

Ultrasonic Vessel Visualization: From Extraction to Perception

Åsmund Rognerud Birkeland



Dissertation for the degree of Philosophiae Doctor (PhD)

Supervised by Ivan Viola
Co-supervised by Trygve Hausken

Department of Informatics
University of Bergen

15th of March 2013

Scientific environment

The work presented in this thesis was conducted as a part of my PhD studies at the Department of Informatics, University of Bergen. In addition, I have been enrolled in the Norwegian Research School of Medical Imaging as well as the ICT Research School at the Department of Informatics, University of Bergen. This work has been part of the VERDIKT funded project *Illustrative Ultrasound* (#1931170). The research has been conducted in collaboration with Christian Michelsen Research and Haukeland University Hospital, with the support of MedViz research network in Bergen, Norway (PK1760-5897- Project 11). Also, we have been in collaboration with Vienna University of Technology, Austria, which I visited during my studies.

ILLUSTRASOUND



Christian Michelsen Research

Medim

Norwegian Research School
in Medical Imaging

Research School In
Information and Communication Technology

ICT

UNIVERSITY OF BERGEN



Acknowledgements

Firstly, I would like to give my outmost thanks to my supervisor, Ivan Viola. For without his support and inspiration, this work would not have been possible. The many scientific and non-scientific discussions has been crucial for a successful completion of my studies. Also, I would like to thank Trygve Hausken, who was my co-supervisor, for his input on medical expertise. Thanks also goes to my friends and colleges in the MedViz network, especially to Odd Helge Gilja, for many fruitful comments and discussions related medical ultrasound. Additional thanks goes to my fellow researchers on the IllustraSound project, Dag Magne Ulvang, Ola Kristoffer Øye and Svein Brekke. A special thanks goes out to my friends and colleges at the Visualization Group at the Department of Informatics here in Bergen. Julius Parulek, Endre Lidal, Andrea Brambilla, Çağatay Turkey, Mattia Natali, Stefan Bruckner, Armin Pobitzer, Jean Paul Balabanian, Ove Daae Lampe, Johannes Kehrer, Paolo Angelelli, Daniel Patel and especially our strong leader Helwig Hauser. You have all made VisGroup a great working environment for me, and I am forever grateful. Veronika Šoltészová, thank you for putting up with me, for providing an artistic eye to my illustrations and renderings, for your helpful comments, and for the huge pink ball that was residing in our office. I am also thankful to Eduard «Der Meister» Gröller from the Vienna University of Technology, for his interesting discussions and input to my research. Also, I would like to thank my colleagues at Meister's research group for making my research stay endurable, and motivating me through the final write-up of my thesis. Not to forget my other friends and colleges in the Vis-community around the world who aided me during my PhD: Timo Ropinski, Andrej Varchola, Roy van Pelt and many more. To my family. I would not be here without your never ending support. My parents, for never doubting in me even when I did. My brother, Tore, who are my great inspiration and role model, followed by his fantastic family, Kjersti and little Margit, who have provided me with sanctuary at Voss for skiing, hiking and beer brewing (mostly tasting). I would like to thank Tore, Martin, Alexey and Gabriel for helping me with spell checking of my thesis, and Raymond for some L^AT_EX-fu. Finally, I thank all my friends who were not previously mentioned. Makalani, Sigurd, Sveinung, Ingunn, Kristin, Juni, Sara and the rest of Morgonsymjng for dei Unge og Vakre, who prevented my physical degradation. Johannes, for making sure I was fed during the long working days. My flat-mates Spiros and Sabrina, for your company and friendship.

Abstract

Ultrasound is one of the most frequently used imaging modalities in modern medicine. The high versatility and availability of ultrasound workstations is applied in various medical scenarios, such as diagnosis, treatment planning, intra-operative imaging, and more. Modern ultrasound workstations provide live imaging of anatomical structures, as well as physiological processes, such as blood flow. However, the imaging techniques have a high presence of noise, a small scan sector, and are much affected by attenuation artefacts. Thus, traditional techniques for segmentation and visualization are not applicable to ultrasound data.

In this theses, we present our latest advancements in segmentation and visualization techniques, tailored specifically for the characteristics of ultrasound data. We present new methods for interactive vessel segmentation for both 3D freehand and 4D ultrasound. By directly involving the examiner in the segmentation approach as well as combining data from different *probe viewpoints*, we are able to obtain 3D models of blood vessels rapidly and robustly.

With the ability of robust vessel extraction, we introduce novel visualization techniques which utilize the previously acquired 3D vessel models. For anatomical imaging, we present a new physics-based approach for volume clipping, enhanced slice rendering and even defining curved Couinaud-surfaces. The technique creates a deformable membrane to adapt to structures in the underlying data, defined either by predefined segmentation, iso-values, or other data attributes.

For functional imaging, medical ultrasound can use the Doppler principle to image blood flow. However, Doppler ultrasound only measures a projected velocity magnitude of the data. In this thesis, we present a technique that uses the direction of the blood vessels in order to reconstruct 3D blood flow from Doppler ultrasound. By extending Doppler ultrasound with this directional information, we are able to apply traditional flow visualization techniques for displaying the blood flow. Finally, we investigated the usage of moving particles as a means to depict velocity in flow visualization. Based on a series of studies targeted for motion perception, we present a new compensation model to correct for distortions in the human visual system. This model can help users to make a more consistent estimation of velocities from evaluating the motion of particles.

List of papers

- (A) **Å. Birkeland** and I. Viola. Ultrasound Painting of Liver Vascular Tree, In *Proceedings of Vision, Modelling and Visualization* (2010). pages 163–170.
- (B) **Å. Birkeland**, S. Bruckner, A. Brambilla and I. Viola. Illustrative Clipping Membrane. In *Computer Graphics Forum* (2012), 31(3) pages 905-914.
- (C) **Å. Birkeland**, V. Šoltészová, D. Hönigmann, O. H. Gilja, S. Brekke, T. Ropinski, and I. Viola. The Ultrasound Visualization Pipeline. In submission to *The Dagstuhl book*, Springer, 2013.
- (D) **Å. Birkeland**, Ç. Turkay, and I. Viola. Perceptually Uniform Motion Space. Submitted to *IEEE Transactions on Visualization and Computer Graphics*, January 2013.
- (E) **Å. Birkeland**, D. M. Ulvang, K. Nylund, T. Hausken, O. H. Gilja and I. Viola. Doppler-based 3D Blood Flow Imaging and Visualization. Submitted to *Spring Conference of Computer Graphics*, March 2013.

The following publication is also related to this thesis:

- (I) I. Viola, **Å. Birkeland**, L. Helljesen, H. Hauser, S. Kotopoulos and K. Nylund. High-Quality 3D Visualization of In-Situ Ultrasonography. In *Proceedings of EUROGRAPHICS (short)* (2013), to appear. Awarded *Dirk Bartz Prize for Visual Computing in Medicine 2013*.

The manuscripts presented in this thesis were written during the PhD studies of the main author. All papers were written in collaboration with Ivan Viola, who is the main supervisor of the main author. Viola contributed with advice and guidance, and also with inspiration to the novel ideas of the research in this thesis. Paper B was coauthored with Stefan Bruckner, who aided with inspiration to the results. He also aided in creating some of the final images, as well as improving the manuscript. This paper was also coauthored by Andrea Brambilla, who helped with collecting related work and formulating the physics in the manuscript. Paper C was coauthored by Veronika Šoltészová, who helped with related work and manuscript regarding pre-processing of data and shading and illumination. Dieter Hönigmann wrote the section regarding transfer-function design in the rendering chapter. Timo Ropinski provided additional input into the shading and illumination section. Svein Brekke helped with the collecting related work as well as input regarding the technology of ultrasound. Odd Helge Gilja aided with his medical expertise in the motivation section of the manuscript. Paper D was

coauthored with Çağatay Turkey, who helped with his expertise in statistics. He also aided in writing and improving the manuscript. Paper E was coauthored by Dag Magne Ulvang who aided with the technological solution for streaming ultrasound data from the ultrasound workstations. Kim Nylund, Trygve Hausken and Odd Helge Gilja provided help with formulating the medical motivation and otherwise improving the final manuscript.

Contents

Scientific environment	i
Acknowledgements	iii
Abstract	v
List of papers	vii
I Overview	1
1 Introduction	3
1.1 Scope and Contributions	5
2 Visual Computing for Medical Ultrasound	7
2.1 Ultrasound Data Characteristics	7
2.2 State of the art in Technologies	10
3 Interactive Vessel Extraction	17
3.1 3D Freehand Ultrasound	19
3.2 4D Ultrasound	21
4 Ultrasound-Tailored Vessel Visualization	25
4.1 Adaptive Volume Clipping	26
4.2 Blood Flow Reconstruction	29
4.3 Perception-based Flow Visualization	33
5 Conclusion and Future Work	37
II Scientific Results	41
A Ultrasound Painting of Liver Vascular Tree	43
A.1 Introduction	43
A.2 Related Work	45
A.3 Ultrasound Painting Methodology	47
A.3.1 Vessel Detection	47
A.3.2 Vessel Tracking	48

A.4	Vessel Operations	49
A.4.1	Point Cloud Editing	50
A.4.2	Point Cloud Triangulation	51
A.5	Results	51
A.6	Conclusion	53
B	Illustrative Membrane Clipping	55
B.1	Introduction	55
B.2	Related Work	57
B.3	Methodology	59
B.3.1	Deformable Membrane	59
B.3.2	Potential field	61
B.3.3	Interaction	61
B.3.4	Visual Mapping	62
B.4	Implementation	64
B.5	Results	64
B.6	Discussion	66
B.7	Summary and Conclusion	68
C	Ultrasound Visualization Pipeline	69
C.1	Introduction	69
C.2	Taxonomy	70
C.3	Pre-processing	72
C.3.1	Reconstruction	73
C.3.2	Data Enhancement	74
C.4	Segmentation	74
C.4.1	Clipping	75
C.5	Registration	76
C.6	Rendering	78
C.6.1	Transfer Function Design	79
C.6.2	Multi-Modal Rendering	80
C.6.3	Shading and Illumination	81
C.7	Ultrasound and Augmented Reality	83
C.8	Summary and Discussion	84
D	Perceptually Uniform Motion Space	87
D.1	Introduction	87
D.2	Related Work	89
D.3	Methodology	90
D.4	Test Design	92
D.5	Initial Study	93
D.6	First Evaluation Study	97
D.7	Final Evaluation Study	99
D.8	Discussion	100
D.9	Summary and Conclusion	102

E Doppler-based 3D Blood Flow Imaging and Visualization	107
E.1 Introduction	107
E.2 Related Work	110
E.3 Methodology	111
E.3.1 Vessel Geometry Extraction	112
E.3.2 Flow Reconstruction	114
E.3.3 Visualization	116
E.4 Results	118
E.5 Evaluation	119
E.5.1 Accuracy	120
E.5.2 User Satisfaction	120
E.6 Conclusion	121

*«Reading goes faster if you don't
sweat comprehension.» –Calvin*

Part I

Overview

Chapter 1

Introduction

In 1895, Wilhelm Röntgen managed to produce an electromagnetic radiation in a specific wavelength, known as X-rays. This discovery would revolutionize modern medicine, with the capability of looking inside the human body, without performing any surgery. A second revolution in medical imaging came in 1942 when Karl Theodore Dussik published the first paper on medical ultrasonics and a radiation free alternative imaging modality was introduced.

With the capability of internal imaging, diagnosis and clinical decision making has improved the success rate in many areas of medicine. Different modalities often serve different purposes. Some modalities are well suited for anatomical imaging and while other can provide functional imaging in physiological processes. In today's medicine, we find a wide range of different imaging modalities. X-ray, Computed Tomography (CT), nuclear Magnetic Resonance Imaging (MRI), and ultrasound, all have their distinct characteristics.

Surgery planning has a great benefit from patient specific anatomical information before the patient reaches the operation table. Proper planning reduces the time spent on the operating table, reduces risk, and reduces recovery time. In brain surgery, it is crucial that the instruments enter the brain at the point with the minimal chance of damage. When performing a liver transplant, vascular structures in the donor must match the ones of the receiver. Similarly, in tumor resection in the liver, it is important not to interfere with the blood supply of the liver, and a complete knowledge regarding the anatomy of the hepatic and portal blood vessels is important. CT is often the modality of choice when acquiring anatomical images. While CT is a very good means for structural information, the radiation can be harmful for the patient, and therefore usage is restricted. The use of virtual surgery planning systems can aid in both patient treatment as well as surgery training [134].

In later years, high frame rate imaging has enabled doctors to investigate functionality in the human body, such as the heart motion, blood flow, and brain activity. 4D MRI provides the ability to measure the actual 3D blood flow over time. Analysis of blood flow is highly useful in the context of diagnosis for cardio-vascular diseases. In intra-operative surgeries as well as in post-operative scenarios, it is important that the blood flow is functioning at a normal level in order to achieve successful treatment of the patient. Fluoroscopy, i.e., a low-intensity X-ray imaging technique, can provide live view of blood flow. However, fluoroscopy is reliant on a contrast agent, which is invasively introduced into the patients blood stream.

When taking a needle biopsy of smaller structures, such as tumors or glands, it is

important to be sure whether or not the sample actually originates from the intended structure. During surgery, live imaging can guide the surgeon to the correct location. For instance, acquiring a biopsy of the thyroid gland, the gland are is completely fixed. Adding pressure on the skincan cause the gland to move and an updated view is required. This can be solved using images from CT or MRI and using a tracked co-registered biopsy needle. However, the drawback with such a solution is that the CT or MRI images have to be acquired prior to the operation. MRI is expensive and in many ways not suited for intra-operative imaging. The harmful radiation from CT makes it ill suited for live imaging. A live updated view would allow the surgeon to evaluate whether there has been any internal movement. Out of the various medical imaging modalities, one imaging modality stands out in versatility, cost, availability, and safety and that is ultrasonography.

Medical Ultrasound

Medical ultrasound is one of the most favoured imaging techniques. One very important factor is portability. Compared to CT or MRI scanners, the ultrasound machine can be brought into the patients room and used bedside. GE Vingmed Ultrasound developed scanner known as the VScan [48], which is of the size of an average smart phone. A doctor can carry the VScan in the pocket, and perform examinations on demand, even outside medical facilities. In addition, the price of acquiring and using an ultrasound machine is a mere fraction of that of an MRI scanner. MRI scanners are very costly to build and require specially trained radiologists to acquire images, and technicians to maintain the machine. This can result in a long patient waiting list to acquire crucial data for a patient. One other important aspect is the health risk. Unlike CT or flouroscopy, there are no proven side effects from ultrasound when used within the FDA regulations.

Medical ultrasound has a wide spread of applications including anatomical imaging, functional imaging, and even treatment when used in targeted drug delivery. The basic ultrasound image is acquired using a method known as B-mode imaging. B-mode imaging results in a gray scale image, where the intensities are given by the returning echo at a given position. B-mode yields high-resolution images, both spatially and temporally. Up until recent years, B-mode ultrasound was only available in 2D. Lately, however, the emergence of 3D and 4D ultrasound has resulted in a great opportunity for visualization research. The opportunity lies within having high resolution live imaging of the human body.

However, due to the characteristics of ultrasound data, creating images which are easy to interpret by an examiner becomes challenging. Unlike CT, the data values from B-mode ultrasound are not consistent to any specific underlying tissue. Attenuation of the ultrasound signal leaves certain areas of the image darker, while others are lit up as high intensity regions. This is a problem for standard volume visualization techniques, such as direct volume rendering, where a transfer function typically maps data intensities to color and opacity. Automatic segmentation techniques also suffer from this limitation. Region growing algorithms would *bleed* into areas, which should not be included. In this case, other modalities, such as CT, are often used for pre-operative planning, because traditional segmentation algorithms behave more robustly. While the

image quality of ultrasound might never reach the same level, the drawbacks of relying on imaging methods which have a low availability or contain a certain risk to the patients health, makes the pursuit of robust segmentation methods worthwhile.

Ultrasound also offers the possibility of imaging physiological processes in real time. The sensitivity of ultrasound allows us to measure the motion of blood. The applications of medical Doppler ultrasound were discussed as early as 1967 [140], and can measure motion towards and away from the ultrasound probe in both 2D and 3D. In this way, the examiner can mentally reconstruct an image of the blood flow. However, this is difficult for a layman and requires extensive training. B–flow imaging, presented in 2000 [183], provides a direct view of the blood flow, but does not allow for velocity measurements. In addition, B–flow requires a very high frame rate and is limited to 2D. Other methods for functional imaging involve contrast enhanced ultrasound (CEUS), which is a good tool for measuring perfusion in selected regions. This involves inserting a contrast agent into the blood stream, and can therefore be considered as invasive.

While ultrasound-based functional imaging has many application areas, there are certain limitations. In blood flow imaging, there is no feasible method for acquiring quantitative information of the real 3D velocities of the blood flow at specific regions. CEUS can measure perfusion, i.e., it can detect how fast blood is entering a specific region. Doppler ultrasound provides only a single component of the 3D velocity and is therefore difficult to interpret. There is a need for non-invasive methods for in-situ 3D flow imaging techniques, which are plausible to use in several scenarios, such as bedside examination and live during surgery.

1.1 Scope and Contributions

While ultrasound is already a highly used imaging modality in medicine, the potential of ultrasound has still not fully been reached. Traditional methods for medical visualization often fail to be applicable to the different types of ultrasound, as they do not consider the particularities of the data. The availability, cost, and non-invasiveness inspires ultrasound visualization research to fill the gaps in medical scenarios normally tackled by other imaging modalities. The main scope of this thesis is medical visualization tailored for ultrasound data. In particular, we focus on the structure and function of blood vessels, which is essential information in many medical scenarios.

Robustness is a key element in segmentation of medical data and therefore ultrasound is often excluded as an option. Exploiting the live aspect of ultrasound, we show how both 2D and 3D B–mode ultrasound can be used for robust interactive segmentation of blood vessels, where the expertise of the user is included in the process. While ultrasound navigation and orientation can be difficult for the user, we show how semantic information, in form of blood vessels, combined with novel methods for data interaction can enable expressive visualization aimed for navigating within the Couinaud segments in the liver. Finally, we show how structural information of blood vessels combined with Doppler ultrasound can aid in the interpretation and numerical estimation of blood flow. The individual contributions presented in this thesis are:

- New methods for interactive segmentation of blood vessels from 2D and 3D B–mode ultrasound.

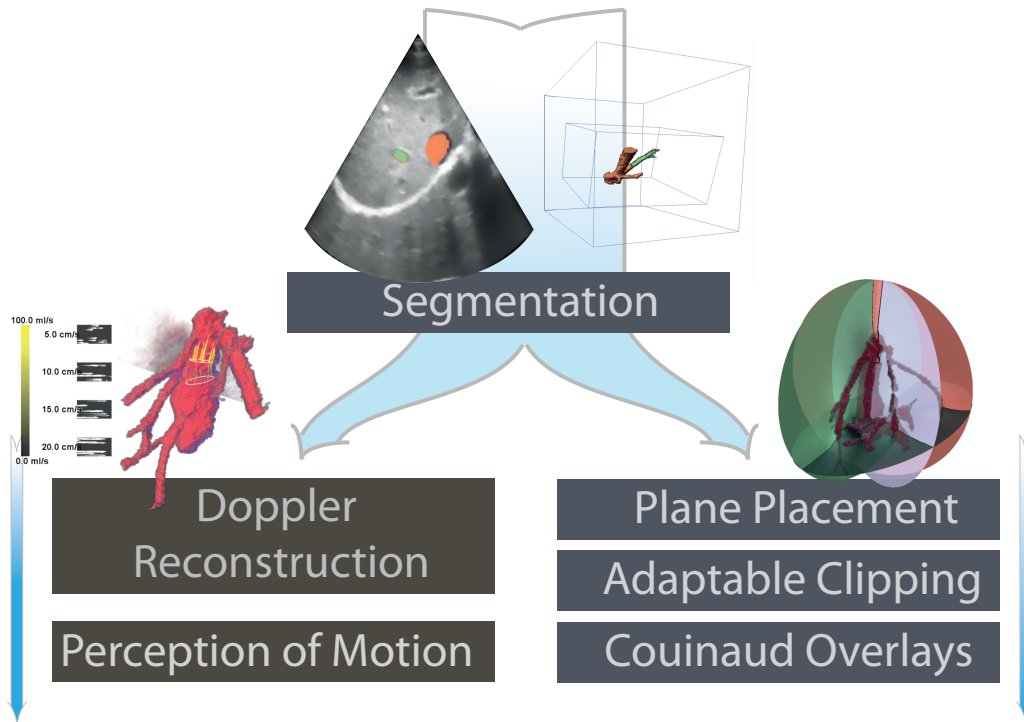


Figure 1.1: In this thesis, we will present new methods for extracting blood vessels from ultrasound data. Then, we show how the newly extracted geometry can be used in new ways for ultrasound visualization.

- Novel utilization of physics simulation to create expressive visualization for medical (ultrasound) data.
- 3D blood flow reconstruction from 3D Doppler ultrasound combined with vessel segmentation.
- New knowledge about human motion perception and new guidelines for using animated particles in visualization based on human perception.

Thesis Structure

This thesis consist of two major parts: The first part summarizes the scientific contributions made to the field of medical visualization, and shows how the presented work extends the current state-of-the-art. The second part consists of the individual publications produced from this thesis.

The outline of this thesis, as depicted in Figure 1.1 is as follows: In the next chapter, we present a more detailed introduction to ultrasound data and show the related work in ultrasound-specific visualization techniques. In Chapter 3, we present our methods for acquiring 3D segmentation of blood vessels and discuss their significance. Chapter 4 contains demonstrative cases where segmentation can be used in combination with live ultrasound as a basis for new visualization techniques. A conclusion and ideas for future work can be found in Chapter 5.

Chapter 2

Visual Computing for Medical Ultrasound

Medical ultrasound distinguishes itself from other medical imaging modalities in more ways than one. For this reason, visualization and image analysis methods which are suitable for CT and MRI, such as thresholding segmentation or ordinary surface illumination can cause problems when applying them to ultrasound data.

In the following section, we discuss in detail the different aspects of medical ultrasound data which makes it stand out from the other modalities. Furthermore, in Section 2.2 we discuss ultrasound aimed methods for segmentation and visualization, and their main difference to more common techniques.

2.1 Ultrasound Data Characteristics

Ultrasound images are formed by emitting pulsed acoustic waves into the soft tissue of the human body, and subsequently measuring the back-scattered waves - the echoes - as a function of time. During pulse propagation, variations in the acoustic impedance of the tissue cause some of the energy in the wave to be reflected, and by assuming a constant speed of sound in the medium and a directional pulse propagation, the measured echo amplitude can be positioned precisely in space. A whole image frame or a data volume is formed by emitting several pulses in different directions, covering a planar or volumetric region in the body.

Different ultrasound transducers can record different types of data. An ultrasound transducer is constructed using elements of piezo-electric materials [7], which can both receive and transmit sound waves. A focused ultrasonic pulse is formed by sending identical pulses from all (or a subset of) the elements, with internal delays designed to form a curved wave-front propagating towards a given point in space. Different probes can be designed by configuring these elements in various ways.

There are two prominent configurations: 1D arrays, where elements are aligned in a row, and 2D arrays, where elements are set together in a matrix. 1D arrays have the ability for live 2D imaging, whereas 2D arrays can acquire 3D volumes in real time. Apart from 2D over time versus 3D over time imaging, different medical areas utilize different configurations. The most simple configuration is a flat 1D array which transmits pulses orthogonally out of the probe. The size and shape of the array determines the usage of the ultrasound probe. A small transducer can be seen as similar to having a small photo sensor in a camera. The large sensor in an SLR camera can achieve better images with less noise, than the small sensor in a pocket camera. This translates to

ultrasound transducers, where a smaller area can contain less elements, which in turn reduces the noise filtering capabilities.

The transducer and attached electronic circuits are packed into a probe unit which is connected with a cable to the scanner. Probes for 2D scanning are roughly of three different kinds: Flat phased arrays, linear arrays and curved arrays [8]. Flat phased arrays (FPA) are designed to steer the propagation angle substantially in order to form a sector scan, and hence the distance between the elements must not be too large. It has a typical aperture of 1.5 to 2cm when designed for a frequency of 3MHz. The main use of FPA probes is transthoracic cardiac scanning, and the small footprint of this probe type makes it possible to reach between the ribs. Special FPA probes for transesophageal scanning exist as well, having even smaller aperture. Linear array (LA) probes are designed for a smaller steering angle and have larger element size and element count. The aperture is typically around 5cm and the frequency around 7-12 MHz. The main use of these probes is scanning vessels and shallow depths of the abdomen. Curved array (CLA) probes also have bigger elements and are designed to transmit beams perpendicular to the curved surface. This gives a wide sector image that covers most of the abdomen in one scan, and a frequency of 3-7 MHz is used to open up for deep scanning of organs and fetuses.

Mechanically swept transducers were introduced in commercial ultrasound machines in 1995, allowing for 3D acquisition and display, which is nowadays common practice in fetal imaging [105]. Since 2002, ultrasound scanners for real-time 3D imaging using fully sampled matrix probes have been commercially available [68]. Most of the matrix probes available so far are of FPA design, i.e, small aperture targeted at transthoracic or transesophageal cardiac imaging. Recently, curved 2D matrix probes have appeared [86], indicating the demand for abdominal and obstetrical live 3D ultrasound imaging.

Since ultrasound uses acoustics rather than photonic measures there are certain limiting factors. In basic ultrasound imaging, the data intensities are recorded by transmitting a pulse in a certain direction and listening for echoes. These echoes are caused by a change in impedance in the medium. For instance, the liver has different impedance than the fat in front. When the pulse hits the boundary, a certain amount of the energy is reflected back, while a remainder propagates through the tissue. This means that the boundaries between two structures with high difference in impedance will be highly visible in the image. However, since the sound waves reflect on boundaries, similar to light reflecting off a surface, the orientation between the ultrasound beam and the surface normal affects the amount of echo transmitted back to the transducer. Hence, the boundary between two structures may have different data values in the same image, and evaluating surface boundaries by gradient change is not always applicable.

As the ultrasound pulse propagates, the energy is attenuated by human tissue. The effect depends on the current depth of the signal and the frequency of the transmitted ultrasound. The amount of attenuation in human tissue is estimated to be between 0.2 - 0.5dB/cm/MHz. As the average attenuation is known, the reflected echoes can be amplified accordingly, also known as depth-gain compensation. The problem with depth-gain compensation is that the compensation model can only be accurate once one would know the actual attenuation. Thus the depth-gain compensation only digitally amplifies the acoustic signal, including noise, for areas deeper inside the body.

The frequency impacts the ultrasound image in two fundamental ways: High fre-

quency gives high spatial resolution, while it also allows for only a low penetration due to a higher attenuation. Therefore, one would like to use as high imaging frequency as possible without losing image quality due to attenuation. For images of the heart or the deep abdomen, where the maximum scanning depth is 10 to 20 cm, a good compromise is in the range of 2-5MHz. In case of scanning shallow vessels such as the carotid artery, one can use as high frequency as 10-15MHz and thereby get very detailed images [7]. Furthermore, since the concentration of acoustic energy is better in the focal depth of the transmitted pulse, the point resolution is not constant over the whole image. Scanners usually let the examiner choose a focus depth in order to maximize image sharpness in the main region of interest within the image frame.

Compared to other medical imaging modalities such as MRI and CT, ultrasound images have a strong presence of noise and artifacts. The noise has mainly two origins: random signal noise (thermal noise and electronic noise) and speckle noise. The random noise is reduced by temporal- and spatial-noise suppression algorithms. The speckle noise, originates from interference between echoes from neighbouring scatterers, and forms a characteristic pattern that persists during several imaging frames [23, 174]. Frequently seen artifacts are reverberations (multiple echoes creating patterns that are misplaced in space), side lobes (acoustic energy refraction one may cause artificial echo patterns in dark areas) and shadows (e.g., ribs occluding the acoustic field and thereby creating hypo-echoic parts of the images). The noise and artifacts, together with phase aberrations [111] (wavefront degradation due to non-uniform speed of sound), decrease the perceived image quality of ultrasound images.

The resolution of ultrasound images can be very high. For high frequency linear arrays, the pixel size lies in the sub-millimeter range. However, there are certain restraints, which are affecting the ultrasound resolution. Since ultrasound is considered to be primarily a live examination technique, the focus is on having a certain level of frame rate. The frame rate is given by how many times one can record the image per second, and the image is constructed by transmitting ultrasound pulses in beams from the transducer. The final resolution of the image is determined by the number of beams used per area. In the traditional approach, only one beam could be formed at a time, because the beam creation time was limited by the speed of sound and the recording depth. This directly links the spatial resolution of the ultrasound image with the frame rate one can achieve.

For 2D ultrasound, the required time for generating one image is low. Thus, the spatial resolution is more limited by the aperture of the transducer than preserving a high frame rate. The typical frame rates for 2D ultrasound range above 60 frames per second. With 3D ultrasound, the amount of beams become so high, that the user must determine the trade-off between frame rate and resolution.

To cover larger areas, transducers are using curved arrays and linear phased arrays. The basic principle is that these types of transducers is in the formation of beams from one point into several directions. This means every array of pixels contains a direction, i.e. the pixels are stored in spherical coordinates. To represent the pixels on the screen in a meaningful manner, the recorded image must be converted back to Cartesian coordinates. This process is called *scan conversion*, and it leads to a reduced spatial resolution with increasing scan depth.

US scanners are not only able to measure the echo amplitude in a spatial region, they can capture the velocity of the scatterers by assessing the Doppler shift induced during

the reflection. To measure the Doppler shift, several pulses have to be sent in the same direction, with a fixed delay between each pulse. A comparison of the delay between reflected pulses and the delay used in transmission shows if the scatterer has moved during the pulse sequence. The corresponding velocity of the scatterer is determined by the delay that gives the maximal echo strength. The ultrasound Doppler principle is frequently used for visualizing and quantifying the velocity field of blood and moving muscles [73].

The diversity of the ultrasound imaging technology provides a great tool for medical diagnostics, but the nature of the data can make it challenging to process. Techniques, which work well for other modalities, have to be substantially adapted to suit the special characteristics of data generated by ultrasound data.

2.2 State of the art in Technologies

Registration

There are certain challenges in ultrasound data visualization which are not that prominent in other imaging modalities. Firstly there is the question of overview. The ultrasound sector is quite small. Compared to CT or MRI, we can only see a fraction of the area, which makes it difficult to get a proper overview. Registration between ultrasound and other imaging modalities can aid in obtaining the context. In addition, information from several spatial positions can be helpful, such as stitching several ultrasound images together, in order to cover a larger area.

Direct image-based registration between ultrasound and CT or MRI can be difficult due to the different nature of the imaging techniques and usually requires a certain level of pre-processing, such as filtering. Different approaches have been proposed. As starting point, the position of the ultrasound probe can be tracked using some form of external tracking, also known as 3D freehand ultrasound. To register the tracking from freehand ultrasound with CT or MRI requires certain calibration [101]. Image registration techniques are then often used for optimizing the position obtained from an external tracking device in freehand ultrasound with a different dataset. Rather than comparing the images directly, Penney et al. used a statistical measure to calculate vessel probability values for both ultrasound images and MR images [123]. The registration is then carried out using the vessel probability images. A similar approach was also presented for CT data, where the probability map was based on bone-to-soft-tissue interfaces [124]. Leroy et al. proposed to highlight the tissue boundaries, before image registrations [95]. However, they reported that the technique was dependent on the choice of ultrasound slices. Ji et al. proposed to create a binary mask to highlight the most prominent features (mostly tumors) and register the features from the ultrasound image with the features in the MR image. Zhang et al. used an local phase image descriptor as a basis for image similarity [186]. However, this was applied to 3D ultrasound registering to a 2D MR image.

With 3D ultrasound, we are not as much reliant on any tracking of the ultrasound probe. However, volume-to-volume registrations also have the problem of different data characteristics. Previous work has described how using landmarks, such as the blood vessels, to register 3D ultrasound with CT or MRI [92, 93, 122]. though, some

reported robust vessel segmentation of ultrasound as an important obstacle. A different approach is the estimation of an ultrasound image for a certain given object or data. Wein et al. proposed to simulate an ultrasound image based on CT data, then performing image registration between the *simulated* and the *real* ultrasound images [177]. Similarly, King et al. used predefined segmentation masks and estimated the probability density map from an ultrasound imaging model [81].

When dealing with ultrasound-to-ultrasound registration, a more basic similarity measure can be used, such as the sum-of-square-differences [121]. Here, tracking of the ultrasound probe was achieved by registering the current ultrasound image with the previous. In order to minimize the error, the registration was performed with several previous ultrasound sectors. The challenge is to keep the registration running at interactive frame rates and a trade-off must be made between performance and accuracy.

Segmentation

Segmentation is the process of marking a certain region in the data as important, typically related to a specific structure. In medical imaging, structures such as blood vessels, bones, tumors, liver, are all important for different scenarios. In medical imaging, the need for extracting such features have fuelled much research in image segmentation.

A typical approach would be to extract areas which have the same data-values. This is suitable for some imaging modalities, such as CT. The data values in CT-data is recorded in a range called the Hounsfield scale. A great advantage of this scale is that the units for certain features are the same for each data type. This means that bone can be extracted by setting a threshold value, where all data points with a certain value are tagged as bone.

As discussed in the Section 2.1, the data values in ultrasound images are not stored in such a format. The multitude of aspects contributing to the inconsistency in the data values, make it inherently difficult to select features by thresholding. This was indicated in early work, stating that more sophisticated methods should be applied when aiming for satisfactory results in ultrasound segmentation [152].

Because the data intervals are an unreliable measure for feature extraction, general automatic techniques are not often seen for ultrasound. For certain cases, the diagnosis or treatment of patient is very reliant on the examiner's knowledge of specific structures [134]. Thus, the segmentation process can be extremely critical. If the doctor or surgeon blindly expects the segmentation mask to be the true representation of the structure, over-segmentation (including parts which should not be there) and under-segmentation (missing regions which should be included) can cause serious consequences. Therefore, it is very important for segmentation techniques to be robust. To increase the robustness, a certain technique is typically tailored for a specific type of tissue.

An extensive survey on ultrasound image segmentation has been presented by Noble and Boukerroui [113] in 2006. Fetal examination has been a driving force in development on ultrasound. While 3D ultrasound was first presented as a means for displaying 3D images of fetuses to the parents, clinical applications of 3D fetal ultrasound has been emerging. With 3D images, a facial recognition can be achieved [40], which can aid the examiner to look for early stages of development deficiencies.




Figure 2.1: Automatic segmentation of the ovarian follicles [33].

The fetus skull is very thin. This allows the ultrasound signal to propagate through the brain, which in turn allows for imaging of the brain. Carneiro et al. have developed an automatic technique for segmenting the brain of a fetus [25]. The method is using a probabilistic model, which combines certain classifiers with contextual information and sequential information. The results had a 80-90 percent success rate. While this is a very high success rate for automatic segmentation of ultrasound, in medical scenarios, this is often not sufficient. Therefore, there are very few automatic segmentation techniques seen in medical work stations. An exception is automatic assessment of the volume of ovarian follicles [33], as shown in Figure 2.1.

Steen et al. [153] introduced a region growing algorithm for extracting blood vessels. The approach based itself on creating a weighted directional graphs, where the voxels constitute as the nodes and edges, having weights indicated by the difference in data intensities. The approach is well suited for segmenting vessels, which have a fairly similar data values. In ordinary B-mode imaging, blood returns very little of the ultrasound pulse, also known as hypo-echoic. This means that the lumen of blood vessels shows up as dark areas. Considering the difference between one voxel and its neighbours, compensates a little for the change in data values, but it does not consider areas of shadow-regions caused by heavy attenuation. In shadow-regions, tissue will appear darker than usual and will have a closer resemblance to hypo-echoic regions. Region-growing techniques will then *bleed out* into these regions causing over-segmentation.

A similar approach was presented by Petersch et al. [127]. Here, a geometric approach was introduced. Instead of only considering the differences in intensities, the method created a distance manifold which also considered the geometric distance between the seeding point to any other point. Then a region growing algorithm was applied to the distance manifold, which extracts everything within a certain distance threshold. Using the geometric distance as an additional constraint can control the bleed out into shadow regions. It is however, dependent on the user to select a proper distance threshold. In addition to a geometric approach, Petersch et al. also presented a different method based on filtering the data using non-linear diffusion and then calcu-

lating the probability for each voxel to be part of the segmentation mask or not. While this provides a fast overview over structures of interest, it allows for minimal control compared to the distance manifold approach.

Most segmentation methods are designed as post-acquisition process, but one of the key elements of ultrasound is the fact that it is used as a live examination modality. Techniques do not take advantage of the fact that the ultrasound machine is continuously acquiring new data. In liver surgery it is important to know the location and extent of the vessels. To support the surgeon, Nowatschin et al. presented a technique for automatically extracting vessels live during surgery [115]. The vessel cross-section is estimated as an ellipse and mapped into a 3D model. While it provided a spatial reference, since the vessel was estimated by ellipses, the technique would not fully capture the actual shape of the vessel.

There is a clear need for segmentation approaches, and while automatic techniques can be considered as a *holy grail* for ultrasound segmentation, there are other aspects which can enable more usable tools for ultrasound segmentation. It is very important to focus on the strengths of the medical ultrasound modality, rather than to only correct for its weaknesses.

Visualization

In 2D ultrasound, the most intuitive method for displaying the recorded image, is to show the scan converted grey scale image. However, when an examiner is analysing the image, two challenges appear: Identifying the tissue type or other structural information (eg. different structures or different parts a structure) in the image, and spatially orient the current 2D image within a larger context, such as other organs or blood vessels. For identifying certain features, the speckle pattern can provide an indicator of what tissue types are present. However, this is not always the case and extensive training is required to correctly identify anatomical structures. For navigation, the doctor must mentally construct the overview by correlating what he/she sees in the image, with the position of the ultrasound probe on the patient, which also requires extensive training. Enhancing the 2D ultrasound image with some higher order information, can aid examiners in navigation and image interpretation.

For instance in liver examinations. The liver is a large, complex organ. Couinaud defined a partitioning scheme of the liver into eight different sections, based on the filtering of blood in standard liver vascular topology. If certain pathologies are detected, it is important to know in which Couinaud-segment the pathology lies. Since the different segments are defined by the blood vessels, the doctors orient themselves by locating the vessels, and having a poster on the wall as a point of reference.

Viola et al. proposed an approach to aid the examiner in identifying the Couinaud-segments when performing liver examinations [169]. By superimposing the currently visible Couinaud-segments onto the 2D ultrasound image, the examiner could immediately detect which segments were visible. The Couinaud-partitioning was pre-defined in a CT dataset, which could be viewed directly using exploded views [20]. CT or MRI can also be used to provide a more structural overview, such as Burns et al.'s illustrative cut-aways combined with 3D freehand ultrasound [24]. Here, a cutaway was generated by the tracked 2D ultrasound plane to show the spatial position of the ultrasound image compared to the CT image.

In 2D ultrasound recordings, the spatial understanding is even more challenging, and the knowing where some features of interest while watching the wider sequence might not be very clear. To improve ultrasound video analysis, Angelelli et al. used a degree-of-interest (DOI) function coregistered with a freehand ultrasound [6]. A interest-function was generated which indicated how much of the DOI-function was intersecting the ultrasound image. The users could then locate interesting parts of the sequence by evaluating the interest-function.

The largest challenge in enhancing 2D ultrasound images, is to the define any higher order information and register the information with the ultrasound images. One solution is to use registration techniques, as discussed above, but this is not always applicable if the information is abstract, such as the Couinaud-partitioning.

3D ultrasound is inherently difficult to display. Unlike 2D ultrasound, the speckle noise is no longer useful and is often filtered out. Advanced filtering approaches aim to reduce both speckle and random noise, while still preserving structures [170]. Traditional shading in direct volume rendering using gradient based normal estimation, has been proven not to be suited for 3D ultrasound and more sophisticated illumination techniques can provide a better perception of the data [171].

Adding functional information, such as Doppler ultrasound adds complexity to the visualization. During Doppler ultrasound imaging, both the Doppler signal as well as the B-mode image is acquired at the same time. This provides the examiner with anatomical information which can be used for navigation. In 2D, the Doppler signal is typically shown as a color-coded overlay on top of the B-mode image. 3D Doppler becomes even more challenging, when dealing with visual clutter from two overlapping volumes.

There have been several sophisticated attempts in combining ultrasound. For instance, Jones et al. proposed surface fitting of the Doppler signal and multi-planar slice rendering [74]. Petersch and Hönigmann proposed to use a combination of silhouette rendering blended with DVR [126] to reduce the emphasis on the B-mode and add focus on the Doppler signal. Other work use data fusion algorithms for DVR [79, 185].

While advanced blending techniques for volume rendering can aid in viewing both B-mode and Doppler ultrasound simultaneously, there are more aspects to consider when visualizing Doppler ultrasound. A major aspect is that Doppler ultrasound represents more than just a structure, it is a projection of the blood flow. The blood flow has a certain direction and speed, which are not depicted in basic rendering techniques. Understanding the blood flow from Doppler ultrasound requires extensive training both from 2D and 3D ultrasound. If the direction of the flow could be indicated, it could aid the doctors to evaluate the blood flow characteristics.

When rendering 3D ultrasound, it is difficult to set opacity levels in a meaningful manner. Interesting features may be occluded by noise or other structures with similar data values. In-vivo ultrasound examinations do not always have the possibility to use segmentation as a means for selective rendering. In such scenarios, clipping can be powerful tool for exploring 3D ultrasound.

A prime example of traditional clipping are found in GE Vingmed Vivid E9 cardiac scanners. Here, a standard clipping plane is available when rendering 4D ultrasound. The user can insert a flat plane into the volume. The plane works as a decider for which part of should be *clipped* away and which parts should be left in the image, to enable direct view of the internal structures.




Figure 2.2: Using MagiCut to Clip the volume generating a clear view to the desired structure [88].

A flat clipping plane is fast and easy to use, and the spatial understanding of the plane is simple. However, the interesting features within the body are rarely planar. Therefore, a clipping surface adapted to the underlying structures could be very suited for clipping away the occluding structures and revealing the features of interest. MagiCut is an example of such clipping, as seen in Figure 2.2. Here we see an example of clipping technique implemented in the GE Voluson machines. This type of clipping is based on techniques found in Sakas et al.'s ultrasound visualization system [141]. Three approaches were presented based on the user selecting the area of interest, either by 2D slice drawing, 3D sketching, or simple point selection.

Sketch-based approaches are a common approach to adaptable clipping geometries [21, 156]. The user has great flexibility for creating the desired clipping. On the other hand, much interaction has to be performed for achieving good clipping for any new feature of interest or for any new dataset. Recent development in medical workstations from GE Voluson, allows for automatic clipping for fetus rendering [49]. This approach is only tailored for fetus examination and it is uncertain if it would provide any meaningful results in cardiac or abdominal examinations.

3D ultrasound for abdominal examinations suffers from the fact that there are no rapid clipping methods which allow the user to explore the data set. Traditional volume rendering falls short, as the density of the volume makes it difficult to set meaningful opacity levels. Automatically defined clipping surfaces are unsuited for abdominal ultrasound. Adaptable clipping planes which interact with the underlying data in could provide the user with tools needed to properly take advantage of the 3D and 4D abdominal ultrasound.

Perceptual Aspects in Context to Ultrasound

Traditionally, visualization has been mostly motivated by depicting the underlying data accurately. However, the final output for a visualization technique is to convey information to the user. While a user can read out numbers accurately, when dealing with more complex representations, the perception might alter the conveyed information. The human visual system should be considered when designing visualization techniques.

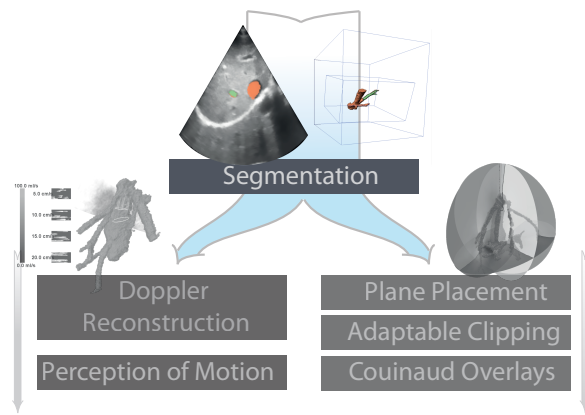
An example of a more perception-driven technique in visualization was the change from a traditional RGB encoding of color to a perceptually-uniform color space, called CIE 1976 (L^* , a^* , b^*) color space (CIELAB). Here, the minimal distance between

two colors is the minimal perceived change in color, also known as the just noticeable difference.

As previously mentioned, the high presence of noise in ultrasound data has produced a challenge for surface shading. Occlusion based illumination has shown great potential in shape depiction. However regions might be severely darkened by shadows, which make it difficult to perceive shape in regions in shadow. Šoltészová et al. performed a study to evaluate the effect of using a chromatic component for depicting shadow [172]. Inspired by illustrators, who use blue colors for depicting shadow, Šoltészová et al. investigated the influence of perceived shape for different chromatic levels of the shadow. From the study they found an improved understanding for shape when using a mix of luminance and chroma color for shadows. They also compared depth perception between occlusion shading to traditional phong shading. They found the use of occlusion shading improved depth perception, and that black shadows had a slight advantage as compared to chromatic shadows. As a conclusion, they recommend using a dark blue color for shading.

Recently, Šoltészová et al. investigated whether there are systematic distortions in our perception when evaluating the shape of a 3D object [173]. From the results, they found trends in estimation error for the slant and constructed a statically-based shading model.

While there now exist models for shape and color perception, there is a missing element in perceptually based models for animated objects. Animated particles are often used in medical flow visualization for depicting both the direction and speed of the flow. The motion of the particles is linearly scaled according to the underlying data. Still, the human visual system might not perceive the speed linearly which can cause an error in the estimation of the user. Previous research has provided certain guidelines for animation [180], but a comprehensive model of motion perception is still missing.



Chapter 3

Interactive Vessel Extraction

When extracting features out of data, automatically or manually, one very important aspect for the user is the reliability of the segmentation. The user must be able to rely on that the extracted segment represents the feature of interest in a satisfactory manner. If the confidence of the segmentation is low, the user must then be informed about problematic areas, where over- or under-segmentation may have occurred.

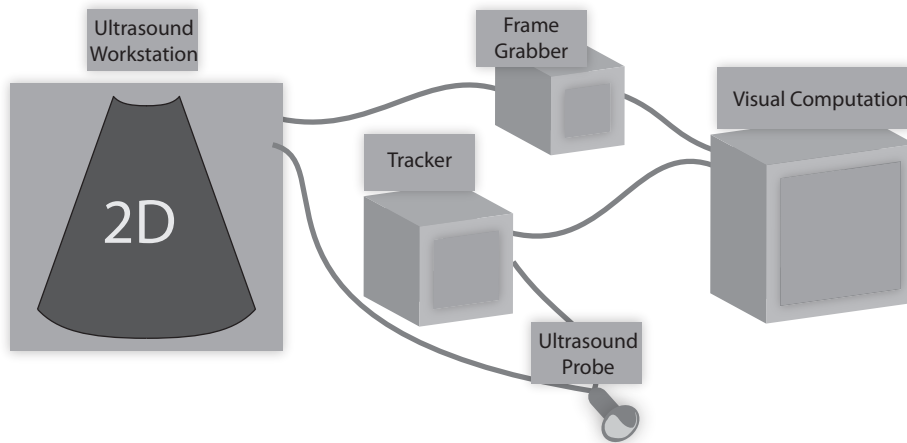
The biggest strength with ultrasound is its availability. Not only the availability of hardware, but also the availability of data. If there is some information missing in the data, certain areas are not covered, or the image quality of the specific section is poor, there is always the possibility to return and the problematic areas and rescan the patient. The reason for this availability is that ultrasound is used interactively, and the examiner can continue the image acquisition until he or she is satisfied with the data.

Most segmentation algorithms for medical data are based on a single recording of the specific patient. For ultrasound segmentation, this can lead to problems since the images are heavily dependent on the *viewpoint* of the ultrasound probe. For example, if one is aiming for vessel segmentation, certain parts might be partially occluded by highly attenuating structures, such as bone. An automated technique might compensate for this to some extent, but if the structures are completely occluded, there will not be any information for an automated approach.


In other modalities, lack of information means the patient has to undergo the complete scanning procedure once again with different parameters. This is cumbersome and also very inconvenient for both the patient and all the medical personnel included. As mentioned, ultrasound is readily available. The lack of information can be solved simply by changing the *viewpoint* of the ultrasound probe and the missing structure comes into view.

Segmentation techniques for ultrasound could benefit greatly from enabling the use of live data. In this chapter, we present new methods which consider segmentation of blood vessels as an iterative procedure. Our techniques allow the user to continuously update the already segmented model with new data. Ultrasound data are streamed live and are co-registered to the currently segmented model.

In order to put everything into a common reference frame, the position of the ultrasound image and the patient must be tracked in space. There are two different categories of tracking-techniques which are suitable for ultrasound tracking: External and internal. External tracking means to keep track of the ultrasound image by knowing the position of the ultrasound probe. Several implementations have been introduced, typically using either optical tracking or magnetic tracking. Internal tracking, also referred to as



(a) 3D freehand ultrasound setup.



(b) 4D ultrasound setup.

Figure 3.1: The different technical solutions for data acquisition and segmentation.

image-based tracking, can be considered as positioning the image by co-registering the image with a common reference frame. This reference frame can be defined by either other imaging modalities, or images captured from a previously recorded ultrasound examination.

The different types of tracking have different scenarios to which they can be applied. Image-based tracking relies on matching one image to the another. If spatial tracking is desired, internal tracking can not be applied to 2D ultrasound, as the out-of-plane motion is not captured. Hence, internal tracking is only suitable for volumetric data. However, using external tracking systems, we can position the probe regardless of the data-type, but we are then unable to compensate for involuntary internal motion caused by for instance breathing or the heart pulse.

We propose two methods for interactive ultrasound segmentation: The first method is designed for live 3D freehand ultrasound (2D ultrasound with external tracking) examinations and is explained in the following section. In Section 3.2, we discuss how we can achieve interactive segmentation of 4D ultrasound. The technical solutions for the two segmentation methods are illustrated in Figure 3.1. 3D freehand ultrasound relies on external tracking, and our initial solution required a frame grabber for streaming the ultrasound images from the ultrasound workstation to the computer. The 4D ultrasound setup was more minimalistic in terms of hardware, but required more costly




Figure 3.2: The Ultrasound Painting Pipeline. The process starts with selecting a vessel in the 2D image., The vessel is then extracted and processed, where the outline is stored in the 3D model, and new seed points are calculated for the next image.

computation. As of now, there are certain areas where 4D ultrasound is not applicable, making both segmentation of 3D freehand ultrasound and 4D ultrasound relevant.

3.1 3D Freehand Ultrasound

2D ultrasound is still the most widely used ultrasound imaging technique. 2D ultrasound offers a higher resolution and records more frames per second than its 3D counterpart. Certain regions of the body, such as the carotid artery, are very shallow. For imaging these structures, only 2D ultrasound is available. Still, navigation and understanding the spatial relations between structures is very difficult. A live vessel extraction method can provide the user with an understanding of the extent of the previously explored vascular structures.

In the general setup, the examiner locates a vessel in the live ultrasound image. The process, as illustrated in Figure 3.2, starts with the examiner selecting a vessel in the ultrasound image. A segmentation technique extracts the vessel, and the outline is stored in a 3D model. In the next frame, the method processes the previously extracted region to track the vessel from frame to frame. The vessel extraction then iterates as the examination continues. As the user sweeps over the vessel, the 3D model is continuously updated, which provides the examiner with a direct view of the progress.

Extracting the vascular structure from 2D ultrasound was done using an adapted version of the soft segmentation for 3D ultrasound introduced by Petersch et al. [127]. First a distance manifold from the current seed point is calculated, and then a region growing from the seed point is performed until a user defined threshold is reached. The need for a user adjusted threshold may be necessary in cases where either the vessel diameter is changing, or the contrast between the vessel and the surrounding tissue is reduced. An example of the vessels extracted by our method are shown in Figure 3.3.

While identifying a vessel in a single image can be performed by the user, we wish to minimize the interaction from the user, by tracking the vessel automatically as the images are streamed into the system. Due to the high frame rate of 2D ultrasound,




Figure 3.3: Three different vessels are tracked and extracted simultaneously using our vessel-tracking method.

we can assume that there will be an spatial overlap between the vessel from frame to frame. In order to provide the region-growing segmentation technique with new seed points for the next frame, we evaluate the shape of the currently extracted segment. The initial seed point is what constitutes as the center of the vessel, i.e., the point with the highest distance to the vessel wall on a vessel cross-section. In addition, we want the vessel-tracking system to be able to detect branching. When sweeping over a vessel, vessel branching will appear as a gradual shape distortion of the cross-section of the vessel. Therefore, we calculate the skeleton of the current vessel region, and use the end points as seed points for the next ultrasound frame. When processing the next frame, we perform the vessel segmentation for each seed point separately and combine the regions which are overlapping. Regions which are not overlapping will be treated as separate branches.

When the examiner is performing the scan, a 2D view shows the current slice, where extracted vessels are highlighted. A 3D view shows the spatial position of the ultrasound image in relation the previously extracted vessels, as shown in Figure 3.4. This enables the user to visually verify the extracted blood vessels, in order to deem the 3D model as a valid representation of the vasculature. Since the technique aims for live segmentation during an examination, the 3D model needs to be updated on-the-fly as new data is extracted. For each vessel region inserted into the model, we sample the outline of the region. Then we utilized an existing triangulation method [22]. This was easily adoptable to handle local parts of the current 3D model, which enabled live triangulation of the vessels.

Other work has tried to tackle the problem of live reconstruction of vasculature using ultrasound. In contrast to Nowatchin's approach [115], our method extracts the shape of the vessel as seen in the image rather than retrieving an estimation using an ellipse. In addition, our method can detect branching of the vessels without any additional interaction from the user. We utilize the expertise of the users, and allow them to correct the 3D model, by cropping unsatisfying results. The iterative approach




Figure 3.4: Liver vessels extracted using the 3D freehand segmentation method. With a few sweeps, the examiner can extract multiple vessels in the liver.

allows the user to find the optimal viewpoint for the ultrasound probe, and continue adding information to the model.

The live extraction technique serves as an aid in understanding the spatial relations of features while performing ultrasound examinations. However, there is often a desire to reuse the extracted 3D models, either for comparison at a later stage, or to utilize it for pre-operative planning. For ultrasound visualization techniques which base themselves on previously extracted semantic information, there must be a possibility to register the semantic information with the new ultrasound data. Using 2D ultrasound, external tracking is the only viable option, which does not compensate for internal motion. For liver examinations this is especially a problem. The liver is positioned right underneath the lungs. When a patient breathes, the liver is pushed down by the expanding lung. While it is possible to perform examinations at breath hold, the patient might not breath in the same amount every time, meaning the liver will be pushed to a different position for every breath. Our 2D ultrasound segmentation method is more difficult to use in areas where much internal motion occurs, and is more suitable in areas where the internal motion is limited, such as for the carotid artery in the neck.

3.2 4D Ultrasound

With the emergence of high quality 4D ultrasound, new opportunities arise for interactive vessel segmentation. Compared to 3D freehand ultrasound, which can only obtain a cross-section of the vessel for each frame, we now can utilize 3D segmentation algorithms. On the other side, 3D segmentation algorithms come at higher computational costs, which must be considered. We propose to use a fast region growing method from user selected positions, and use image-based tracking to stitch together smaller segments. Structures, such as the liver vasculature, often extend beyond a single ultrasound sector. It is therefore still important to allow for using multiple scan sectors.

The pipeline of our technique, depicted in Figure 3.5, starts with the user selecting a vessel in a slice view. From there we use Petersch et al.'s [127] approach for im-




Figure 3.5: The pipeline for 4D ultrasound segmentation. The process starts with the user selecting a vessel. The extracted region is then visually inspected before storing in the segmentation mask. The user then either repositions the probe, or selects a new point within a vessel.

age segmentation to extract a region of the blood vessel. From testing, we achieved better results using a bilateral filter for noise reduction, rather than the Gaussian filter proposed in their paper. In contrast to the technique in Section 3.1, we do not store the currently extracted region automatically and propagate to the next frame, but allow the user to verify the region visually, both in 3D and 2D to determine whether or not the current segmentation mask is valid. If so, the current vessel-region is stored into a separate volume, which will constitute the final segmentation mask.

Since we use 4D ultrasound we are now able to utilize image-based tracking for positioning of the probe. In our segmentation approach we utilize a tracking technique presented by Øye et al. [121]. This method allows real time tracking of the ultrasound probe using the sum-of-square difference for the similarity measure. Tracking can be achieved in two different manners. Sequential tracking, where the current image is registered with the five previous images continuously, which is also known as multiframe registration. The other method is a global tracking scheme, where the current image is registered with a compound volume. The compound volume can be constructed using the sequential tracking method, where an examiner sweeps over the area of interest, and the ultrasound images are stitched together into one large volume, as shown in Figure 3.6.

The key element for interactive segmentation is to provide a visual feedback to the user. Similar to the 3D freehand technique, we provide a 3D view in addition to the slice view, as shown in Figure 3.6. The 3D view show the location of the current ultrasound image, in relation to the previously extracted vessels. Also, the currently extracted vessel region is rendered using a standard volume renderer, so that the user can inspect the current region for over- or under-segmentation.

As the imaging process behind the segmentation technique is similar to that of Petersch et al., our methods main contribution is the possibility to utilize multiple 3D

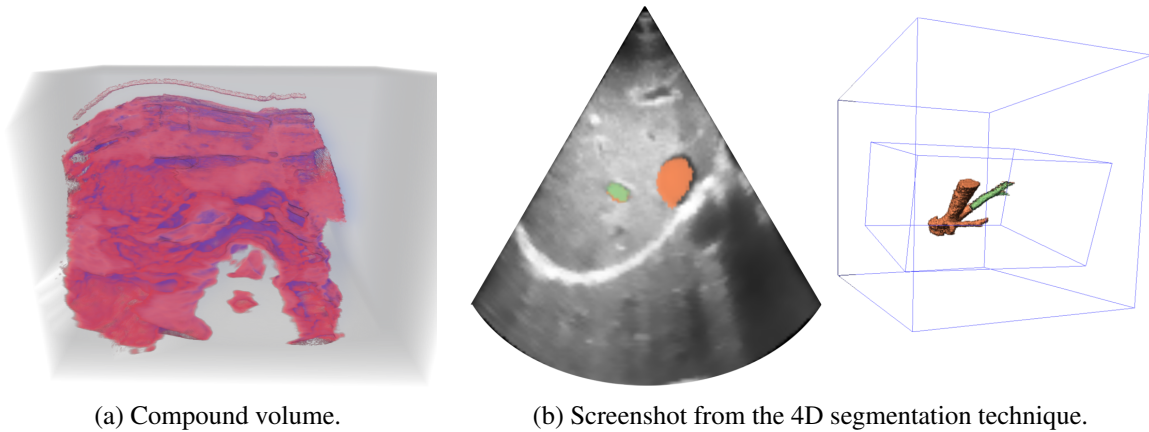


Figure 3.6: (a) shows a compound volume for used for registration. Several 3D ultrasound images are stitched together using image-based tracking [121]. In (b), we see a screen-capture from the 4D ultrasound segmentation tool. The left image shows the B-mode image with the currently extracted region (green) and the segmentation mask (orange) as overlays. To the right, we see the outline of the current 3D ultrasound image as a blue box, compared to the stored segmentation mask. The rendering shows both the segmentation mask in orange and the currently extracted region in green.

ultrasound sectors for extracting the same feature. Utilizing a tracking technique and segmentation technique which both work in real time, the user can verify the result on-the-fly and move the probe to a different location if the image quality is not good enough. The same tracking technique also enables the user to segment structures which extend beyond a single ultrasound sector.

Figure 3.7 shows segmentation masks of blood vessels in the liver taken from three different volunteers. For an experienced user, the average segmentation duration was around five minutes. In the next chapter we discuss how we can utilize these segmentation masks for various visualization techniques targeted for ultrasound data.

The two vessel extraction techniques presented in this chapter both utilize an adapted version of Petersch et al.'s method for image segmentation [127]. While there exist other methods for identifying vessels, the distance-manifold approach has the advantage of controlling the region growing from bleeding into shadow-regions. In addition, a fairly low computational cost makes it ideal for interactive segmentation.

The two segmentation methods have different areas to which they can be applied. 2D ultrasound with high spatial resolution can be used for segmenting out smaller vessels near the surface of the body. Abdominal vasculature suffers from much involuntary motion from respiration. Today's 4D ultrasound probes have a much lower image quality than 2D ultrasound, but can cover larger areas. 3D ultrasound images can also be used for image-based tracking, which compensates for both voluntary and involuntary motion. This has a great impact on reusability, since the current images can be registered with data acquired during a previous ultrasound examination.


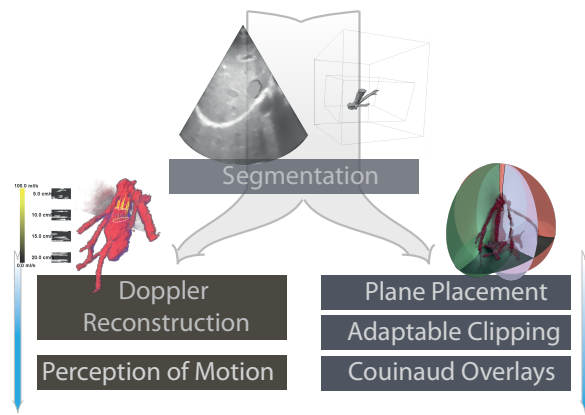


Figure 3.7: Liver vasculature from three healthy volunteers. The segmentation mask was extracted by an experienced user in around five minutes.

Chapter 4



Ultrasound-Tailored Vessel Visualization

With the availability of robust vessel extraction, we can now take advantage of the predefined vascular structures to create expressive visualization techniques. Different approaches utilize predefined structures, such as blood vessels or other features of interest.

For instance, many occlusion solving approaches use predefined segmentation. For instance, *cutaways* is a technique which often relies on segmentation. Inspired by techniques used by illustrators, a feature of interest can be combined with the surrounding context in a meaningful manner [168]. One can also include automatic camera position [166] to provide an optimal view of predefined features. Another approach is cutaways combined with segmentations masks for expressive visualizations [24]. *Ghosting* is a similar technique to cutaways, where the opacity of occluding features is reduced to show the underlying structures [19, 168]. Different approaches translate the occluding data to a non-occluding position such as *exploded views* [20] or illustrative peeling [16].

A different area which can benefit from having predefined vascular structures is the definition of Couinaud-partitioning in the liver. The Couinaud partitioning is mainly defined by four planes, three following the three branches of the hepatic vein and one following the portal vein. Parametrization of the liver vascular tree can be used to define the Couinaud-partitioning. Wand et al. proposed to use predefined vessels to parametrize the branches in order to generate the Couinaud-partitioning [175]. However, this technique is depending on the completeness of the pre-segmented vessels. Manual definition of the Couinaud-planes can be used as a solution when the vessel-segmentation is lacking [120]. Still, it is difficult to comprehend the actual orientation of the vessels by only examining 2D ultrasound. An indication of the 3D extent of the vessels would make the positioning of the planes easier.

In this chapter we will present new ways of utilizing pre-segmented vascular structures, for visualization techniques targeted for ultrasound. In the following section, we present a new method for creating expressive volume visualizations using physics simulation, in a novel concept of *illustrative clipping membrane*. The physics frame work also allows us to utilize extracted vessels to create adaptable Couinaud-partitioning surfaces which follow the shape of the vessels. In Section 4.2, a new method for Doppler visualization method is presented, which uses the direction of the vessels to reconstruct a 3D flow. Finally, in Section 4.3, a statistics-based correction model is introduced. The goal was assess the trends in motion perception to correct for perceptual distortions in speed estimation when using animated objects in flow visualization.

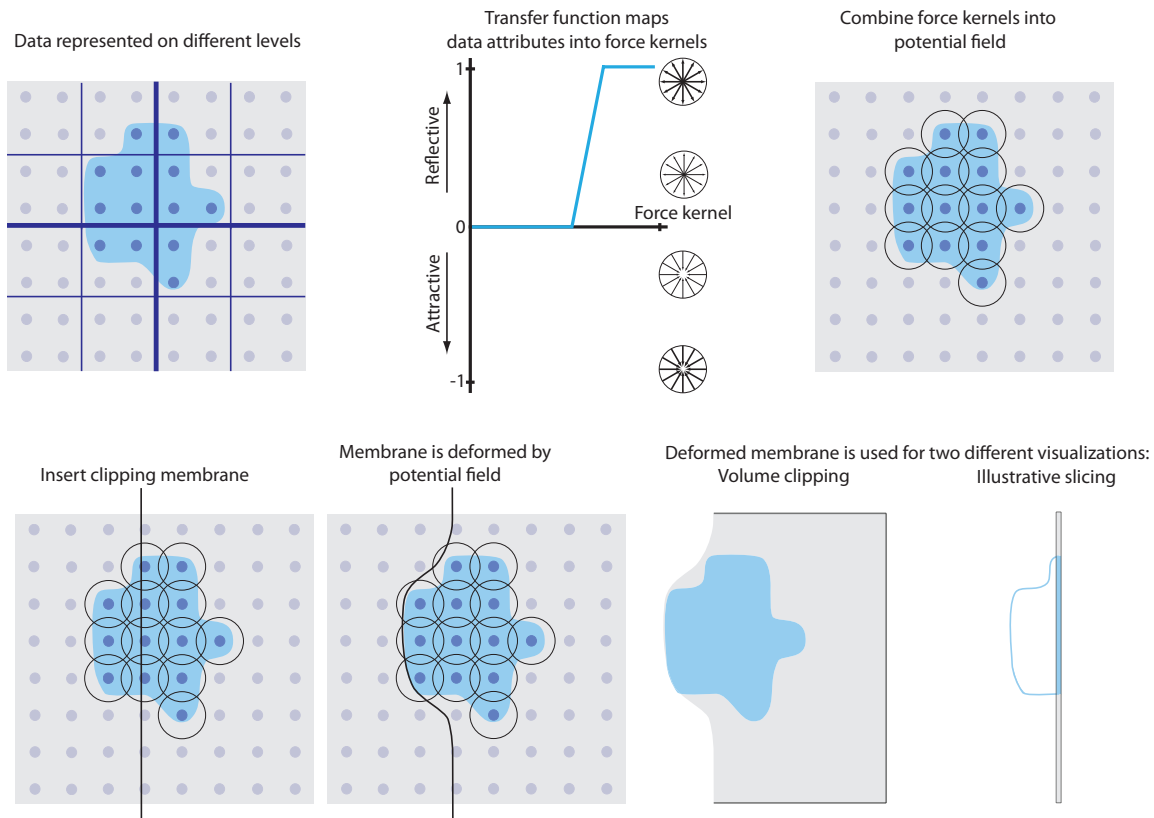


Figure 4.1: The process pipeline for physics-based deformation of the clipping membrane. The transfer function translates the data values to a potential field. Inserting the membrane into potential field, the combined forces repel the membrane to the surface. We can then apply the membrane for both volume clipping and illustrative slicing.

4.1 Adaptive Volume Clipping

Clipping in volume visualization is a powerful tool for solving the problem of occlusion. Clipping is also a concept frequently used by illustrators to show internal structures within their context. However, clipping by illustrators is not done indiscriminately, but rather the metaphorical clipping plane is following the surface of the structures in the vicinity of the plane. Structures are either kept whole or left out completely. In only rare cases is a structure clipped in half, in illustrations.

Traditional clipping is defined by a simple plane. The plane is used in deciding whether or not elements should be visualized by testing whether they are in front or behind a plane. Since structures in volume data are rarely planar, the flat clipping plane would cut through structures and leave smaller fragments in the image. Overall, the clipping might, in cases when it fragments an object into smaller pieces, confuse the viewer.

Inspired by the method used in illustrations, we propose a new method for adaptable clipping by allowing a deformable surface to define which structures to visualize and which not to. The deformation of the surface is caused by the underlying data, preventing the clipping plane to cut through interesting features. The method allows for the user to interactively define which data-types should interact with the structure without the need for any predefined segmentation masks.




Figure 4.2: Regular volume clipping cuts through the horizons in the seismic dataset. Using the deformable membrane, we can get a cut which is aligned with the horizon. As the membrane is pushed through the potential field generated from the data, it *snaps* onto the next horizon.

In essence, we want a surface to behave similar to a cloth, which folds tightly around structures it interacts with. In the same way, we want the surface to interact with different structures in a dataset, e.g. bones, vessels, or muscles. We translate the cloth-metaphor into the domain of volume visualization using a physics-based approach, where we define structure with physical properties similar to a cloth, which we call a deformable membrane. We then translate the data elements into physical actors which can interact with the membrane. Figure 4.1 illustrates the pipeline for our physics-based approach.

In order to translate data elements to physical actors, we apply a technique similar to color transfer functions defining color and opacity in volume rendering. We define physical actors from the voxels by using a potential field transfer function. For each element we insert a *force kernel*, which can interact with the deformable membrane. The membrane, i.e. a mesh where the nodes are connected by springs. The transfer function determines how strongly the force kernels interact with the membrane. The combined force kernels constitute a potential field. The user then inserts the membrane into the potential field and a physics simulation determines the deformation of the surface. Finally, we utilize the membrane in two different manners: volume clipping and illustrative slicing.

As the membrane is deformed to follow structures within the volume, we can now use this membrane to cut through the volume, while we avoid cutting through the features selected by the potential field transfer function. Seismic data is another utilization of ultrasonic imaging for subsurface analysis of geological structures. In seismic datasets, as shown in Figure 4.2, evaluating the shape of the horizons using simple volume rendering can be difficult. Traditional flat clipping tools cut through the horizons and leave only partial information. Applying the adaptable clipping to seismic data, we obtain a much better view of the horizons in the dataset.

Slice rendering is one of the most common tools for analyzing medical imaging. As 3D rendering can provide a good structural overview of the data, occlusion can make it difficult for the viewer to obtain a detailed view of the internal structures. Slice rendering provides great detail without any occluding factors, but the slice provides little information regarding the how structures extend out of the plane. A combination of the strengths from both slice rendering and 3D rendering would have great potential for efficient medical data visualization. One attempt for such a visualization technique is slab rendering, i.e. a smaller subset of the volume is defined by two parallel clipping plane. Similar to regular clipping, slab rendering suffers from indiscriminate clipping through structures. Our deformed membrane folds around structures in the vicinity of a slice. This can then be used as a selective means for adding context into slice rendering.

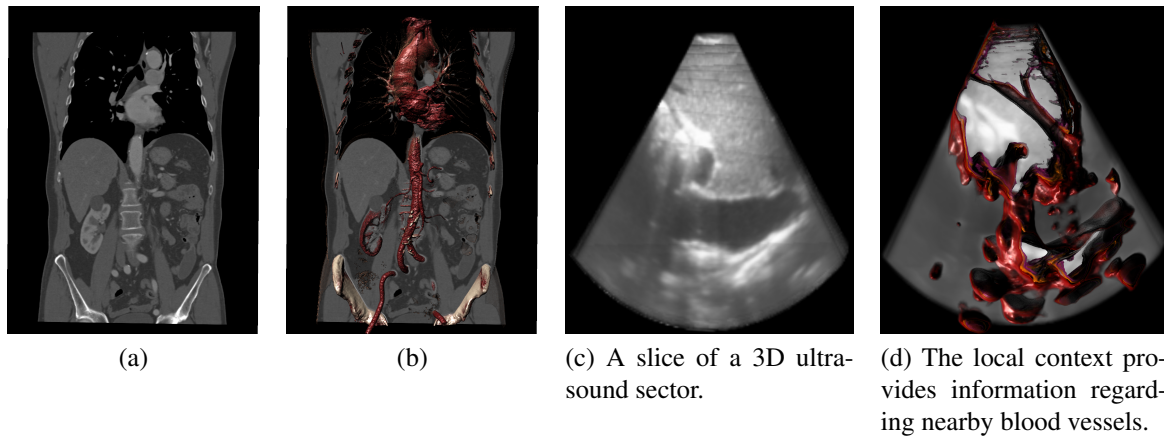


Figure 4.3: Using the deformed membrane, we can add context to slice rendering for different types of medical data. (a) show slice of a CT dataset. In (b), context is added in form of bones and blood vessels. (c) and (d) show a similar case for 3D ultrasound data

4

We suggest to use volume rendering between the deformed membrane and the original slice position. This way we can create expressive slice rendering with a minimal level of occlusion, as seen in Figure 4.3.

One advantage from our clipping technique compared to sketch-based techniques, is that the need for interaction is reduced. After the potential field transfer function is defined, the user can browse through the volume by moving the slice position. The clipping membrane interacts with whatever data elements have been included by the transfer function. Recent work combine volume clipping with *on-the-fly* segmentation [34]. Here, the user starts with a flat clipping plane or a basic slice rendering which intersects the data. The user are then able to *pull* structures out of the flat clipping surface or the 2D slice. The segmented structures are then extruding from the flat cut, or slice. However, unlike our method, they have to define each separate part which should extrude from the surface. As we allow for a more global definition of the interesting features, using the potential field transfer function.

Our physics based approach has additional application areas. The examples shown in Figure 4.2 and 4.3 generate a potential field which to expel the membrane out from the structures. However, we can change the sign of the force kernels to create attractors rather than repellents.

Couinaud Segmentation

Previous work proposed to use enhanced slice rendering with Couinaud-partitioning overlays, for guidance in 2D ultrasound examinations of the liver [120, 169]. This method has been adopted in our system for *high quality 3D visualization of in-situ ultrasonography*. In this case, the Couinaud planes were positioned manually, aided by a 3D visualization of the vascular structures. However, in the reported case we utilized flat planes, whereas the actual vessels, which define the Couinaud segmentation, are curved.

Our system allows us to use predefined structures, such as vessel segmentation masks, to affect the deformation of the membrane. This way, we can create surfaces




Figure 4.4: Using attractors within the vessel, the deformable membrane aligns with the vessels. This can then be used to define the Couinaud-partitioning more accurately than using planar surfaces.

which follow the vessels more accurately. As we want the membrane to be aligned with vessel, we insert attractors at the center-line of the blood vessels. Then as the plane is positioned, the attractors pull the membrane towards the center. In Figure 4.4, we have positioned the three vertical planes aligned with the hepatic branches, in addition to the horizontal plane defined by the portal vein. Unlike Wand et al.'s method [175], our technique will not fail if there are any missing information regarding the blood vessels. If only partial segmentation of the blood vessels is available, we can still utilize this for local deformation of the membrane.

4.2 Blood Flow Reconstruction

While extracted vessels can be used for a multitude of rendering techniques they can also provide us with semantic information. Doppler ultrasound is an imaging modality which can benefit from added semantics regarding the vasculature.

Doppler ultrasound is a frequently used imaging modality for analyzing blood flow in a patient in-situ. It is used for a wide range of medical scenarios, such as diagnosis, prevention and monitoring of diseases. Comprehending the blood flow from the Doppler signal is difficult for a novice user, and much training is required for efficiently utilizing this live imaging modality. In addition, Doctors are trained to use Doppler ultrasound with the angle between the blood vessel and ultrasound probe not exceeding 60° . Thus, requiring the user to be skilled in navigating the ultrasound probe to obtain the right angle between blood vessel and ultrasound probe.

Blood flow is 3D phenomenon, but the measurements from Doppler ultrasound are only velocity projections towards or away from the ultrasound probe, i.e. a single component of the 3D flow. This means the examiner has to mentally reconstruct the actual

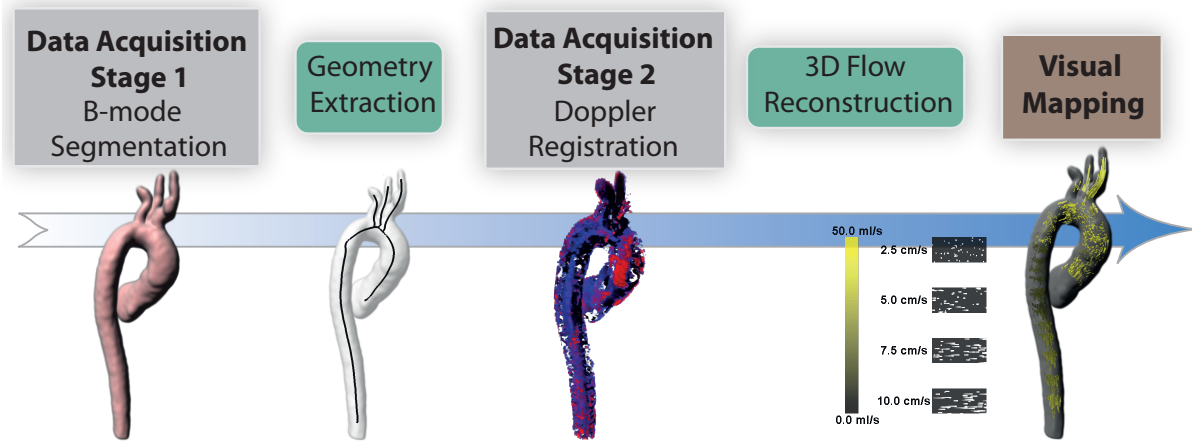


Figure 4.5: 3D blood flow reconstruction starts with acquiring a segmentation mask of desired vessels. From the segmentation mask, we extract information about the geometry in form of the assumed direction of the blood flow. The Doppler signal is registered with the vessel-geometry, which allows us to reconstruct the 3D blood flow on-the-fly. Finally, we visualize the newly generated 3D flow field using a tailored visualization approach.

4

flow, by compensating for the direction of actual flow. To estimate the direction of the actual flow, the examiner must evaluate the structure of the blood vessels simultaneously to interpreting the color encoding of the Doppler signal.

The main goal was to provide a in the usage of Doppler ultrasound in two ways: Firstly, providing a good indicator of the angle between the blood vessel and the ultrasound probe to the help the examiner with navigation. Secondly, relieve the examiner from mentally reconstructing the actual flow by compensating the Doppler signal with the angle between the blood vessel and the ultrasound probe.

To help us with this challenge, we can utilize predefined models of the vasculature for Doppler visualization. The added models have two main purposes. To allow the user to visually navigate the ultrasound probe to an appropriate position, and to add semantic information which we can use to reconstruct the 3D blood flow from the single-component Doppler ultrasound. The reconstruction process is based on the assumption that the blood flow is in general laminar and that blood flows along the direction of the vessel.

The pipeline for this technique consists of two data acquisition stages. In the first stage, we acquire a 3D model of the desired vascular structures from the segmentation process. This is achieved by utilizing the method described in Section 3. From the segmentation mask, we extract the geometric information regarding the blood vessel, based on the vessel center-line. In the second stage, we record the Doppler signal. To register the Doppler signal with the predefined vessels, we use the corresponding B-mode image, which is recorded simultaneously with the Doppler image. The B-mode image is used for tracking the ultrasound probe, using the image based tracking method by Øye et al. [121]. Knowing the position of the ultrasound probe, provides us with the information regarding which direction the Doppler signal is measured. This allows us to back-project the Doppler signal onto the vessel directions, generating an unsteady 3D flow field. Finally, we represent the 3D flow field with a tailored visualization design.

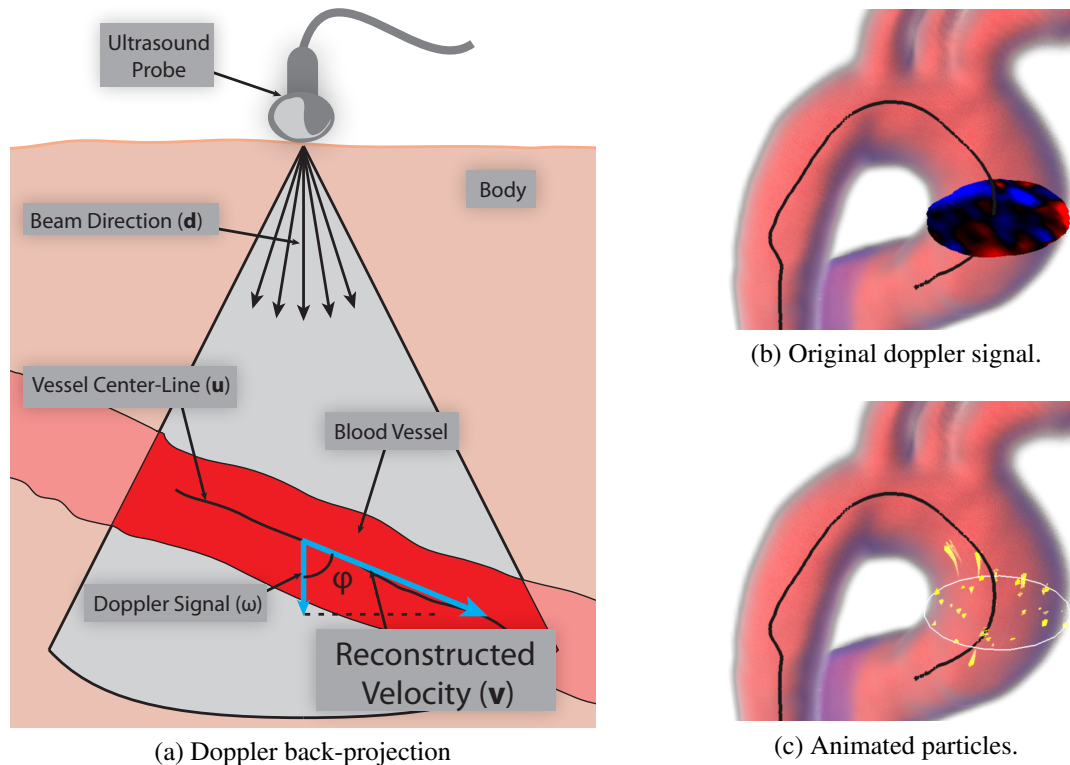


Figure 4.6: The velocity reconstruction is achieved by back-projecting the Doppler signal onto the vessel direction, as shown in (a). We can then translate the Doppler signal in (b) to the moving particles in (c)

A key element of the flow reconstruction is to estimate the direction of the flow. We assume that the flow in blood vessels is mostly laminar, meaning it follows the general direction of the current vessel. We estimate the vessel-direction to be aligned with the direction of the center-line of the vessel.

To extract the center-line, we apply a center-extraction method which is robust to noise [63]. This method allows for rapid extraction of the center-line and simultaneously provides the information of connectivity. In short, it propagates a wave from a point which is known to be a center-line point, namely the point with the largest distance to the vessel-wall. The velocity of the wave is determined by the distance to the wall, meaning the further it is from the wall the faster it propagates. To extract the center-line it then backtracks on the minimal cost path from topological significant positions back to the origin of the wave.

To be able to perform the reconstruction in real time, we need a fast way to acquire flow-direction. From the center-line of the vessel, we then generate a lookup-table, i.e. a vector field, which holds the flow-direction for each point in the vessel. This allows us to quickly extract the flow-direction at any given position.

In Figure 4.6, we see an illustration of the flow-reconstruction method. From the Doppler image, we get a value, ω , stating how fast the blood is moving along the beam-direction, \mathbf{d} . The beam-direction is calculated as the position of the probe to the current point. We then use the estimated flow direction, \mathbf{u} , and calculate the velocity, \mathbf{v} , as follows:




Figure 4.7: The flow inside the inferior vena cava from two different healthy volunteers.

$$\mathbf{v} = \frac{\mathbf{u}}{|\mathbf{u} \cdot \mathbf{d}|} \cdot \omega \quad (4.1)$$

The reconstruction happens in real time, and generates an unsteady flow field. This opens up a visualization challenge for the newly acquired data. To visualize the reconstructed flow field, we utilize a particle-based flow visualization technique. The user selects cross-sections along the center-line which then seed particles into the flow field. The motion of the particles depicts both the speed and direction of the particles. Since we use motion to depict velocity, we have the possibility to encode more information using color. With estimation of the actual blood flow, comes the possibility to calculate more information regarding the blood flow. For instance, by integrating the velocities over a cross-section of the vessel, we can get an estimation of the flux, i.e. blood volume over time. We can then encode both velocity and flux simultaneously, which can be important in transplantation surgery, where blood flow characteristics can help the surgeon to complete the procedure successfully.

Since the reconstruction is an estimation of the flow, we allow the user to examine the original Doppler signal by probing along the center-line of the vessel, as shown in Figure 4.6. The user has then the opportunity to evaluate the detail in the Doppler-image while the particles provide an overview of the flow.

From the users perspective, estimating velocities from the motion of the particles requires a certain point of reference. This could be the actual size of the blood vessels, but relating it to a certain scale might be difficult. When color coding is used, there follows a legend which describes what color relates to which value. Similarly to a color-legend, we propose to use velocity-legends which provide several indicators of discrete velocities, and which screen-space velocity it relates to.

In Figure 4.7, we see the result from two different healthy volunteers. The motion of the particles shows the direction of the blood flow. Using a color-coded Doppler rendering technique [74, 126], this information must be mentally constructed. While there exist other methods to reconstruct the blood flow which use multiple 3D ultrasound probes [57], our technique only requires a single ultrasound probe, reducing the technical requirements. Adding flow visualization, we contribute with a new way of looking at Doppler ultrasound. The extracted vasculature provides an excellent means

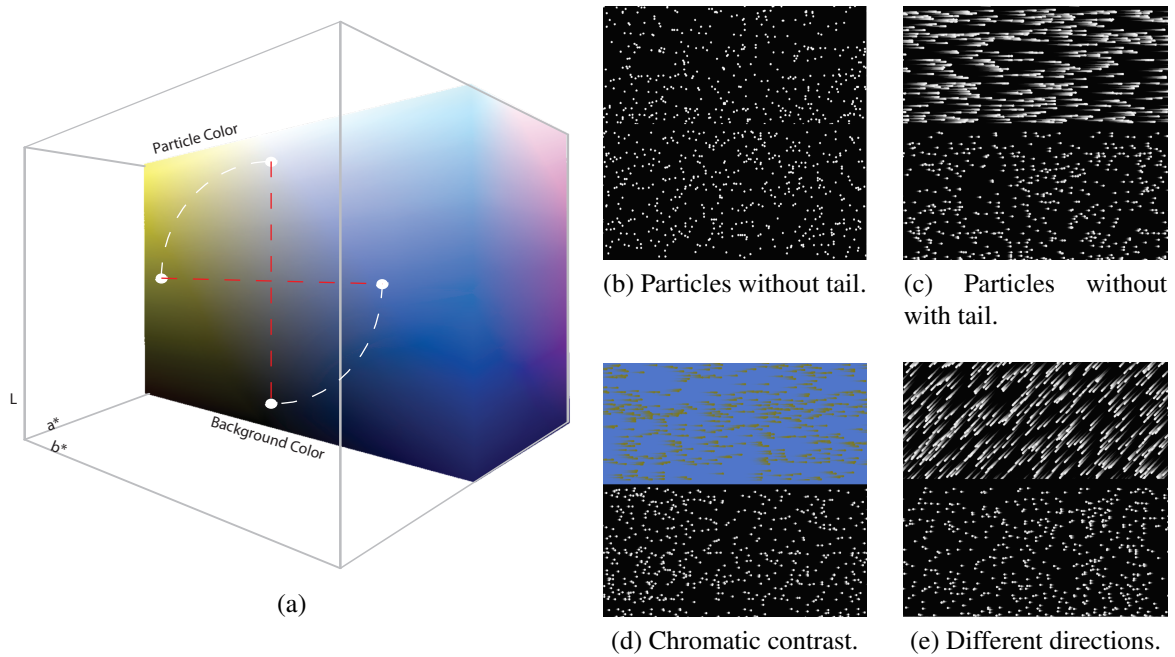


Figure 4.8: (a) illustrates how the contrast-parameter is changed from achromatic to isoluminant. (b), (c), (d) and (e) shows screenshots from the tasks with the different parameters.

for navigation.

Using motion to depict velocity is seen frequently in flow visualization, but there is no clear understanding of how well we can estimate velocity from motion. In the next section, we present our work on analysis of motion perception.

4.3 Perception-based Flow Visualization

A trend in visualization is aiming to use rendering techniques which represent the physical behaviour of the phenomenon. In shading and illumination, this is prevalent with trying to simulate how light propagate through a medium or reflects of a surface. However, as Ben Schneidermann stated, «The purpose of visualization is insight and not pictures», meaning we should focus on the user interpreting the conveyed information correctly, rather than the correctness of the data representation. The idea of simulating physical phenomenon is that this is close to the mental model of the viewer. However, there might be the case that the perception of objects is not well-aligned with the way we represent them. For shading and its erroneous perception, this was demonstrated by Šoltészová et al. [170].

The same analogy can be transferred to the use of animation in visualization. Particularly in flow visualization, animated particles are used as a means for depicting velocity at the current position in the flow field. Here, the velocity of the particle is scaled linearly to the underlying data. Similar to the shape perception, the perception of speed from moving particles, might not correspond to the physical phenomenon of motion.

We have performed a series of perception studies to evaluate the trends in speed estimation based on moving particles. While there are many parameters which can




Figure 4.9: The pipeline for generating the compensation model. First we executed a pilot study to investigate the proper range of each parameter. Then, we executed an initial study which tested the effect of each parameter individually. From the trends in estimation error, we calculated a compensation model, which was tested and refined over two iterations.

influence the perceived speed of particles, we choose a subset of parameters which are seen in most particle-based visualization techniques:

- **Contrast type:** We wanted to investigate if the change in contrast type is affecting the perceived speed for moving particles. The contrast-parameter was then set to be the angle of rotation between two points in CIELAB color space, as shown Figure 4.8.
- **Relative speed-up factor between particles:** As the relative speed difference between two particles is increasing we assume there would be an increase in estimation error.
- **Global scale of velocities:** When watching objects moving very fast in our visual sector, it becomes blurry, for instance when we look out the window of speeding train. We assume that this affects the perception of speed. We included a parameter which globally scales up the velocities of the particles.
- **Direction:** There are asymmetries in the human visual system when it comes to the horizontal and vertical axis. We therefore chose to test for the change in estimation error compared to the direction of the particle motion.

The goal was to create a correction model which can compensate for any systematic errors in our perception. We started by investigating the estimation error when evaluating velocity visually from moving particles. Following the workflow shown in Figure 4.9, the investigation started by creating a study, where the subjects were asked to provide the relative speed-up factor between two sets of perceptual parameters. The trends in estimation error would then provide us with the basis for a compensation model, correction for distortions in the perception. We started with a pilot study to find the appropriate range for each parameter. From there, we continued with an initial round where each parameter was isolated. Based on the trend for each parameter, we constructed a compensation model of the statistically significant parameters. We then performed a second study which tested the effect of the compensation model. The results from the second study showed a great improvement in the impact on the trends in estimation error from each parameter, but there was a constant under estimation. This constant was then included in the compensation model and a final study was performed. From the result of the last study, we found a similar improvement to the parameter impact on estimation error, and also an improvement in estimation.

Perception Studies

After the pilot study, we executed an initial study with 22 people. Here we compared the estimation of relative speed-up factor, which was tested for each parameter at a time. For each parameter we analysed the results for trends. The results showed correlation between estimation error and three of the four parameters: global scale of the velocities, relative speed-up factor, and direction. Interestingly, we did not find a correlation between estimation error and contrast-type, despite the reports in previous work [155, 180]. From the three parameters which seemed to have an effect on estimation error, we created a three dimensional compensation model. Although we could only find a weak correlation between direction and estimation error, we choose to include it in the next study-round.

In the second round, we wanted to test the effect of our compensation model. For this we created a new set of tests combining the three parameters in the compensation model. The new round was executed with ten new participants. From the results, we found that the direction-parameter had no statistically significant effect on the estimation error. However, we managed to reduce the impact of the global scale-parameter and the relative speed-up factor by around 90 percent. While the parameter impact was lowered, we found a constant level of under estimation. We decided to remove the direction-parameter from the compensation model and add a factor to compensate for the constant underestimation and perform an additional study for the corrected compensation model.

The third and last study showed a similar level of improvement of parameter-impact. The average level of underestimation was also reduced a significant amount from the previous study. The total amount of estimation error was -17.8 percent. Figure 4.10 shows the estimation error versus the significant parameters, with and without the compensation model. The regression curves show a clear improvement in impact for the two parameters.

The resulting compensation function is as follows:

$$\begin{aligned}
 E_{\sigma}(\sigma) &= a_{\sigma} \cdot \log(\sigma) + b_{\sigma} \\
 E_s(s) &= a_s s^2 + b_s - (a_s + b) \\
 C_{\sigma}(\sigma) &= 1 + E_{\sigma}(\sigma) + E_{\sigma}(1) \\
 C_{\varepsilon} &= 1 + \varepsilon \\
 C_s(s) &= 1 + E_s(s) \\
 F_c(\sigma, s) &= \frac{1}{C_{\sigma} \cdot C_d \cdot C_{\varepsilon}}.
 \end{aligned} \tag{4.2}$$

Constant	Value
a_s	-0.010414
b_s	0.02680727
a_{σ}	0.48538867
b_{σ}	0.29785412
ε	-0.26

The two functions $E_{\sigma}(\sigma)$ and $E_s(s)$ are the functions describing the trend in estimation error for global scale of velocities, and relative speed-up factor respectively. $C_{\sigma}(\sigma)$ and C_{ε} are their corresponding correction functions. The combined correction function $F_c(\sigma, s)$ is then composed by the two correction functions, as well as correction for the constant, $C_s(s)$, which was discovered in the second study.

This initial investigation of perception of animated particles suggests that there are certain aspects which influence the perceived relative motion. However, contrary to previous work, which indicated the contrast-type would affect the perceived speed, we could not reproduce these results. One reason might be that the distortions from

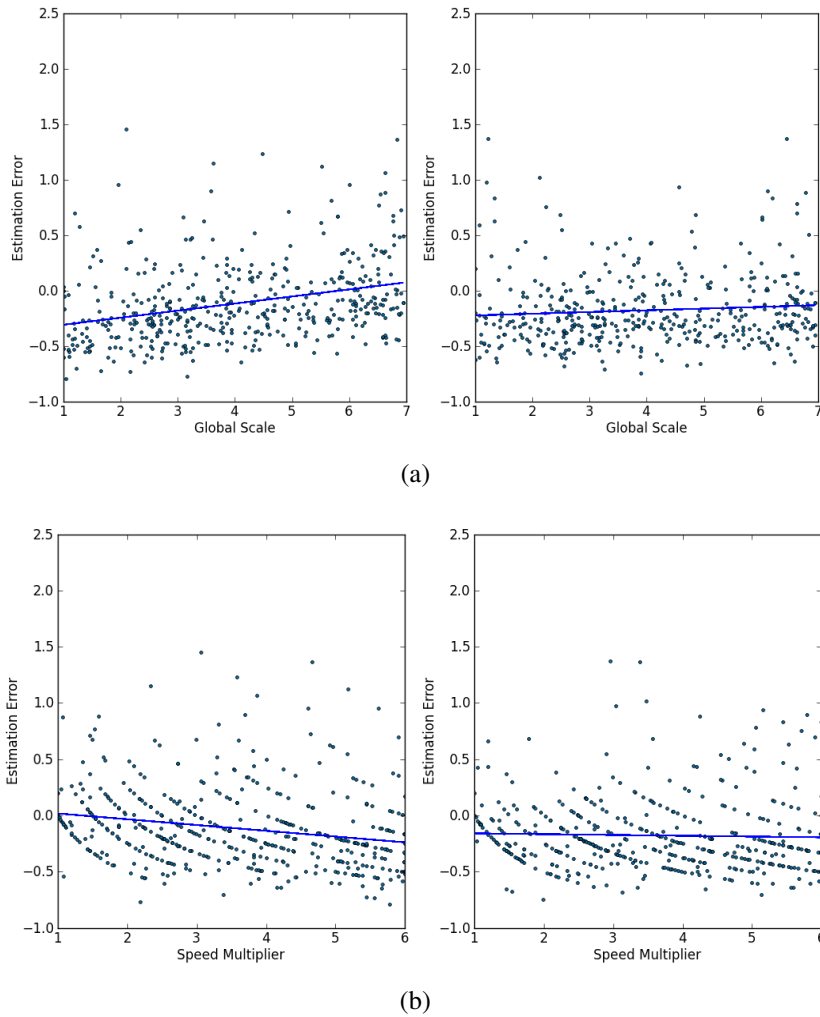


Figure 4.10: The graphs show the regression curves for the two significant parameters. The two scatter plots on the left show the estimation error without the compensation model and the scatter plots on the right show the error when our compensation model was applied.

contrast-type is caused by a continuous change in contrast, which was not present in our study. For instance, a continuous change in luminance can cause a perceived motion in stationary objects [103], and might be the reason for the controversy seen in related work [18, 158].

When one applies the compensation model to a real world scenario, there are certain aspects which should be considered. Firstly, since the compensation model changes the velocity the particle, the spatial vs. temporal position becomes more distorted from the *ground truth*. This means the compensation model should be applied when speed and direction is the more important than the arrival time of the particles. Furthermore, the compensation model also requires a certain velocity to compare the data-velocities to. Therefore, a basis is needed, such as the velocity legends presented in Section 4.2, which provide a screen-space velocity to relate to.

The parameter space for animated objects is large, and more work needs to be done to create the complete perceptually uniform motion space.

Chapter 5

Conclusion and Future Work

The main goal for the work done in this thesis was to develop visualization techniques which would improve the readability and usability of medical ultrasound. We especially targeted for abdominal examinations connected to the liver vascular structures. We have aimed to improve navigation within a complex anatomical feature, by adding structural or other semantic information to 3D freehand and 4D ultrasound. Our work also aimed to improve the readability of Doppler ultrasound for non-expert users by using structural information regarding blood vessels to estimate the direction of the blood flow. Finally, we wanted to provide a mean for using animated particles in visualization in a perceptually correct manner.

Achievements

The contributions from this thesis have been divided into two parts. One part describes targeted extraction of structural information from 3D freehand ultrasound and 4D ultrasound. In this part, we have showed how a tracked 2D ultrasound image can be used to extract vascular structures live during an examination. The extracted blood vessels provide an aid for the user in navigation of the ultrasound image, since it allows the user to see how the vessels in view are extending beyond the 2D plane. However, we found that involuntary motion, such as breathing, combined with an external tracking method makes it difficult to reuse the extracted vessels. This vessel-extraction method is more suited for areas where involuntary motion is not as prominent, such as the carotid artery or vessels in the extremities. In such vessels, the involuntary motion is limited to the pulsation caused by the heart, which is unavoidable. One should also note that imaging of such vessels is primarily done using 2D ultrasound probes with a flat linear array, due to the high spatial and temporal resolution. This makes it clear that 2D ultrasound is still very relevant in medical scenarios.

In addition to 3D freehand ultrasound, we have shown how 4D ultrasound can be used for interactively extracting vessels. By utilizing image-based tracking we have proved to increase the applicability of ultrasound for vessel segmentation in three ways. First, we can compensate for involuntary motion caused by breathing. Second, it also allows us to extract vessels beyond the size of single ultrasound sector by stitching together multiple 3D ultrasound images. Finally, it allows us to compensate for attenuation artifacts, such as shadows cast from bones, by allowing to use ultrasound images from different probe-view points. An added benefit to image-based tracking is the fur-

ther utilization of the extracted features.

The second part is focused on how we can use these extracted vascular structures in new ways for ultrasound visualization. We have shown how a physics-based approach can be used for adding context to slice rendering and for creating feature preserving volume clipping.

For navigation, other work has shown how to display the Couinaud segmentation as overlays on 2D ultrasound images. The position and orientation of the segmentation is defined by four planes which follow the three hepatic branches as well as the portal vein. The physics-based framework offers the ability to use structural information to create deformed surfaces which can follow the vessel more accurately than planes.

Doppler ultrasound is a very powerful imaging modality for live blood flow imaging. However, using only a simple color-coded rendering scheme makes it inherently difficult to read, since the value only depicts one projected component of the 3D blood flow. We have shown that by using predefined vessels, we can make assumptions regarding the direction of the blood flow. This allows us to back-project the Doppler signal onto the assumed flow direction, and reconstruct the actual blood flow velocity. Finally, this allows us to display Doppler ultrasound using a particle-based flow visualization technique, making Doppler ultrasound more readable without extensive training.

Using animated particles for depicting a velocity, we should know how the user will perceive this visually conveyed velocity. Based on a perception study, we have found systematic trends in perceived speed of moving particles. From these trends we presented a compensation model, which can correct for the distortions. This has then been applied to the Doppler ultrasound visualization technique.

Future Work

Building on the new advancements in visualization techniques targeted for ultrasound data presented in this thesis, we can increase the potential usage of ultrasound imaging in several directions.

Less manual vessel segmentation: We would like to investigate further segmentation methods for ultrasound. While we are certain visual verification is a key element to robust ultrasound segmentation, investigations into reducing interaction can be of great benefit for vessel extraction. For instance, we can potentially develop a less manual propagation for our 4D ultrasound segmentation method, where seed-points are generated by processing the already extracted vessels.

Segmentation of other structures: While blood vessels have a quite distinct appearance in ultrasound images, other features such as the liver parenchyma, are not as easy to distinguish from other tissue. The surface of the liver is quite important for medical scenarios, such as liver resection. However, segmentation from ultrasound is difficult using anything other than manual segmentation techniques. Perhaps a physics-based approach can be helpful for rapid approximations of the liver organ and other structures.

Different flow models: Our Doppler reconstruction method assume a laminar flow aligned with the vessel center-line. However, blood flow through vessel branches and highly curved vessel may create a more turbulent flow. Applying a more complex

model for estimating the flow direction or a consecutive flow simulation, might improve the accuracy of the flow reconstruction. This could be calculated using computational fluid dynamics, based on the vessel segmentation, but requires further investigations on how to include the information from Doppler ultrasound.

More complex motion compensation models: Visualizing Doppler ultrasound using animated particles instead of traditional color coding allows us to see more than just the velocity. Our investigations into the perception of moving particles show that the user estimation is not completely accurate, but our compensation model is a starting point for correct usage of moving particles to depict data attributes, such as velocity. We should investigate the effect of more parameters, such as particle size, particle shape, density of particles, and motion direction in 3D, to create a more complete compensation model.

Part II

Scientific Results

Paper A

Ultrasound Painting of Liver Vascular Tree

Åsmund Birkeland and Ivan Viola

University of Bergen

Abstract

In treatment planning and surgical interventions, physicians and surgeons need information about the spatial extent of specific features and the surrounding structures. Previous techniques for extracting features, based on magnetic resonance imaging and computed tomography scans, can be slow and cumbersome and are rarely used by doctors. In this paper we will present a novel approach to extract features from tracked 2D ultrasound, in particular hypo-echoic regions such as blood vessels. Features are extracted during live examinations, removing the need for slow and cumbersome post-scan processes and interaction is based on the natural interaction techniques used by doctors during the examination. The ultrasound probe is utilized as a 3D brush, painting features in a 3D environment. The painting occurs during a regular examination, producing little extra interaction from the doctor. We will introduce a novel approach to extract hypo-echoic regions from an ultrasound image and track the regions from frame to frame. 3D models are then generated by storing the outline of the region as a 3D point cloud. Automatically detecting branching, this technique can handle complex structures, such as liver vessel trees, and track multiple regions simultaneously. During the examination, the point cloud is triangulated in real-time, enabling the doctor to examine the results live and discard areas which are unsatisfactory. To enable modifications of the extracted 3D models, we present how the ultrasound probe can be used as a interaction tool for fast point cloud editing.

A.1 Introduction

In recent years, sophisticated visualization techniques for liver examination and pre-operative planning have been developed to aid in liver treatment, such as tumour resection and living donor liver transplantation. The complexity of the liver and the vascular




Figure A.1: On the right we see an ultrasound image with multiple intersecting vessels. Each vessel is tracked from a single user selected seed-point and stored as a 3D model shown on the right. Complex structures are extracted *live* from tracked 2D ultrasound with little required interaction.

structure within, make it difficult to get a clear view with simple slice based visualization techniques. 3D visualizations of the different structures within the liver can provide a better overview and tools based on 3D images for surgery training and pre-operative planning exist to help teaching surgeons in training handle the complexity of the liver [9]. Many different liver treatment planning techniques are based on a liver segmentation scheme called Couinaud segmentation. In the late '50s Claude Couinaud [28] suggested that the liver can be separated into eight different parts divided by the hepatic and portal veins. For complex procedures, such as liver donation, it is vital to know the location of these veins. Hence there is a high demand for vessel-extraction algorithms in medical imaging.

There are techniques developed for defining the Couinaud segmentation automatically in 3D images of the liver [175]. Since the branching of the portal and hepatic veins essentially define the segments, the vessel tree can be parametrized and the spatial extent of each segment can be calculated. These techniques require a pre-extracted vessel tree from the underlying data. Current vessel extraction algorithms are mostly applied to computed tomography (CT) data due to the consistency of the data values. Since the values are typically given in normalized Hounsfield units, they are stable for a certain tissue type across individuals. This consistency make it easier to automatically extract the blood vessel as one can define which value to include in the segmentation. Although CT produces consistent values with high resolution, it is potentially dangerous to use. CT scans sends out X-rays through the body which produce a large amount of radiation, and are therefore not desirable to use frequently. Magnetic Resonance Imaging (MRI) can also be used for feature extraction. MRI does not produce any radiation and is in that sense regarded as a non-invasive imaging technique, but it requires contrast agents to be able to clearly detect the vessel trees. The contrast agent can be considered invasive as a foreign substance is inserted into the blood stream. MRI scans may also suffer from respiration artefacts, since the patient is breathing during the scan. MRI scans are very expensive and due to the high cost of MRI hardware it is not available in every hospital.

Ultrasound is a modality not so often used for feature extraction. Still, it is one of most frequently used tools for medical imaging. Unlike CT, ultrasound transmits sound

waves into the body and records the echo produced by the internal organs. Sound waves have no proven dangerous side effects when used within the regulations given by Food and Drug Administration, which enables the physician to perform scans as often as required. MRI and CT data are usually examined after acquisition and if vessel-extraction techniques fail due to scanning artefacts in the acquired data, the acquisition process has to start over again. This is costly, time consuming and sometimes dangerous. A major winning factor with ultrasound is *live* examination. The physician is examining the images as they are generated. Because ultrasound has no proven dangers, he is able to continue the examination until he is satisfied with the images generated.

Vessels in the liver can be thin and vessel detection algorithms require a high resolution image. Compared to CT and MRI, 2D ultrasound has a very high temporal as well as spatial image resolution. High resolution images and consistent blood values are a key element for vessel extraction. In ultrasound imaging, blood vessels are hypo-echoic. This means that a blood vessel does not give any significant echo, and blood remains consistently black in the ultrasound image. Still, the ultrasound suffers from a low signal-to-noise ratio and refraction from small structures in the body produce image speckle which have made ultrasound often considered not suitable for feature extraction. With the approach presented in this paper, we aim to change this misconception of ultrasound data.

In this work we propose an interactive approach to extract anatomical structures during live ultrasound examinations. We focus our technique on the complex vascular structures in the liver. We show how hypo-echoic regions, such as blood vessels, are detected from a user selected seed point. For the next frame in the ultrasound scan, new seed points are calculated based on the area of the blood vessel from the previous frame, to enable automatic vessel-tracking. As liver vessel trees are branching out, we propose a technique which automatically detects the separation of a vessel-region into multiple separate vessel branches.

Live acquisition enables the doctor to examine the 3D model as it is generated. Areas with sparse sampling have less certainty for precise feature representation. We apply a segmentation-certainty visualization by highlighting sparse regions. The doctors can then move back to the marked area for re-sampling. In some cases, vessel tracking may be lost. The tracking can be re-obtained by moving the probe into a previously extracted vessel. In some cases the extracted vessel can be less satisfactory. We introduce a technique for editing the generated 3D model, using the ultrasound probe as an intuitive selection tool. The selection is based on a tree-graph of the vessel tree, making it possible to select branches or entire sub-trees of the 3D model.

In the following section, we discuss previous techniques, and in Section [A.3](#) and [A.4](#) we describe the process going from tracked 2D ultrasound to 3D models during live examinations. In Section [A.5](#) we present the results from the proposed technique and in Section [A.6](#), we discuss future improvements and applications.

A.2 Related Work

A driving force pushing research in 3D liver vessel segmentation are emerging new techniques for liver resections, setting new requirements for the preoperative imaging and planning [104, 116]. Virtual liver surgery planning system use high-level image

analysis algorithms and virtual reality technology to help physicians find the best resection plan for each individual patient [134]. In a seminal paper in this field, Selle et al. propose a region growing based approach for liver vessel segmentation based on computed tomography (CT) scans [145]. Their approach comprises region growing for vessel segmentation, skeleton extraction, and transformation into a graph for subsequent shape analysis. Using this information, intra-hepatic vessel systems are identified and visualized with graphics primitives fitted to the skeleton to provide smooth visualizations without aliasing artefacts. Oeltze and Preim propose a refined visualization of vasculature, producing smooth transitions at branchings and closed, rounded ends by means of convolution surfaces [119]. In [97], Li et al. propose an algorithm incorporating both spatial and temporal information of a propagating front to advance the segmenting contour, whereas Beichel et al. employ a vessel enhancement filter for segmentation of liver vessels in CT data [11]. Small branches, however, are difficult to segment due to noise and partial volume effects. A similar strategy to the work presented by Beichel et al. has been suggested in [78]. Recent approaches on automated extraction of liver vasculature include adaptive region growing techniques [146], graph cuts methods [38, 66], segmentation using graphics hardware [37], atlas based methods [83]. See [82] for a comprehensive review on vessel segmentation techniques.

However, there is also research on vessel extraction and tracking techniques from sonographic B-mode data. Gooding et al. propose second-order shape measurements for the detection of mammary ductal structures in 3D breast images [59]. In [115], Nowatschin et al. present a system for automatic analysis of intra-operative B-Mode ultrasound images of the liver. The system is part of an assistance system for open liver surgery. It supports 3D-ultrasound imaging and automatic segmentation of vessel structures in the intra-operative ultrasound images. With this segmentation results, an ultrasound based 3D-model of the vascular structure is extracted and manually registered to a preoperative CT-based model. Anatomical landmarks like bifurcations of vessels are automatically extracted in the ultrasound-based model and used for an automatic registration process in further developments.

Lange et al. propose an approach aiming at an accurate navigation in liver surgery via intra-operative 3D ultrasound [92]. Ultrasound data are augmented with preoperative anatomical models and planning data as an important additional orientation aid for the surgeon. Their approach for fast intra-operative non-rigid registration of the preoperative models to the ultrasound volume is based on the vessel center lines and consists of a combination of the Iterative Closest Point algorithm and multilevel B-Splines. In [123], the authors present a method to register a preoperative MR volume to a sparse set of intra-operative ultrasound slices. The spatial relationship between ultrasound slices is obtained by tracking the probe using an optical tracking system. Their registration algorithm converts the intensity values of the MR and ultrasound images into vessel probability values. The registration is then carried out between the vessel probability images. Zhang et al. [186] propose a method to register a real time 3D ultrasound volume to a 2D cardiovascular MR image utilizing a local phase presentation as an image descriptor. This information is then employed in a multi-scale registration process to estimate the global affine transformation using a differential technique. A similar approach is utilized in [60] for registration of multi-view 3D echocardiographic sequences. W. Wein et al. have done extensive research on CT-ultrasound registration [177, 179]. Rigid transformation between CT and processed 2D ultrasound images has

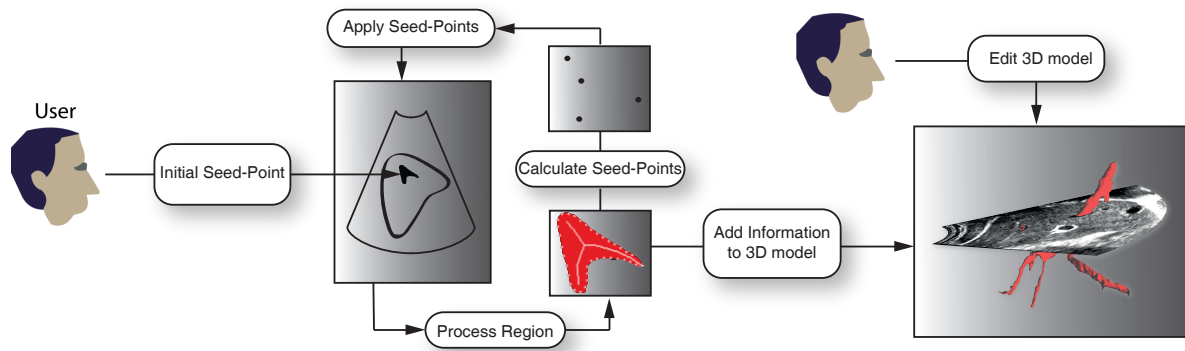


Figure A.2: The user first selects a single point, the system extracts features from the 2D image and maps the outline into a 3D environment. The process then repeats itself until the user is satisfied.

been presented by A Leroy et al.[95]. A. King et al. suggests an approach for rigid registration between 3D ultrasound and MRI data [81].

A.3 Ultrasound Painting Methodology

The painting process starts with the doctor locating a vessel cross-section in the ultrasound image and selecting the vessel with a button click. A distance manifold of the image is generated from the point selected by the doctor. The area of the vessel is then extracted with a region-growing algorithm applied on the distance manifold using the selected point as a seed-point. The technique used for vessel-region-extraction is described in Section A.3.1. In the next step we estimate new seed-points based on this vessel-region and they are used to extract new vessel-regions in the following frame. As the examination continues, the process repeats itself from frame to frame, as shown in Figure A.2. If vessels bifurcate, the system automatically starts tracking both branches without any required interaction from the user. If the vessel moves out of the field-of-view or it becomes too narrow from that particular transducer *viewpoint*, the system may *lose* tracking of the current vessel. Tracking of a vessel can easily be re-obtained by simply moving the ultrasound probe to a previously extracted area. The approach for tracking vessels and detecting vessel bifurcation is described in Section A.3.2.

A.3.1 Vessel Detection

Petersch et al. present a technique for soft segmentation of 3D ultrasound based on non-linear diffusion and distance functions [127]. Our approach use a 2D segmentation algorithm per frame based on the geometric segmentation technique presented by Petersch et al. To reduce the effect of noise and to remove speckle, we apply a Gaussian filter to the ultrasound image before features are extracted. The algorithm is based on geometric distance in an distance manifold of the image from a given seed-point, \mathbf{p} . We then build the distance manifold $\mathbf{d}(\mathbf{p}')$ out of the ultrasound image. The distance manifold regards distance measures equivalent to \mathbb{E}^3 where the difference between the central image intensity, $I(\mathbf{p})$, and the value of the neighbouring sample $I(\mathbf{p}')$ combined




Figure A.3: The distance function is given by the spatial difference and a weighted intensity function.

with a weight w , is regarded as the third coordinate. As shown in Figure A.3, using the 2D vector $(x, y) = \mathbf{p}' - \mathbf{p}$ and the difference in intensity, $\Delta I = I(\mathbf{p}') - I(\mathbf{p})$, we get the following equation:

$$d(\mathbf{p}') = \sqrt{x^2 + y^2 + (\Delta I \cdot w)^2} \quad (\text{A.1})$$

The default weight was set to 1 which turned out to provide the best result for most vessels, but for regions with low contrast the weight can be adjusted up to compensate. To extract the vessel-region we apply a region growing algorithm on the manifold $d(\mathbf{p}')$ with a user-defined cut-off using the current seed-point from the distance manifold. The extracted region is then used as the basis for seed-points in the following frame and the outline is plotted into the 3D model.




Figure A.4: The vessel-extraction technique output a region (a) and the skeleton is then estimated to provide new seed-points for continuous tracking (b).

A.3.2 Vessel Tracking

As the first seed-point is selected by the doctor, the system utilize his visual sense and medical experience to assume that it is a valid hypo-echoic region representing a blood-vessel. To relieve the user from having to select seed-points on every frame during the ultrasound examination, we developed a technique for automatically estimating the next position of the vessel based on the vessel-region extracted from the previous frame. To maintain tracking of the current vessel and to detect the new vessel-branches, the

tracking system makes an educated guess where to insert new seed-points within the next image.

After a vessel-region has been extracted with the technique described in Section A.3.1, we assume that the next frame will contain the same blood-vessel approximately in the same area. The first step is to estimate the center of the previously generated region. For our purpose we found the maximum value in a 2D distance transform of the vessel-region to be a good estimation. In the next frame we then use the estimated center of the vessel-region from the previous frame as the first seed-point for the vessel-extraction algorithm. For a more robust tracking and to detect branching, we calculate several seed-points in the outer region of the vessel-region. These seed-points in the outer areas are calculated from the 2D skeleton of the vessel-region. To extract the skeleton from a vessel cross-section, a morphological thinning process [58] is utilized. This results in a pixel-wide structure which represent the topology of the region, as shown in Figure A.4b. The end points of the skeleton are used as assisting seed-points in the next frame.

From the generated seed-point, several regions are extracted. Some of the calculated seed-points from the previous frame may not hit the vessel in the next image. To ensure that we are tracking a vessel we check if a seed-point is within the value-window of a hypo-echoic region, typically the lower 10% of the entire grey-scale range.

Hypo-echoic regions are not the only dark areas in the image. Certain seed-points will generate vessel-regions that bleed out into dark areas, such as shadows from bones. To detect these regions we look at the topology of the skeleton. If the skeleton contain too many different bifurcations, the region is considered non-valid. The first valid region generated are then merged with all valid overlapping regions. A non-overlapping valid region means that the blood-vessel is branching out. The new vessel-branch is then tracked in the same manner as the current vessel. After each vessel tracked in the current ultrasound image have been extracted, the outline and the center for each region is stored into the 3D model. In the next section we describe how information from the vessel-region is used to generate and handle a 3D model.

A.4 Vessel Operations




Figure A.5: Upper part of the left hepatic vein rendered as a point cloud with normals.

An ultrasound examination is usually a live process and the user examines the ultrasound image as the examination progresses. Since vessel-regions are updated in the 3D model every frame, we needed a fast method for storing and presenting the 3D model.




Figure A.6: During the examination, the user selects a branch, (a), and removes the corresponding points from the point cloud.

From each region we extract the outline sampled into discrete points, and the center point. The outline samples are put into a point cloud and the center point is used as the center line for the 3D point cloud as shown in Figure A.5. For each sampled point, we estimate the surface normal which is used later in triangulation and shading calculation. The point cloud is then triangulated while the points are added providing a surface rendering of the 3D model live during the examination. We have developed a technique based on Buchar et al.'s approach for local triangulation of point clouds [22]. The technique for live triangulation is described in Section A.4.2.

The blood vessels in the liver form a tree. In the next section, we present a technique for editing the point cloud using the branches of the center-line tree as a reference to points in the point cloud. Using a simple tree structure reduces interaction into a basic graph-interaction scheme where the user utilizes the ultrasound probe as a selection tool to prevent the need for sketch-based editing techniques.

A

A.4.1 Point Cloud Editing

Shadows from bones and air pockets appear as low intensity regions. Since the region detection approach is based on hypo-echoic regions with low intensity, the vessel-tracking system can be affected and generate undesirable artefacts. The topology of the vessel-tree is represented by the center-line-tree generated from the center of the vessel-region.

To edit and delete points from the point cloud we utilize the tree structure as a way to isolate points which belong to each vessel-branch. The center points from each extracted vessel region are put together into a tree, where the nodes of the tree act as parents for the surrounding points. The intersection between the ultrasound plane and the center-line branches are highlighted in the ultrasound image and the user marks an intersection point as the root of the center-line-tree. After the center-line-tree has been flipped according to the selected root, the user traverses to the branch which he wishes to remove. The selected sample points are highlighted and can then be removed. In Figure A.6, the user has removed the lower part of the branch due to artefacts generated from noise in the ultrasound image.

A.4.2 Point Cloud Triangulation

Comprehending the shape of distributed points in 3D space can be difficult if the points are rendered as single coloured dots on the screen. Many techniques exist for estimating point normals and point cloud triangulation. Due to live acquisition of data, we based our technique on a local triangulation technique developed by C. Buchar et al. [22]. To provide means of triangulating live while the point cloud size increase, we organize the point cloud into an octree structure. The triangulation within each cell can then be handled separately, considering only the neighbouring cells.

The local triangulation technique require point normals. The normal estimation is based on the eigenvector-vector of the covariance matrix to the *k-nearest-neighbours*. For our approach we chose *k* to be 16, which provided stable results. The normal is estimated by the cross-product of the eigenvectors corresponding to the two largest eigenvalues. The result of a cross-product between two vectors is dependent on the order of the vectors. Using the cross-product can result in normals pointing into the structure. Buchart et al. algorithm for orienting the normals correctly, assumes that on a smooth surface, each point connected by an edge in the mesh will have the approximately the same normal orientation. This require a time-consuming traversal of all the points to ensure correct orientation. Since we have the center-line already computed we can orient a point normal, \mathbf{n} , according to the normalized vector from the center-line to the point, \mathbf{v} using,

$$\mathbf{n} = \begin{cases} \mathbf{n} & \text{if } \mathbf{n} \cdot \mathbf{v} > 0 \\ -\mathbf{n} & \text{if } \mathbf{n} \cdot \mathbf{v} < 0 \end{cases} \quad (\text{A.2})$$

A.5 Results

For the results generated with the presented approach we have three different set-ups, one set-up using pre-acquired ultrasound data and two set-ups for live ultrasound acquisition. For the first set-up we have three datasets from a mechanically tilted curved ultrasound probe and a *GE Logiq E9* ultrasound machine. The images are stored in a stack and the transformation of each slice is calculated based on the number of slices per image and the angle of the tilt. Each dataset has been co-registered manually. The




Figure A.7: Triangulation of the point cloud into a smooth surface and provide a better shape-perception. Regions with sparse sampling are not triangulated correctly and holes appear in the surface.

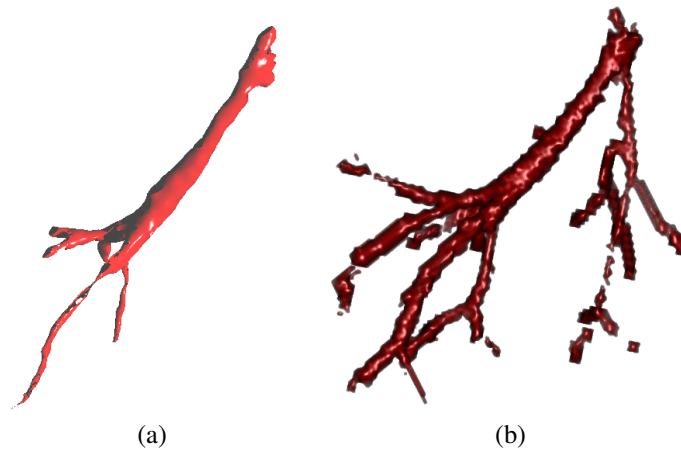


Figure A.8: In (a) we see a smooth surface rendering of a point cloud representing the middle hepatic vein. The point cloud is extracted from ultrasound data from a single sweep mechanically tilted probe. With only one sweep we are able to extract the several branches. In (b) we see the same vein extracted manually from a co-registered MRI dataset.

second set-up is a *Vingmed System Five* ultrasound machine with a curved probe and for the third set-up we have used the *GE Logiq E9* ultrasound scanner also with a curved ultrasound probe. For both live-acquisition set-ups we retrieve the ultrasound image using an *Epiphan VGA LR* frame grabber. Tracking of the ultrasound probe is acquired from an *Ascension Flock of Birds* magnetic tracker system. The prototype was developed in the framework *VolumeShop* created by S. Bruckner [19].

In Figure A.1 we have extracted the middle hepatic vein from the pre-acquired data. This data also have a co-registered pre-segmented MRI scan of the same patient. A comparison with the ultrasound extraction and manual segmentation of the middle hepatic vein can be seen in Figure A.8. In the Figure A.8a, some vessels are not extracted since they are too narrow from provided the ultrasound-probe *viewpoint*. This causes the system to miss the bifurcation and certain small vessels are left out. The vessels in Figure A.8b was extracted by a moderately skilled user in 4 hours, while the vessels in Figure A.8a was extracted in less than 10 seconds.

From the *System Five* set-up, we acquired a bifurcation of a liver vessel, seen in Figure A.9. From the transducer *viewpoint* only two vessels are extracted. Moving the ultrasound-probe to another *viewpoint*, tracking can be restarted and we can extract two more vessels shown in Figure A.10.



Figure A.9: A liver vessel bifurcation extracted with a single sweep using the *System Five* set-up.




Figure A.10: A second sweep from a different angle adding a to new liver vessels to the existing model.

Certain vessels are challenging to get a clear view as the ribs are casting shadows over the liver. Figure A.11 show the portal vein extracted using the *GE Logiq E9*, however certain vessel-branches have a limited length due to the shadow cast over the liver.




Figure A.11: An extraction of the portal vein. The semi-transparent discs are representing the cross-section between the ultrasound image and the center-line tree.

Our medical partners have earlier expressed that previous segmentation processes been too slow for efficient use. They said the results from the static set-up showed great potential for fast feature extraction. The doctors experienced however that respiration movement, heart movement and shadows in the image made tracking difficult for extracting entire vessels trees.

A.6 Conclusion

In this paper we have presented a technique for extracting features from tracked 2D ultrasound images. We focus our work on the vascular structures in the liver. However, the technique is applicable to other anatomical structures appearing as hypo-echoic structures in an ultrasound image, such as the gall bladder or hypo-echoic tumour.

With the presented approach vessels can be extracted *live* during an examination. As the user paints the extracted vessels, visualization of the extracted vessel is shown, providing a *what-you-see-is-what-you-get* approach for feature segmentation. The extraction technique is prone to disturbance from image noise and shadows from bone and air pockets. This can create unwanted artefacts. We have therefore provided a means of editing extracted data by applying an intuitive interaction scheme with the underlying center-line tree of the vessels. A challenge remains to provide a good *viewpoint* for the ultrasound probe to get an unobstructed view of liver vessels.

Since the extraction is performed live, the vessel-tracking is affected by respiratory movement. The lungs work as a piston pushing the liver down as the subject breaths in. The heart is also causing distortions in the vessel tracking. Veins close to the heart are moved back and forth, thus making extraction of vessels close to the heart challenging. We are currently investigating on motion compensating techniques to enable a more stable vessel-tracking system.

Acknowledgements

This work has been carried out within the IllustraSound research project (# 193170), which is funded by the VERDIKT program of the Norwegian Research Council with support of the MedViz network in Bergen, Norway (PK1760-5897- Project 11). We would like to thank anonymous reviewers from a previous version of this manuscript, for a thorough review containing many detailed suggestions for improvement. We would also like to thank our medical partners Trygve Hausken and Odd Helge Gilja at the Haukeland University Hospital for their input and aid with data acquisition. Finally, we would like to thank Helwig Hauser for his advice through the work process for this paper.

Paper B

Illustrative Membrane Clipping

Å. Birkeland¹, S. Bruckner², A. Brambilla¹ and I. Viola^{1,3}

¹University of Bergen, Norway

²Vienna University of Technology, Austria

³Christian Michelsen Research, Norway

Abstract

Clipping is a fast, common technique for resolving occlusions. It only requires simple interaction, is easily understandable, and thus has been very popular for volume exploration. However, a drawback of clipping is that the technique indiscriminately cuts through features. Illustrators, for example, consider the structures in the vicinity of the cut when visualizing complex spatial data and make sure that smaller structures near the clipping plane are kept in the image and not cut into fragments. In this paper we present a new technique, which combines the simple clipping interaction with automated selective feature preservation using an elastic membrane. In order to prevent cutting objects near the clipping plane, the deformable membrane uses underlying data properties to adjust itself to salient structures. To achieve this behaviour, we translate data attributes into a potential field which acts on the membrane, thus moving the problem of deformation into the soft-body dynamics domain. This allows us to exploit existing GPU-based physics libraries which achieve interactive frame rates. For manual adjustment, the user can insert additional potential fields, as well as pinning the membrane to interesting areas. We demonstrate that our method can act as a flexible and non-invasive replacement of traditional clipping planes.

B.1 Introduction

Clipping is an essential mechanism in visualization which resolves spatial occlusion in three- or more dimensions. This basic idea is used in visual depiction in many different variants, such as clipping into half-spaces, near clipping planes, section views, and cutaway-views which can all be considered variants of the concept.

In 3D data visualization, clipping becomes very powerful when making it adjustable by the user or adapting it to data properties. Clipping provides a simple mechanism for revealing otherwise invisible structures. Similarly, slicing uses the same interaction scheme, but shows only a single slice without any context outside of the intersecting plane.

The idea of removing occluding objects has been used in illustration craft for centuries. However, at closer inspection traditional illustrations rarely employ strict planar clipping but actually use hybrid representations where particular structures extend in front of and behind a hypothetical selected plane. Similarly, elongated structures do not appear or disappear from the rendering, as they would when cut by a plane. Instead, they are aligned with the cut plane.

One example, are dense cell illustrations where structures are clearly discernible and emerge out of the hypothetical clipping plane. It is then possible to follow elongated structures even if they spatially overlap. Figure B.1 shows an illustration of a slice through a *Mycoplasma mycoides* cell. As one can see in this illustration, the cell appears to be cut open by a plane revealing the inner structures, yet the structures at the position of the cutting plane are not sliced.

This is practically impossible to achieve with dense 3D visualizations. Objects will likely be intersected by the clipping plane so that only partial structures are depicted. The remaining portion itself may not have a clear meaning to the viewer and may hinder spatial comprehension. The research question arises of how to formalize the illustrative clipping concept in an explorative 3D data visualization scenario with limited semantic information at hand. One way of interpreting illustrative clipping is that the clipping geometry is not planar, but adapts to the structure of the object. The resulting clipping geometry will be a deformed 2D manifold instead of a plane. This manifold preserves the consistency in the structural depiction, provides subtle 3D shape cues and allows the following of elongated structures over the illustration. The resulting clipping membrane reveals otherwise occluded structures but at the same time communicates additional 3D information.

In our approach, the user positions the membrane in the scene and the scene elements act on the clipping geometry to cause the deformation. The membrane is utilized in order to reveal more of its structural properties, and to remove irrelevant information. To avoid cutting through an object when the focus is on its outer shape rather than its inner structures, the object repels the clipping membrane. Other data properties, such as object interfaces, can act attractively, so that the clipping membrane tends to be aligned with these interfaces. For this purpose, our clipping membrane concept is translated to a physics optimization problem. The scene elements generate a potential field where each data element's force is specified by a *potential field transfer function* (PFTF). The resulting potential field is an integral of (signed) forces which act on a deformable cloth object that constitutes the membrane. The membrane is not explicitly visible, but serves as a selection tool for direct volume rendering (DVR).

This idea has a wide range of potential applications. One scenario is the enhancement of traditional 2D slicing with selective depiction of 3D structures, but it also applies to adaptive clipping during mechanic movements of imaged structures in time-dependent data. To achieve interactive performance, the implementation of our clipping membrane concept makes use of NVIDIA's PhysX technology, which executes a physically-based cloth simulation on the GPU.




Figure B.1: Illustration using clipping to reveal the macro-molecular structures inside a *Mycoplasma mycoides* cell, by David S. Goodsell [2], the Scripps Research Institute.

B.2 Related Work

3D region extraction or clipping are common tools in medical workstations, for instance in ultrasound scanners with 3D acquisition capability [88]. Manually defining the clipping geometry has been very popular for creating expressive visualizations [43, 108, 181]. Konrad-Verse et al. showed how manual sketching can define a deformable cutting plane, aiding in virtual resection [87]. Deformed planes are also used in slice-based blood-vessel visualization, reforming a plane to the center of the vessel [75]. However, manually defining complex geometries can be cumbersome. Automatic definition of 3D clipping regions is typically based on semantic information [167].

Selective clipping adjustment has also been used in several visualization techniques. For preoperative planning, Beyer et al. demonstrated the use of clipping for combining the information of multiple imaging modalities [14]. Similarly, Burns et al. have developed a technique for combining 2D ultrasound and 3D CT in a fused visualization [24]. In their approach, the clipping geometry was defined by a 2D ultrasound plane. For additional context, certain features defined by segmentation were not clipped. Inspired by enhanced slice rendering [160] where contextual information of the surrounding areas is merged with the slice, we propose a method for using a deformed geometry to create expressive slice rendering without any required segmentation. Volume splitting [71]

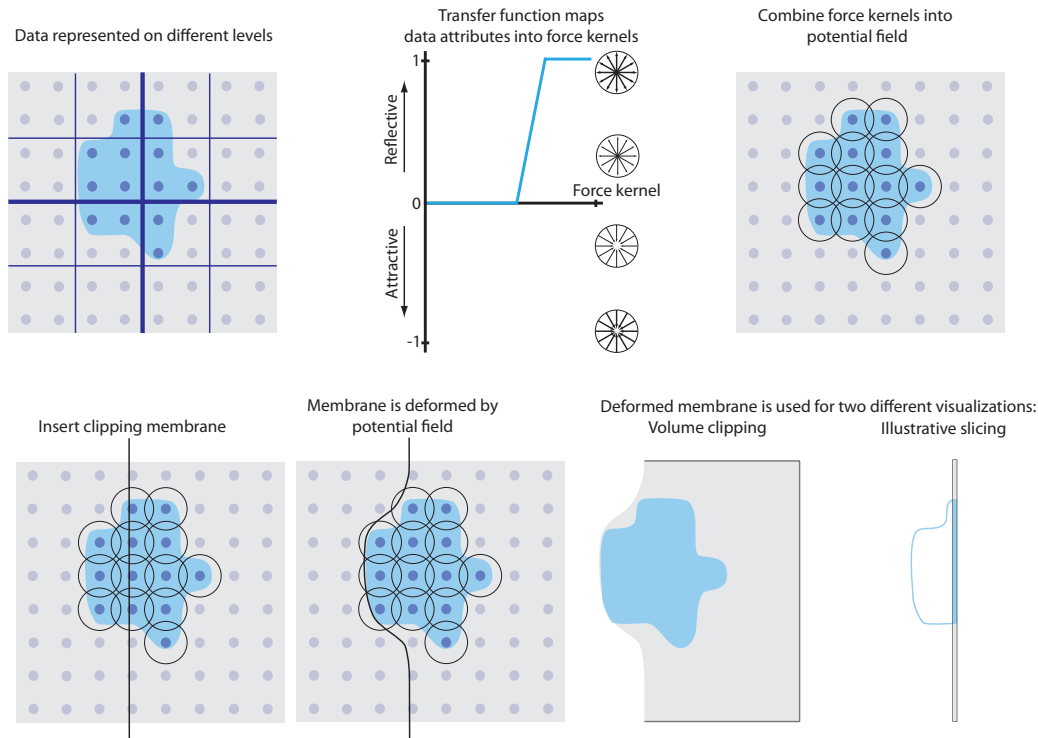


Figure B.2: Elementary building blocks of the illustrative clipping pipeline and the two use-case scenarios. Data samples from a specific level are mapped to forces using a *potential field transfer function*. The forces constitutes a potential field, which interact with a flexible membrane. The deformed membrane is then applied in volume visualization.

can also be seen as an advanced form of clipping, but rather than removing parts of the volume, it translates the clipped parts in order to ensure direct line of sight. An example of this are exploded views [20], where the different parts of the volume are repelled from each other, as well as repelled from the line of sight.

Schultz et al. used deformable clipping geometry to simulate a familiar resection method for nerve tracts in the brain, called Klingler resection [144]. The technique known as Virtual Klingler uses the data intensities to aid in the deformation of the clipping geometry. Schultz’s method was the first one to use forces to deform a clipping geometry. We see this method as a specific application of the general concept presented in this paper. While the Virtual Klingler method was highly focused on producing visualizations similar to Klingler resection, the technique had some limitations. The deformation of the plane was defined by creating a resistance field from MRI datasets. To achieve a deformation, the plane needed to be manually moved through the volume. We propose to use active forces, derived from the data, to deform an elastic geometry employed in volume visualization. In addition, we provide a controlling mechanism connecting data attributes to the strength of the forces and propose interaction schemes as well as visual mapping strategies.

Physically-based deformation of three-dimensional models is a very active research area, especially in the context of computer graphics. A detailed overview of the most recent approaches in this field was presented by Nealen et al. [109]. A notable amount of work has been dedicated to the physical simulation of cloths [69, 102]. In our work, we exploit the features of the NVIDIA PhysX library to simulate an elastic piece of




Figure B.3: Rendering of a medical CT dataset. Compared to the regular slice view (B.3a), adding local DVR provides interesting structures out of the plane (B.3b). Figure B.3c shows a tilted view of the same plane and the grid in Figure B.3d shows the deformed membrane.

cloth which acts as a clipping plane. Details about how the library handles this simulation can be found in Section B.3.1 or in the paper by Müller et al.[107].

B.3 Methodology

The core idea of this paper, is to translate the illustrative clipping concept into an interaction between soft-body dynamics actors, as depicted in Figure B.2. Firstly, the data attributes are translated into physical elements. This is done by creating an attractive or repelling force for each data element, called a *force kernel*. The combined forces from all the data elements construct a *potential field* depending on the data attributes. In addition, the user needs a controlling mechanism for linking data attributes to the strength of forces exerted by an element. We adapted the familiar transfer function concept, into a *potential field transfer function* (PFTF). Data attributes are then mapped by the PFTF to the strength of the force kernels.

Secondly, we need a geometry which can interact with the forces in the potential field. We therefore create a mesh which acts as an elastic membrane and then insert it into the potential field. By applying a soft-body physics simulation, each vertex in the membrane interacts with the nearby force kernels and conforms to the potential field created from the data attributes. As we aim to use the deformed geometry as a means to create expressive volume visualizations, we propose two different methods to exploit the deformable geometry in already existing rendering schemes. In volume clipping, we exchange the regular flat clipping plane with our deformed geometry. As a second utility example, we suggest the use of the deformed geometry for illustrative slice rendering, enhancing it with the local context defined by the deformed membrane.

B.3.1 Deformable Membrane

The clipping membrane is represented as a constrained elastic piece of cloth whose shape and position are automatically determined according to a set of physical properties. The physical simulation is based on *Position Based Dynamics* (PBD) [107], a




Figure B.4: Two different methods for repositioning the membrane: Passive repositioning creates a new membrane at the position of the rigid frame. Active repositioning moves the frame and pulls the membrane through the potential field. We provide a means for interacting locally with the potential field by creating a force kernel with a desired radius to encapsulate interesting structures. In addition, the user can pin the membrane at the current position to prevent the membrane from moving through interesting features.

physics-based approach to simulate soft bodies. Typical physical simulation algorithms are based on Newton’s second law of motion, where every particle is associated with an ordinary differential equation (ODE) which includes the forces acting on the particle, the particle’s velocity and acceleration, and possibly additional constraints. The system of ODEs is usually solved using an integration algorithm that takes into account all these variables. In the PBD approach, the positions of the particles are directly manipulated: in this way overshooting problems (caused by strong accelerations or high velocities) are avoided, and collisions and penetrations can be easily managed.

In PBD, the dynamical system is defined by a set of vertices and a set of constraints. Similarly to the common mass-spring models, every vertex $i \in \{1, \dots, N\}$ has a mass m_i , a position \mathbf{x}_i and a velocity \mathbf{v}_i . A constraint $j \in \{1, \dots, M\}$ instead is described by either an equality $C_j(x_{i_1}, \dots, x_{i_{n_j}}) = 0$ or an inequality $C_j(x_{i_1}, \dots, x_{i_{n_j}}) \geq 0$, where C_j is a real valued function defined over the positions of a subset of the particles. A stiffness parameter $k_j \in [0, 1]$ determines the *strength* of the constraint. Given a time step size Δt , the vertices are initially advected according to the explicit Euler’s integration scheme [62]:

$$\begin{cases} \tilde{\mathbf{v}}_i &= \mathbf{v}_i + \Delta t \frac{\mathbf{f}_{ext}(\mathbf{x}_i)}{m_i} \\ \tilde{\mathbf{x}}_i &= \mathbf{x}_i + \Delta t \tilde{\mathbf{v}}_i \end{cases} \quad (\text{B.1})$$

where $\mathbf{f}_{ext}(\mathbf{x}_i)$ is the force the potential field exerts on particle i . At this point collisions are detected and, for each of them, an additional constraint is generated. Finally, the algorithm iterates through all the constraints and repeatedly modifies the temporary particle positions $\tilde{\mathbf{x}}_i$ trying to satisfy all the equalities and inequalities (taking into account their stiffness k_j). For more detailed information about the general approach, we refer to the article of Müller et al. [107].

In the case of cloth simulation, two types of constraints are used: the *stretching* and *bending*. The stretching constraint is defined over every edge of the cloth mesh. It aims at preserving the original length of the edge, therefore the associated stiffness parameter can be used to control how rigid the cloth is. The bending constraint instead preserves the angle between each pair of adjacent triangles, so its stiffness parameter determines the cloth’s resistance to folding.

To maintain the simple clipping interaction the membrane must retain a certain shape. Without any constraints, the membrane would collapse. In our technique, the membrane is attached to a rigid frame. The rigid rectangular frame ensures that the membrane retains a topologically planar shape.

B.3.2 Potential field

The potential field is essentially a function $\mathbf{f}: \{1, \dots, N\} \rightarrow \mathbb{R}^3$ which associates a force vector to each of the N particles in the dynamical system. In particular, in the NVIDIA PhysX framework, the potential field is given by the sum of forces \mathbf{f}_k exerted by a set of kernels $k \in \{1, \dots, K\}$ on the particles:

$$\mathbf{f}(i) = \sum_{k=1}^K \mathbf{f}_k(i) \quad (\text{B.2})$$

The framework lets the user freely specify how the kernel force functions are computed. Specifically, in order to support datasets based on non-cartesian grids, we decided to treat each of the three components of \mathbf{f}_k separately. The force a kernel k , with position \mathbf{x}_k , exerts on a particle i , with position \mathbf{x}_i , is given by:

$$f_k^j(i) = c \left(1 - \frac{|\mathbf{x}_k - \mathbf{x}_i|}{a_j}\right) \quad (\text{B.3})$$

with $j \in \{0, 1, 2\}$. The parameter c is determined using the PFTF, while the parameter a_j defines the area of effect of the kernel and is proportional to the voxel size in the j^{th} direction.

At this point, building the potential field is straightforward: for every voxel the PFTF is evaluated, then a suitable kernel is generated and placed at the voxel's center. However, even with moderate size datasets, the resulting field would be so complex that the physical simulation would not run at interactive rates. To make the system more scalable, we adopt an octree structure to control kernels' placement: the desired octree depth is specified by the user, then, for every super-voxel at that depth, a kernel is generated. The kernel is placed at the center of the super-voxel and its strength is obtained averaging the PFTF values of all the included voxels. Eventually, the force exerted by the potential field on a certain particle is taken into account during the simulation as the f_{ext} term of Equation B.1. The equilibrium state of the membrane is then determined through the PBD algorithm.

While the potential field push the membrane away from the original slice position, the only force back towards the slice comes from stretching the nodes. This results in the membrane not following the the cavities of the potential field. As a means to force the membrane back towards its original position, we add a global force towards the slice, resulting in a tighter fit around the structures.

B.3.3 Interaction

The main interaction with the potential field is done using a widget for creating a transfer function. Adjusting the transfer function to data with very similar data values can be very challenging. To simplify interaction, we allow the user to sketch directly on the

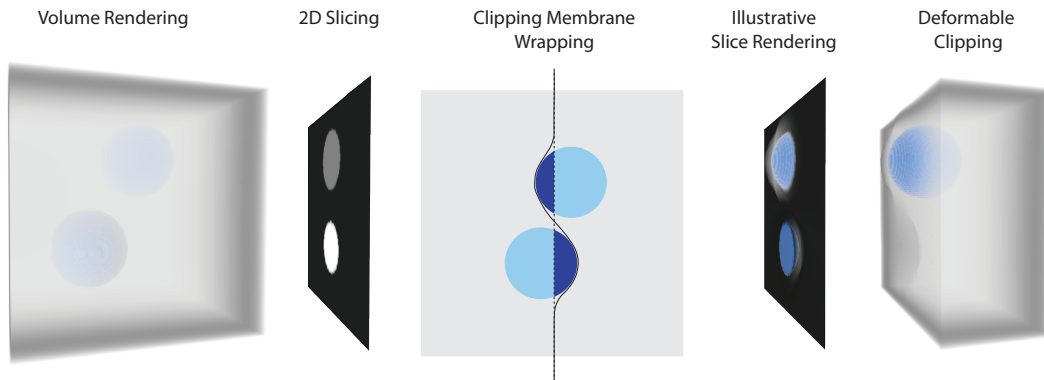


Figure B.5: Traditional explorative tools compared to illustrative clipping applied to a synthetic dataset.

original slice or clipping plane. From the sketch, we define a neighbourhood and create a histogram of the included samples which we use to adjust the potential field transfer function.

Positioning the membrane is essentially as simple as moving an ordinary clipping plane or slice through the volume. We propose two different schemes for positioning the membrane, passive insertion and active movement, as illustrated in Figure B.4. During passive insertion, the membrane is not affected by the force field as it moves. In essence, a membrane is created inside the potential field during every reposition, and it is then deformed by the force kernels it intersects. In active mode, the membrane is moved across the potential field, and it is deformed in real-time during movement. As the membrane moves across the volume, it *collides* with the potential field. While the frame is moved further across the volume, the membrane moves through the kernels from one equilibrium state to another.

For certain structures, very similar data values makes it difficult defining a proper PFTF. A more direct interaction with the membrane can be useful to keep interesting features in view. We have created two methods for manually interacting with the physical model. One method, is where the user indirectly deforms the membrane by inserting additional repulsive or attractive force kernels at the location of the original slice. Then, using a simple click-and-drag motion, the user can add a kernel with the desired radius into the existing potential field at the depth specified by the slice position.

In addition, the user might desire to keep certain structures in view while exploring the rest of the volume. In this case, the membrane can be *pinned* to structures by simply clicking on the membrane. An immovable anchor is then attached to the membrane, keeping the attached elements from moving with the rest of the membrane. This allows the user to manually *fine tune* the clipping. Figure B.4 shows examples of the two interaction schemes. To better illustrate the effect, we removed all other forces from scene.

B.3.4 Visual Mapping

We propose to apply the deformed clipping geometry in two different ways: membrane clipping in DVR and illustrative slicing with local DVR. Figure B.5 illustrates the different rendering methods. In both cases, we use a bounding geometry, for calculating




Figure B.6: A mouse imaged using ultramicroscopy [35]. Adjusting the transparency according to the distance to the original slice. The figures show the effect of adjusting the relative maximum distance for contributing to the image.

the entry and exit points for ray casting.

Regular clipping in ray casting can be achieved using a simple dot product test. With a deformed geometry, it becomes more challenging. Instead of testing which elements are in front or behind the plane, we start the ray casting at the surface of the cut. We then replace one face of the volume's bounding box with the clipping membrane. The new bounding geometry can be used to calculate new entry and exit points from the position of the membrane.

While clipping is a very common technique for volume exploration, domain experts, especially in radiology, tend to use slice rendering as their main visualization tool. Slice rendering provides highly detailed views, with no occlusion. The slice contains little or no context outside of the plane and a challenge rests on the user to comprehend the outer planar spatial extent of structures. Slab rendering can overcome this problem to some extent, by choosing a subset of the volume to be rendered. A slab is a subset of the volume defined by two parallel clipping planes. Similarly, with regular clipping, uninteresting objects may clutter the image and interesting features can be clipped.

To add 3D context to the slice of only those structures the user finds interesting, we propose using the elastic membrane for illustrative slicing, where DVR is only added in between the deformed membrane and the original slice. In this manner, we calculate the entry and exit points for ray casting from the proxy-geometry created by combining the slice and the membrane. The interesting structures are typically in the vicinity of the plane. We prevent clutter from distant structures by modulating the opacity of a certain sample, by the distance to the plane, as depicted in Figure B.6. Blending between the slice and the local DVR, it is important to be able to differentiate between structures in front or behind the plane. We have created a technique for user-selected blending and emphasis. Using two sliders, the user selects the opacity and the emphasis between structures behind and in front of the plane. One slider controls the opacity of the slice from completely opaque to fully transparent. Emphasis on structures behind and in front of the slice can be changed by a slider, ranging from -1 to 1. At -1 the DVR behind the slice is emphasized and the opacity of the DVR in front is set to zero. At 0, both the sides of the slice have equal emphasis and at 1 only the DVR in front are displayed.

An example rendering of a synthetic dataset can be seen in Figure B.5. On the left, regular volume rendering and slice rendering can be seen. Inserting the deformable plane in between the two spheres, it wraps around on either side. The illustrative slice




Figure B.7: Revealing the structure of a seismic horizon in seismic amplitude data. Flat clipping compared to illustrative clipping.

rendering shows the slice blended with the local DVR. Volumetric clipping, on the other hand, completely removes one sphere while the other remains visible.

With regular clipping, understanding the shape of the cut is fairly easy, as the cut is flat. The shape of a deformed surface might be more difficult to interpret. Shading the surface of the cut can aid in understanding the shape of the membrane and aids in understanding how the structures in view spatially correlate.

Typically, shading in ray casting uses the gradients as surface normals. This does not work at the surface of the cut itself as the gradient might be parallel to the tangent of the cut surface. To shade the surface of the cut, we use normals of the membrane in the first hit of the volume rendering, similar to Weiskopf et al. [182]. We render the normals of the deformed membrane into a separate texture which is then used during the ray casting.

B.4 Implementation

The prototype was implemented on a system with an Intel Core i7 3 GHz CPU and a GeForce GTX 580 graphics card, running Windows 7. Our goal was to create a technique which can handle reasonable sized data sets and run both the physics simulation and rendering at interactive frame rates. The physical simulation was more computationally demanding than the CPU could process. Therefore, we chose to implement our technique using the high performance GPU-based physics library *NVIDIA PhysX Engine* [1]. The library provided us with a cloth data structure which we could use as the elastic membrane for our technique, as well as data types for defining force kernels and collision elements in a physics environment constituting our force field.

B.5 Results

Illustrators have the option of selecting which structures they wish to keep in the illustration. In visualization of 3D data without any semantic information, this selection becomes very difficult. Clipping allows the user to spatially select which elements should be included. The membrane can be used to selectively clip seismic datasets. Here, the horizons below the sea bed are of high interest to the geo-scientific domain. However, seismic data has a high presence of noise and is difficult to explore using standard DVR and clipping. Applying a clipping surface adapting to the data attributes reveals more of the horizon compared to a flat clipping plane, as shown in Figure B.7. Here we used only the PFTF to create force kernels from high intensity samples (red).




Figure B.8: The hippocampus in a mouse brain imaged using ultramicroscopy [35]. Moving the slice plane through, the membrane is held back by the nerve cells (red), preserving it in the final image.




Figure B.9: Certain structures are completely lost using basic slice viewing of 3D ultrasound (B.9a). Using the illustrative slicing technique, elongated structures such as the blood vessels are kept in the image (B.9b).

The membrane is then moved smoothly over the volume. Other types of datasets have structures completely surrounded by similar data values. In these situations, creating the membrane inside the data might provide a more suitable option. In Figure B.8, the slice wraps around the nerve cells in the hippocampus of a mouse brain. The elongated structures are then preserved in the final image.

Basic slice viewing is the most common technique for exploring large 3D data. With a combination of a slice rendering and membrane defined local DVR, we can get a better view of the structure close to the slice, such as the blood vessels and bone in Figure B.3. The additional information can provide a better contextual perception. These images were created by drawing a simple sketch on a slice of the main aorta, which then automatically generated the PFTF. 3D ultrasound data is difficult to interpret using DVR due to much noise and similar data values. Local DVR rendering combined with slice rendering, can be a powerful tool for selective display of data and can aid in spatial understanding of dense 3D ultrasound volumes. In Figure B.9, the elongated blood vessel structures are preserved as 3D structures in the image. Capturing the blood vessels required some time to define the correct PFTF, due to many data samples with very similar values.

B.6 Discussion

Our illustrative clipping system is composed of two main elements, i.e., the potential field and the elastic membrane. Their behaviour and their interactions are determined by the physics-based simulator, and the whole process depends on a fair number of parameters which have to be set manually. The optimal configuration is admittedly highly dependent on the purpose of the visualization scenario. We have explored the usage of the various parameters for both the potential field and the membrane, and here we give some guidelines to help the user to setup the system according to his/her needs.

One of the most relevant parameters is the scaling factor c , in Equation B.3, which affects the strength of the force fields. Its value is determined through the PFTF, which has to be set according to scalar values in the dataset under consideration. In general, weak force fields make the membrane behave like a traditional clipping plane, while strong forces better emphasize inner structures, but may cause large oscillations of the membrane's position. All the examples in the paper have been generated with values of c up to 1000 and the membrane has never shown significant instability.

The global structure of the potential field is mainly affected by the positions and the areas of influence of every kernel. The placement is controlled by the octree level parameter: low level values produce dense kernel distributions, which, in turn, make the detection of the structures of interest highly accurate. The hierarchical ordering of voxels allows handling of even large datasets at interactive frame rates; however, in this case, generating a membrane which *wraps* around small structures can be difficult, as details may be lost due to averaging. In practice, the octree level can be used to control the tradeoff between speed and accuracy, and it has to be set, taking into account the dataset's size and resolution, and the available computational power. Figure B.10 shows the effect of using two different levels.

The area of influence of the kernels depends on the parameter a_j . Since the membrane should not intersect structures of interest, this value has to be set so that adjacent force fields overlap. On the other hand, the larger a_j is, the less tightly the membrane will fit to the structures of interest. After extensive testing, we found out that the best results are achieved when adjacent kernels overlap for 50% of their volume, i.e. when $a_j = 2v_j$, with v_j given by the super-voxel size in the j^{th} direction.

Taking now into consideration the elastic membrane, the main concerns are the bending and stretching stiffness parameters, which determine the membrane's resistance to bending and stretching, respectively. In general, low values of these parameters make the clipping plane fit more tightly to the structures in the volume, while values close to 1 make the clipping plane smoother.

The final shape of the clipping plane is heavily affected by the structure of the mesh used to represent the membrane. It is important to set the mesh resolution high enough in order to capture the smallest structures of interest, but, as for the potential field resolution, too high values could compromise the interactivity of the approach. Moreover, increasing the resolution indirectly makes the membrane stiffer, therefore the kernel's strength c should be increased as well. We set an initial resolution of 50×50 , but we let users freely modify this value in order to match the minimum feature size of specific datasets.

Lastly, the PBD algorithm assumes that every particle of the membrane has a certain mass. This value does not need to resemble the actual mass of the simulated piece of

Data (res.)	Hierarchy	# of kernels	FPS w/physics	FPS display only	Figure
Torso (256 × 256 × 556)	3	2,3k	9	15	Figure B.3 and B.10b
Mouse (424 × 279 × 190)	2	15k	3	12	Figure B.6
Seismic (121 × 121 × 121)	2	16k	5	13	Figure B.7
Hippocampus (212 × 191 × 44)	1	16k	3	5	Figure B.8
3D Ultrasound (181 × 245 × 190)	3	5k	10	14	Figure B.9
Torso (256 × 256 × 556)	4	300	12	15	Figure B.10a

Table B.1: Performance results from the technique. The numbers are taken from the images which are included in this paper. FPS w/physics shows the FPS while the physics simulation is running, and FPS display only shows when it is paused. Resolution of the membrane for all images was set to 50×50 .

cloth, it is simply used to control the inertia of the particles. In the implementation of PBD provided by the NVIDIA PhysX framework, this parameter is optional, so we have not specified it. Its influence on the resulting visualization will be the subject of future investigation. The parameter settings for the examples in the paper are summarized in Table B.1. Firstly, in all the cases, the default mesh resolution lead to fairly detailed results with no serious impact over the frame rate, so, unless there are very specific requirements, the average user won't need to modify it. The torso example can be used to compare the effects of different octree levels: a higher value results in higher frame rates, with some minor differences in the rendered images (see the comparison in Figure B.10).

The final results were shown to domain experts in the geological, medical and molecular biology domain. Geo-science experts described a recent work case where several horizons containing high-amplitude anomalies were dipping. The only way to display these surfaces were either to take a time slice (horizontal slice that does not consider the structural dip), or interpret a horizon and then create a surface that is used as a basis for making horizon slices. The second approach would be accurate, but also extremely time-consuming. Our method might be very useful to get a quick overview of high-amplitude anomalies and their extents, with an attractive trade-off between processing time and accuracy.

Medical experts stated that our deformable clipping approach could be interesting to combine with haptics for detecting cirrhosis, where dense *knots* form in the tissue and are otherwise difficult to detect in ordinary scanning. Overall, their comments stated that deformable clipping is promising for quickly gaining an overview of structures hidden in dense data.

Additionally, results from the adaptive seismic dataset clipping were shown to a molecular biologist and artist who is the author of the illustration in Figure 1. Currently, he is using automated methods for stylized rendering of entire molecular structures, but when it comes to revealing internal structures he needs to turn to a hand-craft. After




Figure B.10: Using a higher traversal depth in B.10a for the data hierarchy results in a less close fit compared to B.10b.

we explained that the same approach can be utilized for adaptive clipping in molecular biology as well, he stated that our method could possibly automate the process of selective rendering, which is only possible nowadays by hand-drawing, and could aid in creating highly expressive images.

B.7 Summary and Conclusion

In this paper, we have presented an illustrative clipping technique for volume exploration using a deformable clipping geometry. By using a deformable geometry, we have demonstrated how we can utilize physically simulated soft bodies to create expressive visualizations of volumetric data. The data attributes are translated into a potential field through a simple transfer function, defining how the data interacts with an elastic membrane. In addition, we have shown how the elastic membrane can be utilized in volume rendering, by aiding in removing visual clutter and adding context in well established volume exploration tools. Using a hierarchical structure and the GPU based physics library, interactive frame rates are achieved for both solving the physics and visualizing the data.

Acknowledgements

This work has been carried out within the IllustraSound research project (# 193170), which is funded by the VERDIKT program of the Norwegian Research Council with support of the MedViz network in Bergen, Norway (PK1760-5897- Project 11). In addition, this work has been partially funded by the Austrian Science Fund (FWF) through the ViMaL project (grant no. P21695) and by the Petromaks program of the Norwegian Research Council through the Geoillustrator project (#200512). We would like to thank our medical partners, Odd Helge Gilja and Trygve Hausken, as well as geologist Sten-Andreas Grundvåg, for valuable input. Finally, we would like to thank Laura Nic Lochlainn for aiding in proof reading and improving the final manuscript, Statoil for providing the seismic dataset, and David S. Goodsell for the motivating molecular machines illustration.

Paper C

Ultrasound Visualization Pipeline

Åsmund Birkeland¹, Veronika Šoltészová¹, Dieter Hönigmann², Odd Helge Gilja^{1,3}, Svein Brekke^{1,4}, Timo Ropinski⁶, and Ivan Viola^{1,6}

¹ University of Bergen, Norway

² Technology Transfer, Wiener Neustadt, Austria

³ Haukeland University Hospital, Norway

⁴ International Research Institute of Stavanger, Norway

⁵ Linköping University, Sweden

⁶ Christian Michelsen Research, Norway

C.1 Introduction

Medical ultrasound has a strong impact on clinical decision making and its high significance in patient management is well established [117, 118]. Ultrasonography (US) has in comparison with CT, MRI, SPECT and PET scanning very favourable cost, great availability world-wide, high flexibility, and extraordinary patient friendliness. In addition to these factors, ultrasonography stands out as the imaging method with the highest temporal resolution and also often the best spatial resolution. Furthermore, ultrasonography is a clinical method that easily can be applied bedside, even using mobile, hand-carried scanners [54] and even pocket sized scanners [48], thus expanding the field of applications considerably. However, low signal-to-noise ratio, "shadowing" and the relative small scan sector make ultrasound images very difficult to interpret. Accordingly, improved visualization of the broad spectrum of ultrasound images has a great potential to further increase the impact of ultrasonography in medicine.

As advancement of technology is fertilising and stimulating medical development, there is a continuous need for research and new applications in visualization. Visualization methods have the capacity to transform complex data into graphic representations that enhance the perception and meaning of the data [53]. Ordinary ultrasound scanning produces real-time 2D slices of data, and these dynamic sequences pose in itself a challenge to visualization methods. One example is functional ultrasonography (f-US),

i.e. ultrasound imaging of (patho)physiology and/or organ function, in contrast to conventional imaging of anatomic structures. Using f-US, information on motility, biomechanics, flow, perfusion, organ filling and emptying can be obtained non-invasively [55, 128]. Moreover, the 2D images can be aligned to form 3D data sets. In such cases, 3D visualization provides added value in terms of more holistic understanding of the data. Typical examples are demonstration of complex anatomy and pathology, pre-operative surgical planning or virtual training of medical students. Furthermore, there are now matrix 3D probes on the market that allow real-time 3D acquisition. To benefit from the high temporal resolution, advanced graphics techniques are required in ultrasound visualization, preventing the visualization technique from being the performance *bottleneck*. This opens up new challenges to the visualization community to develop fast and efficient algorithms for rendering on-the-fly.

In addition, co-registration techniques enable use of multi-modal data sets. Fusion imaging, where ultrasound is combined with either CT, MRI, or PET images, allows for more precise navigation in ultrasound-guided interventions. This challenging new arena demands advanced visualization research to enlighten how different data types can be combined and presented in novel ways.

The diversity of the ultrasound imaging technology provides a great tool for medical diagnostics, but the nature of the data can make it challenging to process. Techniques which work well for other modalities are being adapted to suit the special characteristic of ultrasound. In this paper we present an overview of the pipeline for advanced visualization specific to ultrasound data. The paper is divided into the chosen taxonomy, in essence each step of the visualization pipeline; pre-processing, segmentation, registration, rendering and augmented reality.

C.2 Taxonomy

Techniques for ultrasound visualization can be categorized in a variety of ways, e.g. when they were developed, which types of data modalities were utilized, which anatomy the technique was focused on, etc. The development of new ultrasound technology leads to different visualization techniques. The step from 2D ultrasound images to 3D freehand ultrasound (2D ultrasound with position information) revealed new challenges as spatial information could be included to generate volumetric data. The recent development of 2D matrix probes provided again a new challenge of 3D + time (4D) data visualization. Karadayi et al. published a survey regarding 3D ultrasound [76]. This paper has a greater focus on data acquisition and volume handling, but also give a brief overview over visualization of 3D ultrasound data.

Another taxonomic scheme for ultrasound visualization is based on the different types of data the technique utilized. 3D Freehand and 4D ultrasound, pose very different challenges compared to 2D ultrasound or when handling multiple modalities. Blending B-mode ultrasound for tissue and Doppler ultrasound for blood flow can be challenging enough in 2D and even more 3D. An example image is shown in Figure C.1d. In addition to the ultrasound input, the combination of other medical imaging modalities, such as CT or MRI with ultrasound, provide more information, but also more challenges to the visualization researcher.

Different anatomic regions have different characteristics in ultrasound images, as




Figure C.1: Example ultrasound images from the cardiac ((a)), gastric ((b)), fetal ((c)) and Blood flow (d) domain.

can be seen in Figure C.1. For instance, in a liver scan one might look for tumors using a high-resolution abdominal 2D probe. For heart infarctions, the doctor might need to examine the strain in the heart muscle to detect defect muscle tissue. The wide spread in tissue and pathology difference lead to anatomically specific visualization techniques.

In this survey we categorized around 60 papers and from the different categories we generated a parallel-coordinate plot, show in Figure C.2. Looking at the graph, we see an increase in rendering techniques for 3D ultrasound in the last five years. Volume rendering is often considered to be a *solved* problem. However, our study shows that much work dealt with volumetric ultrasound data. Yet, 3D ultrasound rendering can still not be considered a solved problem. The high presence of noise, shadows from hyper-echoic areas and inconsistent data values provide a great challenge to make 3D ultrasound a more easy-to-use tool for examiners.

We also see an absence of augmented reality techniques for 3D ultrasound. Yet another trend is the neglecting of 2D ultrasound from the visualization community. 2D ultrasound is the most used modality by physicians and while presenting the signal data onto the screen is straight forward, understanding what you see is not so trivial. Increasing the readability of 2D ultrasound is mostly worked on in the commercial domain, aiming to give the company an edge over its rivals.

In Figure C.2 we see the categorized papers in a parallel coordinate plot where each axis corresponds to the different taxonomy classification. The third axis (the pipeline axis) is selected as the classification for this survey. Five categories were chosen based on what we recognize as the essential parts in the visualization pipeline for ultrasound data:

- **Pre-processing:** Processing ultrasound data prior to segmentation, registration or rendering.
- **Segmentation:** Extracting features from ultrasound data.
- **Registration:** Combining ultrasound data with other types of medical imaging modalities.
- **Rendering:** Presenting ultrasound data.
- **Augmented Reality:** Mixing ultrasound rendering with natural vision.

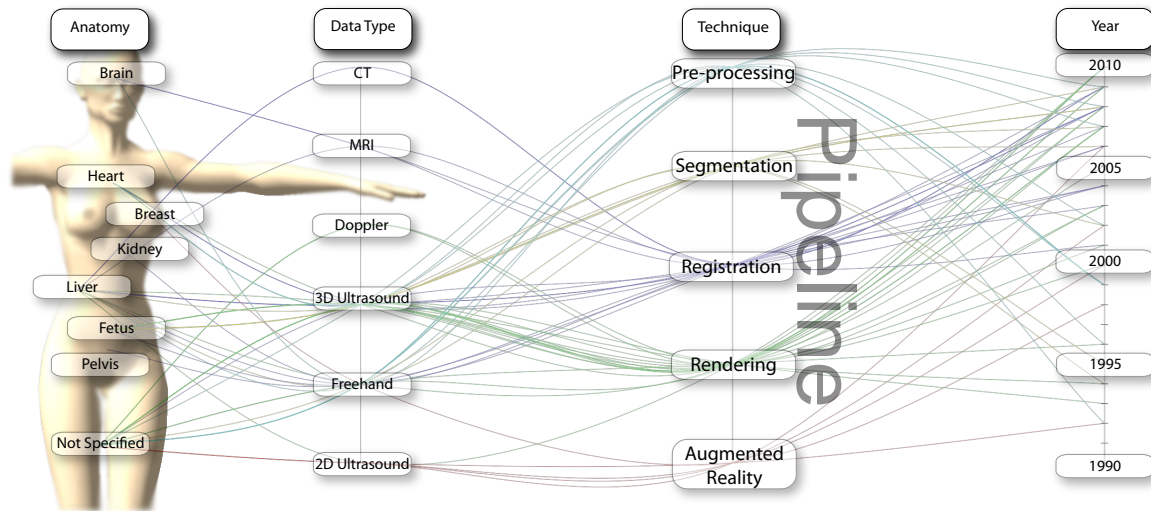


Figure C.2: The different classifications shown in a parallel-coordinate plot. The colors depict which technique a publication has given the most weight.

In the following sections we motivate the need for each of the major topics, focus on significant techniques and how they deal with the characteristics of ultrasound data.

C.3 Pre-processing

3D ultrasound is often employed in clinical diagnostic imaging. If a dedicated 3D probe is unavailable, 3D volumes can be acquired using freehand ultrasound systems; a 2D probe with an attached tracking which places and orients the 2D images in 3D space. Volume compounding consists of two levels: acquisition and reconstruction. Precise reconstruction requires calibration of the tracking system and correction of pressure-induced artifacts from the probe onto the skin.

Ultrasound allows for extracting more information, such as tissue strain. Strain is a tissue-deformation property and can be used to detect functional deficiencies, e.g., from myocardial infarction. Strain determination via tissue tracking is a complex task and can be done by using tissue Doppler[65]. Deprez et al. advanced in 3D strain estimation by providing a better out-of-plane motion estimation [31]. Visualization of strain has however stagnated compared to the development of technology and is mostly depicted by elementary color coding.

For freehand ultrasound systems, it is necessary to calibrate the position and orientation of the 2D image with respect to the tracking sensor. Wein and Khamene proposed to make two perpendicular sweeps through tissue containing well-visible structures [177]. They used an optimization strategy to maximize the similarity between two volumes reconstructed from each sweep.

To achieve the best possible quality of scans, the clinician press the probe against the body. However, the human factor causes a non-constant pressure and different deformations of underlying structures in the body. Prager et al. correlated images in the sequence and used a rigid translation in x and y directions followed by a non-rigid shift in depth z [131].

Ultrasound acquisition takes place in polar coordinates (ϕ, R) for 2D or (ϕ, ψ, R) for 3D. The angles ϕ and ψ correspond to the azimuth and elevation angles of the beam and R is the depth of the tissue boundary which has reflected the echo. In order to use of-the-shelf 3D volume rendering techniques, the grid must be scan-converted to a Cartesian lattice. This can be done as a preprocessing step or on-the-fly directly at the rendering stage.

This section is dedicated to selected methods for *volume reconstruction* from scan-converted freehand ultrasound and for *data enhancement* tailored for ultrasound volumes, which in the pipeline typically follow the reconstruction stage.

C.3.1 Reconstruction

Volume reconstruction from a set of 2D images needs to solve several important problems. Each images must be inserted precisely into the right spatial context, space-filling between individual images is also crucial and the high framerate of 2D ultrasound implies speed requirements.

A detailed categorization of reconstruction algorithms was done by Rohling et al. [136] and Solberg et al. [150]. We adopt the categorization by Solberg et al. into *voxel-*, *pixel-* and *function-based* methods and complete it by recent works.

Voxel-based methods, i.e., *backward compounding*, run through the voxel grid and assign each of them a value estimated by an interpolation method such as the Stradx system [132]. It allows for real-time visualization of freehand ultrasound including plane re-slicing based on nearest-neighbour interpolation and later also for direct volume rendering [131]; They blend images generated by re-slicing as described in their previous work. Gee et al. also used nearest neighbor interpolation for direct plane re-slicing [52]. The reconstructed plane is intended for direct viewing - implying only one re-sampling stage. Linear, bilinear and trilinear interpolation methods have been also used [13, 159]. Recent development by Wein et al. improve both quality and performance by applying a backward-warping paradigm implemented on dedicated graphics hardware [178].

Karamalis et al. used interpolation on the GPU for high-quality volume reconstruction [77]. They select an optimal orientation of reconstruction slices based on the orientation of the scans and reconstruct the volume by following this direction. Each sampling layer is reconstructed from scans which intersect this layer by interpolating intensity values between the intersections. The visualization pipeline includes two re-sampling steps: one during the reconstruction and one while volume rendering.

Pixel-based methods, i.e., *forward compounding*, traverse each pixel of all acquired 2D images and update the value of one or several voxels of the target grid. Gobbi and Peters used splatting as a high-quality interpolation method described a technique in real-time while the data is captured [56].

Function-based methods employ a specific function to interpolate between voxels. In most applications, the shape of the underlying data is not considered. Rohling et al. investigated the quality of interpolation using splines, which is a polynomial function [137]. They compared this technique with other standard methods and showed that it produces more accurate reconstructions.

Tetrahedron-based methods reconstruct 3D model built from tetrahedra using an iterative subdivision of an initial tetrahedron instead of a regular grid [139]. The subdi-

vision terminates, if all tetrahedra contain one data point. Each point is assigned a value which corresponds to the barycentric coordinates of the data point in this tetrahedron. This strategy is adaptive; the model adapts as new data is streamed in.

We listed selected algorithms in categories based on how they were implemented. If choosing a specific algorithm, one must choose between speed and quality. Solberg et al. compared the performance of some of the algorithms [150]. From all listed methods, the radial-based function reconstruction by Rohling et al. [136] delivers reconstructions of the best quality but it is also the most computationally costly. However, the increasingly powerful dedicated graphics hardware for computational acceleration solves this problem.

C.3.2 Data Enhancement

Ultrasound is a challenging modality for visualization due to its natural properties such as low dynamic range, noisiness and speckle [141]. Also, the geometric resolution varies with depth and the tissue boundaries can be several pixels wide depending on their orientation. Tissue boundaries can even disappear if they are parallel to the ultrasound beam. 2D images are preferred without filtering and enhancement. Speckle patterns refer to the texture of the tissue boundary which is a valuable information for clinicians. However, speckle in 3D brings no added value to the visualization and is considered as an artifact same as noise. Therefore, prior to the rendering stage, the 3D data is filtered to enhance its quality.

For a review on early speckle reduction techniques, refer to the survey of Forsberg et al. [41]. Belohlavek et al. [12] uses the *eight hull* algorithm with a geometric filter [29]. Recent techniques are based on region growing [27], adaptive filtering [147], compression techniques [61] and anisotropic diffusion filters [90].

Systems usually employ a blend of image-processing techniques to enhance the data. Sakas et al. listed techniques with a good trade-off between loss of information and quality [141]. They employed Gaussian filters for noise reduction, speckle-removal methods for contour smoothing and median filters for gap closing and noise reduction. Median filters remove small surface artifacts and preserve the sharpness of boundaries. There exist fast implementations where a histogram can be used to keep track of values [70]. Still, they require a more advanced memory management, making them less parallelizable than the evaluation of fast Gaussian filters. Lizzi and Feleppa described a technique to increase the axial resolution by processing the signal in the frequency domain. This resolution gain is especially valuable in ophthalmology when visualizing thin layers within the cornea [100].

C.4 Segmentation

Selecting interesting features to be visualized is important to be able to root out the occluding elements from large datasets.

For most modalities, segmentation can be performed by extracting regions with similar data values. For instance, because of the physical properties of x-rays, the data values in a CT scan are recorded into Hounsfield units which provide a good basis for binary thresholding techniques for certain tissue types. Early work indicated




Figure C.3: Automatic segmentation of the ovarian follicles [33].

that binary thresholding techniques are not very well suited for ultrasound data [152]. More sophisticated techniques are required for satisfactory segmentation. An extensive survey on ultrasound image segmentation was presented by Noble and Boukerroui [113] in 2006. In this section we have focused on significant publication from recent years.

To increase robustness of the ultrasound segmentation, the various approaches are usually tailored for specific anatomies. Carneiro et al. have developed an automatic technique for segmenting the brain of a fetus [25]. By first detecting the cerebellum, the system can *narrow down* the search for other features. On the other hand, segmentation is an extremely critical procedure which may obscure diagnostically relevant aspects of the anatomy under examination. Consequently, fully automatic segmentation techniques have not been implemented in clinical systems so far, with the exception of a method for follicle volumetry [33], as shown in Figure C.3.

A great challenge with ultrasound segmentation is that the data is dependent on many factors. For one, different positions and orientations of the probe, while looking at the same anatomical part, can provide very different images. Hyper-echoic regions cast shadows onto the tissue behind it according to the probe position. This alone, makes ultrasound segmentation data highly uncertain. Most segmentation techniques return a model with no indication of the uncertainty of the result. To compensate for the fuzzy nature of the ultrasound data, Petersch et al. developed a soft segmentation technique for 3D ultrasound data [127]. This technique calculates a probability map for 3D ultrasound data, which in turn can be used to create *soft* representations of the features extracted.

C.4.1 Clipping

Feature extraction can be computationally costly. In-vivo 3D ultrasound examination cannot always afford the extra time necessary to extract the interesting structures. Therefore clipping is commonly used tool in live visualization of 3D ultrasound. Interactively removing regions which are not interesting, the user gets a clear view of the features normally occluded. Sakas et al. developed a clipping tool in their ultrasound




Figure C.4: Using MagiCut to Clip the volume generating a clear view to the desired structure [88].

visualization system [141] which is nowadays a standard feature in commercial 3D ultrasound systems. The user can in-vivo segment the dataset using three approaches. Drawing on one of the three axially-aligned slices, selecting everything along the current axis and within the sketch. Another tool is based on sketching directly on the 3D rendered scene. Each voxel is projected onto the screen and removed if it lies within the marked area. The third clipping tool is based on the distance from a single mouse-click on the view-plane. A hemispherical wave front is propagated from the seed-point and stops when the voxels reach user-specified threshold. Figure C.4 show an example of clipping implemented in the GE Voluson machines [88].

C.5 Registration

Merging ultrasound with other modalities can be very beneficial. While ultrasound provides high resolution images at a high frame-rate, other modalities, such as MRI or CT can provide information complimentary to the ultrasound images. Data registration is the process of transforming different modalities into the same reference frame to achieve as much comprehensive information about the underlying structure as possible. While CT and MRI are typically pre-operative imaging techniques, ultrasound can easily be performed live during surgery. For instance, the radiation from CT is dangerous and the large electro magnets in an MRI scanner require that everything in the room is non-magnetic. Recently Curiel et al. built a non-magnetic ultrasound scanner for proper simultaneous intra-operative imaging [30]. There where some electric interference between the two modalities. Yet, the technique is promising, although availability will most likely be very low.

Nikas et al. published an evaluation of the application of co-registered 2D ultrasound and MRI for intra-operative navigation [112]. Ultrasound based navigation shows promising results due to live acquisition at high frame rates and easy portability. For prostate brachytherapy a combination of ultrasound and co-registered CT can be used, as shown by Fuller et al. [44]. Existing commercial products apply optical tracking for intra-operative navigation during neurosurgery [151]. Figure C.5 shows how ultrasound and CT can be blended together into a single reference frame [24].

Registration can be divided into two different types: Rigid and non-rigid. Rigid registration can be used to quickly obtain a registration between two modalities and is suitable for rigid anatomies such as the skull. A common approach to register two



Figure C.5: Registering ultrasound and CT, see Figure (a), enables advanced visualization techniques to provide a better structural overview [24], as shown in Figure (b).

images is to search for the transformation which minimizes a difference function, for instance sum-of-square-differences. Direct image based registration between ultrasound and CT or MRI can be difficult due to the different nature of the imaging techniques and usually some pre-processing, such as filtering, is required. For instance, an approach presented by Leroy et al. used a gradient-preserving speckle filter and then looked for the similarity in the gradients.

Penney et al. proposed a technique for registering MRI and ultrasound. The system calculates a probability map of each element being a part of a liver-vessel [123]. Later Penney et al. extended their technique for CT-ultrasound registration of the pelvis and femur [124]. The system was validated using cadavers, showing that the registration was accurate to a 1.6mm root-mean-square error on average. A similar technique for the cardiovascular domain was proposed later by Zhang et al.[186].

Combining segmentation with registration, King et al. presented a technique for registering pre-segmented models with ultrasound [81]. The technique predicts the probability that the ultrasound image was produced by the segmented anatomy.

In addition to a rigid transformation, affine registration includes non-uniform scaling which sometimes needs to be applied in order to get a more correct registration. Wein et al. developed an automatic affine-registration technique between CT and ultrasound[177]. To provide a better similarity of the ultrasound an CT, the system creates a simulated ultrasound image out of the CT scan based on the tracked probe position. The simulated ultrasound image is generated using a ray-traced approach to calculate the ultrasound wave reflection and attenuation in the tissue. To simulate tissue specific echogeneity, they apply an angle-independent polynomial function based on which tissue the region corresponds to.

External pressure or different laying positions of the patient when acquiring the images are influential factors. To account for local deformations while imaging soft tissue, a more complex registration is required. Papenberg et al. proposed two approaches for CT ultrasound registration [122] given a set of paired landmarks in both the CT and ultrasound data set. One approach use the landmarks as hard constraints and in the other, the landmarks are considered as soft constraints and are combined with intensity value information, in this case the normalized gradient field. The paper shows a non-rigid registration between the liver vascular structures. The latter technique was later evaluated by Lange et al. [93].

C.6 Rendering

Visual presentation of the data is the last stage of the pipeline before the user. The basic B-mode ultrasound images can be depicted on a screen in a straight-forward manner as varying pixel intensities according to the echo amplitude. Doppler information can be included as well with color-encoded blood-flow direction. Other data, such as tissue strain can also be included into 2D as overlays. Another example of overlays is the *CycleStack Plot* which superimposes the respiratory signal onto a selected feature of interest in the ultrasound image [94]. Doctors use this information to account for the respiration-caused motion of the tumor in order to minimize the damage done by certain tumor treatments.

Freehand ultrasound In Section C.3.1, we discussed how freehand ultrasound systems can be used to create large volumes by putting images into 3D spatial context. Garrett et al. presented a technique for correct visibility ordering of images using a binary positioning tree [45]. Visualization of large volumes leads to visual clutter. Therefore, Gee et al. extended existing re-slicing tools to create narrow-band volumes which contain less elements and are easier to present [51].

3D ultrasound is not as trivial to present due to its natural properties. In an early work, Nelson and Elvis discussed the effect of existing techniques for presenting 3D ultrasound data, such as surface fitting and volume rendering [110]. Later, seven ultrasound-dedicated volume projection techniques were evaluated by Steen and Olstad [152]. They included maximum intensity projection (MIP), average intensity projection (AIP) and gradient magnitude projection (GMP). The techniques were applied to 3D fetal data, where GMP was valued to give the best detail and robustness towards viewing parameters.

Data definition in polar coordinate system is another challenge for ultrasound volume rendering. Kuo et al. presented a technique for quick on-the-fly scan-conversion [91]. To reduce the costs of the functional evaluation of $\tan(\phi)$ and $\tan(\psi)$, the functional values were pre-calculated and stored in a texture as a look-up-table.

Surface Rendering is a common tool for many imaging modalities. In ultrasound, the low signal-to-noise ratio and parallel tissue boundary discontinuities make defining smooth surfaces difficult. Smoothing of a surface can be performed at the rendering stage. Fattal et al. presented an approach to render smooth surfaces from 3D ultrasound [39]. The surface is extracted based on the variational principle. Fuzzy surface rendering is done by a technique called oriented splatting. Oriented splatting creates triangles aligned with the gradient of the surface function, the triangle is then colored with a Gaussian function and rendered in a back-to-front order. Wang et al. proposed an improved surface rendering technique for 3D ultrasound data of fetuses [176]. To remove the noise and to preserve edges, a modified anisotropic diffusion is first applied to the dataset. To enhance low intensities which appear due to signal loss as the sound wave propagates through the tissue, a light absorption function based on the distance from a point is applied to the data. Finally, a texture-based surface rendering is used, where the texture is extracted from images of infants. The textures are warped and blended with the surface of the fetus face. To create smooth surfaces and remove unimportant noise in direct volume rendering, Lim et al. proposed a filtering technique in their GPU based ultrasound rendering framework [98]. This technique employs different sized filters to smooth out the noise.




Figure C.6: The parameter I_{thresh} determines which echo intensity values to render transparent. A user control with immediate feedback, indicating transparent regions in pink, is essential.

C.6.1 Transfer Function Design

For direct volume rendering, the *transfer functions* map ultrasound data, i.e., voxel echogenicity in B-mode imaging, and frequency information in Doppler imaging, onto colors and opacities. Usually, this mapping is based on look-up tables. In color Doppler imaging the commonly used red-to-blue color transfer function encodes direction and velocity of flow, whereas a variety of predefined color maps is in use for B-mode volume rendering. Custom color map editors are available, but hardly ever used. Overall, there is a well-established set of color-maps used in clinical practice.

Different from color transfer functions, where the selection largely depends on the preferences of the sonographer, the proper design of an appropriate *opacity transfer function* (OTF) is crucial: When designing OTFs, the goal is to assign a high opacity to voxels of structures of interest, while mapping all other samples to low opacities, thus avoiding any occlusion of the target structure. Whereas computed tomography allows classification of tissue based on voxel intensities, tissue classification-based transfer functions do not work in B-mode imaging due to the completely different data characteristics: Generally, a high signal intensity arises at a transition *between* tissues of different acoustic properties. Thus, at least in the case of soft tissue structures, we will measure high signal intensity at transitional areas and lower intensity signals within homogeneous tissue. This is the reason for applying monotonically increasing opacity transfer functions in DVR of ultrasound data: The aim is to opacify the tissue transitions in the hope of obtaining a visualization of an entire target structure.

The most commonly used OTF in volume rendering of B-mode data assigns voxels to one of three classes depending on their echogenicity, namely invisible, transparent, and opaque. The corresponding piecewise linear OTF is modified manually by means

of two parameters, namely a threshold intensity I_{thresh} and a transparency value α controlling the increase of opacity for intensities above I_{thresh} . The effect of modifying I_{thresh} is depicted visually on MPR images, see Figure C.6.

The parameters of the OTF affect the rendered image in a substantial way: The lower I_{thresh} , the lower the rendered image's brightness, due to an increasing number of hypoechoic voxels contributing to the image. Furthermore, the OTF affects depth contrast, i.e., the contrast arising from a spatial discontinuity in the target structure, and tissue contrast, i.e., contrast due to different echogenicity of adjacent tissue. See [67] for an evaluation of these effects on linear and parabolic OTFs. On the other hand, any modification of fundamental acquisition parameters, such as, overall gain, or depth gain compensation, and any change of the position of the transducer or the target structure, changes the echogenicity distribution and thus requires modifying the OTF for an optimal visualization. For a real time imaging modality, incessant modification is not feasible. Hence, in clinical practice sonographers use a default OTF providing reasonable visualization in the majority of cases, and hardly ever touch the OTF control panel.

Therefore, there is a need for automatic determination of an optimal OTF for every single acquisition. Due to the distinct characteristics and the real-time nature of ultrasound imaging, most conventional approaches for transfer function design have proven inadequate or require substantial modification in order to be applicable to ultrasound volume imaging. Among the most important advances in transfer function design for CT data is the work by Kindlmann et al. [80] and subsequent work by Kniss et al. [84], introducing the concept of histogram volumes for semi-automatic generation of OTFs for datasets where the regions of interest are boundaries between materials of relatively constant value. In [72], von Jan et al. adapt this approach to ultrasound data and apply it successfully to 3D freehand acquired volumes of hyperechoic structures.

Hönigmann et al. suggest an approach dedicated to the specific problem of rendering hyperechoic structures embedded in hypoechoic fluid [67]. By analyzing so called *tube cores* they yield an estimate for the position of the most prominent tissue transition, in rendering direction. Voxel intensities *prior to* and *at* the detected interface steer the extent of modification of an initial, parabolic OTF in a multiplicative way. In a subsequent publication the authors assess the temporal coherence of the tube core method and conclude that it is sufficiently efficient and robust for online-computation of OTFs for an entire sequence of acquired volumes, if smoothing in the temporal domain is employed [125].

Additional challenges arise when it comes to DVR of multiple datasets.

C.6.2 Multi-Modal Rendering

Multi-modal rendering is meant to bring two or more data sets of the same object, into a single image. Having two or more datasets in the same scene creates a challenge to keep the cluttering of less interesting regions to a minimum from the datasets. For ultrasound, 3D Doppler data can be acquired simultaneously with 3D B-mode data. Jones et al. discuss several approaches to explore and visualize 4D Doppler data [74]. Multi-planar rendering, showing several slices at once, surface fitting of the Doppler data based on YCbCr color scheme values to improve separation between Doppler data and B-mode data. An approach is presented to blend multi-planar slice rendering

into a DVR scene. The DVR is shown very transparent and the slices provide better detail along the perspective. A different way of combining B-mode with Doppler data was presented by Petersch and Hönigmann [126]. They propose a one level composite rendering approach allowing for blending flow and tissue information arbitrarily. Using silhouette rendering for the B-Mode and a mix of Phong shaded DVR and silhouette rendering on color Doppler.

A new technique for blending Doppler and B-mode was introduced by Yoo et al. [185]. Instead of blending two 2D rendered images (post fusion), or a blending the two volumes while rendering (composite fusion), they proposes a way to do both called progressive fusion (PGF). Post fusion has a problem with depth blending and composite fusion will get a too early ray termination. PGF compensates for this by using an if-clause to adjust the alpha-out value in the ray-caster to either composite the Doppler-signal or the B-mode-signal.

Burns et al. applied illustrative cut-aways combined with 3D freehand ultrasound and CT [24]. This provides a better spatial overview for the ultrasound images. To add more information onto the 2D ultrasound image, Viola et al. proposed an approach to enhance the ultrasound image by overlaying higher order semantics[169], in this case in the form of Couinaud segmentation. The segmentation is pre-defined in a CT dataset and visually verified using exploded views. To combine it with ultrasound images, CT dataset is co-registered with the ultrasound using rigid transformation according to user defined landmarks. The different segments are superimposed onto the ultrasound image enabling the user to directly see which segments are being imaged. To improve ultrasound video analysis, Angelelli et al. used a degree-of-interest (DOI) distribution superimposed on the image [6]. The video sequence was presented as a function of time (x-axis), where the y-axis was defined by the amount the current ultrasound image covered the DOI-function.

C.6.3 Shading and Illumination

Light is an indispensable part of scenes we see in real life. Also in computer graphics, light sources and light transport models have to be taken into account, when rendering realistic scenes. In volume graphics, the problem of illumination and light transport has been tackled by a handful of researchers as well.

We distinguish between local and global illumination models. Local illumination models use gradients of the volumetric function instead of surface normals to evaluate the diffuse and specular terms of the Phong illumination model [96]. While local illumination models already reveal structures, global illumination methods result in a more realistic appearance, which further supports spatial perception. While gradient-based local illumination methods are faster to evaluate, gradient computation is sensitive to noise and high frequencies, which are natural properties of ultrasound data.

Recent works show that global illumination models based on gradient-free methods are suitable for rendering ultrasound volumes [138, 171]. Ropinski et al. described a volumetric lighting model which simulates scattering and shadowing [138]. They use slice-based volume rendering from the view of the light source to calculate a light volume and raycasting to render the final image (see Figure C.7b). A perceptual evaluation of the generated images indicates that the proposed model yields stronger depth cues than gradient-based shading. Šoltészová et al. presented a single-pass method for simu-




Figure C.7: (a): Diastole of the aortic valve on a modern ultrasound workstation using color-coding based on depth. (b): Rendering of 3D ultrasound of human heart with shadowing from the work of Ropinski et al. [138] and (c) rendered using the technique presented in the work of Šoltészová et al. [171].

lation of light scattering in volumes [171]. Light transport is approximated using a tilted cone-shaped function which leaves elliptic footprints in the opacity buffer during slice-based volume rendering. They use a slice-based renderer with an additional opacity buffer. This buffer is incrementally blurred with an elliptical kernel, and the algorithm generates a high-quality soft-shadowing effect (see Figure C.7c). The light position and direction can be interactively modified. While these two techniques have been explicitly applied to 3D US data, the application of other volumetric illumination models potentially also improves the visual interpretation of 3D US data. Figure C.8 shows a comparison of six different shading techniques as applied to a 3D US scan of a human heart. While the first row of Figure C.8 shows examples for the already addressed shading techniques, the second row shows three alternative approaches. Figure C.8d incorporates scattering of light in volume data, as proposed by Kniss et al. [85]. Their slicing technique allows textured slices to be rendered from both light and viewing direction simultaneously. By sampling the incident light from multiple directions while updating the light's attenuation map, they account for scattering effects in slice-based volume rendering. Figure C.8e shows the application of the directional occlusion shading technique [143]. This technique constrain the light source position to coincide with the view point. Finally, Figure C.8f shows the application of a technique based on spherical harmonic lighting [99].

Advanced illumination techniques are being now implemented in the commercial ultrasound workstations. Some workstations use additional color coding based on depth. Deeper tissues are coloured with cold tones such as blue while close regions have red and orange tones. This effect has been firstly described by Einthoven [36] and is also referred to as chromostereopsis [5]. Figure C.7a shows a chromatic depth-encoding rendering of a 3D human heart in a modern ultrasound workstation. Recently, advanced rendering has incorporated in the domain of obstetrics. GE Healthcare has implemented a global illumination technique for photo-realistic rendering fetus [49]. Figure C.9 shows example images from the HD*live* system, implemented in GE Voluson E8 Expert workstations.




Figure C.8: Comparison of six volume shading models as applied to a 3D US scan of a human heart.

C.7 Ultrasound and Augmented Reality

Ultrasound is commonly viewed on a separate monitor. Therefore it is difficult to comprehend the spatial relationship between what you see on the monitor and where it is located in the patient's body. Augmented reality can aid the user by for instance super-imposing the ultrasound image onto the body where the ultrasound probe is positioned. Bajura et al. presented a system which linked 3D freehand ultrasound with a head-mounted display (HMD) [10]. The HMD contains a camera, tracker and two displays, one for each eye. The system can then project the tracked ultrasound image onto the tracked camera feed so the user can see where in the body the image actually is positioned.

Combining segmentation, surface rendering and augmented reality, Sato et al. aimed to aid surgeons during breast tumor removal for minimizing risk and maximizing breast preservation [142] by projecting a segmented tumor onto a video feed. The tumor is segmented using a minimal intensity projection based selection of the volume of interest. In the final stage, the tumor is surface rendered and superimposed on the video image.

Stetten et al. show how tomographic reflection can provide a superimposed image onto the body without any tracking systems [154]. The ultrasound probe carries a half-silvered mirror. The mirror reflects the ultrasound image which is shown on a flat panel monitor mounted on the probe. This technique was extended in the Sonic Flashlight [148]. The tomographic reflection showed to increase the localization perception




Figure C.9: Example images from the HDlive system in GE Voluson E8 Expert ultrasound workstation. Image courtesy GE Healthcare, Healthcare Systems, Ultrasound.

compared to conventional ultrasound [184].

Augmented reality shows a great potential benefit in medical ultrasound imaging. Yet, there is a lag from technology development to the actual integration into the every day usage. Sielhorst et al. published a detailed review for advanced medical displays in 2008 [149]. This paper discuss the potential benefit and the increasing use for augmented reality in medical imaging in general. Stating that improvements in both technologies is needed to be able to create a seamless integration into the physicians and surgeons work flow.

C.8 Summary and Discussion

In this paper we have categorized several of the most important works in what constitute the ultrasound visualization pipeline. The pipeline is defined as the five major categories in data processing and rendering. The five categories are pre-processing, segmentation, registration, rendering and augmented reality.

Medical ultrasound data is very different compared to other medical imaging modalities. Techniques for the individual steps in the visualization pipeline are tailored to suit the special nature of the data. For instance, techniques meant for in-vivo use have strong performance requirements to handle the high frame rate of ultrasound images. Segmentation and registration becomes very challenging, due to inconsistent data values for similar tissue. Still, ultrasound remains as one of the most used imaging modalities in medicine. Research in advanced ultrasound visualization techniques focus greatly on 3D ultrasound, but the trend in diagnostics is mostly 2D due to higher frame-rate, high resolution and a minimal requirement for interaction. The temporal and spatial resolution for ultrasound is approaching the physical limits of the speed of sound. It is very important to explore what the strengths and weaknesses the different modalities possess and combine the strengths into the natural work flow of medical personnel.

Acknowledgments

This work has been carried out within the IllustraSound research project (# 193170), which is funded by the VERDIKT program of the Norwegian Research Council with support of the MedViz network in Bergen, Norway (PK1760-5897-Project 11). We would also like to thank Helwig Hauser for invaluable help and fruitful discussions.

C

Paper D

Perceptually Uniform Motion Space

Åsmund Birkeland, Çağatay Turkay, and Ivan Viola

University of Bergen, Norway

Abstract

Flow data is often visualized by animation of particles inserted into the flow field. The velocity of a particle on the screen is typically linearly scaled by the velocities in the flow data. However, the human perception of velocity magnitude in animated particles is not necessarily strictly linear. In this paper we present a study on how different parameters affect relative motion perception. To find a correction model we have investigated the impact of four parameters which are assumed to affect the perceived speed. The four parameters are speed multiplier, difference in direction, different type of contrast and the global scale of the velocities. In addition, we investigate how multiple motion cues affect the speed estimation. We have performed three studies. Firstly, we did a study where the impact of each parameter was analysed. In the results, we noticed a strong trend in *scale* and *multiplier*, a weak in *direction*, and no trends in *contrast type*. Using the trends for the significant parameters, we designed a compensation model, which would adjust the particle speed to compensate for the effect of the three parameters. We then performed a second study to investigate the performance of the compensation model. From the results of the second study we detected a constant estimation error, for which we adjusted the compensation model accordingly for the last study.

D.1 Introduction

The use of motion can be seen in visualization techniques frequently. It has a various set of purposes, for instance depicting a data attribute, attract attention, or convey shape information. An advantage of using motion in visualization is that motion detection is a pre-attentive process in the human cognitive system [114]. Motion can then effectively guide the users' attention to interesting features in the data.




Figure D.1: Process pipeline for achieving perceptually uniform motion space. We started with a pilot study to determine suitable ranges in parameter space. From the pilot study, we designed an initial study to investigate the selected parameters. The output was an initial compensation model. In the evaluation stage, we performed two iterations to determine the efficiency of our model and adapt it to the new results.

Flow visualizations are aimed to provide insight into how a fluid, metaphorically or not, deforms under applied shear stress for a given situation. Unlike data representing physical objects, there are no *real* structures in a flow. Still, when referring to the topology of a flow, there exist abstract structures such as critical points and vortexes. While topology is an important aspect of flow analysis, the velocity magnitude is also important in many cases. One example can be seen in simulation of flow around an air plane wing. The lift from a wing is generated from the low pressure over the wing caused by the difference in velocity of the air moving above and below the wing [3]. In medicine, blood flow analysis is important in prevention, diagnosis, and follow-up monitoring of diseases. In this case, accurate flow measurement can be acquired with a range of techniques such as 4D MRI for 3D flow over time, and B-flow ultrasound. For making correct decisions, multiple features of the flow have to be analysed, such as pressure, vorticity and velocity.

Different visual cues can be utilized for different data attributes. Color coding is a typical method for quantitatively depicting data. For instance, the velocity magnitudes in a flow field can be visualized with color. However, this provides no information regarding the direction of the flow, and other techniques such as glyphs must be applied in addition. The usage of color can be very efficient for a single data attribute. Still, there are other aspects of the flow which are often important for the user. An issue arises when more attributes should be visualized simultaneously. If color is already occupied for velocity magnitude, other means must be utilized.

In engineering, analysing aerodynamics is often done by adding a flow of air around the object of interest. The air flow is as such not visible. To be able to actually see the flow, particles are added into the flow field. The motion of the particles conveys intuitively information regarding both the direction and the velocity magnitude of the flow. This technique has commonly been adopted into flow visualization [4, 161–163].

Investigations into how the human visual system processes motion can be approached from different directions. In the direction of neurophysics, the fundamental laws of nature are applied for investigating the physical connection between motion detection and the neural activity. The approach is more directed towards signal processing, where one tries to investigate how signals from the optic system are processed and transmitted into the visual cortex. This can be done either by applying physical models to simulate the neurological response in the brain, or procedures, where the actual signals in the brain are being monitored, for instance electroencephalography.

Psychophysics is a different direction. Here, investigation is performed by examining how physical stimuli can affect the perception in subjects. For motion detection, this is typically performed by providing a task for the user and based on a certain stimulus. Analysis is then based on examining correlations between the parameters in the stimulus and the response given by the user.

In this paper we have evaluated how the relative motion perception is affected by four distinct parameters: relative speed up factor, global scale of the velocities, chromatic and luminance contrast, and motion orientation. We have performed user studies, where we measured the subjects estimation of relative speed between two separate sets of moving particles. We have discovered significant trends in estimation error for two of the parameters, relative speed factor and global scale of the velocities. The main contribution in this paper is a compensation model for creating a perceptually uniform motion space when using animated particles in visualization. From a series of perception studies we have shown how the compensation model successfully compensates for the effect of selected parameters. From our studies, we have shown how adding multiple visual cues, have a small improvement in perceived speed. In addition, we provide evidence that the effect of chromatic vs. luminance contrast are not as prominent as presented in previous work.

D.2 Related Work

In order to understand how the human visual system is detecting and analysing motion, there are several aspects which come into play. In neuro-physiology, one tries to understand the signal processing used in our visual system to detect motion. For instance, a typical view of problems in vision, is that they can all be modelled as correspondence problems. Finding the correspondence of an object compared to the brain's representation of that same object is used in object recognition. In motion detection, the problem can be modelled as detecting correspondence over time. As a computational problem, motion detection can be seen as detecting changes in a given position. The *Reichardt detector* [133] is an implementation based on finding the correspondence over time.

In contrast to the Reichardt detector, an alternative model for motion detection involves finding change in luminance over time, known as the *temporal derivative*. A more detailed explanation can be found in related literature [42].

From a psychophysics perspective, we find much work in experimentation on the effect of different parameters when users evaluate motion. Experiments suggest that contrast change has an effect on the perception of the given speed [155, 157]. However, there is controversy regarding how this affects the perceived speed in general [18, 158].

Research indicates an impact of color and luminance to motion detection. A continuous change in luminance can create apparent motion of stationary objects [103]. Derrington et al. presented a study on the impact of color to the motion after-effect [32]. Backed up by the claim of color and motion being processed in parallel in the human visual system, evidence has indicated that applying only chromatic contrast compared to luminance contrast will cause a slow down in the perceived velocity [26, 180].

Most work in motion detection results in qualitative statements stating more the effect of a given stimulus. Little information is provided regarding how to compensate for any systematic distortions in the human visual system. In color theory, there ex-




Figure D.2: Screenshots from tasks with different parameter settings. [D.2a](#) shows a basic setup with direction angle set to 0 degrees and the speed multiplier set to 5. [D.2b](#) shows chroma only contrast. In [D.2c](#), the direction of the flow was set to 45 degrees down and to the left. [D.2d](#) shows the basic particles without a tail.

ists a compensation model which creates perceptually uniform color space called CIE 1976 (L^* , a^* , b^*) color space (CIELAB). In the context of motion, a similar model would aid visualization techniques based on moving elements in conveying the correct information according to the underlying data.

We can explain our goal for this paper, by drawing an analogy to color theory. As CIELAB is a color model where the perceptual difference for each step in the color space is uniform, we intend to create a compensation model for moving particles, which can compensate for systematic distortions in the visual system when estimating relative motion. This way, we can apply animated particles in visualization in manner where the perceived information correlates better with the underlying data.

D

D.3 Methodology

In order to create a compensation model for a given set of parameters, we need to find the impact that each parameter would have on the perceived relative motion. For estimating relative motion of animated particles, there are several parameters which can affect how a subject estimates motion. The density of particles, the size of the particles,

contrast level between particle and background, contrast type (luminance or chromatic), relative difference in speed between particles, direction of motion, global scale of the velocities, shape of the particles, all can affect velocity estimation in different ways.

For our compensation model, we first selected a sub-set of the possible parameters to investigate. To assess a proper range for the parameters and the performance of our test-design, we performed a pilot study with a small set of participants. After adjusting the study, we performed a larger study to find the trends for each selected parameter. Using the resulting trends we discovered in the initial study, we created an initial compensation model and started an iterative study process to test and refine the model. An illustration of the process is depicted in Figure D.1.

We started with a study as simple as possible with basic particles with an uniform distribution and size. For a basic set of parameters, we chose four parameters which are seen in frequently in visualization techniques using animated particles. Examples of tasks involving each parameter can be seen in Figure D.2.

Firstly, there is the range of screen-space velocities for the particles. When looking out the window of a fast moving car, it is difficult to clearly see stationary objects, like trees or road signs, close to the car. Based on this effect, we can deduce that there is a speed limit in the visual sector, for which humans can perceive objects clearly. We assumed that speed estimation would be affected by this effect as velocities would approach the limit of the visual system. To investigate this, we chose a parameter which would scale up the speed of the animated particles globally.

Due to the asymmetries in horizontal and vertical vision, we can draw the assumption that there would be a change in the perceived motion when viewing particles moving from left to right rather than moving from top to bottom. In addition, comparing motion when particles are moving in the same direction would be easier than comparing motion in the opposite direction. We chose to test for this effect by introducing a parameter which would set the direction of the particle motion. Direction ranged from 0 degrees to 359, where 0 degrees meant the particles were moving in the same direction.

It has been stated that change in contrast and color can affect the perceived motion [180]. When mixing color coding and motion in visualization one should then be careful about the resulting contrast between the moving particles and the background. Based on the experiments presented in previous work, we assumed there would be a trend in estimation error as the contrast type would change from luminance-only to chroma-only. To adjust the contrast type, generated colors by selecting to points in CIELAB color space with luminance only. To shift the contrast type we rotated the points around the center-point within CIELAB color space up to 90 degrees. 0 degrees would then result in two colors with luminance-only contrast, and 90 degrees would result in two colors with chroma-only contrast. The contrast level would remain the same for all degrees. Figure D.3 illustrates how the colors where selected.

As the difference in speed would increase, we assumed that the perceived difference would not scale linearly. The main parameter tested was speed multiplier between two sets of moving particles.

The final element which we investigated, was the effect of adding additional visual cues which would indicate the speed of the particles. To test if additional visual cues improved the perceived speed, we added comet tails to particles to one half of the tasks generated. The length of the comet tails were linearly scaled by the speed of the

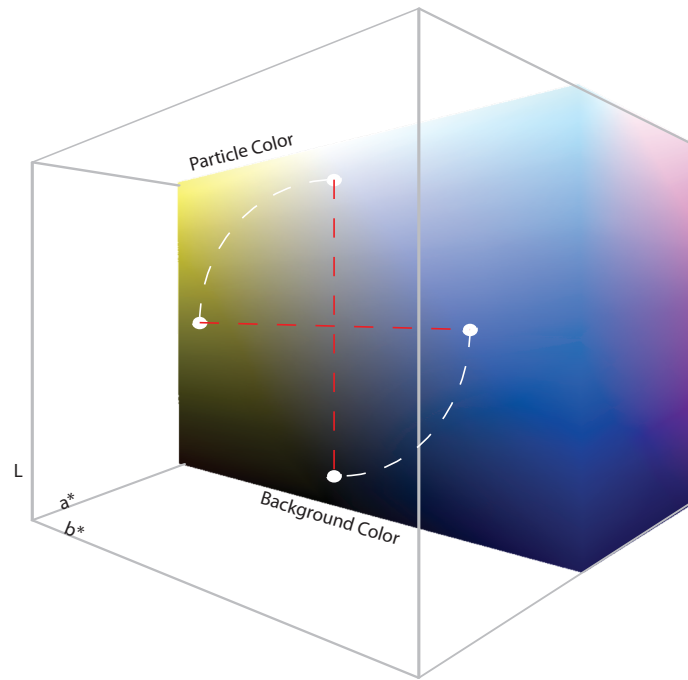


Figure D.3: The contrast-type parameter was set to be the angle of two points circling around a center point in CIELAB color space. 0 degrees meant achromatic contrast, and 90 degrees return isoluminant contrast, maintaining a constant contrast level.

particles and had a linear drop-off in width and opacity.

From the pilot study, we found that estimating speed for particles moving with a velocity over 41.5 degrees per second in the visual sector, would not return any meaningful results. In addition, if the velocity magnitude of one set of particles became more than six times the velocity magnitude of another, the error in estimated speed multiplier became too high to be useful. We set the base speed to 0.83 degrees per second, which constitutes one pixel per frame at 30 frames per second. This prevents any jitter in the movement of the particle. With a base level of 0.83 degrees per second, the range of the global scale parameter was set to 1 to 7.

D.4 Test Design

The aim for the study was to test for perception of relative speed between particle sets. In particle based visualizations, there is a multitude of variations in complexity among particle types, size, density, direction of motion, flow topology and more. To investigate the perception of relative motion, we made a deliberate decision to start with a very basic setup, which would test only one parameter at a time. This way, we can remove any unforeseen effects from other parameters.

In our tests we have two sets of particles displayed on the screen. One set is displayed at the top half of the screen (*test particles*) moving at a velocity \mathbf{v}_t , and one set is displayed at the bottom half (*reference particles*) moving at a velocity \mathbf{v}_r . The user was then asked to provide the multiplication factor, s , which satisfies the equation:

$$|\mathbf{v}_r| \cdot s = |\mathbf{v}_t| \quad (\text{D.1})$$

Change in density as well as a repeating pattern can distort the perception of speed. To compensate for this, the particles for each half were generated randomly under the constraint that the average spatial distribution remained uniform regardless of velocity and direction.

Each user was given a set of 100 tasks. We then selected subjects from various ages, gender, and professions. To avoid any learning curve, we excluded previous participants for the consecutive experiments. The subjects were not informed of the parameter intervals and were told to estimate velocities based on visual impression and not from explicitly timing particles' traversal over the screen (for example by counting seconds and comparing distances).

For every test subject we used a 24 inch screen with a 16:10 aspect ratio. The pixel size was 0.27mm and the canvas dimension was set to 768×768. The user's head was approximately 50 cm away from the screen.

If the eye is fixated on a continuous motion over some time, the eye will *adapt* to the motion. The reason for this is an effect called neural adaptation [164], where the neurons coding the particular motion reduce their responses. This can result in a distortion of our test. To avoid direction fatigue for the base speed, we alternate the direction of the flow for every other test. Furthermore, the subjects were asked to take a two-minute break after each 25 tasks have been completed.

The test design remained unchanged for the three rounds of testing.

D.5 Initial Study

In the first round we wanted to investigate the effect of the selected parameters separately. We generated three types of tasks, one which tested the effect of a global scaling of the velocities. The scaling is added on both sets of the particles. This changes Equation D.1 to:

$$|\mathbf{v}_r| \cdot s \cdot \sigma = |\mathbf{v}_l| \cdot \sigma \quad (\text{D.2})$$

The range of σ was 1 to 7 and the range of s was 1 to 6. For all the tasks involving global scale, the color-parameter remained constant at 0 degrees and direction remained strictly horizontal.

The second type was aimed for contrast type. We generated tests where colors was selected with the scheme explained in Figure D.3. The parameter ranges of 0 to 90 degrees in CIELAB color space. The range of s was 1 to 6. σ was constant at 1 and direction remained strictly horizontal.

Finally we tested and one for the direction of the flow in the range of 0 to 359 degrees direction. For each task we queried the subjects for speed multiplier between the two sets of particles in the range from 1 to 6. σ was set to 1 and color-parameter was set to 0 degrees.

For each parameter tested, we created tests with random configurations within the given parameter space. For instance for contrast-type, a configuration might be as follows: $\sigma = 1$, $s = 2.3$, color-parameter = 90 degrees and direction = 0 degrees. To prevent samples from being too clustered in the parameters space, we constrained a random function to keep the same number of tests within each interval. In total 2220 tasks were generated for the initial study. When the study was completed, outliers were removed by using a Mahalanobis score using the speed multiplier vs. signed estimation

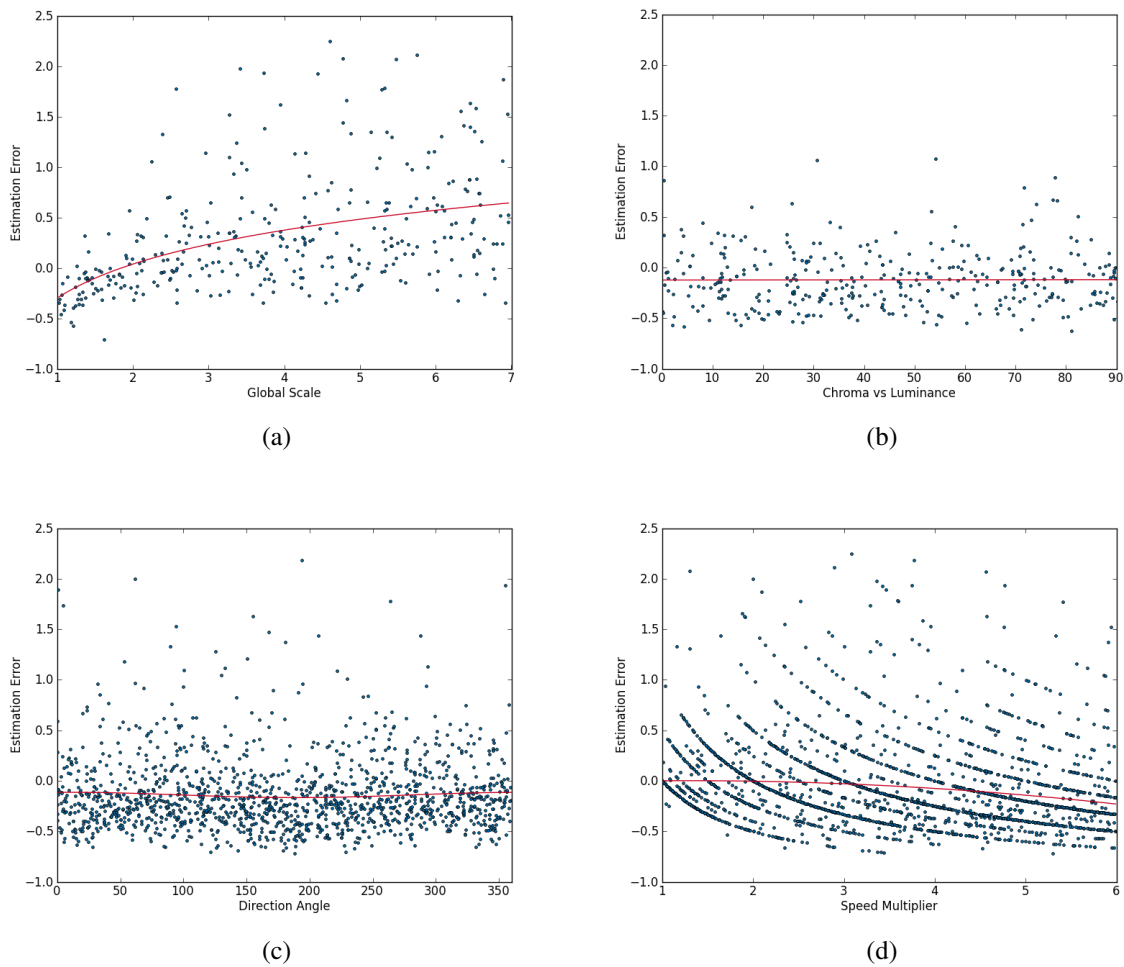


Figure D.4: Scatter plots of the samples generated from the first perceptual study round. The red curve shows the regression curves used for the compensation model. The curve for chroma-type (D.4b) shows no specific trend and thus this parameter was not included in the compensation model.

error. The threshold for outlier removal was 14.2, determined by visual inspection of the scores. This resulted in the removal of 22 outliers. The final sample size was then 2198.

From the results of the initial study we examined the relation between signed error and the parameters, i. e. if test-subjects overestimate or underestimate the *speed*. The trend in signed error is more relevant than unsigned. As there is a general trend in overestimation, we can compensate for this by slowing down the test particles. By this we will achieve a closer match between human perception of speed and the intended information communicated via the visualization. We investigated the estimation trend for each parameter. The significance of the trends were evaluated by binning samples in parameter intervals and performing a *two-sample unpaired t-test with unequal variances* between each bin, and the bin at the initial parameter setting. To provide a more robust significance evaluation, we performed the t-test for different parameter intervals. Table D.6 displays the t-scores for each parameter. The table shows the results from having 10 bins for each parameter. Interval shows the bin size in parameter space. N is

Parameter	tail	Mean Error/Std. Dev.
All	without	0.335 / 0.283
	with	0.299 / 0.271
Global scale	without	0.489 / 0.489
	with	0.375 / 0.432
Contrast type	without	0.253 / 0.182
	with	0.250 / 0.174
Direction	without	0.321 / 0.219
	with	0.296 / 0.240

Table D.1: This table shows the mean error and standard deviation of the estimation error for each task type, compared to the overall. Comparing the tasks where the additional perception cue was taken into account, there is a 3,6 percent decrease in mean error and standard deviation goes down 1.2 percent.

the number of samples in the bin. Mean shows the average in signed estimation error and T-score shows how the bin compares to the top row. DoF indicates the degrees of freedom for the t-test.

We assumed that there would be a trend in estimation error which correlated with the relative motion between the two sets of particles. From Table D.6, we can see there is a trend to continuously underestimate the speed value as the speed multiplier is increasing. Figure D.4d shows a scatter plot of the error in estimation compared to the speed multiplier parameter.

The results from the initial study also justify the assumption that there would be a change in estimation as the global scale of the velocities increase. Unlike the speed multiplier, subjects move from an underestimation to an overestimation. However, the slope for the regression curve for the relative motion factor remained unaffected by the global scale parameter. This lead us to conclude that there was an independence between the trend in error caused by increase in speed multiplier, and increase in global scale.

We had a hypothesis that comparing particles moving in different directions would be more difficult than comparing particles moving in parallel to each other. The difference in direction, however, seems to have a very small impact on the estimation error. From the results we could only find a significant correlation in unsigned error. There was, however, a weak correlation could be seen in the signed estimation error. This lead to the conclusion that direction should be accounted for in our compensation model.

Based on previous experiments in prior work [180], we expected to see a general trend in underestimation when using iso-luminant contrast. However, we were unable to reproduce the reported *slow-down* effect in our experiment. Figure D.4b shows the impact on the estimation error compared from the contrast-type.

In addition, Table D.1 shows the measured effect of the added visual cue. The effect is small and only prominent for particles with higher velocities.

The overall goal was to detect which parameters caused a trend in estimation error, which in turn, could be compensated for in particle-based visualization. From analysing the result, we found strong correlations between speed multiplier and scale

when comparing them with the estimation error. We also found a weak correlation between direction and estimation error. Therefore we chose to include the three parameters into our compensation model.

To create the compensation model we needed a function to describe the trends in the estimation error. For the global scale, we choose a logarithmic function,

$$E_{\sigma}(\sigma) = a \cdot \log(\sigma) + b \quad (\text{D.3})$$

where s is the global scale and a and b are constants defined in Table D.2. The logarithmic function gave a slightly worse fit compared to a second order polynomial, but since the first order derivative became negative at 6.5 in parameter space, we found the second order unsuitable to describe the trend, since the average estimation error was not decreasing (See Table D.6).

To find a function for the effect of direction, we fitted a cosine function,

$$E_d(\alpha) = a \cdot \cos(\alpha + b) + c \quad (\text{D.4})$$

where α is the angular difference between the two sets of particles and a , b and c are defined in Table D.2. Since the trend in error is cyclic, we found the cosine function to be better suited than a higher order polynomial.

Finally, we fitted a second order polynomial function for the speed multiplier parameter which we constrained to be zero when the particles were moving with the same speed,

$$E_s(s) = as^2 + bs - (a + b) \quad (\text{D.5})$$

where s is the speed multiplier, and a and b can be looked up in Table D.2.

Each function provides an estimated error for the given parameter. The final compensation function should provide a scaling function for the velocity to compensate for the total error from all parameters. In addition, the average error at the base level for each parameter would be contained in each function. This was solved by only including the change in error for the scale parameter and the direction parameter. The compensation function combines the error functions in the following manner:

$$\begin{aligned} C_{\sigma}(\sigma) &= 1 + E_{\sigma}(\sigma) + E_{\sigma}(1) \\ C_d(\alpha) &= 1 + E_d(\alpha) + E_d(0) \\ C_s(s) &= 1 + E_s(s) \end{aligned} \quad (\text{D.6})$$

$$F_c(\sigma, s, \alpha) = \frac{1}{C_{\sigma} \cdot C_d \cdot C_s}.$$

In the two following sections we present our evaluation and adjustment of the compensation function.

D

Function	a	b	c
$E_s(s)$	-0.010414	0.02680727	-
$E_d(\alpha)$	0.0261351	-0.16202086	-0.14056624
$E_r(r)$	0.48538867	0.29785412	-

Table D.2: Regression line parameters fitted to test results with/without modification




Figure D.5: A volume rendering for the 3D function defining our compensation model.

D.6 First Evaluation Study

After we built the initial version of the compensation model, we continued the process by evaluating it through a new round of perceptual study. Our aim at this stage was to assess the changes in the results due to the modifications by the compensation model. We then aimed to improve the model further, as a result of the investigation of the results from the new user study.

In this second round of the study, we made tests to evaluate the three parameters, namely, speed multiplier, direction of motion, and, velocity scale. Note that, we left out the contrast parameter at this stage. We randomly created 486 different combinations of these parameters to build the tasks for this round. For each of these parameter combinations, we created two separate types of tasks. In the first task, the final velocity was modified by the compensation model (i.e., experimental group), and, in the second task the final velocity was set without any modification (i.e., control group). This process led to $486 \times 2 = 972$ tasks in total. Similar to the previous round, we used particles with tails for half of the 486 parameter combinations. The tasks were then distributed randomly to 10 subjects who have not taken part in the first part of the test.

We started analysing the results by performing an outlier analysis of the results. Firstly, we removed the corresponding tasks from two specific users since their results exhibit conflicting trends when compared to both the 22 users in the first round and the other 8 users in the second round. Additionally, we removed 16 results after an inspection of their Mahalanobis scores. Here, we use two Mahalanobis scores computed using two sets of variables, 1) speed multiplier vs. estimation error 2) global scale vs. estimation error. The threshold to mark samples as outliers are 8.3 for the first score and 7.8 for the second. After the outlier removal, the remaining set consists of 756 task

Parameter	Modification	Slope	Intercept
Speed multiplier	with	-0.005	-0.241
	without	-0.05	0.120
Global scale	with	0.005	-0.283
	without	0.076	-0.354

Table D.3: Regression line parameters fitted to test results with/without modification

results.

To evaluate the impact of our compensation model, we treat the control and the experimental group separately (with/without modification). We observe the effect of each parameter on the estimation error separately. In order to do that, we plot the estimation error against the three different parameters for both of the experimental and control group. These plots could be seen in Figure D.7. Moreover, we fitted regression lines to each plot (Table D.3) and computed the average estimation error for the two groups of tasks.

For the tasks that are modified with our compensation model, all the three regression lines highlight a very significant result. We observe that our compensation model manages to flatten out the estimation error trends for all the parameters. Specifically, for relative motion the estimation error trend slope goes down from -0.05 to -0.005, and for velocity scale the slope goes down from 0.07 to 0.005. This amounts to an approximately 90-93 percent improvement for the trends in estimation error.

Although the correlation between these parameters and the error is removed, the results show that our modifications lead to an overall underestimation of the velocities. This is clearly seen when the average signed estimation error is observed. The average signed error changed from -0.06 to -0.26. The same observation is also supported by the placement of the regression lines in the second column of Figure D.7, i.e., the regression line is below the x-axis. This observation lead to a modification of the compensation model. We corrected the overall underestimation by inserting a constant ϵ into Equation D.6, where ϵ is equal to the average signed error, i.e., -0.26.

A second modification to our compensation model relates to the *direction* parameter. In order to check for the correlation between estimation error (both signed and unsigned) and the direction of motion parameter, we first group the task results into 10 bins that corresponds to 10 consecutive intervals of the motion direction parameter, i.e., each interval spans $360/10 = 36$. Secondly, we calculate the correlation between the signed/unsigned error and the direction of motion parameter over these 10 intervals. However, there is no significant correlation in any of these intervals. Therefore, we have left out the direction of motion parameter from our compensation model.

As a result of these two modifications, updated compensation model is formulated

with:

$$\begin{aligned}
 C_{\sigma}(\sigma) &= 1 + E_{\sigma}(\sigma) + E_{\sigma}(1) \\
 C_{\varepsilon} &= 1 + \varepsilon \\
 C_s(s) &= 1 + E_s(s) \\
 F_c(\sigma, s) &= \frac{1}{C_{\sigma} \cdot C_d \cdot C_{\varepsilon}}.
 \end{aligned}
 \tag{D.7}$$

These modifications to the compensation model called for a second evaluation round to assess the efficiency of the updated model.

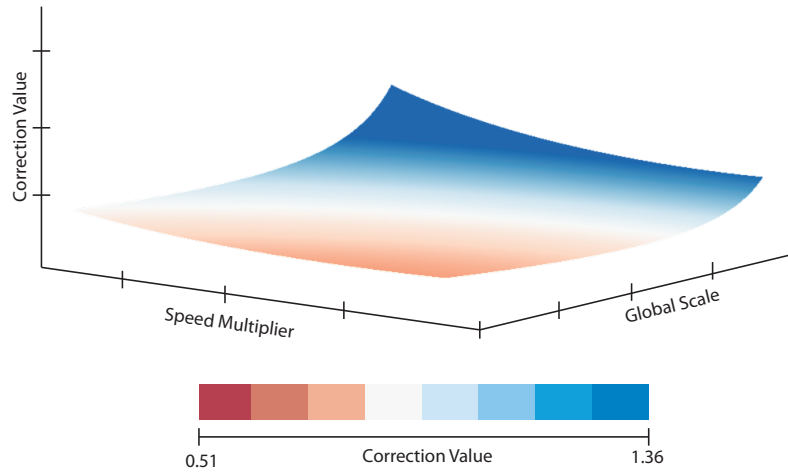


Figure D.6: The final compensation model resulted in a 2D function, depicted here as a height map.

D.7 Final Evaluation Study

From the previous round, we found the compensation function to significantly reduce the effect of the two most prominent parameters, global scale and relative motion. However, the new results indicate that there are no specific trends in error estimation related to the direction of flow. While we improved the impact of global scale and relative motion towards error in estimation, we also discovered a constant underestimation of 26 percent.

In the final round, we wanted to investigate whether the correction based on the constant found in the previous round would have the desired effect. In addition, we removed the compensation for direction as this parameter had no significant impact. The setup for the final study was almost identical to the previous study with 972 tasks, two sets with equal parameters leaving 486 with compensation from Equation D.7, and 486 without any compensation. The tasks were generated using the same constraints from the previous study. Again, we performed the study with ten *new* participants.

From the results we again removed outliers using the Mahalanobis score. We computed the distance for global scale compared to estimation error, and relative motion compared to the estimation error. The distance threshold was determined by visual inspection resulting in a threshold of six for the global scale parameter and ten for the

Parameter	Modification	Slope	Intercept
Speed multiplier	with	0.007	-0.153
	without	-0.051	-0.068
Global scale	with	0.015	-0.240
	without	0.064	-0.37

Table D.4: Regression line parameters fitted to test results with/without modification

relative motion parameter. In addition, samples from one user was deemed unusable, due to having trends conflicting with the other 39 subjects included from all the results. In total 81 samples were removed, leaving 892 samples for analysis.

The results from the final study are shown in Table D.4. The average estimation error for sample without compensation was -11.5 percent. For samples with compensation the average estimation was -17.8 percent. If we compare this to the previous study, we can see that the average estimation error for the samples without compensation, to be 5.5 percent lower than the corresponding set in the previous round. Samples with compensation included was 8.2 percent higher than the previous round. If we assume that the new set of subjects tend to estimate 5.5 percent lower compared to the subjects in the previous study, we can deduce a total improvement of 13.7 percent.

The parameter impact showed similar improvement in the last round, as we can see in Table D.4. In Figure D.8, we can see the regression curves for the tested parameters. The improvement in slope is most prominent for global scale, yet the effect of both parameters have successfully been reduced by more than 75 percent.

Based on the results from the last studie, we present the the final compensation model as follows:

$$\begin{aligned}
 E_{\sigma}(\sigma) &= a_{\sigma} \cdot \log(\sigma) + b_{\sigma} \\
 E_s(s) &= a_s s^2 + b_s - (a_s + b) \\
 C_{\sigma}(\sigma) &= 1 + E_{\sigma}(\sigma) + E_{\sigma}(1) \\
 C_{\varepsilon} &= 1 + \varepsilon \\
 C_s(s) &= 1 + E_s(s) \\
 F_c(\sigma, s) &= \frac{1}{C_{\sigma} \cdot C_d \cdot C_{\varepsilon}}.
 \end{aligned} \tag{D.8}$$

Where the constants can be found in Table D.5.

D

D.8 Discussion

Experiments have shown that there was only a systematic change in estimation error when comparing two of the four parameters selected for our study, the speed multiplier and the global scale of the velocities. Increasing the speed multiplier, i.e. when the difference in speed was increased, the perceived difference seemed to scale at a lower rate. Also, when increasing the global scale of the velocities, the estimation error changed from underestimation at low speed (base-speed lower than approximately 1.66 degrees

Constant	Value
a_s	-0.010414
b_s	0.02680727
a_σ	0.48538867
b_σ	0.29785412
ϵ	-0.26

Table D.5: Constants for Equation D.8.

per second), to overestimation at higher speeds (above 1.66 degrees per second). In this case there seems to be a *sweet spot* where we are most likely to achieve the best estimation, without any compensation. The trend in estimation error compared to speed multiplier remained approximately the same throughout the range of the global scale parameter and should be accounted.

Although, there was no systematic trends in the signed estimation error, when comparing to the direction of the flow, there was a change in the error magnitude. When the difference in direction approached 180 degrees, there was a small increase in error. However, due to the lack of any systematic change, we were unable to correct for this effect.

From the experiments in previous work, it has been reported a slow-down effect when gradually changing from luminance contrast to chromatic contrast. In our experiments, this effect was not reproduced. While we should be careful to rule out this effect completely, we might see that this slow down effect is caused not only by different types of contrast, but the gradual change. George Mather [103] reported that a change in luminance can create an apparent movement of stationary objects, and it could be that this effect is affecting the perceived motion of moving objects.

From the results using multiple visual cues for velocity encoding, we can only see a slight improvement from simple moving particles. The lack of improvement could indicate that the claim of using multiple visual cues improves the subjects' understanding of speed, is not as prominent as previously believed. However, for higher speeds there is an improvement of 11.4 percent, shown in Table D.1. In addition, our results only relates to perceived speed and not to direction, and the additional visual cue provides information not only about the speed but also about the history of motion as well as that of the particle's current direction. In addition, the usage of comet tails enables the encoding of velocities in still images which simple particles do not offer, as can be seen in Figure D.2.

To utilize the compensation model in a real world environment, there are certain aspects which should be considered. Since this scales the velocities of particles based on the underlying data, integrating the particle velocities in order to calculate its position becomes an issue. The reason for this is that the actual motion of a particle in the given flow would move at the velocity given by the flow, and arrival time and position would be distorted by the scaling from the compensation model. This can be resolved by scaling the reference frame. Still, the distortion might lead to a less comprehensible visualization. A more proper usage would be to integrate particles for a short time, in order to encode speed at certain position. For longer temporal integration, other visu-

alization techniques, such as stream lines, would be advised. It is also import to note that visual speed estimation of moving objects will not be 100 percent accurate. Using moving particles should be used more as an overview, and using a compensated model for depicting the velocities would create a better impression of what exists in the underlying data.

D.9 Summary and Conclusion

In this paper we have presented a new perceptually based compensation model for using animated particles in visualization. The compensation model is based on the results from a series of perception studies, investigating the perceived speed of moving particles. The main goal was to assess trends in estimation error based on selected parameters. We chose to test for four different parameters, namely global scaling of the particle velocity, the velocity direction, contrast type (isoluminant vs. achromatic), and speed multiplier. Four rounds of studies where performed: A pilot study, An initial study, which would was aimed for testing each parameter separately. A second study, which would test the compensation model. Finally, the last study, which should test for the improvements from the previous round.

From the results we have extracted the parameters which have an apparent significant effect on the error in estimation. This study evaluated each parameter separately. The results showed that significant trends were only visible in two parameters, global scaling and speed multiplier. A weak trend was found in the direction parameter. Using the trends in estimation error, we constructed a model which can compensate and reduce the effect of each parameter. The improvement was confirmed with a new study where the different parameters where combined. The results from the second study showed a large improvement in the impact factor from the selected parameters. However, the direction parameter was deemed insignificant. In addition, we also found a constant underestimation in speed estimation. Finally, we adjusted our compensation model according to the underestimation constant and performed an evaluation study of the corrected model. This again confirmed the reduction in impact from the significant parameters, as well as it improved the error in estimation compared to the previous study. This work was aimed at 2D flow, and can be used as a starting point for the perception of moving particles in 3D.

The final output from this work is a working compensation model for the perceived speed of moving particles. Based on the global scale of the velocities and their relative speed-up factor, we can create a perceptually uniform motion space for animated particles.

Direction					Global Velocity Scale				
Interval	N	Mean	T-score	DoF	Interval	N	Mean	T-score	DoF
0 - 36	150	-0.126	0.000	298	1.0 - 1.6	40	-0.206	0.000	78
36 - 72	141	-0.155	0.682	289	1.6 - 2.2	34	-0.007	3.903	72
72 - 108	142	-0.110	0.352	290	2.2 - 2.8	35	0.244	5.834	73
108 - 144	148	-0.094	0.722	296	2.8 - 3.4	31	0.304	6.435	69
144 - 180	145	-0.158	0.729	293	3.4 - 4.0	37	0.375	6.023	75
180 - 216	149	-0.220	2.054	297	4.0 - 4.6	33	0.313	6.369	71
216 - 252	146	-0.170	1.058	294	4.6 - 5.2	32	0.425	5.468	70
252 - 288	145	-0.107	0.424	293	5.2 - 5.8	36	0.594	7.804	74
288 - 324	128	-0.106	0.440	276	5.8 - 6.4	28	0.530	8.643	66
324 - 360	144	-0.155	0.745	292	6.4 - 7.0	35	0.617	8.644	73

Speed Multiplier					Chroma vs. Luminance				
Interval	N	Mean	T-score	DoF	Interval	N	Mean	T-score	DoF
1.0 - 1.5	224	0.032	0.000	446	0 - 9	35	-0.104	0.000	68
1.5 - 2.0	213	0.009	0.676	435	9 - 18	47	-0.127	0.344	80
2.0 - 2.5	217	0.011	0.571	439	18 - 27	48	-0.133	0.405	81
2.5 - 3.0	221	-0.016	1.279	443	27 - 36	40	-0.154	0.772	73
3.0 - 3.5	216	-0.015	1.172	438	36 - 45	39	-0.176	1.164	72
3.5 - 4.0	228	-0.099	3.464	450	45 - 54	39	-0.052	0.712	72
4.0 - 4.5	187	-0.135	5.158	409	54 - 63	42	-0.102	0.032	75
4.5 - 5.0	263	-0.112	4.165	485	63 - 72	38	-0.059	0.572	71
5.0 - 5.5	207	-0.153	5.400	429	72 - 81	50	-0.145	0.681	83
5.5 - 6.0	222	-0.200	7.144	444	81 - 90	41	-0.156	0.739	74

Table D.6: Results from the *Two-sample unpooled t-test with unequal variances*. A significant correlation can be seen for the speed multiplier and global velocity scale parameters. For the Direction parameter only a weak correlation can be seen, at around 180 degrees. Unlike previous work, we could not detect any significant trend from change in contrast type. Here the results are collected into 10 bins for each parameter. Interval column indicates the parameter range for each bin. N is the number of samples in each bin. Mean shows the average in signed estimation error and T-score shows how the bin compares to the top row. DoF indicates the degrees of freedom for the t-test.

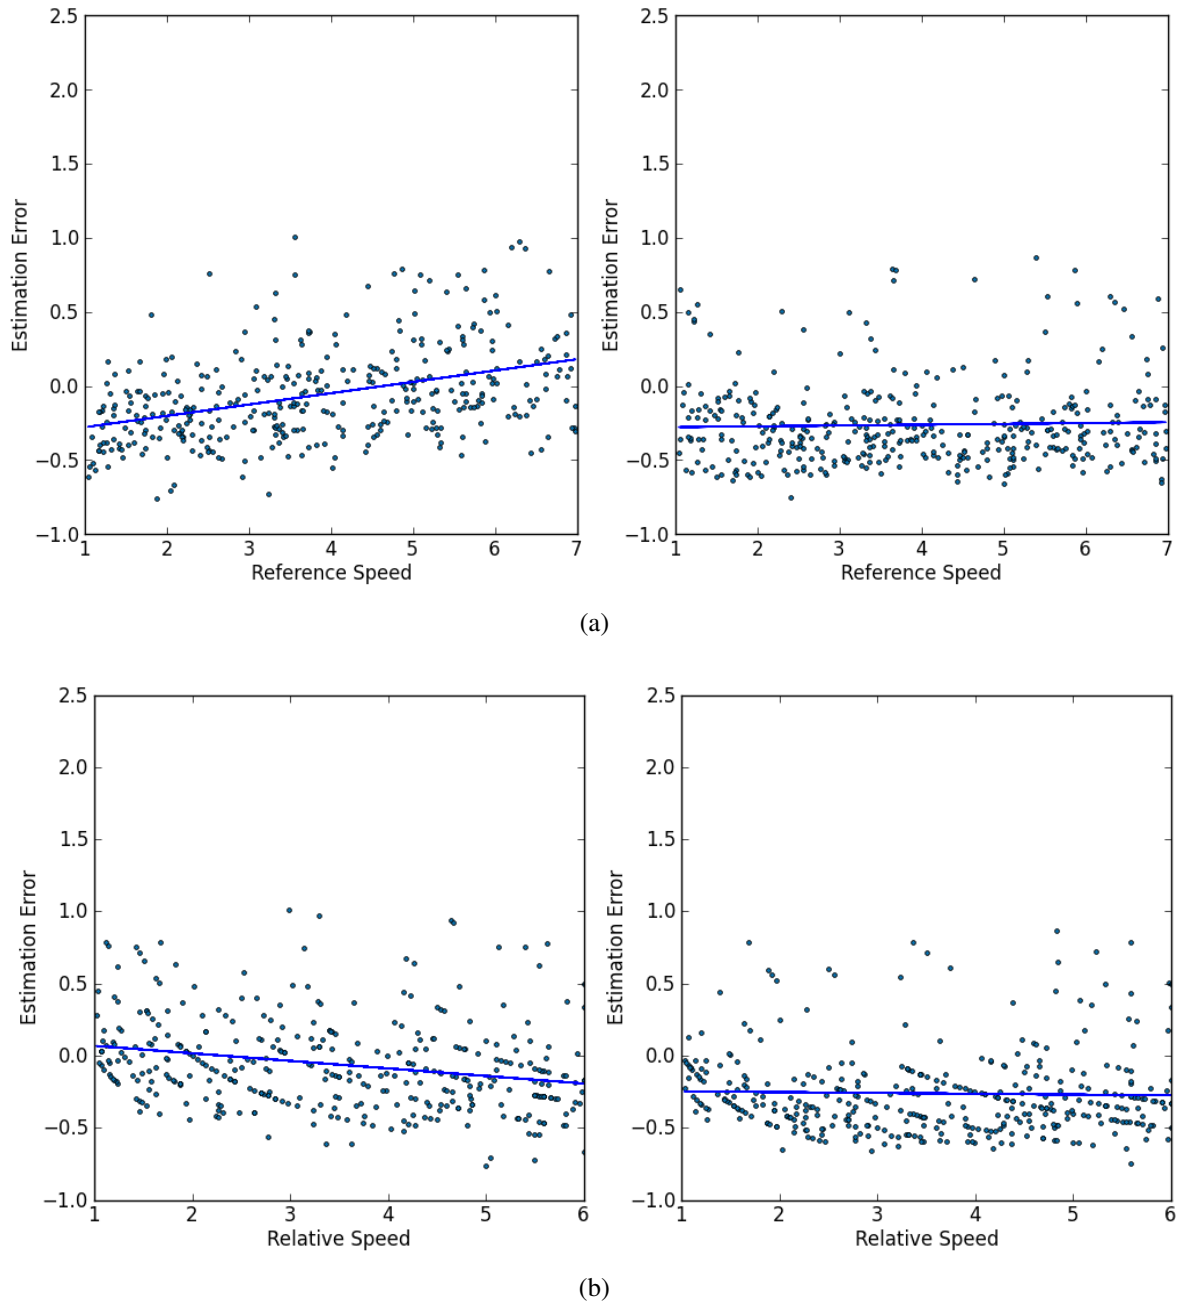

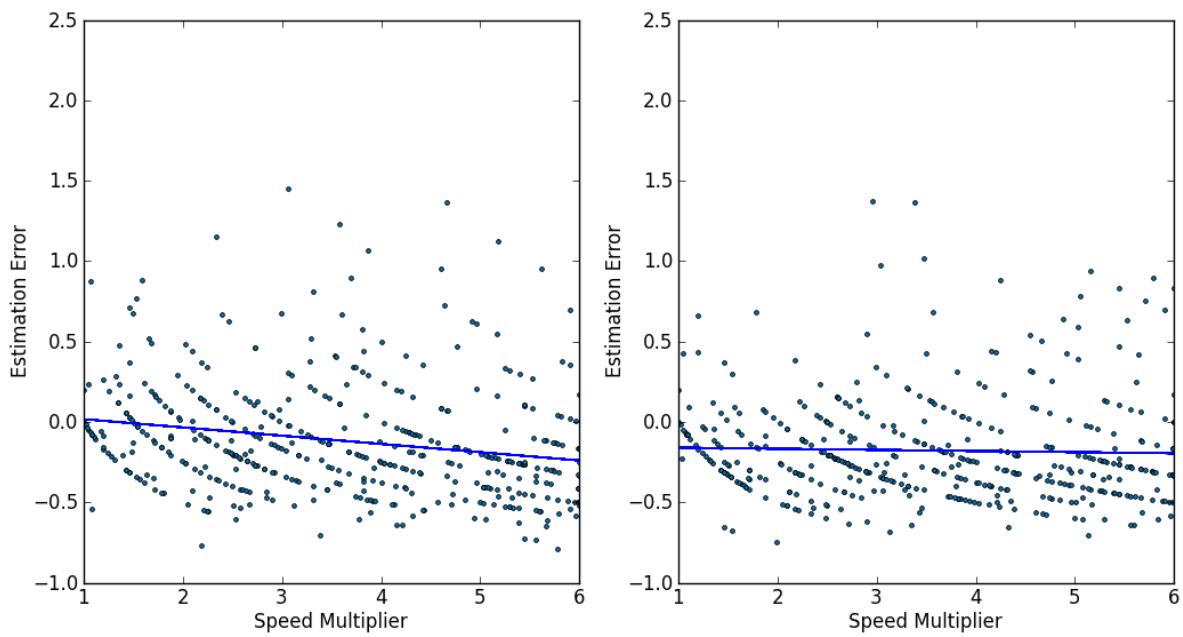


Figure D.7: The regression curves from samples generated in the first evaluation round. The different curves show with (right) and without (left) our compensation model.



(a)



(b)

Figure D.8: The results from the final evaluation round.



D

Paper E

Doppler-based 3D Blood Flow Imaging and Visualization

Åsmund Birkeland¹, Dag Magne Ulvang², Kim Nylund³, Trygve Hausken^{1,3}, Odd Helge Gilja^{1,3} and Ivan Viola⁴

¹ University of Bergen, Norway

² Christian Michelsen Research, Norway

³ Haukeland University Hospital, Norway

⁴ TU Wien, Austria

Abstract

Blood flow is a very important part of human physiology. In this paper, we present a new method for estimating and visualizing 3D blood flow on-the-fly based on Doppler ultrasound. We add semantic information about the geometry of the blood vessels in order to recreate the actual velocities of the blood. Assuming a laminar flow, the flow direction is related to the general direction of the vessel. Based on the center line of the vessel, we create a vector field representing the direction of the vessel at any given point. The actual flow velocity is then estimated from the Doppler ultrasound signal by back-projecting the velocity in the measured direction, onto the vessel direction. Additionally, we estimate the flux at user-selected cross-sections of the vessel by integrating the velocities over the area of the cross-section.

In order to visualize the flow and the flux, we propose a visualization design based on traced particles colored by the flux. The velocities are visualized by animating particles in the flow field. Further, we propose a novel particle velocity legend as a means for the user to estimate the numerical value of the current velocity. Finally, we perform an evaluation of the technique where the accuracy of the velocity estimation is measured using a 4D MRI dataset as a basis for the ground truth.

E.1 Introduction

Blood flow assessment using ultrasonography is important in prevention, diagnosis, and monitoring of diseases. There is a broad spectrum of clinical scenarios, where




Figure E.1: Blood flow in the inferior vena cava at peak velocity. Animated particles depict the blood flow. The velocity legend allows the user to read the particle velocities from the particle motion.

blood flow analysis is significant for patient outcome. One example is the assessment atherosclerosis for the prevention of thromboembolic events, such as myocardial infarction or cerebral stroke. Information about turbulence and velocity profiles in 3D, can be utilized to, diagnose and follow-up patients. Another medical scenario is intra-operative 3D blood flow assessment during a transplantation surgery, where blood flow characteristics could aid the surgeon in completing a successful procedure. In a post-operative scenario 3D blood flow turbulence could be assessed after, for example, coronary artery bypass graft (CABG) operation, and later follow-up and prevention of atherosclerosis in coronary arteries. Furthermore, blood flow analysis can be used in cancer assessment to monitor the response of treatment through changes over time in angiogenesis [129] and/or blood flow to the tumor. Furthermore, in pregnancy, changes in placental blood flow can be monitored in cases of intrauterine growth retardation. Moreover, measurement of flow and perfusion may play an important role also in targeted treatment using ultrasound and microbubbles [128].

The blood flow quantification, is carried out using various imaging modalities. The most advanced non-invasive imaging method for blood flow velocity measurements is the 4D magnetic-resonance blood-flow imaging (4D MRI). This modality measures true velocities of the blood in 3D, delivers the best imaging quality, and is in daily use for imaging heart and aortic arch blood flow. The 4D MRI is at the same time a very costly procedure, which cannot be performed bedside, and the flow analysis is carried out as postprocessing, disallowing in-situ flow visualization and analysis.

A widely used method for blood flow assessment is ultrasound-based Doppler imaging. This imaging modality can convey blood flow in real time. However only projections of velocity vectors, instead of true velocity vectors, are imaged and visualized. The true velocity is unknown, as only its projection onto the ultrasound beam can be measured by means of the Doppler phenomenon. Therefore, the examiner has to mentally reconstruct the blood flow velocity by considering the angle between the imaged blood vessels and the direction of the ultrasound beam. The more the beam is aligned with the vessel, the more the Doppler measurements represents the true blood flow

characteristics. This is, however, often not achievable due to the alignment of the vascular structures.

Recently another modality has been introduced to the clinical arena. It is known as the B-flow technique that visualizes the movement of the blood cells as particles. It provides a good qualitative description of the blood flow velocity, given that the beam direction is aligned with the blood flow. This modality however, does not provide any quantitative measures of velocity, it is only available as 2D, and since it requires to high frequency ultrasound it has fairly low penetration, limiting its use to vascular structures located directly beneath the skin.

There are other imaging modalities that can provide some information about the blood flow in situ. The X-ray based digital subtraction angiography (DSA) can reveal occlusions, and contrast-enhanced ultrasound (CEUS) can convey perfusion and angiogenesis. These techniques however, rather than measuring the flow phenomenon directly, provide information about the structure the flow is embedded in. Moreover, DSA and CEUS techniques involve the use of intravenous contrast agents and thus could be considered as being invasive imaging methods.

From the brief review of available imaging modalities it becomes apparent that there are currently no modalities which provide a satisfactory option for measuring blood flow quantities in-situ over larger areas of a blood vessel non-invasively. 3D Doppler ultrasound is the only modality which has the potential for real-time quantitative blood flow measurements non-invasively, but the estimates of flow information are prone to errors. The true flow direction is missing in 3D Doppler imaging. Assuming laminar flow, this can be estimated by tracing the direction of the vessel. The goal of this project was therefore to develop an imaging and visualization approach that combines the positive properties of 4D MRI and 3D Doppler imaging, which is in short: in-situ 3D blood flow velocity and flux (transport of blood at a given cross-section) quantification.

The main idea of our blood flow reconstruction methodology is based on combination of (a) assumed flow direction with (b) the implied flow orientation and velocity magnitude. This approach is carried-out as a two-stage process. In the first stage the flow direction is estimated. The fundamental component of vascular blood flow is laminar and our flow model also considers only laminar vascular flow. With this assumption we are able to extract the blood flow direction from the vessel geometry, or its respective skeleton. The second stage estimates the blood flow velocity magnitude and its orientation. This information is obtained from the the continuous flow (CF) Doppler measurements, and by backprojecting these onto the assumed flow direction.

Besides the blood flow reconstruction, the second contribution of this paper is a tailor-made visualization design that effectively depicts qualitative and quantitative characteristics of the blood flow dynamics. The third contribution is a preliminary evaluation of selected aspects of the novel imaging and visualization methodology. We perform an in-silico study to investigate the angular certainty ranges and informal evaluation of user satisfaction with the segmentation interaction and the applicability of flow labels for velocity quantification purposes.

E.2 Related Work

There are two major classes of prior work related to proposed concept on blood velocity imaging from the Doppler data and visualization of blood flow. One class deals with the challenge of reconstructing blood flow and the second class deals with the problem of how to effectively visualize measured or simulated blood flow phenomenon.

Imaging: Kripfgans et al. proposed a Doppler-based blood flow reconstruction method, known as *surface integration of velocity vectors* (SIVV) [89]. This work is estimating the flux quantities, without an a-priori knowledge about the vessel orientation. It is based on the Gauss' theorem, which relates the flux through an area with the divergence of the velocity vectors. Authors acquire the Doppler signal from a vascular structure, where an array-of-beams series constitutes one cross-section plane over the vessel. The cross-section consists of small surface areas, perpendicular to the beam. The Doppler values from this array-of-beams series are multiplied with the small surface area of the cross-section corresponding to each individual Doppler value. The total flux is the sum of these multiplications. Later this technique has been extended by an in-vivo study, taking into account the pulsative arterial flow [135]. Both in-vitro and in-vivo studies confirmed a 10% deviation from the gold standard, and thus Doppler-derived flow was considered as a reliable method for estimating flux in blood vessels.

Our approach is in contrast to the SIVV technique focusing on estimation of true blood velocities, instead of flux at the first place. From the reconstructed true velocity values we are able to compute the flux, while SIVV does not offer computing blood velocity without prior knowledge of the flow direction. The aim of our research was to provide the examiner qualitative visualization means conveying the blood flow coupled with quantitative flow velocity magnitudes and flux. The SIVV technique can provide only the quantitative flux information.

In ultrasound workstations, Doppler flow estimates can be improved by manually defining the direction of the flow. Hausken et al. demonstrated how this can be used for quantification of gastric emptying and duodenogastric reflux stroke volumes [64].

Our reconstruction pipeline is based on the vessel geometry information obtained through a rapid segmentation process. For this purpose we have adopted the fuzzy segmentation technique from Petersch et al. [127] for hypoechoic regions, which is based on a non-linear diffusion technique. This 3D segmentation technique has been adapted for live vessel extraction from a position-tracked 2D ultrasound [17]. While such an approach was reported to deliver reasonable results for patients during a breath-hold, it was not robust enough with respect to voluntary and involuntary patient-anatomy movements.

In ultrasound imaging there is a strong need for positioning information, which is robust with respect to various sources of tissue movements. In our pipeline we incorporate recent work where the positioning is based on an image-based registration of 4D ultrasound [121]. Having positioning information based on internal tissue information, the breathing will not affect the precision of positioning, unlike other positioning approaches facilitating external positioning hardware.

Visualization: The output of basic Doppler ultrasound imaging is a signed scalar field depicting the velocity of the blood towards or from the ultrasound probe combined with a B-mode image for structural reference. In a typical ultrasound workstation, 2D Doppler is displayed as a colored overlay over the B-mode with a corresponding leg-

end depicting which color correspond to which velocity. 3D Doppler is rendered using direct volume rendering (DVR), blending both the B-mode and the velocities. Still, 3D Doppler ultrasound data is very dense and it is challenging to combine both structural and flow information in a meaningful manner. However, Petersch and Hönigmann developed a new method for blending 3D B-mode data with 3D Doppler data [126]. They demonstrated a blending technique using a one level composite rendering approach. They combine silhouette rendering for the B-Mode and a mix of Phong shaded DVR and silhouette rendering on color Doppler. Jones et al. discuss several approaches to explore and visualize 4D Doppler data [74]. They apply a surface fitting technique on the Doppler data based on YCbCr color scheme values. This improves the separation between Doppler data and B-mode data. Using multi-planar slice rendering into a DVR scene, the details can be observed in the slice while a semi-transparent DVR provides the context.

While Doppler ultrasound visualization has the challenge of combining two scalar volumes, 4D magnetic resonance imaging (MRI) can provide accurate measurements of the *real* velocities [130]. This introduced new complexity to visualizing blood flow. Van Pelt et al. demonstrated different illustrative techniques for exploring and visualizing blood flow dynamics based on 4D MRI. They evaluated a set of proposed visualization techniques using integral lines and arrows seeded at cross-sections [161]. Later van Pelt et al. presented a technique, using particles as a means for depicting the flow, which was evaluated to enhance the perception of the blood flow recirculation [162]. To compensate for the high density of information van Pelt et al. presented a method for spatio-temporal hierarchical clustering [163], relieving clutter from the image by visualizing the clusters as path arrows.

Computational fluid dynamics (CFD) simulation of blood flow can also provide 3D or 4D flow data. Gasteiger et al. combined illustrative surface rendering and flow visualization for analysing cerebral aneurysms [47]. Recently, Gasteiger et al. presented the *flow lens* metaphor, a focus-and-context based approach for exploration of both 3D and 4D flow datasets from CFD and phase-contrast MRI datasets of aneurysms [46].

Above review of blood flow visualization methods was aimed to provide a short overview on the major trends in the medical-oriented flow visualization research field. A comprehensive survey of the state-of-the-art blood flow visualization techniques is provided in the referred report [165].

E.3 Methodology

There are several stages of reconstruction required before we can visualize 3D blood flow using color Doppler ultrasound. The general pipeline of our approach is depicted in Figure E.2. In the first stage, we start out with creating a segmentation mask of the blood vessels which are of interest. After a suitable segmentation has been acquired, we extract geometry information using a center-line extraction technique. From the geometry information, we generate a vector field, where each vector represents the blood vessel direction at the current location. In the second stage, we start by acquiring 3D Doppler ultrasound. In order to correctly input the Doppler ultrasound into our reconstruction technique, we co-register the Doppler ultrasound with our vessel segmentation mask, using image-based tracking [121]. The co-registered Doppler signal is then

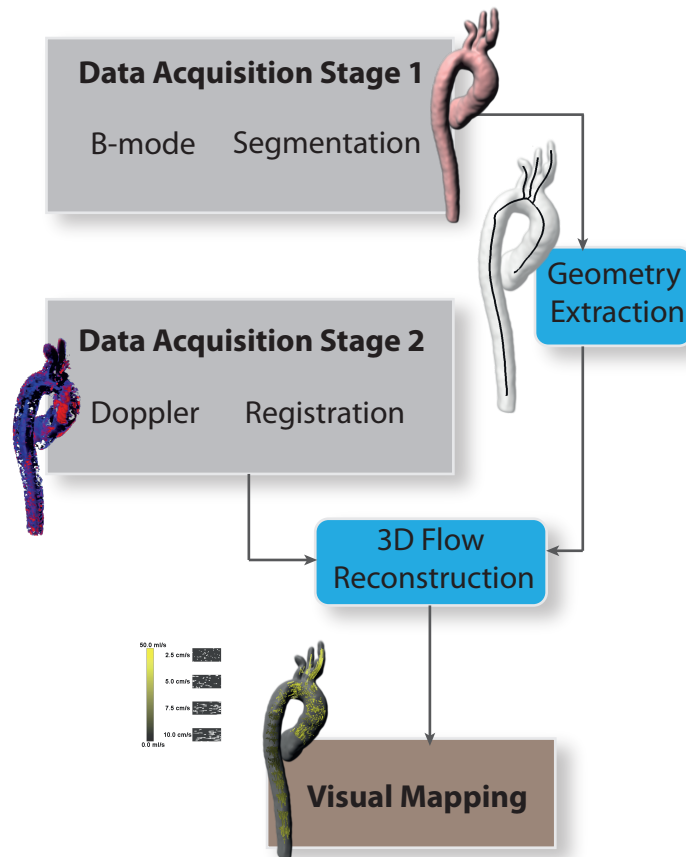


Figure E.2: An overview of our method. Stage 1: Firstly, create a segmentation mask of the blood vessels using B-mode ultrasound. Secondly, acquire geometry information by extracting vessel center-line. Stage 2: Register Doppler ultrasound with previously segmented blood vessels. Secondly, combine the ultrasound information with the geometry information to reconstruct blood flow. Finally, display reconstructed blood flow using illustrative flow rendering. Example images show in-silico example generated from 4D MRI data of the aortic arch.

combined with the vector field, representing the vessel directions. Assuming a laminar flow, we define the flow direction as the direction of the vessels. We reconstruct the 3D blood flow by back-projecting the Doppler information onto the assumed flow direction, generating an unsteady flow field live during an ultrasound examination. Finally, we have designed a tailored visualization mapping for this newly acquired data focused on depicting local velocity and total flow throughput at user-defined cross-sections. The following sections describe the different stages in our pipeline.

E.3.1 Vessel Geometry Extraction

The geometry of the blood vessels is essential for reconstructing the 3D blood flow. While there exist techniques to register pre-segmented vessels from other modalities, such as CT [122], our system does only rely on ultrasound data. Not depending on a multi-modal approach, alleviates the need for a more time-consuming and costly

medical procedure.

To rapidly and robustly extract blood vessels from ultrasound, we have adopted Petersch et al.'s 3D soft segmentation approach [127]. The 3D soft segmentation technique uses a region growing method on a new distance function, which combines the geometric distance to the seeding point and the difference in intensity at a given point and at the seeding point.

The spatial extent of a 3D ultrasound scanning sector is relatively small compared to other modalities. Blood vessels might extend further than what is captured by a single 3D ultrasound sector. In addition, since ultrasound data has a low signal-to-noise ratio and is heavily affected by inconsistent data values, a global segmentation approach is prone to errors. In order to segment larger structures and interactively select which blood vessels to segment live during an examination, we have tailored our method to take into consideration the expertise of the user, allowing the user to locally segment areas, which is suitable through visual feedback. Figure E.3 depicts the size of a segmented vessel using our technique, compared to the size of a 3D ultrasound scanning sector.

In our method, the user applies Petersch et al.'s segmentation method to data acquired during a 3D ultrasound scan. A problem is to combine the segmentation from one time step with the segmentation from the next one. Performing this manually would be very tedious and inefficient. Instead we create a compound volume over the entire area of interest, facilitating the image based tracking [121]. Consequently, we are able to determine the location of ultrasound probe relative to the compound volume.

The compound volume is generated by automatically stitching together several ultrasound sectors from a sweep-scan over a larger area. The composite volume can then work as a global positioning reference for the segmentation process. During the scan, the segmentation from each time frame is directly shown in a 3D view, allowing the user to evaluate the current segment to whether or not it is a desired section of the blood vessel. If the current segment is deemed valid by the user, it is stored in a volume co-registered with the composite volume. The user continues to iterate over this process until a sufficient segmentation mask is acquired.

After a suitable segmentation mask of the vessels has been acquired, we can extract information about the direction of the vessels. Our system assumes that the main flow in blood vessels is laminar, which means the flow vectors in a given cross section are parallel and follow the direction of the vessels. To determine the flow direction, we then need the direction of the vessel itself. The general direction of the vessel corresponds to the direction of its centerline. To robustly extract the center line from ultrasound based data, we require a technique which is not sensitive to noise. We implemented the centerline extraction technique based on level-sets presented by Hassouna et al. [63].

The centerline extraction technique works as follows. First a single centerline point is extracted automatically as the point with the highest distance to the surface. From the starting point a fast wave is propagated throughout the volume. Speed of the front of the wave at a given point is defined as:

$$F(\mathbf{x}) = e^{\alpha D(\mathbf{x})}, \alpha \geq 0, \quad (\text{E.1})$$

where $D(\mathbf{x})$ is the distance to the surface from a point \mathbf{x} and the parameter α is user defined. We followed Hassouna et al.'s advice which states that using a pre-set of




Figure E.3: During the segmentation process, the user can observe the results both in a 2D slice view and a 3D view. The current extracted segment of a hepatic vein is shown in green and the stored segmentation mask is shown in orange.

$\alpha = 25$ provides sufficiently robust results. From Equation E.1, we get a cost function,

$$U(\mathbf{x}) = e^{-\alpha D(\mathbf{x})}, \alpha \geq 0, \quad (\text{E.2})$$

where the minimal cost path between two central points is a centerline. To extract the center-line, we start with the seeding-point and calculate the cost $D(\mathbf{x})$, at each point as the wave propagates. We can then back trace from a topological significant local maxima in the cost function, along the gradient of the cost-function to the starting seeding point, creating a path representing the final centerline.

E.3.2 Flow Reconstruction

Doppler ultrasound is primarily examined as a live stream. In order to achieve live update of the blood flow reconstruction, we generate a vector field where each voxel contains the current vessel direction. Along with the Doppler signal, there is a separate corresponding B-mode image. This enables us to co-register the Doppler ultrasound with the vessel directions by applying the same registration technique used in the segmentation approach.

Doppler ultrasound contains a large amount of noise from movement in surrounding tissue. Before we start the reconstruction we mask out the noise by sampling the the Doppler signal only within the segmented vessel. Doppler shifts are measured relative towards or from the *viewer*. In color Doppler ultrasound, this means the measured direction \mathbf{d} at given point, is the normalized vector from the probe position to that point. The blood flow is denoted as a unit vector \mathbf{u} , which we have estimated to be the direction of the vessel at a given cross-section. If the Doppler ultrasound outputs a value, ω , we can reconstruct the blood flow velocity, \mathbf{v} , using the following equation:




Figure E.4: Reconstructing the actual velocity of the flow by projecting the measured velocity from the Doppler ultrasound onto the vessel direction.

$$\mathbf{v} = \frac{\mathbf{u}}{|\mathbf{u} \cdot \mathbf{d}|} \cdot \omega, \quad (\text{E.3})$$

as illustrated in Figure E.4.

The Doppler ultrasound measurements become more inaccurate the more the blood flow direction becomes orthogonal to the beam direction. Similarly, as we can see in Equation E.3, the flow reconstruction gets more prone to numerical errors as φ approaches 90° , where $\cos \varphi = \mathbf{u} \cdot \mathbf{d}$.

Having reconstructed the flow velocity, we can calculate more aspects of the blood flow. The flux is the volume of flow moving through a cross-section of the blood vessel over time. Since we have the blood vessel already defined, as well as the 3D flow velocities, we estimate the flux, F , by integrating over the area of the cross section, A , using the following equation:

$$F = \int^A \mathbf{v} \approx \sum_i \mathbf{v}_i \quad (\text{E.4})$$

where i is a voxel in the cross-section. The cross-section is defined by inserting plane orthogonal to the center-line. We then sum up the velocities from the samples along the given plane, similarly to the SIVV technique.




Figure E.5: The different legends used for encoding both volume flow and flow-velocity. The users can relate the motion of the particles in the legend to the motion of the particles in the final rendering.

E.3.3 Visualization

The data generated from the two acquisition stages and flow reconstruction presents the challenge of visualizing multiple aspects of flow and the uncertainty of the data, in a meaningful manner. From the first data acquisition stage, we obtain pre-segmented vessels and a compound volume. Then we generate a separate volume containing 3D vectors, representing the direction of the vessels. In the second data acquisition stage, we obtain a live stream of 3D Doppler ultrasound. The Doppler data registered with the vessel directions provide us with unsteady flow data. While the direction of the vessel is constant for each position in the vessel in our model (we abstract from pulsative behaviour of the artery walls close to the heart), the orientation of the flow vector and its length is updated according to the Doppler data, making the flow unsteady. In addition, uncertainty information is updated along the centerline, when the ultrasound probe is moved.

We examined different types of flow visualization mappings, in order to find out which would be suitable for the multiple aspects of our input data. As the flow directions are constant, performing feature analysis, such as vortex detection in a laminar flow would be meaningless. Since we are focused on displaying the velocities and not how the flow evolves over time, methods such as line integral convolution, path lines/surfaces, or streak lines/surfaces are not suitable as they do not depict the current velocity. For a given point in the vessel, the velocity can either be in the direction of the main flow or opposite if the point is in a backwash area. For this reason time-lines and time-surfaces, which connect particles seeded from cross-sections, would be affected by backwash along the vessel-walls, resulting in very skewed lines or surfaces, and would be difficult to comprehend. Texture advection depicts velocity, but is more suited to depict the mixture of the flow, which is not present in laminar flow. This reasoning left particle-rendering as the most suitable means for visual representation of




Figure E.6: A semi-transparent rendering of the current 3D ultrasound cone allows the user to estimate the angle between the probe and the vessel of interest.

reconstructed 3D blood flow.

We consider the blood velocity and the flux as the two main properties in blood flow imaging. Therefore we have chosen to use both moving particles and color to depict the two properties simultaneously. In order to get a meaningful interpretation of the flow, we must control the amount of seeded points. To hinder cluttering by an over-abundance of particles, we only seed out particles at user selected areas. The user can simply select an area along the center-line and particles are seeded randomly in the orthogonal cross-section. The motion of the particles depicts the velocities from the cross-section. To prevent perception of the particles from one seeding location to interfere with particles from another, we give each particle a short life time. To increase the spatial understanding of the seeding planes, we insert a circle at the cross-section. The circle helps the user to understand not only the position, but the orientation of the plane provides a circular indicatrix of the vessel direction. Recent studies have shown a systematic distortion of speed perception of moving particles [15]. They provide a compensation model, which adjusts the velocities in order for the viewer to get a better estimate of the speed encoded in the underlying data. We include this compensation model in our particle animation scheme. Finally, the flux from the seeding cross section is color coded on the particle.

Turbulent flow can occur in blood vessels and this contradicts the assumption of laminar flow. In this case the regions with turbulence will appear as regions where many particles are moving in opposite directions without a distinct pattern, such as a back wash region. In this case it is also important to show the original Doppler signal. We have incorporated a slice viewer which shows a slice of the Doppler signal at the current position of the probes position along the center-line. This enables the user to see the original data at for validation purposes.

Similarly to color coding, velocity encoding requires a means to estimate the numerical value of the particles in the scene. For this we propose to use a novel visual metaphor, a velocity legend. The velocity legend contains a discrete set of numerical flow values, where each value has a small *flow sample* related to it. Figure E.5 shows an example of the velocity legend, where we can see four discrete velocities and the




Figure E.7: Segmentation mask and centerline of the inferior vena cava with hepatic branches including an aberrant vessel entering below. The vessels were segmented using our approach in about four minutes.

corresponding particle motion.

Uncertainty of the data is another characteristic which is important to communicate to the user. When visualizing data, it can be crucial to inform the user about to what degree she can trust the information in the final image. The reconstructed blood flow is based on a simplified flow model compared to using complex simulation, such as computational fluid dynamics. Opacity is often used as a means to convey uncertainty. However, color is already used to depict flux. An ultrasound examiner is trained to set the angle between the probe and the blood vessels to be not exceeding 60° . Adding the current ultrasound sector as an overlay around the blood vessel, as in Figure E.6, the user can estimate the angle between the vessel and the ultrasound probe. Thus, the user can determine how reliable the measurements are.

E.4 Results

We have implemented the prototype in the VolumeShop framework [19]. The prototype was tested on a Windows 7 system, running a Intel Core i7 3 GHz CPU with a GeForce GTX 680 graphics card. 3D B-mode and color Doppler ultrasound were acquired using a GE Vingmed Vivid E9 scanner, using a 4V-D ultrasound probe [50].

As a demonstrational case, we have acquired color Doppler measurements from two healthy volunteers. The subject of analysis was the inferior vena cava and the connecting hepatic veins bringing low oxygenated blood back to the heart. These vascularities are frequently investigated for blood velocities and flux, in situations of portal hypertension (post-hepatic) [106]. From a B-mode ultrasound data stream we have segmented the inferior vena cava and the hepatic veins. The resulting segmentation and the extracted centerline can be seen in Figure E.7. The segmentation process took after a training period, four minutes. The examiner then swept over the inferior vena cava with 3D color Doppler ultrasound. The velocities are reconstructed in real-time at user




Figure E.8: Blood flow reconstructed in the inferior vena cava.

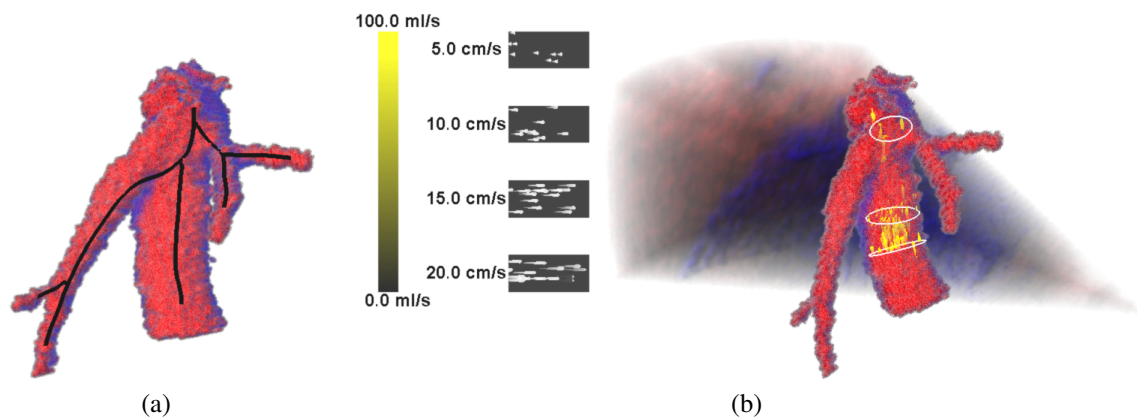


Figure E.9: Blood flow reconstructed in the inferior vena cava.

selected cross sections seeding particles into the newly acquired flow field. In Figure E.8, we show the traced particles from two cross-sections. In Figure E.9, images of the inferior vena cava from a different patient are shown.

E.5 Evaluation

When new data is reconstructed based on a limited amount information, we need an evaluation of the technique to determine its usefulness. We have addressed two aspects: accuracy of the flow reconstruction and user satisfaction of the technique as such. For the accuracy, we compare our results with the actual flow and for the user satisfaction, we evaluate subjectively, how easy users find the implemented prototype to operate.

E.5.1 Accuracy

To compare the difference between the reconstructed flow and the actual flow, we need a way to measure the *ground truth*. In order to measure the accuracy in our 3D flow reconstruction, we wish to remove noise from tissue movement and probe-movement as much as possible. For this reason, we base our accuracy measurement on an in-silico phantom constructed from the flow measured by 4D MRI scans. We have acquired a 4D MRI flow dataset with a corresponding segmentation mask of the aortic arch. The Doppler signal is synthetically generated as the component vector from a predefined virtual probe position. In this way, we can create artificial Doppler signal from multiple directions to compare how the accuracy compares to the angle between the ultrasound-beam direction and the vessel direction. In Figure E.10, rendering of the 4D MRI dataset are shown comparing reconstructed flow (b) with real flow (c).

We extract flux estimates from both the recorded 4D MRI blood flow from one time-step, with the corresponding estimated generated by our reconstruction method. Samples are extracted from uniformly distributed cross-sections along the vessel center-line. In total we extracted 1344 samples using four different simulated probe-positions, getting a average estimation of 91% of actual flow. We assume that the error increases as the angle between the ultrasound beam-direction and the flow direction approaches 90° . In Figure E.11, we collected the error into bins of 2.5° intervals and plotted the mean and standard deviation of our results. From this graph we can see that the precision of our estimated flow is severely lowered when the angle between the beam direction and vessel direction exceeds 55° . Our method tends to slightly underestimate at degrees lower than 55° , but overestimates when the angle increases. We can assume this overestimation causes bias in our calculation of the total average accuracy. If we look at the average below 55° , we get an accuracy of 69%.

The flow in the aortic arch is much more turbulent compared to flow in liver vessels due to lower blood velocities. The high turbulence will cause problems when we are assuming a laminar flow, lowering the accuracy. This would indicate a more robust reconstruction is achieved in vessels, such as the inferior vena cava. With these results from the in-silico experiment, we can conclude that flow estimation is possible in laminar-flow dominated veins. However, Doppler ultrasound acquisition is also affected by an error and further investigation is required determine the actual error in vivo.

E.5.2 User Satisfaction

During the segmentation procedure, we asked our medical partners to evaluate the workflow of the segmentation process. In general, it was found to be a quick and robust tool for vessel segmentation compared to other segmentation techniques based on level-sets or manual segmentation, our segmentation approach makes it possible for a novice user to segment actual vessel trees within an average of 7-10 minutes. Aspects such as a live updated 3D view of the current segment as well as the larger combined segmentation, which gave direct view of the segmentation progress, was much appreciated. Testing proved that it was difficult to extract smaller vessels due to lack of a *zooming* feature, which can be important for micro-vascular perfusion analysis. In certain areas, the data values for the blood vessels were too similar to the data values in

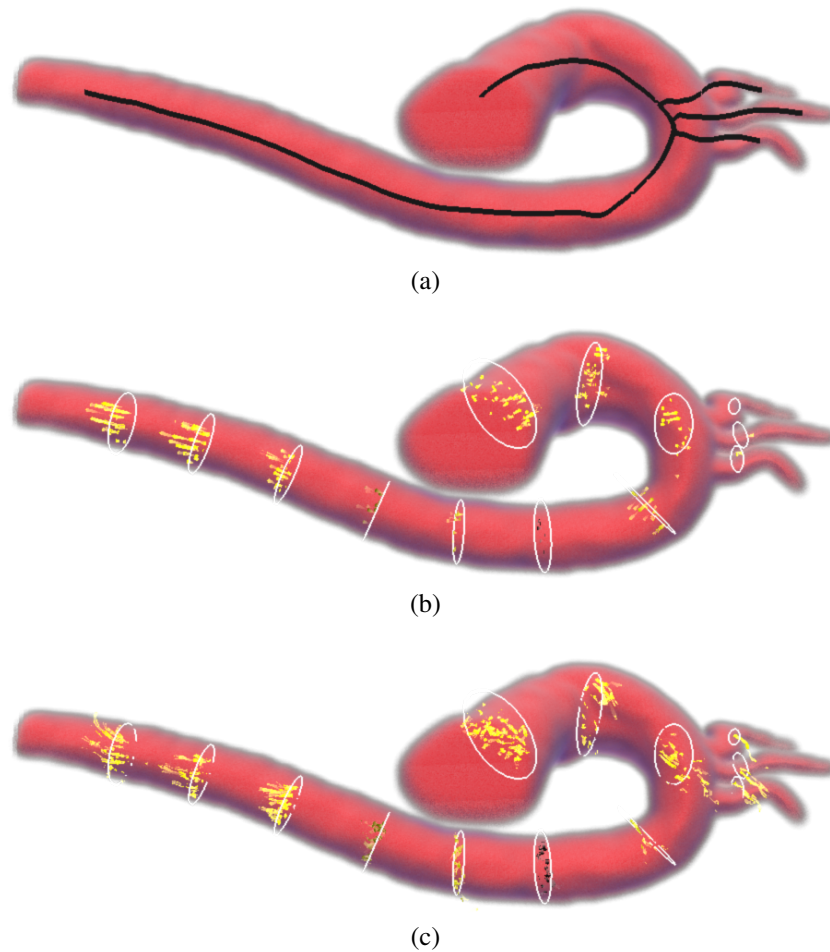


Figure E.10: Renderings from our in-silico experimental set-up. Starting with the segmentation mask of the aortic arch, we generate the centerline (a). Images (b) and (c) show our visualization technique applied to the reconstructed velocities and the velocities from a 4D MRI scan.

the surrounding tissue. In these cases, our system was experiencing difficulties differentiating vessels from other tissue. In such a case, features such as manual editing on the 2D slice as well as directly in 3D view was noted as missing.

For the flow reconstruction mapping, our medical partners stated that the reconstruction method showed great potential for the evaluation of degree of stenosis. Also, stent and graft patency are two interesting indications for this method to be further explored with high probability of usefulness. Further testing would be required for the performance of the velocity legend as a means for estimating velocity in a 3D environment.

E.6 Conclusion

In this paper, we have presented a new method for reconstructing 3D blood flow based on 3D B-mode and 3D color Doppler ultrasound. By segmenting blood vessels from B-mode ultrasound, we anticipate the direction of blood vessels as the direction of the corresponding centerline. In our reconstruction model, we assume that blood flow




Figure E.11: The average estimation in percent compared to the angle between the ultrasound beam direction and the blood flow direction, based on an in-silico experiment from 4D MRI data.

through vessels is in general laminar, and defines the 3D flow direction as the direction of the vessel. Since the Doppler measurements are only relative to the ultrasound probe position, we back project the data onto the 3D flow directions. Finally, we have proposed a visualization design, tailored for the data generated from our reconstruction technique.

We have performed an evaluation of the accuracy of our reconstruction using 4D MRI data of the aortic arch, by simulating the Doppler signal and using our technique to reconstruct the real velocity. The reconstructed data was then compared to the *ground truth* from the 4D MRI dataset, by comparing the flux at cross-sections along the centerline. We have found that our technique can provide a good estimation of the actual flux. However, the accuracy is highly dependent on the orthogonality between the ultrasound beam direction and the vessel direction. While our accuracy measurement is under ideal conditions, more testing is required to see the performance on *real* ultrasound data. While it is difficult to measure the *real* blood flow non-invasively, future work can include manual blood flow estimation, compared to our method.

Based on our in-silico experiments, we can state that a streamlined implementation of our technique can improve the usability of 3D color Doppler ultrasound for 3D blood flow estimation.

*«Some days even my lucky rocketship
underpants won't help» –Calvin*

Bibliography

- [1] Nvidia PhysX Engine. <http://uk.geforce.com/hardware/technology/physx>, 2011. B.4
- [2] Illustration by David S. Goodsell, the Scripps Research Institute. <http://mgl.scripps.edu/people/goodsell>, 2012. B.1
- [3] ABBOTT, I., AND VON DOENHOFF, A. *Theory of Wing Sections: Including a Summary of Airfoil Data*. Dover Books on Physics and Chemistry. Dover Publ., 1959. D.1
- [4] ACEVEDO, D., JACKSON, C. D., DRURY, F., AND LAIDLAW, D. H. Using Visual Design Experts in Critique-Based Evaluation of 2D Vector Visualization Methods. *Transactions on Visualization and Computer Graphics* 14, 4 (2008), 877–884. D.1
- [5] ALLEN, R. C., AND RUBIN, M. L. Chromostereopsis. *Survey of Ophthalmology* 26, 1 (1981), 22–27. C.6.3
- [6] ANGELELLI, P., VIOLA, I., NYLUND, K., GILJA, O. H., AND HAUSER, H. Guided visualization of ultrasound image sequences. In *Proceedings of Visual Computing for Biomedicine* (2010), pp. 125–132. 2.2, C.6.2
- [7] ANGELSEN, B. A. J. *Waves, signals and signal processing in medical ultrasonics*, vol. I. Dept. of Physiology and Biomedical Engineering, Norwegian University of Science and Technology, Trondheim, Norway, 1995. 2.1
- [8] ANGLESEN, B. A. J., TORP, H., HOLM, S., KRISTOFFERSEN, K., AND WHITTINGHAM, T. A. Which transducer array is best? *European Journal of Ultrasound* 2, 2 (1995), 151–164. 2.1
- [9] BADE, R., RIEDEL, I., SCHMIDT, L., OLDHAFFER, K. J., AND PREIM, B. Combining training and computer-assisted planning of oncologic liver surgery. In *Bildverarbeitung für die Medizin* (2006), pp. 409–413. A.1
- [10] BAJURA, M., FUCHS, H., AND OHBUCHI, R. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Proceedings of ACM SIGGRAPH* 26, 2 (1992), 203–210. C.7
- [11] BEICHEL, R., POCK, T., JANKO, C., ZOTTER, R. B., REITINGER, B., BORNIK, A., PALAGYI, K., SORANTIN, E., WERKGARTNER, G., BISCHOF,

- H., AND SONKA, M. Liver segment approximation in CT data for surgical resection planning. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (May 2004), vol. 5370 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, pp. 1435–1446. [A.2](#)
- [12] BELOHLAVEK, M., DUTT, V., GREENLEAF, J., FOLEY, D. A., GERBER, T. C., AND SEWARD, J. B. Multidimensional ultrasonic visualization in cardiology. In *Proceedings of IEEE Ultrasonics Symposium* (1992), vol. 2, pp. 1137–1145. [C.3.2](#)
- [13] BERG, S., TORP, H., MARTENS, D., STEEN, E., SAMSTAD, S., HØIVIK, I., AND OLSTAD, B. Dynamic three-dimensional freehand echocardiography using raw digital ultrasound data. *Ultrasound in Medicine and Biology* 25, 5 (1999), 745–753. [C.3.1](#)
- [14] BEYER, J., HADWIGER, M., WOLFSBERGER, S., AND BUHLER, K. High-Quality Multimodal Volume Rendering for Preoperative Planning of Neurosurgical Interventions. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1696–1703. [B.2](#)
- [15] BIRKELAND, Å., TURKAY, C., AND VIOLA, I. Perceptually uniform motion space. *IEEE Transactions on Visualization and Computer Graphics* (2013), Submitted. [E.3.3](#)
- [16] BIRKELAND, Å., AND VIOLA, I. View-dependent peel-away visualization for volumetric data. In *Proceedings of Spring Conference on Computer Graphics* (2009), pp. 133–139. [4](#)
- [17] BIRKELAND, Å., AND VIOLA, I. Ultrasound painting of vascular tree. In *Proceedings of Vision, Modeling, and Visualization* (2010), pp. 163–170. [E.2](#)
- [18] BLAKEMORE, M. R., AND SNOWDEN, R. J. The Effect of Contrast Upon Perceived Speed: A General Phenomenon? *Perception* 28, 3 (1999), 33–48. [4.3](#), [D.2](#)
- [19] BRUCKNER, S., AND GRÖLLER, M. E. Volumeshop: An interactive system for direct volume illustration. In *Proceedings of IEEE Visualization* (2005), pp. 671–678. [4](#), [A.5](#), [E.4](#)
- [20] BRUCKNER, S., AND GRÖLLER, M. E. Exploded views for volume data. *Transactions on Visualization and Computer Graphics* 12, 5 (2006), 1077–1084. [2.2](#), [4](#), [B.2](#)
- [21] BRUNO, R., RICHARD, B., AND NICOLAS, J. An Interactive Tool To Visualize Three-Dimensional Ultrasound Data. *MMW Fortschritte der Medizin* 149, 15 (2000), 56. [2.2](#)
- [22] BUCHAR, C., BORRO, D., AND AMUNDARAIN, A. GPU local triangulation: an interpolating surface reconstruction algorithm. In *Proceedings of EUROGRAPHICS* (2008), vol. 27. [3.1](#), [A.4](#), [A.4.2](#)

- [23] BURCKHARDT, C. B. Speckle In Ultrasound B-mode Scans. *IEEE Transactions on Sonics and Ultrasonics* 25, 1 (1978), 1–6. [2.1](#)
- [24] BURNS, M., HAIDACHER, M., WEIN, W., VIOLA, I., AND GRÖLLER, E. Feature emphasis and contextual cutaways for multimodal medical visualization. In *Proceedings of EUROGRAPHICS* (2007), pp. 275–282. [2.2](#), [4](#), [B.2](#), [C.5](#), [C.5](#), [C.6.2](#)
- [25] CARNEIRO, G., AMAT, F., GEORGESCU, B., GOOD, S., AND COMANICIU, D. Semantic-based indexing of fetal anatomies from 3-D ultrasound data using global/semi-local context and sequential sampling. *Proceedings of Computer Vision and Pattern Recognition* (2008), 1–8. [2.2](#), [C.4](#)
- [26] CAVANAGH, P., TYLER, C. W., AND FAVREAU, O. E. Perceived Velocity of Moving Chromatic Gratings. *Journal of The Optical Society of America A-optics Image Science and Vision I* (1984). [D.2](#)
- [27] CHEN, Y., YIN, R., FLYNN, P., AND BROCHAT, S. Aggressive region growing for speckle reduction in ultrasound images. *Pattern Recognition Letters* 24, 4-5 (2003), 677–691. [C.3.2](#)
- [28] COUINAUD, C. *Le Foie Etudes anatomiques et chirurgicales*, vol. 1. 1957. [A.1](#)
- [29] CRIMMINS, T. R. The geometric filter for speckle reduction. *Applied Optics* 24, 10 (1985), 1438 – 1443. [C.3.2](#)
- [30] CURIEL, L., CHOPRA, R., AND HYNYNEN, K. Progress in multimodality imaging: Truly simultaneous ultrasound and magnetic resonance imaging. *IEEE Transactions on Medical Imaging* 26, 12 (December 2007), 1740–1746. [C.5](#)
- [31] DEPRez, J., BRUSSEAU, E., SCHMITT, C., CLOUTIER, G., AND BASSET, O. 3d estimation of soft biological tissue deformation from radio-frequency ultrasound volume acquisitions. *Medical Image Analysis* 13, 1 (2009), 116 – 127. [C.3](#)
- [32] DERRINGTON, A., AND HENNING, G. Detecting and Discriminating the Direction of Motion of Luminance and Colour Gratings. *Vision Research* 33, 5-6 (1993), 799 – 811. [D.2](#)
- [33] DEUTCH, T. D., JOERGNER, I., MATSON, D. O., OEHNINGER, S., BOCCA, S., HÖNIGMANN, D., AND ABUHAMAD, A. Automated assessment of ovarian follicles using a novel three-dimensional ultrasound software. *Fertility and Sterility* 92, 5 (2009), 1562–1568. [2.1](#), [2.2](#), [C.3](#), [C.4](#)
- [34] DIAZ, J., MONCLUS, E., NAVAZO, I., AND VAZQUEZ, P. Adaptive cross-sections of anatomical models. *Computer Graphics Forum* 31, 7pt2 (2012), 2155–2164. [4.1](#)
- [35] DODT, H.-U., LEISCHNER, U., SCHIERLOH, A., MAUCH, N. J. C. P., DEININGER, K., DEUSSING, J. M., EDER, M., ZIEGLGÄNSBERGER, W., AND

- BECKER, K. Ultramicroscopy: three-dimensional visualization of neuronal networks in the whole mouse brain. *Nature Methods* 4, 4 (2007), 331–336. [B.6](#), [B.8](#)
- [36] EINTHOVEN, W. Stereoscopie durch Farbendifferenz. *Von Graefes Archiv* 19 (1885), 211–238. [C.6.3](#)
- [37] ERDT, M., RASPE, M., AND SÜHLING, M. Automatic hepatic vessel segmentation using graphics hardware. In *MIAR (2008)*, vol. 5128 of *Lecture Notes in Computer Science*, Springer, pp. 403–412. [A.2](#)
- [38] ESNEAULT, S., LAFON, C., AND DILLENSEGER, J.-L. Liver vessels segmentation using a hybrid geometrical moments/graph cuts method. *Biomedical Engineering, IEEE Transactions on* 57, 2 (feb. 2010), 276–283. [A.2](#)
- [39] FATTAL, R., AND LISCHINSKI, D. Variational classification for visualization of 3D ultrasound data. *Proceedings of IEEE Visualization (2001)*, 403–410. [C.6](#)
- [40] FENG, S., ZHOU, S. K., GOOD, S., AND COMANICIU, D. Automatic fetal face detection from ultrasound volumes via learning 3D and 2D information. *Proceedings of Computer Vision and Pattern Recognition (2009)*, 2488–2495. [2.2](#)
- [41] FORSBERG, F., HEALEY, A. J., LEEMAN, S., AND JENSEN, J. A. Assessment of hybrid speckle reduction algorithms. *Physics in Medicine and Biology* 36, 11 (1991), 1539–1549. [C.3.2](#)
- [42] FRISBY, J. P., AND STONE, J. V. *Seeing : The Computational Approach to Biological Vision*, 2nd ed ed. Cambridge, Mass. : MIT Press, 2010. [D.2](#)
- [43] FUCHS, R., WELKER, V., AND HORNEGGER, J. Non-convex polyhedral volume of interest selection. *Journal of Computerized Medical Imaging and Graphics* 34, 2 (2009), 105–113. [B.2](#)
- [44] FULLER, D. B., JIN, H., KOZIOL, J. A., AND FENG, A. C. CT-ultrasound fusion prostate brachytherapy: a dynamic dosimetry feedback and improvement method. A report of 54 consecutive cases. *Brachytherapy* 4, 3 (januar 2005), 207–16. [C.5](#)
- [45] GARRETT, W. F., FUCHS, H., WHITTON, M. C., AND STATE, A. Real-time incremental visualization of dynamic ultrasound volumes using parallel BSP trees. *Proceedings of IEEE Visualization (1996)*, 235–240,. [C.6](#)
- [46] GASTEIGER, R., NEUGEBAUER, M., BEUING, O., AND PREIM, B. The FLOWLENS: A focus-and-context visualization approach for exploration of blood flow in cerebral aneurysms. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2183–2192. [E.2](#)
- [47] GASTEIGER, R., NEUGEBAUER, M., KUBISCH, C., AND PREIM, B. Adapted surface visualization of cerebral aneurysms with embedded blood flow information. In *Proceedings of EG Workshop on Visual Computing in Biology and Medicine VCBM 2010 (2010)*, pp. 25–32. [E.2](#)

- [48] GE HEALTHCARE. A closer look at GE's pocket-sized vscan ultrasound, 2010. <http://www.gereports.com/a-closer-look-at-ge-pocket-sized-vscan-ultrasound/>. 1, C.1
- [49] GE HEALTHCARE. Voluson E8, HDlive imaging – experience realism with volume ultrasound, 2012. www.gehealthcare.com/ultrasound. 2.2, C.6.3
- [50] GE VINGMED. Vivid E9. <http://www3.gehealthcare.com/ultrasound>, 2008. E.4
- [51] GEE, A., PRAGER, R., TREECE, G., AND BERMAN, L. Narrow-band volume rendering for freehand 3D ultrasound. *Computers & Graphics* 26, 3 (2002), 463–476. C.6
- [52] GEE, A., PRAGER, R., TREECE, G., CASH, C., AND BERMAN, L. Processing and visualizing three-dimensional ultrasound data. *The British Journal of Radiology* 77 (2004), 186–193. C.3.1
- [53] GILJA, O. H., HATLEBAKK, J. G., ØDEGAARD, S., BERSTAD, A., VIOLA, I., GIERTSEN, C., HAUSKEN, T., AND GREGERSEN, H. Advanced imaging and visualization in gastrointestinal disorders. *World Journal of Gastroenterology* 13 (2007). C.1
- [54] GILJA, O. H., HAUSKEN, T., ØDEGAARD, S., WENDELBO, Ø., AND THIERLEY, M. Mobile ultrasonography in a medical department. *Tidsskrift for Norsk Lægeforening* 19 (2003), 270–285. C.1
- [55] GILJA, O. H., HEIMDAL, A., HAUSKEN, T., GREGERSEN, H., MATRE, K., BERSTAD, A., AND ØDEGAARD, S. Strain during gastric contractions can be measured using doppler ultrasonography. *Ultrasound in Medicine and Biology* 28, 11-12 (2002), 1457 – 1465. C.1
- [56] GOBBI, D., AND PETERS, T. Interactive intra-operative 3D ultrasound reconstruction and visualization. *Proceedings of Medical Image Computing and Computer-Assisted Intervention* (2002), 156–163. C.3.1
- [57] GOMEZ, A., SIMPSON, J., YAO, C., SCHAEFFTER, T., AND PENNEY, G. 3d flow reconstruction from multiple registered echo doppler views. In *Proceedings of IEEE International Symposium on Biomedical Imaging: Macro to Nano* (2011), pp. 879–882. 4.2
- [58] GONZALEZ, R. C., AND WOODS, R. E. *Digital image processing*. Upper Saddle River, N.J., 2008. A.3.2
- [59] GOODING, M. J., MELLOR, M., SHIPLEY, J. A., BROADBENT, K. A., AND GODDARD, D. A. Automatic mammary duct detection in 3D ultrasound. In *MICCAI* (2005), vol. 3749 of *Lecture Notes in Computer Science*, pp. 434–441. A.2
- [60] GRAU, V., BECHER, H., AND NOBLE, J. Registration of multi view real-time 3D echocardiographic sequences. *Medical Imaging, IEEE Transactions on* 26, 9 (sept. 2007), 1154 –1165. A.2

- [61] GUPTA, N., SWAMY, M. N. S., AND PLOTKIN, E. Despeckling of medical ultrasound images using data and rate adaptive lossy compression. *Medical Imaging, IEEE Transactions on* 24, 6 (2005), 743–754. [C.3.2](#)
- [62] HAIRER, E., NØRSETT, S. P., AND WANNER, G. *Solving Ordinary Differential Equations I*, 2nd ed. Springer Series in Computational Mathematics. Springer, Berlin Heidelberg, 1993. [B.3.1](#)
- [63] HASSOUNA, M., AND FARAG, A. Robust Centerline Extraction Framework using Level Sets. In *Proceedings of IEEE Computer Vision and Pattern Recognition, CVPR 2005* (2005), vol. 1, pp. 458–465. [4.2](#), [E.3.1](#)
- [64] HAUSKEN, T., LI, X., GOLDMAN, B., LEOTTA, D., ØDEGAARD, S., AND MARTIN, R. Quantification of gastric emptying and duodenogastric reflux stroke volumes using three-dimensional guided digital color Doppler imaging. *European Journal of Ultrasound* 13, 3 (2001), 205–213. [E.2](#)
- [65] HEIMDAL, A., STØYLEN, A., TORP, H., AND SKJÆRPE, T. Real-time strain rate imaging of the left ventricle by ultrasound. *Journal of the American Society of Echocardiography* 11, 11 (1998), 1013 – 1019. [C.3](#)
- [66] HOMANN, H., VESOM, G., AND NOBLE, J. A. Vasculature segmentation of CT liver images using graph cuts and graph-based analysis. In *ISBI* (2008), pp. 53–56. [A.2](#)
- [67] HÖNIGMANN, D., RUISZ, J., AND HAIDER, C. Adaptive design of a global opacity transfer function for direct volume rendering of ultrasound data. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* (2003), 489–496. [C.6.1](#)
- [68] HOUCK, R. C., COOKE, J., AND GILL, E. A. Three-dimensional echo: transition from theory to real-time, a technology now ready for prime time. *Current Problems in Diagnostic Radiology* 34, 3 (May-Jun 2005), 85–105. [2.1](#)
- [69] HOUSE, D. H., AND BREEN, D. E., Eds. *Cloth modeling and animation*. 2000. [B.2](#)
- [70] HUANG, T., YANG, G., AND TANG, G. A fast two-dimensional median filtering algorithm. *Acoustics, Speech and Signal Processing, IEEE Transactions on* 27, 1 (February 1979), 13 – 18. [C.3.2](#)
- [71] ISLAM, S., SILVER, D., AND CHEN, M. Volume Splitting and Its Applications. *IEEE Transactions on Visualization and Computer Graphics* 13, 2 (2007), 193–203. [B.2](#)
- [72] JAN, U., SANDKÜHLER, D., RAUBERGER, M., MATTHIES, H. K., AND OVERHOFF, H. M. Enhanced visualization of ultrasound volumes for diagnostic and therapeutic purposes. In *Proceedings of 4th European Conference of the International Federation for Medical and Biological Engineering, IFMBE Proceedings*. 2009, pp. 689–692. [C.6.1](#)

- [73] JENSEN, J. A. *Estimation of Blood Velocities Using Ultrasound. A Signal Processing Approach.*, 1 ed. Cambridge University Press, Cambridge, United Kingdom, 1996. [2.1](#)
- [74] JONES, M. G., SHIPLEY, J. A., AND ROBINSON, T. M. Visualisation of 4-d colour and power doppler data. *Ultrasound in Medicine and Biology* 29, 12 (2003), 1735–1747. [2.2](#), [4.2](#), [C.6.2](#), [E.2](#)
- [75] KANITSAR, A., FLEISCHMANN, D., WEGENKITTL, R., FELKEL, P., AND GRÖLLER, M. E. CPR - Curved Planar Reformation. Tech. rep., Institute of Computer Graphics and Algorithms, Vienna University of Technology, 2002. [B.2](#)
- [76] KARADAYI, K., MANAGULI, R., AND KIM, Y. Three-dimensional ultrasound: From acquisition to visualization and from algorithms to systems. *IEEE Reviews in Biomedical Engineering* 2 (2009), 23–39. [C.2](#)
- [77] KARAMALIS, A., WEIN, W., KUTTER, O., AND NAVAB, N. Fast hybrid free-hand ultrasound volume reconstruction. In *SPIE Medical Imaging* (2009). [C.3.1](#)
- [78] KAWAJIRI, S., ZHOU, X., ZHANG, X., HARA, T., FUJITA, H., YOKOYAMA, R., KONDO, H., KANEMATSU, M., AND HOSHI, H. Automated segmentation of hepatic vessel trees in non-contrast x-ray ct images. vol. 6512, p. 65123A. [A.2](#)
- [79] KIM, E., MANAGULI, R., AND KIM, Y. New flexible multi-volume rendering technique for ultrasound imaging. In *SPIE Medical Imaging* (2009), vol. 7265, p. 72651D. [2.2](#)
- [80] KINDLMANN, G., AND DURKIN, J. W. Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of IEEE Symposium on Volume Visualization* (1998), pp. 79–86. [C.6.1](#)
- [81] KING, A. P., MA, Y., YAO, C., JANSEN, C., RAZAVI, R., RHODE, K. S., AND PENNEY, G. P. Image-to-physical registration for image-guided interventions using 3-d ultrasound and an ultrasound imaging model. In *Proceedings of International Conference on Information Processing in Medical Imaging* (2009), pp. 188–201. [2.2](#), [A.2](#), [C.5](#)
- [82] KIRBAS, C., AND QUEK, F. A review of vessel extraction techniques and algorithms. *ACM Computing Surveys* 36, 2 (2004), 81–121. [A.2](#)
- [83] KITAGAWA, T., ZHOU, X. R., HARA, T., FUJITA, H., YOKOYAMA, R., KONDO, H., KANEMATSU, M., AND HOSHI., H. Automated segmentation of middle hepatic vein in non-contrast x-ray ct images based on an atlas-driven approach. *SPIE Medical Imaging* (2008). [A.2](#)
- [84] KNISS, J., KINDLMANN, G., AND HANSEN, C. Multidimensional transfer functions for interactive volume rendering. *Transactions on Visualization and Computer Graphics* 8, 3 (2002), 270–285. [C.6.1](#)

- [85] KNISS, J., PREMOZE, S., HANSEN, C., AND EBERT, D. Interactive translucent volume rendering and procedural modeling. In *Proceedings of IEEE Visualization* (2002), pp. 109–116. [C.6.3](#), [C.8d](#)
- [86] KONINKLIJKE, PHILIPS, AND ELECTRONICS. iU22 xMATRIX Ultrasound System. <http://www.healthcare.philips.com>. [2.1](#)
- [87] KONRAD-VERSE, O., LITTMANN, A., AND PREIM, B. Virtual Resection with a Deformable Cutting Plane. In *SimVis* (2004), pp. 203–214. [B.2](#)
- [88] KRETZTECHNIK AG. 3d ultrasound: A dedicated system. *European Radiology* 9 (1999), S331–S333. [2.2](#), [B.2](#), [C.4](#), [C.4.1](#)
- [89] KRIPFGANS, O. D., RUBIN, J. M., HALL, A. L., GORDON, M. B., AND FOWLKES, J. B. Measurement of volumetric flow. *Journal of Ultrasound in Medicine* 25, 10 (2006), 1305–11. [E.2](#)
- [90] KRISSIAN, K., WESTIN, C. F., KIKINIS, R., AND VOSBURGH, K. G. Oriented speckle reducing anisotropic diffusion. *IEEE Transactions on Image Processing* 16, 5 (2007), 1412–1424. [C.3.2](#)
- [91] KUO, J., BREDTHAUER, G. R., CASTELLUCCI, J. B., AND VON RAMM, O. T. Interactive volume rendering of real-time three-dimensional ultrasound images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 54, 2 (2007), 313–8. [C.6](#)
- [92] LANGE, T., EULENSTEIN, S., HUNERBEIN, M., LAMECKER, H., AND SCHLAG, P. M. Augmenting intraoperative 3D ultrasound with preoperative models for navigation in liver surgery. *Proceedings of Medical Image Computing and Computer-Assisted Intervention* (2004), 534–541. [2.2](#), [A.2](#)
- [93] LANGE, T., PAPPENBERG, N., HELDMANN, S., MODERSITZKI, J., FISCHER, B., LAMECKER, H., AND SCHLAG, P. 3D Ultrasound-CT Registration of the Liver Using Combined Landmark-Intensity Information. *International Journal of Computer Assisted Radiology and Surgery* 4 (2008), 79–88. [2.2](#), [C.5](#)
- [94] LEE, T., CHAUDHURI, A., PORIKLI, F., AND SHEN, H. Cyclestack: Inferring periodic behavior via temporal sequence visualization in ultrasound video. In *Proceedings of IEEE Pacific Visualization* (2010), pp. 89–96. [C.6](#)
- [95] LEROY, A., MOZER, P., PAYAN, Y., AND TROCCAZ, J. Rigid registration of freehand 3D ultrasound and CT-scan kidney images. *Proceedings of Medical Image Computing and Computer-Assisted Intervention* (2004), 837–844. [2.2](#), [A.2](#)
- [96] LEVOY, M. Display of surfaces from volume data. *IEEE Computer Graphics and Applications* 8, 3 (1988), 29–37. [C.6.3](#)
- [97] LI, S., WAITE, J. M., LENNON, B. T., STEFANSIC, J. D., LI, R., AND DAWANT, B. M. Development of preoperative liver and vascular system segmentation and modelling tool for image-guided surgery and surgical planning.

- In *Medical Imaging 2008: Visualization, Image-guided Procedures, and Modelling* (2008), vol. 6918, p. 69180C. [A.2](#)
- [98] LIM, S., KWON, K., AND SHIN, B. S. GPU-based interactive visualization framework for ultrasound datasets. *Computer Animation and Virtual Worlds* 20, 1 (2008), 11–23. [C.6](#)
- [99] LINDEMANN, F., AND ROPINSKI, T. Advanced Light Material Interaction for Direct Volume Rendering. In *Proceedings of IEEE/EG International Symposium on Volume Graphics* (2010), pp. 101–108. [C.6.3](#), [C.8f](#)
- [100] LIZZI, F. L., AND FELEPPA, E. J. Image processing and pre-processing for medical ultrasound. In *Proceedings of Imagery Pattern Recognition Workshop* (2000), p. 187. [C.3.2](#)
- [101] LUNN, K. E., HARTOV, A., KENNEDY, F. E., MIGA, M. I., ROBERTS, D. W., PLATENIK, L. A., AND PAULSEN, K. D. 3D ultrasound as sparse data for intra-operative brain deformation model. In *Proceedings of SPIE* (2001), vol. 4325, p. 326. [2.2](#)
- [102] MAGNENAT-THALMANN, N., CORDIER, F., KECKEISEN, M., KIMMERLE, S., KLEIN, R., AND MESETH, J. Simulation of Clothes for Real-time Applications. In *Proceedings of EUROGRAPHICS* (2004). [B.2](#)
- [103] MATHER, G. Luminance Change Generates Apparent Movement: Implications for Models of Directional Specificity in the Human Visual System. *Vision Research* 24, 10 (1984), 1399 – 1405. [4.3](#), [D.2](#), [D.8](#)
- [104] MEINZER, H. P., THORN, M., AND CARDENAS, C. E. Computerized planning of liver surgery—an overview. *Computers and Graphics* 26, 4 (2002), 569 – 576. [A.2](#)
- [105] MERZ, E., BAHLMANN, F., AND WEBER, G. Volume Scanning in the Evaluation of Fetal Malformations - A New Dimension in Prenatal-Diagnosis. *Ultrasound In Obstetrics & Gynecology* 5, 4 (1995), 222–227. [2.1](#)
- [106] MITTAL, P., GUPTA, R., ABD, G. M., AND KALIA³, V. Association between portal vein color doppler findings and the severity of disease in cirrhotic patients with portal hypertension. *Iran J Radiol* 8 (2011), 211–217. [E.4](#)
- [107] MÜLLER, M., HEIDELBERGER, B., HENNIX, M., AND RATCLIFF, J. Position Based Dynamics. *Journal of Visual Communication and Image Representation* 18, 2 (2007), 109 – 118. [B.2](#), [B.3.1](#), [B.3.1](#)
- [108] NAKAO, M., WATANABE, T., KURODA, T., AND YOSHIHARA, H. Interactive 3D Region Extraction of Volume Data Using Deformable Boundary Object. In *Medicine Meets Virtual Reality 13: The Magical Next Becomes the Medical Now*. IOS Press, 2005, pp. 349–359. [B.2](#)
- [109] NEALEN, A., MÜLLER, M., KEISER, R., BOXERMAN, E., AND CARLSON, M. Physically Based Deformable Models in Computer Graphics. *Computer Graphics Forum* 25, 4 (2006), 809–836. [B.2](#)

- [110] NELSON, T. R., AND ELVINS, T. T. Visualization of 3D ultrasound data. *IEEE Computer Graphics and Applications* 13, 6 (1993), 50–57. [C.6](#)
- [111] NELSON, T. R., PRETORIUS, D. H., HULL, A., RICCABONA, M., SKLAN-SKY, M. S., AND JAMES, G. Sources and impact of artifacts on clinical three-dimensional ultrasound imaging. *Ultrasound In Obstetrics & Gynecology* 16, 4 (2000), 374–383. [2.1](#)
- [112] NIKAS, D. C., HARTOV, A., LUNN, K., RICK, K., PAULSEN, K., AND ROBERTS, D. W. Coregistered intraoperative ultrasonography in resection of malignant glioma. *Neurosurgical focus* 14, 2 (2003), e6. [C.5](#)
- [113] NOBLE, J. A., AND BOUKERROUI, D. Ultrasound image segmentation: a survey. *IEEE Transactions on Medical Imaging* 25, 8 (2006), 987–1010. [2.2](#), [C.4](#)
- [114] NOTHDURFT, H.-C. The role of features in preattentive vision: Comparison of orientation, motion and color cues. *Vision Research* 33, 14 (1993), 1937 – 1958. [D.1](#)
- [115] NOWATSCHIN, S., MARKERT, M., WEBER, S., AND LUETH, T. C. A system for analyzing intraoperative B-Mode ultrasound scans of the liver. In *Proceedings of IEEE Engineering in Medicine and Biology Society* (2007), pp. 1346–1349. [2.2](#), [3.1](#), [A.2](#)
- [116] NUMMINEN, K., SIPILÄ, O., AND MÄKISALO, H. Preoperative hepatic 3D models: Virtual liver resection using three-dimensional imaging technique. *European Journal of Radiology* 56, 2 (2005), 179 – 184. *New Frontiers in Diagnostic Imaging*. [A.2](#)
- [117] ØDEGAARD, S., GILJA, O. H., AND GREGERSEN, H. *Basic and New Aspects of Gastrointestinal Ultrasonography*. Advanced Series in Biomechanics. World Scientific, 2005. [C.1](#)
- [118] ØDEGAARD, S., NESJE, L. B., AND GILJA, O. H. *Atlas of Endoscopic Ultrasonography*. Fagbokforlaget, 2007. [C.1](#)
- [119] OELTZE, S., AND PREIM., B. Visualization of vasculature with convolution surfaces: Method, validation and evaluation. *IEEE Transactions on Medical Imaging* 24(4) (2005), 540–548. [A.2](#)
- [120] ØYE, O. K., ULVANG, D. M., GILJA, O. H., HAUSER, H., AND VIOLA, I. Illustrative couinaud segmentation for ultrasound liver examinations. In *Smart Graphics*, vol. 6815 of *Lecture Notes in Computer Science*. 2011, pp. 60–77. [4](#), [4.1](#)
- [121] ØYE, O. K., WEIN, W., ULVANG, D. M., MATRE, K., AND VIOLA, I. Real time image-based tracking of 4d ultrasound data. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention* (2012), pp. 447–454. [2.2](#), [3.2](#), [3.6](#), [4.2](#), [E.2](#), [E.3](#), [E.3.1](#)

- [122] PAPPENBERG, N., LANGE, T., MODERSITZKI, J., SCHLAG, P. M., AND FISCHER, B. Image Registration for CT and Intra-Operative Ultrasound Data of the Liver. In *SPIE Medical Imaging* (2008), vol. 6918, pp. 691–808. [2.2](#), [C.5](#), [E.3.1](#)
- [123] PENNEY, G., BLACKALL, J., HAMADY, M., SABHARWAL, Y., ADAM, A., AND HAWKES, D. Registration of freehand 3D ultrasound and magnetic resonance liver images. *Medical Image Analysis* 8, 1 (2004), 81–91. [2.2](#), [A.2](#), [C.5](#)
- [124] PENNEY, G. P., BARRATT, D. C., CHAN, C. S. K., SLOMCZYKOWSKI, M., CARTER, T. J., EDWARDS, P. J., AND HAWKES, D. J. Cadaver validation of intensity-based ultrasound to CT registration. *Medical Image Analysis* 10, 3 (2006), 385–95. [2.2](#), [C.5](#)
- [125] PETERSCH, B., HADWIGER, M., HAUSER, H., AND HÖNIGMANN, D. Real time computation and temporal coherence of opacity transfer functions for direct volume rendering of ultrasound data. *Computerized Medical Imaging and Graphics* 29, 1 (2005), 53–63. [C.6.1](#)
- [126] PETERSCH, B., AND HÖNIGMANN, D. Blood flow in its context: combining 3D B-mode and color Doppler ultrasonic data. *Transactions on Visualization and Computer Graphics* 13, 4 (2007), 748–57. [2.2](#), [4.2](#), [C.6.2](#), [E.2](#)
- [127] PETERSCH, B., SERRANO-SERRANO, O., AND HÖNIGMANN, D. 3D soft segmentation and visualization of medical data based on nonlinear diffusion and distance functions. In *Transactions on Visualization and Computer Graphics* (2006). [2.2](#), [3.1](#), [3.2](#), [3.2](#), [A.3.1](#), [C.4](#), [E.2](#), [E.3.1](#)
- [128] POSTEMA, M., AND GILJA, O. H. Ultrasound-directed drug delivery. *Current Pharmaceutical Biotechnology* 8 (2007), 355–361. [C.1](#), [E.1](#)
- [129] POSTEMA, M., AND GILJA, O. H. Contrast-enhanced and targeted ultrasound. *World J Gastroenterol* 17 (2011), 28–41. [E.1](#)
- [130] POTCHEN, J. E., HAACKE, M. E., SIEBERT, J. E., AND GOTTSCHALK, A. *Magnetic Resonance Angiography - Concepts & Applications*. Mosby, 1993. [E.2](#)
- [131] PRAGER, R., GEE, A., TREECE, G., AND BERMAN, L. Freehand 3D ultrasound without voxels: volume measurement and visualisation using the Stradx system. *Ultrasonics* 40, 1-8 (2002), 109–15. [C.3](#), [C.3.1](#)
- [132] PRAGER, R. W., GEE, A., AND BERMAN, L. Stradx: real-time acquisition and visualization of freehand three-dimensional ultrasound. *Medical Image Analysis* 3, 2 (1999), 129 – 140. [C.3.1](#)
- [133] REICHARDT, W. Autocorrelation, A Principle for the Evaluation of Sensory Information by the Central Nervous System. In *Principles of Sensory Communications*. John Wiley, 1961, pp. 303–317. [D.2](#)

- [134] REITINGER, B., BORNIK, A., BEICHEL, R., AND SCHMALSTIEG, D. Liver surgery planning using virtual reality. *IEEE Computer Graphics and Applications* 26, 6 (2006), 36–47. [1](#), [2.2](#), [A.2](#)
- [135] RICHARDS, M. S., KRIPFGANS, O. D., RUBIN, J. M., HALL, A. L., AND BRIAN, J. Mean volume flow estimation in pulsatile flow. *Ultrasound in Medicine & Biology* 35, 11 (2009), 1880–1891. [E.2](#)
- [136] ROHLING, R., GEE, A., AND BERMAN, L. A comparison of freehand three-dimensional ultrasound reconstruction techniques. *Medical Image Analysis* 3, 4 (1999), 339–359. [C.3.1](#)
- [137] ROHLING, R., GEE, A., BERMAN, L., AND TREECE, G. Radial basis function interpolation for 3D ultrasound. *Lecture Notes in Computer Science 1613* (1999), 478–483. [C.3.1](#)
- [138] ROPINSKI, T., DÖRING, C., AND REZK-SALAMA, C. Interactive volumetric lighting simulating scattering and shadowing. In *Proceedings of IEEE Pacific Visualization* (2010), pp. 169–176. [C.6.3](#), [C.7](#), [C.8b](#)
- [139] ROXBOROUGH, T., AND NIELSON, G. M. Tetrahedron based, least squares, progressive volume models with application to freehand ultrasound data. In *Proceedings of IEEE Visualization* (2000), pp. 93–100. [C.3.1](#)
- [140] RUSHMER, R., BAKER, D., JOHNSON, W., AND STRANDNESS, D. Clinical applications of a transcutaneous ultrasonic flow detector. *JAMA* 199, 5 (1967), 326–328. [1](#)
- [141] SAKAS, G., SCHREYER, L., AND GRIMM, M. Preprocessing and volume rendering of 3D ultrasonic data. *IEEE Computer Graphics and Applications* 15, 4 (1995), 47–54. [2.2](#), [C.3.2](#), [C.4.1](#)
- [142] SATO, Y., NAKAMOTO, M., TAMAKI, Y., SASAMA, T., SAKITA, I., NAKAJIMA, Y., MONDEN, M., AND TAMURA, S. Image guidance of breast cancer surgery using 3-D ultrasound images and augmented reality visualization. *IEEE Transactions on Medical Imaging* 17, 5 (1998), 681–93. [C.7](#)
- [143] SCHOTT, M., PEGORARO, V., HANSEN, C., BOULANGER, K., AND BOUA-TOUCH, K. A directional occlusion shading model for interactive direct volume rendering. In *Proceedings of EUROGRAPHICS* (2009), pp. 855–862. [C.6.3](#), [C.8e](#)
- [144] SCHULTZ, T., SAUBER, N., ANWANDER, A., THEISEL, H., AND SEIDEL, H.-P. Virtual Klingler Dissection: Putting Fibers into Context. *Computer Graphics Forum* 27, 3 (2008), 1063–1070. [B.2](#)
- [145] SELLE, D., PREIM, B., SCHENK, A., AND PEITGEN, H. Analysis of vasculature for liver surgical planning. *IEEE Transactions on Medical Imaging* (2002). [A.2](#)

- [146] SHANG, Q., CLEMENTS, L., GALLOWAY, R. L., CHAPMAN, W. C., AND DAWANT, B. M. Adaptive directional region growing segmentation of the hepatic vasculature. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (2008), vol. 6914 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*. A.2
- [147] SHANKAR, P. M. Speckle reduction in ultrasonic images through a maximum likelihood based adaptive filter. *Physics in Medicine and Biology* 51, 21 (2006), 5591–5602. C.3.2
- [148] SHELTON, D., STETTEN, G. D., AND CHANG, W. Ultrasound visualization with the sonic flashlight. *Proceedings of ACM SIGGRAPH* (2002), 82–82. C.7
- [149] SIELHORST, T., FEUERSTEIN, M., AND NAVAB, N. Advanced medical displays: A literature review of augmented reality. *Journal of Display Technology* 4, 4 (December 2008), 451–467. C.7
- [150] SOLBERG, O. V., LINDSETH, F., TORP, H., BLAKE, R. E., AND NAGELHUS HERNES, T. A. Freehand 3D ultrasound reconstruction algorithms—a review. *Ultrasound in Medicine and Biology* 33, 7 (2007), 991–1009. C.3.1
- [151] SONOWAND. Sonowand - product description. <http://www.sonowand.com/>, 2011. C.5
- [152] STEEN, E., AND OLSTAD, B. Volume rendering of 3D medical ultrasound data using direct feature mapping. *IEEE Transactions on Medical Imaging* 13, 3 (1994), 517–25. 2.2, C.4, C.6
- [153] STEEN, E. N., OLSTAD, B., BERG, S., MYKLEBUST, G., AND SCHIPPER, K. P. Combined visualization of tumor and vessel geometry. *Proceedings of IEEE Ultrasonics Symposium* (1994), 1645–1648. 2.2
- [154] STETTEN, G. D., CHIB, V. S., AND TAMBURRO, R. J. Tomographic reflection to merge ultrasound images with direct vision. *Proceedings of 29th Applied Imagery Pattern Recognition Workshop* (2000), 200–205. C.7
- [155] STONE, L. S., AND THOMPSON, P. Human Speed Perception is Contrast Dependent. *Vision Research* 32, 8 (1992), 1535 – 1549. 4.3, D.2
- [156] TAN, C., AND LIU, D. C. Interactive Editing for 3D Ultrasound Volume Rendering. *Proceedings of 2nd International Conference on Bioinformatics and Biomedical Engineering* (2008), 2307–2310. 2.2
- [157] THOMPSON, P. Perceived Rate of Movement Depends on Contrast. *Vision Research* 22, 3 (1982), 377–380. D.2
- [158] THOMPSON, P., BROOKS, K., AND HAMMETT, S. T. Speed can go up as well as down at low contrast: Implications for models of motion perception. *Vision Research* 46, 6-7 (2006), 782 – 786. 4.3, D.2

- [159] THUNE, N., GILJA, O. H., HAUSKEN, T., AND MATRE, K. A Practical Method for Estimating Enclosed Volumes Using 3D Ultrasound. *European Journal of Ultrasound* 3, 1 (1996), 83 – 92. [C.3.1](#)
- [160] TIETJEN, C., MEYER, B., SCHLECHTWEG, S., PREIM, B., HERTEL, I., AND STRAUSS, G. Enhancing Slice-based Visualizations of Medical Volume Data. In *Proceedings of the Eurographics / IEEE VGTC Symposium on Visualization* (2006), pp. 123–130. [B.2](#)
- [161] VAN PELT, R., BESCOS, J. O., BREEUWER, M., CLOUGH, R. E., GROLLER, M. E., ROMENIJ, B. T. H., AND VILANOVA, A. Exploration of 4D MRI Blood Flow Using Stylistic Visualization. *Transactions on Visualization and Computer Graphics* 6, 16 (2010), 2153–2162. [D.1](#), [E.2](#)
- [162] VAN PELT, R., BESCOS, J. O., BREEUWER, M., CLOUGH, R. E., GROLLER, M. E., ROMENIJ, B. T. H., AND VILANOVA, A. Interactive Virtual Probing of 4D MRI Blood-Flow. *Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2153–2162. [E.2](#)
- [163] VAN PELT, R., JACOBS, S., TER HAAR ROMENY, B., AND VILANOVA, A. Visualization of 4d blood-flow fields by spatiotemporal hierarchical clustering. *Computer Graphics Forum* 31, 3 (2012), 1065–1074. [D.1](#), [E.2](#)
- [164] VERSTRATEN, F. A. On the Ancient History of the Direction of the Motion Aftereffect. *Perception* 25, 10 (1996), 1177–1187. [D.4](#)
- [165] VILANOVA, A., PREIM, B., VAN PELT, R., GASTEIGER, R., NEUGEBAUER, M., AND WISCHGOLL, T. Visual exploration of simulated and measured blood flow. *Corr* (2012), 20. [E.2](#)
- [166] VIOLA, I., FEIXAS, M., SBERT, M., AND GRÖLLER, M. E. Importance-driven focus of attention. *Transactions on Visualization and Computer Graphics* 12, 5 (2006), 933–940. [4](#)
- [167] VIOLA, I., KANITSAR, A., AND GROLLER, M. E. Importance-Driven Volume Rendering. In *Proceedings of IEEE Visualization* (2004), pp. 139–146. [B.2](#)
- [168] VIOLA, I., KANITSAR, A., AND GRÖLLER, M. E. Importance-driven feature enhancement in volume visualization. *Transactions on Visualization and Computer Graphics* 11, 4 (2005), 408–418. [4](#)
- [169] VIOLA, I., NYLUND, K., ØYE, O. K., ULVANG, D. M., AND GILJA, O. H. Illustrated Ultrasound for Multimodal Data Interpretation of Liver Examinations. In *Proceedings of Eurographics Workshop on Visual Computing in Biomedicine* (2008), pp. 125–133. [2.2](#), [4.1](#), [C.6.2](#)
- [170] ŠOLTÉSZOVÁ, V., HELLJESEN, L. E. S., WEIN, W., GILJA, O. H., AND VIOLA, I. Lowest-variance streamlines for filtering of 3d ultrasound. In *Proceedings of Visual Computing for Biomedicine* (2012), pp. 41–48. [2.2](#), [4.3](#)

- [171] ŠOLTÉSZOVÁ, V., PATEL, D., AND BRUCKNER, S. A multidirectional occlusion shading model for direct volume rendering. *Computer Graphics Forum* 29, 3 (2010), 883–891. [2.2](#), [C.6.3](#), [C.7](#), [C.8c](#)
- [172] ŠOLTÉSZOVÁ, V., PATEL, D., AND VIOLA, I. Chromatic shadows for improved perception. In *Proc. Non-photorealistic Animation and Rendering (NPAR 2011)* (2011), pp. 105–115. [2.2](#)
- [173] ŠOLTÉSZOVÁ, V., TURKAY, C., PRICE, M., AND VIOLA, I. A perceptual-statistics shading model. *IEEE Transaction on Visualization and Computer Graphics* 18, 12 (2012), 2265–2274. [2.2](#)
- [174] WAGNER, R. F., SMITH, S. W., SANDRIK, J. M., AND LOPEZ, H. Statistics of speckle in ultrasound B-scans. *IEEE Transactions on Sonics and Ultrasonics* 30, 3 (1983), 156–163. [2.1](#)
- [175] WAND, B., CHENG, M., HUANG, X., AND JU, Y. A simplified method to segment liver according to couinaud’s classification. In *The 1st International Conference on Bioinformatics and Biomedical Engineering* (2007), pp. 523–526. [4](#), [4.1](#), [A.1](#)
- [176] WANG, S., SUN, Y., AND CHANG, F. Artifact removal and texture-based rendering for visualization of 3D fetal ultrasound images. *Medical & Biological Engineering & Computing* 46, 6 (june 2008), 575–88. [C.6](#)
- [177] WEIN, W., BRUNKE, S., KHAMENE, A., CALLSTROM, M. R., AND NAVAB, N. Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Medical Image Analysis* 12, 5 (2008), 577–85. [2.2](#), [A.2](#), [C.3](#), [C.5](#)
- [178] WEIN, W., PACHE, F., ROEPER, B., AND NAVAB, N. Backward-warping ultrasound reconstruction for improving diagnostic value and registration. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention* (Oct. 2006), Lecture Notes in Computer Science, Springer, pp. 750–757. [C.3.1](#)
- [179] WEIN, W., RÖPER, B., AND NAVAB, N. Automatic registration and fusion of ultrasound with ct for radiotherapy. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention* (2005), J. S. Duncan and G. Gerig, Eds., vol. 3750 of *Lecture Notes in Computer Science*, pp. 303–311. [A.2](#)
- [180] WEISKOPF, D. On the Role of Color in the Perception of Motion in Animated Visualizations. In *Proceedings of IEEE Visualization* (2004), pp. 305–312. [2.2](#), [4.3](#), [D.2](#), [D.3](#), [D.5](#)
- [181] WEISKOPF, D., ENGEL, K., AND ERTL, T. Volume clipping via per-fragment operations in texture-based volume visualization. In *Proceedings of IEEE Visualization* (2002), pp. 93–100. [B.2](#)
- [182] WEISKOPF, D., ENGEL, K., AND ERTL, T. Interactive Clipping techniques for Texture-Based Volume Visualization and Volume Shading. *Transactions on Visualization and Computer Graphics* 9, 3 (2003), 298–312. [B.3.4](#)

- [183] WESKOTT, H. B-flow—a new method for detecting blood flow. *Ultraschall in der Medizin* 21 (2000), 59–65. [1](#)
- [184] WU, B., KLATZKY, R. L., SHELTON, D., AND STETTEN, G. D. Psychophysical evaluation of in-situ ultrasound visualization. *Transactions on Visualization and Computer Graphics* 11, 6 (2005), 684–693. [C.7](#)
- [185] YOO, Y. M., MANAGULI, R., AND KIM, Y. New multi-volume rendering technique for three-dimensional power Doppler imaging. *Ultrasonics* 46, 4 (2007), 313–22. [2.2](#), [C.6.2](#)
- [186] ZHANG, W., NOBLE, J. A., AND BRADY, J. M. Real Time 3-D Ultrasound to MR Cardiovascular Image Registration Using a Phase-Based Approach. In *Proceedings of IEEE International Symposium on Biomedical Imaging: Macro to Nano* (2006), pp. 666–669. [2.2](#), [A.2](#), [C.5](#)