The Ultrasound Visualization Pipeline - A Survey

Å. Birkeland[1], V. Šoltészová[1], D. Hönigmann, O. H. Gilja[1,3], S. Brekke[1], T. Ropinski and I. Viola[1]

[1]University of Bergen, Norway
[2]n22 Research & Technology Transfer, Wiener Neustadt, Austria
[3]National Centre for Ultrasound in Gastroenterology, Haukeland University Hospital, Norway
[4]University of Münster, Germany

Abstract

Ultrasound is the most frequently used imaging modality in medicine. The high spatial resolution, its interactive nature and non-invasiveness makes it the first choice in many examinations. Image interpretation is one of the ultrasounds main challenges. Much training is required to obtain a confident skill level in ultrasound-based diagnostics. State-of-the-art graphics techniques is needed to provide meaningful visualizations of ultrasound in real-time In this paper we present the process-pipeline for ultrasound visualization, including an overview of the tasks performed in the specific steps. To provide an insight into the trends of ultrasound visualization research, we have selected a set of significant publications and divide them into a technique-based taxonomy covering the topics pre-processing, segmentation, registration, rendering and augmented reality. For the different technique types we discuss the difference between ultrasound-based techniques and techniques for other modalities.

1. Introduction

Medical ultrasound has a strong impact on clinical decision making and its high significance in patient management is well established [ØGG05, ØNG07]. Ultrasonography has in comparison with CT, MRI, SPECT and PET scanning very favourable cost, great availability world-wide, high flexibility, and extraordinary patient friendliness. Despite 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 [GHØ*03] and even pocket sized scanners [Hea10], thus expanding the field of applications considerably. 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. Advanced visualization methods have the capacity to transform complex data into graphic representations that enhance the perception and meaning of the data [GHd*07]. 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 [GHWB96, GHH*02, TH98, KGN*, PG07]. Moreover, the 2D images can be aligned to form 3D data sets. Advanced visualization is particularly important when 3D acquisitions are available. Inspection of the 2D slices does not fully transfer the whole information hidden in the images. 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 is 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.

Medical imaging data is indeed becoming more and more complex. Not only is the temporal and spatial resolution ever increasing making the manual slice-by-slice inspection is very time-consuming, but new 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.

1.1. Ultrasound Technology

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.

The entire imaging process is normally performed within an ultrasound scanner, a machine capable of both generating and receiving the acoustic waves, and also for processing and presenting the data for physicians. The signal processing pipeline may vary, but the overall scheme is as shown in Figure 1.

The transmission and reception of the US pulses is performed by a transducer made of piezo-electric material [Ang95]. Modern transducers are made up of a large number of elements, either aligned in a row to perform 2D scanning, or arranged as a matrix to perform 3D scanning. 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.

Ultrasonic waves are attenuated in human tissue, typically by 0.2 - 0.5dB/cm/MHz. Thus the reflected echoes have to be amplified with a factor depending of the traversal depth - a process denoted depth-gain compensation.

The frequency impacts the ultrasound image in two fundamental ways: High frequency gives high spatial resolution and low penetration due to high attenuation. Thus, one would like to use as high imaging frequency as possible without loosing sensitivity due to attenuation. For images of the heart or the deep abdomen, where the maximum scan depth is 10 to 20 cm, a good compromise is in the range of 2-5MHz, while scanning shallow vessels such as carotid, one can use as high frequency as 10-15MHz and thereby get very detailed images. [Ang95] 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. Thus, scanners usually let the examiner choose a focus depth and thereby maximize image sharpness in the main region of interest within the image frame. Compared to competing 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 handled by temporal and spatial noise suppression algorithms. The speckle noise on the other hand, originates from interference between echoes from neighbouring scatterers, and forms a characteristic pattern that persists during several imaging frames [BUR78, WSSL83]. Frequently seen artifacts are reverberations (multiple echoes creating patterns that are misplaced in space), side lobes (acoustic energy in other direction than the focused 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 [NPH*00](wavefront degradation due to non-uniform speed of sound), makes the perceived image quality of ultrasound images “low”.

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: Phased arrays, linear arrays and curved arrays [ATH*95]. 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 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 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 the 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 [MBW95]. Since 2002, US scanners for real-time 3D imaging using fully sampled matrix probes have been commercially available [HCG05]. The matrix probes available so far are of FPA design, i.e, small aperture targeted at transthoracic or transesophageal cardiac imaging. The 3D data are visualized in real-time on the scanner, either as sliced volumes, as direct volume renderings, or both.

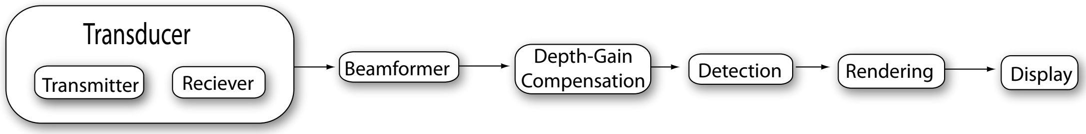

Figure 1: Signal processing steps for 3D ultrasound intensity mapping.

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 pulses. 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 delay giving the maximal echo strength is calculated and yields a corresponding velocity for the scatterer. The ultrasound Doppler principle is frequently used for visualizing and quantifying the velocity field of blood and moving muscles [Jen96].

Techniques which work well for other modalities are being adapted to suit the special characteristic of data generated by sound waves. In this paper we present an overview of the pipeline for advanced visualization specific to ultrasound data.

2. Taxonomy

Techniques for ultrasound visualization can be categorized in a variety of ways, e.g, when they where developed, what types of data modalities is utilized, which anatomy is the technique 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.

Another taxonomic scheme for ultrasound visualization is based on the different types of data the technique utilizes. Different data types require different types of visualization. 3D Freehand and 4D ultrasound, as mentioned above, are very different. 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. For instance, in a liver scan one might look for tumors using a high-resolution gastric 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 difference lead to anatomically specific techniques.

In this survey we have classified around 60 papers with regard to which technique they have been given most weight. In Figure 3 we see the classified papers in a parallel coordinate plot where each axis corresponds to the different taxonomy classification. The second axis (the pipeline axis) is selected as the classification for this survey. Five categories where 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***: Combining ultrasound rendering with natural vision.

In the following sections we present the need for each of

(a) Heart (b) Liver (c) Fetus

Figure 2: Example ultrasound images from the cardiac *(2a), gastric (2b) and fetal (2c) domain.

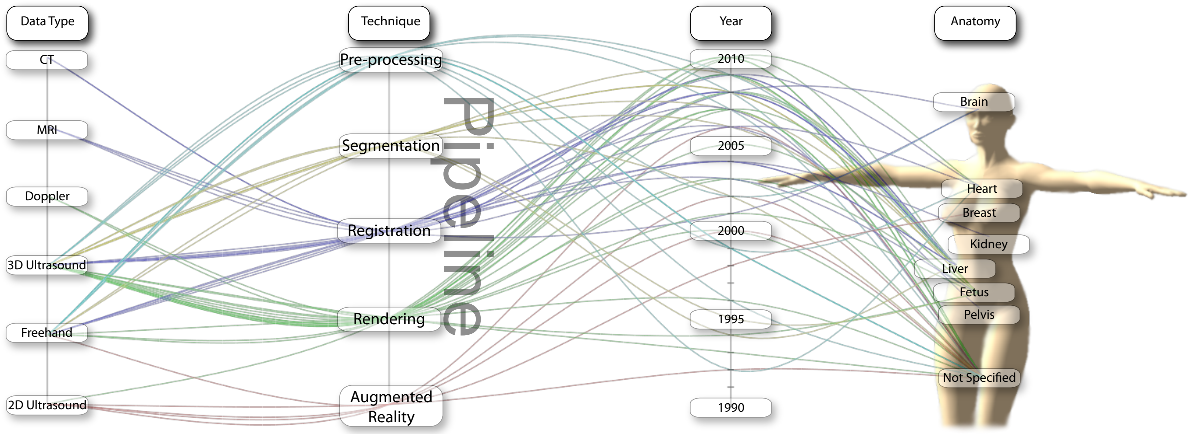

Figure 3: The different classifications shown in a parallel-coordinate plot. The colors depict which technique a publication has given the most weight.

the major topics and recently developed techniques in their respective topic.

3. Pre-processing

3D ultrasound is nowadays often employed in clinical diagnostic imaging. Even though dedicated 3D probes exist, 3D volumes are in many cases acquired by compounding 2D images with spatial interpretation. This technique is called 3D freehand ultrasound. The composition of the compound volume is a two-level process: tracked acquisition and reconstruction. The tracking system delivers spatial position and orientation of the probe which is necessary to place 2D images in space. In order to achieve better results, calibration of the tracking system is necessary as well as correction of pressure-induced artifacts coming from the ultrasound probe onto the skin.

For 3D freehand ultrasound systems, it is necessary to define the position and orientation of the 2D image with respect to the tracking sensor. The problem of calibration is addressed in the work of Wein and Khamene [WK08]. They proposed to make two angular sweeps with approximately perpendicular orientation. The sweep must contain well visible structures. They explained how to use a non-linear optimization strategy to maximize the similarity between two volumes reconstructed from each sweep.

To acquire good quality scans, the clinician must press the probe against the subject's body with some force. The human factor causes that this force is not constant and different deformations of underlying structures in the body occur. Prager et al. applied sequential image correlation techniques [PGTB02]: each 2D image is correlated with the previous and with the next image. They use a rigid translation in x and y directions followed by a non-rigid shift in depth y.

Both 2D and 3D ultrasound is not acquired in Cartesian coordinates but in polar coordinates (ϕ, R) in 2D or (ϕ, ψ, R) in the case of 3D ultrasound. The angles ϕ and ψ correspond to the azimuth and elevation angles of the beam and R is the radius or 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 converted to Cartesian. The *scan-conversion* can be done as a preprocessing step or on-the-fly directly in the rendering stage.

In the first part of this section, we list selected approaches to volume reconstruction from scan-converted 3D freehand ultrasound. Further step in the visualization pipeline of ultrasound volumes is often data enhancement because of its noisy character. The second part of this section is dedicated to data enhancement methods which are tailored for ultrasound volumes.

3.1. Reconstruction

A detailed categorizations of reconstruction algorithms has been done by Rohling et al. [RGB99] and Solberg et al. [SLT*07]. Solberg et al. categorize the algorithms into voxel-based methods, pixel-based methods and function-based methods. We built upon their categorization and complete it by recent works.

Voxel-based methods run through the voxel grid and assign each of them a value estimated by an interpolation method. Stradx [PGB99] is a system that allows for real-time visualization of 3D freehand ultrasound. It supports any plane re-slicing which works essentially as nearest neighbour interpolation. The further development of the Stradx system [PGTB02] includes support for direct volume rendering: they blend images generated by re-slicing as described in their previous work. Gee et al. described a pipeline of processing and visualization for 3D ultrasound [GPT*04]. Their system also includes plane re-slicing with nearest neighbour interpolation of the original data. The reconstructed plane is intended for direct viewing - implying only one re-sampling stage. Linear, bilinear and trilinear interpolation methods have been also used [TGHM96, BTM*99].

Karamalis et al. presented a technique of high-quality volume reconstruction partly on the GPU [KWKN09]. They define the most optimal sampling direction from the spatial position of 2D images and sampling layers which are orthogonal to the chosen sampling direction. Each sampling layer is reconstructed as follows: For each subsequent pair of 2D images, two intersection lines with the current reconstruction layer are found. These lines connect into a quadrilateral and graphics hardware interpolates their grey-values over the quadrilateral. The visualization pipeline includes two resampling steps: one during the reconstruction and one while volume rendering.

Pixel-based methods traverse each pixel of all acquired 2D images and update the value of one or several voxels of the target grid. Gobbi and Peters described a technique where the 3D reconstruction happens in real-time while the data is captured [GP02]. They used splatting as a high-quality interpolation method.

Function-based methods employ a specific function to interpolate between voxels. In most applications, the 3D volume is reconstructed on a regular grid and the techniques are not considering the shape of the underlying data. Rohling et al. investigated the quality of interpolation using splines, which is a polynomial function [RGBT99]. They compared this technique with other standard methods and showed that it produces more accurate reconstructions without visible artifacts.

Tetrahedron-based methods, described in the work of Roxborough and Nielson [RN00], do not reconstruct a volumetric grid but a 3D model built from tetrahedra. The model is created by iterative subdivision of an initial tetrahedra. The subdivision terminates, if the tetrahedra contains 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, meaning the model adapts as new data is streamed in. The complexity of the model conforms the complexity of the data.

We listed and described selected algorithms in categories based on how they are implemented. If choosing a specific algorithm, one must choose between speed and quality. The work of Solberg et al. compares the performance of some of the algorithms on specified hardware. From all listed methods, the pixel nearest neighbor real-time reconstruction described by Gobbi and Peters [GP02]. The radial-based function reconstruction by Rohling et al. [RGBT99] delivers reconstructions of the best quality but at are also the most computationally costly.

3.2. Data Enhancement

Ultrasonic volumes are usually not rendered directly without preprocessing because of their natural properties [SSG95] such as low dynamic range, noisiness and speckle. Also, the geometric resolution varies with depth and fuzzy tissue boundaries which can be several pixels wide and varies with its orientation. Tissue boundaries can even disappear if they are parallel to the ultrasonic beam. These properties pose a challenge to 3D visualization techniques.

In contrast, clinicians prefer to view original 2D images and avoid any filtering and enhancement. E.g., speckle patterns have value in 2D because it refers to the texture of the tissue boundary. However, speckle in 3D brings no added value to the visualization and is considered as artifact same as noise. Therefore, prior to the rendering stage, the 3D data is filtered to enhance its quality. We present speckle reduction techniques separately and followed by dedicated preprocessing for specific applications.

Speckle reduction. For an early review on speckle reduction techniques, see the survey of Forsberg et al. [FHLJ91]. Early work of Belohlavek et al. [BDG*92] uses the *eight hull* algorithm with a geometric filter described by Crimmins [Cri85]. More recently published techniques are based on region growing [CYFB03], adaptive filtering [Sha06], compression techniques [GSP05] and anisotropic diffusion filters [KWKV07].

Applicatio-dedicated enhancement. Generally, systems employ a blending of image-processing techniques to enhance the quality of the data. The work of Sakas et al. [SSG95] lists techniques with a good trade-off between loss of information and visualization quality. They employed Gaussian filters for noise and speckle reduction and speckle removal for contour smoothing and median filters for gap closing and noise reduction. Median filters deliver results of higher quality: they remove small surface artifacts while preserving the sharpness of boundaries. However their evaluation is more costly. Lizzi and Feleppa [LF00] described innovations in ultrasound technology and techniques for image enhancement. They describe a technique to increase the axial resolution by processing the signal in the frequency domain before an image is formed. This resolution gain is especially valuable in opthalmology when visualizing thin layers within the cornea.

4. Segmentation

The amount of data acquired in various research fields is substantially growing. Data size has long overgrown the amount of information we can display on a single 2D screen at once. Selecting interesting features to be visualized is important to be able to root out the occluding elements.

For most modalities, segmentation can be performed by extracting regions with similar data values. For instance, the data values in a CT scan are normalized into Hounsfield units which provide a good basis for binary thresholding techniques. Ultrasound is a non-linear imaging system. This means that the system contains components which are not linear, such as logarithmic compression. The dynamic contrast and logarithmic compression change the image depending on which structures that are visible. A high intensities in hard tissue, such as bone, will reduce intensities in soft tissue regions. Non-linearity makes segmentation a challenging task and early work indicated that binary thresholding techniques are not very suited for ultrasound data [SO94]. More sophisticated techniques are required for satisfactory segmentation. An extensive survey on ultrasound image was been presented by Noble and Boukerroui [NB06] in 2006. In this section we have focused mostly resent publications aimed for advanced visualizations.

Due to the inconsistency in the data values in ultrasound and the difference specific speckle patterns, segmentation techniques are often tailored for specific anatomical structures. Steen et al. [SOB*94] presented a graph-based approach for extracting vessels from 3D ultrasound data. In this case, the volume is considered a directional graph where each voxel is connected to neighbouring voxels. The weight for each edge is the difference between the user selected intensity and the current intensity. Finally, the vessels are extracted with a region growing algorithm applied to the graph. While blood vessels have a clear characteristic in ultrasound images, other work deals with anatomy which is not as clearly defined in the ultrasound image. Carneiro et al. have developed an automatic technique for segmenting the brain of a fetus [CAG*08]. 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 [DID*09], as shown in figure 4. Feng et al. show that facial recognition can be achieved in 3D ultrasound [FZGC09]. First, the technique searches the 3D dataset a set of possible sub-volumes, where the face could be located. Then a 2D profile detection algorithm is applied to find the correct candidate. Segmentation techniques have also been developed which are not anatomically specific. Zhang et al. can extract a 3D point cloud from large 3D ultrasound dataset [ZRP02]. The system looks for clearly visible surfaces which face the current ultrasound-probe viewpoint. This work shows promising result for the water-tank test set-up, but has not been tested on soft tissue. 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 [PSSH06]. This technique defines a probability map for 3D ultrasound data, which in turn can be used to create *soft* representations of the features extracted, as shown in Figure 4.

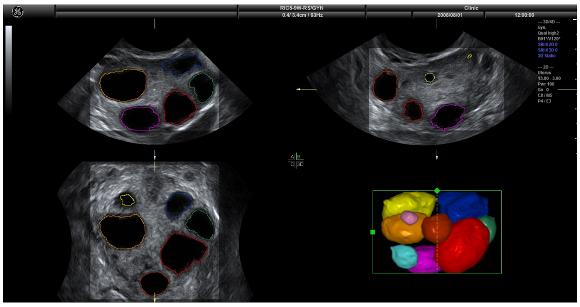

Figure 4: Soft segmentation of ultrasound data.

Segmentation is often a post-acquisition process, but data examination is mostly done live. 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 [NMWL07]. The vessel cross-section is estimated as an ellipse and mapped into a 3D model. Recently, Birkeland and Viola presented a technique for extracting blood vessels live during examination [BV10]. This provides a better spatial understanding of the spatial extent of the vessel tree. Based on Petersch et al.'s method for hypo-echoic region extraction, the technique extracts vessel-regions in 3D freehand ultrasound and maps the outline of the blood vessels into a geometrical 3D model.

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 visualization system [SSG95] 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 axial 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 the 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 5 show and example of clipping implemented in Voluson machines. Bruno et al. presented a framework for selective volume rendering of 3D ultrasound data [BRN00]. The selected volume is defined from user-

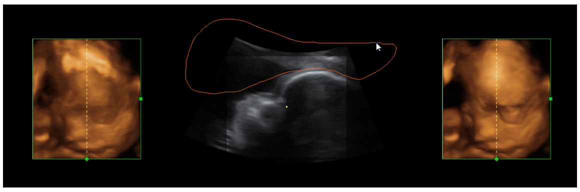

Figure 5: Using MagiCut to Clip the volume generating a clear view to the desired structure [99].

defined sketches on three orthogonal planes. A surface is then generated mathematically and can either be rendered directly or used as a clipping boundary. Another system by Tan and Liu [TL08] also incorporate clipping for editing the visual representation. The paper presents two ways of interactively selecting regions in 3D ultrasound which are to be excluded from the rendering. First method is based on shear-warp, where the user selects a region on the rendered scene, and the projection of the region is removed from the final image. The second method starts with the user selecting a region on a 2D slice of the volume, the system then performs a 1D border detection to define the region which is to be removed. The process then repeats itself slice-by-slice using the previous region as a basis.

5. Registration

Combining both the live image of ultrasound and images extracted previously from 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 pre-operative imaging techniques, ultrasound can be performed live during surgery. The live feed provides the surgeon with a direct view of underlying structures while operating on the patient.

Nikas et al. published an evaluation of the application of co-registered 2D ultrasound and MRI for intra-operative navigation [NHL*03]. 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. [FJKF05].

Registration is also used in simulation and training. Ultrasound guided biopsy is a very challenging task which can take a long time to master. To prevent potential practice on live patients, Dong Ni et al. developed a system for simulating ultrasound guided biopsy [NCQ*09]. First they create a panorama ultrasound dataset by stitching several volumes together. A CT scan is co-registered with the ultrasound data. Two haptic devices are utilized to represent the transducer and the biopsy needle. To add realism and provide the feeling of penetrating through tissue, physics simulation of the needle-tissue penetration is calculated and a feedback is returned through the haptic device. Adding to reality, respiration and deformation are estimated and included into the 3D model.

5.1. Rigid Registration

Rigid registration can be used to quickly get a registration between two modalities and is suitable for rigid anatomies such as the skull. Much work has been done within different modalities, such as CT–X-ray registration [MTLP10]. The imaging system for the modalities discussed in this paper are based on the same principles. 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. To register CT and ultrasound scans of the kidney [LMPT04], Leroy et al. applied a gradient-preserving speckle filter to a gradient similarity evaluation. To avoid unnecessary calculations *shadows* from bone were detected and discarded, as such areas do not contribute to the registration.

Penny 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 [PBH*04]. Later Penney et al. extended their technique for CT-ultrasound registration of the pelvis and femur [PBC*06]. 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. [ZNB].

A method for image based re-registration was developed and evaluated by Ji et al. [JWH*08]. This method is based on a registration using fiducial markers as a starting point and is then registered again in the operating room to compensate for brain shift which can occur after the initial registration. The re-registration technique is based on the normalized mutual information. Combining segmentation with registration, King et al. presented a technique for registering pre-segmented models with ultrasound [KMY*09]. The technique predicts the probability that the ultrasound image was produced by the segmented anatomy.

5.2. Non-Rigid Registration

In addition to a rigid transformation, affine registration includes non-uniform scaling which sometimes need to be applied in order to get a more correct registration. Wein et al. developed an automatic affine-registration technique between CT and ultrasound [WK08]. 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 human body is in general not rigid and can be *deformed* by for instance external pressure or different laying positions of the patient when acquiring the images. To account for local deformations while imaging soft tissue, a more complex registration is required. Blood vessels in the liver can be segmented from different modalities and a vascular tree can be used to create a non-rigid registration as shown by Lange et al. [LEH*04]. First a rigid registrations is calculated, then a the rigid transformation is replaced by a B-spline transformation. Papenberg et al. proposed two approaches for CT ultrasound registration [PLM*08] given a set of paired landmarks in both the CT and ultrasound dataset. 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. [LPH*08].

6. Rendering

Visual depiction of the data is the last step between the data and the user. All the pre-processing, segmentation and registration in the world will be in vain if there is not a method for presenting data to the user. The basic 2D B–mode ultrasound is trivially presented on a 2D screen. Doppler information can be included as well with color according to the blood flow. Other data, such as tissue strain can also be included into 2D. In order to minimize the damage done by certain tumor treatments, the doctors need to account for the motion of the tumor caused by the patients breathing. Lee et al. presented a technique called *CycleStack Plot* [LCPS10]. This technique superimposes the respiratory signal onto a selected feature of interest in the ultrasound image.

3D Freehand ultrasound provide spatial information by including the position of the transducer. Multiple 2D images can be put into the same 3D context and a 3D representation can be generated on-the-fly. Garrett et al. presented a technique for correct visibility ordering of the 2D images using a binary positioning tree [GFWS96]. In Section 3.1 we discussed how 3D freehand ultrasound can be used to create large volumes. Visualization of large volumes easily leads to visual clutter. To limit clutter, Gee et al. extended existing reslicing tools to create narrow volumes [GPTB02]. These volumes contain less elements and are easier to present. However, creating volumes from slices can cause reconstruction artifacts.

3D ultrasound is not as trivial to present. In their early work, Nelson and Elvis discussed the effect of existing techniques for presenting 3D ultrasound data [NE93], such as surface fitting and volume rendering. Later, seven different volume projection techniques applied to 3D ultrasound were evaluated by Steen and Olstad [SO94]. The various techniques included maximum intensity projection (MIP), average intensity projection (AIP) and gradient magnitude projection (GMP). In the study, the techniques were applied to 3D fetal data, where GMP was valued to give the best results regarding detail and robustness towards viewing parameters.

Volume rendering in itself is not trivial when dealing with live 3D data defined in polar coordinate system. Kuo et al. presented a technique for quick scan conversion during rendering [KBCvR07]. The main issue was that the required computing of $tan(\phi)$ and $tan(\psi)$ is very costly. To reduce computation time, the tangens-values are 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 boundaries makes 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 [FL01]. 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 coloured 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 [WSC08]. To remove 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 [LKS08]. This technique employs different sized filters to smooth out the noise.

6.1. Transfer function design

For direct volume rendering, *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

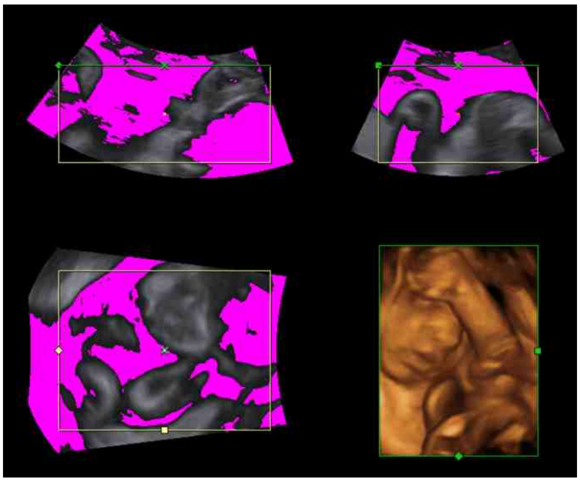

Figure 6: The parameter I_{thresh} determines which echo intensity values to render transparent. A user control with immediate feedback, indicating transparent regions in pink, is essential.

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_{tresh} 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 Fig.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 [HRH03] for an evaluation of these effects on linear and parabolic OTFs. On the other hand, any modification of fundamental acquisition parameters, such as, e.g., 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 optimum visualization. Obviously, 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. [KD98] and subsequent work by Kniss et al. [KKH02], 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 [JSR*09], 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 [HRH03]. 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 [PHHH05].

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

6.2. Multi-Modal Rendering

Multi-modal rendering is meant to bring two or more datasets 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 4-D Doppler data [JSR03]. 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 [PH07]. 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 introduce by Yoo et al. [YMK07]. Instead of blending two 2D rendered images (post fusion), or a blending the two volumes while rendering (composite fusion), it 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 compensate for this by using an if-clause to adjust the alpha-out value in the ray-caster to either the Doppler-signal or the B–mode-signal.

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 [VNOy*08], 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 in the visible. To improve ultrasound video analysis, Angelelli et al. used a degree-of-interest (DOI) function superimposed on the image [AVN*10]. The video sequence was presented as a graph, where the height was defined by the amount the current ultrasound image covered the DOI-function.

6.3. Shading and Illumination

Light is an indispensable part of scenes we see in real life. Therefore also in computer graphics, light sources and light transport models have to be taken into account, when realistically rendering 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 [Lev88]. While local illumination models already reveal structures, global illumination methods result in a more real and plastic look, 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 [RDR10, vPB10]. Ropinski et al. described a volumetric lighting model which simulates scattering and shadowing [RDR10]. 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 7a). 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 simulation of light scattering in volumes [vPB10]. 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 7b). 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 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 8 shows examples for the already addressed shading techniques, the second row shows three alternative approaches. Figure 8d incorporates scattering of light in volume data, as proposed by Kniss et al. [KPHE02]. 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 8e shows the application of the directional occlusion shading technique [SPH*09], which is similar to the technique presented by Šoltészová et al. However, it constrains the light source position to coincide with the view point. Finally, Figure 8f shows the application of a technique based on spherical harmonic lighting [LR10].

Advanced illumination techniques are not yet implemented in the workstations. Current workstations use color coding based on depth. Deeper tissues are colored with cold tones such as blue while close regions have red and orange tones. This effect has been firstly described by Einthoven [Ein85] and is also referred to as chromostereopsis [AR81]. Figure 9 shows a chromatic depth-encoding rendering of a 3D human heart in a modern ultrasound workstation.

7. Ultrasound and Augmented Reality

Ultrasound is commonly viewed on a separate monitor. Therefore it is difficult for non-expert users, such as physicians in training, to comprehend the spatial relationship between what you see on the monitor and where it is located in the patients body. Augmented reality can aid the user by for

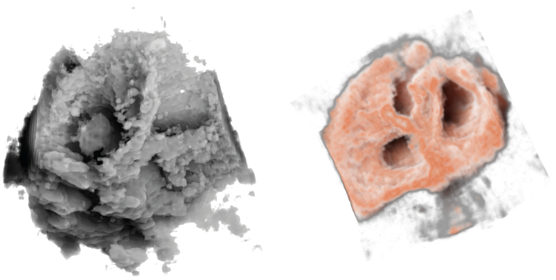

(a) (b)

Figure 7: Rendering of 3D ultrasound of human heart with shadowing: (a) from the work of Ropinski et al. [RDR10] and (b) rendered using the technique presented in the work of Šoltészová et al. [vPB10].

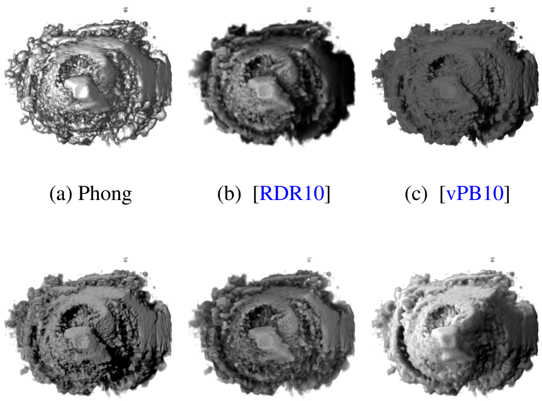

(a) Phong (b) [RDR10] (c) [vPB10]

(d) [KPHE02] (e) [SPH*09] (f) [LR10]

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

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) [BFO92]. 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 [SNT*98]. Combining tracked 2D ultrasound and camera, to project a correct view of a segmented tumor. The tumor is segmented using a minimal intensity projection based selection of the volume of interest. In the

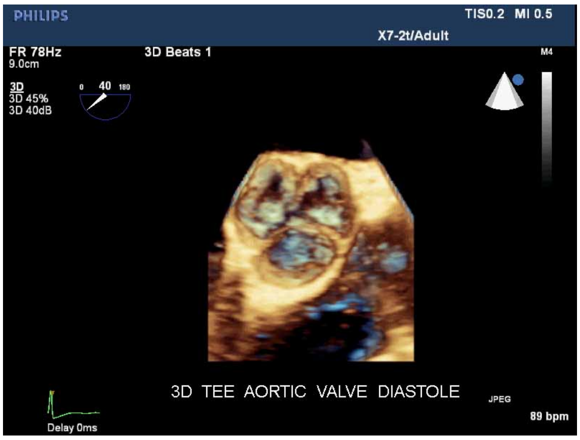

Figure 9: Diastole of the aortic valve on a modern ultrasound workstation using color-coding based on depth.

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 [SCT00]. 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 [SSC02]. The tomographic reflection showed to increase the localization perception compared to conventional ultrasound [WKSS05].

Scaled teleoperation is a technique which is increasingly in use in medicine. The idea is to utilize direct hand-eye coordination when interacting with a remote environment on a different scale. Clanton et al. present a technique for rendering scaled ultrasound combined with a proxy tool [CWC*06]. The master proxy tool is connected to a slave tool scaled down to the original ultrasound scale. The display method is based on the tomographic reflection. The image is shown in a flat panel reflected on a half-silvered mirror to combine the ultrasound with the *patient* underneath.

8. Summary and Discussion

Medical ultrasound data is very different compared to other medical imaging modalities. Techniques for the different steps in the visualization pipeline is specially tailored to suit the different nature of the data acquisition. Techniques meant for in-vivo use have strong performance requirements to handle the high frame rate of ultrasound images. Yet there is a dire need for techniques to improve communication, training etc. Research in advanced techniques focus greatly on 3D ultrasound, but the trend in diagnostics is mostly 2D due to higher frame-rate and resolution. The temporal and spatial resolution for ultrasound is approaching the physical limit. The responsibility now lies on visualization techniques to take it further, combining high resolution 2D and contextual 3D ultrasound.

References

[99] , , : 3d ultrasound: A dedicated system. *European Radiology 9* (1999), S331–S333. 10.1007/PL00014067. 7

[Ang95] ANGELSEN B. A. J.: *Waves, signals and signal processing in medical ultrasonics*, vol. I. Dept. of Physiology and Biomediacl Engineering, Norwegian University of Science and Technology, Trondheim, Norway, 1995. 2

[AR81] ALLEN R. C., RUBIN M. L.: Chromostereopsis. *Survey of Ophthalmology 26*, 1 (1981), 22–27. 10

[ATH*95] ANGLESEN B. A. J., TORP H., HOLM S., KRISTOFFERSEN K., WHITTINGHAM T. A.: Which transducer array is best? *European Journal of Ultrasound 2*, 2 (April 1995), 151–164. 2

[AVN*10] ANGELELLI P., VIOLA I., NYLUND K., GILJA O. H., HAUSER H.: Guided visualization of ultrasound image sequences. pp. 125–132. 10

[BDG*92] BELOHLAVEK M., DUTT V., GREENLEAF J., FOLEY D., GERBER T., SEWARD J.: Multidimensional ultrasonic visualization in cardiology. vol. 2, pp. 1137 – 1145. 5

[BFO92] BAJURA M., FUCHS H., OHBUCHI R.: Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Proceedings of ACM SIGGRAPH 26*, 2 (1992), 210. 11

[BRN00] BRUNO R., RICHARD B., NICOLAS J.-M.: AN INTERACTIVE TOOL TO VISUALIZE THREE-DIMENSIONAL ULTRASOUND DATA. *MMW Fortschritte der Medizin 149*, 15 (april 2000), 56. 6

[BTM*99] BERG S., TORP H., MARTENS D., STEEN E., SAMSTAD S., HØIVIK I., OLSTAD B.: Dynamic three-dimensional freehand echocardiography using raw digital ultrasound data. *Ultrasound in Medicine and Biology 25*, 5 (1999), 745 – 753. 5

[BUR78] BURCKHARDT C.: Speckle IN Ultrasound B-mode Scans. *Transactions on Sonics and Ultrasonics 25*, 1 (1978), 1–6. 2

[BV10] BIRKELAND Å., VIOLA I.: Ultrasound painting of vascular tree. In *Proceedings of Vision, Modeling, and Visualization* (2010). 6

[CAG*08] CARNEIRO G., AMAT F., GEORGESCU B., GOOD S., 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* (juni 2008), 1–8. 6

[Cri85] CRIMMINS T. R.: The geometric filter for speckle reduction. *Applied Optics 24*, 10 (1985), 1438 – 1443. 5

[CWC*06] CLANTON S. T., WANG D. C., CHIB V. S., MATSUOKA Y., STETTEN G. D.: Optical merger of direct vision with virtual images for scaled teleoperation. *Transactions on Visualization and Computer Graphics 12*, 2 (2006), 277–85. 11

[CYFB03] CHEN Y., YIN R., FLYNN P., BROSCHAT S.: Aggressive region growing for speckle reduction in ultrasound images. *Pattern Recognition Letters 24*, 4-5 (2003), 677–691. 5

[DID*09] DEUTCH T. D., I. J., D.O. M., S. O., S. B., D. H., A. A.: Automated assessment of ovarian follicles using a novel three-dimensional ultrasound software. *Fertility and Sterility 92*, 5 (2009), 1562–1568. 6

[Ein85] EINTHOVEN W.: Stereoscopie durch farbendifferenz. *Von Graefes Archiv 19* (1885), 211–238. 10

[FHLJ91] FORSBERG F., HEALEY A. J., LEEMAN S., JENSEN J. A.: Assessment of hybrid speckle reduction algorithms. *Physics in Medicine and Biology 36*, 11 (1991), 1539–1549. 5

[FJKF05] FULLER D. B., JIN H., KOZIOL J. A., 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. 7

[FL01] FATTAL R., LISCHINSKI D.: Variational classification for visualization of 3D ultrasound data. *Proceedings of IEEE Visualization* (2001), 403–410. 8

[FZGC09] FENG S., ZHOU S., GOOD S., COMANICIU D.: Automatic fetal face detection from ultrasound volumes via learning 3D and 2D information. *Proceedings of Computer Vision and Pattern Recognition* (juni 2009), 2488–2495. 6

[GFWS96] GARRETT W., FUCHS H., WHITTON M., STATE A.: Real-time incremental visualization of dynamic ultrasound volumes using parallel BSP trees. *Proceedings of IEEE Visualization* (1996), 235–240,. 8

[GHd*07] GILJA O. H., HATLEBAKK J. G., ØDEGAARD S., BERSTAD A., VIOLA I., GIERTSEN C., HAUSKEN T., GREGERSEN H.: Advanced imaging and visualization in gastrointestinal disorders, 2007. 1

[GHH*02] GILJA O. H., HEIMDAL A., HAUSKEN T., GREGERSEN H., MATRE K., BERSTAD A., ØDEGAARD S.: Strain during gastric contractions can be measured using doppler ultrasonography. *Ultrasound in Medicine and Biology 28*, 11-12 (2002), 1457 – 1465. 1

[GHØ*03] GILJA O. H., HAUSKEN T., ØDEGAARD S., WENDELBO Ø., THIERLEY M.: Mobile ultrasonography in a medical department. *Tidskrift for Norsk Lægeforening 19* (2003), 270–285. 1

[GHWB96] GILJA O., HAUSKEN T., WILHELMSEN I., BERSTAD A.: Impaired accommodation of proximal stomach to a meal in functional dyspepsia. *Digestive Diseases and Sciences 41* (1996), 689–696. 10.1007/BF02213124. 1

[GP02] GOBBI D., PETERS T.: Interactive intra-oeprative 3D ultrasound reconstruction and visualization, 2002. 5

[GPT*04] GEE A., PRAGER R., TREECE G., CASH C., BERMAN L.: Processing and visualizing three-dimensional ultrasound data, 2004. 4

[GPTB02] GEE A., PRAGER R., TREECE G., BERMAN L.: Narrow-band volume rendering for freehand 3D ultrasound. *Computers & Graphics 26*, 3 (juni 2002), 463–476. 8

[GSP05] GUPTA N., SWAMY M. N. S., PLOTKIN E.: Despeckling of medical ultrasound images using data and rate adaptive lossy compression. *Medical Imaging, IEEE Transactions on 24*, 6 (2005), 743–754. 5

[HCG05] HOUCK R. C., COOKE J., 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

[Hea10] HEALTCARE G.: A closer look at ges pocket-sized vscan ultrasound, Sep 13th, 2010. http://www.gereports.com/a-closer-look-at-ges-pocket-sized-vscan-ultrasound/. 1

[HRH03] HÖNIGMANN D., RUISZ J., 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. 9

[Jen96] JENSEN J. A.: *Estimation of Blood Velocities Using Ultrasound. A Signal Processing Approach.*, 1 ed. Cambridge University Press, Cambridge, United Kingdom, 1996. 3

[JSR03] JONES M. G., SHIPLEY J. A., ROBINSON T. M.: Visualisation of 4-D colour and power Doppler data. *Ultrasound in Medicine and Biology 29*, 12 (desember 2003), 1735–1747. 9

[JSR*09] JAN U., SANDKÜHLER D., RAUBERGER M., MATTHIES H. K., OVERHOFF H. M.: Enhanced visualization of ultrasound volumes for diagnostic and therapeutic purposes. In *4th European Conference of the International Federation for Medical and Biological Engineering*, Sloten J., Verdonck P., Nyssen M., Haueisen J., (Eds.), 22 ed., IFMBE Proceedings. Springer Berlin Heidelberg, 2009, pp. 689–692. 9

[JWH*08] JI S., WU Z., HARTOV A., D. P. K., W. R. D.: Mutual-information-based image to patient re-registration using intraoperative ultrasound in image-guided neurosurgery. *Medical Physics 35* (2008). 7

[KBCvR07] KUO J., BREDTHAUER G. R., CASTELLUCCI J. B., VON RAMM O. T.: Interactive volume rendering of real-time three-dimensional ultrasound images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 54*, 2 (februar 2007), 313–8. 8

[KD98] KINDLMANN G., DURKIN J. W.: Semi-automatic generation of transfer functions for direct volume rendering. *Proceedings of IEEE Symposium on Volume Visualization* (1998), 79–86, 170. 9

[KGN*] KUO P., GENTILCORE D., NAIR N., STEVENS J. E., WISHART J. M., LANGE K., GILJA O. H., HAUSKEN T., HOROWITZ M., JONES K. L., RAYNER C. K.: The nitric oxide synthase inhibitor, ng-nitro-l-arginine-methyl-ester, attenuates the delay in gastric emptying induced by hyperglycaemia in healthy humans. 1

[KKH02] KNISS J., KINDLMANN G., HANSEN C.: Multidimensional transfer functions for interactive volume rendering. *Transactions on Visualization and Computer Graphics 8*, 3 (2002), 270–285. 9

[KMY*09] KING A. P., MA Y.-L., YAO C., JANSEN C., RAZAVI R., RHODE K. S., 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* (Berlin, Heidelberg, 2009), Springer-Verlag, pp. 188–201. 7

[KPHE02] KNISS J., PREMOZE S., HANSEN C., EBERT D.: Interactive translucent volume rendering and procedural modeling. pp. 109–116. 10, 11

[KWKN09] KARAMALIS A., WEIN W., KUTTER O., NAVAB N.: Fast hybrid freehand ultrasound volume reconstruction. 5

[KWKV07] KRISSIAN K., WESTIN C. F., KIKINIS R., VOSBURGH K. G.: Oriented speckle reducing anisotropic diffusion. *IEEE Transactions on Image Processing 16*, 5 (2007), 1412–1424. 5

[LCPS10] LEE T.-Y., CHAUDHURI A., PORIKLI F., SHEN H.-W.: Cyclestack: Inferring periodic behavior via temporal sequence visualization in ultrasound video. IEEE, pp. 89–96. 8

[LEH*04] LANGE T., EULENSTEIN S., HUNERBEIN M., LAMECKER H., SCHLAG P.: Augmenting intraoperative 3D ultrasound with preoperative models for navigation in liver surgery. *Proceedings of Medical Image Computing and Computer–Assisted Intervention* (2004), 534–541. 8

[Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications 8*, 3 (1988), 29–37. 10

[LF00] LIZZI F. L., FELEPPA E. J.: Image processing and preprocessing for medical ultrasound. p. 187. 5

[LKS08] LIM S., KWON K., SHIN B.: GPU-based interactive visualization framework for ultrasound datasets. *Computer Animation and Virtual Worlds 20*, 1 (2008), 11–23. 8

[LMPT04] LEROY A., MOZER P., PAYAN Y., 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. 7

[LPH*08] LANGE T., PAPENBERG N., HELDMANN S., MODERSITZKI J., FISCHER B., LAMECKER H., 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. 10.1007/s11548-008-0270-1. 8

[LR10] LINDEMANN F., ROPINSKI T.: Advanced Light Material Interaction for Direct Volume Rendering. In *Proceedings of IEEE/EG International Symposium on Volume Graphics* (2010), pp. 101–108. 10, 11

[MBW95] MERZ E., BAHLMANN F., WEBER G.: VOLUME SCANNING IN THE EVALUATION OF FETAL MALFORMATIONS - A NEW DIMENSION IN PRENATAL-DIAGNOSIS. *Ultrasound In Obstetrics & Gynacology 5*, 4 (APR 1995), 222–227. 2

[MTLP10] MARKELJ P., TOMAŽEVIČ D., LIKAR B., PERNUŠ F.: A review of 3d/2d registration methods for image-guided interventions. *Medical Image Analysis* (April 2010). 7

[NB06] NOBLE J. A., BOUKERROUI D.: Ultrasound image segmentation: a survey. *IEEE transactions on Medical Imaging 25*, 8 (2006), 987–1010. 6

[NCQ*09] NI D., CHAN W. Y., QIN J., CHUI Y.-P., QU I., HO S. S. M., HENG P.-A.: A Virtual Reality Simulator for Ultrasound Guided Organ Biopsy Training. *IEEE Computer Graphics and Applications* (2009), 1–11. 7

[NE93] NELSON T., ELVINS T.: Visualization of 3D ultrasound data. *IEEE Computer Graphics and Applications 13*, 6 (1993), 50–57. 8

[NHL*03] NIKAS D. C., HARTOV A., LUNN K., RICK K., PAULSEN K., ROBERTS D. W.: Coregistered intraoperative ultrasonography in resection of malignant glioma. *Neurosurgical focus 14*, 2 (februar 2003), e6. 7

[NMWL07] NOWATSCHIN S., MARKERT M., WEBER S., LUETH T.: 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. 6

[NPH*00] NELSON T., PRETORIUS D., HULL A., RICCABONA M., SKLANSKY M., JAMES G.: Sources and impact of artifacts on clinical three-dimensional ultrasound imaging. *Ultrasound In Obstetrics & Gynacology 16*, 4 (SEP 2000), 374–383. 2

[ØGG05] ØDEGAARD S., GILJA O. H., GREGERSEN H.: *Basic and New Aspects of Gastrointestinal Ultrasonography*. Advanced Series in Biomechanics. World Scientific, 2005. 1

[ØNG07] ØDEGAARD S., NESJE L. B., GILJA O. H.: *Atlas of Endoscopic Ultrasonography*. Fagbokforlaget, 2007. 1

[PBC*06] PENNEY G. P., BARRATT D. C., CHAN C. S. K., SLOMCZYKOWSKI M., CARTER T. J., EDWARDS P. J., HAWKES D. J.: Cadaver validation of intensity-based ultrasound to CT registration. *Medical Image Analysis 10*, 3 (juni 2006), 385–95. 7

[PBH*04] PENNEY G., BLACKALL J., HAMADY M., SABHARWAL Y., ADAM A., HAWKES D.: Registration of freehand 3D ultrasound and magnetic resonance liver images. *Medical Image Analysis 8*, 1 (2004), 81–91. 7

[PG07] POSTEMA M., GILJA O. H.: Ultrasound-directed drug delivery. *Current Pharmaceutical Biotechnology 8* (2007), 355–361(7). 1

[PGB99] PRAGER R. W., GEE A., BERMAN L.: Stradx: real-time acquisition and visualization of freehand three-dimensional ultrasound. *Medical Image Analysis 3*, 2 (1999), 129 – 140. 4

[PGTB02] PRAGER R., GEE A., TREECE G., BERMAN L.: Freehand 3D ultrasound without voxels: volume measurement and visualisation using the Stradx system. *Ultrasonics 40*, 1-8 (2002), 109–15. 4

[PH07] PETERSCH B., 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. 10

[PHHH05] PETERSCH B., HADWIGER M., HAUSER H., 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. 9

[PLM*08] PAPENBERG N., LANGE T., MODERSITZKI J., SCHLAG P. M., FISCHER B.: Image registration for ct and intra-operative ultrasound data of the liver. Miga M. I., Cleary K. R., (Eds.), vol. 6918, SPIE, p. 691808. 8

[PSSH06] PETERSCH B., SERRANO-SERRANO O., 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). 6

[RDR10] ROPINSKI T., DÖRING C., REZK-SALAMA C.: Interactive volumetric lighting simulating scattering and shadowing. In *Proceedings of IEEE Pacific Visualization* (2010), pp. 169–176. 10, 11

[RGB99] ROHLING R., GEE A., BERMAN L.: A comparison of freehand three-dimensional ultrasound reconstruction techniques, 1999. 4

[RGBT99] ROHLING R., GEE A., BERMAN L., TREECE G.: Radial basis function interpolation for 3D ultrasound, 1999. 5

[RN00] ROXBOROUGH T., NIELSON G. M.: Tetrahedron based, least squares, progressive volume models with application to freehand ultrasound data. pp. 93–100. 5

[SCT00] STETTEN G., CHIB V., TAMBURO R.: Tomographic reflection to merge ultrasound images with direct vision. *Proceedings of 29th Applied Imagery Pattern Recognition Workshop* (2000), 200–205. 11

[Sha06] 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. 5

[SLT*07] SOLBERG O. V., LINDSETH F., TORP H., BLAKE R. E., HERNES T. A. N.: Freehand 3d ultrasound reconstruction algorithms - a review, 2007. 4

[SNT*98] SATO Y., NAKAMOTO M., TAMAKI Y., SASAMA T., SAKITA I., NAKAJIMA Y., MONDEN M., TAMURA S.: Image guidance of breast cancer surgery using 3-D ultrasound images and augmented reality visualization. *IEEE transactions on Medical Imaging 17*, 5 (oktober 1998), 681–93. 11

[SO94] STEEN E., OLSTAD B.: Volume rendering of 3D medical ultrasound data using direct feature mapping. *IEEE transactions on Medical Imaging 13*, 3 (januar 1994), 517–25. 6, 8

[SOB*94] STEEN E., OLSTAD B., BERG S., MYKLEBUST G., SCHIPPER K.: Combined visualization of tumor and vessel geometry. *Proceedings of IEEE Ultrasonics Symposium* (1994), 1645–1648. 6

[SPH*09] SCHOTT M., PEGORARO V., HANSEN C., BOULANGER K., BOUATOUCH K.: A directional occlusion shading model for interactive direct volume rendering. pp. 855–862. 10, 11

[SSC02] SHELTON D., STETTEN G., CHANG W.: Ultrasound visualization with the sonic flashlight. *Proceedings of ACM SIGGRAPH* (2002), 82. 11

[SSG95] SAKAS G., SCHREYER L.-A., GRIMM M.: Preprocessing and volume rendering of 3D ultrasonic data. *IEEE Computer Graphics and Applications 15*, 4 (1995), 47–54. 5, 6

[TGHM96] THUNE N., GILJA O. H., HAUSKEN T., MATRE K.: A practical method for estimating enclosed volumes using 3d ultrasound. *European Journal of Ultrasound 3*, 1 (1996), 83 – 92. 5

[TH98] T. HAUSKEN O. H. GILJA S. Ø. A. B.: Flow across the human pylorus soon after ingestion of food, studied with duplex sonography: Effect of glyceryl trinitrate. *Scandinavian Journal of Gastroenterology 33*, 5 (1998), 484–490. 1

[TL08] TAN C., LIU D. C.: Interactive Editing for 3D Ultrasound Volume Rendering. *Proceedings of 2nd International Conference on Bioinformatics and Biomedical Engineering* (mai 2008), 2307–2310. 7

[VNOy*08] VIOLA I., NYLUND K., Ø YE O., ULVANG D., GILJA O.: Illustrated Ultrasound for Multimodal Data Interpretation of Liver Examinations. *Proceedings of Eurographics Workshop on Visual Computing in Biomedicine* (2008). 10

[vPB10] ŠOLTÉSZOVÁ V., PATEL D., BRUCKNER S.: A multidirectional occlusion shading model for direct volume rendering. *Computer Graphics Forum 29*, 3 (june 2010), 883–891. 10, 11

[WK08] WEIN W., KHAMENE A.: Image-based method for in-vivo freehand ultrasound calibration. In *SPIE Medical Imaging* (2008). 4, 7

[WKSS05] WU B., KLATZKY R. L., SHELTON D., STETTEN G. D.: Psychophysical evaluation of in-situ ultrasound visualization. *Transactions on Visualization and Computer Graphics 11*, 6 (2005), 684–693. 11

[WSC08] WANG S.-R., SUN Y.-N., CHANG F.-M.: Artifact removal and texture-based rendering for visualization of 3D fetal ultrasound images. *Medical & Biological Engineering & Computing 46*, 6 (june 2008), 575–88. 8

[WSSL83] WAGNER R., SMITH S., SANDRIK J., LOPEZ H.: STATISTICS OF SPECKLE IN ULTRASOUND B-SCANS. *Transactions on Sonics and Ultrasonics 30*, 3 (1983), 156–163. 2

[YMK07] YOO Y. M., MANAGULI R., KIM Y.: New multi-volume rendering technique for three-dimensional power Doppler imaging. *Ultrasonics 46*, 4 (november 2007), 313–22. 10

[ZNB] ZHANG W., NOBLE J., BRADY J.:. *Proceedings of IEEE International Symposium on Biomedical Imaging: Macro to Nano*. 7

[ZRP02] ZHANG Y., ROHLING R., PAI D.: Direct surface extraction from 3D freehand ultrasound images. *Proceedings of IEEE Visualization* (2002), 45–52. 6