the frame sampling rate of the ultrasound sweep. They are computationally rather cheap but their reconstruction quality is limited [170].

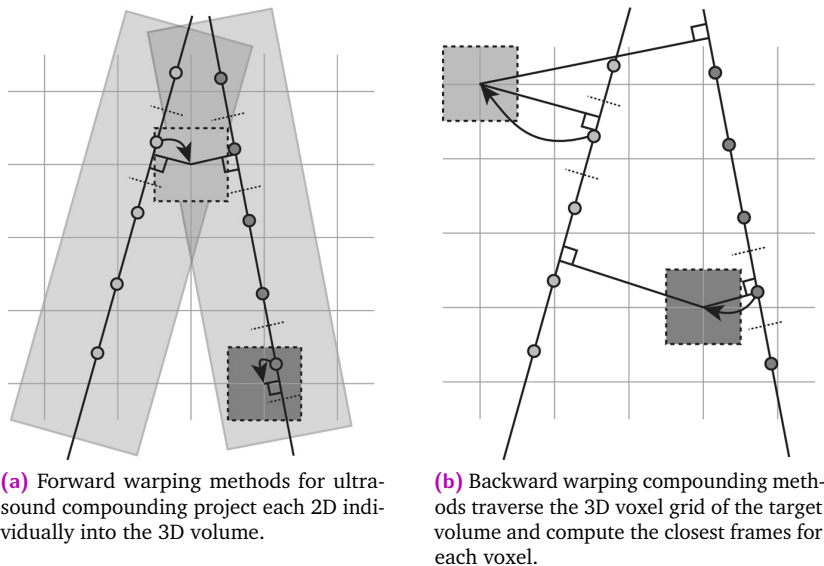


Fig. 7.1. Illustrations comparing forward warping methods with backward warping methods for ultrasound compounding. Images from [170].

Voxel-based methods use the inverse approach and are thus also referred to as *backward warping* methods. As illustrated in Figure 7.1b, they traverse the voxel grid of the target volume and back-project their position into the pixel space of each ultrasound frame to lookup the image intensity. Since multiple pixels may contribute to the final voxel value, they employ a weighting function based on intensity and/or distance. Wein et al. show that voxel-based methods yield superior quality compared to pixel-based methods [199]. Furthermore, backward warping algorithms are capable of reconstructing subsets of the target volume alone (e.g. a multi-planar reconstruction (MPR)), which offers a computational benefit over forward warping methods in such situations). However, their use in real-time interventional imaging is limited as their workflow requires to first acquire the whole sweep before the reconstruction process can start. Karamalis et al. propose a hybrid method combining the advantages of both approaches into a fast GPU reconstruction scheme [76].

Finally, *function-based methods* estimate the coefficients for a set of locally supported basis functions to approximate the input data. These functions are then evaluated on the voxel grid to reconstruct the compounded volume [82, 149, 156]. Function-based techniques can further be extended to more than three dimensions to reconstruct time-varying data such as 3D velocity fields with a flow profile from Doppler ultrasound [207]. While these methods yield 3D ultrasound reconstructions of very high quality, they are currently not feasible for interactive clinical practice due to their high computational costs. In terms of reconstruction quality, recently introduced tensor-based approaches yield results comparable to traditional spline-based methods, while at the same time having a much lower computational complexity [118].




Fig. 7.2. Illustration of function-based methods for ultrasound compounding, which estimate the coefficients for a set of locally supported basis function to approximate the input data. Image from [170].

When browsing the scientific literature on this topic, one realizes that most works on ultrasound compounding make simplified assumptions on the input data and presume constraints such as constant probe pressure and/or constant motion of the ultrasound transducer along a linear path. However, since the ultrasound transducer has to be always in contact with the skin surface, these assumptions do not hold for applications where the anatomy exposes high curvatures, such as MSK. Here, the compounding algorithms do not only need to deal with the inherent probe pressure changes, but may also have inconsistent intensity information for the same reconstructed voxel due to overlapping image frames. Such overlapping frames pose a particular challenge, since ultrasound may yield different information (i.e. image intensities) for the same point within the anatomy if scanned from different perspectives or at different times. This is due to the dynamics and high complexity of the ultrasound image formation being dependent on incident angle, probe pressure and patient positioning (cf. Section 5.1) [2]. Traditional, solely distance-based, compounding methods do not resolve such ambiguities so that their reconstructions may show image artifacts and low continuity of the image, as well as potentially wrong image intensities due to averaging on opposing information.

To the best of our knowledge, there has been very little work on explicitly resolving intensity ambiguities occurring at overlapping frames. The only generic method compensating for probe pressure changes is the work of Treece et al. who use an image-based non-rigid registration technique [184]. By computing the line-wise maximum normalized correlation between two adjacent B-mode images and applying a monotonicity constraint, they estimate the deformation in depth introduced by the probe pressure. To avoid drift in the registration, the registration results are regularized using the tracking data.

7.2 Correlation Terms for Tracked Ultrasound Sweeps

As observed by Housden et al., the correlation between two ultrasound frames depends on both their proximity and their orientation [70]. Therefore, we introduce two correlation terms that we will use in our ultrasound compounding pipeline.




Fig. 7.3. Illustration of the elevational decorrelation curve describing the distance between two adjacent ultrasound frames given their correlation ρ [1, 36, 70].

Technically, the elevational decorrelation curve describing the distance d of two frames given a correlation ρ is of non-linear form (cf. Fig. 7.3) and needs to be carefully calibrated to the used ultrasound transducer [1, 36, 70]. However, since we assume to have tracking information available, we do not need to solely rely on speckle decorrelation and can approximate the elevational speckle decorrelation curve by a linear function, which allows for a generic formulation. For two tracked ultrasound frames i, j with centroids c_i, c_j , we consider their correlation to be zero if their Euclidean distance exceeds a threshold d_{\max} , while their correlation should be maximal if $c_i = c_j$. Therefore, we define

$$D(i, j) := \max\left(\frac{d_{\max} - \|c_i - c_j\|}{d_{\max}}, 0\right), \quad (7.1)$$

which normalizes the correlation to the interval of $[0, 1]$. A value of 1mm for d_{\max} showed excellent results in all of our experiments. While yielding comparable results for standard cases, this distance-based correlation term is not only simpler than the original Gaussian window formulation in [161], but also a more robust generalization since it does also allow for non-homogeneous frame sampling rates.

We model the orientation-based correlation between two ultrasound frames by the cosine distance of their normals. Thus, for two tracked ultrasound frames i, j with normals n_i, n_j , we define the orientation-based correlation term as

$$O(i, j) := \max\left(1 - \frac{2}{\pi} \cdot \operatorname{acos}\left(\frac{n_i \cdot n_j}{\|n_i\| \|n_j\|}\right), 0\right), \quad (7.2)$$

which is a simple extension of the Euclidean dot product formula, normalized to yield results in the interval of $[0, 1]$. In both terms, the maximum function avoids negative correlations.

7.3 Orientation-driven Inter-frame Registration of Ultrasound Sweeps




Fig. 7.4. Reconstruction of an abdominal phantom scan with probe pressure changes: (a) MPR through the compounded volume without applying our pressure compensation technique; (b) the same MPR through the compounded volume with pressure compensation applied. For comparison, (c) shows the same target acquired using constant probe pressure.

To correct for errors and inaccuracies in the tracking data (e.g. due to jitter, inaccurate calibration or patient movement), we use the above framework for performing an intensity-based inter-frame registration, similar to Treece et al. [184]. Since the image changes are minimal for adjacent ultrasound frames in a sweep, a simple and thus real-time capable pixel-wise uphill search evaluating the sum of squared differences (SSD) is sufficient for registering each ultrasound frame to its surrounding frames in terms of in-plane translation and rotation. To effectively compensate for registration drift, we regularize by registering each ultrasound frame to a window of N surrounding frames. Thus, for a given reference patch P and equally sized moving patch P' the windowed SSD (wSSD) is given by

$$\text{wSSD}_{P,P',N}(i) = \sum_{\substack{(p,p') \\ \in P \times P'}} \sum_{n=-N}^N T(i, i+n) \cdot (I_i(p) - I_{i+n}(p'))^2, \quad (7.3)$$

where i is the index of the reference frame and $I_i(p)$ denotes the image intensity of ultrasound frame i at the position p .

The central part in Eq. 7.3 is the correlation term $T(i, j)$, which describes the correlation between frames i and j and weights the surrounding frames' contribution to the similarity measure. As motivated in the previous section, the correlation between the two frames depends on both their proximity and their relative orientation. Therefore, we define the $T(i, j)$ as

$$T(i, j) := D(i, j) \cdot O(i, j). \quad (7.4)$$

To compensate for probe pressure artifacts, we apply the above inter-frame registration technique not only to a single patch, but to a grid of independent patches. We make the simplified assumption that for our applications a deformation model is sufficient, which expects the deformation to be only in the direction of the transducer. Therefore, we allow in-plane translations of the individual patches. After computing the transformation for each patch as above, we regularize the deformation field using a Gaussian convolution, which corresponds to one iteration of a fluid-type demons registration [19, 180, 209] on the above-mentioned grid. We compute the coarse inverse deformation field using fixed-point iteration [37], which is eventually extended to pixel level using bilinear filtering during the compounding process. An exemplary result is displayed in Fig. 7.4.

7.4 An Information-fusion Approach to Ultrasound Compounding

As previously discussed in Section 7.1, compounding ultrasound frames acquired from different viewpoints is not trivial since ultrasound may show different intensities for the same anatomy if it is scanned from a different angle or at a different time. Thus, tortuous acquisition sweeps with overlapping frames are particularly challenging to reconstruct. Traditional compounding schemes based on averaging or distance-based weighting may fail to correctly reconstruct such regions. Fig. 7.5 illustrates the reason for this: Since closer pixels are preferred over pixels being farther away, the compounded voxel intensity mainly depends on the ultrasound frame of highest proximity. If we now consider a neighbor voxel, the closest frame may have a completely different orientation and thus show different information (due to the view dependency of ultrasound). The resulting reconstructions exhibit a low image quality with significant discontinuities in the anatomy. Furthermore, unwanted stripe or pixel artifacts may arise as depicted in Fig. 7.6a.




Fig. 7.5. Illustration of artifacts occurring in distance-weighted compounded regions where multiple ultrasound frames from different directions intersect. With traditional methods, voxel v_1 is mainly based on the information of frame f_1 , while its neighbour voxel v_2 is mainly based on the information in frame f_3 . However, their intensities at this spatial location may significantly differ since the frames travel through different acoustic windows.

Besides promoting artifacts and inhomogenities, distance-based weighting can also lead to incorrect reconstruction results since the distance of the frame to the voxel has no correlation with the amount of information present in this pixel (i.e. level of uncertainty/noise). For instance, it may ignore a pixel being farther away but having low uncertainty and instead prefer a high uncertainty pixel (i.e. noise) because it is closer to the voxel.




Fig. 7.6. Effect of the orientation-driven clustering of ultrasound frames: Both images show reconstructions of a tortuous sweep of upper arm ultrasound. (a) Traditional techniques with distance-only weighting show severe stripe artifacts (yellow arrows) in regions of overlapping frames. (b) Partitioning the sweep into clusters of homogeneous orientation and fusing the clusters based on uncertainty suppresses these artifacts almost entirely. Please mind that both volumes were compounded *without* applying our pressure compensation technique to highlight the beneficial effect of the orientation-driven clustering.

These two issues are the main motivation for our proposed orientation-driven ultrasound compounding technique. It tackles them by partitioning the ultrasound sweep into clusters of similar alignment and then using additional uncertainty information in a two-step fusion approach. We assume that for each ultrasound pixel with intensity I_i , we also have an uncertainty value u_i that we will use for weighting the image intensities during compounding. While the method considers the generation of uncertainty information as black box and is thus independent from it, our implementation uses the method described in Chapter 5 of this thesis.

7.4.1 Clustering of Ultrasound Sweeps by Direction

As a first step, we perform a hierarchical clustering to identify tortuous sweep trajectories and regions of overlapping ultrasound frames. This partitions the ultrasound sweep trajectory into parts where the frames have homogeneous orientation without requiring us to predefine the number of clusters. We apply an average group linkage algorithm on the previously defined orientation term $O(i, j)$ (cf. Eq. 7.2). This yields a set C of sub-sweeps meeting the usual restriction of being contiguous and uniformly oriented.

7.4.2 Compounding of each Cluster

A backward warping algorithm then compounds each cluster $c \in C$ into a 3D volume I_c applying our pressure compensation method as discussed in Section 7.3. Since the ultrasound frames of each cluster are guaranteed to have the same orientation and are thus travelling through the same acoustic window, we can safely assume the uncertainty distribution to be homogeneous within nearby frames. We compute the intensity for voxel x as

$$I_c(x) = \frac{\sum_{i \in S} I_i \cdot d_i^{-\mu}}{\sum_{i \in S} d_i^{-\mu}}, \quad (7.5)$$

where S is the set of frame pixels close to the compounded voxel x , d_i the Euclidean distance of pixel i to the compounded voxel, and $\mu > 1$ a smoothness parameter ensuring that $I_c(x)$ approximates the original data for $d_i \rightarrow 0$ [168]. Wein et al. showed that this inverse distance weighting scheme yields excellent results in terms of image quality [199]. We use the same weighting to propagate the uncertainty to a separate 3D volume U_c :

$$U_c(x) = \frac{\sum_{i \in S} u_i \cdot d_i^{-\mu}}{\sum_{i \in S} d_i^{-\mu}}. \quad (7.6)$$

Thus, for each cluster $c \in C$ we end up with a compounded ultrasound intensity volume I_c and a compounded uncertainty volume U_c .

7.4.3 Uncertainty-based Fusion of the Compounded Clusters

In a final step, our method merges the clusters into the final 3D volume based on the propagated uncertainty values. We want the individual clusters $c \in C$ to contribute in an additive manner weighted by the propagated voxel uncertainties. Therefore, we compute the final intensity I at voxel x by

$$I(x) = \frac{\sum_{c \in C} (1 - U_c(x)) I_c(x)}{\sum_{c \in C} 1 - U_c(x)}. \quad (7.7)$$

Since we are acquiring a continuous sweep of ultrasound frames and apply a regularized inter-frame registration with active pressure compensation, we can safely assume that the individual clusters are well aligned. Hence, an additional 3D-3D registration is not needed.

7.5 Incremental Compounding System

For interventional usage, one does not only require short processing times, but also wants to support an interactive update of the compounded volume during the acquisition to allow for an additional refinement of selected regions. Therefore, we introduce *incremental ultrasound compounding* as a real-time extension of our orientation-driven compounding technique that can be integrated into the interactive ultrasound processing pipeline.

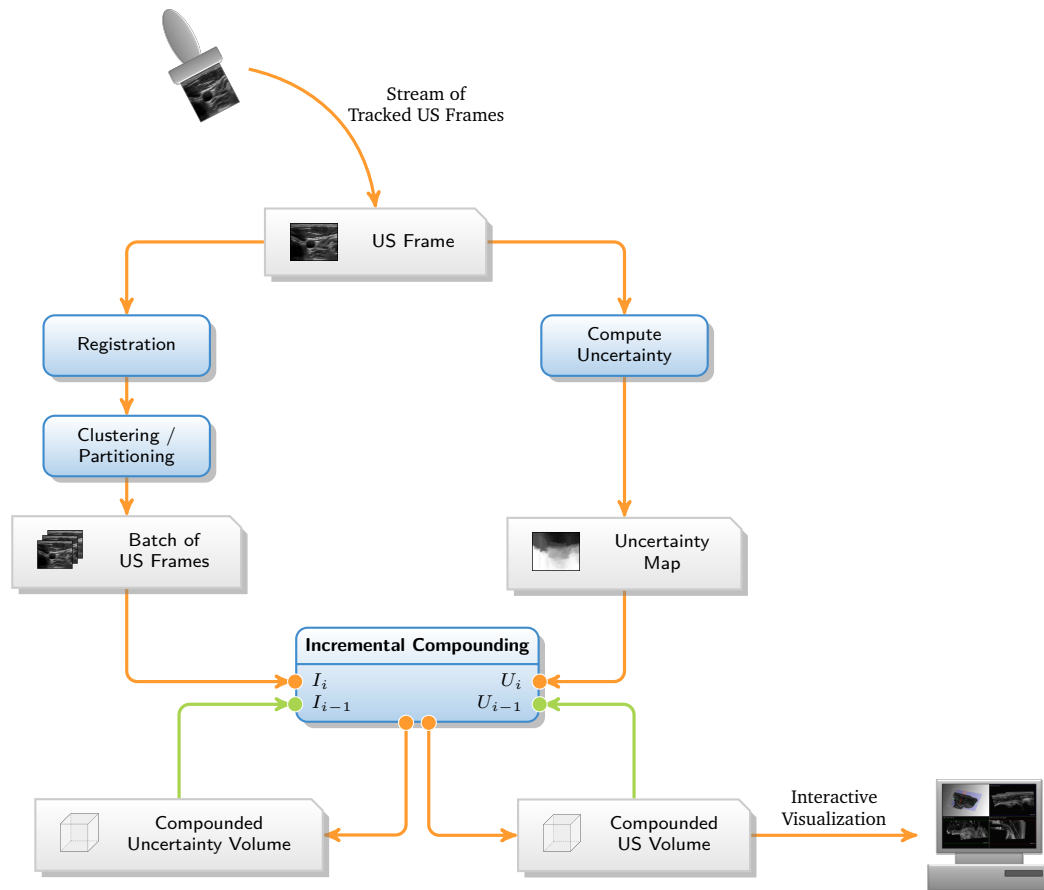


Fig. 7.7. Schematic diagram of the proposed incremental compounding technique. We acquire a constant stream of tracked ultrasound frames and compute the per-pixel uncertainty for each frame. We furthermore perform the online inter-frame registration and partition the stream into small batches of homogeneous orientation. Our incremental compounding then continuously compounds the batches into the accumulated ultrasound intensity and uncertainty volume. After each incremental compounding step the interactive visualization can be updated.

7.5.1 Incremental Compounding Pipeline

As illustrated in Fig. 7.7, our system receives a constant stream of incoming ultrasound frames. The regularized inter-frame registration (cf. Section 7.3) needs only a limited number of frames lookahead (i.e. size of the regularization window) and can hence be performed online requiring only a small number of frames to be buffered. We extend the orientation-driven clustering approach (cf. Section 7.4) to a partitioning into batches of ultrasound frames. In addition to starting new clusters when the correlation term exceeds the threshold, we also start new batches in a regular time interval (e.g. twice every second) to keep their size small and the updates interactive.

Instead of reconstructing a separate volume for each batch, we use a single volume as accumulation buffer and adapt the above technique to an in-place algorithm. The reconstructed voxels of each cluster can be incrementally added to an accumulation buffer by rewriting Eq. 7.7 to a recurrence scheme. Given the accumulated compounded intensity I_{i-1} and


accumulated compounded uncertainty U_{i-1} of the previous runs, and I_c, U_c of the current run, we define the new intensity I_i and uncertainty U_i as

$$I_i = \frac{U_{i-1}I_{i-1} + (1 - U_c)I_c}{U_{i-1} + (1 - U_c)}, \quad (7.8)$$

$$U_i = U_{i-1} + (1 - U_c).$$

This transforms our technique into an online algorithm, significantly lowering the computational complexity and memory footprint and furthermore allowing for an interactive update of the visualization after each batch.

Apart from a worse but in our case neglectable numerical stability, Eq. 7.7 and Eq. 7.8 yield the same reconstruction result. It should be noted that, with the incremental clustering, frames of the same orientation may end up in different clusters compared to the conventional formulation. While technically this may alter the final result, we did not observe any qualitative differences in our experiments.



(a) No pressure compensation applied. **(b)** Inter-frame registration with 3×4 patches yielding a patch size of $10\text{mm} \times 10.5\text{mm}$. **(c)** Inter-frame registration with 4×20 patches yielding a patch size of $7.5\text{mm} \times 2.1\text{mm}$.

Fig. 7.8. Multi-planar reconstructions of in-vivo acquisitions of the Infraspinatus muscle showing the effectiveness of the pressure compensation. Though the sweep was acquired by a professional sonographer trying to maintain constant probe pressure, the result of a traditional compounding technique (a) has an unsteady appearance and shows discontinuities throughout the different muscle layers. Applying our orientation-driven inter-frame registration and pressure compensation even on low resolution grid with $10\text{mm} \times 10.5\text{mm}$ patch size (b) yields a significant increase in image quality. Further increasing the resolution of the deformation field to $7.5\text{mm} \times 2.1\text{mm}$ patch size (c) yields only slight improvements.

7.5.2 Real-time Implementation

The most time-consuming part of our reconstruction pipeline is the incremental compounding of each batch of ultrasound frames into the target volume. To yield high quality results, we employ a backward compounding technique where each voxel of the target volume is compared to each ultrasound frame to find contributing pixels (cf. Section 7.4.2). This process can be parallelized heavily on voxel level. We implemented the entire incremental compounding process on the GPU using OpenGL 4.3. To this end, we store the tracking information as well as other required data in Shader Storage Buffer Objects [12] and the image data of the ultrasound frames in 2D texture arrays to support hardware-accelerated interpolation during sampling. This allows us to yield highly interactive frame rates for the compounding process. This is further accelerated by exploiting that we do not need to traverse the voxels of the entire target volume but can restrict the reconstruction volume of each

pass to the bounding box of the corresponding batch of ultrasound frames. Our information fusion approach using uncertainty information ensures that the additional partitioning of the ultrasound sweep into batches does not impair the reconstruction quality.

7.6 Results and Evaluation

To evaluate our methods, we used an Acuson S2000 ultrasound machine equipped with an Acuson 9L4 linear transducer and Ascension trakSTAR 2 electromagnetic tracking hardware. The system was calibrated using the method described in [198], yielding an experimental calibration error of less than 2mm and 4 degrees. We acquired both phantom and in-vivo data of different anatomies where the in-vivo sweeps were acquired by a professional sonographer. All baseline results for comparison within this section were computed using our own implementation of a state-of-the-art backward compounding technique with inverse distance weighting as proposed by Wein et al. [199].

7.6.1 Parameter Evaluation

To evaluate the effectiveness of our orientation-driven inter-frame registration and pressure compensation, and its dependency on parameters, we applied our technique to the acquired ultrasound sweeps. Figures 7.4 and 7.8 show representative results.

All experiments show a significant improvement in terms of image appearance and continuity of anatomy features when applying inter-frame registration, even with only minimal pressure changes present in the input data. In particular the in-vivo scans show good results already with low resolution deformation fields, such as 3×4 patches for the Infraspinatus data set resulting in a patch size of roughly $10\text{mm} \times 10\text{mm}$. Increasing the deformation field resolution by decreasing the patch size yields only marginal improvements while increasing the computational burden (cf Fig. 7.8). Hence, in our implementation, we use a target patch size of $10\text{mm} \times 10\text{mm}$ as default setting.

7.6.2 Reconstruction Accuracy

To validate the physically correct reconstruction of anatomy, we acquired ultrasound sweeps of an abdominal phantom including a tumor target of spherical shape as depicted in Fig. 7.4. Since the target is positioned relatively close to the surface, it can be scanned from different directions and is prone to deformation. Therefore, this is a valid scenario to evaluate our method. We computed 50 MPRs of arbitrary orientation through the target and compared the maximum diameter with reference measurements acquired from CT. The reconstructed ultrasound volume yielded an average target diameter of 14.63 ± 0.48 mm compared to 14.5 ± 0.84 mm in CT. This proves an excellent physical accuracy of our reconstruction method.

7.6.3 Reconstruction Quality

The qualitative effects of our inter-frame registration and pressure compensation technique can be observed in Fig. 7.4 and 7.8 showing the reconstructions of different ultrasound sweeps. The originally round target of the abdominal phantom (Fig. 7.4c) is severely deformed through probe pressure changes (Fig. 7.4a). Our pressure compensation technique is capable of mostly restoring the original anatomy (Fig. 7.4b). While the original shape is not perfectly restored, there is a clear improvement visible. Also Fig. 7.8 shows significant improvements in terms of continuity of anatomy when enabling our pressure compensation. Compared to [184], there is still room for improvement. However, we want to emphasize that the implemented pressure compensation is of only complementary nature and the focus of this work is the orientation-driven clustering and compounding technique.

Fig. 7.6 demonstrates the effect of our clustering technique when reconstructing a twisted ultrasound sweep of a human arm exhibiting out-of-plane rotations of up to 35 degrees. Due to overlapping frames, the standard compounding in (a) shows artifacts at locations where the frames for neighboring voxels are acquired through different acoustic windows. The reconstruction in (b) removes such artifacts by using our clustering technique to avoid overlapping frames in a single cluster. Additionally, it exploits uncertainty information when fusing the clusters so that unreliable intensities do not bias the final result. Please note that, in order to highlight the effect of the clustering, these volumes were reconstructed without pressure compensation and thus show discontinuities in the anatomy.

To quantitatively evaluate the effectiveness of our orientation-driven compounding technique in terms of ensuring continuity of the anatomy and reducing artifacts in regions of overlapping ultrasound frames, we evaluate a local, scale-invariant entropy measure of the reconstructed sweeps. For every voxel we computed the intensity difference to the neighboring voxels and computed the mean and standard deviation over the entire volume. This measure penalizes high frequency intensity changes such as the mentioned stripe artifacts while preferring smooth and continuous reconstructions of the anatomy, which are desired for clinical and visualization purposes. To avoid biasing the result, we considered only the inner part of the volume and masked the outer faces of the sweep, as they naturally have a high entropy. The results as displayed in Table 7.1 show that our technique is successful in reducing the amount of reconstruction artifacts. While the entropy measures are not directly comparable between

	Baseline [199]	Our Technique
In-vivo leg / linear sweep	2.65 ± 2.42	2.49 ± 2.39
In-vivo arm / twisted sweep	4.59 ± 8.42	2.44 ± 4.62
In-vivo leg / twisted sweep	3.58 ± 5.82	2.34 ± 4.19
In-vivo carotid artery / twisted	7.2 ± 13.35	3.17 ± 8.58
In-vivo infrapinatus / twisted	4.17 ± 7.21	2.99 ± 6.17

Tab. 7.1. Quantification of the local entropy in compounded volumes in terms of mean intensity difference to neighboring voxels.

the data sets, all test data sets with a twisted probe trajectory show a significant drop in the mean difference and standard deviation. As reference, the first line shows the measure for a data set with a strictly linear sweep trajectory. Here, the drop in entropy turns out to be much lower.

For further quantitative evaluation we acquired pairs of overlapping sweeps with perpendicular main trajectory. After compounding the sweeps into separate 3D volumes using our methods, we applied a 3D-3D rigid registration using the tracking data as initialization. Expecting our techniques to yield better matching volumes, we compared the differences of the two volumes in the overlapping region. With the average of the two volumes as expected result for a correct reconstruction, we quantify their difference in Normalized Cross-Correlation (NCC) and log-scale Signal to Noise Ratio (SNR_{dB}), for which define the signal as average of the volumes and the noise as root mean square of the differences (Table 7.2).

	Baseline [199]		Our technique	
	NCC	SNR_{dB}	NCC	SNR_{dB}
Phantom / constant pressure	0.90	19.39	0.94	23.16
Phantom / pressure changes	0.81	13.02	0.94	22.47
In-vivo leg / constant changes	0.72	9.21	0.76	11.69
In-vivo leg / pressure changes	0.67	8.53	0.75	11.03

Tab. 7.2. NCC and log-scale SNR in the overlapping region after registering the two compounded volumes of two sweeps with perpendicular trajectories of the same anatomy.

The sweeps with pressure changes show a significant improvement in terms of increase in both NCC and SNR_{dB} when our technique is applied. Furthermore, when comparing constant pressure with pressure changes, our technique shows significantly less drop of the measures. The slight improvements for the sweeps acquired with constant pressure are mainly due to the inter-frame registration correcting for the tracking error. Since the sweeps are acquired with perpendicular trajectories and the volumes therefore show different interpretations of the underlying data, no algorithm yields a perfect match. Furthermore, the in-vivo sweeps are expected to have lower similarity since they show by far less homogeneous anatomy. Figure 7.9 shows the difference images for the second phantom sweep.




Fig. 7.9. Illustration of our evaluation method: First two images show the MPRs for each perpendicular sweep, the third one shows the color-coded absolute difference of their intensities after 3D-3D rigid registration. (a) traditional backward-compounding [199] fails to align the different structures; (b) our technique yields alignment of all structures. The quantitative results are shown in Table 7.2.

7.6.4 System Performance

For the evaluation of the system performance, we ran our implementation on a workstation with an Intel Core™ i7-4770 CPU and a nVidia GeForce™ GTX 670 GPU. We loaded recordings of different ultrasound acquisitions, streamed them into our incremental compounding pipeline and measured the time required for computing the final reconstructed volume at different resolutions.

Voxel Size	Volume Size	Time	Speed
In-vivo arm			
0.5 mm	160 × 119 × 157	728 ms	687 fps
0.25 mm	320 × 237 × 313	1551 ms	322 fps
0.1 mm	800 × 591 × 781	13499 ms	37 fps
In-vivo leg			
0.5 mm	187 × 138 × 101	389 ms	620 fps
0.25 mm	373 × 276 × 214	906 ms	266 fps
0.1 mm	931 × 689 × 533	8433 ms	29 fps
In-vivo carotid artery			
0.5 mm	90 × 62 × 97	657 ms	743 fps
0.25 mm	179 × 124 × 194	735 ms	664 fps
0.1 mm	447 × 308 × 485	2569 ms	189 fps
In-vivo infraspinatus			
0.5 mm	114 × 104 × 178	609 ms	608 fps
0.25 mm	228 × 207 × 355	1135 ms	326 fps
0.1 mm	570 × 517 × 887	6761 ms	55 fps

Tab. 7.3. System performance of our incremental compounding technique at different target volume resolutions. The timings measure the compounding only (no generation of uncertainty information or visualization).

Table 7.3 shows the timings for four exemplary data sets. The performance mainly depends on the number of frame pixels contributing to a single voxel (i.e. relationship between pixel size, voxel size and frame sampling rate) as well as the number of identified clusters (i.e. tortuosity of the sweep). The results show that for a clinically adequate resolution of 0.25 mm voxel size, our system has enough bandwidth to reconstruct ultrasound sweeps of very high framerate. Even high-quality reconstructions of 0.1 mm voxel size can be compounded at interactive frame rates.

7.7 Conclusion

In this chapter, we presented a novel orientation-driven approach to ultrasound compounding featuring an incremental computation scheme. We developed our methods with the aim to free clinicians from the restrictive scanning protocols of many of today's state-of-the-art methods,

which require the acquisition of linear probe trajectories and constant probe pressure. Our target application was MSK ultrasound, where typical sweeps exhibit pressure changes, back-and-forth and twisting motions. Therefore, we introduced a framework of two correlation terms to capture that the correlation of two frames in an ultrasound sweep depends on their proximity and orientation to each other. This framework is used to perform a complementary pressure compensation as well as a clustering of ultrasound frames by orientation to remove different kinds of artifacts. In the subsequent compounding process, we use an information fusion approach exploiting ultrasound Confidence Maps (cf. Chapter 5). This allows for more accurate reconstructions in regions where we have information from different acoustic windows, since intensities from uncertain regions do not bias intensity information from reliable regions. Finally, we presented an integrated system for incremental compounding that interactively fills the reconstructed volume while acquiring the sweep. This provides the clinician with direct feedback on the ongoing acquisition and allows for an interactive refinement of target regions.

We evaluated the different aspects of our technique using both phantom and in-vivo data, which was acquired by a professional sonographer. The results show that our methods accurately reconstruct the original anatomy and yield superior results than traditional backward compounding methods in terms of image quality, continuity of the anatomy and presence of artifacts. Furthermore, we demonstrated that our implementation of the proposed incremental compounding system has enough bandwidth to reconstruct ultrasound sweeps with more than a hundred frames per second when using a clinically adequate target voxel size of 0.25mm. Even with high resolution volumes of 0.1mm voxel size, we still can yield interactive frame rates to update the compounded volume during the acquisition. Thus, our method can be integrated into the real-time ultrasound processing pipeline.

Predicate-based 3D Ultrasound Visualization

” pre·di·cate, *noun*

1. a) *something that is affirmed or denied of the subject in a proposition in logic*
b) *a term designating a property or relation*
2. *the part of a sentence or clause that expresses what is said of the subject and that usually consists of a verb with or without objects, complements, or adverbial modifiers*

— Merriam-Webster’s English Dictionary

High-quality volumes from tracked B-mode ultrasound sequences, as generated in the previous chapter, need to be presented adequately to the clinician. As discussed in Section 3.2, direct volume visualization techniques, such as ray casting, are today’s state-of-the-art algorithms for the visualization of three-dimensional medical images. However, the quality of their results depends heavily on the used classification scheme. Global, intensity-based 1D transfer functions are their simplest form and usually sufficient for images of high signal-to-noise ratio, such as CT reconstructions. However, with decreasing signal-to-noise ratio, image intensities become a less reliable basis for classification and in particular for 3D ultrasound they often fail to yield meaningful visualizations. Advanced, multi-dimensional transfer functions are capable of providing a more meaningful classification but most users have difficulties handling the increased complexity of the parameter domain [142]. Thus, effective and generic visualization tools are still missing for such data. This may be one of the main reasons why volumetric visualization still has not been fully accepted by most clinicians for their workflow and is, if at all, only used for publications or presentations to patients. Even with CT imaging, where the image intensities (Hounsfield units, cf. Section 3.1) directly correspond to a physical property that allows for direct discrimination of tissue types, many radiologists still prefer to scroll through the stack of 2D slices and make up the 3D model in their minds instead of looking at 3D visualizations.

Thus, in order to better integrate volume rendering into the clinical workflow, it requires for both powerful classification tools and simple and intuitive user interfaces at the same time. Striving for a unified formulation, integrating the evaluation of both local and global features into the classification step, we introduce the generic concept of point predicates for direct volume rendering. These binary-valued functions are defined on every sampling point within the image domain and allow also for an evaluation of high-level non-local information such as texture or even anatomical features and segmentations. We further propose the

predicate histogram as intuitive user interface, which reduces the size and dimensionality of the parameter domain and facilitates its manipulation. Together with the descriptive semantics of predicates, it provides the user with an easy and intuitive widget to setup a relevance-based focus-and-context visualization, ensuring that important anatomies are always visible in the rendering, while context information is preserved in order to provide spatial cues. The final element of our work is a scribble technique, where the user can define focus and context regions directly in image-space.

Since the above-mentioned challenges play a particularly strong part in 3D B-mode ultrasound visualization, we designed our methods in context of this modality and evaluate our methods on various ultrasound volumes of vascular and musculoskeletal applications. Nevertheless, the generic formulation of our technique makes it applicable to any kind of volumetric data.

This work has originally been published in [162].

8.1 Background and Related Work

Our proposed method has three important properties:

1. It is particularly designed for the visualization of ultrasound volumes.
2. It integrates different levels of information, ranging from low-level local image intensities to even anatomical models, into a single consistent formulation that is exposed to the user using an intuitive widget.
3. It exploits illustrative focus-and-context rendering techniques.

In this section, we therefore focus on these aspects.

Challenges in Ultrasound Visualization

As already discussed in Section 3.3, B-mode ultrasound intensities have the special nature of rather showing the changes in physical properties than the physical properties themselves. As a consequence, ultrasound images are fundamentally different to those obtained from tomographic imaging modalities like CT and MR. In addition to this gradient-like nature, ultrasound is a highly directional modality and suffers from a considerable amount of noise. These properties prohibit the straight-forward application of standard classification techniques for direct volume visualization as illustrated in Figure 8.1. In this vascular ultrasound volume the clinician would be interested in examining the carotid artery, which however is hidden by surrounding tissue.

Thus, standard implementations for visualizing volumetric ultrasound use segmentation algorithms to extract the target anatomy, which is then rendered as geometry. However, integrating non-opaque geometry into direct volume rendering is not straight-forward. Furthermore, such renderings do no longer expose the original image data (i.e. ultrasound intensities) and in particular the speckle texture, which is an essential part of ultrasound imaging. Therefore, in lack of better alternatives, clinicians prefer to look at 2D slices extracted from 3D ultrasound volumes. Such multi-planar reconstructions (cf. Section 3.2) allow for the visualization of

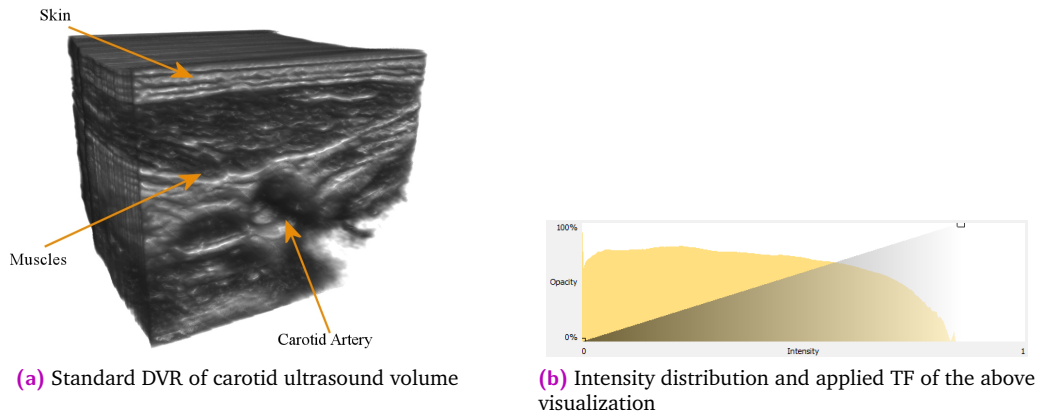


Fig. 8.1. Illustration of occlusion artifacts of when using traditional classification methods: (a) While one can see the outer shell of the different layers, clinical relevant information such as the path and shape of the carotid artery and its bifurcation are hidden. (b) Due to the homogeneous distribution of ultrasound intensities and their highly context-sensitive nature, there is no chance in discriminating relevant features through their intensities.

arbitrary planes, which is not directly possible with 2D B-mode ultrasound. However, MPRs are hardly able to expose spatial context and connectivity information to the clinician.

An overview of state-of-the-art visualization techniques for ultrasound data is presented in Section 3.3.

Focus-and-Context Visualization

One efficient way to tackle occlusion problems in volume visualization are illustrative importance-driven visualization techniques. They extend the classic direct volume rendering compositing scheme toward focus-and-context rendering that defines a focus region, which is shown as foreground, and context regions that are shown with less priority for better spatial understanding. Approaches not relying on transfer functions, such as the Importance-Based Accumulated Transparency Modulation introduced by Wan and Hansen [192], extend the Maximum Intensity Difference Accumulation (MIDA) compositing scheme to allow relevance-based visualization without the need to setup a transfer function for classification. The work of Bruckner et al. allows for context-preserving see-through rendering by evaluating local shading information with two global parameters, which however do not have a direct semantic meaning [26]. The ClearView technique of Krüger et al. exploits curvature information as well as distance metrics to determine sample importances and introduces different shading and compositing techniques to map relevance to optical properties [90]. De Moura Pinto and Freitas introduce a further importance-aware compositing scheme, which is mathematically motivated and justified and integrates very well with the standard direct volume rendering pipeline [119]. Kutter et al. present an integration of state-of-the-art relevance-based volume rendering techniques into a medical augmented reality framework for video see-through [91].

Our Approach

Our approach is inspired by line predicates used in flow visualization, where streamline tracking yields an extensive number of streamlines that represent the global connectivity of

the data very well, but at the same time it greatly suffers from occlusion of important features. Here, line predicates offer an effective technique to filter the flow field for certain features such as vortices or high-velocity jets [15, 86, 155]. However, since line predicates are applied to geometry representations and simply toggle streamline visibility, they are very limited to this specific application. Our approach can thus be seen as a generalization of this idea providing a consistent formulation for relevance-based rendering with particular focus and application on 3D ultrasound visualization.

8.2 Point Predicates as Generic Concept for Classification

Our predicate-based approach is designed to fully integrate into a standard volume rendering pipeline and consists of three steps, as illustrated in the schematic diagram of Figure 8.2. We define a point predicate P as a boolean-valued function f_P on the image domain X augmented with an importance factor κ_P and a color modulation δ_P :

$$P := (f_P : X \rightarrow \{\text{true}, \text{false}\}, \kappa_P, \delta_P). \quad (8.1)$$

After selecting the predicates to apply from the point predicate library, the user specifies the rendering outcome by configuring κ_P and δ_P . This process is supported by the predicate histogram (cf. Section 8.3), which we propose as a user interface for rendering setup. During the classification stage of the rendering, the ray caster applies each predicate to each sampling

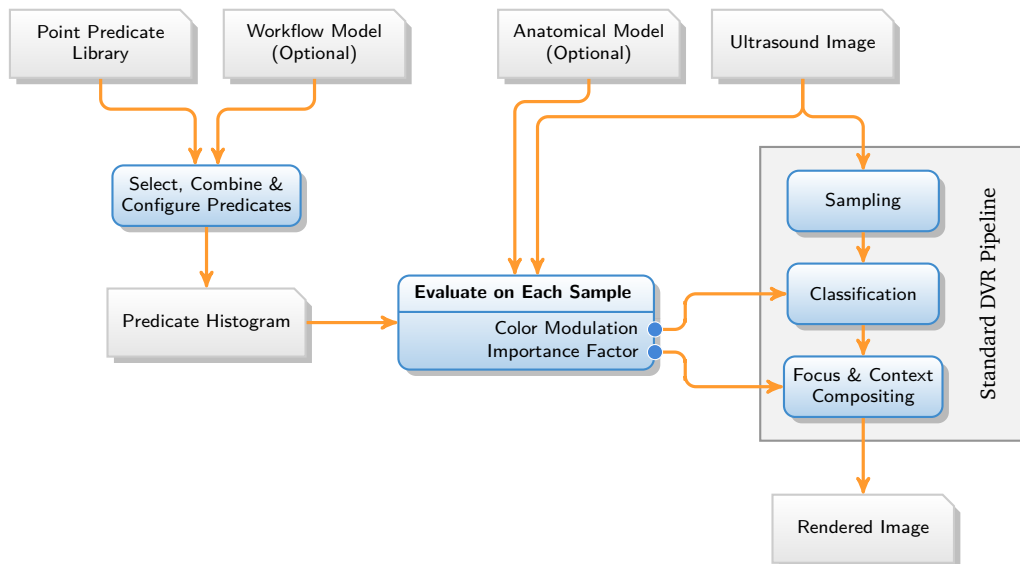


Fig. 8.2. Schematic diagram of the proposed predicate-based rendering technique: The user selects a set of descriptive point predicates from the library and configures their settings, importance factors and color modulation using the intuitive predicate histogram widget. These predicates are then evaluated on each sample during classification and eventually accumulated using relevance-based compositing exploiting the importance factors. While the predicates may evaluate both local and global features in the ultrasound image as well as external information (e.g. an anatomical model), the application-specific selection of the predicates can potentially be learned and supported with a workflow model.

point yielding a color modulation for each sample, which are eventually accumulated using a focus-and-context compositing technique exploiting the predicates' importance factors.

The term point predicate suggests that the predicate can be evaluated at every point within the image domain X . However, it is essential that predicates are not limited to local values but also to features of local environments or even of global nature, such as texture information or anatomical models and segmentations. Furthermore, we do not pose any presumptions on the spatial representation of X .

8.2.1 Predicate Library

Our current implementation consists of a point predicate library based on a variety of methods to evaluate both local and global features in the image.

Range-based predicates apply to intervals of scalar measures in the image such as intensity and gradient magnitude. While their sole expressiveness is rather limited, prone to inter-data set variability, and in particular does not go beyond traditional 1D/2D transfer functions, they are an essential part in combination with the other predicates. For instance, since high curvature regions often carry important visual depth and context cues, assigning high importance values to the gradient magnitude predicates may yield significant visual improvements.

Due to the direction dependency of ultrasound images, *direction-based predicates* are an important factor for our predicate-based 3D ultrasound visualization. They exploit additional information on the ultrasound scanning direction, which we annotate to the compounded ultrasound volume. The gradient angle predicate evaluates the scanline direction in a local context by computing the angle between scanline and the smoothed gradient. Since the reflection of the ultrasound wave depends on the incident angle with the interface, this ultrasound specific point predicate is a powerful tool to highlight or mask certain structures.

As a third group we implemented a variety of predicates based on derived measures. Since they apply a configurable threshold to the derived measure, they can be seen as extension to the range-based predicates.

The *Signal-to-Noise ratio predicate* computes the variance-based SNR of the ultrasound image in a local neighborhood of the sample position and can be used to select regions with certain entropy. We therefore build for each voxel the weighted variance within a uniform $7 \times 7 \times 7$ kernel and perform a windowing on the results. Since an elaborate parameter evaluation is beyond the scope of this work, our implementation leaves the windowing parameters user-controlled.

Furthermore, we implemented a GPU version of Frangi's vesselness filter [52] and integrated it into the *vesselness predicate*. By computing local Hessians for each voxel on multi-level Gaussians of the ultrasound volume and performing an eigenvalue analysis, the Frangi filter provides a scalar measure of presence of tubular structures. Being an integral part of the classification process, this predicate is particularly useful in vascular imaging.

The presence of ultrasound Confidence Maps [77] (cf. Chapter 5) can be exploited with the *confidence predicate*. Their per-pixel signal attenuation estimate can directly be integrated into the volume rendering pipeline. Our experiments however show only limited use of the additional information, which may be an issue with the Confidence Maps themselves. Nevertheless, we think that this is a promising direction for future work to integrate uncertainty visualization into our pipeline.

Finally, we implemented a *label predicate* that evaluates on optionally existing labeling information of the image in two different ways. For our evaluation in Section 8.6 we performed a manual segmentation on the data sets partitioning the ultrasound volumes into semantic layers of anatomy such as skin, muscle and bone. This was performed in the style of Surface Function Actives segmentation [45], which could yield a similar result in a (semi-)automatic fashion [78]. The segmentation layers were parameterized as B-spline surfaces and stored as set of control points in voxel space. With these, the label predicate can then directly determine to which layer the current sample belongs. The second version of the label predicate uses a 3D voxel grid as input data, where each voxel stores a bit mask with the labeling information. With the cost of discretization, this predicate allows the same evaluation of segmentation information but in a more generic way. This yields a very natural way to seamlessly integrate segmentations into the ray casting process and offers a powerful tool to accurately define importance factors or color modulations for specific regions.

As our work focuses on providing a novel approach together with a high-level tool, the presented point predicate library is certainly not complete, but rather provides a proof-of-concept implementation. Future work by the community may yield a variety of further application-dependent predicates that evaluate both local and global features in the image to yield even better classification results.

8.2.2 Predicate Combination

To allow sufficient generality, we keep the individual predicates simple with each just evaluating a single feature in the image domain. We achieve further flexibility by assembling multiple predicates P_1, \dots, P_n into combined *meta predicates*. The boolean formulation of the predicate function thus suggests to use the classic boolean operators for combination:

- NOT:** $\overline{P_1}$ yields the negation of the original predicate and allows to semantically complement the accepted domain of the original predicate.
- AND:** $P_1 \wedge \dots \wedge P_n$ yields a predicate that satisfies the constraints of all predicates P_1 through P_n . Often it is desired to have this operation in order to semantically restrict the accepted domain to the intersection of the accepted domains of each predicate, as also shown in Figure 8.3 where the second and the third predicate from the left are of this type.
- OR:** $P_1 \vee \dots \vee P_n$ yields a predicate that satisfies the constraints of either of the predicates P_1 through P_n . Since our formulation computes the final sample importance by computing a weighted sum of the predicates' importances (cf. Section 8.4), the OR operator can also be obtained by suitable setup of the importance factors.

8.3 The Predicate Histogram

The point predicates to apply are selected from the user or, alternatively, given by a workflow model later potentially learned specifically for the given anatomy and application. Each predicate has a small set of parameters, defining the predicate importance κ and an optional color modulation δ having components hue $\delta^{(H)}$ and saturation $\delta^{(S)}$. Certain predicates may have further optional parameters defining where the predicate yields true (e.g. intensity range).




Fig. 8.3. Our proposed widget shows the predicate histogram and allows an easy and intuitive configuration of the point predicates for relevance-based rendering: The main area shows the distribution of predicate importances and enables easy manipulation using drag and drop. Two sets of controls allow for further manipulation of the currently selected predicate: The sliders on the left set up the color modulation; the controls on the bottom set up optional predicate parameters.

We propose to combine predicate selection and configuration via the *predicate histogram*, a single intuitive widget as shown in Figure 8.3. The heights of the bars directly represent the relative importance weighting of the predicates, while their fill color shows the color modulation. The users can manipulate the importance of each predicate with intuitive drag and drop interaction on the histogram bar and will directly see the results of their actions, as the predicate-based rendering is evaluated in real-time. In order to further reduce the parameter space without losing flexibility, the predicate histogram normalizes the sum of the predicates weights:

$$\sum_j \kappa_j = 1. \quad (8.2)$$

Hence, increasing the importance of one predicate automatically decreases the importances of the other predicates and vice versa.

We place two additional sets of manipulators around the predicate histogram widget, which allow controlling additional parameters of the currently selected predicate: Two vertical sliders on the left allow the user to configure the color modulation in terms of hue and saturation. The resulting color is visualized in the predicate histogram as fill color of the corresponding bar. Furthermore, widgets to control feature specific predicate parameters (such as the range where the predicate yields true) are placed below the histogram in a horizontal layout.

The predicate selection process is currently a fully manual task of explorative nature. However, we argue that achieving satisfying rendering results is significantly more intuitive and easier than with traditional 1D/2D transfer functions, since the parameter space of our point predicate technique is much more expressive due to its semantic nature and with the whole process running in real-time, providing interactive feedback. Furthermore, the predicate histogram reduces the parameter space effectively in size through the normalization of predicate importances (cf. Equation (8.2)).

8.4 Predicate-based 3D Ultrasound Visualization

For classification and compositing, our predicate-based ray caster evaluates the selected predicates on each input sample point s to compute the sample importance $\kappa(s)$ and sample color modulation $\delta(s)$. As we wish the individual importances to contribute in an additive manner, we define $\kappa(s)$ as

$$\kappa(s) = \frac{\sum_{j=1}^n \chi_j(s) I(\kappa_j)}{\sum_{j=1}^n \chi_j(s)}, \quad (8.3)$$

where $I(\kappa_j)$ denotes the impact function for predicate j and $\chi_j(s)$ the characteristic function for predicate j (i.e. indicating whether $f_j(s)$ is true). We would like the following two conditions to hold:

1. A non-linear amplification of importance differences, meaning that $\kappa_i \sim \kappa_j \Rightarrow I(\kappa_i) \sim I(\kappa_j)$ but $\kappa_i < \kappa_j \Rightarrow I(\kappa_i) \ll I(\kappa_j)$. This allows for better usability since the user does not need to be pixel-perfect when assigning the same importance to multiple predicates. This condition can be rewritten to

$$\kappa_i - \kappa_j < I(\kappa_i) - I(\kappa_j), \quad (8.4)$$

which holds for all differentiable I with $I'(x) > 1$.

2. In the case of all predicates having the same importance factor assigned, we wish $\kappa(s)$ to be always 1:

$$\forall i, j : \kappa_i = \kappa_j \Rightarrow \forall s : \kappa(s) = 1. \quad (8.5)$$

Experimentally, we found $I(\kappa_j) := (n \cdot \kappa_j)^2$, where n is the total number of predicates, to yield good results and satisfy both conditions. This gives us the following full definition (cf. Equation (8.3)):

$$\kappa(s) = \frac{\sum_{j=1}^n \chi_j(s) \cdot (n \cdot \kappa_j)^2}{\sum_{j=1}^n \chi_j(s)}. \quad (8.6)$$

It should be noted that one can select very small κ_i, κ_j , so that $n \cdot \kappa_i < 0.5$ and $n \cdot \kappa_j < 0.5$ violating the first condition. However, in such cases there is at least one κ_k with $n \cdot \kappa_k > 1$ due to Equation (8.2). We argue that this effect is even desirable as predicate k shall have the major impact on the visualization and we can neglect the importance difference between predicates i and j .

In order to specify the optical properties of the sample s , we compute the sample color modulation in terms of hue $\delta^{(H)}(s)$ and saturation $\delta^{(S)}(s)$ in HSL color space. The sample

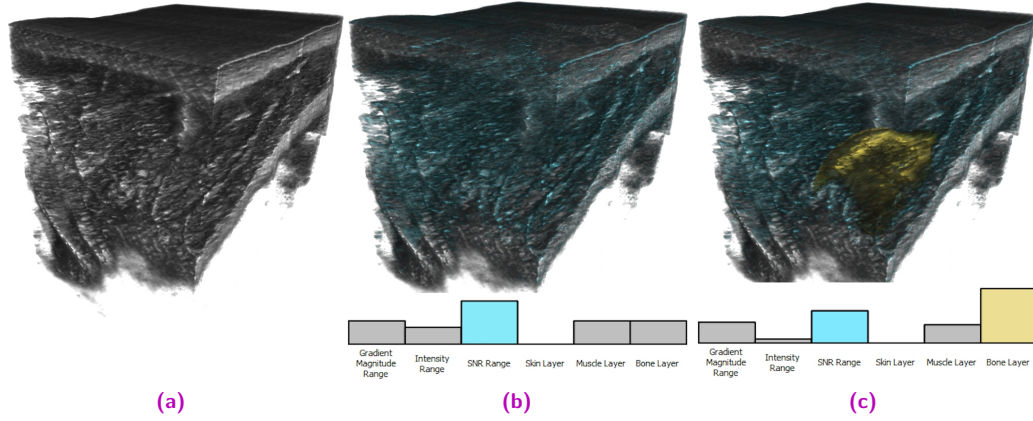


Fig. 8.4. Visualization results for the shoulder data set together with the used predicate histograms: Compared to the standard transfer function-based visualization (a), (b) was generated by exploiting the SNR Range predicate to mark fibrous structures in the muscle layer. The final rendering (c) additionally highlights the bone surface as focus region.

luminance as well as its opacity are set to the original ultrasound intensity. This serves two goals: It reduces the dimensionality of the parameter space, and even more importantly allows to retain the appearance of the original ultrasound data, which we see as an essential part of ultrasound visualization. The color modulation is computed in an additive manner using the predicate contributions

$$\omega_j(s) := \chi_j(s) \cdot (n \cdot \kappa_j)^2 \quad (8.7)$$

as weights. The hue modulation is additionally weighted by the saturation modulation to avoid undesirable effects in cases where large hue shifts with small saturation modulations are combined with small hue shifts with large saturation modulations.

$$\begin{aligned} \delta^{(S)}(s) &= \frac{1}{\sum_{j=1}^n \omega_j(s)} \cdot \sum_{j=1}^n \omega_j(s) \cdot \delta_j^{(S)}, \\ \delta^{(H)}(s) &= \frac{1}{\sum_{j=1}^n \omega_j(s) \delta_j^{(S)}} \cdot \sum_{j=1}^n \omega_j(s) \delta_j^{(S)} \cdot \delta_j^{(H)}. \end{aligned} \quad (8.8)$$

With sample importance, color and opacity set, the samples are passed to a focus and context compositing scheme. While our proposed method should work with most relevance-based compositing techniques, we use the one proposed by De Moura Pinto and Freitas [119], which has various desirable properties:

1. Its well motivated front-to-back recurrence scheme formulation allows straight-forward integration into standard ray casters.
2. In the absence of importance values or with all samples being equally important, it resolves to the standard emission-absorption-based volume rendering integral.
3. Its input parameters are clear and limited to a scalar importance value besides the standard optical properties of color and opacity. We therefore map $\kappa(s)$ to I_s , $\delta(s)$ to C_s and $I(s)$ to α_s as introduced in [119].

8.4.1 Integration into the Standard DVR Pipeline

Integrating the presented predicate-based classification and compositing into a standard GLSL ray caster is straight-forward and needs only minimal changes in the shader code: Instead of a transfer function texture lookup, the shader gathers all necessary input data for the predicates and calls the predicate evaluation function yielding the final sample importance and color modulation. The color modulation is then applied to sample intensity in HSL space, yielding the sample color value and opacity. Apart from setting up the predicate logic, all evaluations and computations are performed on the GPU using a single fragment shader.

The client-side selection and setup is automatically transformed into shader code, for which we use a building block like system similar to the shader templates in [140]. Every point predicate provides three blocks:

1. A *GLSL header* defining uniform names for the predicate parameters and further optional definitions.
2. A *closed GLSL expression* defining how to evaluate the predicate with respect to the input data.
3. A *function to setup the shader*, which is called by the renderer (prior to rendering) and assigns the parameter values to the corresponding uniforms.

The classification scheme from Equations (8.6) and (8.8) is transformed into a single GLSL function combining the evaluation expressions and yielding the final sample intensity and color modulation. Together with the predicates' GLSL headers, the predicate evaluation function is injected into the renderers fragment shader. Meta predicates are easily implemented using a composite pattern.

Potentially expensive recompilation of the shader has to be performed only when the set of selected point predicates changes, as this is the single event changing the definition of uniform variables and predicate expressions. Subsequent rendering setup by configuring the predicates' importance factors, color modulations, and optional parameters is entirely implemented by updating the uniform values, which is even less overhead than an update of the transfer function texture as performed by traditional transfer function-based renderers.

Together with the sample importance they are directly passed to the compositing scheme as follows (the equations are the same as in [119] but adapted to our naming conventions): Let $\kappa(s)$ be the computed sample importance, $C(s)$ be the sample color and $\alpha(s)$ be the sample opacity as we defined in Section 8.4. Then we define the visibility factor $\text{vis}(\kappa(s), \kappa_i)$ of the current sample $\kappa(s)$ compared to the accumulated importance κ_i as

$$\text{vis}(\kappa(s), \kappa_i) = 1 - e^{-\kappa_i \kappa(s)}, \quad (8.9)$$

as well as the modulation factor m as

$$m = \begin{cases} 1 & \kappa(s) \leq \kappa_i, \\ 1 & 1 - \alpha_i \geq \text{vis}(\kappa(s), \kappa_i), \\ \frac{1 - \text{vis}(\kappa(s), \kappa_i)}{\alpha_i} & \text{else.} \end{cases} \quad (8.10)$$

The incremental front-to-back compositing scheme is defined by:

$$\begin{aligned}
C'_{i+1} &= mC_i + (1 - m\alpha_i)C(s), \\
\alpha'_{i+1} &= m\alpha_i(1 - \alpha(s)) + \alpha(s), \\
\alpha_{i+1} &= \alpha_i(1 - \alpha(s)) + \alpha(s), \\
C_{i+1} &= \begin{cases} 0 & \alpha'_{i+1} = 0, \\ \frac{\alpha_{i+1}C'_{i+1}}{\alpha'_{i+1}} & \text{else,} \end{cases} \\
\kappa_{i+1} &= \max\left(\kappa_i, \ln\left(\alpha(s) + (1 - \alpha(s))e^{\kappa_i - \kappa(s)}\right) + \kappa(s)\right).
\end{aligned} \tag{8.11}$$

8.5 Scribble-based Predicate Histogram Setup

As demonstrated in the previous section, our technique integrates very well into the existing volume rendering pipeline. This is also valid for a variety of possible extensions that current state-of-the-art methods provide to facilitate standard transfer function-based classification, such as data driven techniques to generate transfer function presets, shape-based transfer functions, or image-driven transfer function setup where the user can manipulate the classification by direct interaction with the original image. Many of these works can easily be transferred to the presented predicate-based approach for direct volume rendering. As an example, we demonstrate how to integrate a scribble technique into our framework similar to the interface proposed by Tzeng et al. [190] and the stroke-based transfer function design by Ropinski et al. [150].

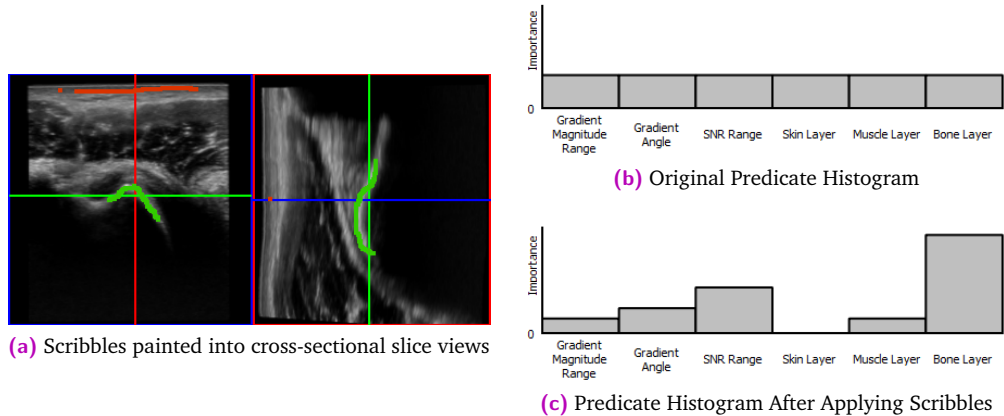


Fig. 8.5. Illustration of scribble-based predicate histogram setup: (a) By delineating image features to show (green scribble along bone structure) and image features to hide (red scribble in skin layer) in the cross-sectional MPRs, the user can configure the predicate histogram directly in the image domain. (b) shows the original predicate histogram. (c) shows the updated predicate histogram after applying the scribbles.

While the predicate histogram widget already provides an intuitive metaphor to setup the predicate configuration, it still requires the user to perform an implicit mapping between the semantic meaning of each predicate and its corresponding features in the image. To further facilitate the predicate importance setup, we implement a scribble technique that offers the user the possibility to directly interact with the ultrasound image by using a painting metaphor

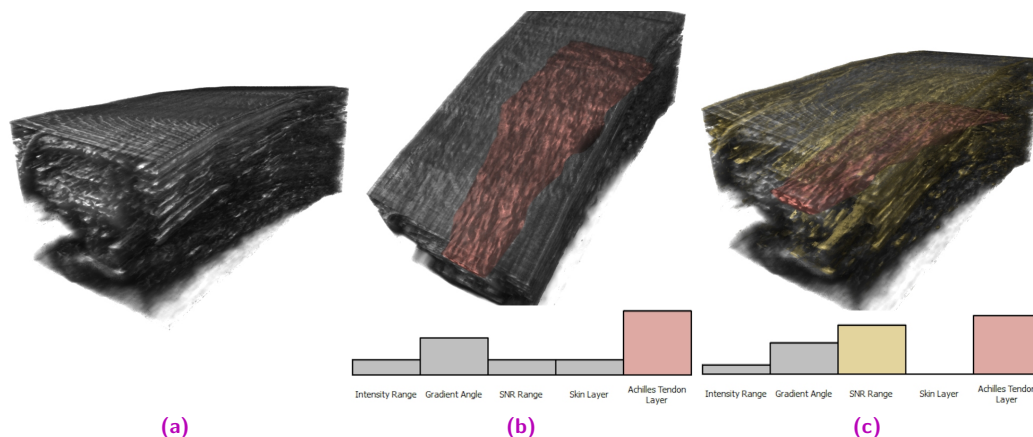


Fig. 8.6. Visualization results for the achilles tendon data set together with the used predicate histograms: Compared to the standard transfer function-based visualization (a), (b) highlights the target anatomy (achilles tendon) through the gradient angle predicate and the corresponding label predicate. (c) Shows the final rendering with highlighted fibrous muscle structures as context information using the SNR range predicate.

to mark focus and context regions in cross-sectional slice views of the volume as illustrated in Figure 8.5.

The scribbles yield a set of positive samples that should be emphasized in the volume rendering and a set of negative samples that should be less prominent. Since emphasizing certain regions relates to increasing the importances of the corresponding predicates and decreasing importances results in decreased visibility, we transform the scribbles into a modulation of the predicate importances κ_i .

For each predicate i we denote the number of positively labeled samples where the predicate yields true by N_i^+ and the number of negatively labeled samples where the predicate yields true by N_i^- . We compute an importance modulation vector m as

$$m_i = \frac{q(N_i^+ - N_i^-)}{\sum_j N_j^+ - N_j^-}, \quad (8.12)$$

where q describes the percentage by which the current predicate histogram is altered through the scribbles, which we empirically set to 0.25. In order to keep the predicate histogram in its normalized state (Equation (8.2)), we first normalize m so that $\sum_j m_j = 0$ before we update each importance κ_i by adding m_i .

8.6 Results and Evaluation

For evaluation we acquired ultrasound sweeps of three distinct anatomies (shoulder, carotid artery, and achilles tendon) using an ACUSON S2000 ultrasound machine equipped with an Acuson 9L4 linear transducer and electromagnetic tracking hardware. These sweeps were reconstructed using the orientation-driven compounding algorithm as presented in Chapter 7.

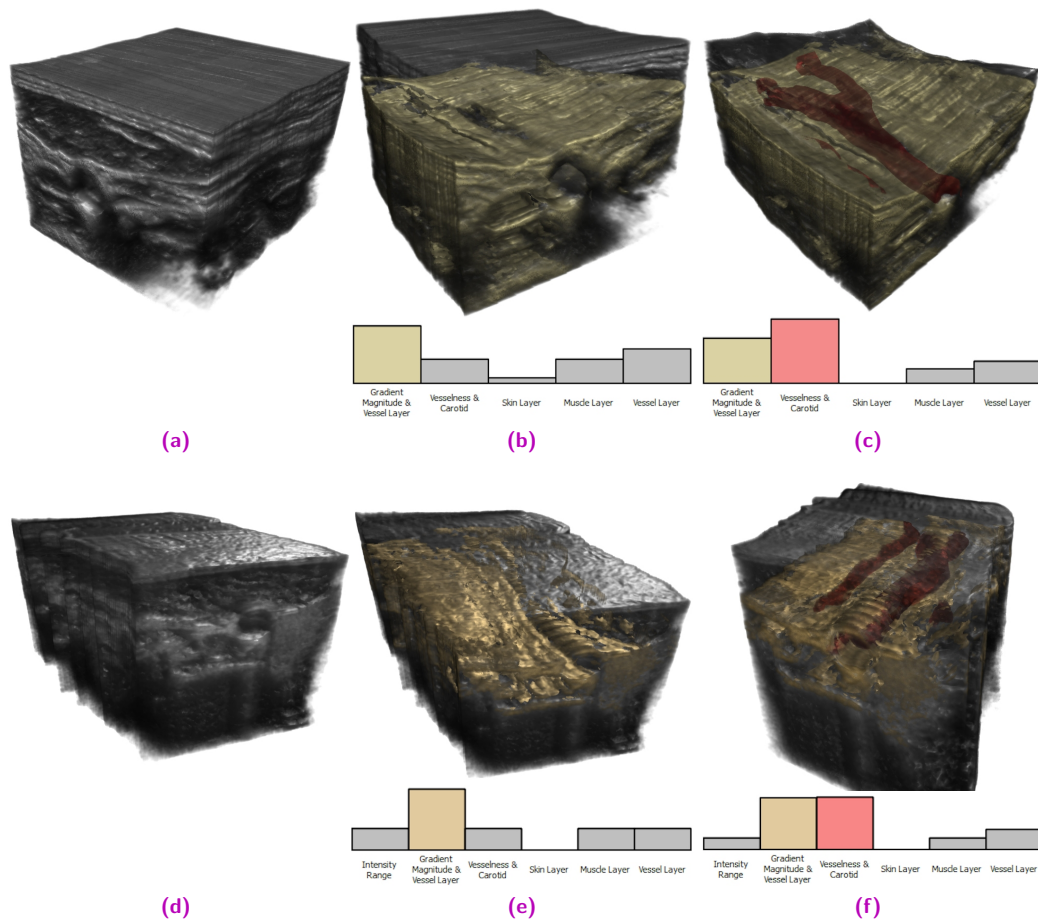


Fig. 8.7. Visualization results for the two carotid artery data sets together with the used predicate histograms: Compared to the standard transfer function-based visualization (a) and (d), the predicate setup in (b) and (e) removes the skin layer and highlights vascular structures. The vesselness predicate in (c) and (f) allows to further show the vessel path.

The resulting volumes have a resolution of 512^3 voxels for the shoulder data set and of 384^3 voxels for the carotid and the achilles tendon data set.

After discussions with clinicians we identified clinically relevant features for each data set: For the shoulder data set, the clinicians were interested in seeing the bone surface in context with the muscle layer. Figure 8.4 shows possible visualization results of our technique in conjunction with the predicate histogram used. For the achilles tendon acquisition, the clinicians need to see the tendon in its whole shape in order to identify possible tears or other lesions. Figure 8.6 displays the results of our visualization technique. Finally, 3D Visualization of carotid ultrasound specifically needs to show the path of the carotid artery and its bifurcation in a spatial context, as shown in Figure 8.7.

To support our predicate-based classification approach, we performed a manual segmentation on the data sets partitioning the ultrasound volumes into semantic layers of anatomy such as skin, muscle and bone (cf. Section 8.2.1 on label predicate).

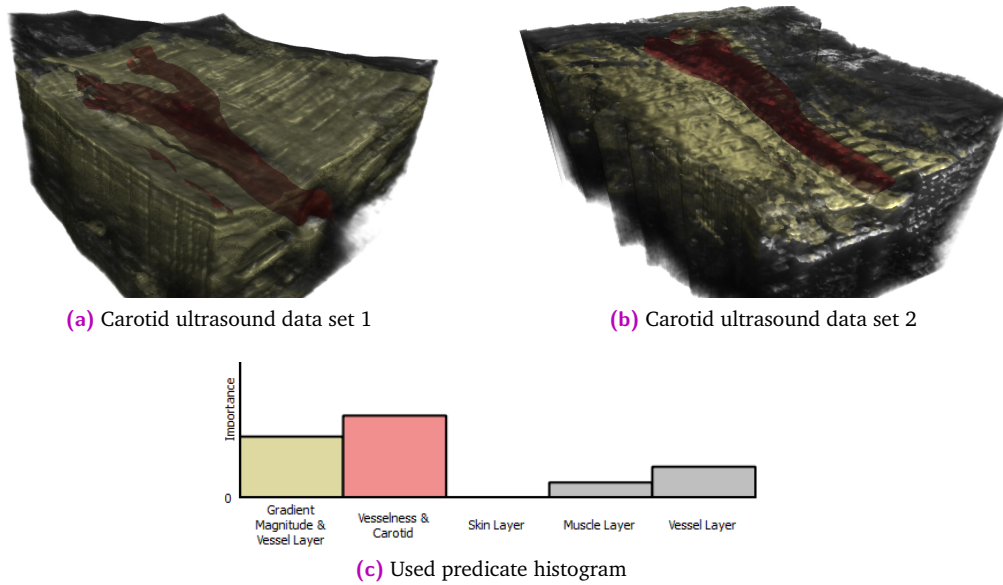


Fig. 8.8. Results suggest a transferability of predicate histograms between different data sets of the same anatomy: (a) A visualization for carotid data set 1 was created using the predicate histogram (b). This was then applied as preset to carotid data set 2 yielding a very viable visualization (c).

For evaluation of the transferability of the predicate histogram setup, we applied the same configuration to both carotid ultrasound data sets. As depicted in Figure 8.8, a single predicate histogram was able to create a proper visualization for both data sets. This suggests that our formulation allows to reuse a predicate histogram, which was created for a specific data set, for a different data set of the same anatomy as valid preset, so that the user does not need to start the exploration and tuning from scratch.

All results presented in this work were acquired in interactive sessions with a nVidia GeForce 670 GTX GPU. While implementing a performance-optimized rendering was not our main goal, our implementation is still capable of generating sufficient frame rates for interactive renderings as shown in Table 1. The decreased performance in comparison to standard 1D transfer function-based direct volume rendering is mainly because the illustrative focus-and-context rendering does no longer allow for early ray termination.

	Achilles Tendon (384 ³ voxels)	Shoulder (512 ³ voxels)
Standard 1D TF-based	102.3 fps	30.4 fps
Predicate-based Rendering	38.9 fps	15.7 fps

Tab. 8.1. Frame rates of our predicate-based approach compared to standard 1D transfer function-based direct volume rendering for different data sets on a GeForce 670 GTX, viewport size 800x600.

8.7 Conclusion

In this chapter, we introduced a novel consistent formulation for predicate-based classification of volumetric image data. It does not only allow for integration of both local and global image information and even anatomical models, but also naturally extends to illustrative focus-and-context visualization. While our technique does not make any a-priori assumptions on the type of input data or its spatial representation, it was designed for the specific application of direct volume visualization to 3D ultrasound volumes, for which it yields superior results in terms of occlusion and distinctly exposing selected image features than traditional global transfer function-based visualization.

To complement our approach, we additionally proposed the predicate histogram as an effective means for reducing the potentially high dimensionality of the predicate configuration domain. It provides the user with an intuitive interface showing an overview of the parameter space, as well as with interaction metaphors to interactively manipulate the visualization result in real-time. Especially in combination with the implemented scribble technique, the predicate histogram can be modified by direct interaction with the rendered ultrasound image. Thus, we obtain an intuitive workflow, which allows also non-expert users to obtain insightful visualizations. The whole framework is designed to seamlessly integrate in the standard volume rendering pipeline without significantly increasing the computational burden and thus allowing for a real-time interaction with the system.

We compared our rendering results only to standard direct volume rendering with a 1D grayscale transfer function, which may raise the question of how our technique compares to other related visualization techniques referenced in Section 8.1. The main issue is that most of these techniques are not applicable to ultrasound data. Ultrasound intensities are of highly context-sensitive nature (cf. Sections 3.1 and 3.3) so that transfer functions solely relying on local information, such as intensity and gradient, neither improve the visual outcome nor are a suitable domain for defining color. Since our results rely to some extent on segmentation output, our approach can be seen as specialization of multi-volume rendering techniques, which could yield similar results. However, we would like to stress that the focus of our work is the full integration of a large variety of classification techniques into a single consistent formulation in combination with an intuitive user interface that allows for highly interactive exploration of the data.

Part III

Conclusion

Discussion and Conclusion

Within this thesis, we presented various novel techniques for advanced ultrasound processing and visualization that were designed to be real-time capable and thereby fit into the ultrasound imaging pipeline. This chapter will summarize and discuss these works and provide ideas and possible directions for future work and improvements.

9.1 Discussion and Future Work

The presented methods can be roughly divided into two parts: After initially proposing techniques that can be applied to 2D B-mode image sequences, we introduced an advanced compounding method to reconstruct 3D ultrasound volumes and presented a generic classification framework to visualize such volumetric data.

Real-time Ultrasound Uncertainty Estimation

The availability of uncertainty information is a crucial part of many of the presented algorithms in this thesis. Therefore, we began in Chapter 5 with presenting a novel method to compute ultrasound Confidence Maps [77] in real-time. By leveraging the dynamic nature of ultrasound acquisitions, we integrated a feedback loop into an iterative solver scheme. Since two consecutive images of an ultrasound sequence are usually very similar, their corresponding Confidence Maps are as well. Thus, we can use the result of one frame as initialization when solving the problem for the subsequent frame. This allows us to yield accurate estimations in the tight time constraints of real-time imaging. Our detailed evaluation has shown that the resulting estimation errors are negligible, particularly for the usage within visual computing algorithms.

Some people argue that ultrasound Confidence Maps are of limited use due to their underlying mathematical model being too simple. However, until today, they are the only generic framework that uses ultrasound-specific constraints to estimate per-pixel signal loss from B-mode images. Thus, bringing this framework into a real-time capable formulation is an important step to make uncertainty information available in the interactive ultrasound processing pipeline.

One major question for future work is how the presented technique can be extended to the third dimension since 2D matrix ultrasound transducers are becoming clinically more significant. While it is absolutely possible to apply the incremental solver scheme also to sequences of 3D volumes, we assume that a practical realization would exceed the computational power of today's hardware when targeting real-time applications due to the heavily increased problem size. However, one should keep in mind that, due to physical constraints (i.e. speed of

sound), 2D matrix probes have a smaller temporal and spatial resolution compared to linear transducers. Thus, a real-time estimation of uncertainty might be possible on low-resolution versions of the image even on today's hardware. A thorough evaluation of the impact of the downsampling on the image quality of the computed 3D Confidence Maps is required in order to get reliable answers.

2D Ultrasound Uncertainty Visualization

One idea for exploiting the previously computed Confidence Maps to improve ultrasound imaging was presented in Chapter 6, which introduced interactive uncertainty visualization techniques for B-mode ultrasound. We deem this beneficial for both educational and clinical applications and therefore proposed three different perceptual mapping schemes using different visual variables to depict per-pixel uncertainty directly in the original B-mode image. These techniques were integrated with the real-time confidence estimation framework from the previous chapter and a clinical ultrasound machine into a fully working clinical setup. The evaluation showed very positive results, in particular for educational applications. The interviewed medical students reported that the detailed visualization of signal loss and resulting uncertainty significantly helps them to gain a better understanding of the complex ultrasound image formation process. Some of the expert sonographers mentioned that the shown uncertainty was somewhat unexpected and that it was missing the detection of reverberation artifacts. Nevertheless, they see great potential in this technique and confirmed that such visualization can actually help in clinical applications, for instance when optimizing the acoustic window.

In general, the response of the test subjects was much better when presenting them the fully integrated system compared to just confronting them with offline results. This underlines the significance of maintaining the interactivity of the ultrasound imaging pipeline, which means that every single step in the processing pipeline has to be real-time-capable.

Future work could investigate how the presented visualization techniques could be improved. For instance, the current mappings make it very difficult to extract exact quantitative information about the local uncertainty from the image since they only use a linear mapping to the visual variables. Techniques, such as isolines could provide the observer with an intuitive and non-intrusive visual cue on where certain uncertainty thresholds are located. However, we actively decided against such visualizations because the underlying Confidence Maps exhibit only relative information about the intensity distribution within the current image. Showing concrete absolute values would run the risk of the observer misinterpreting the visualization. The clinician might understand that everything below a certain threshold is wrong and should be ignored, which is simply false (instead, the information at this location is just less reliable and should be interpreted with care).

Furthermore, the remarks of some of the clinicians that the visualization was sometimes unexpected point out another important detail of our presented technique, which actually aligns with the discussion in the previous section. The mathematical model of ultrasound Confidence Maps is rather simple. The integration of algorithms that can for instance detect reverberation artifacts is therefore a key step to make our technique more robust and clinically relevant. Despite these drawbacks of our current implementation, the evaluation results

clearly show that our visualization techniques provide a valuable addition to traditional plain B-mode ultrasound images.

Advanced Ultrasound Compounding

As of today, the majority of clinical ultrasound acquisitions is performed in 2D mode. However, being restricted to a 2D imaging plane may impose limitations for some applications where the anatomy of interest can not be shown in sufficient detail. One way of extending 2D ultrasound to the third dimension is to equip a traditional transducer with 6 DOF tracking hardware and acquire a continuous sweep, which is then compounded into a 3D volume using the tracking information. Though there have been many works in the past on ultrasound compounding, most of the algorithms used today still impose certain restrictions to the acquisition sweep, such as being formed by a linear trajectory so that the individual frames overlap as little as possible.

Therefore, we present an advanced compounding scheme in Chapter 7 that performs a clustering of the incoming sweep based on the orientation of its frames. This allows us to identify frames with possibly opposing information and ensures that all frames within one cluster are based on the same acoustic window. We then exploit uncertainty information to implement an information fusion approach for compounding. This enables us not only to handle opposing information for the same voxel adequately, but also to implement an incremental compounding scheme. In doing so, we can reconstruct the volume interactively during the acquisition and thereby maintain the real-time nature of ultrasound. The evaluation results demonstrate that our technique reconstructs accurate and high-quality B-mode ultrasound volumes and that the incremental compounding system has sufficient bandwidth for interactive frame rates at clinically relevant voxel sizes.

We see our orientation-driven approach as an important step toward providing clinicians with more freedom regarding 3D freehand ultrasound scanning protocols so that they can focus on *which* images to acquire instead of on *how* they have to acquire them. The main limitation of the presented method is the rather simple pressure compensation technique, which is not always capable of fully compensating for all deformations. Thus, in order to really meet the original goal of completely freeing clinicians from restrictive scanning protocols, one would need to integrate a thorough pressure compensation technique such as the one presented by Treece et al. [184].

Another intriguing idea is to extend the model for uncertainty propagation to 3D. Our proposed technique propagates the uncertainty information from the 2D images to the 3D volume using the same inverse distance weighting that is used for the intensities. It might be interesting to evaluate whether it is beneficial to use a more advanced model to estimate the per-voxel uncertainty also based on the intensity variance of contributing pixels.

Advanced Ultrasound Visualization

The visualization of 3D ultrasound volumes is a challenging task. Traditional classification and visualization techniques as presented in Section 3.2 often yield unsatisfactory results due to the low signal-to-noise ratio of B-mode ultrasound and the fact that the same anatomy may exhibit different intensities and vice versa. Thus, more advanced classification methods

beyond local intensity and gradient are required in order to reliably distinguish different anatomies within the same ultrasound volume.

In order to bring such methods into a unified framework for classification, Chapter 8 introduces the concept of point predicates and shows how they seamlessly integrate into the standard volume rendering pipeline. We further showed that the predicates can be designed and selected to represent semantic layers of the image and introduce the predicate histogram as an intuitive user interface to manipulate the visualization outcome. In combination with an image-based scribble technique, this user interface significantly reduces the parameter domain and general complexity of needed interaction by the user. A qualitative evaluation of our methods on ultrasound volumes of vascular and musculoskeletal anatomies demonstrates that a predicate-based visualization yields superior results in terms of occlusion and distinctly exposing target anatomy than traditional techniques.

A very important aspect is the intuitiveness of our predicate histogram based on the semantic nature of the predicates. Therefore, it is essential to provide the predicates and their settings with easily comprehensible naming. Since all the users needs to do is deciding whether they want to see more or less of a certain feature and adjust the corresponding bar in the predicate histogram accordingly, this can also be considered as an implicit parameter space reduction.

Being a rather fundamental approach in the context of scientific volume visualization and introducing a novel paradigm for its classification stage, the evaluation of our method is limited to qualitative results in terms of case studies. Future work needs to apply the presented concepts to concrete medical applications and can then for instance conduct a task-based quantitative evaluation (cf. Section 2.5).

Reliable and universal classification of anatomies in 3D ultrasound volumes can be considered as one of the remaining "*holy grails*" in medical visualization. Since ultrasound intensities rather describe changes of physical properties than the physical properties of the underlying anatomy themselves, it is indispensable to use non-local features for classification. One promising feature for this could be the speckle texture. Although the texture appearance itself does not necessarily correspond to the underlying tissue microstructure, the local brightness and distribution can [124]. Different statistical models have been presented for ultrasound tissue characterization [167, 179, 210]. Future work could investigate how one can derive a mathematical descriptor from them that can be integrated into the classification scheme.

Though our work was originally designed for the application with volumetric ultrasound, the general concept of predicate-based classification can certainly also be applied to other modalities. Since one can see our method also as a specialization of multi-volume rendering, future work could further investigate whether the predicate histogram provides an intuitive user interface also in this context. Our early experiments in this direction showed promising results.

Finally, the parameter domain of predicates might be well suited for applying machine learning techniques, which could automatically generate application-specific histogram presets. In our evaluation, we showed that one predicate histogram can yield meaningful visualization for multiple data sets of the same anatomy. This suggests that the parameter domain is transferable

between data sets and thus well suited for machine learning approaches. They could generate application-specific workflow models to automatically provide predicate histogram presets.

Part IV

Appendix

CAMPVis – A Game Engine-inspired Framework for Medical Imaging and Visualization

All works discussed in this thesis were implemented in the CAMPVis software framework, which was developed in the course of this PhD. The original idea behind this framework was to create a unified platform for the development of medical imaging and visualization techniques integrating algorithms of the entire visualization pipeline (cf. Section 2.3). It is designed to be used as a rapid prototyping framework in heterogeneous university research environments where many people work on a variety of different topics and to promote synergy effects between them. Furthermore, it should bridge the gap between the development of novel techniques and their deployment into a clinical environment.

As a novel approach, the CAMPVis framework uses software design aspects of modern video game engines. This chapter will discuss the software design decisions that were made based on the identified requirements. Furthermore, we will present important features of CAMPVis as well as two example applications that were not only developed in CAMPVis but later also deployed into a clinical setup.

Parts of this chapter have been published in a Technical Report [160].

A.1 Motivation

Recent advances in the field of medical image computing provides today's clinicians with a large collection of imaging modalities and algorithms for automatic image analysis. However, translating innovations from research into the daily workflow of clinicians is a difficult and time-consuming task since the deployment into a clinical setup poses various challenges. Developing solutions for medical imaging and visualization beyond mere image viewers does usually not yield small and self-contained algorithms but rather an aggregation of many libraries and algorithms. This is mainly due to the visualization pipeline (cf. Section 2.3) being rather long and requiring many preprocessing steps (e.g. image retrieval, registration, filtering, etc.) before one can start with the actual work on the visualization aspects. Thus, this is an extensive field of research and brings together experts and researchers of different disciplines who are working on various aspects of the visualization pipeline.

However, at the same time everybody needs a more or less complete implementation of the entire pipeline in order to implement and evaluate their work: Researchers working on novel processing methods need sophisticated visualization techniques in order to evaluate their

results. At the same, time visualization researchers also require advanced preprocessing algorithms in order to yield high quality input data for their renderings. This mutual dependency offers a large potential for synergy effects when researchers work within the same software framework.

Unfortunately, such highly interdisciplinary work often runs into problems when it comes to sharing a common code base or integrating the work from multiple working groups into a single solution. In particular in environments with limited funding and high employee fluctuation, such as universities, once created software libraries are often abandoned after finishing the project and are hardly designed to be reused by others.

With this motivation in mind, we identified the following list of requirements and design goals focusing on the usage in heterogeneous academic environments:

- *Modern software architecture*: Usage of platform-independent and standard-compliant state-of-the-art techniques. Start mostly from scratch and avoid deprecated interfaces due to forced backwards-compatibility.
- *Bridging the gap between development and deployment*: Focus on research-usage supporting rapid software prototyping, but at the same time allowing for easy transformation of implemented algorithms to end-user products.
- *Sandbox-like environments*: While using the same code base, multiple developers can implement code independently from each other without forcing each other to meet one's code dependencies.
- *Distributed/decentralized computing*: Allow CAMPVis to be run on different devices and support communication between them in order to share computational resources.

In order to meet these requirements, we use the software architecture design of modern video game engines as reference and inspiration. 3D video games have a long tradition in simulating complex environments and providing real-time visualizations. They usually run on a wide range of hardware and massively multiplayer online games (MMOGs) even manage to synchronize the game state over thousands of computers. Hence, video game engines provide a promising approach to integrate the handling of a large amount of data, real-time graphics, interactivity and network computing into a uniform, extensible infrastructure.

A.2 Related Work

As of today, there is a plethora of different software platforms and applications for medical imaging and visualization available. Categorizing them in a structured manner is a challenging task as one can differentiate them along various dimensions.

The specific focus of the software platforms is one dimension to differentiate. Some libraries such as ITK [8] target only the image processing part while others like VTK [9], Voreen [112], Inviwo [177] or ImageVis3D [51], only focus on the visualization part. Most software frameworks, however, try to combine both aspects into a single platform. While 3D Slicer [50] and MITK [206] emphasize the application domain, other frameworks, such as MeVisLab

[147], DeVIDE [16] or XIP [12], use the concept of a data flow network to better support rapid prototyping development.

In 2007, Bitter et al. compared four freely available frameworks for medical imaging and visualization based on ITK [10]. The survey paper of Caban et al. focuses more on the rapid development aspect of such libraries [29].

A.2.1 Entity Component System Architecture

Many modern game engines exploit data-driven programming [3, 96] and implement the *Entity Component System* (ECS) paradigm as main software architecture. While there is no official definition of this paradigm, most approaches show strong similarities in their central design. The main intention of ECS is to yield a cleaner software architecture than classic object-oriented programming (OOP) approaches. Game engines have a large number of different game objects, each of them being formed of multiple aspects, such as physics, player interaction, graphical representation and automation. Trying to model such a complex setup by a traditional OOP class hierarchy will yield a very complicated inheritance graph that is very hard to maintain and extend [108, 13, 14].




Fig. A.1. UML diagram illustrating the Entity Component System software architecture, which is separated into three parts: The data domain storing the system state, the algorithm domain storing the functionality and the EntityManager as database for storing all entities.

An alternative is to follow the common “favor composition over object-inheritance” paradigm [57]. Therefore, the central idea of ECS is to decouple objects from their state and their functionality as illustrated in Figure A.1. This is achieved through introducing three concepts [10, 13]:

Entity The entity is the single general purpose object that stores neither data nor functionality (i.e. methods). Its sole purpose is to provide a tag for each game object.

Component Components are attached to entities and store the raw data but no functionality. Their purpose is to define a certain aspect of the object and how it interacts with the world. Attaching a component to an entity labels the entity to possess this particular aspect. An entity can have multiple components attached and each component can be attached to multiple entities.

System The system defines the actual functionality. Usually, there is one system for each component (aspect) that models and implements the global interaction and functionality of the game.

This concept allows for a very flexible game design where usually many objects of different type share parts of their aspects. Squeezing this into a classic OOP architecture would require a highly fragmented and complex inheritance graph of a large number of very tiny classes and interfaces or a bunch of quite large and partly redundant classes [108].

A.3 CAMPVis Software Architecture

The core of our software framework consists of four main components that interact together in order to generate output (cf. Figure A.2).

- *DataContainers* act as central database storing and managing all non-temporary data that occurs during execution.
- A *Pipeline* defines what computations are actually performed, handles user interaction and provides the output render data.
- *Processors* act as building blocks implementing specific algorithms. Although one could certainly implement all algorithms fully inside a pipeline, we do not encourage this. Instead, encapsulating single algorithms in processors makes it easier to re-use them in other projects.
- *Properties* are used for configuring implemented algorithms.

In the application domain, these components are then exposed through an *OpenGL canvas* that takes care of bringing the corresponding pipeline's output onto the screen, as well as through *PropertyWidgets* that wrap around properties to automatically generate a graphical user interface (GUI).

In this section, we will describe how these components interact with each other in order to implement a variant of the ECS paradigm. Furthermore, we will present architectural software design decisions and discuss how they relate to the initial set of requirements.

A.3.1 Build System

To provide a uniform build process across multiple platforms and compilers, CAMPVis uses the CMake build system [7] where the build process is defined using a scripting language. This allows us to effectively manage and configure the various build options, as well as to scan the file system for available modules to implement our module architecture (cf. Section A.3.5). In a separate step, the build instructions are then transformed to project and make files for the specific target architecture.

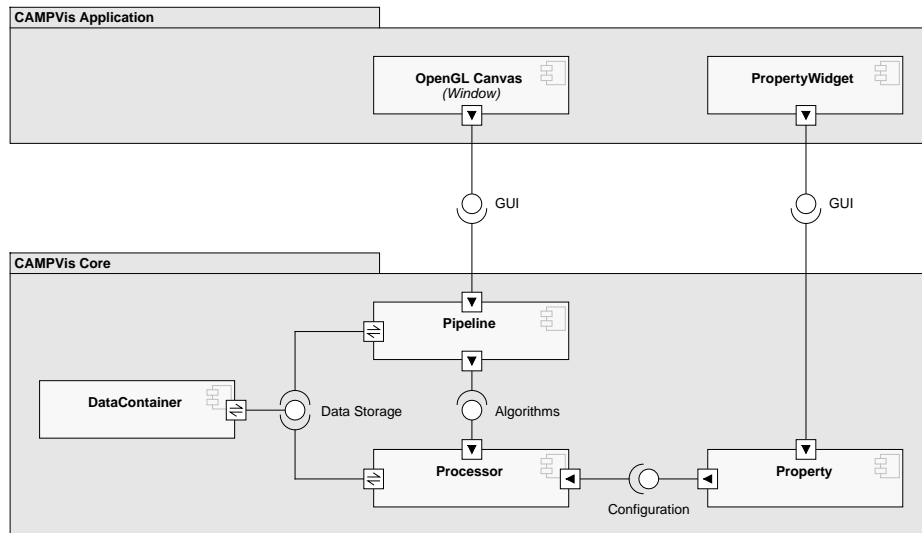


Fig. A.2. UML component diagram providing an overview of the main concepts of CAMPVis and how they interact with each other.

A.3.2 Package Architecture

As introduced in Section A.2.1, the ECS paradigm encourages the separation of data domain and algorithm domain. We transfer this concept to the main CAMPVis package architecture as illustrated in Figure A.3. Driven by our goal to minimize the gap between development and deployment and inspired by Voreen architecture [112], we furthermore separate all GUI toolkit dependent code into a separate package. This allows to easily switch out the rapid-prototyping user interface with an end-user interface suited for clinical setups.

Core Package The core package wraps the data domain and provides the infrastructure for the core system of CAMPVis. These are basic data structures, base classes for processors, pipelines, properties, etc., common tool classes such as type traits or string utilities, as well as common GLSL headers. This package has as few external dependencies as possible and in particular no GUI dependencies.

Modules Package This package wraps the algorithm domain and contains the main functionality of CAMPVis. Modules contain individual processors that implement concrete algorithms, as well as individual pipelines that implement solutions and workflows for concrete applications. This package is also not dependent on GUI libraries. An important aspect of the modules package is that the package actually is a collection of modules, which can individually be selected to be included into the build or not. This allows to easily manage multiple independent modules maintained by different people. If one of them is faulty it can be excluded from the build so that CAMPVis is still compilable (cf. Section A.3.5).

Application Package The application package provides the tools to build an actual application based on CAMPVis. It ties together the core and modules package and

provides the user with a default GUI. This GUI is intended as research interface and hence exposes all internal parameters (i.e. properties) of the individual processors and pipelines. Furthermore, it comes with a convenient debug interface to inspect the contents of the DataContainer. However, the application package is fully optional and can be replaced by own implementations, for instance when integrating CAMPVis into an existing application.

The number of external libraries required for the core package is kept as small as possible. Besides OpenGL 3.3 and a small OpenGL wrapper library (CAMP Graphics Toolkit, cgt), the single other mandatory external library is Intel TBB [5] providing clean interfaces to support multithreading and concurrent algorithms. All other potential dependencies are part of individual CAMPVis modules and thus only required when the corresponding module is enabled.

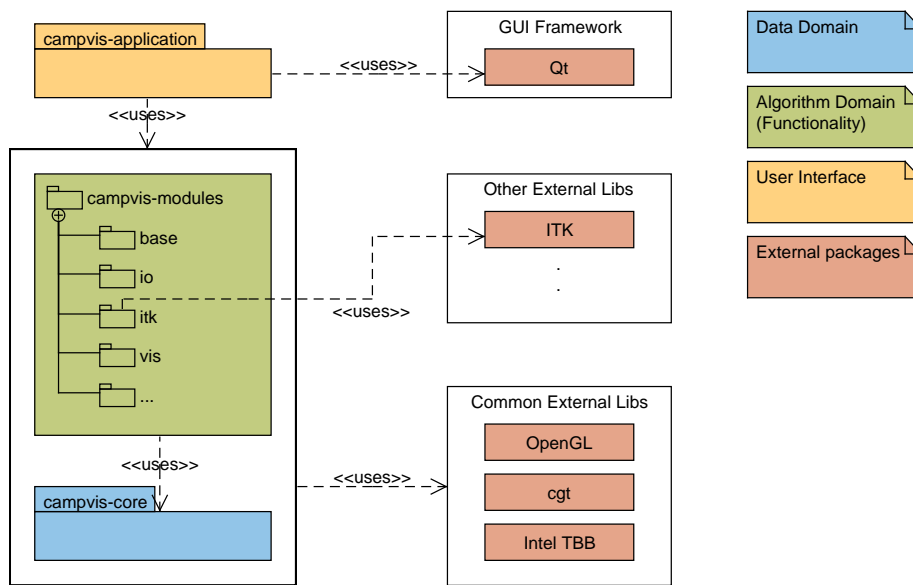


Fig. A.3. UML diagram of the CAMPVis package structure, which features a separation into data domain (core module), algorithm domain (modules package) and user interface (application package).

A.3.3 The Entity Component System for CAMPVis

We use the Entity Component System paradigm as basic architecture for the CAMPVis framework. Its main benefit is the very strict and clear separation of data and algorithms that allows for a great amount of flexibility. However, a software for medical imaging and visualization has a significantly smaller amount of alive objects during runtime than video games and not all of its systems are of fully global nature. Therefore, we applied some modifications to the classic ECS approach presented in Section A.2.1 and developed with the design shown in Figure A.4.

On the data domain, we implement the classic concept of entities, which act as general purpose object and store neither data nor functionality. Each entity has a unique identifier (for which




Fig. A.4. UML diagram illustrating the adapted Entity Component System software architecture for CAMPVis, which also separates the database, the data domain and the algorithm domain from each other.

we use strings in order to facilitate the handling for the user) and is stored and managed by the *DataContainer*, an entity database implemented as map. Therefore, we call our entities *DataHandle*. In the current implementation, every *DataHandle* just points to a single data aspect (e.g. transformation, image, geometry, etc.). However, for the future it is planned to have *DataHandles* aggregating multiple data aspects (components) in order to move this design closer toward the classic ECS architecture.

Regarding the algorithm domain, there is no fixed collection of systems in CAMPVis because of its nature of being a research framework. The necessary systems are rather dependent on the actual project implemented in the CAMPVis platform. Therefore, we decided to adapt the classic ECS model to a platform featuring a library of systems, which behave like building blocks and can be easily assembled together for each individual application. We call these building blocks of systems *processors* as they encapsulate specific algorithms. This approach is very similar to modules in MeVisLab and processors in Voreen. It encourages the developers to break their problems into sub-problems, which facilitates reusing code and provides an excellent rapid-prototyping environment.

However, based on our experience, there are certain limits in terms of generality. Forcing developers to design the processors as generic as possible either leads to a flood of very tiny and specific processors or to massive blobs that do everything alone and try to handle every corner case. Since we consider both these extremes as not desirable, we want to provide as much freedom as possible when combining the processors. Therefore, contrary to MeVisLab or Voreen, CAMPVis does not impose a fixed a-priori structure, such as a data flow network. Instead, CAMPVis offers the very generic concept of a *Pipeline* that coordinates the data domain (*DataContainer*) with the algorithm domain (processors). Every pipeline works on a single *DataContainer* and can aggregate multiple processors (Figure A.5). The evaluation logic can be automated (e.g. to simulate a data flow network) as well as a custom pipeline-specific implementation. This provides maximum freedom to developers by allowing them to implement pipeline-specific code directly in the pipeline, instead of forcing them to extend existing processors or writing new ones for one-time tasks. In ECS terminology, a pipeline represents a system and pipelines are either executed in a continuous (render) loop or event-based on user interaction.

CAMPVis processors are designed to be very loosely coupled: Following the ECS paradigm, a processor only implements functionality and is inherently state-less regarding the data domain. Hence, it neither is aware of other processors it might be collaborating with nor does it know a-priori on which data to work on. Instead, this information is provided when executing




Fig. A.5. Illustration of the CAMPVis pipeline-processor concept. CAMPVis pipelines take care of coordinating the algorithm domain (processors) with the data domain (DataContainer).

the processor by passing a reference to the DataContainer holding the entities. This is also the indented way of coordinating the processors within a pipeline (cf. Figure A.5). During execution, the processor queries the DataContainer for entities with certain components and performs its computations on them. The results are then pushed back into the DataContainer so that they can be used by other processors.

A.3.4 Properties

In order to support the rapid prototyping design goal, CAMPVis also offers a property system to configure processors and pipelines, similar to the one found in ITK, Voreen and Inviwo. Instead of directly configuring a class through primitive local member variables, one should use *Properties*, which wrap around data types and offer various benefits:

- Providing an observer-like behavior
- Providing automatic getter/setter methods
- Taking care of thread-safety
- The application package can provide an automatically generated GUI

A.3.5 Module Architecture

One of the initial requirements was to provide sandbox environments for developers. Since CAMPVis is intended to be used as research platform in a heterogeneous environment, there may be a large number of developers working on different projects at the same time.

Therefore, we developed a module system using CMake build scripts. It parses the file system for available modules and allows the user to select for each module whether it should be included into the build or not. Thereby, each module can be considered as a separate, independent sandbox minimizing possible side effects when having multiple projects sharing the same code base. For instance, if one module requires an external library as dependency, other independent modules are still able to compile and run without it. At the same time however, it is possible to define module dependencies so that developers can easily reuse code.

A.4 Framework Features

The CAMPVis software framework offers various features to programmers to implement their objective. We will present the most important ones in this section.

A.4.1 Signal Manager

For a large-scale software framework, inter-object communication becomes an important issue. Game engines often feature an event system for this purpose, where objects notify potential listeners by sending generic events through a central event manager instance [108]. Since sender and receiver of events do not know each other, the sender object simply sends an event message of certain type and the event manager then takes care of delivering the message to all listeners that registered themselves for this event type. The major advantage of such a system is its simplicity in design, easy decoupling of sending and receiving events (asynchronous messaging) and the fact that all communication runs through a central place, which makes tracking, monitoring and debugging of communication easy.

However, for our targeted software platform for medical imaging, such a system has one disadvantage: The mapping between sender and receiver of a message is solely based on the event type. Since the inheritance graph of our software framework is broader than it is deep (i.e. many communicating classes inherit from the same base class), effective filtering and routing of messages becomes an issue: One solution would be declaring a distinct event type for each subclass. However, since the semantic nature of all those events is the same (only the sender type changes), this approach would not follow clean object-oriented software design. The other solution would be to define the event type in the common base class, so that all child objects send messages using the same type. In such a case, however, the receiving object would receive all messages of all objects and thus needs to filter out the relevant messages.

Therefore, we decided to use the signal-slot pattern for our software framework and enhance it with the central manager part of event systems. As in the traditional signal-slot pattern presented by Qt [11], relationships between senders (signals) and receivers (slots) are defined through connections. Thus, the sender does know which objects would like to receive its messages. The actual processing of the communication is however done by the signal manager, which acts as a central singleton and takes care of the dispatching of messages. This way, we achieve the flexibility of signals and slots in combination with easy tracking and monitoring, as well as with implicit asynchronous messaging.

Our signal-slot API allows for emitting signals in three different ways:

Direct/blocking call Using `signal::triggerSignal()`, the signal will be processed directly in the emitting thread, the call will block until all signal processing has finished.

Asynchronous call Using `signal::queueSignal()`, the signal will be put into the signal manager queue and processed asynchronously in the signal manager thread. Hence, the call will immediately return.

Default call Using `signal::emitSignal()`, the signal will be queued by default, unless the calling thread is also the signal handling thread (i.e. in case of cascading signals). This ensures that cascading signals are processed in a single batch.

Implementation Details We designed a special data structure to store the per-signal connection information. Our `concurrent_pointer_list<T>` data structure is a list-like container that allows for thread-safe bidirectional iteration, insertion and removal of elements. Instead of removing deleted items from the data structure, we mark them internally as `nullptr`. Since connections are stored as pointers to slots and pointer types support atomic operations, this is an effective solution to avoid per-signal mutexes.

Furthermore, we use a memory pool for creating the signal handles, which are relatively tiny objects and created with high frequency from different threads. Relying on the default system allocator here would introduce a performance penalty as it usually employs critical sections around each allocation and deallocation. Implementing a pre-allocated memory pool as custom allocator elegantly circumvents this issue. For minimal efforts, we use `tbb::memory_pool` in `CAMPVis`, which yields a speedup of almost 20% compared to the standard memory allocator.

To facilitate the debugging of sent messages, we implemented a transparent debug layer into our signaling API. When built with the debug flag enabled, all emitted signals will automatically store information on the calling function, file and line together with the sent message (Listing A.1). This compensates for the incomplete call stack information in multi-threaded, asynchronous messages.

```
class signal0 {
    // [...]
    void triggerSignal();
}

#ifdef CAMPVIS_DEBUG
struct signal0_debug_decorator : public signal0 {
    // overload method to store debug information and return this
    signal0& triggerSignal(string caller, string file, int line) {
        storeDebugInfo(caller, file, line);
        return *this;
    }
};

// redefine all necessary symbols
#undef signal0
#define signal0 signal0_debug_decorator

#undef triggerSignal
#define triggerSignal triggerSignal(__FUNCTION__, __FILE__, __LINE__).triggerSignal
#endif
```

Listing A.1 Code excerpt of the debugging hooks for our signal-slot API. Using C macros, we can redefine the signal symbol to a debug implementation that stores information on the calling function, file and line for debugging purposes.

A.4.2 Factory Registration

Another notable feature of the CAMPVis software platform is dynamic module registration. This lets modules register their classes with object factories, so that the core code can generically access their functionality without actually knowing of them at build time. With CAMPVis, this counts for instance for the dynamic registration of image converters, pipelines and property widgets.

CAMPVis combines the C++ static registration idiom with the factory pattern. Using C++ templates and static member variables, we achieve a non-intrusive solution for automatic module registration at static initialization time. Its central piece is the templated Registrar class with two essential parts:

1. a static function to create an object of that type (as it knows the object type through the template),
2. a static integer member variable.

The initialization of the integer member is performed by a function call to the factory singleton's registration method (cf. Listing A.2), which returns an integer to store in the static field (since the actual value does not matter, our implementation returns the sequential number of the just registered object).

```
// The Registrar class takes care of the actual registration.
// Template parameter is the type of the registree
template<typename T>
class Registrar {
    // Static factory method for creating the object of type T.
    static AbstractRegistree* create() {
        return new T(dc);
    }

    // static field stores the result of the registration function
    static const size_t _factoryId;
};

// Static initialization performs the function call to register
// the registree with the factory.
template<typename T>
const size_t Registrar<T>::_factoryId = Factory::register<T>(&Registrar<T>::create);

// In a particular module, define a concrete class to register:
class ConcreteRegistree : public AbstractRegistree {
    // [...]
}

// Explicit template instantiation: instantiate the registrar,
// which performs the registration at static initialization time
template class Registrar<ConcreteRegistree>;
```

Listing A.2 Code excerpt showing the C++ static registration idiom we use for factory registration.

A.4.3 Scripting Layer

CAMPVis has an optional scripting layer, which can be enabled in order to add support for runtime-scripting with the Lua scripting language [4]. This serves various purposes: On the

one hand, it offers a scripting console allowing the user to inspect and modify the data model at runtime through Lua commands. Similar to MATLAB or many game engines, this provides developers with straight-forward interaction with their software at runtime without the need to explicitly program a graphical user interface. On the other hand, the scripting layer allows defining entire pipelines in Lua scripts (cf. Listing A.3). This further accelerates the rapid prototyping nature of CAMPVis since changes to pipelines no longer relate to a recompilation of C++ code but only to updating the script and loading it at runtime.

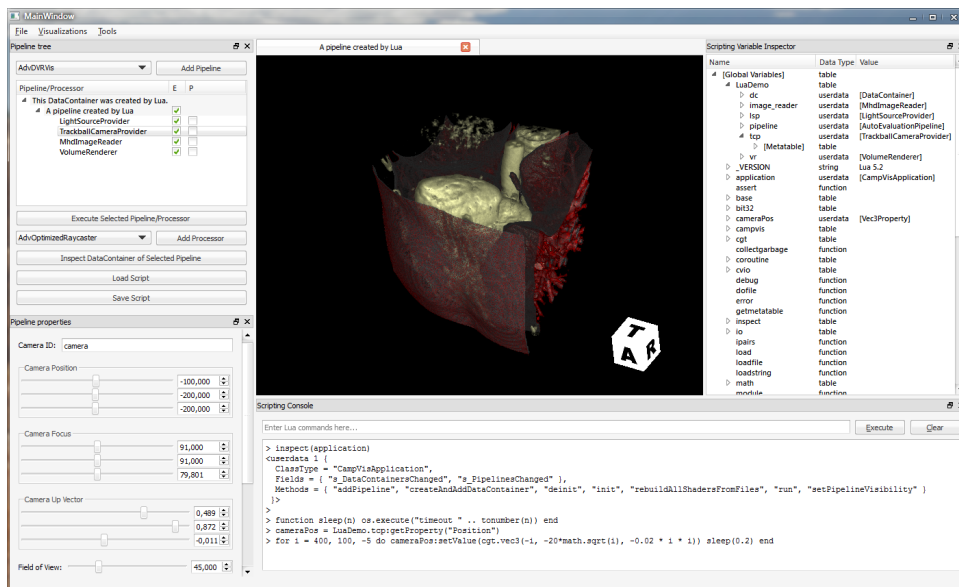


Fig. A.6. Screenshot of CAMPVis with enabled Lua scripting functionality. The scripting console in the bottom dock allows for accessing and manipulating the internal CAMPVis data model at runtime, for instance to modify the properties of processors in a systematic fashion. Furthermore, the widget in the dock on the right-hand side allows to inspect all variables and tables that are present in the Lua virtual machine.

Furthermore, we exploit the scripting support as persistence mechanism: To save the current program state in a file, CAMPVis writes a Lua script holding the necessary code to recreate the pipeline and its state. The program state can later be recovered by simply executing the Lua script.

Implementation Details C++ classes are exposed to the Lua virtual machine by creating appropriate bindings for them. The scripting layer uses SWIG [2] to wrap selected CAMPVis modules and core packages, and generate Lua modules that make them available to scripted pipelines. The process of generating bindings is not completely automated - SWIG must be told what classes to wrap, which of their members to expose, and how to deal with advanced C++ features such as templates. This information is encoded in interface files which are passed to SWIG along with the C++ code to produce Lua modules (cf. Figure A.7).

A.4.4 Network Communication

One of the initial requirements during CAMPVis development was a good support for network communication, which has several use cases. Multi-modal image fusion often has to deal

```

require("application")
require("cvio")
require("vis")

-- All created CAMPVis objects have to be kept alive and must not be garbage-collected
-- by Lua. Thus, we create a global table for this script, where everything resides in
LuaDemo = {}

-- create DataContainer and Pipeline
LuaDemo.dc = application:createAndAddDataContainer("DataContainer Name")
LuaDemo.pipeline = campvis.AutoEvaluationPipeline(LuaDemo.dc, "Pipeline Name")
local pipeline = LuaDemo.pipeline;

-- create the processors we need
local canvas_size = pipeline:getProperty("CanvasSize")
LuaDemo.lsp = base.LightSourceProvider()
LuaDemo.tcp = base.TrackballCameraProvider(canvas_size)
LuaDemo.image_reader = cvio.MhdImageReader()
LuaDemo.vr = vis.VolumeRenderer(canvas_size)

-- register the processors with the pipeline
pipeline:addProcessor(LuaDemo.lsp)
pipeline:addProcessor(LuaDemo.tcp)
pipeline:addProcessor(LuaDemo.image_reader)
pipeline:addProcessor(LuaDemo.vr)

-- setup event listener and register it with the pipeline
LuaDemo.tcp:addLqModeProcessor(LuaDemo.vr)
pipeline:addEventListenerToBack(LuaDemo.tcp)

-- create an init callback function, so that the following code gets called when the
-- pipeline gets initialized by CAMPVis.
local initCallback = function()
    -- set up the processors' properties
    LuaDemo.vr.p_outputImage:setValue("combine")
    pipeline:getProperty("renderTargetID"):setValue("combine")

    LuaDemo.image_reader.p_url:setValue(campvis.SOURCE_DIR
        .. "/modules/vis/sampledData/smallHeart.mhd")
    LuaDemo.image_reader.p_targetImageID:setValue("reader.output")
    LuaDemo.image_reader.p_targetImageID:addSharedProperty(LuaDemo.vr.p_inputVolume)

    -- automatically adjust the camera to the data
    LuaDemo.image_reader.p_targetImageID:addSharedProperty(LuaDemo.tcp.p_image)
end
-- register the callback with the s_init signal
pipeline.s_init:connect(initCallback)

-- Finished creating our LuaDemo -> register it with the CAMPVis application
application:addPipeline(pipeline)

```

Listing A.3 Example Lua script for creating a fully functional volume rendering visualization with CAMPVis.




Fig. A.7. Illustration of the Lua binding generation process using SWIG. In order to expose CAMPVis functionality to the Lua virtual machine, SWIG uses interface files to create Lua bindings for the existing CAMPVis code.

with multiple devices, which have a continuous exchange of data such as images, tracking information or control commands. In cases where mobile devices (e.g. tablet computers) lack the necessary computation power, distributed computing can furthermore provide a solution where the actual computations are performed on a stationary workstation and only the results are streamed to the mobile device.

The original plan was to integrate such a network communication stack directly into the CAMPVis core, as it is done with many video game engines. However, since there are already various established solutions available for streaming medical imaging data over network, we decided against implementing another protocol. Instead, networking support is enabled through CAMPVis modules. The current implementation features wrappers for both CAMPCom [159] and OpenIGTLink [182], two state-of-the-art libraries for real-time streaming of medical imaging data.

A.5 Clinical Application

As discussed in Section A.1, it is important to keep the gap between research/development and deployment small. This is the case in particular in the context of medical imaging and visualization where the ultimate goal is always to improve the work of the clinician and/or the outcome for the patient. Since this can only be evaluated by the domain experts themselves, one needs to implement a working prototype that suffices clinical requirements. To demonstrate the capabilities of CAMPVis in this regard, we will present two case studies as examples where the CAMPVis platform was used for both development and deployment.

A.5.1 Real-time Uncertainty Visualization for 2D B-mode Ultrasound

As first example, we discuss our system for real-time uncertainty visualization for 2D B-mode ultrasound, which was introduced in Chapters 5 and 6. Since the motivation, background and methodology of this project has already been discussed there, this section will focus only on the technical implementation regarding CAMPVis and its deployment into a clinical setup.

The entire system consists of an ultrasound machine and a standard workstation, which are connected through wired network (cf. Figure A.8). The ultrasound system acquires the raw echo data and performs the initial processing into a B-mode image, which is then sent via




Fig. A.8. Deployment of CAMPVis into a clinical setup for real-time B-mode ultrasound uncertainty visualization. The image is acquired by the ultrasound machine (left) and sent to a workstation running CAMPVis (right) through a network connection. CAMPVis interactively performs the processing of the incoming images and routes the resulting images back to the ultrasound machine. This way, our system is minimally intrusive since the domain expert only has to deal with the familiar ultrasound device.

OpenIGTLink to the workstation. There, the CAMPVis application receives the image through the OpenIGTLinkClient processor and performs some further filtering using the corresponding processors. The CudaConfidenceMapsSolver processor computes the Confidence Map, which is eventually fused into the final visualization by the AdvancedUsFusion processor. This rendering is then routed back to the display of the ultrasound machine. Since the new visualization is shown at exactly the same place as a normal ultrasound image, and the user only interacts with familiar hardware, our system has a minimal impact on the evaluation results and we can safely assume that they really represent the added value of the visualization technique itself and are not biased by the quality of the system integration.

Since the presented CAMPVis processing pipeline is entirely composed of processors as building blocks, each of the processing steps can be easily reused in a different project to maximize synergy effects (cf. Section A.1). Furthermore, the asynchronous implementation in the CAMPVis framework allows for a smooth and interactive visualization even if single ultrasound frames should drop (e.g. due to network issues).

A.5.2 Multi-modal Image-guided Prostate Biopsy

As second example, we present an integrated system for multi-modal image-guided prostate biopsy, which was recently developed in our lab and improves the specificity of biopsy results for prostate cancer diagnosis [166, 208]. For this application, the current gold standard is a random biopsy of 10 to 12 samples under trans-rectal ultrasound (TRUS) guidance. However, this approach is prone to a high number of false negative results due to TRUS hardly highlighting suspicious regions. We improve this workflow by additionally acquiring a

combined pre-operative PET-MR image of the anatomy using a ^{68}Ga -labeled PSMA radioactive tracer exhibiting a very high specificity for prostate cancer [46]. After registering the PET-MR images with the tracked ultrasound, our system uses multi-modal visualization for guiding the biopsy toward the highly suspicious regions.

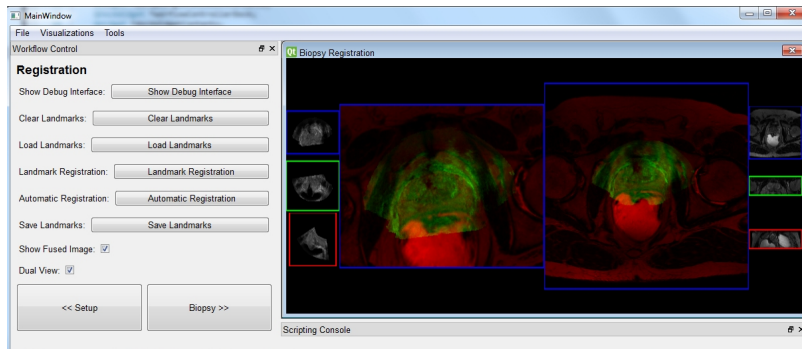


Fig. A.9. Screenshot of the developed multi-modal image-guided prostate biopsy framework [166, 208] implemented in CAMPVis. This image shows the semi-automatic registration step for aligning the tracked ultrasound images with the MR volume, for which the clinician defines a set of corresponding landmarks in each image.

Since the clinical workflow consists of multiple steps, we implemented a custom user interface for the clinician. As a first step, the system is initialized, a fully-automatic machine learning-based registration is performed and all necessary data is read from the disk. In the second step, CAMPVis presents a slice-based multi-modal visualization to the clinician, in which they evaluate the result of the automatic MR-TRUS registration. If the images are not aligned well, the clinician can define a set of anatomical landmarks in each modality in order to yield a better registration. The final step then uses OpenIGTLink to stream the ultrasound image and tracking information to CAMPVis, similar as with the presented ultrasound uncertainty visualization system. A real-time multi-modal visualization of the data the guides the clinician. One window shows the live B-mode ultrasound image and a second window shows the corresponding MPR in the PET-MR volumes, both featuring a needle guide to target suspicious regions appearing with high intensity in the PET data.

Since most functionality is encapsulated in processors, many parts of this project's code can easily be reused. For instance, the `OpenIGTLinkClient` processor is the same as the one used for the uncertainty visualization system and the `BiopsyMprRenderer` processor of the third workflow step is also used in a similar fashion in other CAMPVis-based projects. As shown in [208], our system does not only yield significantly improved clinical results in terms of biopsy specificity but also got very good feedback from our clinical partners in terms of usability and effectiveness.

A.6 Conclusion

We presented a new software framework targeted for both development of medical imaging and visualization techniques in heterogeneous research groups as well as for deployment into a clinical setup. As novel approach, the system architecture and design was inspired by

modern video game engines, which are very effective in handling large amounts of data in an interactive real-time environment.

The development of CAMPVis was funded by the the “SoftwareCampus” program of the German Federal Ministry of Education and Research (BMBF, Förderkennzeichen 01IS12057). The CAMPVis software framework is be published under the permissive Apache License, Version 2.0 to prevent any licensing issues in case parts of the code will make its way into commercial products.

List of Authored and Co-authored Publications

2016

- [66] Katharina Hofschien, Timo Geissler, Nicola Rieke, **Christian Schulte zu Berge**, Nassir Navab, and Stefanie Demirci. “Image Descriptors in Angiography”. *Proceedings on Workshop Bildverarbeitung für die Medizin (BVM)*, 2016, Berlin, DE.

2015

- [163] **Christian Schulte zu Berge**, Denis Declara, Christoph Hennesperger, Maximilian Baust, and Nassir Navab. “Real-time Uncertainty Visualization for B-Mode Ultrasound”. *Proceedings on Scientific Visualization (IEEE Vis)*, 2015, Chicago, USA.
- [164] **Christian Schulte zu Berge**, Jakob Weiss, and Nassir Navab. “Schematic Electrode Map for Navigation in Neuro Data Sets”. *Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*, 2015, Chester, UK.
- [208] Oliver Zettinig, Amit Shah, Christoph Hennesperger, Matthias Eiber, Christine Kroll, Hubert Kübler, Tobias Maurer, Fausto Milletari, Julia Rackerseder, **Christian Schulte zu Berge**, Enno Storz, Benjamin Frisch, and Nassir Navab. “Multimodal Image-guided Prostate Fusion Biopsy Based on Automatic Deformable Registration”. *International Journal of Computer Assisted Radiology and Surgery* 10.12, (2015).

2014

- [162] **Christian Schulte zu Berge**, Maximilian Baust, Ankur Kapoor, and Nassir Navab. “Predicate-based Focus-and-Context Visualization for 3D Ultrasound”. *IEEE Transactions on Visualization and Computer Graphics* 20.12, (2014).
- [207] Oliver Zettinig, Christoph Hennesperger, **Christian Schulte zu Berge**, Maximilian Baust, and Nassir Navab. “3D Velocity Field and Flow Profile Reconstruction from Arbitrarily Sampled Doppler Ultrasound Data”. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2014, Boston, USA.
- [166] Amit Shah, Oliver Zettinig, Tobias Maurer, Cristina Precup, **Christian Schulte zu Berge**, Jakob Weiss, Benjamin Frisch, and Nassir Navab. “An Open Source Multimodal Image-guided Prostate Biopsy Framework”. *Clinical Image-Based Procedures – Translational Research in Medical Imaging (MICCAI-CLIP)*, 2014, Boston, USA.

- [160] **Christian Schulte zu Berge**, Artur Grunau, Hossain Mahmud, and Nassir Navab. CAMPVis – A Game Engine-inspired Research Framework for Medical Imaging and Visualization. *Technical Report, Technische Universität München, 2014.*
- [194] Xiang Wang, **Christian Schulte zu Berge**, Stefanie Demirci, Pascal Fallavollita, and Nassir Navab. Improved interventional X-ray appearance”. *IEEE Symposium on Mixed and Augmented Reality (ISMAR), 2014, Munich, DE.*
- [161] **Christian Schulte zu Berge**, Ankur Kapoor, and Nassir Navab. “Orientation-driven Ultrasound Compounding Using Uncertainty Information”. *Information Processing in Computer-Assisted Interventions (IPCAI), 2014, Fukuoka, JPN.*


2011

- [44] Stefan Diepenbrock, **Christian Schulte zu Berge**, Klaus Hinrichs, Lydia Wachsmuth, and Cornelius Faber. “DTI Visualization Using the Voreen Framework”. *Poster at ESMRMB Congress, 2011, Leipzig, DE.*

Abstracts of Publications not Discussed in this Thesis

Schematic Electrode Map for Navigation in Neuro Data Sets

C. Schulte zu Berge, J. Weiss, N. Navab



(a) Our proposed schematic electrode map with the electrode glyph showing an overview over the depth electrode configuration. Furthermore, it serves as navigation tool for linked multi-modal visualizations, e.g. (b) definition of the 2D MPR plane; (b) camera placement or DTI fiber filtering in 3D volume rendering.

Neuro resection surgery is one of the last resorts when treating epilepsy patients where conservative treatment shows no effect on seizure reduction. However, due to the severity of the surgery, the resection planning has to be as precise as possible in order to avoid harming any critical anatomy. The tight time constraints in clinical routine demand for a highly optimized workflow. In this work, we therefore introduce a novel visualization in order to simplify the navigation in the complex multi-modal neuro data sets and support the clinician with the planning procedure. We propose a schematic electrode map based on a force-directed graph model providing an intuitive overview over the topology of the implanted depth electrode configuration. To further facilitate the planning workflow, our carefully designed electrode glyph supports different scalar, nominal and binary annotations augmenting the view with additional information. Brushing and linking techniques allow for easy mapping of the EEG data to the corresponding anatomy, as well as for straight-forward navigation within the visualization of the anatomical and functional imaging modalities in order to identify the origin and spread of the seizure. Our results show that the proposed graph layouting method

successfully removes occlusions of the projected electrodes while maintaining the original topology of the depth electrode configuration. Initial discussions with clinicians and the application to clinical data further show the effectiveness of our methods.

Eurographics Workshop on Visual Computing for Biology and Medicine (2015)

Image Descriptors in Angiography

K. Hofschen, T. Geissler, N. Rieke, C. Schulte zu Berge, N. Navab, S. Demirci

Despite recent advances in the field of image-guided interventions (IGI), the bottleneck for Angiography/X-ray guided procedures in particular is accurate and robust 2D-3D image alignment. The conventional, straight-forward parameter optimisation approach is known to be ill-posed and less efficient. Retrieval-based approaches may be of superior choice here. However, this requires salient and robust image features, which can handle the difficulties of Angiographic images such as high level of noise and contrast variance. In this paper, we investigate state-of-the-art features of the field of computer vision regarding the applicability and reliability in the challenging scenario of Angiography.

Workshop Bildverarbeitung fuer die Medizin (2016)

Multimodal Image-Guided Prostate Fusion Biopsy based on Automatic Deformable Registration

O. Zettinig, A. Shah, C. Hennesperger, M. Eiber, C. Kroll, H. Kübler, T. Maurer, F. Milletari, J. Rackerseder, C. Schulte zu Berge, E. Storz, B. Frisch, N. Navab

Purpose. Trans-rectal ultrasound (TRUS) guided random prostate biopsy is, in spite of its low sensitivity, the gold standard for the diagnosis of prostate cancer. The recent advent of PET imaging using a novel dedicated radiotracer, ^{68}Ga -labeled PSMA (Prostate Specific Membrane Antigen), combined with MRI provides improved pre-interventional identification of suspicious areas. This work proposes a multimodal fusion image-guided biopsy framework that combines PET-MRI images with TRUS, using automatic segmentation and registration, and offering real-time guidance. *Methods.* The prostate TRUS images are automatically segmented with a Hough transform based random forest approach. The registration is based on the Coherent Point Drift algorithm to align surfaces elastically and to propagate the deformation field calculated from thin plate splines to the whole gland. *Results.* The method, which has minimal requirements and temporal overhead in the existing clinical workflow, is evaluated in terms of surface distance and landmark registration error with respect to the clinical ground truth. Evaluations on agar-gelatin phantoms and clinical data of 13 patients confirm the validity of this approach. *Conclusion.* The system is able to successfully map suspicious regions from PET/MRI to the interventional TRUS image.

International Journal of Computer Assisted Radiology and Surgery, 10.12 (2015)

3D Velocity Field and Flow Profile Reconstruction from Arbitrarily Sampled Doppler Ultrasound Data

O. Zettinig, C. Hennersperger, C. Schulte zu Berge, M. Baust, N. Navab

With the need for adequate analysis of blood flow dynamics, different imaging modalities have been developed to measure varying blood velocities over time. Due to its numerous advantages, Doppler ultrasound sonography remains one of the most widely used techniques in clinical routine, but requires additional preprocessing to recover 3D velocity information. Despite great progress in the last years, recent approaches do not jointly consider spatial and temporal variation in blood flow. In this work, we present a novel gating- and compounding-free method to simultaneously reconstruct a 3D velocity field and a temporal flow profile from arbitrarily sampled Doppler ultrasound measurements obtained from multiple directions. Based on a laminar flow assumption, a patch-wise B-spline formulation of blood velocity is coupled for the first time with a global waveform model acting as temporal regularization. We evaluated our method on three virtual phantom datasets, demonstrating robustness in terms of noise, angle between measurements and data sparsity, and applied it successfully to five real case datasets of carotid artery examination.

International Conference on Medical Image Computing and Computer Assisted Interventions (2014)

An Open Source Multimodal Image-guided Prostate Biopsy Framework

A. Shah, O. Zettinig, T. Maurer, C. Precup, C. Schulte zu Berge,
J. Weiss, B. Frisch, N. Navab

Although various modalities are used in prostate cancer imaging, transrectal ultrasound (TRUS) guided biopsy remains the gold standard for diagnosis. However, TRUS suffers from low sensitivity, leading to an elevated rate of false negative results. Magnetic Resonance Imaging (MRI) on the other hand provides currently the most accurate image based evaluation of the prostate. Thus, TRUS/MRI fusion image-guided biopsy has evolved to be the method of choice to circumvent the limitations of TRUS-only biopsy. Most commercial frameworks that offer such a solution rely on rigid TRUS/MRI fusion and rarely use additional information from other modalities such as Positron Emission Tomography (PET). Other frameworks require long interaction times and are complex to integrate with the clinical workflow. Available solutions are not fully able to meet the clinical requirements of speed and high precision at low cost simultaneously. We introduce an open source fusion biopsy framework that is low cost, simple to use and has minimal overhead in clinical workflow. Hence, it is ideal as a research platform for the implementation and rapid bench to bedside translation of new image registration and visualization approaches. We present the current status of the framework that uses pre-interventional PET and MRI rigidly registered with 3D TRUS for prostate biopsy guidance and discuss results from first clinical cases.

Workshop on Clinical Image-based Procedures: Translational Research in Medical Imaging (2014)

Improved Interventional X-ray Appearance

X. Wang, S. Demirci, C. Schulte zu Berge, P. Fallavollita, N. Navab

Depth cues are an essential part of navigation and device positioning tasks during clinical interventions. Yet, many minimally-invasive procedures, such as catheterizations, are usually performed under X-ray guidance only depicting a 2D projection of the anatomy, which lacks depth information. Previous attempts to integrate pre-operative 3D data of the patient by registering these to intra-operative data have led to virtual 3D renderings independent of the original X-ray appearance and planar 2D color overlays (e.g. roadmaps). A major drawback associated to these solutions is the trade-off between X-ray attenuation values that is completely neglected during 3D renderings, and depth perception not being incorporated into the 2D roadmaps. This paper presents a novel technique for enhancing depth perception of interventional X-ray images preserving the original attenuation appearance. Starting from patient-specific pre-operative 3D data, our method relies on GPU ray casting to compute a colored depth map, which assigns a predefined color to the first incidence of gradient magnitude value above a predefined threshold along the ray. The colored depth map values are carefully integrated into the X-Ray image while maintaining its original grey-scale intensities. The presented method was tested and analysed for three relevant clinical scenarios covering different anatomical aspects and targeting different levels of interventional expertise. Results demonstrate that improving depth perception of X-ray images has the potential to lead to safer and more efficient clinical interventions.

International Symposium on Mixed and Augmented Reality (2014)

Bibliography

- [1] N. Afsham, M. Najafi, P. Abolmaesumi, and R. Rohling. “A generalized correlation-based model for out-of-plane motion estimation in freehand ultrasound”. In: *Medical Imaging, IEEE Transactions on* 33.1 (2014), pp. 186–199 (cit. on p. 78).
- [2] J. E. Aldrich. “Basic physics of ultrasound imaging”. In: *Critical care medicine* 35.5 Suppl (2007), S131–7 (cit. on pp. 49, 63, 77).
- [3] C. M. Barnum and L. A. Palmer. “More Than a Feeling: Understanding the Desirability Factor in User Experience”. In: *CHI '10 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '10. Atlanta, Georgia, USA: ACM, 2010, pp. 4703–4716 (cit. on p. 17).
- [4] S. R. Barry. *Fixing my gaze: a scientist's journey into seeing in three dimensions*. Basic Books, 2009 (cit. on p. 45).
- [5] P. J. Basser and C. Pierpaoli. “Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI”. In: *Journal of Magnetic Resonance* 213.2 (2011), pp. 560–570 (cit. on p. 30).
- [6] N. Bell and M. Garland. *Efficient sparse matrix-vector multiplication on CUDA*. Tech. rep. Nvidia Technical Report NVR-2008-004, Nvidia Corporation, 2008 (cit. on p. 54).
- [7] J. Benedek and T. Miner. “Measuring Desirability: New methods for evaluating desirability in a usability lab setting”. In: *Proceedings of Usability Professionals Association 2003* (2002), pp. 8–12 (cit. on p. 16).
- [8] J. Beyer, M. Hadwiger, S. Wolfsberger, and K. Buhler. “High-Quality Multimodal Volume Rendering for Preoperative Planning of Neurosurgical Interventions”. In: *Visualization and Computer Graphics, IEEE Transactions on* 13.6 (2007), pp. 1696–1703 (cit. on p. 32).
- [9] Å. Birkeland, V. Soltészová, D. Hönigmann, et al. “The Ultrasound Visualization Pipeline - A Survey”. In: *CoRR* abs/1206.3975 (2012) (cit. on pp. 22, 34, 35).
- [10] I. Bitter, R. Van Uitert, I. Wolf, L. Ibanez, and J.-M. Kuhnigk. “Comparison of four freely available frameworks for image processing and visualization that use ITK”. In: *Visualization and Computer Graphics, IEEE Transactions on* 13.3 (2007), pp. 483–493 (cit. on p. 119).
- [11] L. M. Blaha, D. L. Arendt, and F. Mohd-Zaid. “More bang for your research buck: toward recommender systems for visual analytics”. In: *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*. ACM, 2014 (cit. on p. 17).
- [12] J. Bolz, P. Daniell, C. Riccio, G. Sellers, B. Merry, and J. Kessenich. “ARB Extension #137: ARB_shader_storage_buffer_object”. In: (2012) (cit. on p. 84).
- [13] R. Borgo, J. Kehrer, D. H. S. Chung, et al. “Glyph-based Visualization: Foundations, Design Guidelines, Techniques and Applications”. In: *Eurographics 2013 - State of the Art Reports*. Ed. by M. Sbert and L. Szirmay-Kalos. The Eurographics Association, 2012 (cit. on p. 65).

- [14] D. Borland and R. M. T. II. “Rainbow Color Map (Still) Considered Harmful”. In: *IEEE Computer Graphics and Applications* 27.2 (2007), pp. 14–17 (cit. on p. 46).
- [15] S. Born, M. Pfeifle, M. Markl, M. Gutberlet, and G. Scheuermann. “Visual analysis of cardiac 4D MRI blood flow using line predicates”. In: *Visualization and Computer Graphics, IEEE Transactions on* 19.6 (2013), pp. 900–912 (cit. on p. 94).
- [16] C. P. Botha and F. H. Post. “Hybrid Scheduling in the DeVIDE Dataflow Visualisation Environment.” In: *SimVis*. 2008, pp. 309–322 (cit. on p. 119).
- [17] H Brandl, A Gritzky, and M Haizinger. “3D ultrasound: a dedicated system”. In: *European radiology* 3.9 (1999) (cit. on pp. 34, 35).
- [18] R. Brecheisen, A. Vilanova, B. Platel, and B. T. H. Romeny. “Parameter sensitivity visualization for DTI fiber tracking”. In: *Visualization and Computer Graphics, IEEE Transactions on* 15.6 (2009), pp. 1441–1448 (cit. on p. 63).
- [19] M. Bro-Nielsen and C. Gramkow. “Fast fluid registration of medical images”. In: *Visualization in Biomedical Computing*. Springer. 1996, pp. 265–276 (cit. on p. 80).
- [20] E. Brookner. *Tracking and Kalman filtering made easy*. A Wiley-Interscience publication. Wiley, 1998 (cit. on p. 55).
- [21] A. Brooks, J. Connolly, and O. Chan. *Ultrasound in emergency care*. John Wiley & Sons, 2008 (cit. on p. 50).
- [22] M. A. Brown and R. C. Semelka. *MRI: basic principles and applications*. John Wiley & Sons, 2011 (cit. on p. 21).
- [23] R. Brown. “Animated Visual Vibrations As an Uncertainty Visualisation Technique”. In: *Proceedings of the 2Nd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia*. GRAPHITE ’04. Singapore: ACM, 2004, pp. 84–89 (cit. on p. 65).
- [24] E. Bruce Goldstein. *Sensation and Perception, Ninth Edition*. Cengage Learning, 2014 (cit. on pp. 39, 41, 43, 45).
- [25] S. Bruckner and M. E. Gröller. “Instant Volume Visualization using Maximum Intensity Difference Accumulation”. In: *Computer Graphics Forum* 28.3 (2009), pp. 775–782 (cit. on p. 26).
- [26] S. Bruckner, S. Grimm, A. Kanitsar, and M. E. Gröller. “Illustrative context-preserving volume rendering”. In: *EuroVis*. 2005, pp. 69–76 (cit. on p. 93).
- [27] M. Burns, M. Haidacher, W. Wein, I. Viola, and E. Groeller. “Feature Emphasis and Contextual Cutaways for Multimodal Medical Visualization.” In: *EuroVis*. Vol. 7. Citeseer. 2007, pp. 275–282 (cit. on p. 38).
- [28] J. Butter, T. H. Grant, M. Egan, et al. “Does ultrasound training boost Year 1 medical student competence and confidence when learning abdominal examination?” In: *Medical Education* 41.9 (2007), pp. 843–848 (cit. on p. 63).
- [29] J. J. Caban, A. Joshi, and P. Nagy. “Rapid development of medical imaging tools with open-source libraries”. In: *Journal of Digital Imaging* 20.1 (2007), pp. 83–93 (cit. on p. 119).
- [30] F. Capra. *Learning from Leonardo: Decoding the Notebooks of a Genius*. Berrett-Koehler Publishers, 2013 (cit. on p. 11).
- [31] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in information visualization: using vision to think*. Morgan Kaufmann, 1999 (cit. on p. 13).
- [32] S. Carpendale. “Evaluating Information Visualizations”. English. In: *Information Visualization*. Ed. by A. Kerren, J. Stasko, J.-D. Fekete, and C. North. Vol. 4950. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2008, pp. 19–45 (cit. on p. 15).

- [33] M. Cercignani and M. A. Horsfield. “The physical basis of diffusion-weighted MRI”. In: *Journal of the Neurological Sciences* 186, Supplement 1 (2001), S11–S14 (cit. on p. 30).
- [34] P. Chatelain, A. Krupa, and N. Navab. “Optimization of ultrasound image quality via visual servoing”. In: *IEEE Int. Conf. on Robotics and Automation, ICRA’15*. 2015 (cit. on p. 53).
- [35] C. Chen. “Information visualization”. In: *Wiley Interdisciplinary Reviews: Computational Statistics* 2.4 (2010), pp. 387–403 (cit. on pp. 11, 12).
- [36] J.-F. Chen, J. B. Fowlkes, P. L. Carson, and J. M. Rubin. “Determination of scan-plane motion using speckle decorrelation: theoretical considerations and initial test”. In: (1997) (cit. on p. 78).
- [37] M. Chen, W. Lu, Q. Chen, K. J. Ruchala, and G. H. Olivera. “A simple fixed-point approach to invert a deformation fielda”. In: *Medical physics* 35.1 (2008), pp. 81–88 (cit. on p. 80).
- [38] R. S. Cobbold. *Foundations of biomedical ultrasound*. Oxford University, 2007 (cit. on p. 50).
- [39] D. M. Coppola, L. E. White, D. Fitzpatrick, and D. Purves. “Unequal representation of cardinal and oblique contours in ferret visual cortex”. In: *Proceedings of the National Academy of Sciences* 95.5 (1998), pp. 2621–2623 (cit. on p. 42).
- [40] J. Cox, V. Gerth, and W. Holmes. “The Programmed Console: An Aid to the Radiologist in Treatment Planning”. In: *First international Conference on The Use of Computers In Therapeutic Radiology, Cambridge, England*. 1966 (cit. on p. 19).
- [41] C. Crassin. “GigaVoxels: A Voxel-Based Rendering Pipeline For Efficient Exploration Of Large And Detailed Scenes”. PhD thesis. UNIVERSITE DE GRENOBLE, 2011 (cit. on p. 26).
- [42] C. Crassin, F. Neyret, S. Lefebvre, and E. Eisemann. “GigaVoxels: Ray-guided Streaming for Efficient and Detailed Voxel Rendering”. In: *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games. I3D ’09*. Boston, Massachusetts: ACM, 2009, pp. 15–22 (cit. on p. 26).
- [43] J. Curie and P. Curie. “Développement, par pression, de l’électricité polaire dans les cristaux hémihédres à faces inclinées”. In: *Comptes Rendus* 91 (1880), pp. 294–295 (cit. on p. 50).
- [44] S. Diepenbrock, C. Schulte zu Berge, K. H. Hinrichs, L. Wachsmuth, and C. Faber. *DTI Visualization Using the Voreen Framework*. Poster at ESMRMB Congress. 2011 (cit. on p. 136).
- [45] Q. Duan, E. D. Angelini, and A. F. Laine. “Surface function actives”. In: *Journal of Visual Communication and Image Representation* 20.7 (2009), pp. 478–490 (cit. on p. 96).
- [46] M. Eiber, S. G. Nekolla, T. Maurer, G. Weirich, H.-J. Wester, and M. Schwaiger. “68Ga-PSMA PET/MR with multimodality image analysis for primary prostate cancer”. In: *Abdominal imaging* (2014), pp. 1–3 (cit. on p. 132).
- [47] K. Engel, M. Hadwiger, J. Kniss, C. Rezk-Salama, and D. Weiskopf. *Real-time volume graphics*. CRC Press, 2006 (cit. on p. 25).
- [48] M. Eyer, M. Brandestini, D. Phillips, and D. Baker. “Color digital echo/Doppler image presentation”. In: *Ultrasound in medicine & biology* 7.1 (1981), pp. 21–31 (cit. on p. 22).
- [49] R. Fattal and D. Lischinski. “Variational classification for visualization of 3D ultrasound data”. In: *Visualization, 2001. VIS ’01. Proceedings*. 2001, pp. 403–410 (cit. on p. 36).
- [50] A. Fedorov, R. Beichel, J. Kalpathy-Cramer, et al. “3D Slicer as an image computing platform for the Quantitative Imaging Network”. In: *Magnetic resonance imaging* 30.9 (2012), pp. 1323–1341 (cit. on p. 118).
- [51] T. Fogal and J. Krüger. “Tuvok, an Architecture for Large Scale Volume Rendering.” In: *Proceedings of the 15th International Workshop on Vision, Modeling, and Visualization*. 2010, pp. 139–146 (cit. on p. 118).
- [52] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. “Multiscale vessel enhancement filtering”. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI’98*. Springer, 1998, pp. 130–137 (cit. on p. 95).

- [53] R. M. Friedhoff and T. Kiely. “The eye of the beholder”. In: *Computer Graphics World* 13.8 (1990), p. 46 (cit. on p. 7).
- [54] M. Friendly. *Milestones in the history of thematic cartography, statistical graphics, and data visualization*. 1995 (cit. on p. 10).
- [55] M. Friendly and D. J. Denis. “Milestones in the history of thematic cartography, statistical graphics, and data visualization”. In: (2001) (cit. on p. 9).
- [56] R. Fuchs and H. Hauser. “Visualization of Multi-Variate Scientific Data”. In: *Computer Graphics Forum* 28.6 (2009), pp. 1670–1690 (cit. on p. 31).
- [57] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design patterns: elements of reusable object-oriented software*. Pearson Education, 1994 (cit. on p. 119).
- [58] R. Gasteiger, M. Neugebauer, O. Beuing, and B. Preim. “The FLOWLENS: A Focus-and-Context Visualization Approach for Exploration of Blood Flow in Cerebral Aneurysms”. In: *Visualization and Computer Graphics, IEEE Transactions on* 17.12 (2011), pp. 2183–2192 (cit. on p. 29).
- [59] L. Grady. “Random Walks for Image Segmentation”. In: *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 28.11 (2006), pp. 1768–1783 (cit. on p. 52).
- [60] N. Gupta, M. Swamy, and E. Plotkin. “Despeckling of medical ultrasound images using data and rate adaptive lossy compression”. In: *Medical Imaging, IEEE Transactions on* 24.6 (2005), pp. 743–754 (cit. on p. 35).
- [61] M. Haidacher, S. Bruckner, A. Kanitsar, and M. E. Gröller. “Information-based Transfer Functions for Multimodal Visualization”. In: *Proceedings of the First Eurographics Conference on Visual Computing for Biomedicine*. EG VCBM’08. Delft, The Netherlands: Eurographics Association, 2008, pp. 101–108 (cit. on p. 32).
- [62] P. K. J. Han, W. M. P. Klein, and N. K. Arora. “Varieties of Uncertainty in Health Care: A Conceptual Taxonomy”. In: *Medical Decision Making* 31.6 (2011), pp. 828–838 (cit. on p. 62).
- [63] L. Hao, C. Healey, and S. Bass. “Effective Visualization of Temporal Ensembles”. In: *Visualization and Computer Graphics, IEEE Transactions on* 22.1 (2016), pp. 787–796 (cit. on p. 63).
- [64] C. Hennersperger, D. Mateus, M. Baust, and N. Navab. “A Quadratic Energy Minimization Framework for Signal Loss Estimation from Arbitrarily Sampled Ultrasound Data”. English. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2014*. Ed. by P. Golland, N. Hata, C. Barillot, J. Hornegger, and R. Howe. Vol. 8674. Lecture Notes in Computer Science. Springer International Publishing, 2014, pp. 373–380 (cit. on p. 51).
- [65] G. T. Herman. *Fundamentals of computerized tomography: image reconstruction from projections*. Springer Science & Business Media, 2009 (cit. on p. 20).
- [66] K. Hofschien, T. Geissler, N. Rieke, C. Schulte zu Berge, N. Navab, and S. Demirci. “Image Descriptors in Angiography”. In: *Proc. of Workshop Bildverarbeitung für die Medizin (BVM 2016)*. Springer, 2016, pp. 283–288 (cit. on p. 135).
- [67] W. F. Holmes. “External Beam Treatment-Planning with the Programmed Console”. In: *Radiology* 94.2 (1970), pp. 391–400 (cit. on p. 19).
- [68] D. Honigmann, J. Ruisz, and C. Haider. “Adaptive Design of a Global Opacity Transfer Function for Direct Volume Rendering of Ultrasound Data”. In: *Proceedings of the 14th IEEE Visualization 2003 (VIS’03)*. VIS ’03. IEEE Computer Society, 2003, pp. 64– (cit. on p. 36, 37).
- [69] G. N. Hounsfield. “Computerized transverse axial scanning (tomography): Part 1. Description of system”. In: *The British Journal of Radiology* 46.552 (1973), pp. 1016–1022 (cit. on p. 20).
- [70] R. J. Housden, A. H. Gee, G. M. Treece, and R. W. Prager. “Sensorless reconstruction of unconstrained freehand 3D ultrasound data”. In: *Ultrasound in medicine & biology* 33.3 (2007), pp. 408–419 (cit. on p. 78).

- [71] D. Jönsson, M. Falk, and A. Ynnerman. “Intuitive Exploration of Volumetric Data Using Dynamic Galleries”. In: *Visualization and Computer Graphics, IEEE Transactions on* 22.1 (2016), pp. 896–905 (cit. on p. 28).
- [72] D. Jönsson, E. Sundén, A. Ynnerman, and T. Ropinski. “A Survey of Volumetric Illumination Techniques for Interactive Volume Rendering”. In: *Computer Graphics Forum* 33.1 (2014), pp. 27–51 (cit. on p. 26).
- [73] A. Kanitsar, R. Wegenkittl, D. Fleischmann, and M. Groller. “Advanced curved planar reformation: flattening of vascular structures”. In: *Visualization, 2003. VIS 2003. IEEE*. 2003, pp. 43–50 (cit. on p. 25).
- [74] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and M. Groller. “CPR - curved planar reformation”. In: *Visualization, 2002. VIS 2002. IEEE*. 2002, pp. 37–44 (cit. on p. 25).
- [75] D. Kao, J. Dungan, and A. Pang. “Visualizing 2D probability distributions from EOS satellite image-derived data sets: a case study”. In: *Visualization, 2001. VIS '01. Proceedings*. 2001, pp. 457–589 (cit. on p. 65).
- [76] A. Karamalis, W. Wein, O. Kutter, and N. Navab. “Fast hybrid freehand ultrasound volume reconstruction”. In: *SPIE Medical Imaging*. International Society for Optics and Photonics. 2009, pp. 726114–726114 (cit. on p. 76).
- [77] A. Karamalis, W. Wein, T. Klein, and N. Navab. “Ultrasound confidence maps using random walks”. In: *Medical Image Analysis* 16.6 (2012), pp. 1101–1112 (cit. on pp. 49, 51, 53, 57, 96, 109).
- [78] A. Katouzian, E. D. Angelini, B. Sturm, and A. F. Laine. “Brushlet segmentation for automatic detection of lumen borders in ivus images: A comparison study”. In: *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*. IEEE. 2012, pp. 242–245 (cit. on p. 96).
- [79] M. Kersten-Oertel, S.-S. Chen, and D. Collins. “An Evaluation of Depth Enhancing Perceptual Cues for Vascular Volume Visualization in Neurosurgery”. In: *Visualization and Computer Graphics, IEEE Transactions on* 20.3 (2014), pp. 391–403 (cit. on p. 26).
- [80] G. Kindlmann and C.-F. Westin. “Diffusion Tensor Visualization with Glyph Packing”. In: *Visualization and Computer Graphics, IEEE Transactions on* 12.5 (2006), pp. 1329–1336 (cit. on pp. 30, 65).
- [81] D. Kleffner and V. Ramachandran. “On the perception of shape from shading”. English. In: *Perception & Psychophysics* 52.1 (1992), pp. 18–36 (cit. on p. 42).
- [82] T. Klein, M. Hansson, and N. Navab. “Modeling of multi-view 3D freehand radio frequency ultrasound”. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2012*. Springer, 2012, pp. 422–429 (cit. on p. 77).
- [83] J. Kniss, G. Kindlmann, and C. Hansen. “Multidimensional transfer functions for interactive volume rendering”. In: *Visualization and Computer Graphics, IEEE Transactions on* 8.3 (2002), pp. 270–285 (cit. on p. 27).
- [84] J. L. Knudtson, J. M. Dort, S. D. Helmer, and R. S. Smith. “Surgeon-performed ultrasound for pneumothorax in the trauma suite”. In: *The Journal of trauma* 56.3 (2004), pp. 527–530 (cit. on p. 64).
- [85] V. D. Köchli and B. Marincek. *Wie funktioniert MRI?: eine Einführung in Physik und Funktionsweise der Magnetresonanzbildgebung*. Springer-Verlag, 2013 (cit. on p. 21).
- [86] B. Kohler, R. Gasteiger, U. Preim, H. Theisel, M. Gutberlet, and B. Preim. “Semi-automatic vortex extraction in 4D PC-MRI cardiac blood flow data using line predicates”. In: *Visualization and Computer Graphics, IEEE Transactions on* 19.12 (2013), pp. 2773–2782 (cit. on p. 94).

- [87] J. Kretschmer, G. Soza, C. Tietjen, M. Suehling, B. Preim, and M. Stamminger. “ADR - Anatomy-Driven Reformation”. In: *Visualization and Computer Graphics, IEEE Transactions on* 20.12 (2014), pp. 2496–2505 (cit. on p. 25).
- [88] K. Krissian, C.-F. Westin, R. Kikinis, and K. G. Vosburgh. “Oriented speckle reducing anisotropic diffusion”. In: *Image Processing, IEEE Transactions on* 16.5 (2007), pp. 1412–1424 (cit. on p. 35).
- [89] J. Krüger and R. Westermann. “Acceleration Techniques for GPU-based Volume Rendering”. In: *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*. VIS '03. IEEE Computer Society, 2003, pp. 38– (cit. on p. 25).
- [90] J. Krüger, J. Schneider, and R. Westermann. “Clearview: An interactive context preserving hotspot visualization technique”. In: *Visualization and Computer Graphics, IEEE Transactions on* 12.5 (2006), pp. 941–948 (cit. on p. 93).
- [91] O. Kutter, A. Aichert, C. Bichlmeier, et al. “Real-time volume rendering for high quality visualization in augmented reality”. In: *International workshop on Augmented environments for Medical Imaging including Augmented Reality in Computer-aided Surgery (AMI-ARCS)*. 2008 (cit. on p. 93).
- [92] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. “Empirical Studies in Information Visualization: Seven Scenarios”. In: *Visualization and Computer Graphics, IEEE Transactions on* 18.9 (2012), pp. 1520–1536 (cit. on p. 15).
- [93] K. Lawonn, S. Glaber, A. Vilanova, B. Preim, and T. Isenberg. “Occlusion-free Blood Flow Animation with Wall Thickness Visualization”. In: *Visualization and Computer Graphics, IEEE Transactions on* 22.1 (2016), pp. 728–737 (cit. on p. 29).
- [94] K. Lawonn, N. Smit, B. Preim, and A. Vilanova. “Illustrative Multi-volume Rendering for PET/CT Scans”. In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. Ed. by K. Bühler, L. Linsen, and N. W. John. The Eurographics Association, 2015 (cit. on p. 32).
- [95] S. M. Lehar. *The world in your head: A gestalt view of the mechanism of conscious experience*. Psychology Press, 2003 (cit. on pp. 41, 42).
- [96] E. Lengyel. *Game Engine Gems 2*. CRC Press, 2011 (cit. on p. 119).
- [97] M. Levoy. “Display of surfaces from volume data”. In: *Computer Graphics and Applications, IEEE* 8.3 (1988), pp. 29–37 (cit. on p. 27).
- [98] F. Lindemann and T. Ropinski. “About the Influence of Illumination Models on Image Comprehension in Direct Volume Rendering”. In: *Visualization and Computer Graphics, IEEE Transactions on* 17.12 (2011), pp. 1922–1931 (cit. on p. 26).
- [99] Y. Livnat. “2 - Accelerated Isosurface Extraction Approaches”. In: *Visualization Handbook*. Ed. by C. D. H. R. Johnson. Butterworth-Heinemann, 2005, pp. 39–IV (cit. on p. 25).
- [100] W. E. Lorensen and H. E. Cline. “Marching Cubes: A High Resolution 3D Surface Construction Algorithm”. In: *SIGGRAPH Comput. Graph.* 21.4 (1987), pp. 163–169 (cit. on p. 25).
- [101] C. Lundström, P. Ljung, A. Persson, and A. Ynnerman. “Uncertainty Visualization in Medical Volume Rendering Using Probabilistic Animation”. In: *Visualization and Computer Graphics, IEEE Transactions on* 13.6 (2007), pp. 1648–1655 (cit. on pp. 62, 65).
- [102] A. MacEachren, R. Roth, J. O’Brien, B. Li, D. Swingley, and M. Gahegan. “Visual Semiotics & Uncertainty Visualization: An Empirical Study”. In: *Visualization and Computer Graphics, IEEE Transactions on* 18.12 (2012), pp. 2496–2505 (cit. on pp. 65, 68).
- [103] J. Marks, B. Andalman, P. A. Beardsley, et al. “Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation”. In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '97. ACM Press/Addison-Wesley Publishing Co., 1997, pp. 389–400 (cit. on p. 28).

- [104] P. Matthies, J. Gardiazabal, A. Okur, J. Vogel, T. Lasser, and N. Navab. “Mini gamma cameras for intra-operative nuclear tomographic reconstruction”. In: *Medical image analysis* 18.8 (2014), pp. 1329–1336 (cit. on p. 23).
- [105] N. Max. “Optical models for direct volume rendering”. In: *Visualization and Computer Graphics, IEEE Transactions on* 1.2 (1995), pp. 99–108 (cit. on p. 26).
- [106] B. H. McCormick. “Visualization in Scientific Computing”. In: *SIGBIO Newsl.* 10.1 (1988), pp. 15–21 (cit. on p. 19).
- [107] “Definition of Visualization”. In: *SIGGRAPH Comput. Graph.* 21.6 (1987). Ed. by B. H. McCormick, T. A. DeFanti, and M. D. Brown, pp. 3–3 (cit. on p. 7).
- [108] M. McShaffry. *Game coding complete*. Cengage Learning, 2012 (cit. on pp. 119, 120, 125).
- [109] T. Merčun. “Evaluation of information visualization techniques: analysing user experience with reaction cards”. In: *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*. ACM, 2014 (cit. on p. 17).
- [110] C. Metz. “Digitally Reconstructed Radiographs”. MA thesis. Utrecht University, 2005 (cit. on p. 26).
- [111] J. Meyer-Spradow, L. Stegger, C. Doring, T. Ropinski, and K. Hinrichs. “Glyph-Based SPECT Visualization for the Diagnosis of Coronary Artery Disease”. In: *Visualization and Computer Graphics, IEEE Transactions on* 14.6 (2008), pp. 1499–1506 (cit. on p. 33).
- [112] J. Meyer-Spradow, T. Ropinski, J. Mensmann, and K. Hinrichs. “Voreen: A rapid-prototyping environment for ray-casting-based volume visualizations”. In: *Computer Graphics and Applications, IEEE* 29.6 (2009), pp. 6–13 (cit. on pp. 118, 121).
- [113] O. Michailovich and A. Tannenbaum. “Despeckling of medical ultrasound images”. In: *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on* 53.1 (2006), pp. 64–78 (cit. on p. 35).
- [114] N. Milickovic, D. Baltas, S. Giannouli, M. Lahanas, and N. Zamboglou. “CT imaging based digitally reconstructed radiographs and their application in brachytherapy”. In: *Physics in Medicine and Biology* 45.10 (2000), p. 2787 (cit. on p. 26).
- [115] M. Mon-Williams and J. R. Tresilian. “Some Recent Studies on the Extraretinal Contribution to Distance Perception”. In: *Perception* 28.2 (1999), pp. 167–181 (cit. on p. 43).
- [116] K. Moreland. “Diverging Color Maps for Scientific Visualization”. English. In: *Advances in Visual Computing*. Ed. by G. Bebis, R. Boyle, B. Parvin, et al. Vol. 5876. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2009, pp. 92–103 (cit. on p. 46).
- [117] S. Mori and P. C. M. van Zijl. “Fiber tracking: principles and strategies – a technical review”. In: *NMR in Biomedicine* 15.7-8 (2002), pp. 468–480 (cit. on p. 30).
- [118] O. Morozov, M. Unser, and P. Hunziker. “Reconstruction of large, irregularly sampled multidimensional images. A tensor-based approach”. In: *Medical Imaging, IEEE Transactions on* 30.2 (2011), pp. 366–374 (cit. on p. 77).
- [119] F. de Moura Pinto et al. “Importance-aware composition for illustrative volume rendering”. In: *Graphics, Patterns and Images (SIBGRAPI), 2010 23rd SIBGRAPI Conference on*. IEEE, 2010, pp. 134–141 (cit. on pp. 93, 99, 100).
- [120] T. Munzner. “A Nested Model for Visualization Design and Validation”. In: *Visualization and Computer Graphics, IEEE Transactions on* 15.6 (2009), pp. 921–928 (cit. on pp. 15, 16).
- [121] T. Munzner. “Process and Pitfalls in Writing Information Visualization Research Papers”. English. In: *Information Visualization*. Ed. by A. Kerren, J. Stasko, J.-D. Fekete, and C. North. Vol. 4950. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2008, pp. 134–153 (cit. on p. 15).

- [122] J. A. Noble. “Ultrasound image segmentation and tissue characterization”. In: *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine* 224.2 (2010), pp. 307–316 (cit. on pp. 49, 51).
- [123] J. A. Noble, N. Navab, and H. Becher. “Ultrasonic image analysis and image-guided interventions”. In: *Interface Focus* 1.4 (2011), pp. 673–685 (cit. on pp. 23, 49, 51, 75).
- [124] J. Noble and D. Boukerroui. “Ultrasound image segmentation: a survey”. In: *Medical Imaging, IEEE Transactions on* 25.8 (2006), pp. 987–1010 (cit. on pp. 35, 112).
- [125] H. Obermaier, K. Joy, et al. “Future challenges for ensemble visualization”. In: *Computer Graphics and Applications, IEEE* 34.3 (2014), pp. 8–11 (cit. on p. 63).
- [126] S. Oeltze, A. Hennemuth, S. Glaßer, C. Kühnel, and B. Preim. “Glyph-Based Visualization of Myocardial Perfusion Data and Enhancement with Contractility and Viability Information”. In: *Proceedings of the First Eurographics Workshop on Visual Computing for Biomedicine*. Ed. by C. Botha, G. Kindlmann, W. Niessen, and B. Preim. Eurographics Association, 2008, pp. 11–20 (cit. on p. 33).
- [127] A. Oliva and A. Torralba. “The role of context in object recognition”. In: *Trends in cognitive sciences* 11.12 (2007), pp. 520–527 (cit. on p. 42).
- [128] J. Ophir, S. K. Alam, B. S. Garra, et al. “Elastography: imaging the elastic properties of soft tissues with ultrasound”. In: *Journal of Medical Ultrasonics* 29.4 (2002), pp. 155–171 (cit. on p. 22).
- [129] D. Patel, M. Haidacher, J.-P. Balabanian, and E. Groller. “Moment curves”. In: *Visualization Symposium, 2009. PacificVis '09. IEEE Pacific*. 2009, pp. 201–208 (cit. on p. 28).
- [130] R. van Pelt, J. Olivan Bescos, M. Breeuwer, et al. “Exploration of 4D MRI Blood Flow using Stylistic Visualization”. In: *Visualization and Computer Graphics, IEEE Transactions on* 16.6 (2010), pp. 1339–1347 (cit. on p. 29).
- [131] R. van Pelt, J. Olivan Bescos, M. Breeuwer, et al. “Interactive Virtual Probing of 4D MRI Blood-Flow”. In: *Visualization and Computer Graphics, IEEE Transactions on* 17.12 (2011), pp. 2153–2162 (cit. on p. 28).
- [132] T. Pfaffelmoser, M. Reitingner, and R. Westermann. “Visualizing the positional and geometrical variability of isosurfaces in uncertain scalar fields”. In: *Computer Graphics Forum*. Vol. 30. 3. Wiley Online Library. 2011, pp. 951–960 (cit. on p. 62).
- [133] C. Plaisant. “The Challenge of Information Visualization Evaluation”. In: *Proceedings of the Working Conference on Advanced Visual Interfaces. AVI '04*. Gallipoli, Italy: ACM, 2004, pp. 109–116 (cit. on p. 15).
- [134] K. N. Plataniotis and A. N. Venetsanopoulos. *Color image processing and applications*. Springer, 2000 (cit. on p. 67).
- [135] K. Pöthkow and H.-C. Hege. “Positional uncertainty of isocontours: Condition analysis and probabilistic measures”. In: *Visualization and Computer Graphics, IEEE Transactions on* 17.10 (2011), pp. 1393–1406 (cit. on p. 62).
- [136] J.-S. Prassni, T. Ropinski, J. Mensmann, and K. Hinrichs. “Shape-based transfer functions for volume visualization”. In: *Pacific Visualization Symposium (PacificVis), 2010 IEEE*. 2010, pp. 9–16 (cit. on p. 27).
- [137] J. S. Praßni, T. Ropinski, and K. Hinrichs. “Uncertainty-aware guided volume segmentation”. In: *Visualization and Computer Graphics, IEEE Transactions on* 16.6 (2010), pp. 1358–1365 (cit. on p. 63).
- [138] B. Preim and C. Botha. *Visual Computing for Medicine: Theory, Algorithms, and Applications*. The Morgan Kaufmann Series in Computer Graphics. Elsevier Science, 2013 (cit. on p. 19).

- [139] B. Preim, A. Baer, D. Cunningham, T. Isenberg, and T. Ropinski. “A Survey of Perceptually Motivated 3D Visualization of Medical Image Data”. In: *Computer Graphics Forum*. Vol. 35. 3. 2016 (cit. on p. 19).
- [140] P. Rautek, S. Bruckner, and M. E. Gröller. “Interaction-dependent Semantics for Illustrative Volume Rendering”. In: *Proceedings of the 10th Joint Eurographics / IEEE - VGTC Conference on Visualization*. EuroVis’08. Eindhoven, The Netherlands: Eurographics Association, 2008, pp. 847–854 (cit. on pp. 28, 100).
- [141] M. Rautenhaus, C. M. Grams, A. Schäfler, and R. Westermann. “Visualization of 3D ensemble weather forecasts to predict uncertain warm conveyor belt situations”. In: *EGU General Assembly Conference Abstracts*. Vol. 17. 2015, p. 11908 (cit. on p. 63).
- [142] C. Rezk-Salama, M. Keller, and P. Kohlmann. “High-Level User Interfaces for Transfer Function Design with Semantics”. In: *Visualization and Computer Graphics, IEEE Transactions on* 12.5 (2006), pp. 1021–1028 (cit. on pp. 27, 28, 91).
- [143] C. Rezk-Salama, P. Hastreiter, J. Scherer, and G. Greiner. “Automatic Adjustment of Transfer Functions for 3D Volume Visualization”. In: *In Proc. Workshop Vision, Modeling, and Visualization (VMV. a, 2000)*, pp. 357–364 (cit. on p. 28).
- [144] C. Rieder, T. Kroeger, C. Schumann, and H. K. Hahn. “GPU-based real-time approximation of the ablation zone for radiofrequency ablation”. In: *Visualization and Computer Graphics, IEEE Transactions on* 17.12 (2011), pp. 1812–1821 (cit. on p. 62).
- [145] A. Rind, W. Aigner, M. Wagner, S. Miksch, and T. Lammarsch. “User Tasks for Evaluation”. In: *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*. ACM. 2014 (cit. on p. 15).
- [146] G. Ristovski, T. Preusser, H. K. Hahn, and L. Linsen. “Uncertainty in medical visualization: Towards a taxonomy”. In: *Computers & Graphics* 39 (2014), pp. 60–73 (cit. on p. 62).
- [147] F. Ritter, T. Boskamp, A. Homeyer, et al. “Medical image analysis”. In: *Pulse, IEEE* 2.6 (2011), pp. 60–70 (cit. on p. 119).
- [148] W. C. Röntgen. “Ueber eine neue Art von Strahlen”. In: *Annalen der Physik* 300.1 (1898), pp. 1–11 (cit. on p. 20).
- [149] R. Rohling, A. Gee, and L. Berman. “A comparison of freehand three-dimensional ultrasound reconstruction techniques”. In: *Medical image analysis* 3.4 (1999), pp. 339–359 (cit. on p. 77).
- [150] T. Ropinski, J. Prassni, F. Steinicke, and K. Hinrichs. “Stroke-based Transfer Function Design”. In: *Proceedings of the Fifth Eurographics / IEEE VGTC Conference on Point-Based Graphics*. SPBG’08. Los Angeles, CA: Eurographics Association, 2008, pp. 41–48 (cit. on pp. 28, 101).
- [151] T. Ropinski and B. Preim. “Taxonomy and Usage Guidelines for Glyph-based Medical Visualization”. In: *Simulation and Visualization 2008 (SimVis 2008), 18-29 Februar 2008, Magdeburg*. Ed. by H. Hauser, S. Straßburger, and H. Theisel. SCS Publishing House e.V., 2008, pp. 121–138 (cit. on p. 33).
- [152] T. Ropinski, M. Specht, J. Meyer-spradow, K. Hinrichs, and B. Preim. “Surface glyphs for visualizing multimodal volume data”. In: *In Proceedings of the 12th International Fall Workshop on Vision, Modeling, and Visualization (VMV07. 2007)*, pp. 3–12 (cit. on p. 33).
- [153] T. Ropinski, F. Steinicke, and K. Hinrichs. “Visually Supporting Depth Perception in Angiography Imaging”. English. In: *Smart Graphics*. Ed. by A. Butz, B. Fisher, A. Krüger, and P. Olivier. Vol. 4073. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2006, pp. 93–104 (cit. on pp. 37, 46).
- [154] G. Sakas, L.-A. Schreyer, and M. Grimm. “Preprocessing and volume rendering of 3D ultrasonic data”. In: *Computer Graphics and Applications, IEEE* 15.4 (1995), pp. 47–54 (cit. on pp. 34, 36).

- [155] T. Salzbrunn, C. Garth, G. Scheuermann, and J. Meyer. “Pathline predicates and unsteady flow structures”. In: *The Visual Computer* 24.12 (2008), pp. 1039–1051 (cit. on p. 94).
- [156] J. M. Sanches and J. S. Marques. “A multiscale algorithm for three-dimensional free-hand ultrasound”. In: *Ultrasound in medicine & biology* 28.8 (2002), pp. 1029–1040 (cit. on p. 77).
- [157] J. Sanyal, S. Zhang, G. Bhattacharya, P. Amburn, and R. Moorhead. “A User Study to Compare Four Uncertainty Visualization Methods for 1D and 2D Datasets”. In: *Visualization and Computer Graphics, IEEE Transactions on* 15.6 (2009), pp. 1209–1218 (cit. on p. 65).
- [158] K. Scanlan. “Sonographic artifacts and their origins”. In: *AJR. American journal of roentgenology* 156.6 (1991), 1267–1272 (cit. on pp. 49, 63).
- [159] A. Schoch, B. Fuerst, F. Achilles, S. Demirci, and N. Navab. “A Lightweight and Portable Communication Framework for Multimodal Image-Guided Therapy”. In: (2013) (cit. on p. 130).
- [160] C. Schulte zu Berge, A. Grunau, H. Mahmud, and N. Navab. *CAMPVis – A Game Engine-inspired Research Framework for Medical Imaging and Visualization*. Tech. rep. Technische Universität München, 2014 (cit. on pp. 117, 136).
- [161] C. Schulte zu Berge, A. Kapoor, and N. Navab. “Orientation-driven ultrasound compounding using uncertainty information”. In: *Information Processing in Computer-Assisted Interventions*. Springer, 2014, pp. 236–245 (cit. on pp. 76, 78, 136).
- [162] C. Schulte zu Berge, M. Baust, A. Kapoor, and N. Navab. “Predicate-based Focus-and-Context Visualization for 3D Ultrasound”. In: *Visualization and Computer Graphics, IEEE Transactions on* 20.12 (2014), pp. 2379–2387 (cit. on pp. 92, 135).
- [163] C. Schulte zu Berge, D. Declara, C. Hennersperger, M. Baust, and N. Navab. “Real-time Uncertainty Visualization for B-Mode Ultrasound”. English. In: *Scientific Visualization 2015, Proceedings on* (2015) (cit. on pp. 49, 61, 135).
- [164] C. Schulte zu Berge, J. Weiss, and N. Navab. “Schematic Electrode Map for Navigation in Neuro Data Sets”. English. In: *Eurographics Workshop on Visual Computing for Biology and Medicine, VCBM 2015, Chester, UK*. Eurographics Association, 2015 (cit. on p. 135).
- [165] P. Sereda, A. Bartroli, I. W. O. Serlie, and F. Gerritsen. “Visualization of boundaries in volumetric data sets using LH histograms”. In: *Visualization and Computer Graphics, IEEE Transactions on* 12.2 (2006), pp. 208–218 (cit. on p. 27).
- [166] A. Shah, O. Zettinig, T. Maurer, et al. “An open source multimodal image-guided prostate biopsy framework”. In: *Clinical Image-Based Procedures. Translational Research in Medical Imaging*. Springer, 2014, pp. 1–8 (cit. on pp. 131, 132, 135).
- [167] M. P. Shankar. “Ultrasonic tissue characterization using a generalized Nakagami model”. In: *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on* 48.6 (2001), pp. 1716–1720 (cit. on p. 112).
- [168] D. Shepard. “A two-dimensional interpolation function for irregularly-spaced data”. In: *Proceedings of the 1968 23rd ACM national conference*. ACM, 1968, pp. 517–524 (cit. on p. 82).
- [169] B. Shneiderman. “The eyes have it: a task by data type taxonomy for information visualizations”. In: *Visual Languages, 1996. Proceedings., IEEE Symposium on*. 1996, pp. 336–343 (cit. on p. 8).
- [170] O. V. Solberg, F. Lindseth, H. Torp, R. E. Blake, and T. A. N. Hernes. “Freehand 3D ultrasound reconstruction algorithms—a review”. In: *Ultrasound in medicine & biology* 33.7 (2007), pp. 991–1009 (cit. on pp. 76, 77).
- [171] V. Soltészova, A. Birkeland, I. Viola, and S. Bruckner. “Visibility-Driven Processing of Streaming Volume Data”. In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. Ed. by I. Viola, K. Buehler, and T. Ropinski. The Eurographics Association, 2014 (cit. on p. 36).

- [172] V. Soltészová, L. E. S. Helljesen, W. Wein, O. H. Gilja, and I. Viola. “Lowest-Variance Streamlines for Filtering of 3D Ultrasound”. In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. Ed. by T. Ropinski, A. Ynnerman, C. Botha, and J. Roerdink. The Eurographics Association, 2012 (cit. on p. 35).
- [173] J. Stasko. “Value-driven Evaluation of Visualizations”. In: *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*. ACM. 2014 (cit. on p. 16).
- [174] E. Steen and B. Olstad. “Volume rendering of 3D medical ultrasound data using direct feature mapping”. In: *Medical Imaging, IEEE Transactions on* 13.3 (1994), pp. 517–525 (cit. on p. 36).
- [175] R. A. Steenblik. “The Chromostereoscopic Process: A Novel Single Image Stereoscopic Process”. In: vol. 0761. 1987, pp. 27–34 (cit. on p. 46).
- [176] S. S. Stevens. *On the theory of scales of measurement*. 1946 (cit. on pp. 13, 14).
- [177] E. Sundén, P. Steneteg, S. Kottravél, et al. “Inviwo – An Extensible, Multi-Purpose Visualization Framework”. In: (2015) (cit. on p. 118).
- [178] N. Svakhine, D. Ebert, and W. Andrews. “Illustration-Inspired Depth Enhanced Volumetric Medical Visualization”. In: *Visualization and Computer Graphics, IEEE Transactions on* 15.1 (2009), pp. 77–86 (cit. on p. 26).
- [179] J. M. Thijssen. “Ultrasonic speckle formation, analysis and processing applied to tissue characterization”. In: *Pattern Recognition Letters* 24.4 (2003), pp. 659–675 (cit. on p. 112).
- [180] J.-P. Thirion. “Image matching as a diffusion process: an analogy with Maxwell’s demons”. In: *Medical image analysis* 2.3 (1998), pp. 243–260 (cit. on p. 80).
- [181] J. Thomson, E. Hetzler, A. MacEachren, M. Gahegan, and M. Pavel. “A typology for visualizing uncertainty”. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by R. F. Erbacher, J. C. Roberts, M. T. Gröhn, and K. Börner. Vol. 5669. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. 2005, pp. 146–157 (cit. on p. 49).
- [182] J. Tokuda, G. S. Fischer, X. Papademetris, et al. “OpenIGTLink: an open network protocol for image-guided therapy environment”. In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 5.4 (2009), pp. 423–434 (cit. on pp. 69, 130).
- [183] M. Tory and T. Möller. “Rethinking visualization: A high-level taxonomy”. In: *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on*. IEEE. 2004, pp. 151–158 (cit. on p. 13).
- [184] G. M. Treece, R. W. Prager, A. H. Gee, and L. Berman. “Correction of probe pressure artifacts in freehand 3D ultrasound”. In: *Medical Image Analysis* 6.3 (2002), pp. 199–214 (cit. on pp. 77, 79, 86, 111).
- [185] L. A. Treinish et al. “Why Should Engineers and Scientists Be Worried About Color?” In: *IBM Thomas J. Watson Research Center, Yorktown Heights, NY* (2009) (cit. on p. 46).
- [186] A. Treisman. “Preattentive processing in vision”. In: *Computer Vision, Graphics, and Image Processing* 31.2 (1985), pp. 156–177 (cit. on p. 39).
- [187] J. R. Tresilian, M. Mon-Williams, and B. M. Kelly. “Increasing confidence in vergence as a cue to distance”. In: *Proceedings of the Royal Society of London B: Biological Sciences* 266.1414 (1999), pp. 39–44 (cit. on p. 43).
- [188] E. R. Tufte and P. Graves-Morris. *The visual display of quantitative information*. Vol. 2. 9. Graphics press Cheshire, CT, 1983 (cit. on p. 11).
- [189] E. R. Tufte and E. Weise Moeller. *Visual explanations: images and quantities, evidence and narrative*. Vol. 36. Graphics Press Cheshire, CT, 1997 (cit. on p. 12).
- [190] F.-Y. Tzeng, E. B. Lum, and K.-L. Ma. “A Novel Interface for Higher-Dimensional Classification of Volume Data”. In: *Proceedings of the 14th IEEE Visualization 2003 (VIS’03)*. VIS ’03. IEEE Computer Society, 2003, pp. 66– (cit. on p. 28, 101).

- [191] A. Varchola. “Live Fetoscopic Visualization of 4D Ultrasound Data”. PhD thesis. Institute of Computer Graphics and Algorithms, Vienna University of Technology, 2012 (cit. on p. 37).
- [192] Y. Wan and C. Hansen. “Fast Volumetric Data Exploration with Importance-Based Accumulated Transparency Modulation.” In: *Volume Graphics*. 2010, pp. 61–68 (cit. on p. 93).
- [193] S.-R. Wang, Y.-N. Sun, and F.-M. Chang. “Artifact removal and texture-based rendering for visualization of 3D fetal ultrasound images”. In: *Medical & biological engineering & computing* 46.6 (2008), pp. 575–588 (cit. on p. 36).
- [194] X. Wang, C. Schulte zu Berge, S. Demirci, P. Fallavollita, and N. Navab. “Improved interventional X-ray appearance”. In: *Mixed and Augmented Reality (ISMAR), 2014 IEEE International Symposium on*. IEEE. 2014, pp. 237–242 (cit. on p. 136).
- [195] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. “Image quality assessment: from error visibility to structural similarity”. In: *Image Processing, IEEE Transactions on* 13.4 (2004), pp. 600–612 (cit. on p. 56).
- [196] M. O. Ward. “A Taxonomy of Glyph Placement Strategies for Multidimensional Data Visualization”. In: *Information Visualization* 1.3-4 (2002), pp. 194–210 (cit. on p. 33).
- [197] C. Ware. *Information Visualization: Perception for Design*. 3rd ed. Morgan Kaufmann Publishers Inc., 2012 (cit. on p. 40).
- [198] W. Wein and A. Khamene. “Image-based method for in-vivo freehand ultrasound calibration”. In: *Medical Imaging*. International Society for Optics and Photonics. 2008, 69200K–69200K (cit. on p. 85).
- [199] W. Wein, F. Pache, B. Röper, and N. Navab. “Backward-warping ultrasound reconstruction for improving diagnostic value and registration”. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2006*. Springer, 2006, pp. 750–757 (cit. on pp. 76, 82, 85–87).
- [200] T. Wendler, K. Herrmann, A. Schnelzer, et al. “First demonstration of 3-D lymphatic mapping in breast cancer using freehand SPECT”. In: *European journal of nuclear medicine and molecular imaging* 37.8 (2010), pp. 1452–1461 (cit. on p. 23).
- [201] M. N. Wernick and J. N. Aarsvold. *Emission tomography: the fundamentals of PET and SPECT*. Academic Press, 2004 (cit. on p. 23).
- [202] C.-F. Westin, S. Maier, B. Khidhir, P. Everett, F. Jolesz, and R. Kikinis. “Image Processing for Diffusion Tensor Magnetic Resonance Imaging”. English. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI’99*. Ed. by C. Taylor and A. Colchester. Vol. 1679. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 1999, pp. 441–452 (cit. on p. 30).
- [203] C.-F. Westin, S. Maier, H. Mamata, A. Nabavi, F. Jolesz, and R. Kikinis. “Processing and visualization for diffusion tensor MRI”. In: *Medical Image Analysis* 6.2 (2002), pp. 93–108 (cit. on p. 30).
- [204] J. J. Wild and J. M. Reid. “Application of Echo-Ranging Techniques to the Determination of Structure of Biological Tissues”. In: American Association for the Advancement of Science. 1952 (cit. on p. 50).
- [205] D. Williams, S. Grimm, E. Coto, A. Roudsari, and H. Hatzakis. “Volumetric Curved Planar Reformation for Virtual Endoscopy”. In: *Visualization and Computer Graphics, IEEE Transactions on* 14.1 (2008), pp. 109–119 (cit. on p. 25).
- [206] I. Wolf, M. Vetter, I. Wegner, et al. “The medical imaging interaction toolkit”. In: *Medical image analysis* 9.6 (2005), pp. 594–604 (cit. on p. 118).

- [207] O. Zettinig, C. Hennemersperger, C. Schulte zu Berge, M. Baust, and N. Navab. “3D Velocity Field and Flow Profile Reconstruction from Arbitrarily Sampled Doppler Ultrasound Data”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Springer, 2014, pp. 611–618 (cit. on pp. 77, 135).
- [208] O. Zettinig, A. Shah, C. Hennemersperger, et al. “Multimodal image-guided prostate fusion biopsy based on automatic deformable registration”. In: *International journal of computer assisted radiology and surgery* 10.12 (2015), pp. 1997–2007 (cit. on pp. 131, 132, 135).
- [209] D. Zikic, M. Baust, A. Kamen, and N. Navab. “A general preconditioning scheme for difference measures in deformable registration”. In: *Computer Vision (ICCV), 2011 IEEE International Conference on*. IEEE, 2011, pp. 49–56 (cit. on p. 80).
- [210] Y Zimmer, S Akselrod, and R Tepper. “The distribution of the local entropy in ultrasound images”. In: *Ultrasound in medicine & biology* 22.4 (1996), pp. 431–439 (cit. on p. 112).

Online Resources

- [@1] Arno Klein. *Information Visualization Database*. 2015. URL: <http://www.infovis.info/index.php?words=konya> (visited on Aug. 27, 2015) (cit. on p. 10).
- [@2] Beazley, D. *Simplified Wrapper and Interface Generator (SWIG)*. 2015. URL: <http://swig.org/> (visited on Nov. 18, 2015) (cit. on p. 128).
- [@3] Bilas, Scott. *A Data-Driven Game Object System*. 2002. URL: <http://gamedevs.org/uploads/data-driven-game-object-system.pdf> (visited on Nov. 18, 2015) (cit. on p. 119).
- [@4] Celes, W. and de Figueiredo, L. H. *Lua (Programming Language)*. 2015. URL: <http://www.lua.org/> (visited on Nov. 18, 2015) (cit. on p. 127).
- [@5] Intel Corporation. *Threading Building Blocks*. URL: <https://www.threadingbuildingblocks.org/> (visited on Nov. 18, 2015) (cit. on p. 122).
- [@6] Jona Lendering. *Livius.org - Articles on ancient history*. 2015. URL: <http://www.livius.org/pen-pg/peutinger/map.html> (visited on Aug. 27, 2015) (cit. on p. 10).
- [@7] Kitware, Inc. *CMake: Cross Platform Make*. 2015. URL: <http://cmake.org/> (visited on Nov. 18, 2015) (cit. on p. 120).
- [@8] Kitware, Inc. *Insight Segmentation and Registration Toolkit (ITK)*. 2015. URL: <http://itk.org/> (visited on Nov. 18, 2015) (cit. on p. 118).
- [@9] Kitware, Inc. *Visualization Toolkit (VTK)*. 2015. URL: <http://vtk.org/> (visited on Nov. 18, 2015) (cit. on p. 118).
- [@10] Martin, Adam. *Entity Systems are the future of MMOG development*. 2007. URL: <http://t-machine.org/index.php/2007/09/03/entity-systems-are-the-future-of-mmog-development-part-1/> (visited on Nov. 18, 2015) (cit. on p. 119).
- [@11] Qt Company Ltd. *Signals & Slots in Qt*. URL: <http://doc.qt.io/qt-5/signalsandslots.html> (visited on Nov. 18, 2015) (cit. on p. 125).
- [@12] The Open XIP Project. *Extensible Imaging Platform (XIP)*. 2015. URL: <http://www.openxip.org/> (visited on Nov. 18, 2015) (cit. on p. 119).
- [@13] M. West. *Evolve Your Hierarchy: Refactoring Game Entities with Components*. 2007. URL: <http://cowboyprogramming.com/2007/01/05/evolve-your-heirachy/> (visited on Nov. 18, 2015) (cit. on p. 119).
- [@14] K. Wilson. *Game Object Structure: Inheritance vs. Aggregation*. 2002. URL: <http://gamearchitect.net/Articles/GameObjects1.html> (visited on Nov. 18, 2015) (cit. on p. 119).

List of Figures

2.1	Redrawing of the Çatalhöyük map, the oldest known map of a town.	10
2.2	First sheet of the Peutinger map considered as the first example of an abstract map.	10
2.3	Examples of Leonardo da Vinci’s recordings on his observations of the human anatomy.	11
2.4	Charles Minard’s visualization of Napoleon’s Russian campaign.	11
2.5	Illustration of John Snow’s cholera epidemic map of 1854.	12
2.6	Illustration of the Visualization Pipeline.	13
2.7	The four-level hierarchical model for visualization creation.	15
2.8	Illustration of an approach for structured evaluation of visualization based on Munzner’s hierarchical framework.	16
3.1	Analog X-Ray image showing the forearm and wrist.	20
3.2	Illustration of different cross-sectional slices through a CT volume.	21
3.3	Illustration of different MRI protocols applied to brain imaging.	22
3.4	Example ultrasound images showing different clinical applications.	22
3.5	Illustration of multi-planar reformations (MPR) used for neuro surgery planning.	24
3.6	Illustration of advanced reformation techniques embedding non-flat surfaces into a 2D image.	25
3.7	Illustration of image order based ray casting for direct volume rendering.	26
3.8	Example of a simple 1D intensity transfer function editor.	27
3.9	Illustration of dynamic galleries as proposed by Jönsson et al.	28
3.10	Comparison of the construction of streamlines, pathlines, and streaklines.	29
3.11	Examples of state-of-the-art flow visualization for vasculature.	29
3.12	Examples of standard diffusion tensor visualization techniques.	30
3.13	Examples of different glyph visualization techniques for DTI data.	31
3.14	Illustration of slice-based multi-modal visualization of brain PET-MR imaging.	31
3.15	Examples for state-of-the-art multi-modal volume rendering.	32
3.16	Illustration of supersphere and supertorus glyphs.	33
3.17	Comparison of different glyph placement strategies.	33
3.18	Illustration of the MagiCut technique.	35
3.19	Exemplary results of the lowest-variance-based filtering method for ultrasound volumes.	35
3.20	Illustration state-of-the-art classification results to 3D fetal ultrasound.	36
3.21	Comparison of an image acquired during fetoscopy with a fetoscopic rendering of 3D ultrasound.	37
3.22	Relevance-based multi-modal ultrasound visualization as presented by Burns et al.	38
4.1	Demonstration of preattentive visual processing in comparison to focused attention visual processing.	40

4.2	Typical response times for outlier search tasks.	40
4.3	Illustration of grouping and segregation during the perceptual organization stage.	41
4.4	Illustration of three of the four key gestalt principles.	42
4.5	Illustration of the light-from-above assumption.	43
4.6	Illustration of different monocular depth cues.	44
4.7	Illustration of chromadepth and pseudo-chromadepth coding of depth.	46
4.8	Examples of different color maps and when they are used best.	46
5.1	Illustration of a longitudinal wave exhibiting compression and rarefaction.	50
5.2	Illustration of ultrasound reflection and transmission at tissue interfaces.	50
5.3	Illustration of Confidence Maps on liver ultrasound.	51
5.4	Illustration of the diffusion graph network for computing Confidence Maps.	52
5.5	Illustration of the proposed incremental computation scheme for Confidence Maps.	54
5.6	Processing pipeline of our reference implementation for real-time Confidence Maps.	54
5.7	Solver runtime per frame with different configurations of PCG iterations and resample scales.	55
5.8	Evaluation of the image quality with respect to the resample scale in terms of structural similarity.	56
5.9	Error in terms of structural similarity over time for different PCG iteration counts.	57
5.10	Qualitative results of the presented incremental solver scheme.	58
6.1	Illustration of uncertainty visualization for radio-frequency ablation as proposed by Rieder et al.	62
6.2	Uncertainty visualization of segmentation algorithm results as proposed by Prassni et al.	63
6.3	Illustration of the effect of transducer repositioning on the image quality.	64
6.4	Schematic diagram of the different proposed uncertainty visualization schemes.	66
6.5	Illustration of the color overlay mapping scheme applied to kidney ultrasound.	67
6.6	Illustration of the chroma mapping scheme applied to kidney ultrasound.	68
6.7	Illustration of the fuzziness mapping scheme applied to kidney ultrasound.	69
6.8	User study results on the educational value of our technique to ultrasound novices.	70
6.9	User study results on the perception and clinical value for ultrasound novices.	71
6.10	Questionnaire results on expert sonographers' uncertainty perception.	72
7.1	Illustrations comparing forward warping methods with backward warping methods for ultrasound compounding.	76
7.2	Illustration of function-based methods for ultrasound compounding.	77
7.3	Illustration of the elevational decorrelation curve for ultrasound images.	78
7.4	Reconstruction of an abdominal phantom scan with probe pressure changes.	79
7.5	Illustration of artifacts occurring in distance-weighted compounded regions where multiple ultrasound frames from different directions intersect.	80
7.6	Effect of the orientation-driven clustering of ultrasound frames.	81
7.7	Schematic diagram of the proposed incremental compounding technique.	83
7.8	Illustration of the effectiveness of the pressure compensation applied to Infraspina-tus ultrasound data.	84
7.9	Illustration of our method to evaluate the pressure compensation.	87
8.1	Illustration of occlusion artifacts of when using traditional classification methods.	93

8.2	Schematic diagram of our predicate-based rendering technique.	94
8.3	Screenshot of the proposed predicate histogram widget.	97
8.4	Visualization results for the shoulder data set.	99
8.5	Illustration of scribble-based predicate histogram setup.	101
8.6	Visualization results for the achilles tendon data set.	102
8.7	Visualization results for the two carotid artery data sets.	103
8.8	Illustration of the transferability of predicate histograms between different data sets of the same anatomy.	104
A.1	UML diagram of the Entity Component System software architecture.	119
A.2	UML component diagram providing an overview of the main concepts of CAMPVis.121	
A.3	UML diagram of the CAMPVis package structure.	122
A.4	UML diagram of the adapted Entity Component System software architecture for CAMPVis.	123
A.5	Illustration of the CAMPVis pipeline-processor concept.	124
A.6	Screenshot of CAMPVis with enabled scripting functionality.	128
A.7	Illustration of the Lua binding generation process using SWIG.	130
A.8	Deployment of CAMPVis into a clinical setup for real-time B-mode ultrasound uncertainty visualization.	131
A.9	Screenshot of the developed multi-modal image-guided prostate biopsy framework.132	

List of Tables

2.1	List of scales (data domains) and their corresponding properties.	14
5.1	Aggregated performance results in terms of runtime and error with different configurations for the number of iterations.	57
6.1	Quantitative evaluation results of our user study with ultrasound novices. . . .	70
6.2	Questionnaire results on expert sonographers evaluating the clinical value of our uncertainty visualization technique.	72
7.1	Quantification of the local entropy in compounded volumes in terms of mean intensity difference to neighboring voxels.	86
7.2	NCC and log-scale SNR in the overlapping region after registering the two compounded volumes of two sweeps with perpendicular trajectories of the same anatomy.	87
7.3	System performance of our incremental compounding technique at different target volume resolutions.	88
8.1	Frame rates of our predicate-based approach compared to standard 1D transfer function-based direct volume rendering.	104

