Ultrasound Guided Surgery:
Image Processing and Navigation

Thomas Langø

Preface

This thesis is submitted to the Norwegian University of Science and Technology (NTNU), for partial fulfillment of the requirements of the degree "Doktor Ingeniør". The research work was performed at the Department of Ultrasound, SINTEF Unimed, in the period 1996-2000 in collaboration with Trondheim University Hospital, particularly the Departments of Neurosurgery, Vascular Surgery, and Laparoscopic Surgery, the Department of Computer and Information Science, NTNU, and the Department of Mathematical Sciences, NTNU.

My supervisor has been Åge Grønningsæter at the Department of Ultrasound, SINTEF Unimed (now CEO at MISON AS), while professor Bjørn A. J. Angelsen at the Department of Physiology and Biomedical Engineering, NTNU, has been the administrative supervisor.

Acknowledgements

Contents

List of papers

Paper A – Novel probe calibration methods for 3D freehand ultrasound
Thomas Langø, Frank Lindseth, Jon H. Kaspersen and Åge Grønningsæter.
Submitted to Computer Aided Surgery, October 2000.

Paper B - Wavelet-based edge detection in ultrasound images
Jon H. Kaspersen, Thomas Langø and Frank Lindseth.
Accepted for publication in Ultrasound in Medicine and Biology, August 2000.

Paper C - Navigation of probes and surgical instruments in ultrasound-guided surgery
Thomas Langø, Steinar Ommedal, Geirmund Unsgård and Åge Grønningsæter.
Submitted to Computer Aided Surgery, October 2000.

Paper D - Accuracy assessment of a prototype ultrasound-based neuronavigation system
Åge Grønningsæter, Frank Lindseth, Thomas Langø and Geirmund Unsgård.
Submitted to Neurosurgery, April 2000.

Paper E - Bayesian 2D deconvolution: A model for diffuse ultrasound scattering
Oddvar K. Husby, Torgrim Lie, Thomas Langø, Jørn Hokland and Håvard Rue.
Accepted for publication in IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, April 2000.

Paper F - Bayesian 2D deconvolution: Effect of using spatially invariant ultrasound point spread functions
Thomas Langø, Torgrim Lie, Oddvar K. Husby and Jørn Hokland.
Accepted for publication in IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, April 2000.

Appendix
Paper G - SonoWand®, an ultrasound-based neuronavigation system
Åge Grønningsæter, Atle Kleven, Steinar Ommedal, Tore E. Årseth, Torgrim Lie, Frank Lindseth, Thomas Langø and Geirmund Unsgård.
Accepted for publication in Neurosurgery, August 2000.

Introduction

This thesis contains seven separate papers (one as an Appendix), which are complete with abstracts and references. They have all been submitted for publication to scientific journals. A summary for each paper is given below. Since three of the papers focus on methods that traditionally are not common in ultrasound, introductions are given on the topics of image processing using wavelets, the main subject in paper B, and Bayesian inference, the theme in papers E and F. The section of concluding remarks and possible future directions is followed by a list of the main references used during this research project and preparation of the papers.

The papers, posters, abstracts, and patent applications submitted and published during this work are also listed in the References [52-54,60,70,79,86-92,107].

1. Background

This thesis is a part of the ongoing research activity on ultrasound guided surgery at the Department of Ultrasound, SINTEF Unimed, in Trondheim, Norway. In 1995, the Ministry of Health and Social Affairs established a center for advanced use of three-dimensional (3D) ultrasound in surgery in Trondheim; the center consisted of a close collaboration between the Department of Ultrasound at SINTEF Unimed, the University Hospital in Trondheim (RiT), and the Medical Faculty at the Norwegian University of Science and Technology (NTNU). The strength and extent of this collaboration has increased since and resulted in the establishment of the company MISON in 1998.

The research activity has concentrated on ultrasound guided neurosurgery, vascular surgery, and laparoscopic surgery. The projects have included research and development of equipment and technology where ultrasound is used as the main intraoperative imaging modality. Preoperative data from modalities such as magnetic resonance imaging (MRI) and computed tomography (CT) have been used for planning and in combination with intraoperative ultrasound images to provide an overview of the anatomy during navigation of minimally invasive surgical procedures.

Ultrasound is expected to be an increasingly important factor in surgery for many reasons. The image quality of modern ultrasound equipment can compete with MRI and CT imaging in certain situations. Another feature of ultrasound is its real time capabilities, making it suitable as an intraoperative imaging modality that can cope with the problems that arise from changes in the anatomy during surgery. Furthermore, ultrasound has the potential of generating not only real time 2D and repetitive 3D, but also real time 3D images, which are expected to dominate in second or third generation ultrasound-based navigation systems. The 3D imaging capability of ultrasound represents a valuable supplement or alternative to interventional MRI because of lower investment and running costs as well as more convenient and portable equipment. However, preoperative MRI and/or CT data will still remain valuable to the surgeon in the planning procedure and during surgery to get an overview of the anatomy of the patient.

Ultrasound-guided surgery

In image-guided surgery, images are used to guide the surgical tool to the lesion inside the patient, requiring (in theory) only a narrow channel for the surgical tool through the normal tissue. Using ultrasound as the main intraoperative imaging modality allows the surgeon to monitor the procedure with the aid of real time 2D imaging. In addition, repetitive 3D freehand ultrasound acquisition is exploited in order to get updated 3D visualizations of the anatomical area of interest.

The process of 3D freehand ultrasound acquisition consists of two steps: scanning and reconstruction. Prior to scanning, a position sensor is attached to the probe for tagging each image with position and orientation in space. A crucial step prior to reconstruction is the calibration of the ultrasound probe to which a position sensor is attached. The calibration procedure establishes a transformation matrix, given by three translation offsets and three angle offsets, defining the position and orientation of the image plane relative to the position sensor attached to the probe. An accurate calculation of this matrix is important in order to obtain a reconstruction that preserves true anatomical shape and size in freehand 3D ultrasound scanning and for navigation involving the real time 2D image. A crude estimate of the transformation may be found by external measurements of the probe case and sensor. However, this estimate will not necessarily be exact since the origin of the ultrasound image system is located inside the housing of the probe and may vary slightly from one probe to another. Furthermore, the position of the center transducer element (origin) and the direction for pulse transmission may vary due to variations in the production of the probe. In addition, for magnetic-based position sensors the origin for the coils in the sensor is not well defined. This indicates that a better approach to finding the calibration matrix would be to image a phantom with known physical properties and dimensions. In this way we can make sure we are using a more precise position for the ultrasound image plane origin.

During our research in Trondheim, we have experienced that a valuable technique during removal of tumors in neurosurgery is to guide the surgical tool into position using real time 2D ultrasound imaging from a separate burr hole in the skull (or alternatively one larger burr hole for both the tool and probe) [52]. In addition, real time ultrasound 2D imaging is used to monitor the shrinkage of the tumor or hematoma during surgery. However, one major limitation with this technique is the difficulty of continuously obtaining a longitudinal cross section view of the surgical instrument, or the distal part of it in the real time ultrasound image. It is of crucial importance to obtain such a view in order to be able to localize the tip of the instrument in relation to important anatomical structures and thereby perform biopsy and minimally invasive surgery with high precision and accuracy. From the probe calibration procedure, we know the location of the 2D image, and by performing a calibration of the surgical tool with an attached position sensor, we are able to calculate the relative position and orientation between the image and the distal part of the tool. This information can in turn be used by the surgeon to adjust either the probe or the tool to obtain a cross section view of the tool in the real time 2D ultrasound image.

An important parameter for the user of a computer-based navigation system is the overall clinical accuracy of the system and the procedure, i.e. a value telling the surgeon how accurately he/she can navigate the tip of a surgical instrument inside the patient's body by means of image information provided on the navigation monitor. The better the accuracy, the more delicate and precise work the surgeon can perform based on image information. Numerical data about the overall clinical accuracy is practically impossible to calculate or

measure since these values will depend not only on the navigation system itself, but also on patient preparation procedures, system set-up, and system handling. Nevertheless, by estimating values for the accuracy of separate steps in the procedure and by performing experiments with phantoms, we believe it is possible to get a close estimate of the accuracy that can be expected under real conditions during surgery.

In order to produce renderings for visualization, e.g. stereoscopic views, we depend on images where there is a high contrast between the object we want to visualize and its surroundings. Unfortunately, ultrasound images often appear noisy and blurred. These drawbacks have previously prevented satisfactorily restorations. The most common approach to ultrasound image restoration seems to be attempting to estimate the point spread function (*psf*) as close to its true shape as possible and then use this in the deconvolution process [1,73,74,77,111,112,135-137]. Wiener filtering [14,72,135] is widely used for deconvolution in ultrasound scanners to provide partial restoration of the reflectance. Nevertheless, the restored images still suffer from significant blur and speckle that adversely affect the observer's ability to interpret the images.

2. Aims of the study

The overall goal of this study was to develop techniques and equipment for ultrasound guided minimally invasive surgery and implement these for use in a prototype neuronavigation system [52].

Several methods for probe calibration have been presented in the literature, most of them described and referred to in [113] and [89]. However, we felt that the existing techniques were either not accurate enough or the necessary equipment was unnecessarily complicated. Therefore, we wanted to develop new and simpler (more automatic) methods for more accurate estimation of the probe calibration transformation, i.e. more precise determination of the image scan plane position and orientation in space in relation to a position sensor attached to the ultrasound probe.

We also wanted to develop techniques for calibration of surgical instruments in order to know the position and orientation of the tip portion of the tool in relation to an attached position sensor. This task demanded a different approach compared to the probe calibration since the different instruments have extremely different shapes and sizes. In addition, we wanted to be able to perform the tool calibrations quickly in the operating room under sterile conditions as soon as a sensor was attached to the tool since the specific situation during the surgical procedure determines the way the sensor is attached.

Furthermore, we wanted to develop a method for guiding surgical tools during navigation in relation to the real time 2D ultrasound image in order to obtain a longitudinal cross-sectional view of the tip portion of the instrument on the ultrasound scanner. The information could include the distance between the tip of the tool and the image plane and the angle between the tip portion of the tool and the image plane. These pieces of information may be calculated from the position readings and the calibration transformations described earlier.

In addition to developing the techniques described above, we needed procedures and equipment for analyzing the accuracy of the methods. Many methods might be investigated

separately, but for techniques involved in a navigation system for surgery it is important to estimate the overall accuracy of the navigation procedures.

Finally, in order to improve renderings for visualization we wanted to investigate new methods for removing the blur and the characteristic speckle pattern found in medical ultrasound images. The aim was to develop a method that increases the contrast in ultrasound images and hence improves visualizations of medical ultrasound data.

3. Summary of papers

Paper A – Novel probe calibration methods for 3d freehand ultrasound

Two probe calibration methods are described in Paper A. The first method we developed is a 2D-based membrane alignment technique. The real time 2D image scan plane is aligned with a thin and rigid membrane oriented approximately perpendicular to the bottom floor of a water tank. The membrane has five distinct corners that are marked in the acquired 2D image. The position vector of the same five points are found in the global coordinate system using a pre-calibrated pointer. The position and orientation of the sensor on the probe in this particular situation, i.e. when the image and the physical membrane are believed to be coplanar, is also acquired. We are then able to calculate the transformation between the sensor on the probe and the image plane origin. Since the probe and sensor constitute one rigid body, this transformation is applicable in general.

From our initial experiences we have found the newer and more automatic freehand-based membrane scan method to be easiest to use. It involves the use of an easy-to-build phantom and, most importantly, it is automatic once the scanning has been performed. The only input to the procedure is an initial guess solution of the transformation matrix and the velocity of sound. The procedure is performed by acquiring a sequence of images of a plane membrane in a water bath oriented approximately horizontal from a range of different angles and distances, covering the six degrees of freedom for movement in space. Next, the membrane (a fine mesh of woven nylon) is detected automatically (as a thin line) in all images using an algorithm based on the wavelet transform (see Paper B). These edge detections are used to find the best straight line fit to represent the membrane; points from these lines are then picked and used in a system of equations that incorporate the position data from the sensor attached to the probe. By making sure we acquire more images than there are unknown parameters we get an over-determined set of N equations with M unknowns ($N>M$). The equations are set up from the transformation of a point from the image system (the membrane) via the sensor on the probe, to the global coordinate system, and back to the membrane system. The unknowns in this transformation, i.e. the calibration matrix, are found by an iterative algorithm, requiring only an initial guess solution as input. A patent application has been filed for the main ideas of this new automatic method [87].

The accuracy of the two methods are compared using a special precise wire phantom and the conclusion is that the *Freehand Membrane Scan Method* performs slightly better than the *Membrane Alignment Method*. The latter method depends on how accurately the user is able to align the ultrasound scan plane with a thin membrane oriented perpendicular at the bottom of a water bath. This can be a difficult task given the thickness of the ultrasound beam. Furthermore, the user must manually mark points in both the image and in the phantom using

a pre-calibrated pointer. This means that the older method is subject to great inter-operator variability.

Paper B - Wavelet-based edge detection in ultrasound images

This paper presents the automatic edge detection method used in the automatic freehand probe calibration method described in Paper A. The method was especially developed for detecting the kind of edges encountered in freehand scanning of a thin planar membrane in a water bath (see sample images in Paper A and B). However, the method proved very robust and may be used for other purposes as well. A reliable and automatic edge detection is of great importance, for instance as an initialization step in various active contour algorithms. Furthermore, reliable segmentation algorithms can be used in conjunction with 3D ultrasound imaging, tissue characterization, and various techniques to provide quantitative information about acoustic properties of tissue and volume measurements.

The algorithm is based on the continuous wavelet transform (CWT). One beam at a time is processed by the detection procedure. In summary, the beam is first transformed into the wavelet domain by using the CWT. Next, the edge is localized by determining the position of the maximum value for the squared CWT at the most energetic scale. The most energetic scale is established by computing the power spectrum based on the CWT modulated by the square of the scale value, i.e. the energy contribution at each scale as a function of the scale. The position for the maximum from this energy distribution is selected as the most energetic scale. One of the main advantages of the algorithm is that, in its simplest form, it is independent of any threshold. For automatic edge detection methods this is of great advantage and importance in order to minimize complex input parameters. In addition, an edge detection method for ultrasound images needs to be robust with respect to noise. We tested this quality by artificially adding noise to some sample images, with good results.

An introduction to wavelet-based image analysis is presented in section 4.

Paper C - Navigation of probes and instruments in ultrasound-guided surgery

We have developed a navigator module that can run in conjunction with the main navigation software, which manages various types of 2D/3D medical data and a position sensor system in combination with an ultrasound scanner for image-guided surgery. The main ultrasound-based navigation system (including its software) is described in [52], which is included in the Appendix (Paper G). The real time navigator described in Paper C gives the surgeon feedback about the position and orientation of the surgical tool relative to the ultrasound image plane in space. In the implementation presented in this paper, the navigator consists of a small graphic bar display that may be placed anywhere on the navigation monitor or on the ultrasound scanner monitor. The information presented in this graphic bar is the shortest distance between the tip of the tool and the image plane and the angle between the distal part of the tool and the image plane. The calculations are based on position measurements from the sensors attached to the probe and the tool. In our initial tests we have used bright colors (different for distance and angle, see Fig. 5 in Paper C) that are easily discernable to the surgeon. The main idea is that the surgeon should be able to understand the spatial relation between the distal part of the tool and the real time 2D ultrasound image with just a quick glance at the graphic navigator bar. The information in the graphic bar display is updated in

real time, which makes it a simple task to adjust the position of the ultrasound probe in order to achieve a longitudinal cross section view of the tool.

In order to calculate parameters describing the spatial relationship between the image plane and the tool, it is necessary to calibrate the tool. This means finding the transformation from the position sensor attached to the tool to the distal part of the tool and the tool tip. In this paper we describe a method for performing tool calibrations under sterile conditions. Both the technique and equipment for this tool calibration technique are further described in a patent application from our group [90], while the corresponding transformation for the ultrasound probe with an attached position sensor is described in Paper A.

Initial trials with the navigator module in the operating room setting during resection of brain tumors have been promising. Adjusting the probe during insertion of the tool was simplified by the use of the navigator since the surgeon got immediate feedback on the spatial relation between the tool and image.

The principle of the tool navigator can also be exploited with respect to 3D visualization. A 3D volume with position information can be acquired and transferred to a visualization computer. The tool navigator information can then be utilized as the surgeon inserts the tool into the anatomical area of interest by drawing a graphic representation of the tool in the 3D volume being visualized on the monitor, e.g. stereoscopically. This information on the representation of the tool (see Fig. 6 in Paper C) will then tell the surgeon where the tool is located in relation to the real time 2D image plane shown on the ultrasound scanner. The 2D image plane may serve as a supplement to the 3D scene, providing updated and detailed image information while the navigator parameters projected onto the tool representation in the 3D scene provide a link between the visualizations.

Paper D - Accuracy assessment of a prototype ultrasound-based neuronavigation system

The accuracy of the probe calibration was evaluated in Paper A. In image-guided surgery there are numerous other sources that contribute to the inaccuracy. This inaccuracy is the difference between the point that the surgeon believes he/she is pointing at from the images on the screen and the actual point, i.e. the location of the tip of the pointer or surgical tool on or inside the patient. Some groups have described and estimated error sources associated with neuronavigation [24,57,59,76,123,125,128,129,141], but we believe that this list is not complete. Paper D supplements the description and gives a more detailed overview of the reduced error chain associated with ultrasound-based neuronavigation compared to the chain of possible error sources associated with conventional MRI-based neuronavigation.

The overall accuracy of our prototype navigation system [52] is measured by a controlled water tank experiment using the same wire phantom as used in Paper A. Preliminary quantification of the system resolution is also presented. The accuracy is evaluated by measuring the differences between wire crossings in the image data (3D acquisition) and in global coordinates. In paper D, we found the wire crossings in global coordinates by using a pre-calibrated pointer while in paper A we utilized the physical measurements from the design and construction of the phantom to define the wire crossings. This means that the accuracy analysis in Paper D includes another possible error source, i.e. the pointing accuracy, which includes both the pointer calibration and the position measurement.

Furthermore, six separate 3D acquisitions and calibrations were performed in Paper A and only one in Paper D.

The major results from Paper D indicate that the overall clinical accuracy associated with ultrasound-based navigation is similar to the accuracy that can be expected from conventional MRI-based navigation systems. Ultrasound-based navigation does not require patient registration as MRI-based navigation does. Patient registration is the procedure that finds the correlation, i.e. matches the position and orientation, between the preoperative images and (MRI/CT) and the patient prior to surgery. Thus, ultrasound is less susceptible to user- and procedure-dependent errors.

Paper E - Bayesian 2D deconvolution: A model for diffuse ultrasound scattering

This is the first of two papers concerning the use of Bayesian Image Restoration methods to improve the image quality of ultrasound images. In order to produce renderings for visualization, e.g. stereoscopic views for tool navigation in ultrasound guided surgery (see Paper C), we depend on images where there is high contrast between the object we want to visualize and its surroundings. Unfortunately, ultrasound images often appear noisy and blurred. Another problem is caused by the speckle effect. Due to the coherent detection in ultrasound imaging, constructive and destructive interference between reflections from small, closely located objects give rise to multiplicative noise, called speckle. This gives the ultrasound images a grainy appearance, and makes high quality 3D visualization difficult.

In order to remove both blur and speckle, we have developed a new statistical model for diffuse scattering in 2D ultrasound radio frequency images by incorporating both spatial smoothness constraints and a physical model for diffuse scattering. The modeling approach is Bayesian in nature, and we use Markov chain Monte Carlo methods to obtain the restorations. The results from restorations of some real and simulated radio frequency ultrasound images are presented, and compared with results produced by Wiener filtering.

The results demonstrate the potential usefulness of Markov random field models in restoration of medical ultrasound images. The restoration of the simulated images produced results that are more representative of the original object, i.e. the true tissue distribution, than the observed ultrasound images are. Also, in the case of the real image, the results obtained are significantly better than those produced by Wiener filtering. In both cases, the speckle patterns are efficiently removed, without the loss of important details or the introduction of artificial structures.

An introduction to Bayesian inference in image analysis is presented in section 5.

Paper F - Bayesian 2D deconvolution: Effect of using spatially invariant ultrasound point spread functions

The Bayesian Image Restoration method presented in paper E is based on knowledge of the ultrasound imaging system, including the size and shape of the *psf*. In a real situation, however, one can not assume that the *psf* is accurately known. Moreover, the *psf* may vary substantially across the imaging area due to the natural broadening of the ultrasound wave-

field, and this variation is enhanced as the ultrasound pulse is distorted by tissue inhomogeneities.

Consequently, we studied the robustness of this restoration method when using incorrect *psf*'s during the restoration process. We used simulated ultrasound images in order to compare restoration results with the true images. The true images were first degraded with a known model *psf* followed by adding Gaussian noise. Then, during restoration the *psf* was systematically varied. We found that small variations in one of the parameters characterizing the *psf*, i.e. less than ±25% change of transducer center frequency, beam width, or pulse length, still gave satisfactory results. When altering the *psf* more than this, the restorations were not acceptable. The method proved most sensitive to changes in frequency. Particularly, increasing the frequency by more than 25% resulted in poor restorations. Thus, 2D Bayesian restoration using a fixed *psf* may yield acceptable results as long as the true variant *psf*'s do not differ too much from the estimated fixed *psf* used in the restoration.

4. Wavelet-based image processing

Singularities, which are discontinuities in functions, and irregular structures often contain the most important information in intensity signals. In images, these discontinuities in intensity provide the locations of object contours, which are particularly useful for recognition purposes.

The Fourier transform (FT) has long been a popular tool for analyzing singularities. The FT is global in the sense that it provides a description of the overall regularity of signals. It is not well adapted for finding the location and spatial distribution of singularities. A serious drawback is the loss of time information when the FT is used to transform a signal to the frequency domain. Furthermore, in order to obtain an ordinary spike signal with positive width $T \sim \varepsilon$ such as $\delta_\varepsilon(t)$ (Dirac delta), we must combine exponentials in a frequency band of finite width $\Omega \sim 1/\varepsilon$. The narrower the spike, the more high-frequency exponentials that must enter into its construction in order to achieve the desired sharpness of the spike. But since these exponentials oscillate forever, we need even more exponentials in order to cancel the previous ones before and after the spike occurs. This shows the immense inefficiency of ordinary Fourier analysis in dealing with local behavior. It is clear that this inefficiency is due entirely to its nonlocal nature. After all, we are trying to compose local variations in time by using nonlocal exponentials.

In an effort to correct these deficiencies, Gabor [41] adapted the FT to analyze only small sections of the signal at a time, a technique called *windowing* the signal. This Short-Time Fourier Transform (STFT) maps a signal into a 2D function in time and frequency. The STFT represents a compromise between the time- and frequency-based views of a signal. It provides some information about when and at what frequency a certain signal event occurs. However, the information retrieved has a limited precision, given by the size of the window. Once the size of the window is chosen, that window is the same for all frequencies. This means that the windowed Fourier transform only partially solves the problem of composing local variations of signals. Many signals require a more flexible approach, one where the window size can be varied to view more accurately either the time or frequency contents.

Wavelet analysis [22,27,75,116,117]represents the next logical step, namely a windowing technique with variable sized regions. Wavelet analysis allows the use of long time intervals

where precise low frequency information is needed, and shorter intervals where information about high frequency content is needed. By decomposing the signal into elementary building blocks that are well localized in space and frequency, the wavelet transform (WT) can characterize the local regularity of signals. This is a major advantage of the WT and the most interesting dissimilarity between the FT and the WT: the individual wavelet functions are localized in space and hence are appropriate tools for analyzing a localized area of a larger signal.

Just as the STFT uses modulation in the time domain to translate the window in frequency, the wavelet transform uses scaling in the time domain to scale the window in frequency. From a physical standpoint, the scaling of frequencies might seem more natural: in music, going up an octave always involves doubling the frequency, rather then shifting it by a constant additive term.

The figure below illustrates the viewing of a signal in the different domains; time, frequency, short time Fourier, and wavelet.
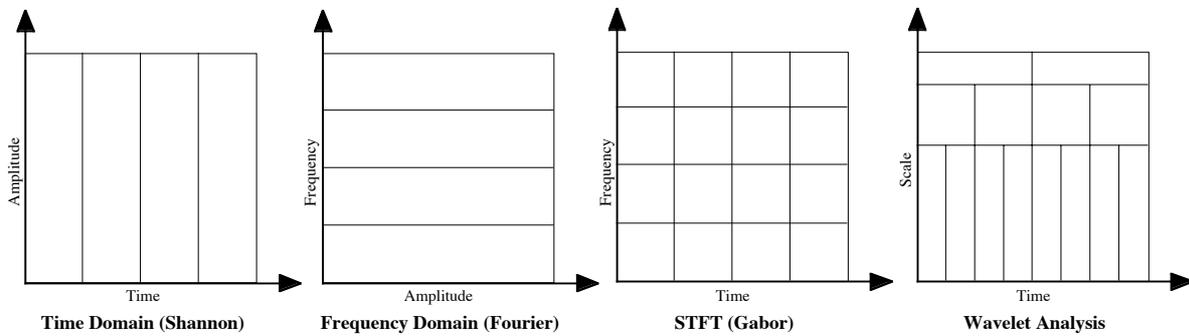

Fig. 4.1. Illustration of the manner in which a signal is viewed in different domains: time, frequency, short time Fourier, and wavelet analysis.

The basis of the WT is that the signal to be analyzed is compared with a function Ψ, often referred to as the analyzing wavelet or mother wavelet. The wavelet transforms do not have a single set of basis functions like the FT, which utilizes only the sine and cosine functions. Instead, wavelet transforms have an infinite set of possible basis functions providing access to information that can be obscured by other time-frequency methods such as Fourier analysis. Dilations by the scaling parameter a and translations by the shift parameter τ of the mother wavelet define an orthogonal basis, the wavelet basis [102] given by

$$\Psi_{\tau,a}(t) = \frac{1}{\sqrt{a}} \Psi\left(\frac{t-\tau}{a}\right) \tag{4.1}$$

where t represents time for the function to be analyzed or transformed. The WT is defined by [102]

$$W_f(\tau,a) = \frac{1}{\sqrt{a}} \int_{-\infty}^{\infty} f(t)\Psi\left(\frac{t-\tau}{a}\right) dt \tag{4.2}$$

where $f(t)$ is the signal to be analyzed. The wavelet transform may be visualized as a microscope used on the signal $f(t)$. The scale parameter a then represents the magnification, and the translation parameter τ represents the location chosen for the observation.

The mother wavelet can be any function, real or complex, satisfying the admissibility condition [102] expressed as

$$C_\Psi = \int_{-\infty}^{\infty} \frac{|F_\Psi(\omega)|^2}{\omega} d\omega < +\infty \qquad (4.3)$$

where C_Ψ is known as the admissibility constant, and $F_\Psi(\omega)$ is the Fourier transform of $\Psi(t)$. This condition implies that

$$\int_{-\infty}^{+\infty} \Psi(t)dt = 0 \qquad (4.4)$$

In mathematics, the local regularity of signals is often measured with Lipschitz exponents [102]. The theory of Lipschitz exponents and WT can be used to localize and characterize singularities of signals. For a well written description of this theory and its application to singularity detection, see the paper by Mallat and Hwang [102]. Their theorems demonstrate that the WT is particularly well adapted for estimating the local regularities of signals. Nevertheless, in order to detect singularities and characterize their Lipschitz exponents, their theorems indicates that we have to measure the decay of the absolute value of the WT in a whole 2D neighborhood of a point in the scale space, which requires a lot of computations. The estimation of Lipschitz exponents up to a maximum value n requires a wavelet with at least n vanishing moments. A wavelet is said to have n vanishing moments if for all positive integers $k<n$, it satisfies

$$\int_{-\infty}^{+\infty} t^k \Psi(t)dt = 0 \qquad (4.5)$$

Wavelets with more vanishing moments have the advantage of being able to measure the Lipschitz regularity up to a higher order, but a higher number of vanishing moments also increases the number of maxima lines in the plot of the absolute value of the WT as a function of time and scale. In order to minimize the amount of computations, a minimum number of maxima necessary to detect the irregularities of the signal is desired. This means that the choice of wavelet should be one with as few vanishing moments as possible, but with enough moments to detect the Lipschitz exponents of highest order that we are interested in. Another related property that influences the number of modulus maxima is the number of oscillations of the wavelet. A wavelet with n vanishing moments, has at least $n+1$ local extrema.

5. Bayesian image analysis

This presentation is a simplified introduction to Bayesian image analysis, mainly for researchers working in the field of ultrasound imaging who are not familiar with the statistics involved in Bayesian inference. Recommended reading for this theory is Robert [119] and Winkler [143].

The Bayesian approach to image restoration was first introduced by Geman and Geman [44]. In Bayesian inference there are two types of information: prior knowledge about the probabilistic generating mechanism (cause) which degrades the original scene, and empirical or observed data (effect). These are both sources of uncertainty or randomness, since empirical data are mostly distorted or degraded ideal data and prior information is usually incomplete. Another aspect making Bayesian inference different from other methods, such as inverse filtering, is the replacement of rigid constraints by weaker ones. In other words, instead of classifying the data into allowed and forbidden states, the data are weighted by an acceptance function quantifying the degree to which they are desired or not. This leads to a probability measure on the set of objects in the data set (e.g. pixels in an image), called the prior distribution or simply the prior. This weak constraint measure is consistently combined with the data or observation modifying the prior into the posterior distribution. In this distribution, the expectations incorporated in the prior compete with faithfulness to the data. By integrating an appropriate rule of choice or decision, a solution to the inverse problem may be found, i.e. an image in proper balance with both the prior anticipations and loyalty to the observed data.

A central topic in Bayesian image analysis is sampling, i.e. drawing random values from a distribution with known parameters. This is opposite from the situation regularly encountered in statistics where an estimate of a distribution is sought. In Bayesian inference we have the distribution and seek a restoration as close to the ideal as possible by sampling candidates from the distribution. The posterior distribution for the unknown image can be found by utilizing Bayes' formula [6], which can be deduced from the equation for conditional probability

$$p(X|Y) = \frac{p(X,Y)}{p(Y)} = \frac{p(Y|X)p(X)}{p(Y)} \propto p(Y|X)p(X) \qquad (5.1)$$

In this equation, X and Y represent the unknown ideal image and the observed image respectively. X and Y are considered stochastic variables. $p(X,Y)$ is the probability for both X and Y, while $p(X|Y)$ means the probability for X given Y; this distribution thus works conditionally on the observations. The last transition can be performed since $p(Y)$ may be regarded as a constant (Y is considered known). Eq. 5.1 gives the probability or weighted acceptance for each candidate solution image X. The distribution $p(X)$ is the prior distribution or a model for the known image. The function $p(Y|X)$ in Eq. 5.1 might seem odd since we do not know X, but by using Y as the known parameter in this equation we are utilizing the likelihood function [143]

$$L(X|Y) = \prod_i p(Y = y_i) = \prod_i p(Y_i|X) \qquad (5.2)$$

The prior distribution $p(X)$ incorporates typical features of the ideal image such as the fact that neighbor pixels most likely have the same value, i.e. smooth homogenous regions

separated by sharp edges. At the same time, it is possible to allow for smoother transitions between regions. Undoubtedly the most critical, challenging and criticized characteristic of Bayesian analysis is the choice of the prior distribution. Moreover, this choice is highly problem dependent. In practice, the available prior information is seldom precise enough to result in an exact determination of the prior, meaning that other probability distributions may be compatible with the same information. An approximation might be necessary, which may drastically alter the subsequent inference.

In order to find a distribution for the likelihood function, we must know something about the degrading mechanisms involved in the imaging process. Assuming a spatially invariant *psf*, the formation of an ultrasound image can be modeled as a 2D convolution (blurring) of the *psf*, *h*, and the unknown sharp image, *x*, and the addition of independent normally distributed white Gaussian noise η with variance τ^2

$$y_{i,j} = \sum_{k,l} h_{k,l} x_{i-k,j-l} + \eta_{i,j} = \{h * x\}_{i,j} + \eta_{i,j} \tag{5.3}$$

$$\eta_{i,j} \sim N(0, \tau^2) \tag{5.4}$$

Both the *psf* and the variance are assumed known. The indexes *i,j* refer to position in the image pixel lattice, while $\eta_{i,j}$ represents independent Gaussian variables with mean zero and common variance τ^2. If we for the moment consider X known, then $\{h * x\}_{i,j}$ simply shifts the expectation value for the pixel distribution

$$p(Y_{i,j}|X) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2\sigma^2}\left(Y_{i,j} - \{h * x\}_{i,j}\right)^2\right] \tag{5.5}$$

This equation tells us how we can expect to observe degraded pixel values given the ideal image *X*. If we further consider pixel noise to be non-correlated between pixels, we get the following distribution for the likelihood function in Eq. 5.1 and 5.2

$$p(Y|X) = \prod_{i,j} p(Y_{i,j}|X) = \frac{1}{\left(\sqrt{2\pi}\sigma\right)^{N\cdot M}} \exp\left[-\frac{1}{2\sigma^2}\sum_{i,j}\left(Y_{i,j} - \{h * x\}_{i,j}\right)^2\right] \tag{5.6}$$

where *M·N* is the dimension of the image. By inserting this likelihood function and a proper prior function into Eq. 5.1 we may calculate the posterior distribution for comparing candidate solution images. For the sake of simplicity a plain smoothing prior is used in this introduction. Such a function may be given by

$$p(X) \propto \exp\left[-\beta \sum_{i,j}\left\{\left(X_{i,j} - X_{i+1,j}\right)^2 + \left(X_{i,j} - X_{i,j+1}\right)^2\right\}\right] \tag{5.7}$$

where β is an input parameter used to control the weighting of the punishment for neighbor pixels that have large intensity differences. If β is large, then images with large homogenous regions will be more likely than images with a lot of edges. When $\beta=0$, all candidate images will be equally likely. The more complicated prior that we adopted in our research can be found in Eq. 4, Paper E.

In order to compare candidate solution images we need to sample entire images, or update all pixel values iteratively. A popular algorithm for this purpose is the Gibbs sampling procedure introduced by Geman and Geman [44]. This algorithm starts by initializing the image X with random values, then, iteratively for all pixels, selecting a candidate $X_{i,j}$ from the distribution

$$p\left(X_{i,j}\middle|X_{k,l\neq i,j},Y\right) \tag{5.8}$$

In other words, we pick a new value for the current pixel by assuming, wrongfully, that all the other pixels are correct. The distribution is called the conditional pixel distribution. In order to sample entire images we must first determine the conditional pixel distribution in Eq. 5.8 and then implement an algorithm for sampling single values for individual pixels. Such sampling algorithms are based on the random function and we have used the popular Metropolis-Hastings algorithm [56] in our implementations in papers E and F.

Utilizing the first half of Eq. 5.1 with $X \rightarrow X_{i,j}$ and $Y \rightarrow X_{k,l\neq i,j}$ we may write

$$p\left(X_{i,j}\middle|X_{k,l\neq i,j},Y\right) = \frac{p\left(X|Y\right)}{p\left(X_{k,l\neq i,j}|Y\right)} \tag{5.9}$$

since Y is given anyway. The marginal distribution of $X_{i,j}$ with respect to the distribution $p\left(X_{k,l\neq i,j}|Y\right)$ is given by

$$p\left(X_{k,l\neq i,j}|Y\right) = \int_{-\infty}^{+\infty} p\left(X_{i,j}=x,X_{k,l\neq i,j}|Y\right)dx \tag{5.10}$$

We are only interested in X_{ij} since this is the pixel to be sampled. All the other components of Eq. 5.9 are either known or constants. Hence, we extract the elements of the exponential function containing X_{ij} (assuming proper normalization of the prior in Eq. 5.7)

$$p\left(X|Y\right) = \exp\left[-\sum_{i,j}\left\{\frac{1}{2\sigma^2}\left(Y_{i,j}-\{h*x\}_{i,j}\right)^2 + \beta\left(X_{i,j}-X_{i+1,j}\right)^2 + \beta\left(X_{i,j}-X_{i,j+1}\right)^2\right\}\right] \tag{5.11}$$

By rearranging this equation, collecting everything that does not contain $X_{i,j}$, and inserting the result of this into Eq. 5.9 and using Eq. 5.10, we achieve the following expression for the conditional posterior distribution

$$p\left(X_{i,j}\middle|X_{k,l\neq i,j},Y\right) \propto \exp\left[-\frac{1}{2\sigma^2}\left(X_{i,j}-\mu_{i,j}\right)^2\right] \tag{5.12}$$

which is recognized as the Normal distribution, i.e. X may be considered a Gaussian field, with expectation and variance given by

$$\mu_{i,j} = \frac{\frac{1}{2\sigma^2}\sum_{m,n}\Gamma_{m,n}h_{m-i,n-j} + \beta\left(X_{i-1,j} + X_{i,j-1} + X_{i+1,j} + X_{i,j+1}\right)}{\frac{1}{2\sigma^2}\sum_{m,n}h^2_{m-i,n-j} + 4\beta} \tag{5.13}$$

$$\sigma^2_{i,j} = \frac{1}{2\left(\frac{1}{2\sigma^2}\sum_{m,n}h^2_{m-i,n-j} + 4\beta\right)} \tag{5.14}$$

where

$$\Gamma_{m,n} = Y_{m,n} - \sum_{(k,l)\neq(m-i,n-j)}h_{k,l}X_{m-k,n-l} = Y_{m,n} - \sum_{k,l}h_{k,l}X_{m-k,n-l} + h_{m-i,n-j}X_{i,j} \tag{5.15}$$

The last transition in Eq. 5.15 allows us to store the values of $\Gamma_{m,n}$ in a 2D table, reducing the double summation in Eq. 5.13 to a single sum over the values of $\Gamma_{m,n}$. For every pixel updated, we have to update as many values in the table as there are elements in h, the *psf* matrix. Accordingly, we reduce the effective neighborhood size of X from the support of h convolved with h to only two times the support of h.

Sampling from the distribution in Eq. 5.12-15 is straightforward using one of the standard algorithms for sampling Gaussian random variables. Several samplers are described in Robert [119], Winkler [143], and Ripley [118]. A popular algorithm for sampling from a Normal distribution is the Box-Müller [11] sampling algorithm. This method provides two independent normal observations out of two uniform random variables. If we first generate randomly u_1 and u_2 in (0,1), then

$$x_1 = \mu + \sigma\sqrt{-2\log u_1}\cos(2\pi u_2)$$
$$x_2 = \mu + \sigma\sqrt{-2\log u_1}\sin(2\pi u_2) \tag{5.16}$$

are two independent normal distributed variables with expectation value μ and variance σ^2. Referring to the Gibbs sampler described earlier, we may now set the current $X_{i,j}$ to x_1 according to Eq. 5.16 before sampling the next pixel. Thus we have an algorithm for updating entire images.

Presenting the results from this type of restoration of images poses some difficulties. One method would be to view the restoration process taking place in real time, i.e. see the updating of the image as the iteration process proceeded, and then simply stop it when the image is good enough. However, most often the mean of a certain number of the last few images is presented. A popular estimator is the Mean of Posterior Marginal (MPM) estimator [143], which calculates the mean pixel value for all or some of the image samples, i.e. a mean of the images calculated pixel by pixel. The name MPM is a result of calculating the mean of corresponding pixels from samples created by a marginal posterior pixel distribution. A disadvantage of this estimator is the fact that we get a smoothing of the edges, which may be the opposite of what we set out to achieve. Therefore, another estimator might be better, namely the Maximum Posterior (MAP) estimator [143] for images. This estimator, which is an optimization technique, selects the most probable image, meaning that it finds the sample image with the greatest value for *p(X|Y)*. It may be viewed as a minimization task in that what

could be regarded as the energy function inside the exponential in Eq. 5.12 should be minimized in order to retrieve the maximum probability for *p*. In other words, the computation of MAP estimators for Gibbs fields amounts to the minimization of energy functions. A simple modification of the Gibbs sampler yields an algorithm that, theoretically, finds the minimum on the spaces. This modification consists of introducing a temperature parameter, *T*, such that

$$p\left(X_{i,j}\middle|X_{k,l\neq i,j},Y\right)^{\frac{1}{T}} \propto \exp\left[-\frac{1}{T}E\left(X_{i,j},\mu_{i,j}\right)\right] \tag{5.17}$$

E represents the energy function (the exponent in Eq. 5.12). The parameter T is decreased during the iteration procedure. For small values of *T* the function *E/T* has the same minima as *E* alone but the minima are much deeper. In the literature this scheme for finding the best restoration is called Simulated Annealing (SA) [143] due to the introduction of a cooling schedule as a decreasing sequence of positive numbers *T(n)*, *n* being the number of steps in the annealing. In practice, the temperature schedule might be piecewise constant, i.e. *T* is kept constant for a while and then decreased again. The SA algorithm was inspired by statistical physics, where large systems tend toward states of minimum energy (ground states) if cooled down carefully. Finding the minima using SA is achieved only if 1/T is increased at most logarithmically. Otherwise, the restoration may be trapped in local minima. This explains why the term annealing, which means controlled cooling, is used instead of freezing. Furthermore, the iteration chain visits each single minimum several times and since the algorithm visits global minima but does not detect them, there is no obvious criterion for when to stop the algorithm. The same reasoning explains why, in general, almost sure convergence cannot be expected. We tested SA in our implementations but decided not to use this in the publication for the reasons just mentioned and the fact that the restorations using this cooling scheme did not turn out significantly better than without it. Finally, the processing time was notably increased by introducing SA in the implementations.

6. Concluding remarks and future directions

6.1. Probe calibration

Paper A describes two probe calibration methods. The methods establish the transformation between a position sensor attached to the ultrasound probe and the image scan plane. This transformation is necessary in order to be able to reconstruct freehand 3D ultrasound acquisitions and apply the navigator described in Paper C. Several publications describe methods for finding this transformation [5,29,67,82,89,94,95,110,113,114]. However, only a few papers have been solely devoted to describing new or existing techniques for probe calibration [89,113,114]. Most articles involving this area of research are mainly concerned with such problems as volume measurements and reconstruction issues [67,82,110], compound ultrasound scanning [5], accuracy of ultrasound scanning, and accuracy of position sensing [29,94,95]. In these publications, probe calibration is merely a necessity in order to perform the main objective of the work. A probe calibration is required regardless of what type of position sensor equipment one uses in order to reconstruct volumes obtained by freehand scanning. Most of the existing methods for probe calibration are limited in accuracy, ease of use, and performance time, and some need complicated equipment such as specially built calibration phantoms. Most of these issues have been addressed in our new methods, especially for the freehand-based scanning method.

The freehand-based scanning method proved more accurate (less bias) than the other 2D based alignment method. Furthermore, the freehand method is more automatic and the equipment is easy and cheap to make. The accuracy for both methods was evaluated and compared with a specially built wire phantom. The methods and equipment at this stage are both such that the calibration should be performed prior to surgery for each combination of ultrasound probe and position sensor. This approach is made possible by gluing an adapter to the probe casing, ensuring accurate repetitive attachment of the sensor to the probe, even through sterile draping.

Further research on this project includes testing the algorithm for other position sensor systems, particularly the new position sensing system that will be used in the SonoWand system (described in the Appendix in Paper G). In addition, evaluation of the accuracy with the new position sensor system is necessary. Making the method more automatic is also of interest, e.g. eliminating the need for an initial guess to the calibration transformation, adding an algorithm that lets the user know when the probe has been moved sufficiently, and a signal showing the accuracy of the iteration as more equations are added during scanning. The last addition means that the processing time for the edge detection procedure must be significantly reduced.

It would be of interest to make the equipment and software such that the entire probe calibration procedure could be performed in the operating room under sterile conditions. The equipment would then have to be altered significantly, e.g. sealing the phantom inside a chamber with an elastic surface membrane that is possible to sterilize. With such a method the surgeon would be able to attach any sensor to any probe and then perform the calibration and even be presented with numbers referring to the accuracy.

6.2. Wavelet-based edge detection

There are enormous amounts of publications on edge detection in general. Even for ultrasound images several different methods are being investigated. However, only very few research groups are using the wavelet transform approach in their work on edge detection for in vivo ultrasound images [68,126,127].

The edge detection method presented in Paper B was developed especially for the automatic probe calibration method presented in Paper A. The method was based on the continuous wavelet transform using one particular mother wavelet function, i.e. the Mexican hat wavelet. The method proved to be robust and useful for other purposes, e.g. as an initialization step in an active contour algorithm for segmentation, which is a current research project at the Ultrasound Department.

The method is subject to further studies, in particular with respect to decreasing the processing time, i.e. optimization of the algorithm, comparing the results for different mother wavelet functions, and testing the algorithm for in vivo images.

6.3. Instrument calibration and navigation

In order to know the location and orientation of the tip portion of a surgical instrument, we developed a calibration routine for instruments with an attached position sensor as we did for the combination of an ultrasound probe and a position sensor. A technique for performing this is described in Paper C. After having developed methods for calibrating both ultrasound probes and instruments, we set out to find a method for providing feedback to the surgeon about the spatial relation between the ultrasound image plane and the distal part of the surgical tool. Two parameters, the distance and angle between the image plane and the distal part of the tool, could easily be calculated from position measurements and the calibration transformations. These parameters were continuously updated in a graphic bar display using bright colors that could easily be seen and interpreted by the surgeon from just a quick glance during the insertion of the surgical tool.

Further development of this method involves optimizing the algorithm in order for it to run in conjunction with the main navigation software. Also, the calculation of the parameters for presentation involves reading the position data continuously for the position sensors on the probe and the tool. By averaging these position readings, a smoother transition in the updating of the graphic bar display is achieved. Furthermore, finding suitable colors (angle and distance information) for use in a navigation system is subject to further study. In addition, determining a geometry and location for the navigator bar display on the navigation monitor (or ultrasound scanner monitor) is subject to investigation.

As mentioned in the paper, further enhancement of this method might be the direct display of this information (angle and distance) onto a graphic representation of the tool in a stereoscopic image of the anatomical area of interest. To achieve this, the volume is first acquired without the instrument present. Next, as the tool is inserted into the body, a graphic representation is rendered in the stereoscopic view with information about the distance and angle projected onto the tool representation. This information is, as before, data on the spatial relation between the tool tip and the real time 2D image on the scanner monitor. The 3D volume is utilized for guidance during insertion of the tool, i.e. to provide an overview of the anatomical area of interest. An advantage with this approach is that we avoid reverberations from the instrument, which might obscure the view of important anatomical structures in the real time 2D view. Furthermore, this technique could prove to be an alternative, or low cost option, to real time 3D ultrasound, which might prove to be expensive for first generation systems.

6.4. Accuracy of ultrasound-based image-guided surgery

As mentioned earlier, a few groups have described and tried to estimate the accuracy of navigation systems for surgery [24,57,59,76,123,125,128,129,141]. An alternative to finding the overall clinical accuracy of a system is to estimate the accuracy of the separate steps or procedures involved, e.g. the accuracy of position sensing systems [29,81,82,95,122].

Paper A revealed that the probe calibration was a critical factor in the accuracy of the reconstructed volumes from 3D ultrasound scans. In Paper D, the overall accuracy and resolution of a prototype ultrasound-based navigation system was described. The results indicated that the overall clinical accuracy of MRI- and ultrasound-based navigation systems

is similar. However, ultrasound-based navigation eliminates the process of preoperative imaging and patient registration.

Further research should be invested in trying to decrease the errors and quantify the accuracy for the different procedures involved in ultrasound-guided surgery as explained in Paper D. Especially the bias that was found in the accuracy is subject to further study. This error might be the result of a systematic error in the probe calibration (as indicated in Paper A) or a time lag in the position tagging of the images when the ultrasound probe is moved during the scanning of the wire phantom. Nevertheless, we believe that this error could be reduced and the accuracy test should be performed again with the new position system that will be used in SonoWand® [52]. In this regard, it is of general interest to develop robust, accurate, and simple methodology (including equipment) for estimation of the overall accuracy of navigation systems.

As can be seen from the comparison of the results in Paper A and Paper D, there is quite a large difference between the reported values. This difference needs to be commented upon. We believe that the main reason for this difference is the fact that six volumes were used to evaluate the accuracy in Paper A, while only one volume was used in Paper D. In addition, different settings on the ultrasound scanner and reconstruction resolution may have contributed to some of the difference. Finally, in Paper D the observers pinpointed the wire crosses in space, while in Paper A the wire crosses were pinpointed only in image space (and also by different observers). Nevertheless, the difference may be the result of some systematic error not yet identified.

6.5. Ultrasound image restoration

The image restoration method presented in papers E and F attempted to remove blur and speckle noise in the ultrasound images through a new method based on Bayesian statistical methods, and thus bring out details that otherwise would have been lost. This was done to improve renderings for visualization, e.g. stereoscopic views for tool navigation in image-guided surgery. In such visualizations, we depend on images with high contrast between the object we want to visualize and its surroundings. We believe that the results of our work are very promising. In the case with the real image (Paper E), the results obtained were significantly better than the results produced by Wiener filtering. The speckle pattern was efficiently removed, while the important details were kept and no artificial structures were introduced.

Nevertheless, further development is needed to make this a practical method. It would be of great interest to test our image restoration method in combination with existing *psf* estimation methods in restoration of real ultrasound images. A few attempts on estimation of 2D *psf* exist [1,73,74,77,111,112,135-137], but this is a field that needs further study.

Another important factor for further research is the reduction of processing time, as the current method is far too slow. The processing time should at least be reduced down to practical levels for post-processing of the data, thus further optimization of the algorithm is necessary. The image restoration method is also dependent on an estimate of the *psf* of the ultrasound system. In paper F we demonstrate that the method is fairly insensitive to errors in the estimated *psf*. This is an important property of our method, as the *psf* can probably never be expected to be accurately known. Furthermore, the *psf* may vary considerably across the

imaging area and also be different from image to image due to tissue-dependent attenuation and distortion of the ultrasound pulse. Still, poor estimates of the *psf* affect the quality of the restorations. Thus, effort should be put into the estimation of the *psf* in order to achieve truly good restorations, especially in larger images.

7. References

1. Abeyratne U. R., Petropulu A. P. and Reid J. M. Higher order spectra based deconvolution of ultrasound images. IEEE Trans Ultrason Ferroelect Freq Contr:42:1064-1075:1995.

2. Adam D., Hareuveni O. and Sideman S. Semiautomated border tracking of cine echocardiographic ventricular images. IEEE Trans Med Imag MI-6:266-271:1987.

3. Angelsen B. A. J. Waves, signals and signal processing in medical ultrasonics. Trondheim, Norway, Dept. of Physiology and Biomedical Engineering, The Norwegian University of Science and Technology (NTNU), 1995.

4. Bamber J. C. and Tristam M. Diagnostic ultrasound. The Physics of Medical Imaging. S. Webb. Bristol, Great Britain, IOP Publishing Ltd, 319-388:1988.

5. Barry C. D., Allott C. P., John N. W., Mellor P. M., Arundel P. A., Thomson D. S. and Waterton J. C. Three-dimensional freehand ultrasound: image reconstruction and volume analysis. Ultrasound Med Biol 23:1209-1224:1997.

6. Bayes T. An essay towards solving a problem in the doctrine of chances. Phil Trans Roy Soc 53:370-418:1763.

7. Besag J. Spatial interaction and the statistical analysis of lattice systems. Journal of the Royal Statistical Society, Series B 36:1974.

8. Besag J. and Green P. Spatial statistics and Bayesian computation. Journal of the Royal Statistical Society, Series B 16:395-407:1993.

9. Besag J., Green P., Higdon D. and Mengersen K. Bayesian computation and stochastic systems. Statistical Science 10:3-66:1995.

10. Besag J. and Kooperberg C. On conditional and intrinsic autoregressions. Biometrika 82:733-746:1995.

11. Box G. E. P. and Muller M. A note on the generation of random normal variates. Ann Math Statist 29:610-611:1958.

12. Brooks S. P. and Roberts G. O. Assessing convergence of Markov chain Monte Carlo algorithms. Statistics and Computing 8:319-335:1998.

13. Bucholz R. D., Yeh D. D., Trobaugh J., McDurmott L. L., Sturm C. D., Baumann C., Henderson J. M., Levy A. and Kessman P. The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device. Lecture Notes in Computer Science Proceedings of the 1st Joint Conference on Computer Vision, Virtual Reality, and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery, CVRMed-MRCAS'97, Grenoble:1997.

14. Burch S. F. Comparison of SAFT and two-dimensional deconvolution methods for the improvement of resolution in ultrasonic B-scan images. Ultrasonics 25:259-266:1987.

15. Canny J. A computational approach to edge detection. IEEE Trans Pattern Anal Machine Intell PAMI-8:679-698:1986.

16. Carr J. Surface reconstruction in 3D medical imaging. Ph.D. thesis, University of Canterbury, Christchurch, New Zealand, 1996.

17. Chalana V., Linker D. T., Haynor D. R. and Kim Y. A multiple active contour model for cardiac boundary detection on echocardiographic sequences. IEEE Trans Med Imag 15:290-298:1996.

18. Charbonnier P., Blanc-Feraud L., Aubert G. and Barlaud M. Deterministic edge-preserving regularization in computed imaging. IEEE Trans Image Proc 6:298-311:1997.

19. Chassat F. and Lavallée S. Experimental protocol of accuracy evaluation of 6-D localizers for computer-integrated surgery: Application to four optical localizers. Lecture notes in Computer Science, Medical Image Computing and Computer-Assisted Intervention - MICCAI'98, Springer-Verlag, Berlin, Germany, 1496:277-284:1998.

20. Chou W. S., Wu C. M., Chen Y. C. and Hsieh K. S. Detecting myocardial boundaries of left ventricle from a single 2DE image. Pattern Recogn 23:799-806:1990.

21. Chu C. H., Delp E. J. and Buda A. J. Detecting left ventricular endocardial and epicardial boundaries by digital two-dimensional echocardiography. IEEE Trans Med Imag 7:81-90:1988.

22. Chui C. K. An introduction to Wavelets. Boston, USA, Academic Press, 1992.

23. Cohen F. S., Georgiou G. and Halpern E. J. WOLD decomposition of the backscatter echo in ultrasound images of soft tissue organs. IEEE Trans Ultrason Ferroelect Freq Contr 44:460-472:1997.

24. Comeau R. M., Sadikot A. F., Fenster A. and Peters T. M. Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. Med Phys 27:787-800:2000.

25. Craig J. Introduction to robotics: Mechanics and control. Reading, Massachusetts, USA, Addison-Wesley, 1989.

26. Czerwinski R. N., Jones D. L. and O'Brien W. D. Jr. An approach to boundary detection in ultrasound imaging. IEEE Ultrasonics Symp:951-955:1993.

27. Daubechies I. Ten lectures on Wavelets. Philadelphia, Pennsylvania, Society for Industrial and Applied Mathematics, 1992.

28. Demigny D., Lorca F. G. and Kessal L. Evaluation of edge detectors performances with a discrete expression of Canny's criteria. IEEE Int Conf Image Proc:169-172:1995.

29. Detmer P. R., Bashein G., Hodges T., Beach K. W., Filer E. P., Burns D. H. and Strandness D. E. J. 3D ultrasonic image feature localization based on magnetic scanhead tracking: In vitro calibration and validation. Ultrasound Med Biol 20:923-936:1994.

30. Dias J. M. B. and Leitao J. M. N. Wall position and thickness estimation from sequences of echocardiographic images. IEEE Trans Med Imag 15:25-38:1996.

31. Diethrich E. B. Endovascular treatment of abdominal aortic occlusive disease: The impact of stents and intravascular ultrasound imaging. Eur J Vasc Surg 7:228-236:1993.

32. Drebin R., Carpenter L. and Hanrahan P. Volume rendering. Proceedings of the SIGGRAPH '88: ACM Computer Graphics, Atlanta, GA, USA, 22:65-74:1988.

33. Duck F. A. Physical properties of tissue. A comprehensive reference book. London, UK, Academic Press Ltd., 1990.

34. Edwards C. H. and Penny D. E. Calculus and analytic geometry. Englewood Cliffs, N.J., USA, Prentice-Hall, 1990.

35. Elvins T. T. A survey of algorithms for volume visualisation. Computer Graphics 26:194-201:1992.

36. Farge M., Kevlahan N., Perrier P. and Goirand E. Wavelets and turbulence. Proc IEEE, 84:639-669:1996.

37. Feng J., Lin W.-C. and Chen C.-T. Epicardial boundary detection using fuzzy reasoning. IEEE Trans Med Imag 10:187-199:1991.

38. Fischler M. A. and Bolles R. C. Random Sample Consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM 24:381-395:1981.

39. Foley J. D., vanDam A., Feiner S. K. and Hughes J. F. Computer graphics: Principles and practice. Reading, Massachusetts, USA, Addison-Wesley, 1990.

40. Friedland N. and Adam D. Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated annealing. IEEE Trans Med Imag 8:344-353:1989.

41. Gabor D. Theory of communication. J IEE 93:429-457:1946.

42. Geman D. and Reynolds G. Constrained restoration and the recovery of discontinuities. IEEE Trans Pattern Anal Machine Intell 14:367-383:1992.

43. Geman D., Reynolds G. and Yang C. Stochastic algorithms for restricted image spaces and experiments in deblurring. Markov Random Fields: Theory and Application. R. Chellappa and A. Jain. New York, USA, Academic Press:1993.

44. Geman S. and Geman D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. IEEE Trans Pattern Anal Machine Intell PAMI-6:721-741:1984.

45. Geman S. and McClure D. E. Statistical methods for tomographic image reconstruction. 46th Sess Int Stat Inst Bulletin ISI, 52:1987.

46. Georgiou G. and Cohen F. S. Statistical characterization of diffuse scattering in ultrasound images. IEEE Trans Ultrason Ferroelect Freq Contr 45:57-64:1998.

47. Geyer C. and Thompson E. A. Constrained Monte Carlo maximum likelihood for dependent data. Journal of the Royal Statistical Society, Series B 54:657-699:1992.

48. Gilks W. R., Richardson S. and Spiegelhalter D. J. Markov chain Monte Carlo in practice. London, Chapman & Hall, 1996.

49. Gonzales R. C. and Woods R. E. Image segmentation. Digital Image Processing. Massachusetts, USA, Addison-Wesley, 413-481:1993.

50. Goodman J. W. Statistical properties of laser speckle patterns. Laser Speckle and Related Phenomena. J. C. Dainty. Berlin, Germany, Springer Verlag:1975.

51. Green P. J. Reversible jump MCMC computation and Bayesian model determination. Biometrika 82:711-732:1995.

52. Grønningsæter Å., Kleven A., Ommedal S., Aarseth T. E., Lie T., Lindseth F., Langø T. and Unsgård G. SonoWand, an ultrasound-based neuronavigation system. Accepted for publication in Neurosurgery:2000.

53. Grønningsæter Å., Lie T., Kleven A., Mørland T., Langø T., Unsgård G., Myhre H. O. and Mårvik R. Initial experience with stereoscopic visualization of 3D ultrasound data in surgery. Surg Endosc (in press):1999.

54. Grønningsæter Å., Lindseth F., Langø T. and Unsgård G. Accuracy assessment of a prototype ultrasound-based neuronavigation system. Submitted to Neurosurgery:2000.

55. Hagan M. T., Demuth H. B. and Beale M. Neural network design. Boston, Massachusetts, USA, PWS Publishing Co., 1996.

56. Hastings W. K. Monte Carlo simulation methods using Markov chains and their applications. Biometrika 57:97-109:1970.

57. Hata N., Dohi T., Iseki H. and Takakura K. Development of frameless and armless stereotactic neuronavigation system with ultrasonographic registration. Neurosurgery 41:608-613:1997.

58. Heath M., Sarkar S., Sanocki T. and Bowyer K. Comparison of edge detectors: A methodology and initial study. IEEE Comp Soc Conf CompVision Pattern Recogn:143-148:1996.

59. Helm P. A. and Eckel T. S. Accuracy of registration methods in frameless stereotaxis. Comput Aided Surg 3:51-56:1998.

60. Hernes T. A. N., Lindseth F., Ommedal S., Gronningsaeter A., Lie T., Langø T., Kleven A. and Mårvik R. 3D visualization techniques in minimally invasive ultrasound guided surgery. EAES, Linz, Germany:1999.

61. Higdon D. M. Auxiliary variable methods for Markov chain Monte Carlo with applications. J Am Statistical Association 93:585-595:1998.

62. Higdon D. M., Bowsher J. E., Johnson V. E., Turkington T. G., Gilland D. R. and Jaszczak R. J. Fully Bayesian estimation of Gibbs hyperparameters for emission computed tomography data. IEEE Trans Med Imaging 14:516-526:1997.

63. Hirschberg H. and Unsgaard G. Incorporation of ultrasonic imaging in an optically coupled frameless stereotactic system. Acta Neurochir Suppl (Wien) 68:75-80:1997.

64. Hokland J. H. and Kelly P. A. Markov models of specular and diffuse scattering in restoration of medical ultrasound images. IEEE Trans Ultrason Ferroelect Freq Contr 43:660-669:1996.

65. Hough P. V. C. Methods and means for recognizing complex patterns. USA. Patent 3,069,654. 1962.

66. Hudgins L. Wavelet analysis of atmospheric turbulence. Ph.D. thesis, University of California, Irvine, USA, 1992.

67. Hughes S. W., D'Arcy T. J., Maxwell D. J., Chiu W., Milner A., Saunders J. E. and Sheppard R. J. Volume estimation from multiplanar 2D ultrasound images using a remote electromagnetic position and orientation sensor. Ultrasound Med Biol 22:561-572:1996.

68. Hunter I. A., Soraghan J. J. and McDonagh T. Fully automatic left ventricular boundary extraction in echocardiographic images. Proc Computers in Cardiology, Vienna:741-744:1995.

69. Hurn M. and Jennison C. Extension of Geman and Reynolds' approach to constrained restoration and the recovery of discontinuities. IEEE Trans Pattern Anal Machine Intell 18:657-662:1996.

70. Husby O. K., Lie T., Langø T., Hokland J. and Rue H. Bayesian 2D restoration: A model for diffuse ultrasound scattering. Accepted for publication in IEEE Trans Ultrason Ferroelect Freq Contr:2000.

71. Iraca D., Landini L. and Verrazzani L. Power spectrum equalization for ultrasonic image restoration. IEEE Trans Ultrason Ferroelect Freq Contr 36:216-222:1989.

72. Jain A. K. Fundamentals of digital image processing. Englewood Cliffs, N.J., Prentice-Hall, 1989.

73. Jensen J. A. Estimation of pulses in ultrasound B-scan images. IEEE Trans Biomed Eng 10:164-172:1991.

74. Jensen J. A. and Leeman S. Nonparametric estimation of ultrasound pulses. IEEE Trans Biomed Eng 41:929-936:1994.

75. Kaiser G. A friendly guide to wavelets. Boston, Birkhäuser, 1994.

76. Kall B., Goerss S. J., Stiving S. O., Davis D. H. and Kelly P. J. Quantitative analysis of a noninvasive stereotactic image registration technique. Stereotact Funct Neurosurg 66:69-74:1996.

77. Kåresen K. F. and Bolviken E. Blind deconvolution of ultrasonic traces accounting for pulse variance. IEEE Trans Ultrason Ferroelect Freq Contr 46:564-573:1999.

78. Kaspersen J. H. A study of coherent structures using Wavelet analysis. Ph.D. thesis, Norwegian University of Science and Technology (NTNU), Trondheim, Norway, 1996.

79. Kaspersen J. H., Langø T. and Lindseth F. Wavelet-based edge detection in ultrasound images. Accepted for publication in J Ultrasound Med Biol:2000.

80. Kaus M., Steinmeier R., Sporer T., Ganslandt O. and Fahlbusch R. Technical accuracy of a neuronavigation system measured with a high-precision mechanical micromanipulator. Neurosurgery 41:1431-1437:1997.

81. Khadem R., Yeh C. C., Sadeghi-Tehrani M., Bax M. R., Johnson J. A., Welch J. N., Wilkinson E. P. and Shahidi R. Comparative tracking error analysis of five different optical tracking systems. Comp Aided Surgery 5:98-107:2000.

82. King D. L., King D. L. and Shao M. Y.-C. Evaluation of in vitro measurement accuracy of a three- dimensional ultrasound scanner. Ultrasound Med 10:77-82:1991.

83. Klinger J. W., Vaughan C. L., Fraker T. D. and Andrews L. T. Segmentation of echocardiographic images using mathematical morphology. IEEE Trans Biomed Eng 35:925-934:1988.

84. Koivukangas J., Louhisalmi Y., Alakuijala J. and Oikarinen J. Ultrasound-controlled neuronavigator-guided brain surgery. J Neurosurg 79:36-42:1993.

85. Koivukangas J., Ylitalo E., Alasaarela E. and Tauriainen A. Three-dimensional ultrasound imaging of brain for neurosurgery. Ann Clin Res 18 Suppl. 47:65-72:1986.

86. Langø T., Lie T., Husby O. K. and Hokland J. Bayesian 2D deconvolution: Effect or non-stationary ultrasound pulses. Accepted for publication in IEEE Trans Ultrason Ferroelect Freq Contr:2000.

87. Langø T., Lie T., Lindseth F., Kleven A. and Gronningsaeter A. Probe calibration. UK. Patent Application 31.15.69543. 1999.

88. Langø T., Lindseth F., Gronningsaeter A., Ommedal S. and Unsgård G. Instrument navigation in ultrasound guided surgery. CARS 2000 (ISCAS), San Francisco, USA:2000.

89. Langø T., Lindseth F., Kaspersen J. H. and Gronningsaeter A. Novel probe calibration methods for 3D freehand ultrasound. Submitted to Comput Aided Surg:2000.

90. Langø T., Lindseth F., Ommedal S., Kleven A., Unsgård G. and Gronningsaeter A. Tool navigator. UK. Patent Application 31.15.69544. 1999.

91. Langø T., Mørland T., Lie T., Kleven A., Sæther O. D., Myhre H. O., Lundbom J. and Gronningsaether A. Virtual reality tour of an artery by stereoscopic viewing of 3D ultrasound images. Oral presentation at the First International Conference on Ultrasound in Vascular Disease, Malmö, Sweden :1996.

92. Langø T., Ommedal S., Unsgard G. and Gronningsaeter A. Navigation of probes and surgical instruments in ultrasound guided surgery. Submitted to Comput Aided Surg:2000.

93. Lawson C. L. and Hanson R. J. Solving least squares problems. Englewood Cliffs, NJ, USA, Prentice-Hall, 1974.

94. Leotta D. F., Detmer P. R., Gilja O. H., Jong J.-M., Martin R. W., Primozich J. F., Beach K. W. and Strandness E. D. Three-dimensional ultrasound imaging using multiple magnetic tracking systems and miniature magnetic sensors. IEEE Ultraso Symp, Seattle, USA, 2:1415-1418:1995.

95. Leotta D. F., Detmer P. R. and Martin R. W. Performance of a miniature magnetic position sensor for three-dimensional ultrasound imaging. Ultrasound Med Biol 23:597-609:1997.

96. Levoy M. Display of surfaces from volume data. IEEE Comp Graph Appl 8:29-37:1988.

97. Liandrat J. and Moret-Bailly F. The wavelet transform: some applications to fluid dynamics and turbulence. Eur J Mech, B/Fluids 9:1-19:1990.

98. Linker D. T., Pearlman A. S., Lewellen T. K., Huntsman L. H. and Moritz W. E. Automated endocardial definition of 2-D echocardiograms: A comparison of four standard edge detectors and improved thresholding techniques. Proc Computers in Cardiology:395-398:1982.

99. Loupas T., Pye S. D. and McDicken W. N. Deconvolution in medical ultrasonics: Practical considerations. Phys Med Biol 34:1691-1700:1989.

100. Maciunas R. J., Galloway R. L. and Latimer J. W. The application accuracy of stereotactic frames. Neurosurgery 35, No. 4:682-695:1994.

101. Mailloux G. E., Langlois F., Simard P. Y. and Bertrand M. Restoration of the velocity field of the heart from two-dimensional echocardiograms. IEEE Trans Med Imag 8:143-153:1989.

102. Mallat S. and Hwang W. L. Singularity detection and processing with wavelets. IEEE Trans Inform Theory 38:617-643:1992.

103. Mardia K. V., Kent J. T. and Bibby J. M. Multivariate analysis, Academic Press, 1979.

104. Mengersen K. L., Robert C. P. and Guihenneuc-Jouyaux C. MCMC convergence diagnosics: A review. Paris, France, CREST, INSEE:1998.

105. Meunier J., Bertrand M., Mailloux G. E. and Petitclerc R. Local myocardial deformation computed from speckle motion. Proc Comput Cardiology, Washington DC:133-136:1988.

106. More J. J. The Levenberg-Marquardt algorithm: Implementation and theory. Numerical Analysis, Lecture Notes in Mathematics. A. Watson, Springer Verlag, 105-116:1977.

107. Mørland T., Langø T., Grønningsæter Å., Lundbom J. and Sæther O. D. Ultrasound imaging during endovascular treatment of AAA. Oral presentation at the First International Conference on Ultrasound in Vascular Disease, Malmö, Sweden:1996.

108. Ødegård L. Phase aberration correction in medical ultrasound imaging. Ph.D. thesis, Norwegian University of Science and Technology, Trondheim, 1995.

109. Ødegård L., Halvorsen E., Ystad B., Torp H. G. and Angelsen B. Delay and amplitude focusing through the body wall; A simulation study. IEEE Ultrason Symp, 2:1411-1414:1996.

110. Ohbuchi R., Chen D. and Fuchs H. Incremental volume reconstruction and rendering for 3-D ultrasound imaging. SPIE - The International Society for Optical Engineering, Chapel Hill, NC, USA, 1808:312-323:1992.

111. Petropulu A. P. and Abeyratne U. R. System reconstruction from higher order spectra slices. IEEE Trans Sign Proc 45:2241-2251:1997.

112. Pozidis H. and Petropulu A. P. System reconstruction based on selected regions of discretized higher order spectra. IEEE Trans Sign Proc 46:3360-3377:1998.

113. Prager R. W., Rohling R. N., Gee A. H. and Berman L. Automatic calibration of 3-D free-hand ultrasound. Cambridge, UK, Cambridge University Department of Engineering:1997.

114. Prager R. W., Rohling R. N., Gee A. H. and Berman L. Rapid calibration for 3-d freehand ultrasound. Ultrasound Med Biol 24:855-869:1998.

115. Rabben S. I., Torp A. H., Støylen A., Slørdahl S., Bjørnstad K., Haugen B. O. and Angelsen B. Semiautomatic contour detection in ultrasound M-mode images. Ultras Med Biol 26:287-296:2000.

116. Rao R. M. and Bopardikar A. S. Wavelet transforms. Introduction to theory and applications. Massachusetts, USA, Addison-Wesley, 1998.

117. Rioul O. and Vetterli M. Wavelets and signal processing. IEEE Sign Proc Magazine 84:14-38:1991.

118. Ripley B. D. Stochastic simulation. New York, USA, Wiley, 1987.

119. Robert C. P. The Bayesian choice: A decision-theoretic motivation. New York, USA, Springer-Verlag, 1994.

120. Robert C. P. Convergence control methods for MCMC algorithms. Statistical Science 10:231-253:1995.

121. Roberts G. O., Gelman A. and Gilks W. R. Weak convergence and optimal scaling of random walk metropolis algorithms. Ann Appl Prob 7:110-120:1997.

122. Rohling R. N. and Peters T. M. Comparison of relative accuracy between an optical and a mechanical position tracker for image guided neurosurgery. J Image Guided Surgery 1:30-34:1994.

123. Sandemann D. R., Patel N., Chandler C., Nelson R., Coakham H. B. and Griffith H. B. Advances in image-directed neurosurgery: preliminary experience with ISG Viewing Wand compared with the Leksell G frame. British J Neurosurgery 8:529-544:1994.

124. Santis A. D. and Sinisgalli C. A Bayesian approach to edge detection in noisy images. IEEE Trans Circuits Syst I: Fund Theory Appl 46:686-699:1999.

125. Schaller C., Meyer B., van Roost D. and Schramm J. Image guided microsurgery with a semifreehand neuronavigational device. Comput Aided Surg 2:162-171:1997.

126. Setarehdan S. K. and Soraghan J. J. Cardiac left ventricular volume changes assessment by long axis echocardiographical image processing. IEE Proc Vis Image Signal Process, 145:203-212:1998.

127. Setarehdan S. K. and Soraghan J. J. Automatic cardiac LV boundary detection and tracking using hybrid fuzzy temporal and fuzzy multiscale edge detection. IEEE Trans Biomed Eng 46:1364-1378:1999.

128. Sipos E. P., Tebo S. A., Zinreich S. J., Long D. M. and Brem H. In vivo accuracy testing and clinical experience with the ISG viewing wand. Neurosurgery 39:194-204:1996.

129. Smith K. R., Frank K. J. and Bucholz R. D. The NeuroStationTM - A highly accurate, minimally invasive solution to frameless stereotactic neurosurgery. Comput Med Imaging Graph 18:247-256:1994.

130. State A., Chen D. T., Tector T., Brandt A., Chen H., Ohbuchi R., Bajura M. and Fuchs M. Case study: Observing a volume rendered fetus within a pregnant patient. IEEE Visualization:364-368:1994.

131. Steen E. and Olstad B. Volume rendering of 3D medical ultrasound data using direct feature mapping. IEEE Trans Med Imaging 13:517-525:1994.

132. Sumanaweera T. S., Adler J. R., Napel S. and Glover G. H. Characterization of spatial distortion in magnetic resonance imaging and its implications for stereotactic surgery. Neurosurgery 35:696-704:1994.

133. Sun Y. N. and Tsai C. T. Segmentation of echocardiograms using a neural network. Microprocessing Microprogram 35:791-798:1992.

134. Tanner M. A. Tools for statistical inference. Methods for the exploration of posterior distributions and likelihood functions. New York, USA, Springer-Verlag, 1993.

135. Taxt T. Restoration of medical ultrasound images using two-dimensional homomorphic deconvolution. IEEE Trans Ultrason Ferroelect Freq Contr 42:543-554:1995.

136. Taxt T. Comparison of cepstrum-based methods for radial blind deconvolution of ultrasound images. IEEE Trans Ultrason Ferroelect Freq Contr 44:666-674:1997.

137. Taxt T. and Frolova G. V. Noise robust one-dimensional blind deconvolution of medical ultrasound images. IEEE Trans Ultrason Ferroelect Freq Contr 46:291-299:1999.

138. Thompson A. M., Brown J. C., Kay J. W. and Titterington D. M. A study of methods of choosing the smoothing parameter in image restoration by regularization. IEEE Trans Pattern Anal Machine Intell 13:326-339:1991.

139. Trobaugh J. W., Richard W. D., Smith K. R. and Bucholz R. D. Frameless stereotactic ultrasonography: Method and applications. Comput Med Imaging Graph 18:235-246:1994.

140. Udupa H. K. and Odhner D. Shell rendering. IEEE Comp Graph Appl 13:58-67:1993.

141. Villalobos H. and Germano I. Clinical evaluation of multimodality registration in frameless stereotaxy. Comput Aided Surg 4:45-49:1999.

142. Wagner R. F., Insana M. F. and Brown D. G. Statistical properties of radio-frequency and envelope-detected signals with application to medical ultrasound. J Opt Soc Am 4:910-922:1987.

143. Winkler G. Image analysis, random fields and dynamic Monte Carlo methods. Berlin, Springer-Verlag, 1995.

144. Xin Y., Boekhaerts P., Bister M., Taeymans Y. and Cornelis J. Left ventricular boundary detection in echocardiographical images using artificial neural networks. 6th European Signal Processing Conf 1199-1202:1992.

145. Xu C. and Prince J. L. Snakes, shapes, and gradient vector flow. IEEE Trans Image Proc 7:359-369:1998.

146. Zong X., Laine A. F. and Geiser E. A. Speckle reduction and contrast enhancement of echocardiograms via multiscale nonlinear processing. IEEE Trans Med Imaging 17:532-540:1998.

Paper A

Novel Probe Calibration Methods for 3D Freehand Ultrasound

Thomas Langø[1], Frank Lindseth[1], Atle Kleven[2] and Aage Gronningsaeter[2]

[1] SINTEF Unimed, Ultrasound, 7465 Trondheim, Norway
[2] MISON AS, 7489 Trondheim, Norway

ABSTRACT

The process of 3D freehand ultrasound acquisition consists of two steps: scanning and reconstruction. Prior to scanning, a position sensor is attached to the probe for tagging each image with its position and orientation in space. A crucial step prior to reconstruction and visualization is calibration of the ultrasound probe with the attached position sensor, i.e. determining the position and orientation of the scan plane relative to the position sensor. We propose a method for probe calibration that involves the use of an easy to build and easy to use calibration phantom, comprising a water tank and a thin planar nylon mesh membrane. The submerged membrane is scanned by freehand, covering all degrees of freedom for movement in space, and the following calculations are performed automatically, including the edge detections. The technique is denoted the *Freehand Membrane Scan Method*. The accuracy of the method is assessed by 3D scanning of a phantom with 27 precision made thread crosses in known locations, then compared to another method developed earlier by our research group, the *Membrane Alignment Method*.

INTRODUCTION

The cost and flexibility of freehand 3D ultrasound imaging using conventional 2D probes and a position sensing system ensures that it will remain a popular choice as a medical imaging modality. The process of freehand 3D imaging consists of two steps: scanning and reconstruction. Prior to the acquisition of the ultrasound images, a position sensor is attached to the ultrasound probe. Measurements of the position and orientation of this sensor enable us to track the position and orientation of the ultrasound image in space. The set of images and their relative positions from the scan process are then used to fill a voxel array, and finally the array for visualization purposes. A crucial step in this process is the calibration of the ultrasound probe to which a position sensor is attached. Calibration of an ultrasound probe signifies the procedure that is performed to establish the transformation matrix that describes the position and orientation of the ultrasound image plane relative to the position sensor attached to the probe. In other words, the calibration determines the transformation needed to continuously track the spatial location of the ultrasound image in the coordinate system of the position sensing system, so-called global coordinates. The transformation itself is a 2D matrix defined by three translation offsets, and three angle offsets. Finding this transformation matrix is one of the most challenging and critical tasks regarding accuracy in ultrasound imaging based on acquisitions of freehand 3D ultrasound volumes. An accurate calculation of this matrix is of crucial importance for a reconstruction that preserves true anatomical shape and size in freehand ultrasound 3D scanning. The calibration parameters can be crudely estimated by external measurements of the probe case and sensor. However, this estimate will not necessarily be exact since the origin of the ultrasound image system is located inside the housing of the probe and may vary slightly from one probe to another. Furthermore, the position of the center transducer element (origin) and the direction for pulse transmission may vary due to variations in the production of the probe. In addition, the origin for the coils in magnetic sensors is not well defined. This indicates that a better approach to finding the calibration matrix would be to image some phantom with known physical properties and dimensions. In this way we can make sure that we are using the correct position of the ultrasound image plane origin. Measurements and calculated parameters from a 3D scan of the phantom can then be used to perform the calculation of the transformation matrix values.

Existing calibration methods
Few publications have been devoted solely to the investigation of probe calibration.[25, 26] Other publications containing descriptions of calibration methods are mainly concerned with volume measurements and reconstruction issues,[15, 18, 24] compound ultrasound scanning,[2] accuracy of ultrasound scanning, and accuracy of position sensing.[7, 21, 22] In these publications, probe calibration is merely a necessity to perform the main objective of the work. A probe calibration is required regardless of the type of position sensor equipment one uses in reconstructing volumes obtained by freehand scanning. Most of the existing methods for probe calibration are limited in accuracy, ease of use, and performance time, and some need complicated equipment such as precision-made calibration phantoms.

Cross-wire phantom
This method consists of scanning the intersection of two crossed wires, or some other small object, such as a suspended bead or a pinhead, in a water bath from different directions and manually detecting the crossing point or small object in each image afterwards. The accuracy of this method depends to a large degree on the accuracy of the manual detection in the

images. Furthermore, the need to manually detect the points of interest makes it difficult to automate this method. Several publications contain descriptions of this calibration method.[2, 7, 18, 21, 22, 25-27]

Three-wire phantom

This technique involves scanning three wires that are submerged in a water bath and crossed in such a manner that they constitute a right-handed orthogonal coordinate system. As for the cross-wire technique, the wires have to be marked manually in each image since automatic detection of points in ultrasound images has proven to be very difficult due to the presence of different kinds of noise. In addition to the accuracy of the manual detection in the images, the accuracy of this method depends on whether the wires are precisely orthogonal and on the straightness of the wires. The advantage of this method over the cross-wire technique is mainly that it is easier to scan the length of a wire than it is to keep the image focused at a center point from various directions. Nevertheless, one must keep track of which wire is being scanned, i.e. the x-, y-, or z-axis, to obtain a correct set of equations. A more detailed description of this method may be found in other publications.[4, 26]

Single-wall Method

The main idea with this method is to let the bottom of the water tank constitute the *xy*-plane of the reconstruction system. The bottom of a water bath is imaged from different directions and the line in each image defining the bottom is detected. Points from this edge are used as definitions of the phantom in making the set of equations that are to be solved for the unknowns (see section on *Solving system of equations* in *Materials and Methods*). One problem with this method is the specular reflection, that causes a low intensity to be returned to the transducer at oblique scan angles. This can to a certain degree be compensated for by roughening the surface of the bottom. Another problem is that it is difficult to determine the true position of the floor in the images based on the reflected signal intensity. This is due to the strong reverberations from the bottom, which appear as a tail in the reflected intensity signal. These drawbacks will be further discussed later. See Prager et al.[25, 26] for a more detailed description of the *Single-wall Method.*

Cambridge Phantom Method

The first step of the *Cambridge Phantom Method* is to attach an ultrasound probe in a clamp, which is placed over a thin bar, which in turn is attached between two wheels. The probe is attached in such a way that a clear image of the top of the bar can be seen on the scanner. Prager et al.[25, 26] also describe a separate technique and equipment to use for making sure that the scan plane and bar are aligned. Next, the phantom is immersed in a water bath, the clamp is placed over the bar, and the bar is scanned with the probe from all possible angles, subject to the constraints imposed by the clamp and the bar-wheels combination. The edges representing the bar in the images are detected automatically. If the probe is aligned perfectly so that the scan plane and bar are coplanar, the image should contain one single thin line. The calculations after the scanning are similar to the procedure outlined later. The *Cambridge Phantom Method* depends on a precision-made phantom. In addition, the alignment of the probe scan plane with the bar is subject to variation and the edge detection method requires several input parameters. See Prager et al.[25, 26] for a more detailed description of this method.

THEORY

A two-dimensional, four by four matrix defines the mathematical transformation we are seeking:[6, 11]

$$M_{j \leftarrow i} = T(t_x, t_y, t_z) \cdot R_z(\alpha) \cdot R_y(\beta) \cdot R_x(\gamma)$$

$$= \begin{bmatrix} 1 & 0 & 0 & t_x \\ 0 & 1 & 0 & t_y \\ 0 & 0 & 1 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \cos\theta & -\sin\theta & 0 & 0 \\ \sin\theta & \cos\theta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \cos\theta & 0 & \sin\theta & 0 \\ 0 & 1 & 0 & 0 \\ -\sin\theta & 0 & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$= \begin{bmatrix} \cos\alpha \cdot \cos\beta & \cos\alpha \cdot \sin\beta \cdot \sin\gamma - \sin\alpha \cdot \cos\gamma & \cos\alpha \cdot \sin\beta \cdot \cos\gamma + \sin\alpha \cdot \sin\gamma & t_x \\ \sin\alpha \cdot \cos\beta & \sin\alpha \cdot \sin\beta \cdot \sin\gamma + \cos\alpha \cdot \cos\gamma & \sin\alpha \cdot \sin\beta \cdot \cos\gamma - \cos\alpha \cdot \sin\gamma & t_y \\ -\sin\beta & \cos\beta \cdot \sin\gamma & \cos\beta \cdot \cos\gamma & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (1)$$

Index j represents the system to which we are transforming position coordinates with index i. The transformation is written as homogeneous coordinates[11] to handle translation, rotation, and scaling with one single matrix multiplication. Eq. 1 can be deduced by multiplying the four matrices: T; translation in 3D space, R_x; rotation about the x-axis, R_y; rotation about y-axis, and R_z; rotation about z-axis, respectively. The rotation between the two coordinate systems i and j is effectuated by first rotating through γ about the x-axis, then through β about the y-axis, and finally through α about the z-axis. The translation is given by the parameters t_x, t_y, and t_z. The particular coordinate systems of interest are the ultrasound image coordinate system denoted by index us, the local coordinate system of the position sensor attached to the probe indicated by index lo, the global coordinate system referred to by index gl, and the membrane system signified by index m. The different systems along with the definitions of the matrix transformations between them are summarized in Fig. 1. For simplicity, in this paper we use the term sensor to refer to both passive receivers, such as reflecting spheres, and active sensors, such as light emitting diodes.
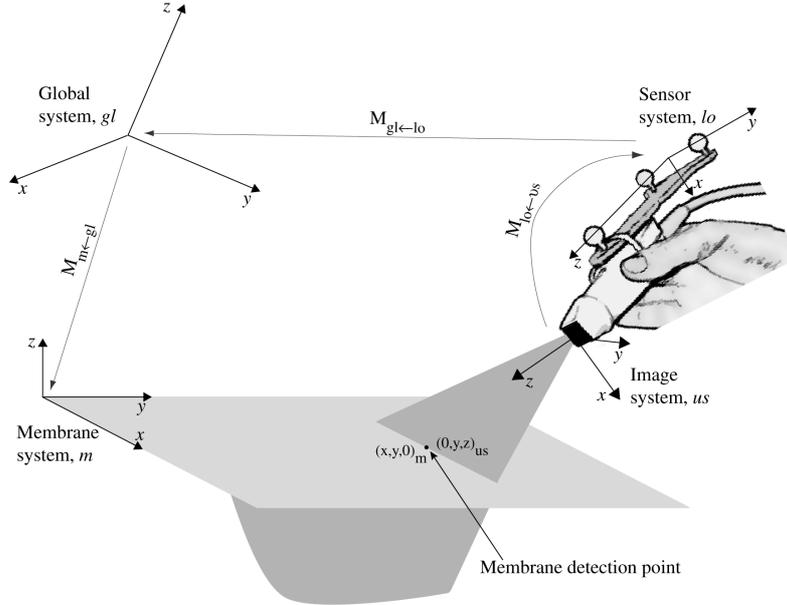
Fig. 1. An illustration of the different coordinate systems used throughout the paper. The multiplication of any of the matrices $M_{j \leftarrow i}$ with a position vector $[x_i, y_i, z_i]$ results in the definition of that position in the system j. The probe calibration matrix is denoted $M_{lo \leftarrow us}$, while $M_{gl \leftarrow lo}$ is the transformation measured by the position sensing system.

During the reconstruction process, all the ultrasound images have to be located correctly in the reconstructed volume. In total, the transformation of image plane pixel positions $[0, y, z]$ (Fig. 1) to the membrane system is handled by a single matrix multiplication. First, the pixel coordinate is transformed to the sensor coordinate system by the calibration matrix. Second, the position system reading for that image is used for transformation from the local to the global coordinate system, often referred to as the reference system. Finally, the transformation to the membrane system is calculated. The total transformation can thus be given by the following multiplication:

$$
\begin{bmatrix} x_m \\ y_m \\ z_m \\ 1 \end{bmatrix} = M_{m \leftarrow gl} \cdot M_{gl \leftarrow lo} \cdot M_{lo \leftarrow us} \cdot \begin{bmatrix} x_{us} \\ y_{us} \\ z_{us} \\ 1 \end{bmatrix} \tag{2}
$$

We acquire digital ultrasound data directly from the ultrasound scanner (via an Ethernet link) and perform the scan conversion after edge detection (see *Materials and Methods*). To ensure correct 2D scan conversion[1] it is important that we use the correct velocity of sound in water. In this manner, we achieve correct geometric images and distances and coordinates may be extracted directly from the scans in millimeters. From the phantom set-up and coordinate definitions, one or more of the parameters x_m, y_m, or z_m will be known and the corresponding row(s) from Eq. 2 yield equation(s) incorporating the unknown parameters and measurements. For example, for the cross-wire phantom described above, all the coordinate parameters on the left side are equal to zero since we are imaging the origin of the wire system:

$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = M_{m \leftarrow gl} \cdot M_{gl \leftarrow lo} \cdot M_{lo \leftarrow us} \cdot \begin{bmatrix} 0 \\ y_{us} \\ z_{us} \\ 1 \end{bmatrix} \tag{3}$$

The parameter x_{us} is equal to zero due to our definition of the image coordinate system (Fig. 1). For the system in Fig. 1 only z_m will be equal to zero on the left hand side of Eq. 3 since the membrane constitutes the xy-plane (compare Eq. 3 to Eq. 14). By performing m scans from different directions, we achieve a system of $\Lambda \cdot m$ non-linear homogeneous equations, where Λ is either equal to 1, 2, or 3, depending on the specifics of the phantom and coordinate definitions. For the case in Eq. 3 $\Lambda=3$. The equations may be written as

$$F(\vartheta,\varphi) = 0 \tag{4}$$

where ϑ represents the measured parameters while φ are the unknowns. Hence, we are looking for the set of parameters that yields $F \rightarrow 0$. The system of equations, which is over-determined by making sure Λm is greater than the number of unknowns n, can be solved with several different iterative methods, for example the robust Levenberg-Marquardt algorithm.[23] Briefly explained, at iteration k, an update $\Delta\varphi$ to the current estimate of φ_k is obtained from the first order Taylor expansion of Eq. 4

$$F(\vartheta,\varphi) \approx F(\vartheta,\varphi_k) + \frac{\partial F(\vartheta,\varphi_k)}{\partial \varphi}(\varphi - \varphi_k) = 0 \tag{5}$$

which further can be written as

$$\Delta F = -F(\vartheta,\varphi_k) = \frac{\partial F(\vartheta,\varphi_k)}{\partial \varphi}(\varphi - \varphi_k) = J(\varphi - \varphi_k) = J\Delta\varphi \tag{6}$$

where J represents the Jacobian, i.e. the gradient matrix of F. Each step of the Levenberg-Marquardt algorithm yields the updated parameters

$$\varphi_{k+1} = \varphi_k + \left(J^T J + \varepsilon I\right)^{-1} J^T \Delta F \tag{7}$$

where I is the identity matrix and upper index T represents the transpose of a matrix. ΔF and J are evaluated at each step of the algorithm for the current estimate φ_k, and the process continues until the corrections $\Delta\varphi$ are sufficiently small.

MATERIALS AND METHODS

All calibrations were performed with a 4-8 MHz phased array probe on a high-end digital scanner, the System FiVe (GE Vingmed Ultrasound, Horten, Norway), and an optical position sensor system, the Vector Vision (BrainLab, Munich, Germany). The sensor tracking error is claimed to be less than 1 mm for the position of one sphere. We calibrated the probe with a particular sensor in the laboratory, using a special adapter glued on to the probe casing that ensured repetitive and precise attachment between the two. This meant that we only had to

calibrate each combination of probe and sensor once. See Fig. 1 for an illustration of the probe with an attached sensor and Gronningsaeter et al. (2000) for a more detailed description and a discussion regarding the reproducibility of removing and replacing the tracking device. We do not investigate this reproducibility in this paper.

We do not use frame-grabbed (video) data in our experiments, but instead non-scan converted digital data from the ultrasound scanner. Incorporated in these data is information from the scan process, such as start angle, angle increment (before next beam), scan depth, etc. This allows us to treat scaling (mm/pixel in the images) issues slightly different than e.g. Prager et al.[25, 26]

Three different phantoms were used in this study. The first was a thin jagged membrane oriented vertically at the bottom of a water bath (Fig. 2) and was used in the first method for probe calibration described in the next section. The second phantom was a thin planar nylon mesh membrane stretched across a rectangular frame in a water bath (Fig. 3). This phantom was used in the new calibration method presented later. Finally, we developed a special wire phantom for investigation of reconstruction accuracy in 3D freehand ultrasound scans. This water-filled phantom consists of 27 precise and well-defined points, 18 thin cotton threads crossing in three layers, inside a 5x5x5 cm cube (Fig. 4). A position adapter was attached to the phantom, making it possible (from physical measurements) to calculate the positions of all 27 crossings relative to the position of the adapter with an expected accuracy of less than ±0.1 mm.
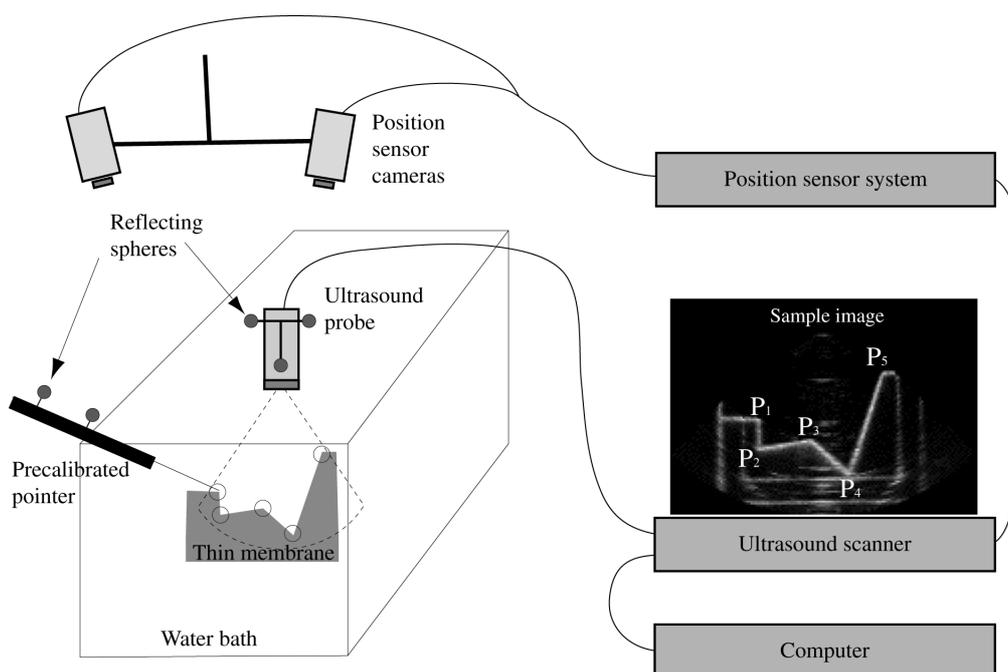

Fig. 2. Set up for the *Membrane Alignment Method* for probe calibration. The image is tagged with the position readings for the sensor frame attached to the probe in the ultrasound scanner and transferred to an external computer for further processing. The pre-calibrated pointer is used to mark the membrane corner points in the global system.
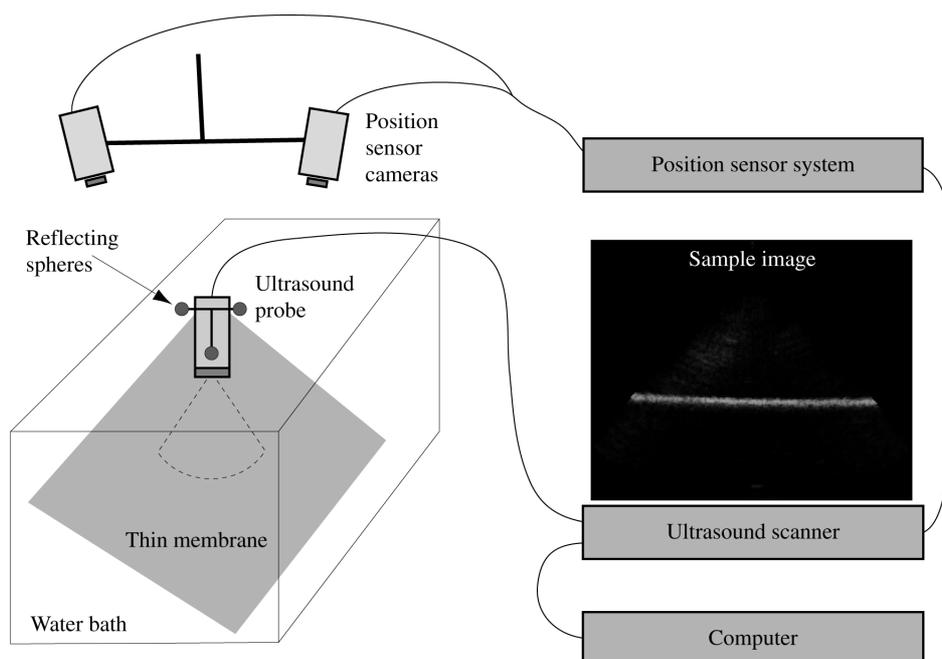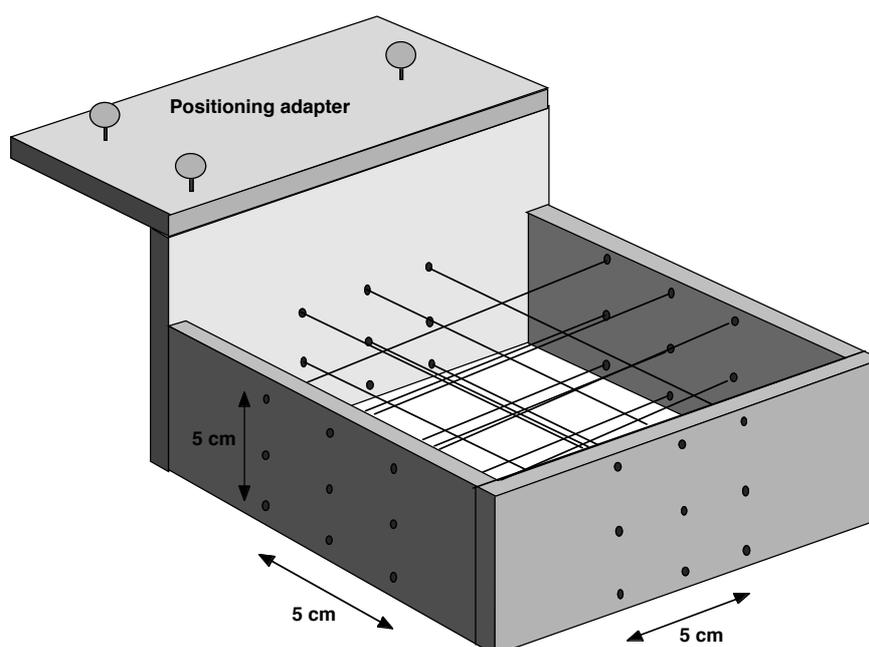
Fig. 3. Set-up for the new *Freehand Membrane Scan Method* for probe calibration. Each image is tagged with the position readings for the sensor frame attached to the probe in the ultrasound scanner and transferred to an external computer for further processing.

Fig. 4. Wire phantom used for measurements in the accuracy analysis. Eighteen cotton threads form a cube of 5x5x5 cm with 27 thread crossings.

Membrane Alignment Method

This is a method developed by our group of researchers previously and used thus far in our studies involving freehand 3D acquisitions. The technique was briefly described by Berg et al.[3] A schematic overview is shown in Fig. 2. The idea is to align the image plane from the ultrasound probe by 2D real time scanning with a submerged, extremely thin, but stiff and jagged membrane oriented approximately perpendicular to the bottom of a water tank. When the image is such that the membrane and the scan plane can be assumed to be coplanar, the

positions of the sensor on the probe and the positions of the membrane are sampled. The positions of the membrane corners are found by pointing at them with a pre-calibrated pointer. The ultrasound image of the membrane is transferred to the computer and the five corners are marked by pointing at them manually with the mouse. These image points are denoted P_{us}, a five column position matrix, each column representing the position of a corner of the membrane in Fig. 2. The corresponding phantom points marked by the pre-calibrated pointer in the global position sensor system are represented by P_{gl}. The pseudo-inverse matrix for transformation between the two systems, i.e. between the image system and the global system, is given by:[13]

$$M_{gl \leftarrow us} = P_{gl} \cdot P_{us}^T \cdot \left(P_{us} \cdot P_{us}^T \right)^{-1} \tag{8}$$

where the upper index T means transposed, lower index gl refers to the global coordinate system, often named patient or the reference system, and the matrix $M_{gl \leftarrow us}$ transforms the points sampled in the image system to the global coordinate system such that

$$P_{gl,us} = M_{gl \leftarrow us} \cdot P_{us} \tag{9}$$

$M_{gl \leftarrow us}$ is transformed to homogeneous coordinate matrix representation[11]

$$M_{gl \leftarrow us} = \begin{bmatrix} p_{x1} & p_{y1} & p_{z1} & t_x \\ p_{x2} & p_{y2} & p_{z2} & t_y \\ p_{x3} & p_{y3} & p_{z3} & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{10}$$

by requiring a right handed coordinate system, i.e. $\vec{p_z} = \vec{p_x} \times \vec{p_y}$. The upper right 3 by 3 matrix is the rotational transformation (see Eq. 1). The following matrix multiplication then gives us the desired calibration transformation

$$M_{lo,probe \leftarrow us} = M_{lo,probe \leftarrow gl} \cdot M_{gl \leftarrow us} = \left(M_{gl \leftarrow lo,probe} \right)^{-1} \cdot M_{gl \leftarrow us} \tag{11}$$

where $M_{gl \leftarrow us,probe}$ is the position sensor reading for the frame attached to the probe as the image and membrane are assumed coplanar. The deviation between the image sampled points P_{us} and the points sampled by the pointer in the global coordinate system P_{gl} (Eq. 9) may be defined as a measure of the error and estimated as

$$\delta_1 = \sqrt{\frac{\sum_{i=1}^{5} \left| P_{i,gl,us} - P_{i,us} \right|^2}{5}} \tag{12}$$

where index i varies from 1 to 5 representing the number of each of the corners of the jagged membrane in Fig. 2. Eq. 12 may be seen as a root mean square error.

A point $\left[x_{us}, y_{us}, z_{us} \right]^T$ in the ultrasound image may be transformed into the global coordinate system by

$$\begin{bmatrix} x_{gl,us} \\ y_{gl,us} \\ z_{gl,us} \\ 1 \end{bmatrix} = M_{gl \leftarrow lo} \cdot M_{lo,probe \leftarrow us} \cdot \begin{bmatrix} x_{us} \\ y_{us} \\ z_{us} \\ 1 \end{bmatrix} \qquad (13)$$

This transformation $M_{lo,probe \leftarrow us}$ is applicable in general since the probe and sensor constitutes a rigid body. We evaluate the accuracy of the probe calibration in a separate section later with a purpose built phantom.

A problem with this method is that it depends on how accurately the user can align the probe image plane and the membrane, which can be difficult given the thickness and widening of the ultrasound beam as a function of depth. In addition, the accuracy of pointing at the membrane corners in the ultrasound images is subject to error. Furthermore, the membrane may have to be built in more than one size given the fact that different probes have highly varying array sizes and maximum scan depths. The process is time consuming and the inter-operator variability may be large. We compare calibrations from this method with the new procedure described in the next section by using a purpose built phantom for accuracy evaluation.

Freehand Membrane Scan Method
We have developed a new method for probe calibration that is based on elements from the previously described methods. In particular, the *Single-wall Method* is similar to this new method. The major improvements are that we use a thin nylon mesh membrane instead of the bottom of the water tank, an automatic edge detection procedure, and only one input parameter. The method is described in detail below.

Scanning
The first step of the procedure is to perform an acquisition of a set of 2D images of the submerged membrane in the water bath, as shown in Fig. 3.[19] For the calibrations performed in this study, we simply used the thin mesh of woven nylon from stockings stretched across a rectangular frame submerged in a water bath. The water was vacuumed so that it contained as few air bubbles as possible to ensure high quality images of the membrane. The probe was then moved about during acquisition within a range of angles and distances covering the six degrees of freedom for movement in space (Fig. 5).
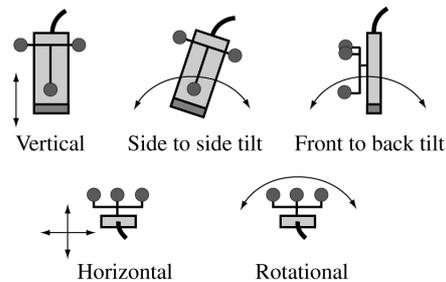

Fig. 5. Illustration of the minimal sequence of motions for calibrating the probe. All six degrees of freedom in the calibration matrix (three translations and three rotations) must be exercised to identify the calibration parameters from the set of equations.

Moving the probe may be accomplished in two ways: 1) Moving the probe from position to position and then keep it still during acquisition, and 2) Moving the probe continuously and acquire all images during the process. The accuracy of four different optical localizers was evaluated and compared by Chassat and Lavallée (1998). They concluded that the dynamic accuracy, i.e. how accurately the systems track a moving sensor, in general for optical position sensing systems seems to be worse than the static accuracy. Due to their findings and some initial tests performed in our laboratory, we have chosen to present results based only on acquisition method 1. The problem of dynamic accuracy is further elaborated in the discussion.

Each acquired image is tagged with position data for the attached sensor and all acquired data are stored for further processing. A program reads the position data continuously during scanning and calculates the motion, providing feedback to the user about the overall movement in millimeters and degrees.

Using u and v as coordinates for the images of the membrane and setting the membrane plane equal to the xy-plane in the membrane system (Fig. 1), a point located on the membrane in an image will be given by the transformation

$$
\begin{bmatrix} x \\ y \\ 0 \\ 1 \end{bmatrix} = M_{m \leftarrow gl} \cdot M_{gl \leftarrow lo} \cdot M_{lo \leftarrow us} \cdot \begin{bmatrix} 0 \\ u \\ v \\ 1 \end{bmatrix} \tag{14}
$$

The z-coordinate of the membrane and the x-coordinate of the ultrasound image are equal to zero due to the way we have defined our coordinate systems (Fig. 1).

Edge detection
Next we need to locate the membrane in the images, i.e. estimate pairs of *(u,v)*-values for insertion in Eq. 14 above. We assume proper settings on the scanner and set-up for the water tank and membrane such that the walls and the bottom of the water tank do not show up in the images. Nevertheless, to obtain a robust edge detection, we have developed and implemented a new wavelet-based edge detection method.[16] In short, each ultrasound beam is transformed into the wavelet domain using the continuous wavelet transform (CWT). As the CWT preserves both scale and time information, it is possible to separate the signal into a number of scales. The edge or discontinuity is located by finding the position of the maximum of the CWT at the most energetic scale. The most energetic scale is found by computing the power spectrum based on the wavelet transform. We use the Mexican hat as the analyzing wavelet, often referred to as the mother wavelet. It should be emphasized that this detection method is not depending on any threshold. Furthermore, this method only finds one edge per scan line (or beam), and hence is especially adapted to edge detection problems of the kind described in this paper. Some sample edge detections based on the CWT method are presented in Fig. 6. The images have not been scan-converted (see next section) to more easily distinguish and visualize the individual beams. Hence, the edges are not straight lines but rather curved. For a more detailed description of this algorithm, see Kaspersen et al.[16]
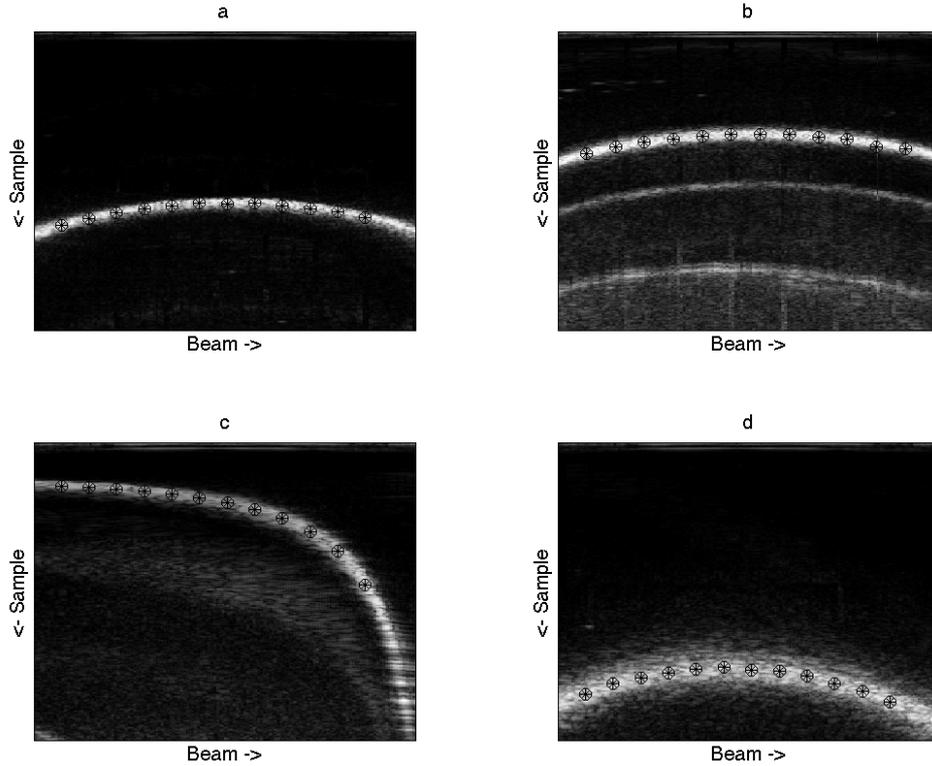
Fig. 6. Edge detections using the wavelet-based method are shown as small black circles (black crosses inside), superimposed on non-scan converted images of the membrane. The detections were performed at every twentieth beam. (a) The probe is approximately normal to the membrane. (b) Same as (a) but the membrane is positioned deep in the water tank so that we get reverberations from the bottom. (c) The probe is tilted approximately 60° about the image *x*-axis relative to the situation in (a). (d) The probe is tilted approximately 60° about the image *y*-axis.

Scan conversion

The images containing the edge detections (Fig. 6), i.e. the ultrasound scan line images, are geometrically distorted images since the beam is actually scanned within a sector. To generate the true geometric shape, a scan converter algorithm is used. This scan converter procedure needs the velocity of sound as input to correctly position the different samples along each beam. We have implemented the formula[8]

$$c_{water} = \sum_{i=1}^{5} k_i T^i \tag{15}$$

for the velocity of sound in water based on measurement of the temperature, where the coefficients are given by

i	k_i (m/s)
0	1402.39
1	5.03711
2	$-5.80852 \cdot 10^{-2}$
3	$3.34199 \cdot 10^{-4}$
4	$-1.47800 \cdot 10^{-6}$
5	$3.14643 \cdot 10^{-9}$

The experiments were all carried out in water at a room temperature of 22°C. Consequently, the value for the velocity of sound was set to 1488 m/s. Assuming the scan conversion is correct, we achieve values for *u* and *v* in millimeters and we omit scale parameters to *u* and *v*.

Nytt: The scaling calibration parameters (see section on *Scan conversion* in *Materials and Methods*) are only valid for the medium in which the calibration process is performed.

Straight line fitting

The next step of the process is to determine the straight line that represents the membrane from each of the scan-converted and edge-detected images. Given the quality of ultrasound images in general, our algorithm for this step should disregard pixels located far from the actual position of the membrane, i.e. we assume that there are some incorrect edge pixel candidates from the edge detection due to acoustical noise. A least squares fit algorithm will not work since all the edge detections contribute to the solution, and not only the points comprising the straight line. Instead we fit a line between random pairs of edge pixel candidates, making sure we do not use the same pixel candidate more than once. Keep in mind that there is only one edge pixel candidate per beam in the image at this point. Next, each of the line segments are given a score according to how far from this line segment the other edge pixel candidates are located, i.e. all of the candidates except the two constituting the current line segment. The score of the line segment is increased every time another candidate is located no more than some threshold value away from the line segment. The principle is illustrated in Fig. 7, where it is easy to see that for a small enough threshold value, line segment *b* (defined by points α and β) will receive many more points than line segment *a* (defined by points δ and γ).
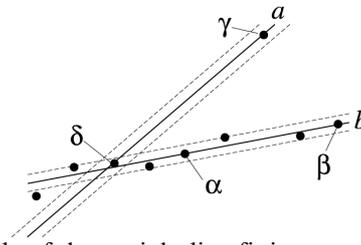

Fig. 7. Illustration showing the principle of the straight line fitting procedure. Line *a* will be discarded due to a large distance between the line segment given by point δ and γ and each of the other edge candidates.

The dotted lines represent the extent of the threshold value. The threshold parameter is the only input to the method and should be small to ensure high accuracy. We have used one pixel unit throughout the trials in this study.

The distance between a point and a line is given by[9]

$$d = \frac{\left|\overrightarrow{AP} \times \overrightarrow{AQ}\right|}{\left|\overrightarrow{PQ}\right|} \tag{16}$$

In this equation $\overrightarrow{PQ}$ represents a vector along the line, i.e. the two edge candidates in question, and A is one of the other edge candidates. The line segment with the highest score is assumed to be a correct representation of the membrane line in the image. Finally, two pixels are selected from the resultant line and used as values for *u* and *v* in Eq. 14. We simply divide

the line in three segments (number of beams divided by three) and select the two points from the straight line fit at the first and last beam number comprising the middle segment. Two points uniquely define a line in space and, moreover, we double the number of equations (see next section) to be solved for the unknowns. See Fischler and Bolles[10] for a more detailed description of the idea of giving scores to line segments in line fitting.

Solving system of equations

From the zero component of Eq. 14 we obtain an equation with nine unknowns for each pixel point (u,v). The unknowns are two rotation angles, β_m and γ_m, and one translation offset, $t_{m,z}$, in the matrix $M_{m \leftarrow gl}$, and all rotation angles and translation parameters, α_c, β_c, γ_c, $t_{c,x}$, $t_{c,y}$, and $t_{c,z}$ from the calibration matrix (Eq. 1). The rotation α_m about the *z*-axis of the membrane and the translation $t_{m,x}$ and $t_{m,y}$ of the membrane in its *xy*-plane do not affect the *z*-component of the equation. This may be seen from the geometry in Fig. 1. Consequently, we need at least nine independent equations to solve the problem. The resulting equations obtained from the line fittings are not necessarily independent unless the probe is moved appropriately during the scanning of the membrane (Fig. 5). As mentioned above we can write the resultant equations for two pixels on each line since two points uniquely define a line, yielding two equations for every 2D scan of the membrane. With *M* images, this yields *m*=2*M* equations, each with nine unknown parameters. Hence, the minimum number of images, *M*, needed to solve the system of equations is five. However, we use more than five images to obtain an over-determined set of equations.

This set of equations is solved with the Levenberg-Marquardt algorithm[23] as described in the theory section. Like other iterative algorithms, the Levenberg-Marquardt method needs a starting point or an initial guess solution to the problem. This guess could be any rough estimate of the true solution, but the closer it is to the true solution, the quicker the convergence. The initial estimate in a probe calibration can be found from a physical measurement with a ruler of the translation and approximate knowledge about sensor orientation relative to the ultrasound image plane. The initial estimate for the parameters β_m, γ_m, and $t_{m,z}$, in the matrix $M_{m \leftarrow gl}$, is found by using some of the edge detections together with the initial estimate for the probe calibration matrix, $M_{lo \leftarrow us}$. Two edge detections from two images, i.e. four points, are transformed to global coordinates by using Eq. 14 (without $M_{m \leftarrow gl}$). The initial estimate of the calibration matrix is used for $M_{lo \leftarrow us}$ in the equation. Each of these two sets of two points, i.e. two vectors defining the membrane in two separate images, are then used to set up a coordinate system of the membrane in the global coordinate system. The *z*-axis (Fig. 1) is found from the cross product of the two vectors, while the origin may be any of the four points. One of the vectors (two edge detections from one of the images) may define the *y*-axis, while the last axis, the *x*-axis, is found from the cross product between the *y*- and *z*-axis. Hence, we have an estimate of the matrix describing the membrane in the global coordinate system, i.e. $M_{m \leftarrow gl}$.

Accuracy analysis

The quality of the probe calibration methods was determined through the following accuracy analysis: The probe was calibrated six times using both the *Membrane Alignment Method* and the *Freehand Membrane Scan Method*. The probe was moved from position to position and held still during acquisition of each 2D frame acquired for the latter method. A corresponding 3D acquisition of the wire phantom (Fig. 4) was performed for each calibration.[12] The

volumes were scan converted utilizing the calibration matrix and the velocity of sound in water. The position of the crosses were pinpointed in each data set using the mouse, while the calculated positions of the wire crosses (from physical measurements) relative to the attached position sensor (Fig. 5) were used as the global coordinates for the crosses. We then compared the positions of the points in image and global space coordinates by calculating residual vectors. These vectors are found by calculating the differences between the points from the two data sets, i.e. the differences between each of the six sets of 27 points from image space and the one set of 27 points from physical measurements, i.e. a total of 162 points. These vectors and their lengths are given by

$$R_i = I_i - G_i = \begin{bmatrix} x_i \\ y_i \\ z_i \end{bmatrix}$$

$$r_i = \|R_i\| = \sqrt{x_i^2 + y_i^2 + z_i^2}$$

(17)

The endpoints of these error vectors describe a cloud in 3D space, while the length of each vector specifies the distance in millimeters between two corresponding measurements. The offset vector is found by calculating

$$\rho_R = \frac{1}{162} \sum_{i=1}^{162} R_i = \begin{bmatrix} \rho_x \\ \rho_y \\ \rho_z \end{bmatrix}$$

(18)

If the length of the vector found from Eq. 18 is significantly different from zero, we probably have some systematic error in the calibration transformation matrix. The spread can be defined by the following formula

$$\sigma_R = \frac{1}{162} \begin{bmatrix} \sqrt{\sum_{i=1}^{162} \left(R_{ix} - \rho_x\right)^2} \\ \sqrt{\sum_{i=1}^{162} \left(R_{iy} - \rho_y\right)^2} \\ \sqrt{\sum_{i=1}^{162} \left(R_{iz} - \rho_z\right)^2} \end{bmatrix} = \begin{bmatrix} \sigma_x \\ \sigma_y \\ \sigma_z \end{bmatrix}$$

(19)

Furthermore , we define a sphere with radius *err* centered at the origin. The percentage of points relative to the total falling inside the sphere for different values of *err* can be accumulated and presented as

$$p = \frac{\sum_{i=1}^{162} n_i}{162} \cdot 100, \text{ where } n_i = \begin{cases} 1, \text{ if } r_i \leq err \\ 0, \text{ if } r_i > err \end{cases}$$

(20)

Accuracy analysis from iteration parameters
By calculating the value for *F* in Eq. 4 for all equations, inserting the parameter values for the unknowns from the solution, we can get an indication of how well the transformations, including the calibration matrix, describe the transition from the membrane in the images to the membrane in global space (Fig. 1). We performed this calculation for both acquisition methods for all the calibration results of the *Freehand Membrane Scan Method* and computed the root mean square value of the results (Eq. 4):

$$\delta_2 = \sqrt{\frac{\sum_{i=1}^{N}\left[F_i\left(\vartheta,\varphi_{solution}\right)\right]^2}{N}} \tag{21}$$

where *N* is the total number of equations used in a calibration matrix estimation.

To check that the equations are significantly independent, and hence that all the unknowns are identifiable, a rank determination may be carried out. This is accomplished by evaluating the Jacobian matrix, e.g. at the initial estimate of the calibration matrix. To determine that this matrix is of full rank, a singular value decomposition (SVD)[14, 17] is performed. This is achieved by letting

$$J = UDV^T \tag{22}$$

where *U* and *V* are unitary matrices and

$$D = diag\left(\mu_1,...,\mu_6\right) \tag{23}$$

is the matrix of ordered singular values, which are the non-negative square roots of the eigenvalues of $D^* \cdot D$. The determinant, or the so-called *Jacobian*,[9] is said to have full rank if all μ_i are nonzero. The identifiable nature of the different unknowns can be measured by the condition number for *D*. This is simply equal to the ratio of largest μ_i-value to the smallest:

$$\kappa = \frac{\max\left(\mu_1,...,\mu_6\right)}{\min\left(\mu_1,...,\mu_6\right)} \tag{24}$$

In other words, this last parameter indicates whether or not there are enough independent equations, i.e. whether we have moved the probe sufficiently and acquired enough images in total. Consequently, by performing edge detections and iterations during acquisition (starting after five or more images have been stored) this parameter, κ, may be used to provide feedback to the user about the status of the current data set.

RESULTS

A plot of the residual vectors given by Eq. 17 is shown in Fig. 8a for the *Membrane Alignment Method* and in Fig. 8b for the *Freehand Membrane Scan Method*. The figures indicate that 68% of the residuals fall inside a sphere of radius *err* = 6.0 mm (Eq. 20) for the *Membrane Alignment Method* and *err* = 4.4 mm for the *Freehand Membrane Scan Method*. This means that we can expect 68% of all points that are marked in an ultrasound volume to correspond within 4.4 mm to the true point in 3D space. This result is slightly worse than what was found by Gronningsaeter et al. (2000).
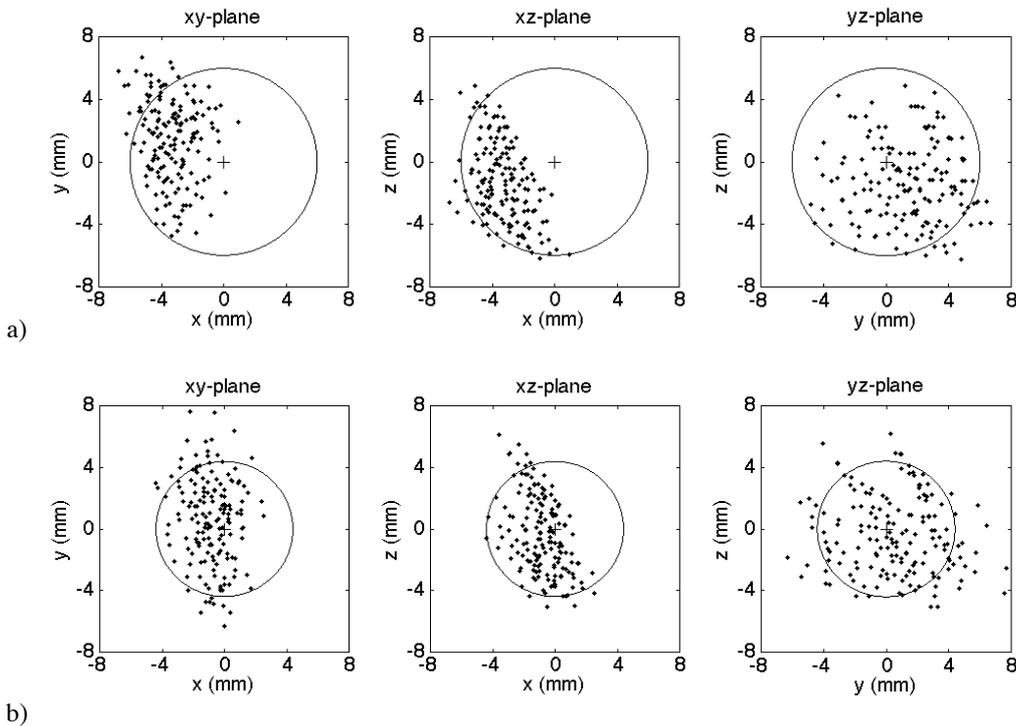

Fig. 8. Two-dimensional plots (projections) of the residual vectors (Eq. 17) based on the volumes scan converted with the calibrations from (a) the *Membrane Alignment Method*, and (b) the *Freehand Membrane Scan Method*. These plots indicate the bias ρ_R (Eq. 18) and the spread σ_R (Eq. 19). The circles (actually projections of a sphere) indicate that 68% of the samples fall inside a sphere of radius (a) *err* = 6.0 mm and (b) *err* = 4.4 mm (Eq. 20).

Table 1 summarizes the results for the root mean square measures (Eq. 12 and 21), average length of residual vectors (Eq. 17), bias vectors (Eq. 18), and *err* values (Eq. 20) for both methods. It can be seen that the offset vector is smaller for the *Freehand Membrane Scan Method* (1.10 mm) than for the *Membrane Alignment Method* (3.81 mm), although the residual vector average length is not significantly different. The spread also seems to be similar from the plots in Fig 8a-b.

	Mean(r)± STD(r) (mm) Eq. 17	$|\rho_R|$ (mm) Eq. 18	err (mm) when $p = 68\%$ Eq. 20	'RMS'
Membrane Alignment Method	5.28±1.43	3.81	6.0	$\delta_1 = 0.34$ (Eq. 12)
Freehand Membrane Scan Method	3.83±1.56	1.10	4.4	$\delta_2 = 0.63$ (Eq. 21)

Table 1. The various measures indicating the accuracy of *Freehand Membrane Scan Method*. See Eq. 12 and Eq. 18-21 for an explanation of the parameters.

The condition number (Eq. 21), which measures the independence of the equations, for the *Freehand Membrane Scan Method* varied between 446 and 643. The number of equations was 48, i.e. 24 images were acquired for each calibration, while the total number of iterations needed to achieve convergence varied from 40 to 80. Compared to the results of Prager et al.[25, 26] our values for the condition number and number of iterations were slightly higher. The reason for this difference is addressed in the discussion. We did not experience that the iteration stopped at local minima. However, this happened in some cases when we tested the *Freehand Membrane Scan Method* with continuous acquisitions, i.e. when the images were acquired with a continuous probe movement. This is further discussed in the next section.

DISCUSSION

We have presented new methods for ultrasound probe calibration. The *Freehand Membrane Scan Method* is easy to use, requires a minimum of input from the user, i.e. the velocity of sound and an initial estimate of the calibration parameters. In addition, the equipment needed to carry out the procedure is simple and inexpensive to make. Furthermore, we have presented a method for estimating the accuracy of probe calibrations by use of a special wire phantom. The accuracy for both the *Membrane Alignment Method* and the *Freehand Membrane Scan Method* were evaluated.

It is difficult to compare results on the performance of calibration methods from different publications because they depend on many factors such as the type of accuracy estimated, the method for calculating the accuracy, range of scanning angles and positions, type of position sensing device, as well as the quality of the ultrasound equipment - the scanner, probe, software optimizations, etc. Nonetheless, we can compare the results from the two methods presented in this paper. It can be seen from these results that the *Freehand Membrane Scan Method* resulted in a lower bias than the *Membrane Alignment Method*. The spread is approximately the same for both methods. Furthermore, it is believed that the *Freehand Membrane Scan Method* yields a lower inter-operator variability due to the alignment procedure for the *Membrane Alignment Method*. It can further be seen from the plots of the residual vectors (Eq. 17, Fig. 8a-b) that the spread is not spherical for the methods (except for perhaps the *yz*-plane projection) indicating that there might be a potential for further improvements.

There are several possible sources of error in the presented methods. A common source for both methods is the position sensing system. As mentioned previously, the accuracy of four different optical localizers was evaluated and compared.[5] The authors concluded that the dynamic accuracy, i.e. how accurately the systems track a moving sensor, seemed to be worse

then the static accuracy. During initial tests, we found that acquisitions for the *Freehand Membrane Scan Method* performed with continuous probe movement had poor results. The *rms* values (Eq. 21) were larger than the *rms* values obtained from the data sets acquired when the probe was held still in each position. The set of equations produced by continuous probe movement also resulted in a larger condition number, demonstrating that even though there were more equations (more images), these were not independent enough. We even experienced some iterations that terminated because they got stuck in local minima.

To illustrate this possible lag in position tagging due to limited dynamic properties in the position system, we have plotted F (Eq. 4) at the initial guess for a shorter sequence of images from a continuous, but rather fast movement, approximately along the positive then the negative y-axis of the membrane (Fig. 9a).

a)   b)
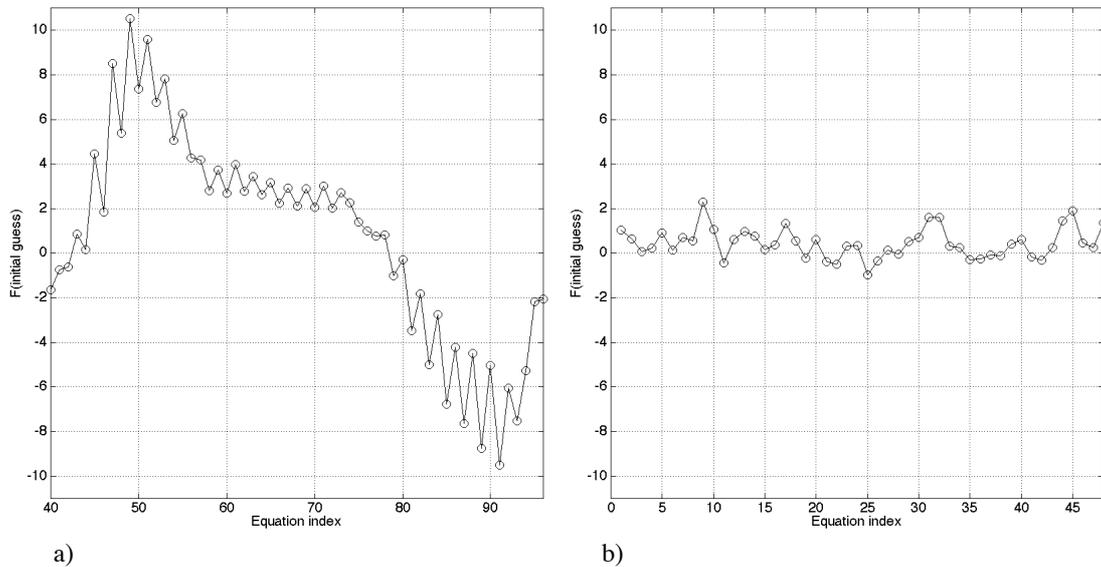
Fig. 9. Illustration of the possible time lag in position tagging during continuous acquisition. (a) Plot of how well the initial guess parameters and the position data describes the transformation of a point on the membrane in the image via the global coordinate system to the corresponding point on the physical membrane (Eq. 4 and 14). The result should be close to zero as explained in Eq. 4 and it represents the residual as explained prior to Eq. 21. The probe was moved rather fast in this test, and the movement was along the positive and negative y-axis of the membrane coordinate system in Fig.1. (b) Corresponding plot for data set from acquisition with the probe held still for each position.

For comparison, we have shown another plot of F (Eq. 4) at the initial guess for a set of equations obtained from acquisition with the probe held still in each position approximately along the same path as for the continuous acquisition (Fig. 9b). These values should all be close to zero since Eq. 4 describes the transformation from the membrane in the image, via the global position sensing origin, and back to the same point on the membrane in global coordinates. The plot in Fig. 9a clearly shows that either the dynamic accuracy is poorer than the static accuracy or the internal position tagging in the scanner is poor. It should be mentioned that these tests were performed with a software version for the Vector Vision position system that has since been updated and the company claims that the problem with dynamic inaccuracy has been solved.

Another factor possibly contributing to the high root mean square values for the continuous probe movement acquisition, is the fact that the membrane seemed to move slightly during

these acquisitions, especially when the probe was moved in the ±*z*-axis direction of the membrane system, i.e. perpendicular to the membrane. The membrane moved much in the same manner as the surface of a drum moves, due to pressure waves from probe movement in the water. This effect was observed in only a few of the scans during processing of the images. The movement can be diminished by using a stiffer membrane or stretching it until it becomes less prone to vibrating movement due to water pressure waves.

A disadvantageous effect in the *Single-wall Method* is the difficulty of localizing the true edge (bottom of the water tank) at oblique scan angles.[26] Due to the finite thickness and widening of the ultrasound beam with depth, the peak intensity position, which is defined as the position of the membrane in the image, is difficult to determine. Since we are using a thin membrane instead of the bottom of the water tank, we achieve a much narrower line in the images, even at extreme tilts about the image *y*-axis (Fig. 6d). The effect of scanning a plane surface (tank bottom or membrane) at an oblique angle is illustrated in Fig. 10.
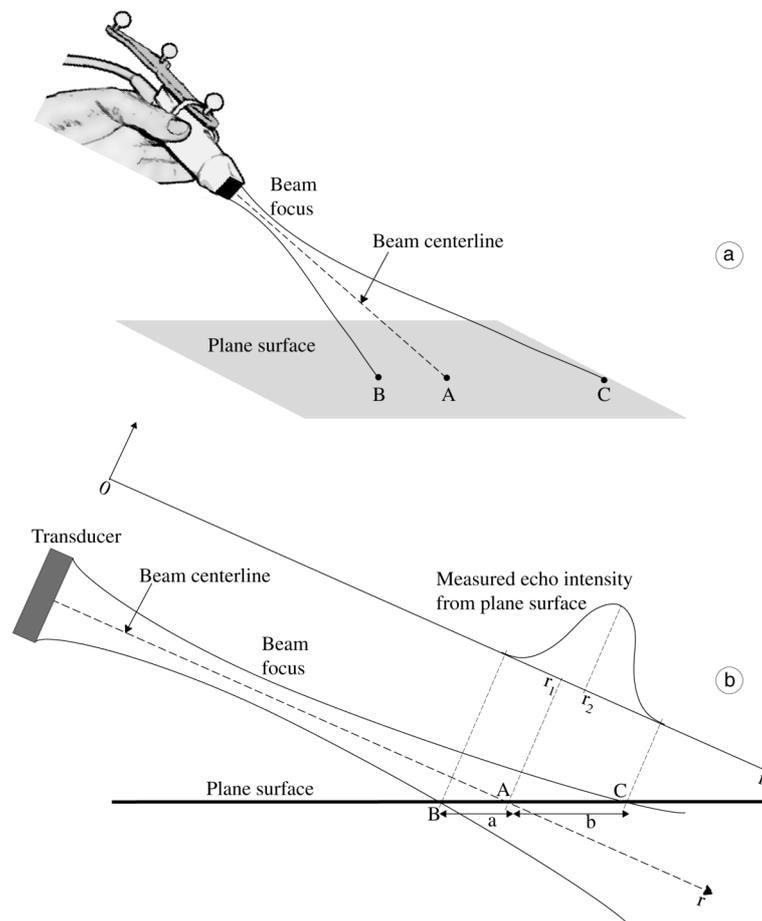

Fig. 10. Illustration of the effect occurring when scanning a flat surface at oblique angles. (a) Due to the beam thickness, point *B* is encountered earlier than point *A*, and the echo from point *B* does not resemble the true distance to the surface with the probe in this position. (b) The estimated distance to the surface based on the measured peak intensity, r_2, does not necessarily correspond to the true distance r_1. However, this error is believed to be small for the *Freehand Membrane Scan Method* compared to the single-wall method. Only the peak position and start plus end positions of the intensity plot are of interest here and hence ewe have not attempted to plot a true intensity profile from a reflected beam, which is typically more asymmetric than depicted in the figure.

The point B in Fig. 10a is encountered earlier than point A, while in Fig. 10b the distance a is shorter than b due to the widening of the beam. By letting r_1 represent the true distance to the surafce and r_2 the estimated distance to the center of the intensity peak, the degree of error, i.e. the difference between r_1 and r_2 depends on the surface, the specific probe (the beam profile), and the magnitude of the tilt angle. We believe that this error is insignificant for the new proposed method based on the results shown in Fig. 6, especially Fig. 6d. Furthermore, the fact that we use a meshed nylon membrane means that even for scans at very oblique angles will give intensity beams like the one depicted in Fig. 10b. The function F (Eq. 4) evaluated at the solution for the images like Fig. 6d did not give values significantly different than F evaluated at points from other images. Nevertheless, this effect is subject to further research. For comparison, in Fig. 11 we have plotted beams from similar membrane images to those in Fig. 6a and 6d and also beams from images of the water tank bottom.
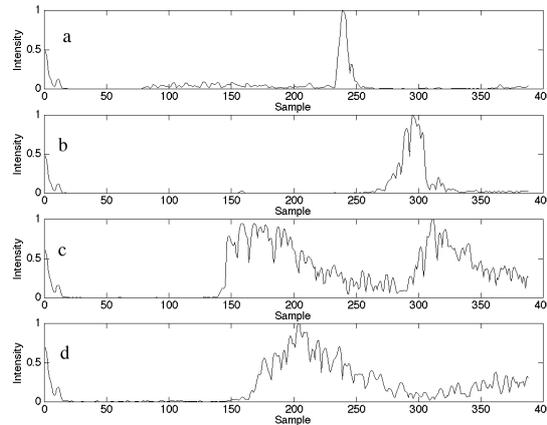

Fig. 11. Intensity plots of beams from images similar to those in Fig. 6. (a) Probe approximately perpendicular to thin membrane. (b) Probe at an oblique angle to the membrane. (c) Probe approximately perpendicular to the bottom of the water tank. (d) Probe oriented at an oblique angle to the bottom. Sample number is along the x-axis, while normalized intensity is given by the y-axis.

It can be seen in Fig. 11c that strong reverberations occur when the bottom is scanned at a perpendicular orientation. In addition, the beams in Fig. 11c and 11d are much wider than the corresponding beams from the images of the membrane (Fig. 11a and 11b). The energy reflected back to the probe at oblique scan angles is smaller when scanning the smooth bottom of the water bath compared to the reflected energy from the thin membrane with a mesh structure. This effect can be somewhat compensated for by roughening the floor of the water bath as suggested by Prager et al.[26]

The accuracy of the calibration depends to some extent on the numerical methods employed, i.e. the Levenberg-Marquardt iteration procedure in our case, and the preparation of the data prior to running this iteration, i.e. the line fitting algorithm. A large condition value (Eq. 24) indicates that the problem is ill-conditioned and that at least one parameter cannot be accurately identified. Having removed parameters we know are not identifiable, there are two main reasons for a large condition number. The first is that the motion of the probe may not have covered all the necessary degrees of freedom. This can be solved by simply extending the image data set. An advantage for magnetic sensors in this regard is that they may be moved more freely than an optical sensor due to the limiting line of sight to the sensor cameras for the optical system. Hence, the condition number should be smaller for a calibration performed with a magnetic sensor, given that the sensors are equally accurate. Secondly, we have degeneracy in the axis representation of rotation, i.e. multiple triples of α,

β, and γ that produce the same transformation matrix. Effectively, one of the angles becomes unidentifiable. However, this degeneracy has no effect on the calibration result since any of the triples solutions are equally valid. The degeneracy may be removed by mounting the frame askew on the probe relative to the scan plane, such that all the rotation angles are different from $\pm 0°$ or $\pm 90°$.

The different unknown parameters have different units, millimeters and radians, and hence different effects on the residuals produced by Eq. 4. Parameter scaling involves modifying Eq. 6 to account for this fact. This is accomplished by letting

$$\Delta F = J \Delta \varphi = J \left(H H^{-1} \right) \Delta \varphi = \left(J H \right) \left(H^{-1} \Delta \varphi \right) \tag{25}$$

where H is a scaling matrix chosen such that JH is better conditioned than J alone. After convergence, the solution is found by unscaling. The scaling matrix can be found by column scaling[20]

$$H = diag \left(h_1, ..., h_6 \right)$$

$$h_i = \begin{cases} \left\| j_i \right\|, \ if \ \left\| j_i \right\| \neq 0 \\ 1, \ if \ \left\| j_i \right\| = 0 \end{cases} \tag{26}$$

where j_i is the ith column of J, evaluated at the initial guess solution. This scaling ensures that each parameter has the same effect on ΔF. We believe that such an implementation would further improve our method.

CONCLUSION

We have reviewed methods for probe calibration and found most of the existing techniques to have drawbacks. In particular, the *Cambridge Phantom Method*, which is close to automatic, depends on a precision-made phantom, several adjustable input parameters, and an alignment procedure that again requires extra equipment. We have suggested an improved method, the *Freehand Membrane Scan Method*. This new method incorporates an automatic edge detection procedure, a simple and easy to use phantom (a stretched nylon membrane in a water tank), and only one input parameter, i.e. the velocity of sound. A phantom and a method for investigating the accuracy of calibration methods has been demonstrated and used to compare the *Membrane Alignment Method* and the *Freehand Membrane Scan Method*. The accuracy analysis showed that the latter method resulted in a lower bias. 68% of all points marked in an ultrasound 3D freehand scan of the wire phantom were within 4.4 mm of the true location, compared to 6.0 mm for the *Membrane Alignment Method*.

ACKNOWLEDGEMENTS

REFERENCES

1. Angelsen BAJ. Waves, signals and signal processing in medical ultrasonics. Trondheim, Norway, Dept. of Physiology and Biomedical Engineering, The Norwegian University of Science and Technology (NTNU), I, 1995.

2. Barry CD, Allott CP, John NW, Mellor PM, Arundel PA, Thomson DS, and Waterton JC. Three-dimensional freehand ultrasound: image reconstruction and volume analysis. Ultrasound Med Biol, 1997;23:1209-1224.

3. 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 Med Biol, 1999;25:745-753.

4. Carr J. Surface reconstruction in 3D medical imaging. Ph.D., University of Canterbury, Christchurch, New Zealand, 1996.

5. Chassat F, and Lavallée S. Experimental protocol of accuracy evaluation of 6-D localizers for computer-integrated surgery: Application to four optical localizers. Lecture notes in Computer Science, Medical Image Computing and Computer-Assisted Intervention - MICCAI'98, Springer-Verlag, Berlin, Heidelberg, 1998;1496:277-284.

6. Craig J. Introduction to robotics: Mechanics and control. Reading, Massachusetts, Addison-Wesley, 1989.

7. Detmer PR, Bashein G, Hodges T, Beach KW, Filer EP, Burns DH, and Strandness DEJ. 3D ultrasonic image feature localization based on magnetic scanhead tracking: In vitro calibration and validation. Ultrasound Med Biol, 1994;20:923-936.

8. Duck FA. Physical properties of tissue. A comprehensive reference book. London, Academic Press Ltd., 1990.

9. Edwards CH, and Penny DE. Calculus and analytic geometry. Englewood Cliffs, N.J., USA, Prentice-Hall, 1990.

10. Fischler MA, and Bolles RC. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM, 1981;24:381-395.

11. Foley JD, vanDam A, Feiner SK, and Hughes JF. Computer graphics : Principles and practice. Reading, Massachusetts, Addison-Wesley, 1990.

12. Gronningsaeter A, Lindseth F, Langø T, Kleven A, Ommedal S, Lie T, and Unsgård G. Accuracy assessment of a prototype ultrasound-based neuronavigation system. Submitted to Neurosurgery, 2000.

13. Hagan MT, Demuth HB, and Beale M. Neural network design. Boston, Massachusetts, PWS Publishing Co., 1996.

14. Hollerbach JM, and Wampler CW. The calibration index and taxonomy for robot kinematic calibration methods. Int J Robotics Research, 1996;15:573-591.

15. Hughes SW, D'Arcy TJ, Maxwell DJ, Chiu W, Milner A, Saunders JE, and Sheppard RJ. Volume estimation from multiplanar 2D ultrasound images using a remote electromagnetic position and orientation sensor. Ultrasound Med Biol, 1996;22:561-572.

16. Kaspersen JH, Langø T, and Lindseth F. Wavelet-based edge detection in ultrasound images. Accepted for publication in J Ultrasound Med Biol, 2000.

17. Kincaid D, and Cheney W. Numerical analysis: Mathematics of scientific computing. Pacific Grove, CA, Brooks/Cole Publishing Co., 1991.

18. King DL, King DL, and Shao MY-C. Evaluation of in vitro measurement accuracy of a three-dimensional ultrasound scanner. Ultrasound Med, 1991;10:77-82.

19. Langø T, Lie T, Lindseth F, Kleven A, and Gronningsaeter A. Probe calibration. UK. Patent Application 31,15,69543. 1999.

20. Lawson CL, and Hanson RJ. Solving least squares problems. Englewood Cliffs, NJ, Prentice-Hall, 1974.

21. Leotta DF, Detmer PR, Gilja OH, Jong J-M, Martin RW, Primozich JF, Beach KW, and Strandness ED. Three-dimensional ultrasound imaging using multiple magnetic tracking systems and miniature magnetic sensors. IEEE Ultraso Symp, Seattle, 1995;2:1415-1418.

22. Leotta DF, Detmer PR, and Martin RW. Performance of a miniature magnetic position sensor for three-dimensional ultrasound imaging. Ultrasound Med Biol, 1997;23:597-609.

23. More JJ. The Levenberg-Marquardt algorithm: Implementation and theory. Numerical Analysis, Lecture Notes in Mathematics. A Watson, Springer Verlag, 1977;105-116.

24. Ohbuchi R, Chen D, and Fuchs H. Incremental volume reconstruction and rendering for 3-D ultrasound imaging. SPIE - The International Society for Optical Engineering, Chapel Hill, NC, USA, 1992;1808:312-323.

25. Prager RW, Rohling RN, Gee AH, and Berman L. Automatic calibration of 3-D free-hand ultrasound. Cambridge, Cambridge University Department of Engineering, 1997.

26. Prager RW, Rohling RN, Gee AH, and Berman L. Rapid calibration for 3-D freehand ultrasound. Ultrasound Med Biol, 1998;24:855-869.

27. State A, Chen DT, Tector T, Brandt A, Chen H, Ohbuchi R, Bajura M, and Fuchs M. Case study: Observing a volume rendered fetus within a pregnant patient. IEEE Visualization, 1994:364-368.

Paper B

Wavelet-Based Edge Detection in Ultrasound Images

Jon Harald Kaspersen [†], Thomas Langø [†] and Frank Lindseth

SINTEF Unimed, Ultrasound, 7465 Trondheim, Norway

ABSTRACT

We introduce a new wavelet-based method for edge detection in ultrasound images. Each beam that is analyzed is first transformed into the wavelet domain using the continuous wavelet transform (CWT). As the CWT preserves both scale and time information, it is possible to separate the signal into a number of scales. The edge is localized by first determining the scale at which the power spectrum, based on the wavelet transform, has its maximum value. Next, at this scale we find the position of the peak for the squared CWT. This method does not depend on any threshold, once the range of scales have been determined. We suggest a range of scales for ultrasound images in general. Sample edge detections are demonstrated in ultrasound images of straight and jagged edges of simple structures submerged in water bath, and of an abdominal aorta aneurysm phantom.

[†] These authors contributed equally to this article.

INTRODUCTION

Ultrasound imaging is an extremely valuable technique in many areas of medical imaging. Nevertheless, the usefulness of conventional two-dimensional (2-D) B-scan imaging is limited by the presence of different types of noise. One of the main problems in ultrasonic imaging is the strong, tissue dependent dampening of the transmitted ultrasonic pulse, the point spread function or psf. Moreover, there is distortion of the pulse as it propagates through tissues with local variations in sound velocity, multiple reflections or false echoes, and reverberations, as well as a high angle dependency in reflections from layered structures or plane surfaces. The result is that the received echo intensities are strongly dependent on the acquisition conditions, and similar tissues may appear very different in different parts of the same image. Furthermore, due to the coherent detection in ultrasound imaging, constructive and destructive interference between reflections from small, closely localized objects give rise to speckle, which modulates the gray level images, making the intensity vary apparently randomly with position. This gives the ultrasound images the well-known grainy appearance, complicating e.g. edge detection and volume visualization of ultrasound data. The frequency and size of the transmitted pulse roughly determine the spatial variation of the speckle. Therefore, with increasing frequency, and thus in general shorter pulses, the speckle becomes finer grained. Furthermore, differences in speed of sound throughout tissue results in a widening of the psf (Ødegård et al. 1996), i.e. the size of the psf increases as a function of depth in the image. The degree of widening depends on the type of tissue, speed of sound, smooth or rugged tissue boundaries etc., in addition to the frequency and shape of the transmitted pulse.

These drawbacks have prevented successful development of accurate and reliable edge detection and segmentation algorithms for ultrasound images. A reliable and automatic edge detection is of great importance in many applications in ultrasound image processing. In new methods for probe calibration, e.g. in (Prager et al. 1998) and in (Langø et al. 2000), automatic edge detection in a series of two-dimensional (2-D) ultrasound images is a crucial step in the procedure. Probe calibration is the process of determining the position and orientation of the ultrasound image plane relative to a position sensor attached to the ultrasound probe. The main goal of the calibration algorithm is the matching between straight edges in the 2-D ultrasound image planes and those in a three-dimensional (3-D) object. The calibration transformation is an important component in a navigation system based on intraoperative ultrasound for neurosurgery (Gronningsaeter et al. 2000). The edge detection method described in this paper is especially developed for the purpose of automatic edge detection in these ultrasound probe calibration algorithms, and it is used in the calibration technique described in (Langø et al. 2000).

Furthermore, a robust edge detection procedure may be used as an initialization step in various active contour algorithms (Xu and Prince 1998). These algorithms seek the entire edge for closed structures/objects in images by minimizing energy functions. In order to obtain a starting point for the curve to be fitted to the edge, a few edge detections in various directions from a point inside the closed contour is necessary. A robust and automatic edge detection routine is of great importance in this regard. We show some initial results from edge detection in ultrasound images of an aorta phantom.

Reliable segmentation algorithms can also be used in conjunction with 3-D ultrasound imaging in order to find the surface of objects and also for various techniques to provide

quantitative information about properties and measures of tissue. In measurement applications for instance (volume, area, etc.) it can be time consuming to manually mark the edges.

Conventional edge detection techniques include both spatial (mask) and frequency (operator) filters such as the Canny method (Canny 1986), gradient or derivative methods such as the Sobel operator (Gonzales and Woods 1993), the "stick" operator (Czerwinski et al. 1993), the Hough transform (Gonzales and Woods 1993; Hough 1962), graph-theoretic techniques (Gonzales and Woods 1993), and simple derivative methods, e.g. using the maximum of the derivative of the beam signal array. Such methods, however, are not appropriate for ultrasound images as shown in (Linker et al. 1982). They concluded that the standard sequence of edge detection used in image processing, consisting of smoothing, edge operator, and applying a threshold is less suited to ultrasound images than to standard density images. Similar results showing the inadequacy of conventional edge detection techniques for echocardiographic data may be found in (Chu et al. 1988) and (Adam et al. 1987).

More complicated algorithms have been introduced for border detection, e.g. a range of statistical methods such as the Bayesian approach in (Santis and Sinisgalli 1999) or Simulated Annealing in (Friedland and Adam 1989), Markov random field (Dias and Leitao 1996; Friedland and Adam 1989), multiple active contour model (Chalana et al. 1996), artificial neural networks (Hunter et al. 1995; Sun and Tsai 1992; Xin et al. 1992), relaxation principle (Chou et al. 1990), optical flow (Mailloux et al. 1989; Meunier et al. 1988), fuzzy logic (Feng et al. 1991), and morphological filters (Klinger et al. 1988). Most of these methods have been developed for 2-D echocardiographic data. Another more recent semiautomatic contour detection method for M-mode images combining both tissue Doppler and gray-scale data is described in (Rabben et al. 2000).

In order to be able to compare different edge detection methods, a scheme must be established first, for instance following the approaches suggested by (Demigny et al. 1995; Heath et al. 1996). Nevertheless, it is difficult to perform comparison tests since most edge detection methods depend on threshold values and hence the result of a comparison might depend on the specific values implemented. Furthermore, edge detection in ultrasound images is further complicated by the presence of the different kinds of noise mentioned above. Nevertheless, a few publications have been devoted to edge detection or segmentation of ultrasound data using the wavelet transform or multiscale approach, e.g. the work found in (Setarehdan and Soraghan 1999) (Setarehdan and Soraghan 1998). Their method were also specifically developed for echocardiographic data.

In this paper we present a new algorithm based on the continuous wavelet transform (CWT) for detection of an edge in an ultrasound image. The method was, as mentioned before, specifically developed for detecting straight edges in ultrasound images of objects in a water bath for ultrasound probe calibration (Langø et al. 2000). The detection procedure works on one beam at a time, but the beams do not necessarily have to equal the scan beams. The beam to be analyzed may be cut from anywhere in an image and oriented arbitrarily. In summary, the beam is first transformed into the wavelet domain by using the CWT. Next, the edge is localized by determining the position of the maximum value for the squared CWT at the most energetic scale. The most energetic scale is established by computing the power spectrum based on the CWT modulated by the square of the scale value, i.e. the energy contribution at each scale as a function of scale. The position for the maximum from this energy distribution is taken as the most energetic scale. One of the main advantages of the algorithm is that it, in its simplest form, is independent of any threshold. For automatic edge detection methods this

is of great advantage and importance in order to minimize complex input parameters. In addition, an edge detection method for ultrasound images needs to be robust with respect to noise.

MATERIALS AND METHODS

Data acquisition

The images used in this study were acquired using a 4-8 MHz phased array probe on a high-end digital ultrasound scanner (System Five, GE Vingmed Sound Inc., Horten, Norway). The objects scanned were a thin planar woven nylon mesh submerged in a water tank, a thin jagged membrane aligned perpendicular at the bottom of the tank, and an abdominal aortic aneurysm phantom. An illustration of the set-up is shown in Fig. 1 with the thin nylon mesh membrane.
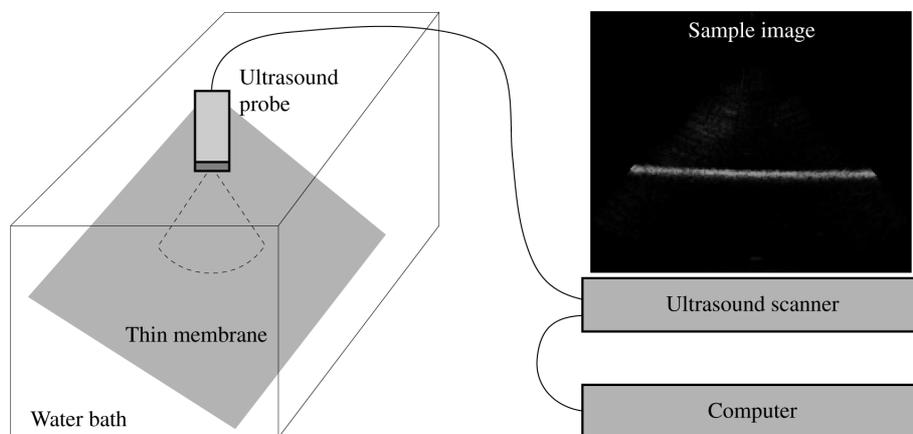

Fig. 1. Data acquisition set-up illustrating the imaging of the thin planar nylon mesh membrane submerged in a water bath, the ultrasound probe, and a sample image of the membrane as appearing on the scanner.

The water in the tank was either fresh from the tap, resulting in an excess of tiny air bubbles, or vacuumed tap water for very clear images. The scanned images were transferred to an external computer (Power Macintosh from Apple Inc., CA, USA) for processing. No image filters were applied prior to edge detection in order to test the robustness of the method with respect to noise. Furthermore, we used non-scan converted (Angelsen 1995) images since this makes it easier to visualize the individual beams from sector scan probes. It also makes it easier to comment on the results by referring to individual beams in the images. Hence a scan of a planar structure will not be seen as a straight line in the images below, but rather curved.

Wavelet theory

In contrast to the Fourier transform, which is global and provides a description of the overall regularity of a signal, the wavelet transform provides both time and scale information of the signal and is well adapted for finding the location and spatial distribution of singularities (Mallat and Hwang 1992). If, $W_n(\tau,a)$ is the wavelet coefficient of ultrasound beam number n at position τ and scale a, the square of the wavelet coefficients, $[W_n(\tau,a)]^2$, is proportional to the energy in the signal, as will be shown later. An edge is often characterized as a jump in intensity over a distance and can therefore be thought of as an energetic large-scale structure.

The basis of the wavelet transform is that the signal, which is to be analyzed, is compared with a function Ψ, often referred to as the mother wavelet. The wavelets are translated in time, or space, by a variable shift τ and scaled by an amplitude parameter, the scale a, so that (Rioul and Vetterli 1991)

$$\Psi_{\tau,a}(t) = \frac{1}{\sqrt{a}}\Psi\left(\frac{t-\tau}{a}\right) \tag{1}$$

where t represents time or sample number along the beam to be transformed. The wavelet transform is defined by

$$W_f(\tau,a) = \frac{1}{\sqrt{a}}\int_{-\infty}^{\infty} f(t)\Psi\left(\frac{t-\tau}{a}\right)dt \tag{2}$$

where $f(t)$ is the signal to be analyzed. The wavelet transform may be seen as a microscope used to visualize a signal at various magnifications and/or locations. The scale parameter, a, then represents the magnification and the translation parameter, τ, represents the location chosen for the observation.

The mother wavelet can be any function, real or complex, satisfying the "admissibility" condition (Chui 1992) expressed as

$$C_\Psi = \int_{-\infty}^{\infty}\frac{\left|F_\Psi(\omega)\right|^2}{\omega}d\omega < +\infty \tag{3}$$

where C_Ψ is known as the admissibility constant, and $F_\Psi(\omega)$ is the Fourier transform of $\Psi(t)$. This condition implies that (Mallat and Hwang 1992)

$$\int_{-\infty}^{+\infty}\Psi(t)dt = 0 \tag{4}$$

In singularity detection, i.e. detection of function discontinuities, one of the most important qualities for the wavelet is the number of vanishing moments (Mallat and Hwang 1992). According to (Mallat and Hwang 1992), a wavelet with one vanishing moment is sufficient in detecting function discontinuities. However, as mentioned in (Liandrat and Moret-Bailly 1990), every wavelet family, i.e. the choice of mother wavelet, is efficient for discontinuity detections. A common mother wavelet with two vanishing moments is the second derivative of the Gaussian distribution, often referred to as the Mexican hat wavelet. We chose this wavelet due to the fact that it is a real function (many wavelets are complex functions) and the ease of implementation, i.e. it is readily available in many available software products. The Mexican hat wavelet reacts to the second derivative of a signal (Liandrat and Moret-Bailly 1990). A consequence is that a zero crossing of the CWT using the Mexican hat as the mother wavelet would correspond to an edge or discontinuity (inflection point) in the signal to be analyzed. Applying the zero crossing detection scheme means that both the shift from black to white and the shift from white to black (a white edge on black background) would be detected. Therefore we found it convenient to identify the maximum or local maximum of the squared CWT (see later). In an ultrasound image of a thin membrane this results in the

detection of the center of the membrane. The Mexican hat wavelet is given by (Daubechies 1992)

$$\Psi(t) = \frac{2}{\sqrt{3}}\pi^{-\frac{1}{4}}\left(1-t^2\right)e^{-\frac{1}{2}t^2}$$

(5)

By defining the scales as

$$a(j) = 2^{j/N}$$

(6)

where j and N are integer values, we obtain the Mexican hat wavelet for different scale values. Fig. 2 shows a sample plot of the Mexican hat for $\tau=256$ and $a=2^{20/4}$.
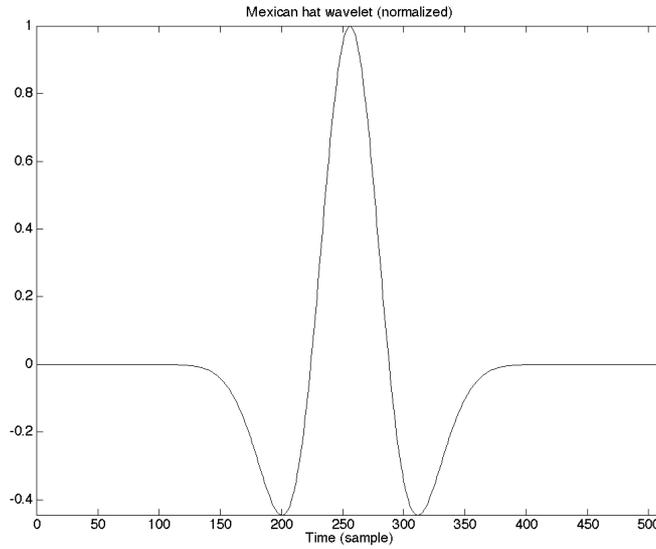

Fig. 2. Plot of the normalized Mexican hat wavelet (see Eq. 5) with $\tau=256$ and $a=2^{20/4}$.

It can be argued mathematically that the Mexican hat wavelet gives a "measure" of the second derivative of the analyzed signal (Farge et al. 1996; Liandrat and Moret-Bailly 1990). In this study the Mexican hat will be used as the mother wavelet and the admissibility constant for this wavelet can be obtained analytically as (Kaspersen 1996)

$$C_\Psi = \int_0^\infty \frac{|F_\Psi(\omega)|^2}{\omega}d\omega = \int_0^\infty 2\pi\omega^3 e^{-\omega^2}d\omega = \pi\Gamma(2) = \pi$$

(7a)

where Γ is the gamma function, defined as

$$\Gamma(x) = \int_0^\infty e^{-t}t^{x-1}dt = \int_0^1\left(\ln\frac{1}{t}\right)^{x-1}dt$$

(7b)

It is possible to compute the energy spectrum for a time series t by means of wavelets. The power spectrum measures the amount of energy as a function of frequency, or scale, in a time series. In (Hudgins 1992) the author showed that the wavelet power spectrum at a scale a is the weighted average of the Fourier power spectrum over all frequencies. The weighting

function is the Fourier power spectrum of the mother wavelet at scale a. The wavelet scale dependent spectrum is defined by

$$S_f(a) = \int\limits_{-\infty}^{\infty} \left| W_f(\tau,a) \right|^2 d\tau \tag{8}$$

and is a measure of the distribution of energy of the signal $f(t)$ as a function of scale. The energy of the signal $f(t)$ is given by

$$
\begin{aligned}
E_f &= \int \left| f(t) \right|^2 dt = \iint \left| W_f(\tau,a) \right|^2 \frac{d\tau da}{a^2} \\
&= \int \left[\int\limits_{-\infty}^{+\infty} \left| W_f(\tau,a) \right|^2 d\tau \right] \frac{da}{a^2} \\
&= \int \frac{S_f(a)}{a^2} da
\end{aligned}
\tag{9}
$$

In other words, the spectrum in Eq. 8 divided by the square of the scale is the energy contribution for that scale, while the square of the wavelet coefficients represents the local energy content. Eq. 9 is utilized in the proposed detection procedure below.

Detection algorithm
Prior to processing we subtract the mean of each beam from every beam value to suppress the possible linear bias in the energy of the signal. Two effects can explain this bias. First, the time gain compensation technique in ultrasound in which the intensity of deeper situated structures are increased in order to compensate for the weaker echoes, causes increased intensity also for the noise from deeper layers in the image. Second, from ultrasound theory we know that the beam thickness increases as a function of depth, and with approximately equally sized noise structures, e.g. tiny air bubbles in water evenly distributed in the imaged medium, we get larger echoes, or smeared out representations from the deeper situated noise. Therefore, along a chosen beam the intensity distribution will be subject to a certain linear bias, and bubbles may cause an increased intensity as a function of depth.

The first step in the edge detection algorithm is the calculation of the CWT scalogram, which shows the CWT as a function of both scale and sample for one beam (see Box 1 for a summary of the algorithm with examples). We use the middle beam for this purpose from the ultrasound image throughout this paper. The input to this calculation is the beam, the mother wavelet (Eq. 5), and the scale array. We chose the scale array

$$
\begin{aligned}
a(j) &= 2^{j/N} \\
j &= 1,...,40 \\
N &= 4
\end{aligned}
\tag{10}
$$

which was adopted from turbulence theory (Kaspersen 1996). This range of scales should be sufficient to cover most structures encountered in ultrasound imaging. However, the scale array may be seen as a threshold parameter for the procedure. We elaborate on this in the discussion.

Next, the wavelet spectrum in Eq. 8 was calculated (see Box 1), i.e. the sum of the squared CWT along each scale in the scalogram. This integral was modulated by multiplying it with the inverse of the square of the scale according to Eq. 9. The result is the energy contribution at a particular scale as a function of scale. The scale to perform detections on is then given by the scale value at which we find the maximum of the spectrum modulated by the square of the scale, i.e.

$$\max\left[\frac{S_f(a)}{a^2}\right] = \max\left[\frac{1}{a^2}\int_{-\infty}^{\infty}\left|W_f(\tau,a)\right|^2 d\tau\right] \tag{11}$$

This optimal scale was then used to calculate the mother wavelet, i.e. the Mexican hat at this scale. Alternatively, the Mexican hat for this scale can be retrieved from the earlier calculations if it was stored. Further, this Mexican hat was used to calculate the CWT of the current beam. The edge was then found as the maximum of the squared CWT at this scale (see Box 1).

At this stage we might calculate the CWT and find the edge for every beam constituting the ultrasound image. However, as this calculation is somewhat time consuming and computationally expensive, depending on the application, it might be sufficient to process only a certain percentage of the beams. Since we wanted to develop an edge detection routine for a probe calibration method (Langø et al. 2000) that uses a planar membrane in a water bath, we already knew that the edge was straight. Hence, we found it sufficient to detect the edge in approximately every twentieth beam and estimate the correct line from these detections.
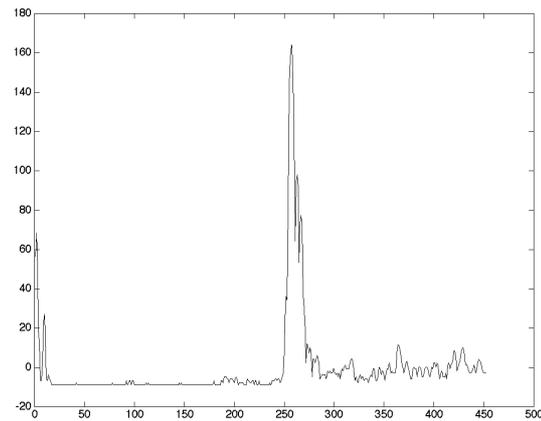
Further details of the underlying theory and discussion of the background theory for the detection procedure can be found in (Kaspersen 1996). This detection algorithm will be denoted as the wavelet peak energy method or simply the WPE method.

Box 1. Detection algorithm.

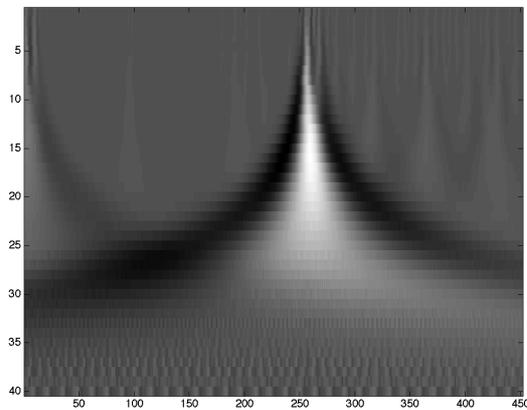Set signal *f(t)* equal to middle beam (subtracted the mean value calculated from the same beam).
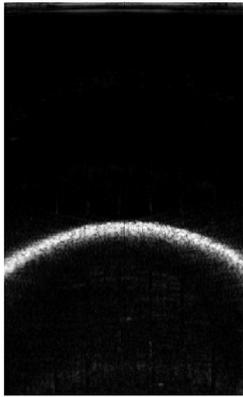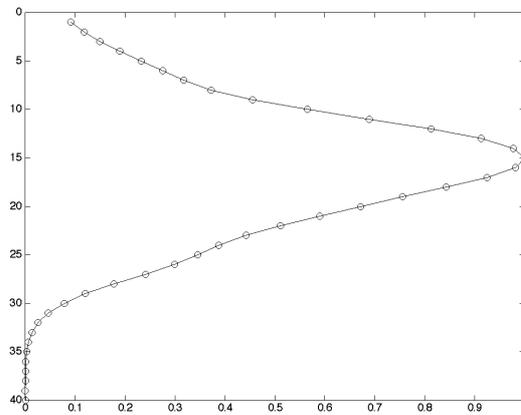
Ultrasound image.   *Middle beam.*

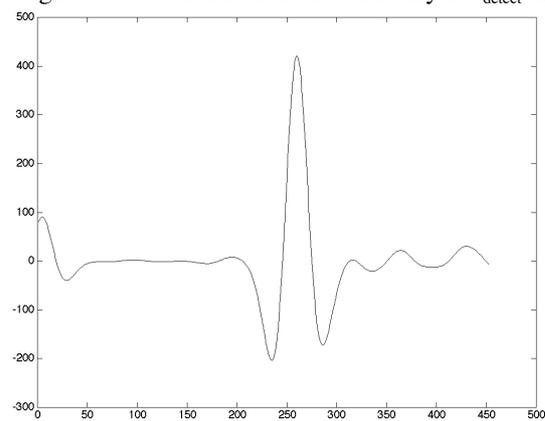Calculate $\dfrac{S_f(a)}{a^2}$ using Eq. 2 inserted into Eq. 8.

CWT as a function of both scale and beam.   *Normalized power spectrum modulated by $1/a^2$.*

Calculate $a_{\det ect} = \max\left[\dfrac{S_f(a)}{a^2}\right]$ from Eq. 11, Mexican hat for a = a_{detect} from Eq. 5 and 6, and CWT of signal for current Mexican hat and only at a_{detect}. The edge is found at the maximum of the CWT.

CWT as a function of sample at the detection scale (one row from the CWT plot above).

RESULTS

The exact value for the estimated optimal scale in all the figures below may be found by counting the circles from the top in the modulated spectrum (plots in *c*), starting at 1, until reaching the horizontal line indicating the maximum of the modified spectrum.

Fig. 3 shows edge detection for a very simple case. During acquisition the probe was oriented approximately perpendicular to the thin nylon mesh membrane submerged in a water bath. The water was vacuumed in order to remove air bubbles, i.e. to obtain a low noise level. Fig. 3a shows a plot of the normalized middle beam of the ultrasound image in Fig. 3d, indicated with a white vertical line. In Fig. 3b the scalogram is shown, i.e. a visualization of the CWT as a function of sample and scale, while the wavelet spectrum $S_f(a)$ modulated by a^2, is given in Fig. 3c. The resultant detections superimposed on the original image is shown in Fig. 3d.
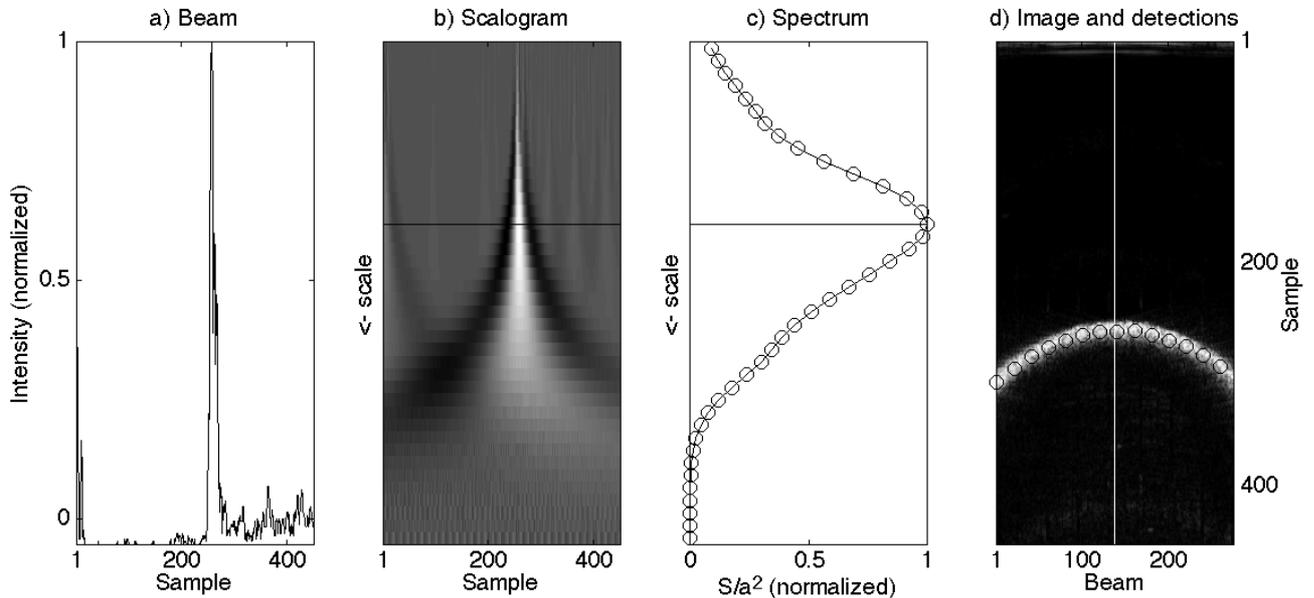

Fig. 3. Acquisition with the probe oriented approximately normal to a thin planar nylon mesh membrane in a vacuumed water tank. a) Intensity plot of beam indicated by vertical white line in d. b) CWT scalogram for beam as a function of scale and sample. The finer scales are at the top and varies linearly along the vertical from 1 to 2^{10}. c) Integrand of Eq. 9, i.e. contribution to total energy at each scale as a function of scale. d) Ultrasound image with detections superimposed as small black circles for every twentieth beam.

The large-scale energetic structures are clearly shown in Fig. 3b. By comparing Fig. 3b with Fig. 3c it can be seen that the peak of the spectrum corresponds to the scale where the edge is the dominant structure in the CWT scalogram. Hence, when we look for an edge in an image, i.e. the middle of a black-white-black area in our case, we search for the local energetic structures in each beam. For a simple case as in Fig. 3d, we see from Fig. 3b that a range of scales around the optimal will find the edge properly. However, for more complex images, this may not be the case, and it is of crucial importance to find the correct scale.

Fig. 4 shows a similar case as in Fig. 3, except that we have used fresh water containing a lot of tiny air bubbles, representing excessive noise in this case. We can see from the scalogram in Fig. 4b that there are several structures at higher frequencies but that the most energetic is

the membrane, and hence the detections are correct. The spectrum shows that there is a fair amount of energy in the bubbles, represented by the lower peak in the upper part of Fig. 4c. In addition, there is a slight amount of energy at the lower frequency end of the spectrum, representing the smeared out noise in the deeper areas of the image.
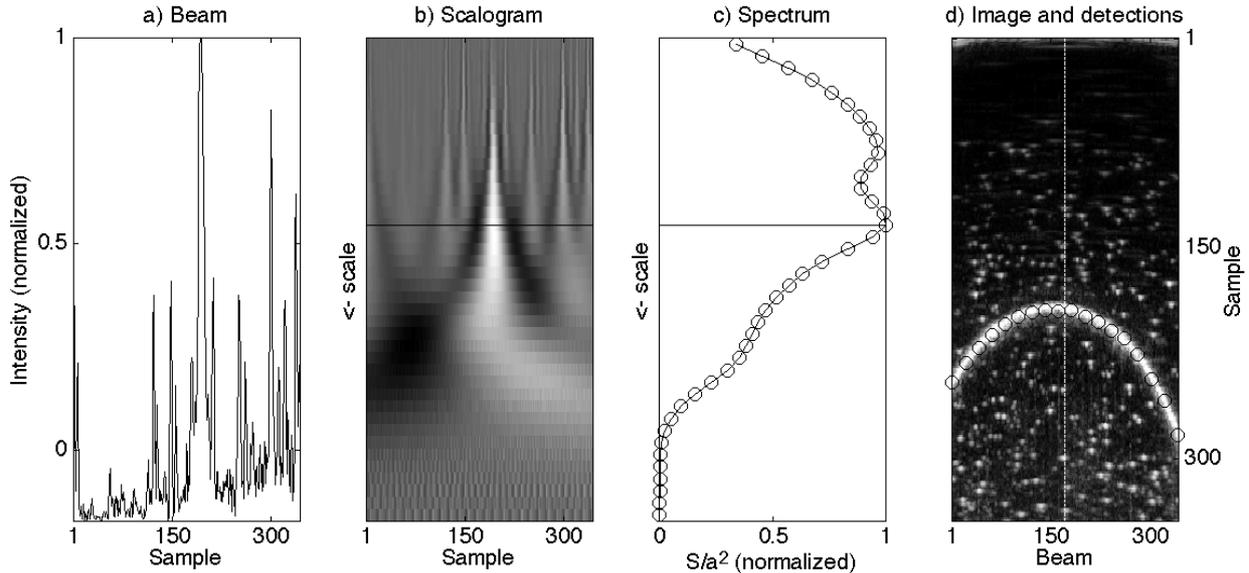

Fig. 4. Acquisition with the probe oriented approximately normal to a thin planar nylon mesh membrane in a tank with fresh tap water, resulting in a lot of tiny air bubbles. a) Intensity plot of beam indicated by vertical white line in d. b) CWT scalogram for beam as a function of scale and sample. The finer scales are at the top and varies linearly along the vertical from 1 to 2^{10}. c) Integrand of Eq. 9, i.e. contribution to total energy at each scale as a function of scale. d) Ultrasound image with detections superimposed as small black circles for every twentieth beam.

In Fig. 5, we have analyzed a different beam from the one processed in Fig. 4d in more detail by separating the signal representing the membrane and the remaining part, which is mainly high frequency noise, probe casing reverberations at the very top of the image, and tiny air bubbles in the water. The selected beam contains an intensity peak from a bubble that is larger than the intensity peak from the membrane. Fig. 5a and Fig. 5e show respectively the beam intensity plots for the entire beam and the entire beam except the signal from the membrane, which was set to zero. Fig. 5b and Fig. 5f show the corresponding CWT scalograms, Fig. 5c and Fig. 5g show the spectrums, and Fig. 5d and Fig. 5h show the ultrasound images with detections superimposed as small white squares with black crosses inside for every twentieth beam. It can be see that the maximum peak in the spectrum is shifted one scale up when the membrane contribution is removed from the signal, i.e. when the high frequency contribution of the bubbles are the most energetic structures left in the signal. Even though the intensity peak of the one bubble is higher than the intensity peak of the membrane for the selected beam, the most energetic structure in the spectrum is the membrane due to the width of this intensity peak. In addition, when the membrane is removed from the signal, only seven of the eighteen beams used for edge detection hit the membrane, while for Fig. 5a-d, when the membrane signal is included in the signal, we get twelve correct detections. The estimated scales were a=7 and a=6. In other words, the scales are approximately estimated to the same value for the entire signal and for only the bubbles plus noise signal.
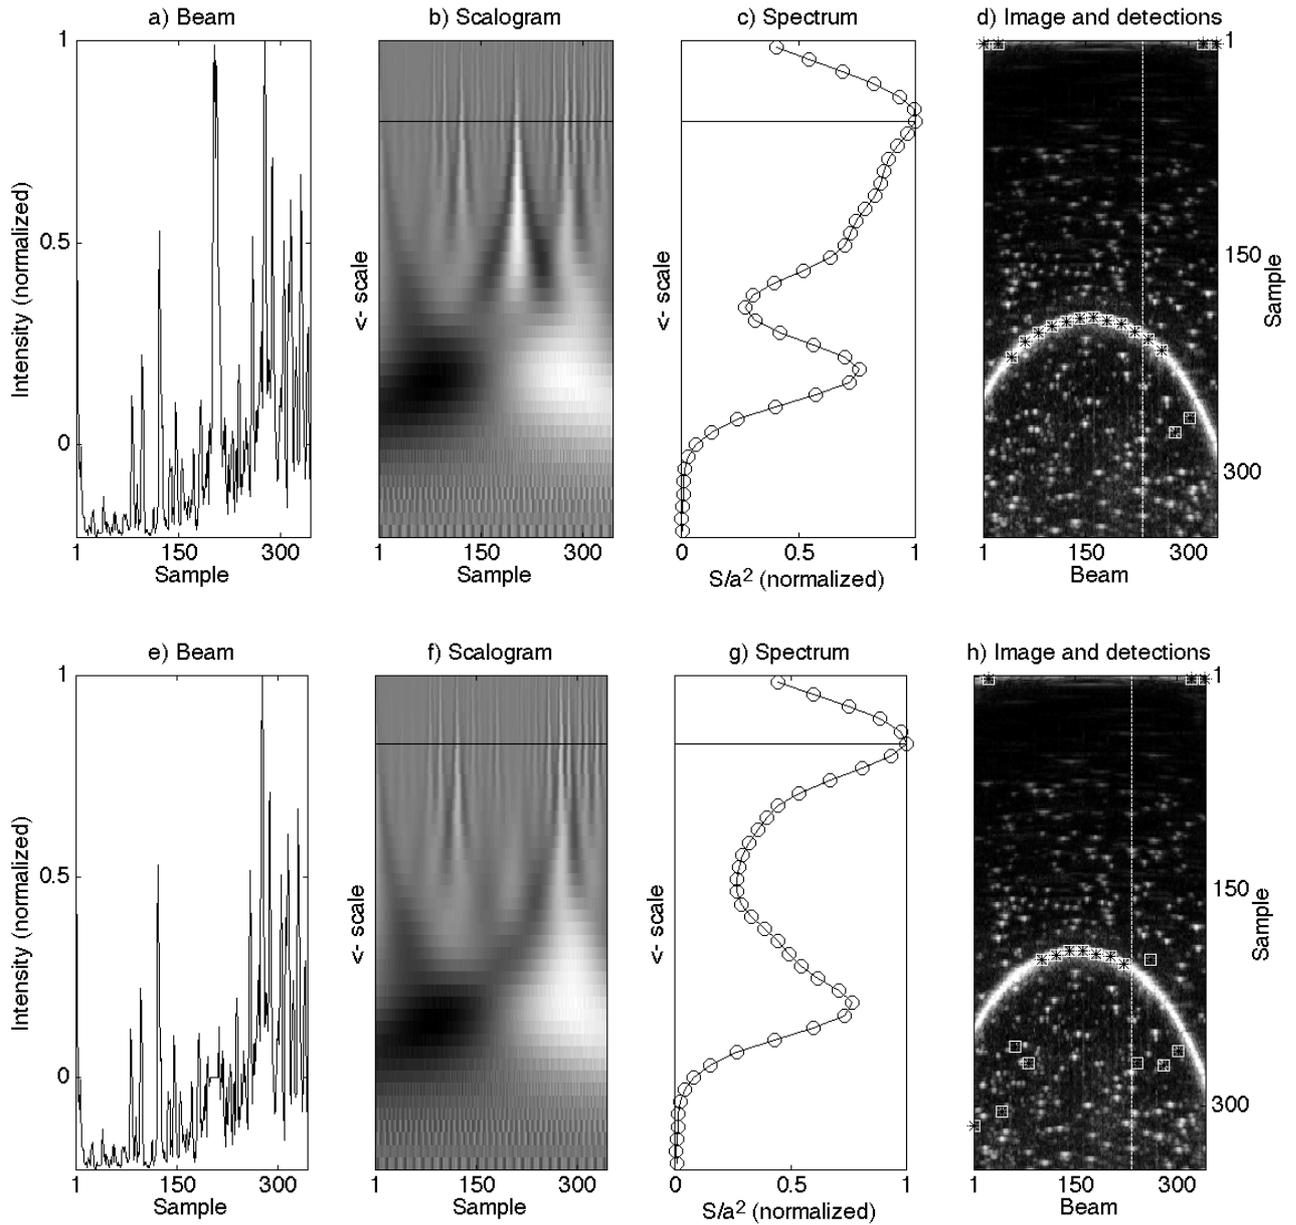
Fig. 5. Acquisition with the probe oriented approximately normal to a thin planar nylon mesh membrane in a tank with fresh tap water, resulting in a lot of tiny air bubbles. a) Intensity plot of beam indicated by vertical white line in d. b) CWT scalogram for beam as a function of scale and sample. The finer scales are at the top and varies linearly along the vertical from 1 to 2^{10}. c) Integrand of Eq. 9, i.e. contribution to total energy at each scale as a function of scale. d) Ultrasound image with detections superimposed as small white squares with black crosses inside for every twentieth beam. e-h: As a-d, except the intensity peak representing the membrane has been removed from the signal (set to zero).

The image in Fig. 6d was acquired with the membrane lowered in the water bath leading to more and stronger reverberations from the bottom of the tank. In this case there is a slight ambiguity of the location of the membrane. Nevertheless, the most energetic structure is the real membrane. The reverberations, especially the deepest one, can be seen on both the scalogram and in the spectrum as a small increase at lower frequencies than the real membrane.

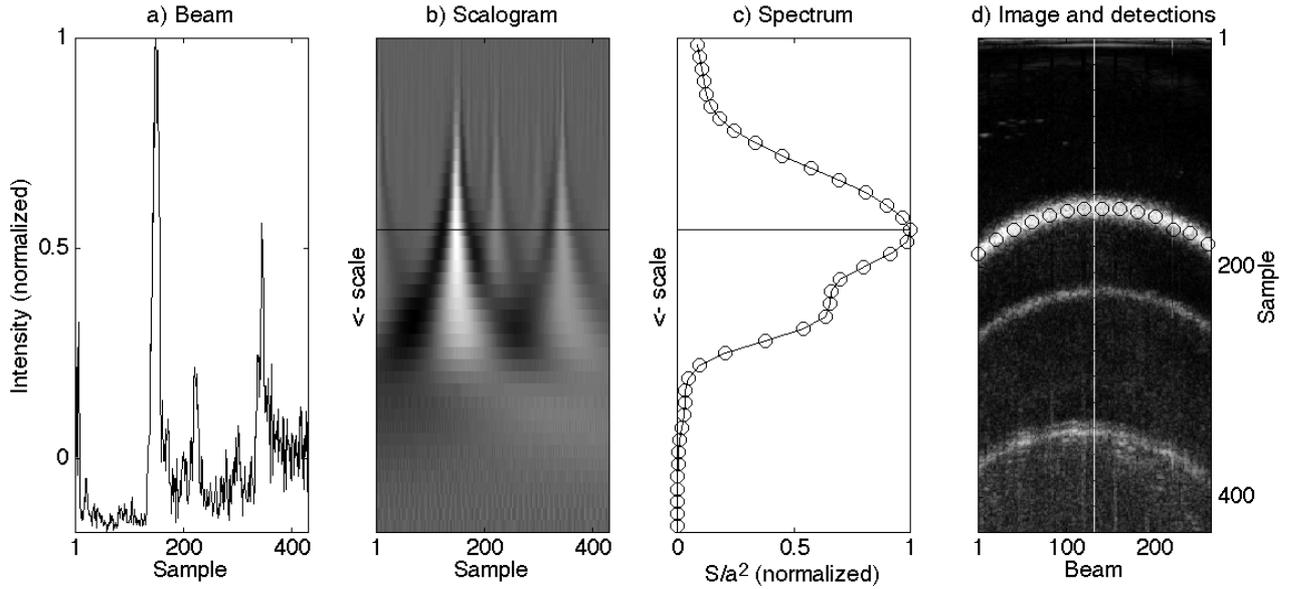a) Beam b) Scalogram c) Spectrum d) Image and detections

Fig. 6. Acquisition with the probe oriented approximately normal to a thin planar nylon mesh membrane in a vacuumed water tank. The membrane (and probe) was lowered in the tank resulting in reverberations from the bottom of the tank. a) Intensity plot of beam indicated by vertical white line in d. b) CWT scalogram for beam as a function of scale and sample. The finer scales are at the top and varies linearly along the vertical from 1 to 2^{10}. c) Integrand of Eq. 9, i.e. contribution to total energy at each scale as a function of scale. d) Ultrasound image with detections superimposed as small black circles for every twentieth beam.

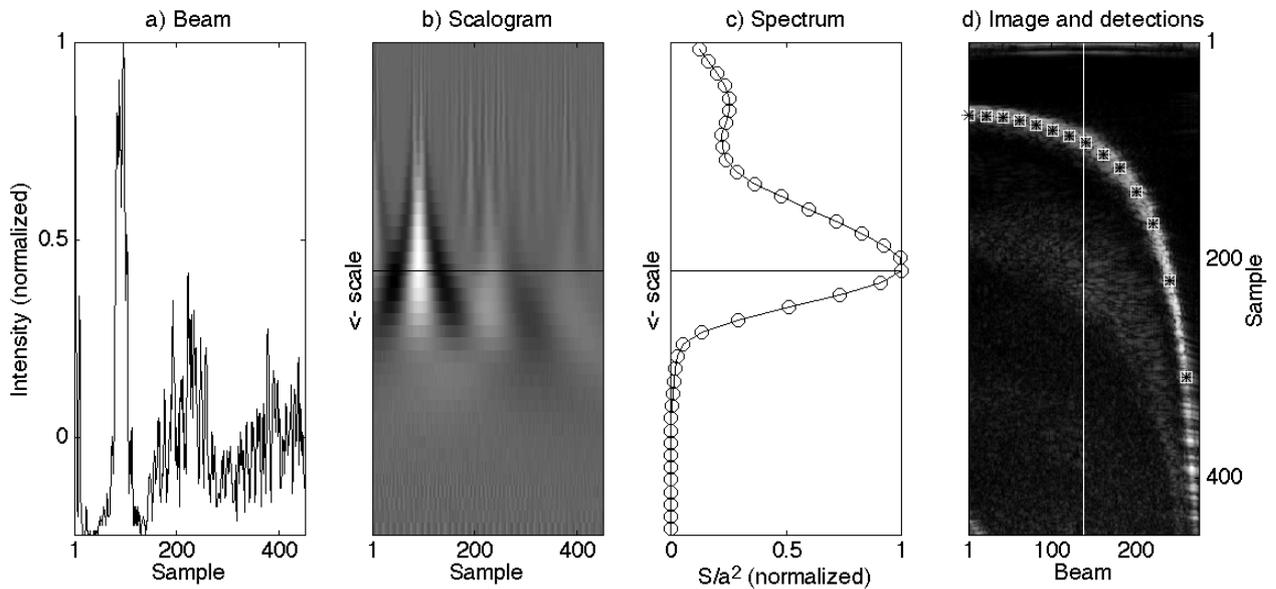a) Beam b) Scalogram c) Spectrum d) Image and detections

Fig. 7. Acquisition with the probe tilted about the image plane normal axis (relative to the perpendicular orientation in Fig. 3). The object is a thin planar nylon mesh membrane in a vacuumed water tank. a) Intensity plot of beam indicated by vertical white line in d. b) CWT scalogram for beam as a function of scale and sample. The finer scales are at the top and varies linearly along the vertical from 1 to 2^{10}. c) Integrand of Eq. 9, i.e. contribution to total energy at each scale as a function of scale. d) Ultrasound image with detections superimposed as small white squares with black crosses inside for every twentieth beam.

The image in Fig. 7d was acquired with the probe tilted about the axis normal to the image plane. We can see that the edge is almost discontinuous at the far right. The high frequency

reverberation noise has a much lower intensity than the membrane, and hence is not detected. Furthermore, the lower frequency noise in the lower parts of the image has much too low intensity to even be seen after the modification of the spectrum with the inverse of the squared scale.

In Fig. 8d the probe was tilted about the lateral axis of the image. The wide definition of the membrane in the image is due to the finite beam width in ultrasound. We see that the peak in the spectrum is at a much lower frequency than compared to the case in Fig. 3c. There is also some low frequency noise in this image, both in front of and at the far side, the radial direction, of the membrane.
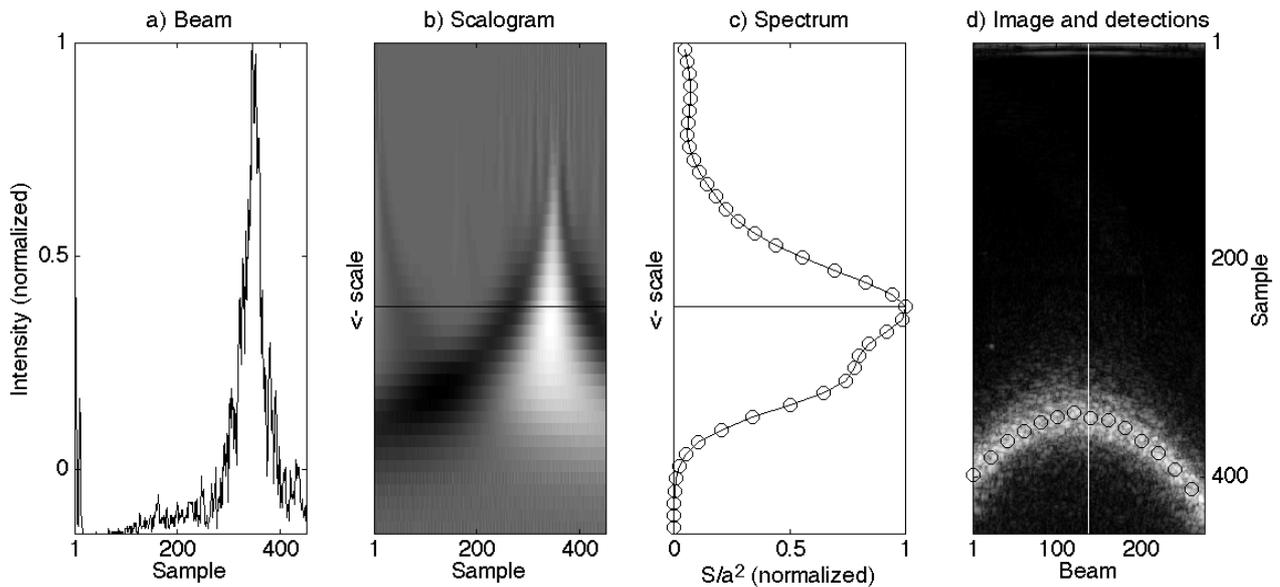

Fig. 8. Acquisition with the probe tilted about the image plane lateral axis (relative to the perpendicular orientation in Fig. 3). The object is a thin planar nylon mesh membrane in a vacuumed water tank. a) Intensity plot of beam indicated by vertical white line in d. b) CWT scalogram for beam as a function of scale and sample. The finer scales are at the top and varies linearly along the vertical from 1 to 2^{10}. c) Integrand of Eq. 9, i.e. contribution to total energy at each scale as a function of scale. d) Ultrasound image with detections superimposed as small black circles for every twentieth beam.

We tested the algorithm for images tilted about the normal and lateral image axis with tiny air bubbles in the water, as for Fig. 4d, and these also produced correct edge detections.

The image shown in Fig. 9d and Fig. 9h was acquired with the probe oriented approximately perpendicular to the water tank bottom. The object scanned was a thin jagged membrane oriented approximately coplanar with the ultrasound image scan plane. In Fig. 9a-d, the middle beam, indicated with vertical white line in Fig. 9d, was used to estimate the detection scale, while in Fig. 9e-h, beam number 25 was used. We can see that the reverberations in the radial direction in the middle of the image from the bottom floor of the water tank causes the scale to be underestimated for the case in Fig. 9a-d. Consequently, a few false detections appear in this case. This fact can be understood by looking at the scalogram, where for a-d there is a certain ambiguity about the location of the jagged membrane, especially for scales less than about 13. The scalogram for beam 25 (Fig. 9f) however, is unambiguous regarding the location of the jagged membrane, and all detections represents the correct edge (Fig. 9h).
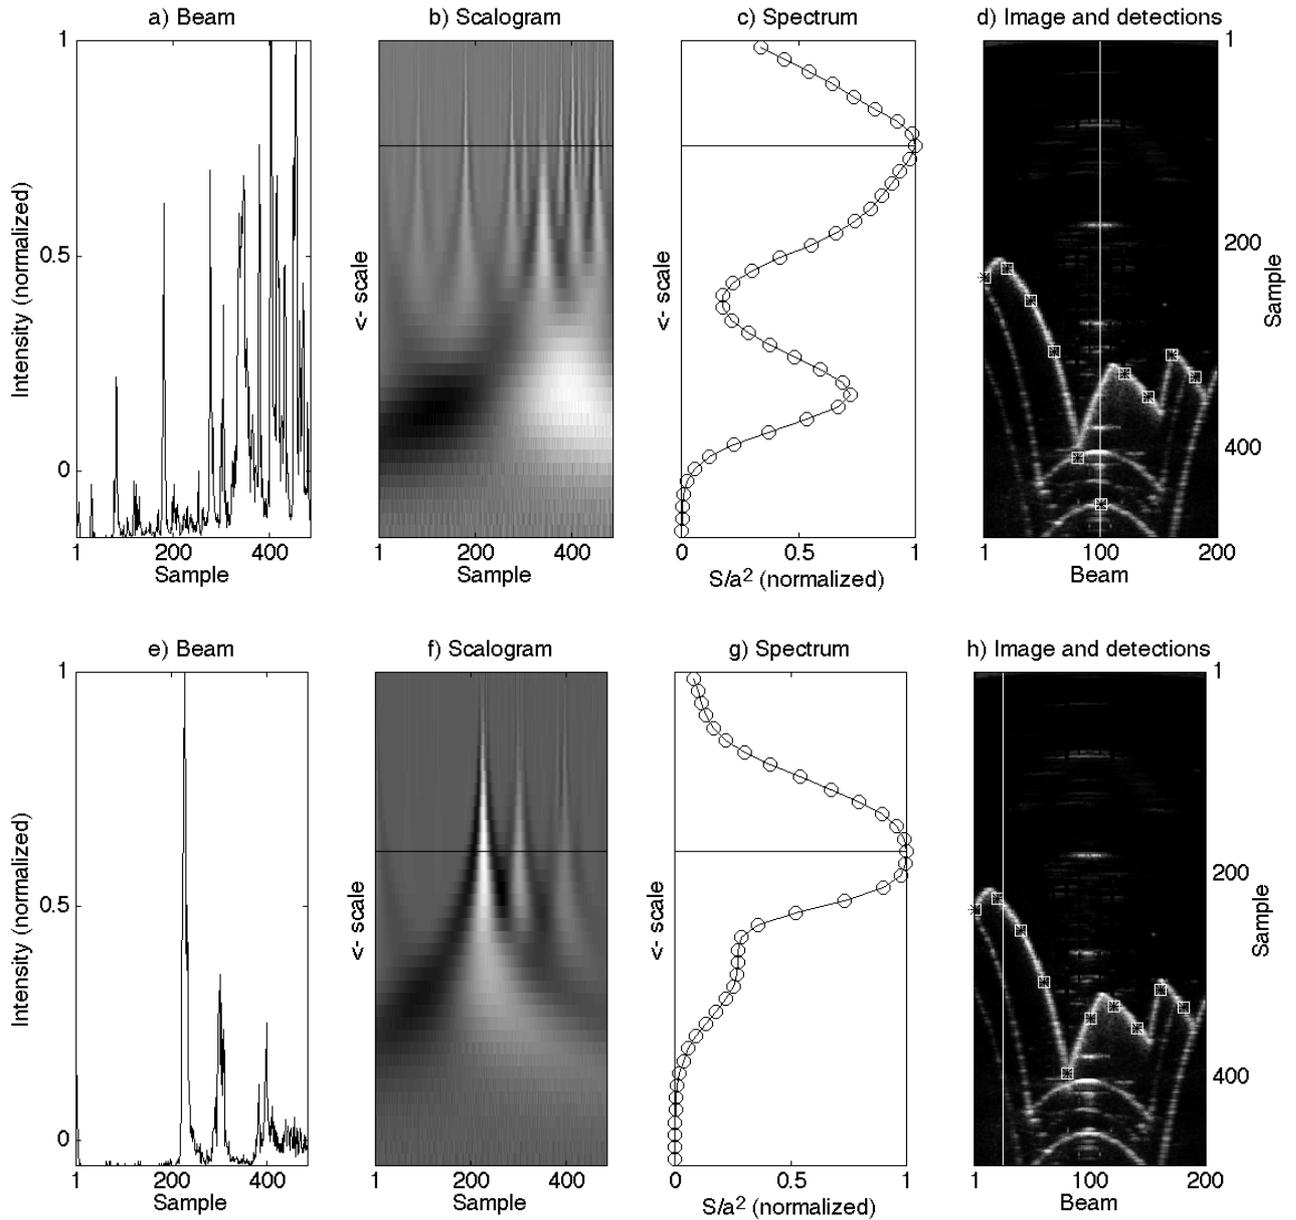
Fig. 9. Results using an ultrasound image of a thin jagged membrane aligned vertically in a water bath and approximately coplanar with the image plane. a) Intensity plot of beam indicated by vertical white line in d. b) CWT scalogram for beam as a function of scale and sample. The finer scales are at the top and varies linearly along the vertical from 1 to 2^{10}. c) Integrand of Eq. 9, i.e. contribution to total energy at each scale as a function of scale. d) Ultrasound image with detections superimposed as small white squares with black crosses inside for every twentieth beam. e-h: As a-d, except beam 25 was used for processing.

Finally, we tested the algorithm for an image of an abdominal aorta aneurysm phantom (Fig. 10). This image was scan converted prior to edge detection since it illustrates the use of our algorithm as an initialization procedure in an active contour algorithm. The analyzed beams are not the ultrasound scan beams in this case, but rather the dotted lines indicated in Fig. 10a. The seed point, which has to be set manually, is shown as a solid white dot in the middle of Fig. 10a. Edge detections are specified with large white circles, while the initial contour is simply the connected closed loop between the edge detections. The image to the right shows only the result after the initialization procedure and a few iterations of a snake algorithm (Xu and Prince 1998). The next step for this algorithm could be to calculate the area inside the

closed loop. By repeating this for several scans from a three-dimensional acquisition of the aneurysm, it is possible to estimate the volume of the aneurysm.
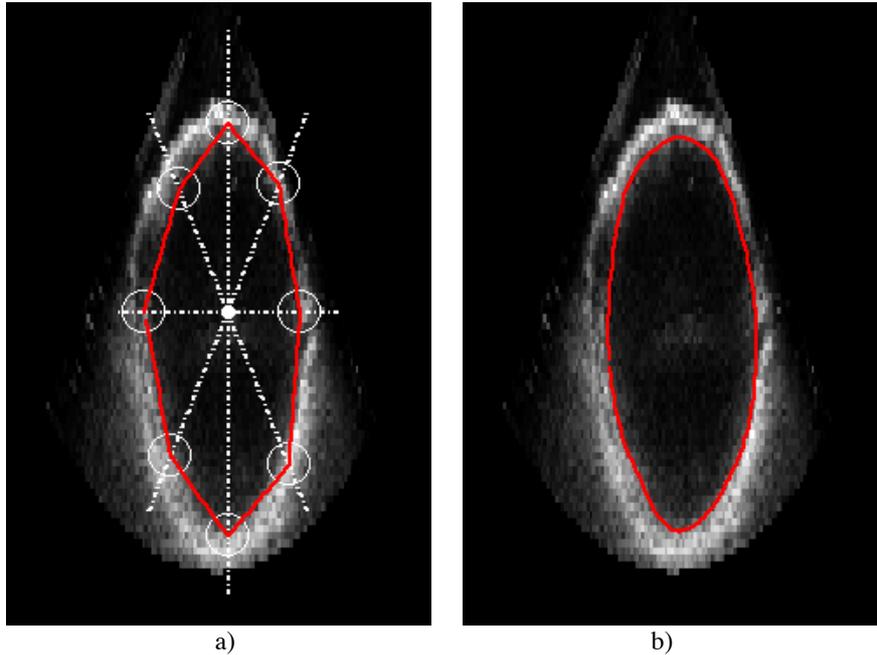

Fig. 10. Cross sectional ultrasound scan of an abdominal aortic aneurysm phantom (scan converted image). a) The seed point of an active contour algorithm (solid white dot) is located approximately in the middle of the aneurysm. The dotted rays in white in eight different directions from seed point are beams processed for edge detection. The edge detections are at the centers of the large white circles. b) The connected white line is the result after one iteration of an active contour algorithm when starting with the initialization from a.

DISCUSSION

We have demonstrated a novel approach to edge detection, and shown several examples of the quality of this method. The proposed technique is well suited to detect edges in ultrasound images of varying quality. Though the method was particularly developed for a probe calibration method (Langø et al. 2000), we found it to be useful in initial trials as an initialization method for an active contour algorithm.

Several interesting general observations of the detection algorithm can be made on the basis of the results above. First, it is found that the algorithm correctly detects the membrane for almost every beam used. The false detections are made at beams in which two or more occurrences of structures of approximately the same size and intensity are found. This might be the case if for instance a noise structure is located in front of the membrane, hence reducing the pulse intensity reaching the membrane and the echo from it. Nevertheless, it can be seen that there are no false detections because of small-scale high intensity noise, which is very important for a robust edge detector in ultrasound imaging. We deliberately acquired some images while the water bath still had a significant amount of tiny air bubbles (fresh tap water) to get images with a high level of noise, simply to test the robustness of the method. Furthermore, we tilted the probe about the image plane normal (Fig. 7d) and the image plane lateral axis (Fig. 8d) to acquire wide and oblique oriented edge images. Even for those images, the detections were good. These cases with the probe oriented at an oblique angle relative to the membrane are important cases in probe calibration (Langø et al. 2000) (Prager

et al. 1998), where it is of interest to scan the object (membrane) from various distances and angles.

Moreover, as in Fig. 5h, if the edge detection is based on the scale being found from just one beam, the result may not be very good. A simple solution to this problem is to estimate the scale from several beams, using the value that occurs most often. We tried this with great success, but this approach increases processing time considerably.

The accuracy of the method is difficult to quantify due to the fact that we do not know the correct position of the edges. Therefore, a certain degree of observer input would have to be used for evaluating the results. During development of the probe calibration method (Langø et al. 2000) we tested the edge detection method described in this paper and a simple derivative method, which found an edge candidate by simply calculating the maximum of the derivative curve for each beam. In order to be able to estimate the edge from the first peak of the derivative maximum, we had to remove the reverberations at the very top of the image, which are the result of reverberations in the probe casing. The methods were tested by calculating the edges for approximately 200 images that were very much similar to the images in Fig. 3, 6, 7, and 8. The evaluation or comparison showed that there were many more poor edge detections for the derivative method. While the edge detections using the wavelet method resulted in a smooth connected line (curve in the non-scan converted images), the derivative method yielded lines made up of points that were scattered about the true line. In addition, when estimating the straight line from the detections the derivative method yielded erroneous lines in approximately 20% of the images, while we observed none for the wavelet detection technique. The erroneous lines were mainly a result of noise being detected as an edge candidate. An more detailed accuracy analysis of the probe calibration method utilizing this method may be found in (Langø et al. 2000).

If the membrane should happen to not be present in parts of the image, e.g. because of a large probe tilt angle, which was almost the case to the far right in Fig. 7d, then the edge detections in those part of the image most likely will represent noise. For the probe calibration techniques in (Langø et al. 2000; Prager et al. 1998) this will not be a problem since we have prior knowledge of the edges being straight, and thus these false detections can easily be excluded from further processing by utilizing a smart straight line fitting algorithm, e.g. as in (Langø et al. 2000).

Nevertheless, several improvements can be implemented to further develop the algorithm. First of all, it would make the method more robust if we estimated the scale based on several beams instead of just one, as mentioned above. However, as more beams are included for scale estimation prior to the edge detection itself, the method becomes more time consuming. For probe calibration (Langø et al. 2000; Prager et al. 1998) performed prior to surgery in the laboratory, the time issue is not a critical factor. On the other hand, for real time processing, e.g. segmentation of data during surgery for visualization, limiting the time spent on image processing is of crucial importance.

We implemented the procedure partly in LabVIEW (National Instruments Inc., CA, USA) and partly in C code. All of the steps in the algorithm could easily be implemented in C code, thereby speeding it up somewhat. Furthermore, the algorithms themselves could be optimized, and this is subject to further investigation. The approximate time for processing (all calculations in Box 1) a 276 beams by 452 samples ultrasound image, only detecting the edge in every twentieth beam, was 90 seconds on a Power Macintosh G4 450 MHz computer

(Apple Inc., CA, USA). A normal data set for performing a probe calibration (Langø et al. 2000) contains about 20 or 250 images depending on the method of acquisition.

An obvious speed increase implementation would be to store the CWT in a table the first time it is calculated. When the optimal scale is found, the CWT for this scale is simply retrieved from the stored CWT table instead of recalculating it. After squaring this 1-D table, the maximum is obtained as position for the edge as before. This approach would be even more rewarding if using several beams to find the optimal scale. Another speed increase could be achieved by calculating the spectrum by using the Fast Fourier Transform (FFT). The spectrum might, however, be slightly different, but qualitatively, the same characteristics of the signal should be available from a spectrum calculated with the FFT. Once the optimal scale has been found from the FFT spectrum, the edge detection is performed as described by calculating the CWT for this scale. The use of the FFT is a subject of interest in a further study of this method.

The only input parameter to the procedure, apart from the mother wavelet function, is the scale array to use for calculating the CWT. This array may be viewed as a threshold parameter. Nevertheless, we believe that by including such a large scale range, most cases interesting to edge detection in ultrasound imaging is covered. A plot of the Mexican hat wavelet for a very small scale ($a=2^{10/4}$) and a very large scale ($a=2^{30/4}$) together with one of the beams from our study should illustrate this (Fig. 11).
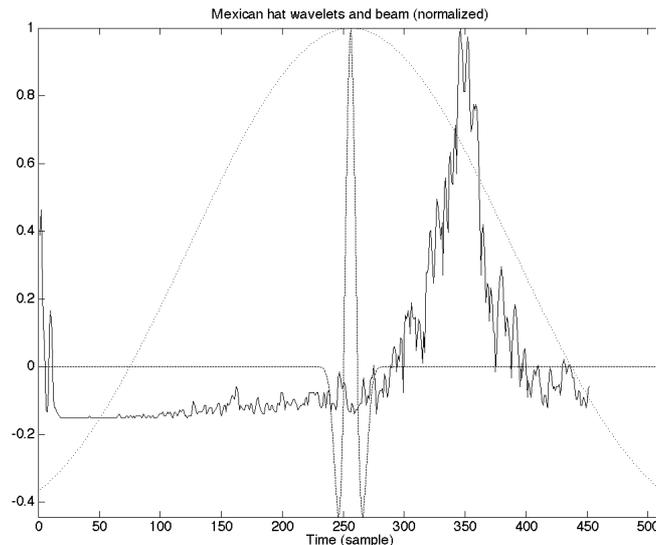

Fig. 11. Plots showing normalized Mexican hat wavelets for $a=2^{10/4}$ (dashed) and $a=2^{30/4}$ (dotted) and ultrasound beam (solid). $\tau=256$ in both cases. The ultrasound beam is the middle beam from the image in Fig. 8d.

A filter, such as a simple median type filter, could be applied to the images prior to edge detection in order to remove high frequency noise. This will of course slow down the algorithm, but improve the results for images with a low signal to noise ratio, especially important for in vivo ultrasound data.

If a segmented image, a probe calibration matrix, a volume estimate etc. is the goal of a procedure, then it would be sufficient to perform edge detection prior to scan conversion of the images and scan convert only the edge image. This could speed up the scan conversion process considerably.

We believe that further research should be invested in this method. Particularly, other mother wavelet functions (Daubechies 1992; Rao and Bopardikar 1998) could be tested and the results should be compared to edge detection on the same images using other existing edge detection methods, such as those mentioned in the introduction.

ACKNOWLEDGEMENTS

REFERENCES

Adam D, Hareuveni O, and Sideman S. Semiautomated border tracking of cine echocardiographic ventricular images. IEEE Trans Med Imag, 1987;MI-6:266-271.

Angelsen BAJ. Waves, signals and signal processing in medical ultrasonics. Trondheim, Norway, Dept. of Physiology and Biomedical Engineering, The Norwegian University of Science and Technology (NTNU), I, 1995.

Canny J. A computational approach to edge detection. IEEE Trans Pattern Anal Machine Intell, 1986;PAMI-8:679-698.

Chalana V, Linker DT, Haynor DR, and Kim Y. A multiple active contour model for cardiac boundary detection on echocardiographic sequences. IEEE Trans Med Imag, 1996;15:290-298.

Chou WS, Wu CM, Chen YC, and Hsieh KS. Detecting myocardial boundaries of left ventricle from a single 2DE image. Pattern Recogn, 1990;23:799-806.

Chu CH, Delp EJ, and Buda AJ. Detecting left ventricular endocardial and epicardial boundaries by digital two-dimensional echocardiography. IEEE Trans Med Imag, 1988;7:81-90.

Chui CK. An introduction to Wavelets. Boston, Academic Press, 1, 1992.

Czerwinski RN, Jones DL, and O'Brien Jr. WD. An approach to boundary detection in ultrasound imaging. IEEE Ultrasonics Symp, 1993:951-955.

Daubechies I. Ten lectures on Wavelets. Philadelphia, Pennsylvania, Society for Industrial and Applied Mathematics, 61, 1992.

Demigny D, Lorca FG, and Kessal L. Evaluation of edge detectors performances with a discrete expression of Canny's criteria. IEEE Int Conf Image Proc, 1995:169-172.

Dias JMB, and Leitao JMN. Wall position and thickness estimation from sequences of echocardiographic images. IEEE Trans Med Imag, 1996;15:25-38.

Farge M, Kevlahan N, Perrier P, and Goirand E. Wavelets and turbulence. Proc IEEE, 1996;84:639-669.
Feng J, Lin W-C, and Chen C-T. Epicardial boundary detection using fuzzy reasoning. IEEE Trans Med Imag, 1991;10:187-199.

Friedland N, and Adam D. Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated annealing. IEEE Trans Med Imag, 1989;8:344-353.

Gonzales RC, and Woods RE. Image segmentation. Digital Image Processing. Massachusetts, Addison-Wesley, 1993;413-481.

Gronningsaeter A, Kleven A, Årseth TE, Langø T, Lindseth F, Lie T, and Unsgård G. SonoWand, An Ultrasound based Neuronavigation System. Submitted to Neurosurgery, 2000.

Heath M, Sarkar S, Sanocki T, and Bowyer K. Comparison of edge detectors: A methodology and initial study. IEEE Computer Society Conf CompVision Pattern Recogn, 1996:143-148.

Hough PVC. Methods and means for recognizing complex patterns. USA. Patent 3,069,654. 1962.

Hudgins L. Wavelet analysis of Atmospheric Turbulence. Ph.D., University of California, Irvine, USA, 1992.

Hunter IA, Soraghan JJ, and McDonagh T. Fully automatic left ventricular boundary extraction in echocardiographic images. Proc Computers in Cardiology, Vienna, 1995:741-744.

Kaspersen JH. A study of coherent structures using Wavelet analysis. Ph.D., Norwegian University of Science and Technology (NTNU), Trondheim, Norway, 1996.

Klinger JW, Vaughan CL, Fraker TD, and Andrews LT. Segmentation of echocardiographic images using mathematical morphology. IEEE Trans Biomed Eng, 1988;35:925-934.

Langø T, Lindseth F, Kaspersen JH, Kleven A, and Gronningsater A. Novel probe calibration methods for 3D freehand ultrasound. Submitted to Comput Aided Surg, 2000.

Liandrat J, and Moret-Bailly F. The wavelet transform: some applications to fluid dynamics and turbulence. Eur J Mech, B/Fluids, 1990;9:1-19.

Linker DT, Pearlman AS, Lewellen TK, Huntsman LH, and Moritz WE. Automated endocardial definition of 2-D echocradiograms: A comparison of four standard edge detectors and improved thresholding techniques. Proc Computers in Cardiology, 1982:395-398.

Mailloux GE, Langlois F, Simard PY, and Bertrand M. Restoration of the velocity field of the heart from two-dimensional echocardiograms. IEEE Trans Med Imag, 1989;8:143-153.

Mallat S, and Hwang WL. Singularity detection and processing with wavelets. IEEE Trans Inform Theory, 1992;38:617-643.

Meunier J, Bertrand M, Mailloux GE, and Petitclerc R. Local myocardial deformation computed from speckle motion. Proc Comput Cardiology, Washington DC, 1988:133-136.

Ødegård L, Halvorsen E, Ystad B, Torp HG, and Angelsen B. Delay and amplitude focusing through the body wall; A simulation study. IEEE Ultrasonics Symp, 1996;2:1411-1414.

Prager RW, Rohling RN, Gee AH, and Berman L. Rapid calibration for 3-d freehand ultrasound. Ultrasound Med Biol, 1998;24:855-869.

Rabben SI, Torp AH, Støylen A, Slørdahl S, Bjørnstad K, Haugen BO, and Angelsen B. Semiautomatic contour detection in ultrasound M-mode images. Ultras Med Biol, 2000;26:287-296.

Rao RM, and Bopardikar AS. Wavelet Transforms. Introduction to theory and applications. Massachusetts, Addison-Wesley, 1998.

Rioul O, and Vetterli M. Wavelets and signal processing. IEEE Sign Proc Magazine, 1991;84:14-38.

Santis AD, and Sinisgalli C. A Bayesian approach to edge detection in noisy images. IEEE Trans Circuits Syst I: Fund Theory Appl, 1999;46:686-699.

Setarehdan SK, and Soraghan JJ. Cardiac left ventricular volume changes assessment by long axis echocardiographical image processing. IEE Proc Vis Image Signal Process, 1998;145:203-212.

Setarehdan SK, and Soraghan JJ. Automatic cardiac LV boundary detection and tracking using hybrid fuzzy temporal and fuzzy multiscale edge detection. IEEE Trans Biomed Eng, 1999;46:1364-1378.

Sun YN, and Tsai CT. Segmentation of echocardiograms using a neural network. Microprocessing Microprogram, 1992;35:791-798.

Xin Y, Boekhaerts P, Bister M, Taeymans Y, and Cornelis J. Left ventricular boundary detection in echocardiographical images using artificial neural networks. 6th European Signal Processing Conf, 1992:1199-1202.

Xu C, and Prince JL. Snakes, shapes, and gradient vector flow. IEEE Trans Image Proc, 1998;7:359-369.

Paper C

Navigation of Probes and Surgical Instruments in Ultrasound-Guided Surgery

Thomas Langø [1], Steinar Ommedal [1], Geirmund Unsgård [2] and Aage Gronningsaeter [3]

[1] SINTEF Unimed, Ultrasound, 7465 Trondheim, Norway
[2] Department of Neurosurgery, University Hospital and Medical Faculty, Norwegian University of Science and Technology, 7491 Trondheim, Norway
[3] MISON, 7489 Trondheim, Norway

ABSTRACT

We describe novel methods for calibrating and navigating surgical instruments during ultrasound-guided surgery. The methods allow us to provide the surgeon with information about the position and orientation of the ultrasound image plane and the surgical tool, relative to each other in space. This means that the surgeon can adjust either the orientation of the ultrasound probe or the surgical tool in order to obtain an optimal view of the tool at all times. The method requires a position sensor on both the tool and the probe; it also requires that both the probe and tool have been calibrated. The probe calibration procedure establishes the position and orientation of the ultrasound image relative to the position sensor attached to the probe. Similarly, the tool calibration calculates the tip location and the orientation of the tip portion of the tool in space relative to the origin of the sensor attached to the tool. It is a prerequisite that the tool has a straight and rigid tip portion. Initial trials with the navigator module in conjunction with a prototype neuronavigation system have been performed during removal of brain tumors.

INTRODUCTION

Image-guided surgery is commonly conducted by the use of preoperative images, such as magnetic resonance images (MRI) or computerized tomography (CT) data. These preoperative images can be very accurate provided no significant changes occur during surgery. However, during surgery many different factors may affect tissue movement, and hence cause changes that are not reflected in the images acquired prior to the surgical procedure. In neurosurgery these changes are often referred to as the brain shift problem.[1, 8-10, 13] This movement or shifting of anatomical structures as the procedure progresses is due to the removal of tumor tissue, to drainage of cerebrospinal fluid, and to gravity as the patient might be positioned differently than was the case during acquisition of the preoperative images. In order to continuously work with updated images during the surgical procedure, intraoperative imaging modalities have been introduced. These include real time 2D ultrasound imaging, open MRI, and interventional MRI. Some operating theaters use a combination of preoperative MRI and intraoperative ultrasound where the ultrasound images are used to identify and quantify the brain shift.[3] In recent years another possibility has been presented, namely real time 2D ultrasound in combination with repetitive 3D ultrasound acquisitions using a position sensing system to track the position and orientation of the images. The preoperative MRI or CT images are used mainly for preoperative planning and to get an overall view of the anatomical area of interest. A major advantage of ultrasound-guided navigation in surgery is the ability to do repetitive 3D imaging during the surgical procedure, and thus continuously work with a data set that has recently been updated according to possible changes in the brain anatomy.[5]

With the SonoWand® system (MISON Inc., Trondheim, Norway)[5] the surgical procedure is guided with regularly acquired 3D volumes from a separate burr hole in the skull, and used in addition to preoperative image data such as MRI and CT. Real time 2D and 3D ultrasound images are used to monitor the shrinkage of the tumor or hematoma. A major limitation with this technique is that it is difficult to obtain a longitudinal cross section view of the tip portion of the surgical instrument at all times in the real time 2D image. Such a view is of crucial importance to know where the tip of the instrument is located relative to important anatomical structures.

This problem is also relevant for ultrasound image-guided biopsies in general. During biopsies guided by real time imaging it is important know if the distal part (including the tip) of the biopsy needle or forceps is located in the 2D image. If it is not, it would be valuable to know how to adjust the instrument or the imaging probe in order to obtain a visualization of the tip of the tool in the region of interest.

We have developed a navigator module that may run as a stand-alone application or in conjunction with the neuronavigation system SonoWand®. This module gives the surgeon feedback on the position and orientation of the surgical tool relative to the ultrasound data. The calculations are based on position measurements from the sensors attached to the probe and tool. In addition, we need the calibration transformations for both probe and tool. These calibrations establish the transformations between the sensor on the probe and the image plane and between the sensor on the tool and the tip portion of the tool. The tip portion of the tool is represented by a line in space, i.e. a 3D Cartesian coordinate system with its origin at the tip and aligned so that one of the main axis points along the tip direction. The scan plane coordinate system is set up with the origin at the top center of the image with the z-axis

pointing along the middle of the image in the radial direction. See Fig. 1 for an overview of the coordinate systems involved and the transformations between them.
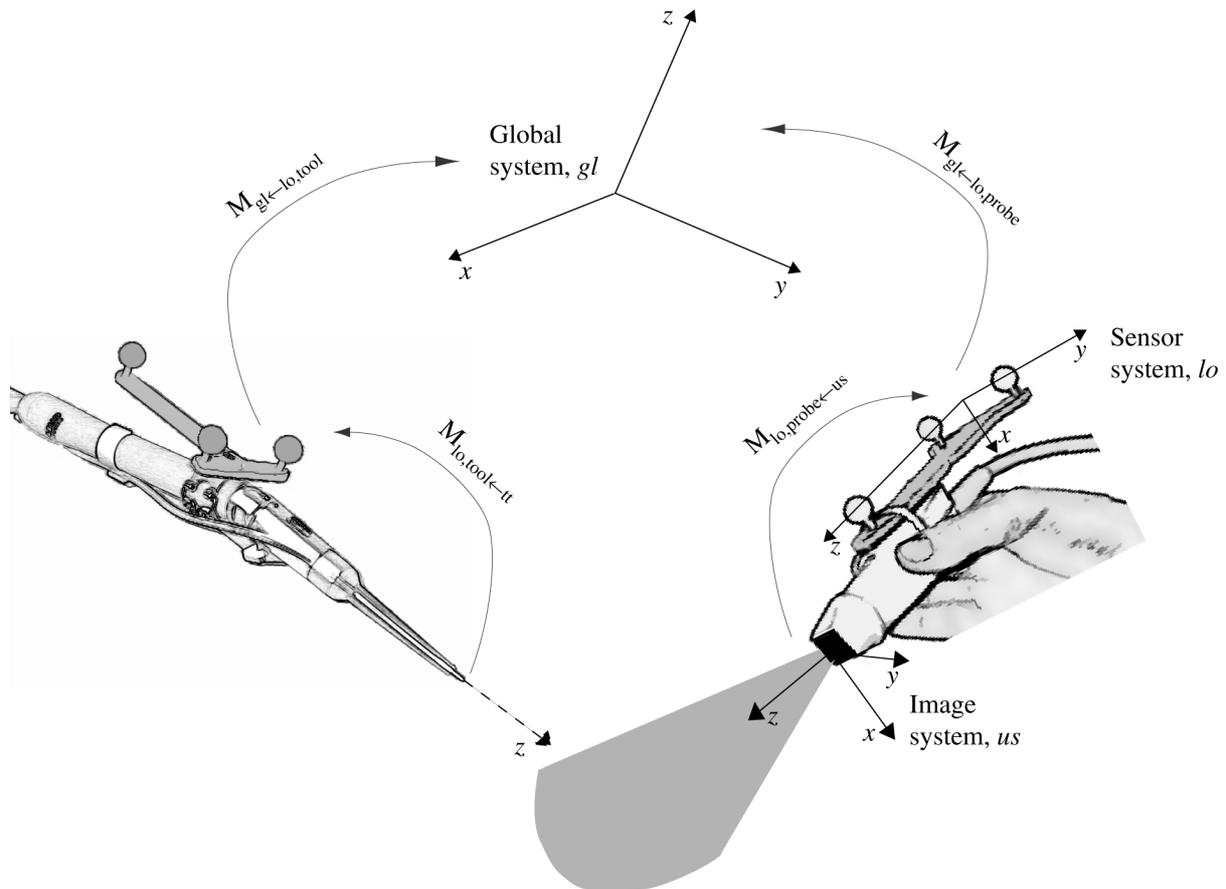

Fig. 1. Illustration of the coordinate systems and transformations. The probe calibration matrix is denoted $M_{lo,probe \leftarrow us}$, while $M_{lo,tool \leftarrow tt}$ is the tool calibration matrix. $M_{gl \leftarrow lo,probe}$ and $M_{gl \leftarrow lo,tool}$ are the transformations measured by the position sensing system. The figure also illustrates the attachment of the sensor to the probe with a glued on adapter and a strap to tighten the sensor to the adapter.

METHODS

In this study we used an optical position sensor system (VectorVision from BrainLAB Inc., Munich, Germany). The system consists of two cameras that emit infrared light and register the reflected light from small spheres arranged geometrically on a frame (see the frame attached to the probe in Fig. 1) in addition to a processing unit. Although the reflecting spheres represent a passive position and orientation measurement device, for simplicity we refer to these as sensors throughout the paper. The ultrasound probe was a 4-8 MHz phased array probe connected to a high performance digital scanner (System-fiVe from GE Vingmed Ultrasound Inc., Horten, Norway). The tools used were two kinds of biopsy forceps and a CUSA ultrasound aspirator (System 2000 from Valleylab, Boulder, CO, USA). A more detailed description of the methods and equipment described in this paper can also be found in a filed patent application.[12]

Ideally, one unique position sensor should be permanently attached to one unique ultrasound probe before calibration. However, this option represents a practical problem since most ultrasound probes can not be sterilized. Both devices could be covered with a sterile drape

during the surgical procedure, but this might hinder free sight between the cameras and the sensor. We have solved this dilemma by using a special adapter glued on to the probe casing that ensures repetitive and precise attachment between the probe and a specific sensor, even through a sterile drape. The specific combination of probe and sensor need only be calibrated once in the laboratory. From our experience, we believe that the variability in the calibration results due to this solution is negligible, or at least less than the fluctuations in the position readings.

Calibration of ultrasound probe

The calibration procedure determines the transformation matrix describing the position and orientation of the ultrasound image plane origin relative to the origin of the position sensor attached to the probe. The transformation itself is a 3 by 3 rotation matrix and a displacement vector combined in a 4 by 4 matrix in homogeneous coordinates[4] so that any transformation can be handled by simple matrix multiplication. An accurate calculation of this matrix is of crucial importance for a reconstruction that preserves true anatomical shape and size in freehand ultrasound 3D scanning and also for navigation in relation to the 2D real time image. The calibration parameters can be crudely estimated by external measurements of the probe case and sensor. However, this estimate will not necessarily be exact since the origin of the ultrasound image system is located inside the housing of the probe and may vary slightly from one probe to another. Furthermore, the position of the center transducer element (origin) and the direction for pulse transmission may vary due to variations in the production of the probe. In addition, for position sensing systems based on magnetic fields, the origin for the coils in the sensors is not well defined. This indicates that a better approach to finding the calibration matrix would be to image a phantom with known physical properties and dimensions. For the calibrations used in the tool navigator we have used a membrane alignment method.[11]

In short, the main idea of the method is to align the ultrasound image plane with a submerged, thin (but stiff), and jagged membrane in a water tank. When the image is such that the membrane and the scan plane can be assumed to be coplanar, the positions of the sensor on the probe and the positions of the membrane are sampled. The positions of the membrane corners are sampled using a pre-calibrated pointer for the Vector Vision system. In addition, the corresponding points in the image are marked manually using a mouse, i.e. the points' coordinates in the ultrasound image system. The pseudo-inverse matrix[7] for transformation between the two systems can then be calculated. For a more detailed description of this method, see Langø et al. 2000.[11]

A newer method for probe calibration[11] proved to be more accurate than the method used for the initial calibrations here. The new method is based on a free hand scan of a thin membrane oriented approximately horizontal in a water bath. The calculations after this image acquisition are automatic and hence the method is simpler to use than the other 2D based method. This new method for probe calibration will be used in future development and testing of the tool navigation method described in this paper.

A point $\left[x_{us}, y_{us}, z_{us} \right]^{T}$ in the ultrasound image, represented in global coordinates, is given by

$$\begin{bmatrix} x_{gl,us} \\ y_{gl,us} \\ z_{gl,us} \\ 1 \end{bmatrix} = M_{gl \leftarrow lo,probe} \cdot M_{lo,probe \leftarrow us} \cdot \begin{bmatrix} x_{us} \\ y_{us} \\ z_{us} \\ 1 \end{bmatrix} \quad (1)$$

where

$$M_{lo,probe \leftarrow us} = \begin{bmatrix} x_1 & y_1 & z_1 & t_x \\ x_2 & y_2 & z_2 & t_y \\ x_3 & y_3 & z_3 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2)$$

denotes the probe calibration matrix and $M_{gl \leftarrow lo,probe}$ is the position measurement from the sensor attached to the probe. The symbols t_x, t_y and t_z denotes the displacement vectors or translation offset, while x_i, y_i and z_i denotes the rotation matrix components. The calibration matrix is applicable in general since the probe and sensor constitute a rigid body. The manner of writing the transformation matrix shown in Eq. 2, i.e. homogenous coordinates[4], is done in order to be able to handle both rotation and translation in a single matrix multiplication.

Calibration of a surgical tool
Next, we need a calibration for the surgical tool in order to know the position and orientation of its tip and distal part respectively. Under sterile conditions, the attachment of the adapter to the tool may change from operation to operation, depending on specific situation. Therefore, it should be possible to perform the tool calibration quickly. We have developed a special sterile calibration device for this purpose (Fig. 2). In order to calibrate the instrument with this method, the tip portion of the instrument needs to be straight and rigid. The instrument as a whole may, however, be bent, as the CUSA in Fig. 1. Furthermore, the sensor can be attached anywhere on the instrument, not necessarily on the straight distal part.

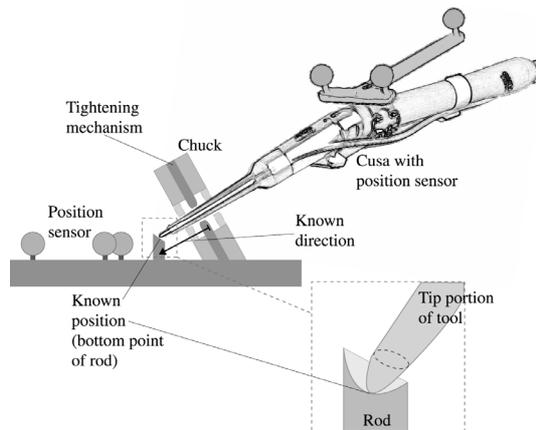Fig. 2. Calibration board for tool calibration. A chuck mechanism ensures that the instrument is aligned with the direction between the center of the chuck and tip position, parameters known from exact measurements in the making of the calibration board. The positions of the sensor on the instrument and the reference position sensor on the board are sampled and the calibration transformation is calculated.

As illustrated in Fig. 2, the straight tip portion is inserted through a hole in a chuck mechanism until it touches the rod (see zoom box). Both the "bottom" point of the rod and the direction through the center of the chuck to the rod are known relative to the reference position sensor frame attached to the device. The objective of the chuck is to hold tools of different thickness aligned along a known direction. Letting the end of the tool rest on the rod provides a known origin or end point for the tool tip. The distance between the rod and the center of the chuck for our device is approximately 4 cm, meaning that we are able to calibrate the CUSA and other regularly used tools such as different sized biopsy forceps. By sampling the position of the sensor that is attached to the instrument and the position of the reference sensor that is attached to the calibration device, we can calibrate the instrument. The result of the instrument calibration is the transformation matrix $M_{lo,tool \leftarrow tt}$, where tt represents the tip of the tool and lo represents the sensor coordinate system. The matrix is written in the same manner as Eq. 2. The tip of the tool in the global system is then given by

$$
\begin{bmatrix} x_{gl,t} \\ y_{gl,t} \\ z_{gl,t} \\ 1 \end{bmatrix} = M_{gl \leftarrow lo,tool} \cdot M_{lo,tool \leftarrow tt} \cdot \begin{bmatrix} x_{tt} \\ y_{tt} \\ z_{tt} \\ 1 \end{bmatrix} \tag{3}
$$

The matrix $M_{gl \leftarrow lo,tool}$ is the measured position of the sensor which is attached to the tool. The error in the precision-made calibration device itself is assumed to be less than ±0.1 mm.

Position and orientation measurements
The position and orientation of both the ultrasound image plane and the distal part of the tool are now described in global coordinates. By multiplying the calibration matrices with the continuously updated position readings for the sensors attached to the probe and the tool, we achieve matrices describing these parameters in global coordinates:

$$
\begin{aligned}
M_{gl \leftarrow us} &= M_{gl \leftarrow lo,probe} \cdot M_{lo,probe \leftarrow us} \\
M_{gl \leftarrow tt} &= M_{gl \leftarrow lo,tool} \cdot M_{lo,tool \leftarrow tt}
\end{aligned} \tag{4}
$$

This knowledge can now be exploited to calculate essential information to be presented to the surgeon. Parameters of interest are the shortest distance between the ultrasound image plane and the tip of the tool, and the angle between the distal part of the tool and the image plane. By continuously updating these parameters and presenting them in a simple manner to the surgeon, he or she is able to easily adjust either the tool, the probe, or both in order to obtain a longitudinal cross section view of the instrument's distal part in the real time 2D ultrasound image.

Distance between tool tip and ultrasound plane
The shortest distance between the tool tip and the image plane is found from the perpendicular projection onto the plane from the tip. Let A be a point at a distance d from the plane, which is spanned out by the three points P, Q, and R (Fig. 3).
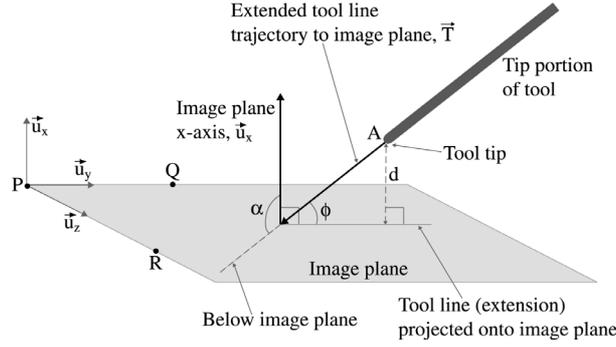
Fig. 3. Illustration of the distance d (Eq. 5) between the tool tip and ultrasound image plane and the angle ϕ (Eq. 6-9) between the distal part of the tool and the image plane. The other parameters are explained in the text.

The perpendicular distance d can be calculated as[2]

$$d = \frac{\left|\overrightarrow{AP} \cdot \overrightarrow{AQ} \times \overrightarrow{AR}\right|}{\left|\overrightarrow{PQ} \times \overrightarrow{PR}\right|} \tag{5a}$$

The point A defines the tool tip in the global coordinate system, P is the image plane origin, Q is a point along the y-axis of the image plane, and R is a point along the z-axis of the image plane. All these data; A, P, Q, and R are readily available from the position readings of the position sensing system and the calibration matrices (Eq. 4).

In Fig. 3 $\vec{u}_x$ is a unit vector along the x-axis of the image plane (perpendicular to the plane) and is given by

$$\vec{u}_x = \frac{\overrightarrow{PR} \times \overrightarrow{PQ}}{\left|\overrightarrow{PR} \times \overrightarrow{PQ}\right|} \tag{5b}$$

Angle between tool and ultrasound plane
The angle of interest is the angle ϕ between the distal part of the tool and the line represented by the perpendicular projection of the extended tool line $\vec{T}$ onto the image plane as shown in Fig. 3. This angle, which is between 0° and 90°, can be found from the angle α in Fig. 3 given by the dot product

$$\vec{T} \cdot \vec{u}_x = \left|\vec{T}\right| \cdot \left|\vec{u}_x\right| \cdot \cos\left\{\angle\left(\vec{T},\vec{u}_x\right)\right\} = \left|\vec{T}\right| \cdot \cos\{\alpha\} \tag{6}$$

There are three possible situations: I) the tool is pointing completely or partially along the negative x-axis ($\alpha > 90°$, which is the case in Fig. 3):

$$\phi = \alpha - \frac{\pi}{2} = \arccos\left\{\frac{\vec{T} \cdot \vec{u}_x}{\left|\vec{T}\right|}\right\} - \frac{\pi}{2} \tag{7}$$

or II) the tool is pointing completely or partially along the positive x-axis ($\alpha < 90°$):

$$\phi = \frac{\pi}{2} - \alpha = \frac{\pi}{2} - \arccos\left\{\frac{\vec{T} \cdot \vec{u}_x}{|\vec{T}|}\right\} \tag{8}$$

or III) the tool is pointing perpendicularly at the *x*-axis ($\alpha = 90°$):

$$\phi = 0 \tag{9}$$

i.e., the tool is parallel with the image plane. This means that it is possible to tell whether the tool is pointing towards the sensor side of the probe or vice versa. We are now able to differentiate among five different spatial relationships between the image and tool as illustrated in Fig. 4. The ultrasound image plane is represented by the dotted line, i.e. a side view along the negative y-axis of the image plane (see Fig. 1). The arrows represent the distal part of the tool.

Ultrasound image plane

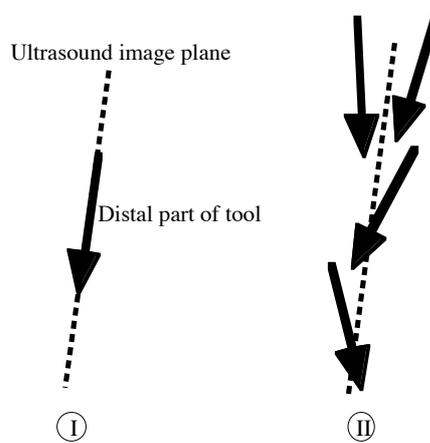Distal part of tool

Ⅰ   Ⅱ

Fig. 4. (I) The optimal relative orientation of the surgical tool and the ultrasound scan plane. (II) Four different orientations of the instrument with respect to the image plane. In both figures the arrows represent the distal part of the surgical tool and the dotted line represents the ultrasound image plane. The image plane is normal to the paper in this figure.

Method I: 2D graphic color bar
In order to present this information to the surgeon we suggest a graphic bar display located either on the ultrasound scanner monitor or on a separate navigation monitor, on which the preoperative images and the ultrasound 3D data may also be presented. The graphic display is continuously updated from the calculations described above. Fig. 5 shows different examples of the graphic bar display with descriptions of the situations below each case in the manner illustrated in Fig. 4. The vertical dark part in the middle is the distance indicator. This indicator has a constant color and grows in width from the middle to one side according to the distance measure *d* (Eq. 5) and the sign of the dot product in Eq. 6. The panels are shown here in gray scale but color is recommended. The distance indicator was dark blue in our initial tests.
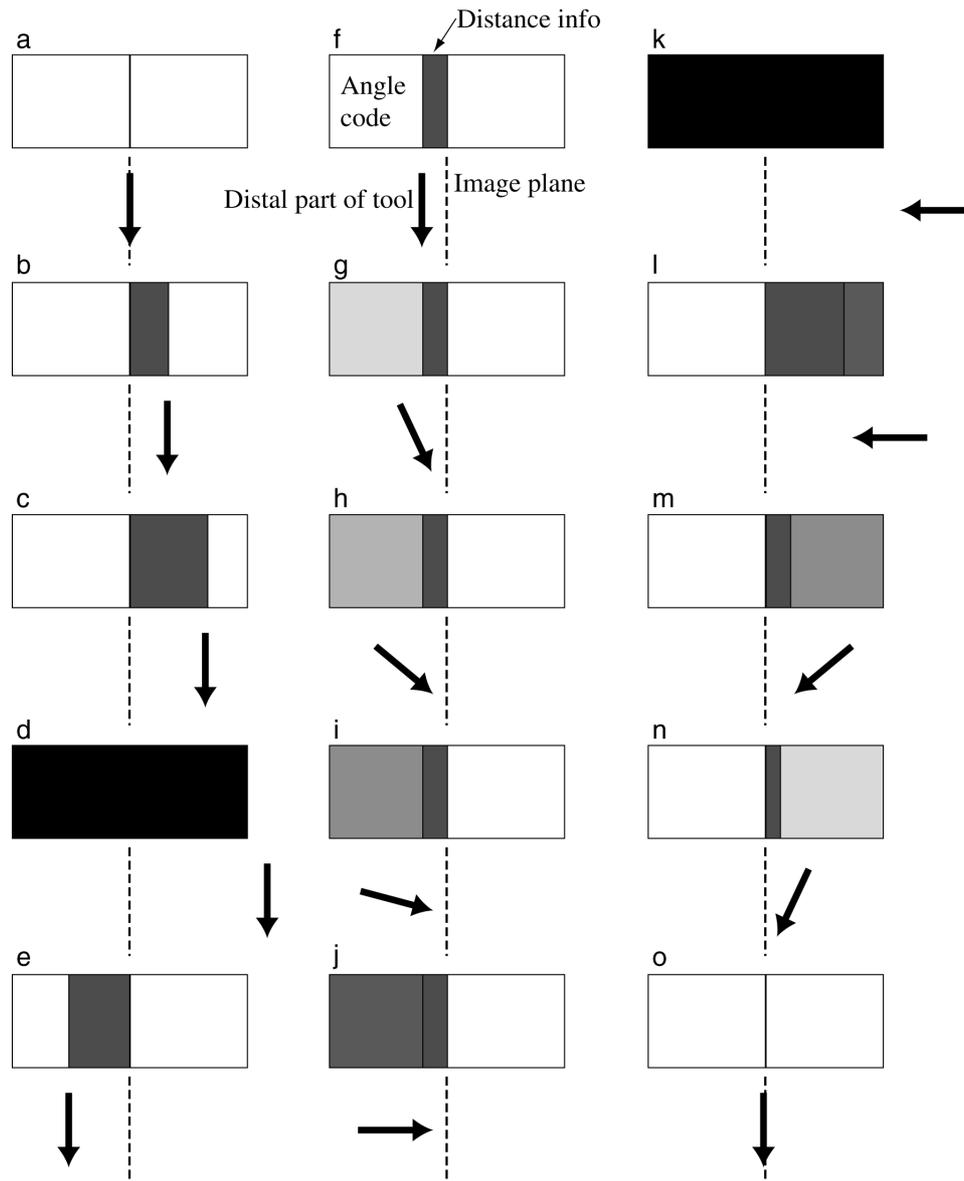
Fig. 5. Examples of the graphic bar display versus tool alignment (see Fig. 4). The darker bar in the middle indicates the distance parameter d (Eq. 5). One of the areas to the side is used to code information about the angle ϕ (Eq. 7-9). A constant color is used for the distance, e.g. blue, while a color range is used for the angle information, e.g. from light yellow (almost perfect alignment between tool and image) to dark red (tool perpendicular to image plane). A white rectangle with a thin black line in the middle corresponds to perfect alignment, i.e. the distal part of the tool is in the image plane (case a and o). See text for description of the cases a through o.

One of the areas to the side of the distance indicator is used to indicate the angle (Eq. 7-9). Which side is used depends on the sign of the dot product in Eq. 5, which indicates which side of the probe-sensor system the distal part of the tool is pointing towards. The angle information is given by the color of the area. In our initial tests we have used a color scale ranging from light yellow to dark red for this information. When the tool and image plane are optimally aligned the area is white; as the angle increases, the color changes through yellow, orange, and into dark red as the angle approaches 90°. Different color scales may be used and selected prior to running the navigator module.

In Fig. 5a-e the distal part of the tool is parallel with the image plane, but at different distances. In Fig. 5d the tool is moved beyond the preset maximum distance and the entire bar is black. In Fig. 5e the bar grows towards the opposite side compared to Fig. 5a-d since the tool is located on the other side of the probe-sensor system. In Fig. 5f-i the tool tip is located a constant distance from the image plane but tilted from a parallel orientation relative to the image plane in Fig. 5f, to a perpendicular orientation in Fig. 5i. In Fig. 5k-o the tool is moved from being beyond the maximum distance (Fig. 5k) to an optimal orientation in Fig. 5o via some intermediate orientations (Fig. 5l-n).

We have implemented the option of setting a maximum distance and angle to be indicated by the navigator module. The operator can at any time change these values. If the distance is greater than this maximum, the navigator indicates that the tool and image are too far away from each other to be of interest for navigation (e.g. the entire bar is black). The entire set-up may also be flipped if the surgeon should happen to switch position from one side of the patient to the other so that left and right in the display keeps its meaning.

The navigator can be exploited in three different ways. The first and maybe the most obvious to the clinician, is to attach the ultrasound probe on the patient in a position that provides a good view of the anatomical structure(s) of interest and then use the navigator as feedback in order to adjust the instrument during insertion into the body. The second approach is to insert the tool in a certain position due to for instance limitations in the anatomy, and then adjust the ultrasound probe according to the feedback from the navigator until an optimal longitudinal view of the instrument is achieved. The last method is to allow both the instrument and the ultrasound probe to be adjusted during insertion of the instrument into the body, and also during the procedure. This could be advantageous if the tool needs to be moved a lot.

The graphic bar should be small enough not to obstruct any important visualizations on the screen, but at the same time not so small that the operator can not see the necessary information with a quick glance at the screen or in his side vision.

FURTHER DEVELOPMENTS

In this section we suggest further developments and possible implementations of the navigator information. These implementations have not been tested yet.

Method II: Acoustic feedback
The graphic display bar may be presented to the surgeon in combination with an audible feedback. This might save time in that the surgeon does not have to spend time looking at the bar. The acoustic feedback needs to be of such a kind that it does not irritate or annoy the staff present in the operating room. When the ultrasound probe and the surgical tool are too far away from each other to be interesting for navigation purposes there should be no sound. In the tabular overview below, a suggestion for audio feedback is described. The distance parameter d controls the repetition rate for the *on/off* of the sound signal, while the pitch (or frequency) during *on* of the sound signal describes the angle parameter ϕ.

Sound $d \leftrightarrow$ repetition rate on/off $\phi \leftrightarrow$ pitch	Description
Silence	Tool and probe are so far away from each other that the navigation feedback is unnecessary. For example, this is the case when tool is removed from the surgical area for whatever reason.
Highest repetition rate, highest pitch	Worst case. The tip portion is at an approximately 90° angle to the image plane and the tip is located at the maximum distance (provided by the user) from the image plane.
Medium repetition rate, medium pitch	Tool tip is far away and tip portion of tool is at a significant angle relative to image plane in space.
High repetition rate, low pitch	Tip portion of tool is close to parallel with the image plane, but the tool tip is at a significant distance from the image plane in space.
Low repetition rate, high pitch	Tool tip is close to being in the image plane, but the angle between the tip portion of the tool and the image plane is significant.
Low repetition rate, low pitch	Situation is close to optimal, i.e. the tip portion of the tool line is close to being in the image plane. For the optimal situation there should be a distinguishable lowest pitch and frequency.
Lowest repetition rate, lowest pitch	Perfect alignment. Sound should be easily distinguishable from the previous one.

When the situation reaches the worst case, maximum limits, the system could present a continuous distinct signal for 2 seconds and then be silent until the tool and image are within the limits again. The difference between minimum and maximum repetition rate and between minimum and maximum pitch should not be too large. The volume and the tones should be adjustable, just like the color scale for Method I. Furthermore, it should be possible to set the maximum values so small for the audible feedback that no sound is made unless the tool tip and image plane are very close to the defined optimal situation. The surgeon will then be able to concentrate on the images and the procedure and use the audio feedback for fine tuning of the probe and/or tool position.

Method III: 'Off-line' 3D navigation
As described in the introduction, repetitive 3D acquisitions using a conventional ultrasound probe with an attached position sensor allows the surgeon to operate according to recently updated images from the anatomical area of interest. These regularly updated 3D volumes may be visualized stereoscopically[6], i.e. providing depth perception by simultaneously displaying two images from slightly different angles, one for each eye. This view provides the surgeon with easily interpretable images and orientation is simpler than with conventional 2D images. Nevertheless, the real time 2D ultrasound image provides important detailed information and an up to date view of the anatomical area of interest. Hence, a combination of these two views, the stereoscopic view of the recently acquired 3D volume and the real time 2D image, could prove useful to the surgeon.

Utilizing the information provided by the tool navigator, it is possible to provide the surgeon with data about the relation between the real time 2D ultrasound image and the distal part of the tool in the 3D view. A graphic representation of the tool is presented in the 3D view with superimposed graphic information (angle and distance indicators in Fig. 6) about the distance and angle between the real time image plane and the distal part of the tool. This is schematically illustrated in Fig. 6. The distance and angle parameters are calculated as before. It is also possible to let the tip of the tool control the front plane of the stereoscopic view, i.e. where the volume is cut. This makes the stereoscopic front plane view and the real time 2D

view correspond when a cross sectional view of the tool is obtained in the real time 2D image. In this situation, the stereoscopic view provides depth information and the real time 2D image shows the detailed and updated scene.
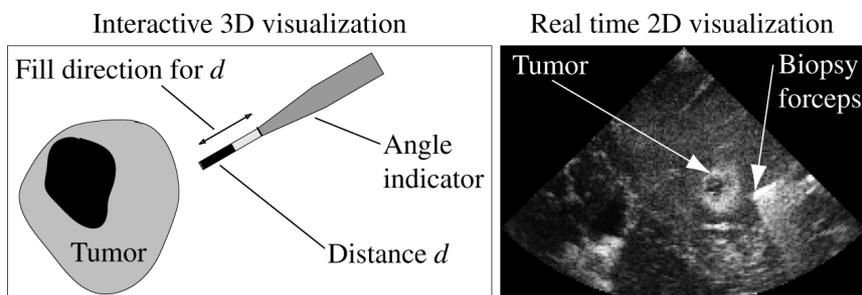

Fig. 6. Illustration of the off-line 3D tool navigation feature with a biopsy forceps and a brain tumor. A graphic representation of the distal part of the tool is shown in the 3D view, which could be stereoscopic (the part marked 'tumor' to the left). Distance and angle information is coded into the tool representation. The black part of the tip shrinks (the white increases) as the tip gets closer to the real time 2D image shown to the right. The other part of the tool representation indicates the angle with a gray scale or color range.

The 3D volume might be displayed in manners other than stereo, e.g. as three orthogonal slices with the tool position and orientation information coded into a line representation. The color of the tip of the representation of the tool might indicate the distance, while the rest of the line describes the angle between the 2D real time image and the surgical instrument.

INITIAL RESULTS AND DISCUSSION

The navigator module has proven useful in initial tests during removal of brain tumors guided by ultrasound from a separate burr hole in the skull. The surgeon first calibrated the CUSA and then inserted the instrument into the tumor. Next, the probe was adjusted according to the tool navigator graphic bar display in order to obtain a longitudinal cross sectional view in the real time 2D image of the distal part of the CUSA. The navigator has not been implemented into the navigation system SonoWand® system.[5] Hence, the tool navigator was run as a stand alone application. We believe this module will be of greater value when used in combination with navigation software.

We expect that the method will simplify and improve the procedure of taking biopsies by providing useful documentation and information about the location and orientation of the biopsy forceps in relation to the real time 2D image on the scanner and the 3D volume used for guidance. By storing the information from the tool navigator along with the navigation images during a biopsy, this information may be reviewed when an analysis of the biopsy sample has been performed to decide further recessing etc. In Fig. 7 we have shown an overview of the situation from a biopsy during removal of a brain tumor. The figure also shows the separate burr hole for the ultrasound probe and the 2D real time image has been projected onto the illustration. The 3D ultrasound volumes are also acquired from this burr hole by sweeping the probe and tracking its position with the attached sensor.
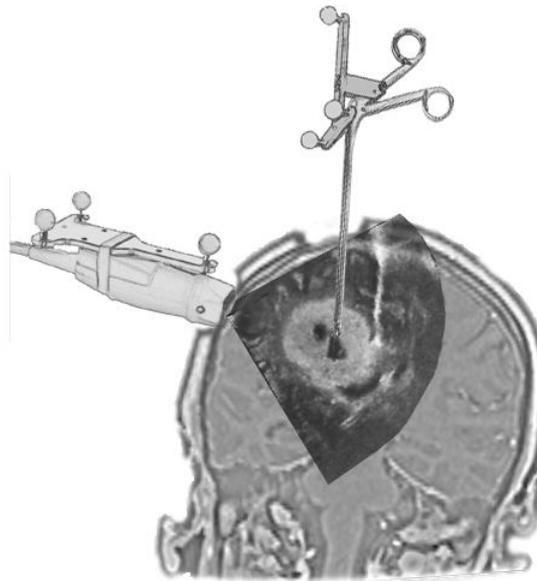
Fig. 7. Illustration of the situation in tool navigation during removal of a tumor in the brain. The figure shows a biopsy forceps, the ultrasound probe, and the real time 2D image plane superimposed on a MR image of the head.

The tool calibration accuracy for the CUSA was checked just after calibration in the operating room by simply pointing the tool at some steady sterile point and moving it about (with the tip steady) while sampling a given number of positions for the attached sensor frame. The movement approximately covered more or less the upper half of a hemisphere, the tip of the tool being the center. The spread of the tip location, calculated with the calibration matrix, then gave an indication of the accuracy. The exact number varied somewhat but was strongly dependent on the extent of the movement while keeping the tip steady, as well as the steadiness of the tip. Nevertheless, the movement during this accuracy test was more extensive than what was the case during insertion of the CUSA into the anatomical area of interest. This was done to obtain the worst case values for the CUSA tip definition. With approximately 100 samples, we achieved a mean value for the radial standard deviation of the tip point definition of approximately 1.0 mm. The minimum and maximum values for this deviation varied, but these numbers strongly depend on whether the user keeps the tip in the defined position while moving the rest about.

The tool navigation feature may lose its value with the introduction of real time 3D scanning. Today a modern high-end ultrasound scanner is capable of making approximately 20 high quality, wide sector images per second. By reducing the sector width and the spatial resolution, it is probably possible to achieve as many as several hundred scans per second (depending on scan depth). This means that a limited 3D sector may be scanned with several 3D volumes per second. The technique will require an advanced scanner with electronic beam steering in two directions, or a motorized probe with mechanical scan plane steering. Both techniques are expected to be available within a few years. Real time 3D imaging will make it possible to see the moving surgical instrument directly in the scene in relation to the surrounding structures. Nevertheless, for 2D scanners, which we believe will be used for several years to come, we expect our navigator module to represent a valuable and less expensive alternative to real time 3D ultrasound. Furthermore, with high demands for resolution and scan depth, one will probably have to cope with repetitive 3D acquisitions as

described above. In addition, should the reverberations from the tool degrade the image too much, the off-line 3D navigation method described above will be useful.

In conclusion, we have developed a method for calculating the distance and angle between the real time 2D ultrasound image plane and the distal part of a surgical tool. This information is coded into a graphic color display bar, which is updated in real time from position sensor measurements and calibration transformations for the probe and tool. A method and equipment for determining the tool calibration under sterile conditions was presented and proved accurate in initial tests. The tool navigator has been tested during removal of brain tumors and proved to be useful. We believe that further developments, including the suggestions put forward in this paper, are of interest and that this module could prove valuable, for instance for performing biopsies. Finally, we believe that this module should be incorporated into an ultrasound-based navigation system as an extra feature providing useful additional real time information to the images themselves.

REFERENCES

1. Bucholz RD, Yeh DD, Trobaugh J, McDurmont LL, Sturm CD, Baumann C, Henderson JM, Levy A, and Kessman P. Correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device. Lecture Notes in Computer Science Proceedings of the 1997 1st Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery, CVRMed-MRCAS'97, Grenoble, 1997;1205:459-466.

2. Edwards CH, and Penny DE. Calculus and analytic geometry. Englewood Cliffs, N.J., USA, Prentice-Hall, 1990.

3. Erbe H, Kriete A, Jödicke A, Deinsburger W, and Böker DK. 3D-ultrasonography and image matching for detection of brain shift during intracranial surgery. Computer Assisted Radiology, CAR '96, Paris, France, Elsevier Science, 1996:225-230.

4. Foley JD, vanDam A, Feiner SK, and Hughes JF. Computer graphics: Principles and practice. Reading, Massachusetts, Addison-Wesley, 1990.

5. Gronningsaeter A, Kleven A, Årseth TE, Langø T, Lindseth F, Lie T, and Unsgård G. SonoWand, An ultrasound-based neuronavigation system. Accepted for publication in Neurosurgery, 2000.

6. Gronningsaeter A, Lie T, Kleven A, Mørland T, Langø T, Unsgård G, Myhre HO, and Mårvik R. Initial experience with stereoscopic visualization of 3D ultrasound data in surgery. Surgical Endosopy, 1999;In press.

7. Hagan MT, Demuth HB, and Beale M. Neural network design. Boston, Massachusetts, USA, PWS Publishing Co., 1996.

8. Hata N, Dohi T, Iseki H, and Takakura K. Development of a frameless and armless stereotactic neuronavigation system with ultrasonographic registration. Neurosurgery, 1997;41:608-613.

9. Hirschberg H, and Unsgaard G. Incorporation of ultrasonic imaging in an optically coupled frameless stereotactic system. Acta Neurochir Suppl Wien, 1997;68:75-80.

10. Koivukangas J, Louhisalmi Y, Alakuijala J, and Oikarinen J. Ultrasound-controlled neuronavigator-guided brain surgery. J Neurosurg, 1993;79:36-42.

11. Langø T, Lindseth F, Kaspersen JH, Kleven A, and Gronningsaeter A. Novel probe calibration methods for 3D freehand ultrasound. Submitted to Computer Aided Surg, 2000.

12. Langø T, Lindseth F, Ommedal S, Kleven A, Unsgård G, and Gronningsaeter A. Tool navigator. UK. Patent Application 31,15,69544. 1999.

13. Trobaugh JW, Richard WD, Smith KR, and Bucholz RD. Frameless stereotactic ultrasonography: Method and applications. Comput Med Imaging Graph, 1994;18:235-246.

Paper D

Accuracy Assessment of a Prototype Ultrasound-Based Neuronavigation System

Aage Gronningsaeter[1], Frank Lindseth[2], Thomas Langø[2], Geirmund Unsgård[3]

[1] MISON AS, 7489 Trondheim, Norway.
[2] SINTEF Unimed, Ultrasound, 7465 Trondheim, Norway.
[3] Department of Neurosurgery, University Hospital and Medical Faculty, Norwegian University of Science and Technology (NTNU), 7006 Trondheim, Norway.

ABSTRACT

OBJECTIVES: We are developing an ultrasound-based neuronavigation system. The SonoWand® system is described in a companion paper, and the objective of this paper is to present the initial results from our assessment of the *overall accuracy* and resolution of a prototype version. In addition, we provide a detailed list of error sources associated with ultrasound-based navigation and conventional MRI-based navigation.

METHODS: A water-filled phantom with 27 precise and well-defined points was scanned by ultrasound, and 10 operators independently pinpointed these points in the 3D image data set using a mouse. In addition, the same operators pinpointed the 27 physical points in the phantom using a pointer device. The *overall accuracy* of the system was estimated through a statistical analysis of these 270+270 measurements.

RESULTS: 68 % of the physical phantom measurements deviated less than 2.77 mm from the corresponding points as pinpointed in the 3D image volume (95 % less than 3.76 mm). Included in this result is a systematic bias error of 2.07 mm. Simulations show that the error would drop from 2.77 mm to 1.7 mm if this systematic error could be eliminated.

CONCLUSIONS: The *overall accuracy* (laboratory test) of the prototype is measured to 2.7 mm, 1.7 mm is expected in the product version. Ultrasound-based navigation does not require patient registration as MRI-based navigation does. Thus, ultrasound is less susceptible to user- and procedure-dependent errors. These facts indicate that our prototype has an *overall clinical accuracy* that is similar to conventional navigation systems.

INTRODUCTION

We are developing an ultrasound-based neuronavigation system that differs from conventional neuronavigation systems in that it consists of a high-performance ultrasound scanner with built-in neuronavigation. This single-rack concept saves space in the operating room compared to two-rack solutions. The system can function as a conventional ultrasound scanner, as a conventional neuronavigation system based on MRI or CT-images, or more important, as a combined system where the features of preoperative MRI and intraoperative 3D ultrasound are utilized. There is a close integration between the ultrasound scanner and the navigation system that enables convenient use of high quality 3D ultrasound. A more detailed description of the system is provided in a companion paper (3). To our knowledge, this is the first system that enables the surgeon to navigate directly by means of *intraoperative 3D* ultrasound. The more common strategy is to display a *real-time 2D* ultrasound image and perform a direct comparison with the corresponding cross section through a preoperative MRI or CT-volume (1, 2, 4, 6, 9, 16).

A preliminary prototype system has been subject to clinical trials since November 1997. The initial experience with the system clearly indicates substantial clinical value of high quality 3D ultrasound in neuronavigation (publication in progress). However, the surgeon's confidence in the system in terms of *overall clinical accuracy* (geometric) and anatomical/pathological representation (spatial/contrast resolution and tissue differentiation) determines the delicacy, precision and extent of the work he/she can do. The *overall clinical accuracy* of the system is a measure of how accurately a pointer can be navigated in the field, including deep inside the brain, by means of the image information that is presented on the navigation monitor. Unfortunately, this accuracy is extremely difficult or impossible to measure since we can not find/place a large number of fixed and well-defined points inside the brain, localize these points with the pointer and on the monitor, and measure the error. The best we can do is to measure the *overall accuracy* of the system in a controlled laboratory setting, investigate and understand all the factors that make a clinical setting different from a laboratory setting, and finally compensate relevant factors as best we can.

The purpose of this study was to perform a preliminary measurement of the prototype system in order to get a first impression about the potential *overall accuracy* and thus, indirectly, the potential *clinical value* of ultrasound-based neuronavigation. We have also carried out a detailed description of the error sources associated with ultrasound- and MRI-based neuronavigation in order to better understand the differences between *overall accuracy* and *overall clinical accuracy*.

MATERIAL AND METHODS

Navigation equipment

The prototype system consists of a high-performance ultrasound scanner (System-fiVe, GE Vingmed Ultrasound, Horten, Norway), a medium cost computer (Daystar Genesis MP900, Daystar, USA) for image processing and navigation, and an optical positioning system (3D digitizer or tracking system). We have used the camera unit of the VectorVision system (BrainLAB, München, Germany) as the positioning system in the two-rack prototype. However, the Polaris (Northern Digital Inc., Ontario, Canada) system will be used in the single-rack product that is currently under commercialization (MISON, Trondheim, Norway) under the brand name SonoWand®.

We have developed our own navigation software that can import MRI-data, perform patient registration using fiducials or anatomical landmarks, and present the images in the same manner as conventional neuronavigation systems. In addition, the system can measure the position of the ultrasound probe and tag all the images from the ultrasound scanner with

positions from the optical system. A special algorithm reconstructs a 3D volume from, typically 200, ultrasound B-scans. This 3D volume is used by the navigation system in a similar manner as MRI or CT data. Special efforts have been made in order to improve the quality of the ultrasound images, thus making it possible to navigate directly by means of ultrasound.

In vitro experiment

An in vitro experiment was conducted in order to get an impression of the *image resolution* as well as about the *overall accuracy* of our ultrasound-based navigation system. A phantom with 27 well-defined points was scanned by 3D ultrasound. Ideally, a large number of operators should then have localized the 27 physical phantom points using a pointer that was steered in position using the crosshairs on the navigation monitor. When a point is localized, the distance between the true phantom point and the pointer tip should be measured, thus representing the overall error. However, this distance is difficult to measure in practical terms. We therefore conducted the experiment slightly differently:

Ten operators pinpointed the 27 points in image space using a mouse, yielding 270 image-point measurements. The same operators pinpointed the 27 physical points in the phantom using a pointer device, yielding 270 corresponding phantom points. The difference between these two data sets describes a residual vector data set which represents a good estimate of the overall error associated with ultrasound-guided surgery.

A frame was made of aluminum in a precise milling machine as illustrated in Figure 1. Eighteen spring-loaded cotton wires (diameter 0.2mm) formed a cube with 27 wire crossings inside a 5x5x5 cm cube. A position adapter was attached to the phantom, making it possible to calculate the positions of all 27 crossings relative to the position of the adapter with an expected accuracy of less than ±0.1 mm.
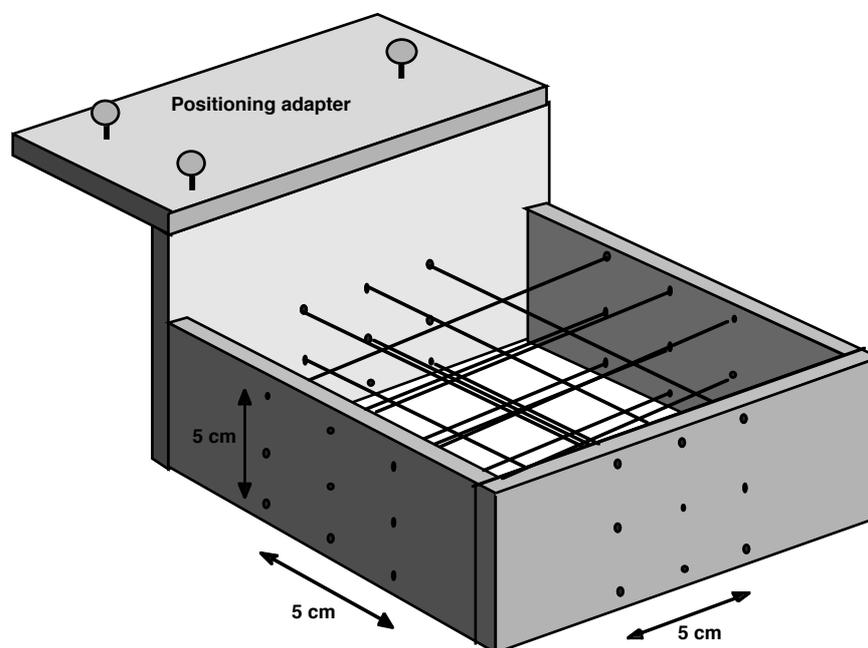
FIGURE 1.   Wire phantom used for measurement of the accuracy associated with ultrasound-guided neuronavigation. Eighteen cotton threads form a cube of 5x5x5 mm with 27 wire crossings.

The phantom was immersed in fresh water (23 ˚C) and ≈200 images were acquired during a free-hand probe translation. The camera system tracked the positions of the probe and measured the position of the phantom. A 3D volume was reconstructed into a cuberille with

0.5x0.5x0.5 mm voxel size using 1490 m/s for the velocity of sound. The center of the cube was located at a depth of approximately 6 cm in the image.

Ten different operators independently pinpointed the 27 points in the ultrasound volume by pointing at the center of the crossings using the mouse. All of these *image points* I_i (i=1-270) were all transformed to a common coordinate system denoted *camera space* through an inverse probe-to-scanplane coordinate transform.

The water tank was emptied, and the same ten operators were asked to physically pinpoint all the 27 wire crosses using a standard navigation pointer. The result was 270 corresponding *phantom points* P_i (i=1-270), also registered in camera space.

Statistics

A residual or error vector was calculated by simply subtracting all the corresponding points in the two data sets:

$$R_i = I_i - P_i = \begin{bmatrix} x_i \\ y_i \\ z_i \end{bmatrix}.$$

The endpoints of the error vectors R_i describe a cloud of points in 3D space. A projection of the 270 points into the x-z plane is shown in Figure 2(a) and the plot indicates that the error has a certain bias (the cloud is not centered in the origin) and spread. As pointed out by Maciunas et al., measurements are considered accurate if the *bias* (or mean value) of the error is low and if the measurements are *precise*, i.e. the error has low spread around the mean value (10).
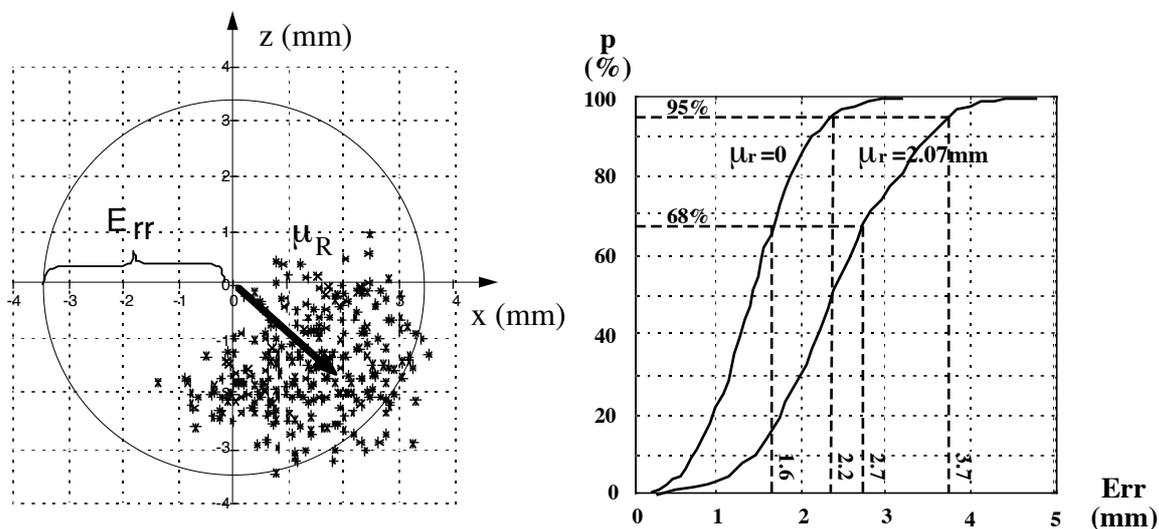
FIGURE 2. a) Two dimensional plot (projection) of all the 270 residual vectors gives an indication of the bias μ_R and the spread σ_R. The circle (actually a sphere) is used to quantify system accuracy, i.e. the per cent p of the samples that fall inside the sphere of radius Err.
b) Plot of p versus sphere radius Err for the actual measurements and for a hypothetical case where the bias is set to zero.

The bias vector is defined by

$$\mu_R = \frac{1}{270}\sum_{i=1}^{270} R_i = \begin{bmatrix} \mu_x \\ \mu_y \\ \mu_z \end{bmatrix} = \begin{bmatrix} 1.37 \text{ mm} \\ 0.45 \text{ mm} \\ -1.48 \text{ mm} \end{bmatrix},$$

and its length is given by

$$\mu_r = \|\mu_R\| = 2.07 mm.$$

This non-zero bias indicates that the system and/or the operators are subject to a systematic error. The spread can be defined by the standard deviation:

$$\sigma_R = \frac{1}{270} \begin{bmatrix} \sqrt{\sum_{i=1}^{270}(R_{ix} - \mu_x)^2} \\ \sqrt{\sum_{i=1}^{270}(R_{iy} - \mu_y)^2} \\ \sqrt{\sum_{i=1}^{270}(R_{iz} - \mu_z)^2} \end{bmatrix} = \begin{bmatrix} \sigma_x \\ \sigma_y \\ \sigma_z \end{bmatrix} = \begin{bmatrix} 1.05\, mm \\ 0.70\, mm \\ 0.86\, mm \end{bmatrix},$$

and these numbers describe the spread of the error components around their bias values. The length of the residual vector is given by:

$$r_i = \|R_i\| = \sqrt{x_i^2 + y_i^2 + z_i^2},$$

and specifies the distance (in millimeters) between two corresponding measurements. Note that the length of the residual vector describes a Rayleigh distribution rather than the more common Gaussian distribution. This means that the bias and variance parameters are less intuitive to interpret. We therefore introduce a more intuitive measure of the accuracy of a navigation system, namely the probability p that the pointer will hit within a certain distance Err from the desired target (as indicated by the crosshairs on the monitor). An approximation to this probability is found through the following numerical calculation:

We define a sphere with radius Err centered in the origin. The number of points falling inside the sphere is accumulated for different values of the sphere radius Err, and presented as a percentage of the total number of points:

$$p = \left(\frac{\#(r_i \le Err)}{270}\right)100.$$

RESULTS

Overall accuracy, ultrasound

The results of the statistical analysis of the in vitro experiment are presented in Table 1 and Figure 2(b). The error Err, or difference between image points and phantom points, was less than 2.77 mm in 68 % of the measurements, and less than 3.76 mm in 95 % of the cases. The length of the bias vector was $\mu_r = 2.07$ mm which indicates that there might be a systematic error in at least one of the processes. A simulation was performed in order to see the result of

a hypothetical case with no systematic error, see Table 1 for μ_r=0. The bias vector was subtracted from all the measurements prior to the statistical analysis, yielding a pure *spread* where 68 % of the values lay within a radius of 1.69 mm.

p	Err μ_r =2.07	Err μ_r =0
50 %	2.35 mm	1.38 mm
68 %	2.77 mm	1.69 mm
95 %	3.76 mm	2.28 mm

TABLE 1.   Results of statistical analysis of the in vitro experiment show that 95 % of the 270 image points are less than 3.76 mm from the corresponding phantom points.

Image resolution

The spatial resolution of an imaging system is defined as the smallest possible distance between two small target points which enables the operator to discriminate them as two distinct points. This parameter is practically difficult to measure in an experiment, but the data set from the wire phantom experiment contains information which gives an impression about the spatial resolution of the 4-8 MHz phased array probe. Three orthogonal images through the data set reveal the smeared-out echoes from the wires in the radial, lateral and azimuthal (perpendicular to the scan plane) directions. The width of these echoes (here defined as the width where the intensity has dropped to approximately 1/3 of the maximum value) will be used as an approximation to the spatial resolution. The result is presented in Table 2. The spatial resolution is typically 0.8 mm in the radial direction, 2 mm in the lateral direction and 2.5 mm in the azimuthal direction in the most relevant operating area for the probe.

Depth	Radial	Lateral	Azimuthal
$\approx$ 18 mm	0.7 mm	1.0 mm	1.9 mm
$\approx$ 41 mm	0.8 mm	2.0 mm	2.4 mm
$\approx$ 68 mm	0.9 mm	3.4 mm	5.4 mm

TABLE 2.   The spatial resolution of the 4-8 MHz phased array probe is approximated by the observed pulse widths from the wire phantom data set.

DISCUSSION

Overall accuracy

The overall accuracy of the system is characterized by an error number Err. This number states that the pointer will, in 68% of the cases, be no more than 2.77 mm away from the desired target (as determined by the navigation images) in the laboratory setting. If this result also applies in a clinical setting, high confidence can be put in the system when applying it for surgery of intracerebral brain lesions.

The systematic bias error μ_r = 2.07 mm is most likely caused by a known lag in the synchronization between the images and the positioning system in the prototype system. This limitation will probably be eliminated in the product version of the system that is under development; thus, the error can possibly be brought down from 2.77 mm to approximately 1.7 mm. If this can be achieved, other error sources such as steadiness of hand and the surgeon's ability to pinpoint a small structure, will probably be dominating.

Ideally, a large number of scans and measurement sequences should have been performed. However, the statistical analysis in the experiment was based on one single 3D scan of the wire phantom. This fact indicates that the error may have been underestimated. On the other hand, the experiment was carried out by letting ten operators pinpoint 27 wire crossings using a pointer. Since this process does not apply in a clinical setting, the error may have been overestimated.

The spatial resolution of the imaging system does not necessarily directly influence the overall accuracy of the system. However, limited resolution causes the objects to be smeared out and appear fuzzy in the images, and this in turn, can reduce the precision in placement of the crosshairs. The spatial resolution of a 3D ultrasound system is best in the focal region of the probe. For the phased array probe applied in this study, this region ranges from approximately two to five centimeters, see Table 2. The resolution will gradually degrade both proximal and distal to this region.

Probe calibration

One of the most critical tasks with respect to overall accuracy is the process of ultrasound probe registration or calibration, i.e. finding a transformation matrix between the position sensor (which is attached to the probe housing) and the ultrasound scan plane (the image plane can neither be seen, nor touched). This calibration can be done once in the factory or laboratory under controlled conditions so that the result does not depend on an arbitrary operator. Ideally, one unique position sensor should be permanently attached to one unique ultrasound probe before calibration in the factory. However, this option represents practical difficulties in that most ultrasound probes cannot be sterilized, while the positioning sensor can. Both devices can alternatively be covered by sterile drape, but this can easily hinder free sight to the cameras.

To our knowledge, there is still no good solution to this problem for optical tracking systems, but current strategies include two methods: I) calibration of any probe to any position sensor inside the operating room using a special ultrasound calibration phantom and algorithm or II) calibration of one particular probe to one particular position sensor in the laboratory or factory, using a special adapter that ensures repetitive and precise attachment between the two, even through a sterile probe drape. The latter option allows neither replacement of the probe nor the positioning adapter without returning both to the factory for new calibration. Our system is based on method II in order for us to have better control of the accuracy.
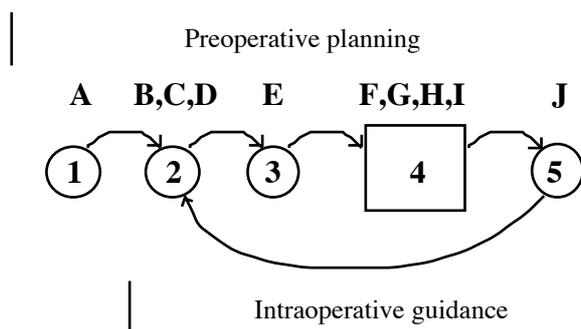
The users of neuronavigation systems would benefit from specifications of *overall clinical accuracy* of a specific system. Unfortunately the vendors can, at best, only provide reliable specifications on *overall accuracy* (laboratory measurements). Clinical settings involve a number of user- and hospital-dependent actions that are not in control of the vendor. In order to gain a better understanding of the relation between laboratory and clinical settings, we will in the following describe the most relevant error sources associated with ultrasound- and MRI-based navigation:

Error sources associated with ultrasound-based navigation

This section, with reference to Figure 3, gives a detailed description of the major actions required to prepare and perform an ultrasound-based neuronavigation case. The error sources associated with each action are described, and attempts are made to quantify the errors.

Preoperative planning: Preoperative planning is normally performed using preoperative MRI-data (not necessary, but convenient). The craniotomy is planned, and in some cases a separate mini-craniotomy is made for the ultrasound probe. The skull is opened and the first 3D ultrasound scan is acquired for further planning purposes. Before ultrasound imaging starts, we attach a positioning adapter to the ultrasound probe housing, action item number

{1}, see Figure 3. Proper design of the adapter ensures a small or negligible repeatability error associated with this process, even through the sterile drape, error source number {A}.

FIGURE 3.   The error chain associated with ultrasound-based neuronavigation.

Action items ultrasound	Error sources ultrasound	Error (mm)		
		<1	1-2	>2
1 = mount position sensor on probe	A = sensor attachment repeatability error	x		
2 = acquire 3D ultrasound data	B = probe sensor tracking error	x		
	C = sensor to scan plane transformation error (factory probe calibration)	x	x	x
	D = synchronization error between pos. data and images	x	x	
3 = load images, reconstruct volume	E = geometric errors in 3D scanconversion	x	x	
4 = point and reconstruct images with colored cross	F = reconstruction and interpolation error	x		
	G = quantization error, colored cross	x		
	H = interpretation error	x		
	I = positioning system error, pointer tip	x		
5 = surgery and brain movements	J = brain shift error	x		

A 3D ultrasound volume is then acquired using the optical system to track the position and orientation of the probe ≈10 times per second {2}. The probe is tilted ≈90 degrees in ≈15 seconds in order to acquire ≈200 images from the volume of interest in the brain. This process is subject to errors in tracking the probe position sensor {B}, errors in estimating the scan plane orientation relative to the position sensor {C} (laboratory calibration) and errors in the synchronization between the positioning data and the images {D}.

The images with position tags are transferred to the navigation computer and reconstructed/ scan-converted into a regular 3D volume {3} with a certain grid resolution. This process is subject to geometric errors {E} caused by: I) the difference in real velocity of sound in the brain relative to what is assumed in the algorithm, II) uncertainty in the discrete representation of the images due to the selected resolution in the 3D grid, and III) uncertainty in the positioning data due to the slow sampling rate (currently ≈10 Hz) relative to the continuous probe movement.

Image-guided planning or surgery can then be performed by applying a pointer or another calibrated surgical instrument at the skin surface or within the brain. The navigation system will reconstruct a 2D image from the 3D volume and draw a crosshairs in the images at a location given by the pointer {4}. This process is subject to a reconstruction and interpolation error from 3D to 2D data {F}, a quantization error in the positioning of the crosshairs in the images {G} and an operator-dependent interpretation error when the surgeon is supposed to position the crosshairs at a desired point in image space {H}. The pointer tip is now supposed

to be located exactly at the spot on the patient which is indicated by the crosshairs in the images. However, this will rarely happen due to an error in measuring and estimating the pointer tip position {I} as well as the contribution from all the other error sources that are listed above. The overall error will not be equal to the sum of all the individual errors, but rather be given by the contribution from all errors in a three-dimensional stochastic process which can not easily be quantified.

Intraoperative guidance: An important advantage of ultrasound-guided navigation is the ability to do repetitive 3D imaging during surgery {2} and thus work with a data set that has recently been updated according to possible changes in the brain anatomy. The error or discrepancy between the images and the anatomy will therefore be small or negligible {J}.
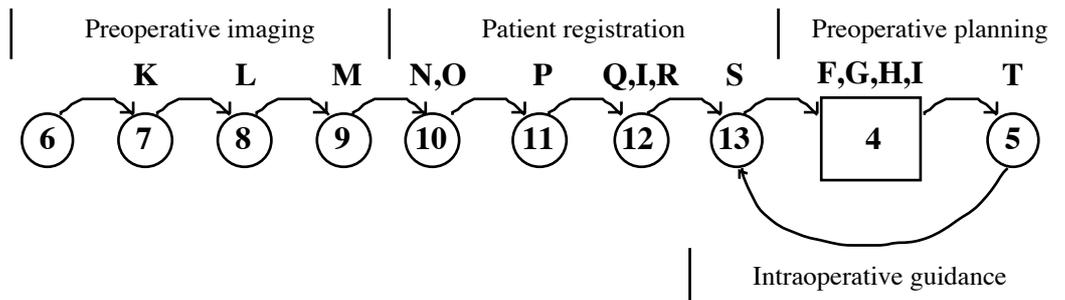
Error sources associated with MRI-based navigation

Preoperative imaging: Fiducials are glued to the skin, action item number one {6}, see Figure 4. The patient is then positioned in the MRI scanner in the supine position {7}. The normal procedure in our hospital is to stabilize the head with bitemporal paddings and a strap across the forehead. The fiducials can easily slide several millimeters during the procedure and cause a significant error, labeled {K}. We normally place five fiducials on the patient, but we avoid the occipital region of the head and exclude the paddings and the strap in order to minimize this error. A 3D MRI scan is performed {8}, and a digital data set is acquired. Inhomogeneities of the magnetic field and nonlinear gradients cause geometric errors {L} between the true anatomy and image information. We have no documentation for our MR-system, but values in the order of 2.1 mm can be expected according to Sumanaweera et al. (15). The selected slice thickness and slice distance will also affect the accuracy (10). The images are transferred to the navigation system and organized into a regular 3D volume of digital numbers {9}. If the original images have a different pixel resolution and image distance than the specified voxel resolution of the regular 3D volume (normally not the case), this process will introduce a quantization error {M}.

Patient registration: Next, the patient is placed in the Mayfield frame on the operating table {10}. Brain shift due to gravity can occur {N}, especially if the head orientation is different from that in the MRI scanner. This effect is often accounted for in functional stereotaxy. The skin and fiducials can also slide and deviate several millimeters from their original position in the MRI scanner {O} due to gravity and/or manipulation during Mayfield-frame placement. This effect is probably most pronounced in elderly people.

The next step in the procedure is to register the position of the fiducials in the 3D data set by a manual or automatic method {11}. This process is subject to operator errors and possibly algorithm/quantization errors depending on the slice distance and the zoom factor (magnification) of the images in which the operator is supposed to pinpoint the fiducials {P}. Patient registration is then performed by pointing at the fiducials with a pointer {12}. The fiducials/skin can again slide several millimeters (5) due to the applied pointer pressure {Q}, and the positioning system (3D digitizer or tracking system) is subject to a certain error in the measurement and calculation of the pointer tip position {I} (8), see Figure 3. The operator normally attempts to point at the center of the fiducial, but this procedure is also subject to a pointing error {R} that can depend on the fiducial type being used.

A matching algorithm is then applied in order to find a best match, i.e. a coordinate transform between patient space and image space {13}. A variety of algorithms are available for this task, and different systems may cause different results, given the same input data. This variance can be associated with an error {S}. Phantom tests and clinical studies have shown that methods based on fiducials perform better than a surface fitting algorithm (5, 13, 17).

Action items MR	Error sources MR	Error (mm)		
		<1	1-2	>2
6 = glue fiducials to skin				
7 = position patient in MR-scanner	K = skin/fiducial slide	x	x	x
8 = perform 3D MR-scan	L = geometric distortion in MR data	x	x	x
9 = load images, generate volume	M= quantization error in volume reconstruction	(x)		
10 = position patient in Mayfield frame	N = brain shift due to gravity	x	x	
	O = skin/fiducial slide due to gravity	x	x	x
11 = register fiducials in MR data set	P = fiducial pinpointing error	x	x	
12 = touch fiducials with pointer	Q = skin/fiducial slide due to pointing	x	x	x
	I = positioning system error, pointer tip	x		
	R = pointing error	x	x	
13 = match images to patient	S = matching algorithm error	x	x	
4 = point at patient and reconstruct images with colored cross	F = reconstruction and interpolation error	x		
	G = quantization error, colored cross	x		
	H = interpretation error	x		
	I = positioning system error, pointer tip	x		
5 = surgery and brain movements	O = discrepancy between anatomy and images	x	x	x

FIGURE 4.   The error sources associated with MRI-based neuronavigation is longer than that of ultrasound due to the need for preoperative imaging and patient registration.

Preoperative planning: Surgical planning can then start by moving the pointer over the skin surface. Again the position of the pointer/instrument determines which images are displayed on the monitor and the position of the crosshairs in the images {4}. The error sources associated with this process are the same as for ultrasound-based navigation {F}, {G}, {H} and {I}.

Intraoperative guidance: The last action item on the list is the actual surgery in the brain {5}. Most procedures will cause a certain brain shift and possibly a deformation of normal and pathological structures. These changes may easily cause a discrepancy between image space and patient space in the order of several millimeters or even centimeters {T}.

Overall accuracy versus overall clinical accuracy

The *overall clinical accuracy* of a navigation system will be determined by the contribution from all the individual error sources described in Figure 3 and 4. The net effect will not be the sum of all error sources, rather a stochastic contribution from all terms. Ultrasound-based navigation does not require patient registration, thus this modality is less susceptible to user- and procedure dependent error sources. Error sources H and J are affected by the ultrasound

user, the remaining 8 are in control of the system vendor. In MRI-based navigation, the user and procedure affect error sources K, L, N, O, P, Q, R, H and O, while the remaining 5 error sources are in control of the vendor. This indicates that the *overall clinical accuracy* (deep in the brain) of an ultrasound-based system may lie closer in value to the *overall accuracy* (laboratory test) than is the case with conventional MRI-based systems.

Accuracy measurements in the literature

The literature is sparse on accuracy measurements of ultrasound-based systems. Hata et al. report an rms-error of 3.1 mm at 10 mm depth from the transducer (4), while Comeau et al. report a ≤1.3 mm error between pixels in ultrasound and MRI (2). Both systems are two-dimensional.

More information is available on MRI-based navigation systems, and the studies span over a broad range of different navigation equipment, imaging and registration techniques as well as setup and measurement techniques. A common way of measuring the *overall accuracy* has been to apply a rigid and precise phantom/head model in the laboratory, typical results are ≈2mm (5), >1.3mm (7) and ≈1mm (14). Other groups have measured the *overall accuracy* in a clinical setting where fiducials and/or the skin surface have been applied, results: 2.1-3.1mm (13), 1.6-4.7mm (17) and <2mm (11). These results are probably close to the true *overall clinical accuracy* numbers. However, a perfect match on the skin surface does not necessarily mean there is a perfect match deep inside the brain. Rotation errors in the registration procedure may be difficult to discover, and may result in significant errors at the skull base. Schaller et al. have actually measured inside the brain using landmarks such as the internal table of the skull, the falx, the tentorium, or the clinoid processes (12). They report errors in the order of 3mm. This approach is interesting, but it is difficult to obtain a precise measurement in more than one or two dimensions, and the number of measurements in one patient will be sparse.

CONCLUSIONS

The *overall accuracy* (laboratory test) of our prototype ultrasound-based navigation system is measured to 2.7 mm with a potential for improvement to 1.7 mm when the optical positioning system is changed and better adapted to the navigation system. Ultrasound-based navigation does not require patient registration. Thus, this modality is less susceptible to user- and procedure-dependent error sources associated with patient registration. The *overall clinical accuracy* (deep in the brain) of an ultrasound-based system is therefore expected to lie closer in value to the *overall accuracy* (laboratory test) than is the case with conventional MRI-based systems (brain shift not encountered). Numbers from the literature indicate that the *overall accuracy* and *overall clinical accuracy* values for MRI-based systems lie in the area 1-2 mm and 2-4 mm respectively. Thus, these initial results justify the commercialization of our ultrasound-based neuronavigation system in terms of accuracy considerations.

Acknowledgments

References

1. Bucholz RD, Yeh DD, Trobaugh J, McDurmont LL, Sturm CD, Baumann C, Henderson JM, Levy A, Kessman P: The Correction of Stereotactic Inaccuracy Caused by Brain Shift Using an Intraoperative Ultrasound Device. Lecture Notes in Computer Science Proceedings of the 1997 1st Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery, CVRMed-MRCAS'97, Grenoble, 1997.

2. Comeau RM, Sadikot AF, Fenster A, Peters TM: Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. Med. Phys. 27 (4):787-800, 2000.

3. Gronningsaeter A, Kleven A, Ommedal S, Aarseth TE, Lindseth F, Langø T, Unsgård G: SonoWand®, an Ultrasound-Based Neuronavigation System. Neurosurgery Vol. ?, No. ?:??-??, 2000.

4. Hata N, Dohi T, Iseki H, Takakura K: Development of a frameless and armless stereotactic neuronavigation system with ultrasonographic registration. Neurosurgery 41:608-613, 1997.

5. Helm PA, Eckel TS: Accuracy of Registration Methods in Frameless Stereotaxis. Comput Aided Surg 3:51-56, 1998.

6. Hirschberg H, Unsgaard G: Incorporation of ultrasonic imaging in an optically coupled frameless stereotactic system. Acta Neurochir Suppl (Wien) 68:75-80, 1997.

7. Kall B, Goerss SJ, Stiving SO, Davis DH, Kelly PJ: Quantitative Analysis of a Noninvasive Stereotactic Image Registration Technique. Stereotact Funct Neurosurg 66:69-74, 1996.

8. Kaus M, Steinmeier R, Sporer T, Ganslandt O, Fahlbusch R: Technical accuracy of a neuronavigation system measured with a high-precision mechanical micromanipulator. Neurosurgery Vol. 41, No. 6:1431-1437, 1997.

9. Koivukangas J, Louhisalmi Y, Alakuijala J, Oikarinen J: Ultrasound-controlled neuronavigator-guided brain surgery. J Neurosurg 79:36-42, 1993.

10. Maciunas RJ, Galloway RL, Latimer JW: The Application Accuracy of Stereotactic Frames. Neurosurgery 35, No. 4:682-695, 1994.

11. Sandemann DR, Patel N, Chandler C, Nelson RJ, Coakham HB, Griffith HB: Advances in image-directed neurosurgery: preliminary experience with the ISG Viewing Wand compared with the Leksell G frame. Br J Neurosurg 8:529-544, 1994.

12. Schaller C, Meyer B, van Roost D, Schramm J: Image guided microsurgery with a semifreehand neuronavigational device. Comput Aided Surg 2:162-171, 1997.

13. Sipos EP, Tebo SA, Zinreich SJ, Long DM, Brem H: In Vivo Accuracy Testing and Clinical Experience with the ISG Viewing Wand. Neurosurgery 39, No. 1:194-204, 1996.

14. Smith KR, Frank KJ, Bucholz RD: The NeuroStationTM - a highly accurate, minimally invasive solution to frameless stereotactic neurosurgery. Computerized Medical Imaging and Graphics Vol. 18, No. 4:247-256, 1994.

15. Sumanaweera TS, Adler JR, Napel S, Glover GH: Characterization of Spatial Distortion in Magnetic Resonance Imaging and Its Implications for Stereotactic Surgery. Neurosurgery Vol. 35, No. 4:696-704, 1994.

16. Trobaugh JW, Richard WD, Smith KR, Bucholz RD: Frameless stereotactic ultrasonography: Method and applications. Comput Med Imaging Graph Vol. 18, No.4:235-246, 1994.

17. Villalobos H, Germano IM: Clinical Evaluation of Multimodality Registration in Frameless Stereotaxy. Comput Aided Surg 4:45-49, 1999.

Paper E

Bayesian 2D Deconvolution: A Model for Diffuse Ultrasound Scattering

Oddvar Husby, Torgrim Lie, Thomas Langø, Jørn Hokland and Håvard Rue

ABSTRACT

Observed medical ultrasound images are degraded representations of the true acoustic tissue reflectance. The degradation is due to blur and speckle, and significantly reduces the diagnostic value of the images. In order to remove both blur and speckle we have developed a new statistical model for diffuse scattering in 2D ultrasound radio-frequency images, incorporating both spatial smoothness constraints and a physical model for diffuse scattering. The modeling approach is Bayesian in nature, and we use Markov chain Monte Carlo methods to obtain the restorations. The results from restorations of some real and simulated radio-frequency ultrasound images are presented, and compared with results produced by Wiener filtering.

O. Husby and H. Rue are with the Department of Mathematical Sciences, NTNU, 7491 Trondheim, Norway. (O. HUsby is corresponding author, e-mail: Oddvar.Husby@math.ntnu.no). T. Lie and T. Langø are with SINTEF Unimed, Ultrasound, 7465 Trondheim, Norway. J. Hokland is with the Department of Computer and Information Sciences, NTNU, 7491 Trondheim, Norway.

I. Introduction

Ultrasound B-scan images are degraded representations of the true acoustical scatterers in the imaged tissue. The degradation introduces image artifacts (speckle) and limits the spatial resolution, hence reducing the diagnostic value of the images. The dominant degradation effects are blur and speckle due to what may be approximated as a convolution of the true reflectance field r with the imaging system point spread function h, with the addition of noise introduced in the image formation process. We assume the point spread function to be spatially invariant [1], [2] , and thus the observed signal y can be modeled as

$$y_{i,j} = \{h * r\}_{i,j} + \eta_{i,j} = \sum_{k,l} h_{k,l} r_{i-k,j-l} + \eta_{i,j}, \quad (i,j) \in \mathcal{I} \tag{1}$$

where $\eta_{i,j}$ are independent Gaussian variables with mean zero and common variance τ^2, and $\mathcal{I}$ is a rectangular lattice with $n_r \times n_l = n$ sites. The indices i and j represent radial and lateral location, respectively, while k and l are local coordinates for the point spread function centered at (i,j). Note that the assumption of spatially invariant point spread function is made out of computational convenience only, since our focus is on the modeling of the image formation process. We study the effect of this assumption in [3] , and refer to e.g. [4], [5] for other discussions of the effect of the point spread function.

Due to the limited bandwidth of the point spread function and the presence of additive noise, perfect deconvolution is impossible. As an alternative, statistical restoration of ultrasound images attempts to remove the effect of the degradation by estimating the true reflectance image r. A simple and commonly used technique is Wiener filtering [6], [1] ; other methods include multiscale wavelet analysis [7] , but even with these methods the restored images suffer from blur and speckle.

We propose a new method for deconvolution of ultrasound images, based on Markov random field models [8] and Bayesian statistical methods [9], [10] . These methods have become increasingly popular over the last decade or two, and represent a general methodology for restoration of indirectly observed measurements. The most important feature is the ability to model the randomness and uncertainty inherent in the system. Specifically, in dealing with ultrasound the methodology allows us to use stochastic models of the imaged tissue and the image formation process, thus incorporating *a priori* knowledge and measurement error in a way not possible in deterministic restoration methods. Also, we are able to quantitatively assess the errors in our restorations, producing a *distribution* over the image space rather than a single filter-image restoration.

Our Markov random field model is based on the standard model for diffuse scattering [11], [12] , where a point scatterer within one resolution cell is modeled to be a Gaussian random variable with mean zero and variance depending on the tissue type within the cell. Our idea is to let a spatially varying variance field represent the *underlying structures* of the imaged region; based on the assumption that this field is the *sole* descriptor needed to obtain faithful image restorations. The argument for the scattering model in [12] goes as follows: Assuming that the concentration of scatterers is large relative to the wavelength of the pulse, the reflectance $r_{i,j}$ at each location (i,j) is a Gaussian random variable with mean zero and variance $\sigma_{i,j}^2$. The reflectances at different points are assumed to be uncorrelated, and thus the reflectances have the

conditional distributions (i.e. the distribution of $r_{i,j}$ given that the value of $\sigma_{i,j}^2$ is known)

$$p(r_{i,j} \mid \sigma_{i,j}^2) = \frac{1}{\sqrt{2\pi\sigma_{i,j}^2}} \exp\left\{-\frac{r_{i,j}^2}{2\sigma_{i,j}^2}\right\}, \ \forall \ (i,j) \in \mathcal{I}. \tag{2}$$

Note that the radio frequency echo in general consists of a diffuse and a coherent scatter component [13] , and that our model is valid only when the coherent component is negligible and when the number of diffuse scatterers is large enough for the law of large numbers to apply. This might not be the case for all tissue types; however, there exists statistical tests for identifying the regions of an image for which the diffuse scattering assumptions are true [14] . For the reminder of this paper we will assume these assumptions to hold. Furthermore, we see that the estimation of the variance field is essential; given the variance field, the reflectance can be simulated by drawing from the distribution (2).

The restoration can now be formulated as a Bayesian estimation problem. We define an appropriate prior distribution $p(\sigma^2)$ for the variance field, trying to model the spatial properties inherent in the imaged regions. Using Bayes theorem, the prior distribution is then combined with the likelihood model $p(y \mid \sigma^2)$ to obtain the posterior distribution $p(\sigma^2 \mid y, h)$ from which inferences can be made:

$$p(\sigma^2 \mid y, h) \propto p(y \mid \sigma^2, h)p(\sigma^2). \tag{3}$$

We define the prior model in Section II, selecting a distribution from the family suggested in [15] , where it was used for tomographic reconstruction. These models have the property that they are able to recover discontinuities in piecewise smooth images. This property is essential in our application, as discontinuities may represent interfaces between regions of different tissue types. Following the definition of the prior we compute the full posterior distribution. The distribution is analytically intractable; hence we have to rely on iterative stochastic techniques to make inferences. Markov chain Monte Carlo methods [9], [16], [17] are well suited for sampling Markov random field models, but we argue that when the model is formulated in terms of σ^2 only, these techniques can be very inefficient. However, in Section III we show that this problem can be solved by augmenting the model with the reflectance field r. We show that y and σ^2 are conditionally independent given r, and thus that r can be viewed as an auxiliary variable [18], [19] decoupling y and σ^2. When sampling from the joint distribution of r and σ^2 this fact can be utilized to construct an algorithm with better computational properties.

A related approach can be found in [20] , who use a discrete Markov model for the variance field. However, their discrete model made the method unstable. Furthermore, it is unclear how well a discrete model describes the underlying structure of e.g. human tissue, and also how well it can account for the radial reduction in intensity seen in ultrasound images. Our main contribution is the formulation of the continuous variance field, which from a computationally viewpoint is far better, making sampling and estimation easier, faster and more robust. Specifically, we avoid the problems of choosing the number of levels in the discrete variance field, and of estimating the variance at each level. We also feel that a continuous model is a better descriptor of the imaged regions, since it is reasonable to assume that scattering properties varies slightly within tissue regions. Finally, radial reduction in intensity can only be explained by a continuous model.

We conclude the paper with examples and a brief discussion in Sections IV and V.

II. BAYESIAN MODEL FORMULATION

A. Introduction

With a slight abuse of notation we will use $p(\cdot)$ to represent both the distribution and the density of the random variables. Recall that the images are defined on the rectangular lattice $\mathcal{I}$, using the polar coordinate representation of B-scan ultrasound images. All images are defined on the domain $\mathcal{I}$, letting e.g. $\sigma^2 = \left(\sigma^2_{1,1}, \ldots, \sigma^2_{n_r,n_l}\right)^T$, where n_r and n_l represent the number of samples in the radial and lateral directions, respectively. For any subset $\mathcal{A} \subset \mathcal{I}$ we write $r_{\mathcal{A}} = \{r_{i,j}\,;\,(i,j) \in \mathcal{A}\}$ and $r_{-\mathcal{A}} = \{r_{i,j}\,;\,(i,j) \in \mathcal{I} \setminus \mathcal{A}\}$, and write $p(r_{\mathcal{A}})$ and $p(r_{-\mathcal{A}})$ for the corresponding densities. Our approach is based on Markov random field (MRF) models. One of the key features of Markov random fields is that each site (i,j) has associated with it a neighborhood $\partial(i,j)$, such that the conditional probability of the field at the site (i,j), given the values at all other sites, only depend on the values of the field in $\partial(i,j)$ [8]. We write $(i,j) \sim (k,l)$ if the sites (i,j) and (k,l) are neighbors.

B. Prior distribution for the variance field

The crucial part of our model is the formulation of the prior distribution for the variance field σ^2. We let the prior be of the form

$$p(\sigma^2) \propto \exp\left\{ -\beta \sum_{(i,j)\sim(k,l)} w_{ij,kl}\phi\left(\ln\left(\sigma^2_{i,j}/\sigma^2_{k,l}\right)\right) - \sum_{(i,j)} \psi(\sigma^2_{i,j}) \right\}, \qquad (4)$$

where the first sum is over all pairs of neighboring pixels in the lattice $\mathcal{I}$. $\phi(\cdot)$ is a spatially structured interaction term, and $\psi(\cdot)$ is an unstructured level term. The choice of the prior model is justified from physical considerations about the scattering medium, and is explained below.

In the diffuse scattering model the variance depends on the acoustical properties of the scattering medium [11], and it is thus reasonable to assume the following:
• Variances tend to be approximately constant in regions of homogeneous tissue.
• Abrupt changes in variance may occur at interfaces between different tissue regions, as the scattering intensity varies between tissue types.
In other words, we believe the field to be piecewise planar with each subregion corresponding to the different tissue types in the imaged region. Hence prior distributions with convex interaction functions will be unsuited for our purpose, since they tend to over-smooth images, favoring gradual changes in intensity. We will instead use an interaction function of the form

$$\phi(u) = \frac{\left|\frac{u}{\delta}\right|^{\gamma}}{1 + \left|\frac{u}{\delta}\right|^{\gamma}}, \qquad (5)$$

as suggested in [15]; this model is known to provide for recovery of discontinuities in blurred and noisy images. The non-interpolating property of the model is due to its concavity and finite limit as $u \to \infty$, favoring large intensity jumps instead of gradual changes; see Figure 1. We refer to [21] and [22] for a further discussion of the properties of this class of models.

In addition to the spatially structured term (5), the unstructured term is chosen from the conjugate family of inverse Gamma distributions; and can be seen as a prior on the level of the variances. The parameter ϵ of the inverse Gamma distribution is chosen small to make the dis-
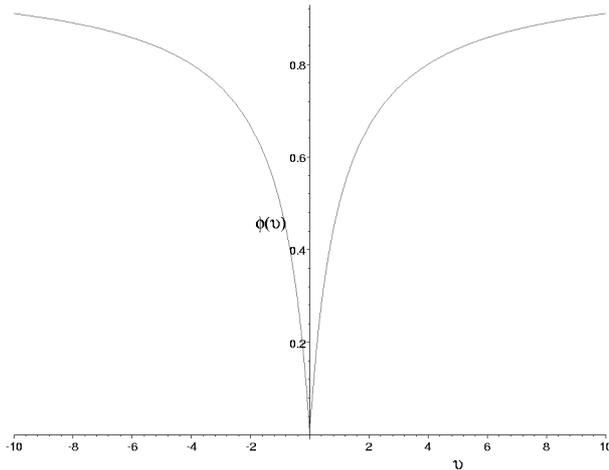
Fig. 1. Plot of the interaction function $\phi(u)$ with $\delta = 1$ and $\gamma = 1$.

tribution sufficiently vague. Combining the two terms, we get the following prior density:

$$p(\sigma^2) \propto \exp\left\{ -\beta \sum_{i,j \sim k,l} w_{ij,kl}\, \phi\left(\ln\left(\sigma_{i,j}^2/\sigma_{k,l}^2\right)\right) - \sum_{i,j} \left(\epsilon/\sigma_{i,j}^2 + (\epsilon+1)\ln\sigma_{i,j}^2\right) \right\}$$

$$= \exp\left\{ -\Phi(\sigma^2) - \Psi(\sigma^2) \right\}, \tag{6}$$

where β and ϵ are positive parameters, and $w_{ij,kl}$ is a weight depending on the neighborhood system. In the experiments we use a second order neighborhood system with $w_{ij,kl} = 1$ for orthogonal nearest neighbors, and $w_{ij,kl} = 1/\sqrt{2}$ for diagonal nearest neighbors. The parameter β determines the strength of the interaction, while δ in Eq. (5) determines the size of the intensity jumps allowed. The parameter γ is set equal to 1 in all experiments, since for $\gamma \leq 1$ sharp boundaries are favored over gradual transitions [21] , while larger values of γ permit more variations within homogeneous regions.

C. Posterior distribution

As noted in Section I, it is possible to formulate the model in terms of the variance field σ^2 only. Given an estimate $\hat{\sigma}^2$, the reflectance field r can be simulated readily using Eq. (2). However, in Appendix A we show that the posterior density $p(\sigma^2 \mid y)$ has poor sampling properties, leading to computational problems. This deficiency is amended by considering the augmented model (σ^2, r) as follows.

C.1 Model augmentation

The idea behind model augmentation is to introduce auxiliary variables that allow us to construct sampling algorithms that in some sense are easier to work with. The variable z of interest is augmented by a variable u, and to generate realizations from $p(z)$ one specifies the conditional distribution $p(u \mid z)$ and builds $p(z, u) = p(u \mid z)p(z)$, which is then sampled from. The idea is to specify u in such a way that the joint sampling of (z, u) is in some sense easier than the sampling of z alone. Note that $p(z, u)$ maintains the marginal distribution $p(z)$.

In our setting we let r be the auxiliary variable. The augmented model is then (σ^2, r) with distribution

$$p(\sigma^2, r \mid y, h) \propto p(y \mid r, h)p(r, \sigma^2) \propto p(y \mid r, h)p(r \mid \sigma^2)p(\sigma^2). \tag{7}$$

Note that σ^2 and y are conditionally independent given r, thus r *decouples* σ^2 and y, making the variance field σ^2 conditionally independent of the blur. This is important because the blur causes long-range dependency, making unconditional sampling of σ^2 very difficult.

C.2 Posterior density

We now proceed to find the joint density for σ^2 and r. According to Eq. (1), the density for the observed image y can be written as

$$p(y \mid r, h) = \prod_{(i,j)\in\mathcal{I}} p(y_{i,j} \mid r, h) \propto \prod_{(i,j)\in\mathcal{I}} \exp\left\{ -\frac{1}{2\tau^2}\left(y_{i,j} - \{h * r\}_{i,j} \right)^2 \right\}, \tag{8}$$

where the convolution between h and r can be considered a constant shifting of the expectation of y, and is as defined in equation (3).

By inserting the expression for $p(\sigma^2)$, $p(r \mid \sigma^2)$, and $p(y \mid r, h)$ in Eq. (7), the joint density for σ^2 and r can be written as

$$p(\sigma^2, r \mid y, h) \propto p(y \mid r, h)p(r \mid \sigma^2)p(\sigma^2)$$

$$\propto \exp\left\{ -\sum_{i,j}\left(\frac{1}{2\tau^2}\left(y_{i,j} - \{h * r\}_{i,j} \right)^2 + \frac{r_{i,j}^2 + 2\epsilon}{2\sigma_{i,j}^2} + \left(\epsilon + \frac{3}{2} \right)\ln\sigma_{i,j}^2 \right)\right.$$

$$\left. -\beta \sum_{i,j\sim k,l} w_{ij,kl}\phi\left(\ln\left(\sigma_{i,j}^2/\sigma_{k,l}^2 \right) \right) \right\}. \tag{9}$$

We will later show that our sampling algorithm will use the conditional distribution of $r_{i,j}$ and $\sigma_{i,j}^2$ at single pixels (i, j), and from Eq. (9) we arrive at

$$p(\sigma_{i,j}^2 \mid \sigma_{\partial(i,j)}^2, r) \propto \exp\left\{ -\left(\epsilon + \frac{r_{i,j}^2}{2} \right)\frac{1}{\sigma_{i,j}^2} - \left(\epsilon + \frac{3}{2} \right)\ln\sigma_{i,j}^2 - \beta\sum_{(k,l)\in\partial(i,j)} w_{ij,kl}\phi\left(\ln\sigma_{i,j}^2/\sigma_{k,l}^2 \right) \right\}$$

$$\tag{10}$$

$$p(r_{i,j} \mid r_{-(i,j)}, \sigma^2, y, h) \propto \exp\left\{ -\frac{1}{2\tau^2}\sum_{m,n}\left(y_{m,n} - \sum_{(k,l)\neq(m-i,n-j)} h_{k,l}r_{m-k,n-l} - h_{m-i,n-j}r_{i,j} \right)^2 - \frac{r_{i,j}^2}{2\sigma_{i,j}^2} \right\}.$$

$$\tag{11}$$

By performing the multiplication of the quadratic parenthesis and rearranging the terms, Eq. (11) can be written as:

$$p(r_{i,j} \mid r_{-(i,j)}, \sigma^2, y, h) \propto \exp\left\{ -r_{i,j}^2\left(\frac{1}{2\tau^2}\sum_{m,n} h_{m-i,n-j}^2 + \frac{1}{2\sigma_{i,j}^2} \right) + 2r_{i,j}\left(\frac{1}{2\tau^2}\sum_{m,n}(h_{m-i,n-j}\Upsilon_{m,n}) \right) \right\}$$

$$\tag{12}$$

$$\propto \exp\left\{ -r_{i,j}^2 g_1(\sigma^2, h, \tau) + 2r_{i,j}g_2(r_{-(i,j)}, y, h, \tau) \right\},$$

where

$$\Upsilon_{m,n} = y_{m,n} - \sum_{(k,l) \neq (m-i,n-j)} h_{k,l} r_{m-k,n-l} = y_{m,n} - \sum_{k,l} h_{k,l} r_{m-k,n-l} + h_{m-i,n-j} r_{i,j}. \quad (13)$$

Concerning implementation, note that we can store the values of $\Upsilon_{m,n}$ in a two-dimensional table, reducing the the double summation in Eq. (11) to a single sum over the values of $\Upsilon_{m,n}$. For every pixel updated, we have to update as many values in the table as there are elements in the point spread function h. Thus we reduce the effective neighborhood size of r from the support of h convolved with h to only two times the support of h.

From Eq. (12) one can recognize that the conditional distribution for $r_{i,j}$ given $y, r_{-(i,j)}$ and σ^2 to be a Gaussian distribution with mean

$$\mu_{i,j} = \frac{g_2}{g_1} = \frac{\frac{1}{\tau^2} \sum_{m,n} h_{m-i,n-j} \Upsilon_{m,n}}{\frac{1}{\tau^2} \sum_{m,n} h_{m-i,n-j}^2 + \frac{1}{\sigma_{i,j}^2}}, \quad (14)$$

and variance

$$\kappa_{i,j}^2 = \frac{1}{2 g_1} = \left[\frac{1}{\tau^2} \sum_{m,n} h_{m-i,n-j}^2 + \frac{1}{\sigma_{i,j}^2} \right]^{-1}, \quad (15)$$

and thus the conditional sampling of $r_{i,j}$ is straightforward using one of the standard algorithms for sampling Gaussian random variables.

III. THE SAMPLING ALGORITHM

A. Metropolis-Hastings algorithm

Given the above model we want to use the *posterior means* as estimators of the variance and reflectance fields. The densities in (10,12) are analytically intractable, hence we have to rely on Markov chain Monte Carlo methods as follows: We construct an ergodic Markov chain $\sigma^2(0), r(0), \sigma^2(1), r(1), \ldots$ converging in distribution to $p(\sigma^2, r \mid y, h)$. This is done as follows: For each $k = 1, 2, \ldots$ we generate $\sigma^2(k)$ from $r(k-1)$ using $p(\sigma^2(k) \mid r(k-1))$, and then we generate $r(k)$ from $\sigma^2(k)$ using $p(r(k) \mid \sigma^2(k), y, h)$. After a burn-in of t_0 iterations an estimate of the posterior means can be found as

$$\widehat{E(\sigma^2 \mid r)} = \frac{1}{t} \sum_{i=t_0+1}^{t_0+t} \sigma^2(i) \qquad \widehat{E(r \mid \sigma^2, y, h)} = \frac{1}{t} \sum_{i=t_0+1}^{t_0+t} r(i). \quad (16)$$

Several methods exist for constructing the Markov chain $\sigma^2(0), r(0), \sigma^2(1), r(1), \ldots$; we refer to [23] for a overview of the theory and methodology related to Markov chain Monte Carlo methods. Among the most common methods is the Metropolis-Hastings algorithm [16], which is the one we have used in the experiments. Assuming we want to sample from a distribution $p(z)$, the algorithm is as follows: If the current state of the Markov chain is z, a new state z' is drawn from a proposal distribution with density $q(z, z')$. The new state is then accepted with probability

$$\alpha(z, z') = \min \left\{ 1, \frac{p(z') q(z', z)}{p(z) q(z, z')} \right\}, \quad (17)$$

or otherwise the old state is retained. This algorithm is usually based on a succession of random local changes, since global changes tend to have small probability of being accepted. Usually, updating is most conveniently done at one site at the time, requiring access to the full conditionals $p(z_{i,j} \mid z_{-(i,j)})$ and relying on the Markov property $p(z_{i,j} \mid z_{-(i,j)}) = p(z_{i,j} \mid z_{\partial(i,j)})$ to do efficient updating. The attractiveness of this approach comes from the fact that $p(z_{i,j} \mid z_{\partial(i,j)})$ usually has a small support, and hence can be computed efficiently.

We will follow this approach, sampling from the conditional densities given in eqs. (10) and (11). This sampling scheme is efficient for the following reasons.

- $p(\sigma_{i,j}^2 \mid \sigma_{-(i,j)}^2, r)$ is a distribution with small support: The neighborhood is of the same size as specified in the prior (6). Because of this local Markov property, sampling is eased. We again stress that *given r*, the variance field σ^2 is conditionally independent of the point spread function h.

- $p(r_{i,j} \mid r_{-(i,j)}, \sigma^2, y, h)$ is a Gaussian distribution, which is straightforward to sample from. Details of the sampling are given in the next section.

B. Sampling algorithm

We now proceed to describe some details related to the sampling scheme given in section III-A; this part may be skipped at a first reading.

B.1 Sampling the variance field

The choice of proposal distribution $q(\cdot)$ poses a small problem. Since variance is a scale parameter, we chose to propose an new state by scaling the existing state with a random number $u \sim U[1/2, 2]$. However, when working with the model, we found that the restorations often contained outlying pixel values in otherwise smooth regions. While not discussed in the original paper by [21], this is a feature of the model, as shown by [24]. The reason is as follows: Let $r_{i,j}^0$ be the true pixel value at site (i, j), and assume that $\tilde{r}_{i,j}$ is the marginal minimizer of (8). Then [24] show that the difference $\tilde{r}_{i,j} - r_{i,j}^0$ is Gaussian distributed with variance $(\frac{1}{\tau^2} \sum_{k,l} h_{k,l}^2)^{-1}$. This variance can get extremely large when the blurring is severe, or when there is much noise. Then because of the concavity of the roughness penalty $\phi(\cdot)$, outlying pixel values are allowed to occur. [21] does not encounter this problem because they use a modified Gibbs sampler supported on a interval around the current values of the state and its neighbors. The algorithm is formalized in [25], restricting the image space to those images in which no pixel values differs from its neighbors by more than a threshold γ_0. However, as noted in [24], this restriction can be removed by embedding the modified Gibbs sampler in a Hastings algorithm.

In the present context, with σ^2 taking values on the whole of $\mathbb{R}^+$, Gibbs sampling is costly to implement. Since the unsatisfactory reconstructions represent local, coordinate-wise minima, a better approach would be to design the algorithm to effectively escape such local minima. We have chosen a new proposal in which, letting $\tilde{\sigma}_{i,j}^2$ be the median of the neighbors of the current state, a new state is proposed by scaling $\tilde{\sigma}_{i,j}^2$ with a random number $u \sim U[1/2, 2]$. Thus if the current state is an outlying pixel value, large jumps can be made. The choice of the median is important; choosing the mean could lead to undesired smoothing at region boundaries.

The solution suggested in [24], where the parameter β is increased to enforce a higher degree of smoothness, would also partly eliminate the problem, but at the cost of slower convergence.

B.2 Sampling the reflectance field

As customary when sampling from Gaussian distributions, we choose a Gaussian distribution centered at the current state as a proposal distribution for $r_{i,j}$. Following [26] , the standard deviation of the proposal distribution is scaled so that the acceptance ratio (17) is about 0.25. This means that we accept about one out of four proposed moves.

C. Diagnosing convergence

When designing algorithms, assessing the convergence of the Markov chain is an important issue. Quantitative convergence rates are in most cases hard to find, but there exists a number of convergence diagnostics [27], [28], [29] which can be used as a guide to the convergence of the chain. In this paper we have assessed the rate of convergence by monitoring the functional

$$f(\sigma^2) = \beta \sum_{i,j \sim k,l} w_{ij,kl} \phi(\ln \sigma_{i,j}^2 / \sigma_{k,l}^2), \tag{18}$$

using as burn-in the approximate number of iterations needed until all appear to have converged. To further verify that the algorithm produces consistent results, we have started sampling with over-dispersed initial states. Figure 2 shows the result from such an experiment. We start two different runs in initial σ^2-fields with constant levels of 10 an 0.1, respectively. Figures 2(a) and 2(b) shows the trace of the runs at two different locations, while 2(c) show traces for the functional f. Clearly, this is a subjective method, but the plots do at least indicate that the estimates produced by the algorithm are independent of the initial values used.

D. Parameter selection

The free parameters are $\beta, \delta, \epsilon, \gamma$ in the prior, the noise variance τ^2 and the point spread function h. As explained in Section II-B we use $\gamma = 1$ in all the experiments. For a discussion of the effects of changing γ, see e.g. [15], [21] . The parameter ϵ in the inverse Gamma distribution is also regarded as fixed; we use the value 0.00001 in the experiments shown, but changes of several orders in magnitude do not seem to have any significant effect on the results. In a considerably simpler setting than ours, [21] and [24] find conditions on β and δ so that the estimators satisfy certain desirable properties. In our case this approach does not seem feasible, and so β and δ is chosen by hand. As pointed out in [21] the value of δ should be chosen approximately the size of the discontinuities considered important, and should thus be set in each specific application. We use $\delta = 2$ in our experiments, but the results seems stable over a wide range of deltas, as was shown quantitatively in [24] . The last prior parameter, β, is chosen by trial and error, and is set to $\beta = 4$. The results do to some extent depend on the value of this parameter, as does the speed of convergence. However, it is hard to do this kind of modeling without having to choose at least one parameter. In particular, one has to chose a parameter balancing the degree of fidelity to the data to the smoothing imposed by the prior distribution.

Following [20] we use a sine modulated circular Gaussian point spread function (see [3]). The radial and lateral standard deviations were estimated from the power spectrum of the observed images, while the noise variance is estimated from the power spectrum outside the acoustic bandwidth of the observed images.
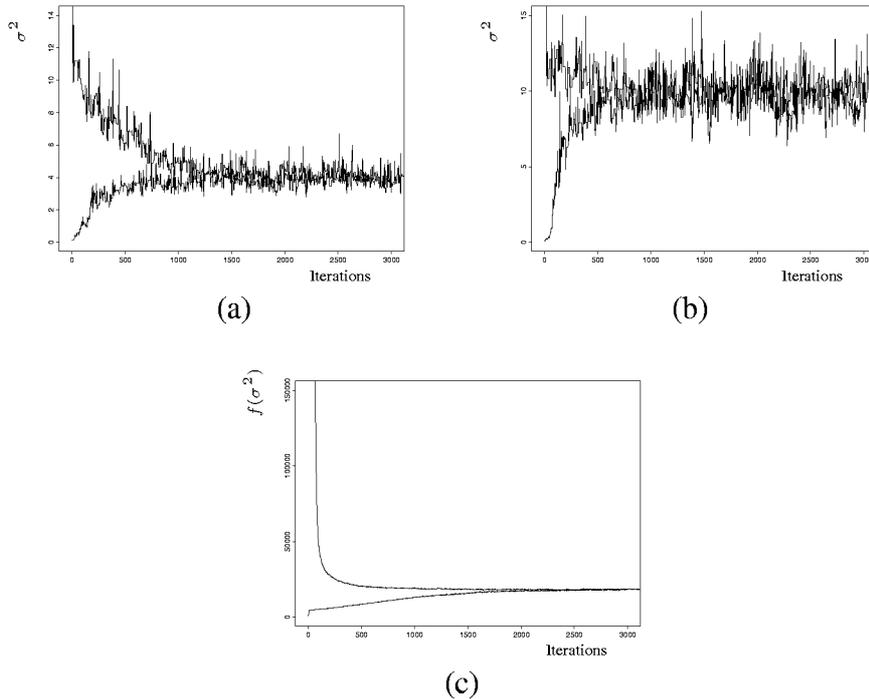
Fig. 2. Traceplots of two Markov chains started in different initial states. (a) and (b) show plots of the variance at two different sites, while (c) shows a traceplot of the functional f.

IV. EXPERIMENTS

The model was tested on a simulated region image, and on a real ultrasound image; see [3] for more experiments. We implemented the algorithm in LabVIEW, and the tests were run on a Power Macintosh 8600/25 computer, using approximately 50 minutes to process each image. A ten-fold increase in speed can be expected by implementing the routines in C. Furthermore, we claim by no means optimality of the algorithm, as our emphasis was on producing best possible restorations. The processing was done in polar coordinates, but the images were log-compressed and scan-converted into Cartesian coordinates prior to display.

A. Simulated images

To produce the simulated images we make a region image [3] and use it as a variance field for the diffuse scatterers $r_{i,j}$, which are then sampled. The data y is made by convolving the ideal image with a parametric point spread function and adding white Gaussian noise.

Figure 3a shows a 128×128 simulated image of diffuse scatterers with three region types having variances $0, 9$, and 27. Figure 3b shows the simulated data produced by blurring and adding white Gaussian noise with variance 0.5. The restoration in figure 3c is visually very close to the true image in 3a, although there is some smoothing due to the choice of estimator.

B. Real image

The ultrasound image in figure 4 was acquired using a Vingmed Sound CFM 750 scanner with a special 16 MHz A/D-converter and a 5 MHz 15 mm annular array probe. 64 beams of 128 samples were recorded within one focal zone, and the image represent an area of approximately
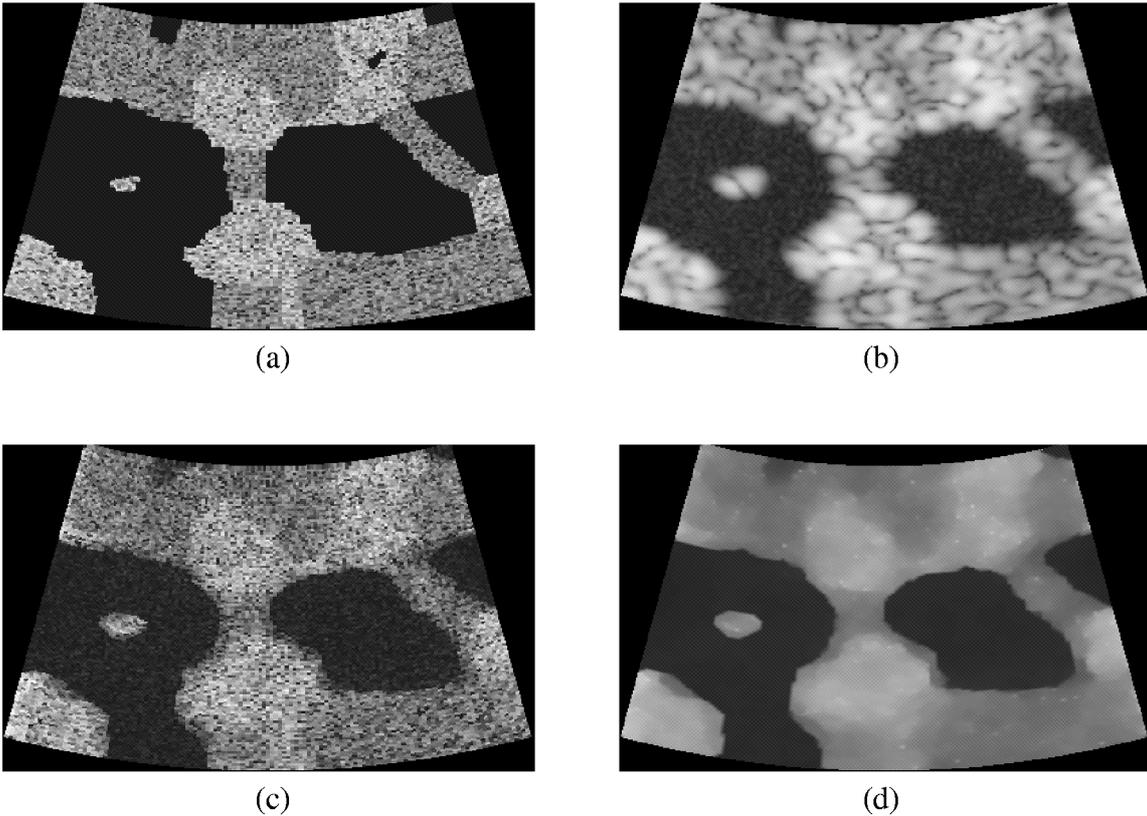
(a)

(b)

(c)

(d)

Fig. 3. Three region simulation of ultrasound image. (a) Simulated 128×128 image of diffuse scatterers r. (b) The data y produced by blurring and adding noise. (c) Estimate of reflectance field r. (d) Estimate of the variance field σ^2.

1.5×0.7 cm. Figure 4a shows a homogeneous medium of diffuse scatterers submerged in water. In 4a, the lower area representing the diffuse scatterers shows much speckle, which can also be seen in Figure 4d. The image in Fig. 4d was estimated using the Wiener filter $\frac{H^*(u,v)}{|H(u,v)|^2+k}$ in the frequency domain, with the regularization factor k tuned for optimal visual result. In our realization of r (Figure 4b) the speckle is removed, and the image shows two relatively homogeneous regions with a smooth interface. We used a sample from the posterior distribution as a realization, since the model states that $r_{i,j} \mid \sigma^2_{i,j} \sim N(0, \sigma^2_{i,j}), \forall (i,j) \in \mathcal{I}$, and the posterior mean therefore is not meaningful as an estimator. Even though the lower region of the image should be homogeneous, the point spread function has introduced artifacts that our model cannot remove. Figure 4c shows the estimated variance field, an image which in fact summarizes all relevant information in a precise fashion. The homogeneity of the two regions, and the border between them are clearly seen, without disturbing noise and image artifacts.

V. Discussion

This paper describes a model for the scattering distribution of homogeneous tissue, and introduces smoothness constraints through the use of the model in [21] for the underlying scattering characteristics. The model is used in a new algorithm for restoration of ultrasound images.

We believe that the results demonstrate the potential usefulness of Markov random field models in restoration of medical ultrasound images. The restoration of the simulated images pro-
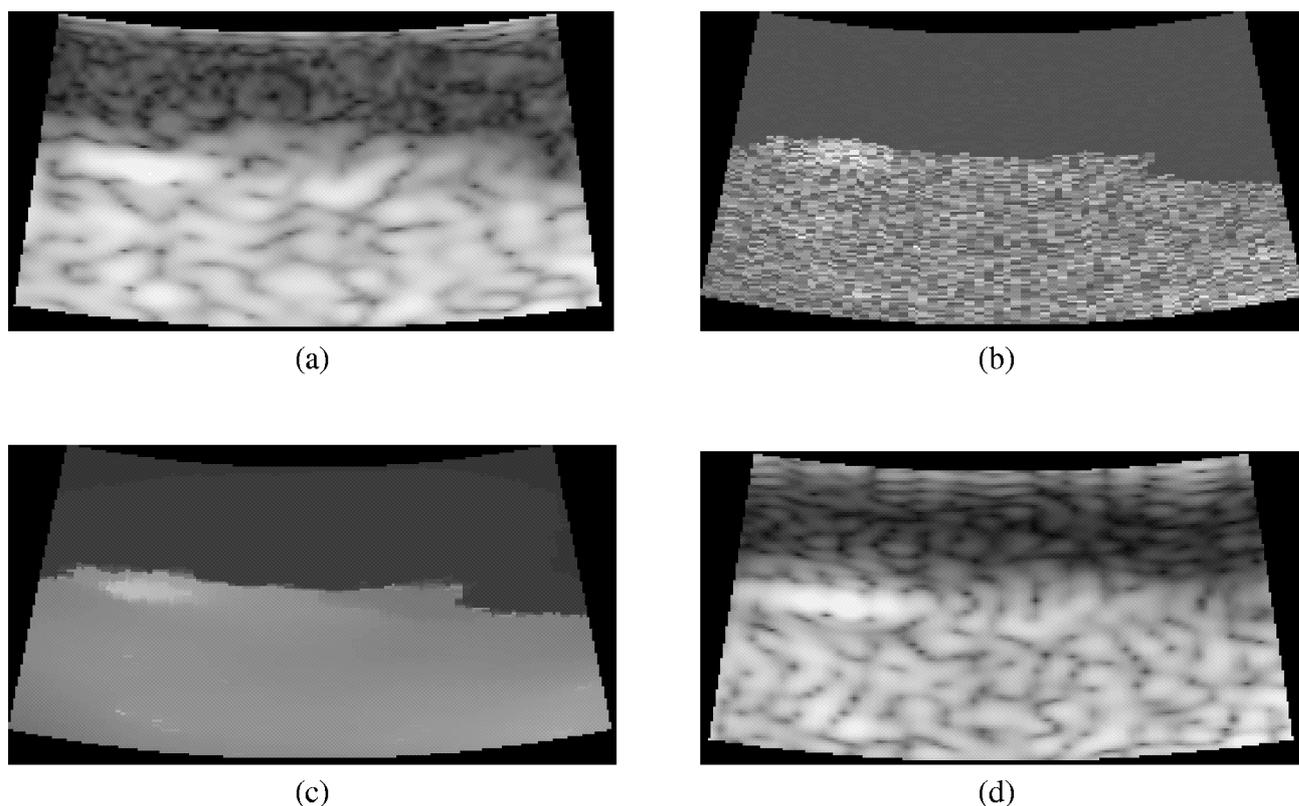
(a)

(b)

(c)

(d)

Fig. 4. Ultrasound image of a homogeneous medium submerged in water. (a) Observed 64 × 128 image y (b) A sample from the posterior distribution for r given y and σ^2. (c) Estimate of the variance field. (d) Wiener filtering of y.

duced results that are more representative of the true object than the observed ultrasound images. Also, in the case of the real image, the results obtained are significantly better than the one produced by Wiener filtering. In both cases the speckle patterns are efficiently removed, while at the same keeping important details and not introducing artificial structures.

Our prior model favors homogeneous areas with sharp transitions between regions, which is a reasonable characteristic for a model representing ultrasound images since tissues are mostly homogeneous with sharp transitions between different tissue types. This characteristic of the prior distribution, which can be understood by studying the plot in figure 1, is a great advantage to the restoration procedure. A disadvantage is that we may loose smaller structures. However, these structures are to a certain degree, depending on their size and the size of the degrading point spread function size, lost anyway in the blurring process.

The removal of speckle is efficient with our method, and we believe this results in visually better images. Speckle removal is important in 3D visualization systems, where the data in a 3D volume is mapped to a 2D image through a process called rendering. In order to make the 2D image understandable, it is important that only significant structure information participates in the rendering process, and that irrelevant information (i.e. noise and speckle) is excluded. It is important to notice that if a certain speckle pattern is interpreted (by the model/algorithms) as a real structure when it in fact is a homogeneous area, this can lead to misinterpretation by the end observer. This is a crucial point if the clinician wishes to use this restoration method to

study and analyze details from a volume scan. In other words, if this method proves useful for real ultrasound data, a surgeon, e.g. during removal of a tumor in the brain, could perform a 3D scan and then use the method to restore the reflection image of a smaller part of the volume to investigate the remains of the tumor. It would then be crucial that the method does not produce false structures due to the speckle pattern.

Further problems have to be solved. As mentioned in Section I our model has limitations, and it is important to consider how to deal with specular structure and quasiperiodic components. In [20] the specular structures were incorporated in the Markov field model, without significantly improving the results. Because of the continuous variance field, we believe that our model has the flexibility to at least partly account for the effects of specular structures. The presence of quasiperiodic scatterers will cause problems, but these occur only in some applications. Secondly, it is well known that point spread functions in ultrasound imaging do not have Gaussian shapes, and that they vary both radially and laterally [5], [30] . As we show in [3] , poor estimates of the point spread function effects the quality of the restorations. Thus effort should be put into the estimation of the point spread function in order to get truly good restorations, especially in larger images. Some attempts on estimation of 2D point spread functions exists [1], [5] , but this is a field that needs further study. From a practical point of view, the efficiency of the algorithm should be improved. Most medical applications require real-time processing or at least on-line response times, and thus further optimization of the algorithm should be investigated. Note that when used on a sequence of images, the proceeding estimates could be used as initialization for the current frame, hopefully leading to a considerable speedup. Also, when doing real-time display of sequences, samples from the posterior distribution could be used instead of the sample means (16).

ACKNOWLEDGMENTS

Appendix

I. Alternative model formulation

We derive the posterior distribution $p(\sigma^2 \mid y, h)$, showing that our model can indeed be formulated in terms of σ^2 only. Using (3) we need the likelihood, which, using (1) and (2), is found by integrating out r:

$$p(y \mid \sigma^2) = \int p(y \mid r) p(r \mid \sigma^2) \mathrm{d}r. \tag{19}$$

For clarity of exposition we define matrices H and Λ as follows. Λ is a $n \times n$ diagonal matrix with entries $\sigma_{1,1}^2, \sigma_{1,2}^2, \ldots, \sigma_{n_r,n_l}^2$ on the diagonal, and H is a $n \times n$ band matrix representing the point spread function h. Letting $\Sigma_s = \tau^2 I$ we get

$$
\begin{aligned}
p(y \mid \sigma^2) &= \int p(y \mid r) p(r \mid \sigma^2) \mathrm{d}r \\
&\propto \int \exp\left\{ -\frac{1}{2}\left((y - Hr)^T \Sigma_s^{-1} (y - Hr) - r^T \Sigma^{-1} r \right) \right\} \mathrm{d}r \\
&= \exp\left\{ -\frac{1}{2} y^T \Sigma_s^{-1} y \right\} \int \exp\left\{ -r^T \left(H^T \Sigma_s^{-1} H + \Sigma^{-1} \right) r \right. \\
&\qquad \left. + 2 r^T H^T \Sigma_s^{-1} y \right\} \mathrm{d}r \\
&= \exp\left\{ -\frac{1}{2} y^T \Sigma_s^{-1} y \right. \\
&\qquad \left. + \frac{1}{2} y^T \Sigma_s^{-1} H \left(H^T \Sigma_s^{-1} H + \sigma^{-1} \right)^{-1} H^T \Sigma_s^{-1} y \right\} \\
&= \exp\left\{ -\frac{1}{2} y^T \left(\Sigma_s^{-1} - \Sigma_s^{-1} H \left(H^T \Sigma_s^{-1} H + \Sigma^{-1} \right)^{-1} H^T \Sigma_s^{-1} \right) y \right\} \\
&= \exp\left\{ -\frac{1}{2} y^T \left(\Sigma_s + H\Sigma H^T \right)^{-1} y \right\},
\end{aligned}
$$

where the last line comes from the identity

$$
\begin{aligned}
A^{-1} &- A^{-1} B \left(C^{-1} + D A^{-1} B \right)^{-1} D A^{-1} \\
&= (A + BCD)^{-1};
\end{aligned}
$$

see e.g. [31] p. 459.

The posterior density is then found by combining the prior density (6) with the density of the likelihood model (8):

$$
\begin{aligned}
p(\sigma^2 \mid y) &\propto p(y \mid \sigma^2) p(\sigma^2) \\
&\propto \frac{1}{|H\Sigma H^T + \tau^2 I|^{1/2}} \exp\left\{ -\frac{1}{2} y^T \left(H\Sigma H^T + \tau^2 I \right)^{-1} y - \Phi(\sigma^2) - \Psi(\sigma^2) \right\}. \tag{20}
\end{aligned}
$$

However, note that the matrix $Q = \left(H\Sigma H^T \right)^{-1}$ in general is full; this can be seen heuristically by expanding Q as a power series. Letting $q_{ij,kl}$ be the elements of Q, it is a well known fact

[32], [33] that

$$E\left(Y_{i,j} \mid y_{-(i,j)}, \sigma^2\right) = -\sum_{k,l \neq i,j} \frac{q_{ij,kl}}{q_{ij,ij}} y_{k,l}, \tag{21}$$

implying that the neighborhood of site (i, j) consists of all sites (k, l) such that $q_{ij,kl} \neq 0$. Hence since Q is full the local Markov property fails to hold since the neighborhood of each site is simply the whole image. This means that conditional on the data y, any two sites in the variance field σ^2 are directly dependent, which is undesirable for at least two reasons. First, updating is time-consuming because for each site that is updated, the full density has to be computed in (17). Secondly, because of the global dependency, the chain will move slowly in the sample space, resulting in slow convergence of the Markov chain. Thus this formulation of the model, while valid and instructive, leads to computational problems. However, this deficiency can be amended by considering the augmented field (σ^2, r), as shown in section II-C.

REFERENCES

[1] T. Taxt, "Restoration of medical ultrasound images using 2-dimensional homomorphic deconvolution," *IEEE Trans. on Ultrason. Ferroelec. Frec. Contr.*, vol. 42, no. 4, pp. 543–554, June 1995.

[2] D. Iraca, L. Landini, and L. Verrazzani, "Power spectrum equalization for ultrasonic image restoration," *IEEE Trans. on Ultrason. Ferroelec. Frec. Contr.*, vol. 36, no. 2, pp. 216–222, march 1989.

[3] T. Langø, T. Lie, O. Husby, and J. Hokland, "Bayesian 2D deconvolution: Effect of using spatially invariant ultrasound point spread functions," *IEEE Trans. on Ultrason. Ferroelec. Frec. Contr.*, 2000, Following article.

[4] T. Taxt and G. Frolova, "Noise robust one-dimensional blind deconvolution of medical ultrasound images," *IEEE Trans. on Ultrason. Ferroelec. Frec. Contr.*, vol. 46, no. 2, pp. 291–299, march 1999.

[5] J.A. Jensen and S. Leeman, "Nonparametric estimation of ultrasound pulses," *IEEE Transactions on Biomedical Engineering*, vol. 41, no. 10, pp. 929–936, 1994.

[6] A.K. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, 1989.

[7] X. Zong, A.F. Laine, and E.A. Geiser, "Speckle reduction and contrast enhencement of echocardiograms via multiscale nonlinear processing," *IEEE Transaction on Medical Imaging*, vol. 17, no. 4, pp. 532–540, August 1998.

[8] G. Winkler, *Image Analysis, Random Fields and Dynamic Monte Carlo Methods*, Springer, Berlin, 1995.

[9] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 6, pp. 721–741, 1984.

[10] J. Besag, P. Green, D. Higdon, and K. Mengersen, "Bayesian computation and stochastic systems (with discussion)," *Statistical Science*, vol. 10, no. 1, pp. 3–66, 1995.

[11] J.W. Goodman, "Statistical properties of laser speckle patterns," in *Laser Speckle and Related Phenomena*, J.C. Dainty, Ed. Springer Verlag, Berlin, 1975.

[12] R.F. Wagner, M.F. Insana, and D.G. Brown, "Statistical properties of radio-frequency and envelope-detected signals with application to medical ultrasound," *J. Opt. Soc. Am, A*, vol. 4, no. 5, pp. 910–922, 1987.

[13] F.S. Cohen, G. Georgiou, and E.J. Halpern, "WOLD decomposition of the backscatter echo in ultrasound images of soft tissue organs," *IEEE Trans. on Ultrason. Ferroelec. Frec. Contr.*, vol. 44, no. 2, pp. 460–472, March 1997.

[14] G. Georgiou and F.S. Cohen, "Statistical characterization of diffuse scattering in ultrasound images," *IEEE Trans. on Ultrason. Ferroelec. Frec. Contr.*, vol. 45, no. 1, pp. 57–64, January 1998.

[15] S. Geman and D.E. McClure, "Statistical methods for tomographic image reconstruction," in *Proc. 46th Sess. Int. Stat. Inst. Bulletin ISI*, 1987, vol. 52.

[16] W. K. Hastings, "Monte Carlo simulation methods using Markov chains and their applications," *Biometrika*, vol. 57, pp. 97–109, 1970.

[17] P. J. Green, "Reversible jump MCMC computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.

[18] J. Besag and P. Green, "Spatial statistics and Bayesian computation," *Journal of the Royal Statistical Society, Series B*, vol. 16, pp. 395–407, 1993.

[19] D.M. Higdon, "Auxiliary variable methods for Markov chain Monte Carlo with applications," *Journal of the American Statistical Association*, vol. 93, no. 442, pp. 585–595, June 1998.

[20] J.H. Hokland and P.A. Kelly, "Markov models of specular and diffuse scattering in restoration of medical ultrasound images," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 43, no. 4, pp. 660–669, July 1996.

[21] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 3, pp. 367–383, 3 1992.

[22] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Transaction on Image Processing*, vol. 6, no. 2, pp. 298–311, february 1997.

[23] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, *Markov Chain Monte Carlo in Practice*, Chapman & Hall, 1996.

[24] M. Hurn and C. Jennison, "An extension of Geman and Reynolds' approach to constrained restoration and the recovery of discontinuities," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 6, pp. 657–662, June 1996.

[25] D. Geman, G. Reynolds, and C. Yang, "Stochastic algorithms for restricted image spaces and experiments in deblurring," in *Markov Random Fields: Theory and Application*, R. Chellappa and A. Jain, Eds. Academic Press, New York, 1993.

[26] G.O. Roberts, A. Gelman, and W.R. Gilks, "Weak convergence and optimal scaling of random walk metropolis algorithms," *Ann. Appl. Prob.*, vol. 7, pp. 110–120, 1997.

[27] S. P. Brooks and G. O. Roberts, "Assessing convergence of Markov chain Monte Carlo algorithms," *Statistics and Computing*, vol. 8, no. 4, pp. 319–335, 1998.

[28] C. P. Robert, "Convergence control methods for MCMC algorithms," *Statistical Science*, vol. 10, no. 3, pp. 231–253, 1995.

[29] K. L. Mengersen, C.P. Robert, and C. Guihenneuc-Jouyaux, "MCMC convergence diagnosics: a 'reviewwww'," Tech. Rep., CREST, INSEE, Paris, May 1998.

[30] L.A. Ødegård, *Phase aberration correction in medical ultrasound imaging*, Ph.D. thesis, The Norwegian Institute of Technology, 1995.

[31] K.V. Mardia, J.T. Kent, and J.M. Bibby, *Multivariate Analysis*, Academic Press, 1979.

[32] J. Besag, "Spatial interaction and the statistical analysis of lattice systems," *Journal of the Royal Statistical Society, Series B*, vol. 36, pp. 192–236, 1974, with discussion.

[33] J. Besag and C. Kooperberg, "On conditional and intrinsic autoregressions," *Biometrika*, vol. 82, no. 4, pp. 733–746, december 1995.

Paper F

Bayesian 2-D Deconvolution: Effect of using Spatially Invariant Ultrasound Point Spread Functions

Thomas Langø [§], Torgrim Lie [§], Oddvar Husby and Jørn Hokland

ABSTRACT

Observed ultrasound images are degraded representations of the true tissue reflectance. The specular reflections at boundaries between regions of different tissue types are blurred, and the diffuse scattering within homogenous regions causes speckle because of the oscillating nature of the transmitted pulse. To reduce both blur and speckle, we have developed algorithms for restoration of simulated and real ultrasound images based on Markov random field models and Bayesian statistical methods. The algorithm is summarized here, while a more detailed description can be found in our companion paper [1]. Because the point spread function (*psf*) is unknown, we investigate the effects of using incorrect frequencies and sizes for the model *psf* during the restoration process. First, we degrade the images either with a known simulated *psf* or a measured *psf*. Then, we use different *psf* shapes during restoration to study the robustness of the method. We found that small variations in the parameters characterizing the *psf*, less than ±25% change in frequency, width, or length, still yielded satisfactory results. When altering the *psf* more than this, the restorations were not acceptable. The restorations were particularly sensitive to large increases in the restoring *psf* frequency. Thus, 2-D Bayesian restoration using a fixed *psf* may yield acceptable results as long as the true variant *psfs* have not varied too much during imaging.

T. Langø and T. Lie are with SINTEF Unimed, Ultrasound, 7465 Trondheim, Norway (e-mail: thomasl@unimed.sintef.no).

O. Husby is with the Department of Mathematical Sciences, NTNU, 7491 Trondheim, Norway.

J. Hokland is with the Department of Computer and Information Sciences, NTNU, 7491 Trondheim, Norway.

§ T. Langø and T. Lie contributed equally to this article.

I. INTRODUCTION

Ultrasound images often appear noisy and inhomogeneous. The main problems in ultrasonic imaging is the strong, tissue-dependent dampening of the transmitted ultrasonic pulse (the point spread function or *psf*), distortion of the pulse as it propagates through tissues with local variations in sound velocity, multiple reflections or false echoes, and reverberations, as well as the high angle dependency in reflections from layered structures or plane surfaces. The result is that the received echo intensities are strongly dependent on the acquisition conditions, and similar tissues may appear very different in various parts of the same image. Because of the coherent detection in ultrasound imaging, constructive and destructive interference between reflections from small, closely localized objects (scatterers) give rise to speckle, which modulates the gray level images, making the intensity vary apparently randomly with position. This gives the ultrasound images the well-known grainy appearance, complicating visual interpretation and volume visualization of ultrasound data [2]. The spatial variation of the speckle is roughly determined by the frequency and size of the transmitted pulse. Therefore, with increasing frequency and, thus, in general shorter pulses, the speckle becomes finer grained. Furthermore, differences in speed of sound throughout tissue result in a widening of the *psf* [3], meaning that the size of the *psf* increases as a function of depth in the image, and it may become severely degraded as it traverses the tissue. The degree of degrading depends on the type of tissue, speed of sound, smooth or rugged tissue boundaries etc., in addition to frequency and shape of the transmitted pulse.

Medical ultrasound images are approximately modeled as a convolution between the tissue reflectance and the imaging system *psf* [4-8]. The *psf* is spatially variable, and, particularly, the limited bandwidth of the *psf* and the additive noise make perfect deconvolution impossible. Wiener filtering [9, 10] is widely used for deconvolution in ultrasound scanners to provide partial restoration of the reflectance. However, the restored images still suffer from significant blur and speckle, which adversely affects our ability to interpret the images.

To improve restorations, we have developed algorithms for restoration of speckle-laden ultrasound images based on Markov random field models [11] and Bayesian statistical methods [12]. The Bayesian framework differs from most conventional methods of image restoration in at least two respects. First, it requires information about the probabilistic mechanism that degrades the original scene, the *psf* and noise in the case of ultrasound images. Second, weaker constraints replace rigid ones. With this methodology, we can model randomness and uncertainty inherent in the system. This means that instead of classifying the data, the current status is weighted by an acceptance function, quantifying the degree to which it is correct or not. The choice of the prior probability distribution, which characterizes the spatial smoothness constraint, is problem dependent and one of the main challenges in Bayesian image analysis and restoration.

Our implementation is based on a model for diffuse scattering, published by Wagner *et al*. [13]. In their paper, the authors assume that the concentration of scatterers is large relative to the wavelength of the pulse, thereby modeling a point reflector within one resolution cell as a Gaussian random variable with zero mean and variance depending on the tissue type within the cell. The reflectance at different resolution cells is assumed to be non-correlated.

In this paper, our main focus is to determine to what extent it is necessary to have a good approximation of the stationary *psf* in order for restoration to yield satisfactory results. Such knowledge is desirable because the *psf* is spatially variable and unknown in ultrasound imaging. Satisfactory methods for estimating non-stationary *psfs* are difficult to find, although some publications claim that this is possible, and attempts have been made [4, 7, 8, 14-19]. We believe it would be of great importance and advantage to have a method of restoration that works well even when the *psf* is poorly estimated and spatially invariant. We use simulated ultrasound images to compare restorations with true images. The true images are first degraded with a known model *psf* or a measured *psf*, in addition to additive Gaussian noise. A related approach to this can be found in [20], where a discrete Markov model for the tissue region field was used, as opposed to our model with the employment of a continuous

tissue region field, which we believe is a better descriptor of acoustic properties of human tissue. The use of a discrete Markov field in [20] seemed to make their estimates unstable.

II. SIMULATED ULTRASOUND IMAGES

To test our restoration algorithms on ultrasound images, we model an observed ultrasound image. Assuming a spatially invariant *psf*, the formation of an ultrasound image can be modeled as a 2-D convolution, i.e., blurring [4-6], between the *psf* and the reflectance image and addition of white Gaussian noise. This process can be expressed mathematically as

$$y_{i,j} = \sum_{k,l} h_{k,l} r_{i-k,j-l} + \eta_{i,j} = \{h * r\}_{i,j} + \eta_{i,j} \tag{1}$$

where $y_{i,j}$ is the observed image value at radial pixel position i and lateral pixel position j on a rectangular polar-coordinate lattice, h represents the *psf*, η is the additive Gaussian noise, and $r_{i,j}$ is the reflectance value. We assume that the noise is independent and normally distributed with variance equal to τ^2, that is

$$\eta_{i,j} \sim N(0, \tau^2)$$

We use a cosine modulated Gaussian function as our *psf* (Fig. 1a), mathematically given by

$$h_{k,l}(f, a, b) = \cos(\pi \cdot f \cdot l) \cdot e^{-\frac{a \cdot k^2 + b \cdot l^2}{2}} \tag{2}$$
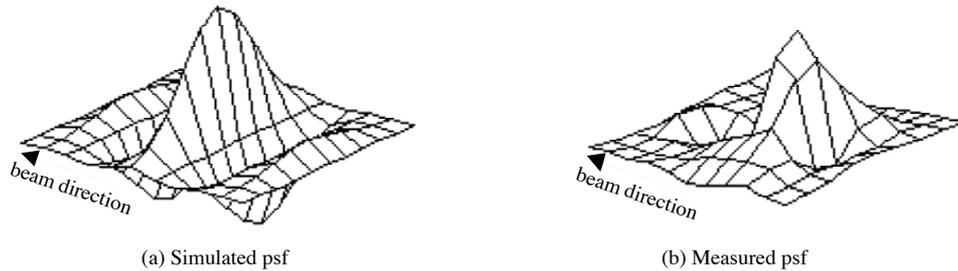

(a) Simulated psf   (b) Measured psf

Fig. 1. a) Plot of the *psf* given in Eq. 2; b) plot of a measured *psf*. The measured *psf* was obtained by recording the back-scattered signal from a tiny air bubble in a water bath with a 2.5-MHz phased array probe.

The number of oscillations and size of the *psf* in Eq. 2 may be altered by changing the parameters f, a, and b, which are important to resemble real *psfs* that vary with depth, frequency, and type of tissue it has traversed. In Eq. 2, k and l are indices. An example of a measured *psf* is shown in Fig. 1b. This *psf* was obtained by measuring the back-scattered echo from a tiny air bubble in a water bath. The *psf* in Eq. 2 is normalized with respect to the total energy, such that $\Sigma h^2 = 1$. Because the *psf* has both positive and negative values, some values may be close to zero, the number being dependent on the sampling frequency and the *psf* itself. By changing the structure of the *psf* from a 2-D array to a 1-D structure (Fig. 2), where the *psf* values and their positions are stored at each index, we can remove the zeros from the calculations and, thereby, reduce the amount of calculation. This new structure is better suited to represent an irregularly shaped, non-rectangular *psf*, because it is easier to remove superfluous elements. Furthermore, we have implemented the option of setting a rejection level, *i.e.,* removing a certain amount of the smallest values from the *psf* matrix to further reduce the amount of calculation.
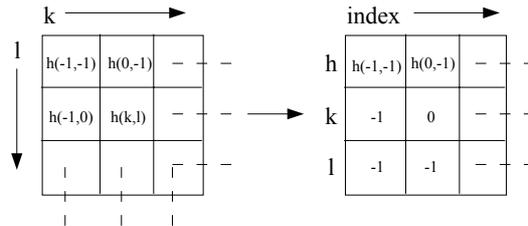
Fig. 2. Data structure for the representation of the *psf*.

We generate a tissue region field image with three different classes: variance levels of 0, 9, and 27. The reason for including the value 0 for the variance is the fact that ultrasound images of water give zero reflection and, thereby, black homogenous regions. The region field image is obtained by Gibbs sampling with inversion [21] from a discrete distribution imposing a strong smoothness constraint of a first-order neighborhood system [20]. The resultant tissue region field is shown in Fig. 3a (note that a hand-drawn image could also have been used). This image has several features that make it suited for the use here; it has both large and small homogenous regions and small bright structures in large dark regions and vice versa. This makes it possible to study various effects in the restoration process.

We use the region image as a variance field to construct the reflection image, which describes the Gaussian distribution of diffuse scatterers in the tissue, *i.e.*, local differences in acoustic impedance. Assuming that the scatterers are non-correlated, the probability distribution for a pixel in the reflectance image is only dependent on the corresponding pixel from the variance image (independent of the rest of the variance image and the reflectance image). The reflectance image in Fig. 3b is obtained by using the Box-Müller method [21] for sampling.

After convolution with the *psf*, addition of white Gaussian noise using the Box-Müller method, log-compression, and scan-conversion into Cartesian coordinates, the reflectance image is transformed into a simulated ultrasound image. Note that the images in Fig. 3a and 3b were also scan-converted. We simulated two different ultrasound images: one degraded with the simulated *psf* and one degraded with the measured *psf*, shown in Fig. 3c and Fig. 3d respectively.
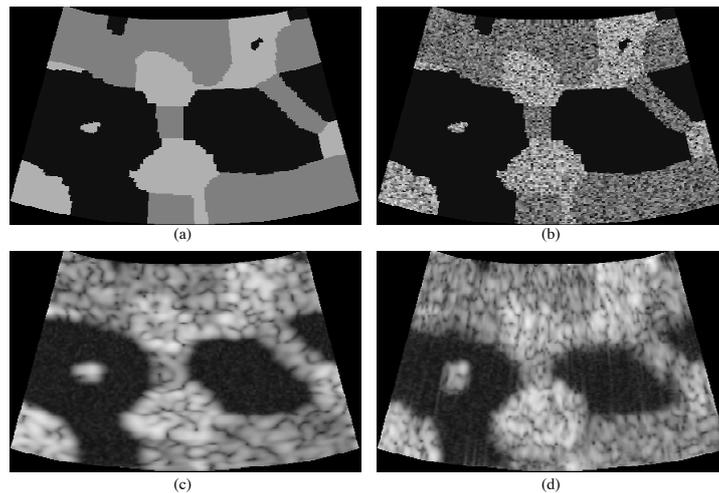

Fig. 3. Three-region simulation of ultrasound images. a) True variance image or region field σ^2; b) simulated distribution of diffuse scatterers in tissue, the reflectance image r; c) simulated ultrasound image from convolution of the reflectance image and the simulated *psf* (Eq. 2, Fig. 1a) and addition of white Gaussian noise ($\tau^2 = 0.3$); d) simulated ultrasound image from convolution of the reflectance image (b) and the measured *psf* (Fig. 1b) and addition of white Gaussian noise ($\tau^2 = 0.3$). The images were log-compressed and scan-converted prior to display to resemble real ultrasound images. The rejection level before degrading with the *psfs* was set to 0.01, meaning that values below 1% of the maximum value in the *psf* matrices were rejected from the calculations. The slight white striped pattern appearing in image d is a consequence of the measured *psf* being discontinuous at the edges (see Fig. 1b).

III. Model formulation, Image sampling, and implementation

For a detailed explanation of the theory relevant to the implementation of our method, including deductions of the distributions given subsequently, we refer to [1]. To implement the method, we need to know the posterior distributions of the variance image σ^2 and the reflectance image r, given the observed ultrasound image y. The posterior joint distribution of σ^2 and r can be found from statistical considerations and the use of Bayes' rule [1, 11]. However, the crucial part of our model is the formulation of the prior distribution for the variance field σ^2. The diffuse scattering model incorporates a variance depending on the acoustical properties of the tissue, and, thus, variances tend to be approximately constant in homogenous tissue regions, while sharp changes in variance occur at interfaces between different tissues. In addition, the level of the reflected intensity varies from one type of tissue to the next. To meet these demands (see article [1]), we used a distribution for the whole image given by

$$p\left(\sigma^2\right) \propto \exp\left\{-\beta \sum_{i,j\sim k,l} w_{ij,kl} \frac{\left|\ln\sigma_{i,j} - \ln\sigma_{k,l}\right|}{\delta + \left|\ln\sigma_{i,j} - \ln\sigma_{k,l}\right|} - \sum_{i,j}\left[\frac{\varepsilon}{\sigma_{i,j}^2} + (\varepsilon + 1)\ln\sigma_{i,j}^2\right]\right\} \quad (3)$$

where β, δ, and ε are positive parameters, $w_{ij,kl}$ is a weight factor depending on the neighborhood system [1], and the notation $i,j\sim k,l$ designates a summation over the neighbors to i,j. These parameters can be varied depending on the specifics of the prior knowledge to be imposed. β characterizes the overall strength of the prior model relative to the observed image and r field, and δ scales the difference between neighbor pixel values.

The reflectance field (Fig. 3b) was modeled as a non-correlated Gaussian field with zero mean and variance given as a function of the region field, which models tissue type. Hence, the probability distribution for a pixel is given by

$$p\left(r_{i,j}|\sigma_{i,j}^2\right) = \frac{1}{\sqrt{2\pi\sigma_{i,j}^2}}\exp\left\{-\frac{r_{i,j}^2}{2\sigma_{i,j}^2}\right\} \propto \exp\left\{-\frac{r_{i,j}^2}{2\sigma_{i,j}^2} - \ln\sigma_{i,j}\right\} \quad (4)$$

Eq. 1 can be employed to obtain the distribution for the data y, given the *psf* and the reflectance image [1]. To apply Gibbs sampling, we need the conditional distributions of $r_{i,j}$ and $\sigma_{i,j}^2$ at single pixel positions i,j. The results are (see [1])

$$p\left(\sigma_{i,j}^2|\sigma_{\Gamma}^2,r,y,h\right) \propto \exp\left\{-\beta \sum_{k,l\sim m,n} w_{kl,mn} \frac{\left|\ln\sigma_{k,l} - \ln\sigma_{m,n}\right|}{\delta + \left|\ln\sigma_{k,l} - \ln\sigma_{m,n}\right|} - \frac{r_{i,j}^2 + 2\varepsilon}{2\sigma_{i,j}^2} - (2\varepsilon + 3)\cdot\ln\sigma_{i,j}\right\}$$
$$(5)$$

$$p\left(r_{i,j}|r_{\Gamma},\sigma^2,y,h\right) \sim N\left(\mu_{i,j},\kappa_{i,j}^2\right) \quad (6)$$

where the index Γ represents all pixels except i,j, and $r_{i,j}$ has a mean given by

$$\mu_{i,j} = \frac{\dfrac{1}{\tau^2}\sum_{m,n}\left(h_{m-i,n-j}\cdot\Upsilon_{m,n}\right)}{\dfrac{1}{\tau^2}\sum_{m,n}h_{m-i,n-j}^2 + \dfrac{1}{\sigma_{i,j}^2}} \quad (7)$$

where

$$\Upsilon_{m,n} = y_{m,n} - \sum_{k \neq m-i, l \neq n-j} h_{k,l} r_{m-k,n-l} \tag{8}$$

and variance

$$\kappa_{i,j}^2 = \left[\frac{1}{\tau^2} \sum_{m,n} h_{m-i,n-j}^2 + \frac{1}{\sigma_{i,j}^2} \right]^{-1} \tag{9}$$

where τ^2 is the variance of the added Gaussian noise (see Fig. 3). The value for the noise level was set to approximately 5% of the energy in the normalized modeled *psf* ($\tau^2 = 0.3$). There exist several Markov chain Monte Carlo methods, or updating procedures, and among the most commonly used are variations of the Metropolis [22] and Gibbs samplers [11]. We sample σ^2 and r images, alternating between σ^2 and r, using the Metropolis-Hastings algorithm [23, 24]. For a detailed description of the image sampling procedure used in our implementation, we refer to the Appendix and [1]. As explained in [1], with a certain probability (5 to 10%) we pick the median of the eight nearest neighbors instead of the current pixel. By picking candidates in this manner, we achieve that single high intensity pixels, most likely noise, in otherwise homogenous areas disappear. Moreover, we pick the pixel to be updated randomly from the image.

IV. RESULTS

We present the results as sample number 500 from the estimations of the reflectance image. The images were log-compressed and scan-converted to resemble real ultrasound images. The results are shown in Fig. 5 and Fig. 6. Moreover, we have calculated the pixel classification errors for the variance image restorations compared with the region image in Fig. 3a, also for restoration sample number 500. The results are shown in Fig. 7. We have shown the true variance image in Fig. 4, marked with numbers referring to interesting structures in the image. For the results in Fig. 5, Fig. 6, and Fig. 7 the prior model parameters were $\beta = 4.5$ and $\delta = 1$. When changing one of the parameters in the *psf* (frequency, lateral length, or radial width), the others are kept equal to the degrading *psf*.
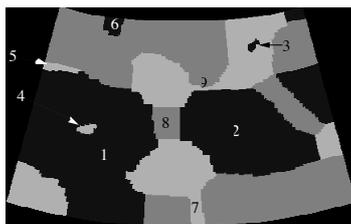

Fig. 4. The true tissue region variance field image from Fig. 3a. The numbers from 1 through 9 are referred in the text.

A. Restoration of the reflectance image from the image degraded with a simulated psf

Fig. 5a shows the data y, *i.e.,* the ultrasound image. The estimate obtained when using the same *psf* during restoration and degradation is shown in 5b. In Fig. 5c and 6a, 5e and 6c, and 5g and 6e, we have decreased the frequency, lateral width, and radial length of the restoration *psf*, respectively. The change in parameters is smaller in Fig. 5 (±10% for the frequency, width, and length) compared with Fig. 6 (±25%).

The bright spot 4 in Fig. 4 becomes wider laterally when the restoring *psf* is increased in the lateral direction (Fig. 5f) and it becomes larger in the radial direction by using a restoring *psf* that has increased length (Fig. 5h). The spot marked 3 disappears in the convolution process

and, thus, is not recovered in the restoration. In total, it seems that a reduction of the width and length of the *psf* results in better restorations compared with a corresponding increase.

The effects of changing the frequency are subtler. Nonetheless, it can be seen from Fig. 5d and Fig. 6b that increasing the frequency results in overall brighter images with less contrast. In addition, spot 4 in Fig. 4 becomes larger, especially when increasing the frequency by as much as 25%. An increase in frequency also causes smearing or blurring of the edges in the restorations, which can be seen in Fig. 5d and Fig. 6b. Decreasing the frequency does not seem to cause edge errors to the same degree.
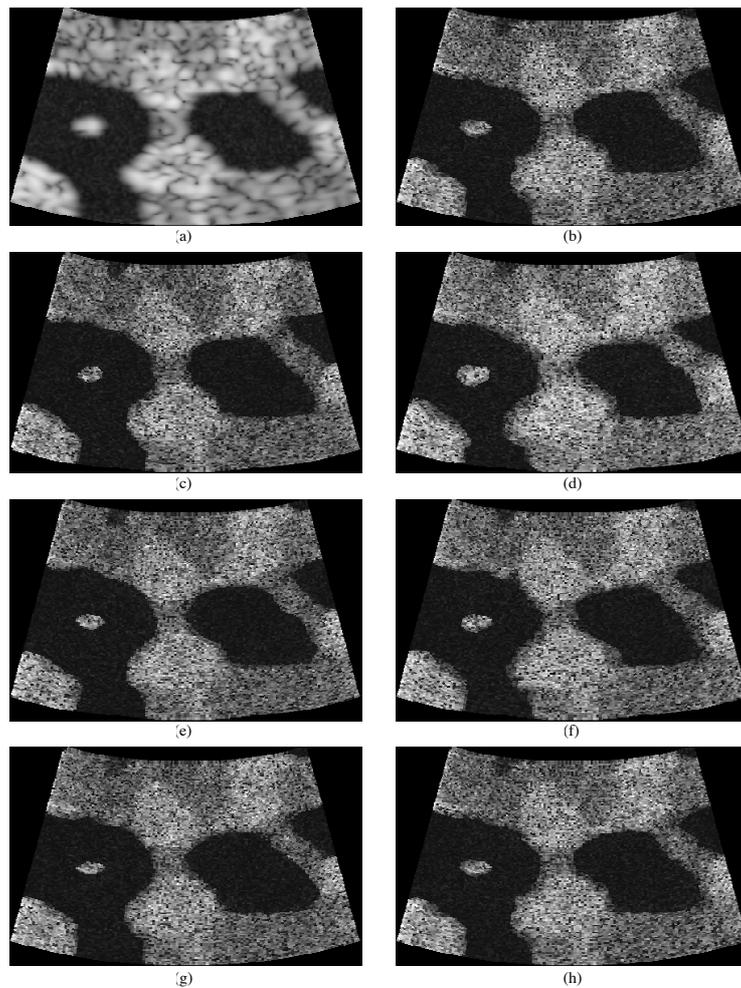

Fig. 5. Results from the restoration of the reflectance image. a) Simulated ultrasound image, b) result when using the same *psf* during restoration as degradation, c) result when decreasing the restoring *psf* frequency with 10%, d) result when increasing the restoring *psf* frequency by 10%, e) result when decreasing the lateral width of the restoring *psf* by 10%. (f) Result when increasing the lateral width of the restoring *psf* by 10%, g) result when decreasing the radial length of the restoring *psf* by 10%, and h) result when increasing the radial length of the restoring *psf* by 10%.
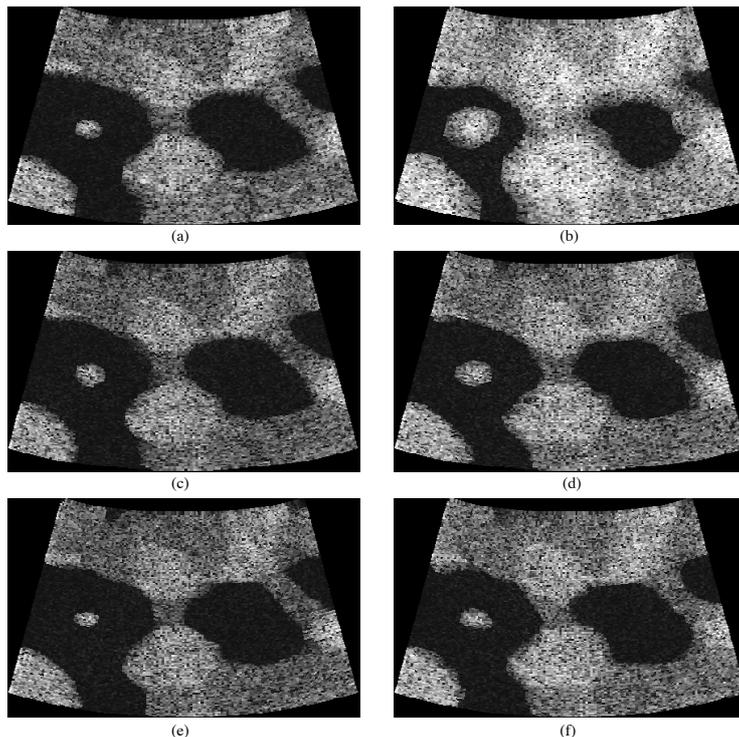
Fig. 6. Results from the restoration of the reflectance image. See Fig. 5a and 5b for the simulated ultrasound image and restoration using the same *psf* as for degradation. a) Result when decreasing the restoring *psf* frequency by 25%, b) result when increasing the restoring *psf* frequency by 25%, c) result when decreasing the lateral width of the restoring *psf* by 25%, d) result when increasing the lateral width of the restoring *psf* by 25%, e) result when decreasing the radial length of the restoring *psf* by 25%, and f) result when increasing the radial length of the restoring *psf* by 25%.

B. *Estimation of the variance image from the image degraded with a simulated psf*

Fig. 7 shows the percent error pixel classifications relative to the original variance image in Fig. 3a. A pixel is considered misclassified if the σ^2 value (see section II and Fig. 3) is closer to a different class than the corresponding pixel in the original image (Fig. 3a). Keep in mind that this classification is only used to estimate the error in the restored images, the restoration process itself does not attempt to assign image pixels to discrete classes. Using the same *psf* during reconstruction as for degrading results in a total error of approximately 10% pixel error classifications. This may give the impression of a large error. Therefore, we have included two sample restorations to give an idea of the visual result. In addition, most errors occur at the edges. It can be seen that increasing the frequency by as much as 25% results in an almost binary image (see Fig. 8b). However, a similar reduction in frequency (Fig. 8a) seems to produce acceptable results, *i.e.*, approximately 14% error. Hence, increasing the frequency of the restoring *psf* has the most severe effect on the results. Furthermore, the overall plot is fairly horizontal, implying that slight changes in the *psf* produce acceptable results. Results for the average of the last 100 iterations were only slightly different than the numbers in Fig. 7, differing by a maximum of 0.9%.
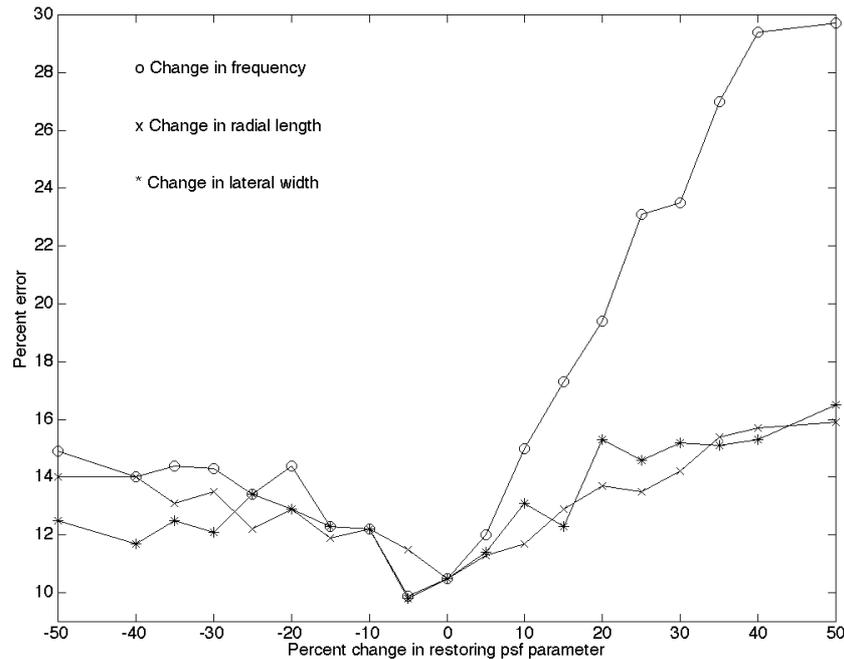
Fig. 7. Percent error pixel classifications relative to the region image in Fig. 3a as a function of percent change in the restoring *psf* parameters from 0 to ± 50% for each. Only one parameter (frequency, length, or width) is changed; the others are kept equal to the degrading *psf*. A pixel is considered misclassified if the σ^2 value is closer to a different class than the corresponding pixel in the original image (Fig. 3a). The estimated variance images used are sample number 500.
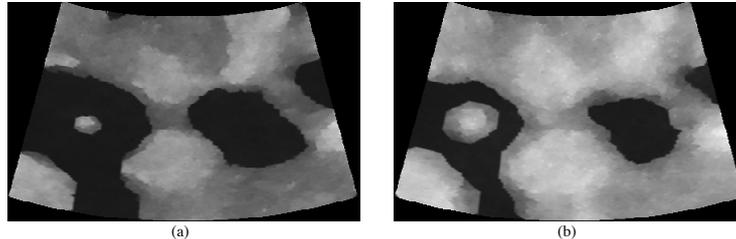

Fig. 8. Sample results from the estimation of the variance image. a) Result when decreasing the restoring *psf* frequency by 25% and b) result when increasing the restoring *psf* frequency by 25%.

C. Restoration of the reflectance image from the image degraded with a measured psf

Fig. 9 shows sample restorations of the simulated ultrasound image degraded with the measured *psf*. The parameter settings for the estimation of the reflectance image from the ultrasound image degraded with the measured *psf* was $\beta = 5$ and $\delta = 1.5$. We used a slightly higher level of noise in this restoration, $\tau^2 = 0.5$ instead of 0.3, which was approximately 10% of the energy in the measured *psf*. We did this to reduce the influence of the restoration *psf* (see Eq. 6, 7, and 9). Nevertheless, the results in Fig. 9 are worse than the results in Fig. 5-8. We did not parameterize the measured *psf*. Instead, we performed the restorations with the same *psf* as we did for the image degraded with the simulated *psf*, except it was shorter in the radial and lateral directions. With a parameterized version of the measured *psf*, the restorations in Fig. 9 might improve slightly. However, we did not investigate this matter because the main idea of this paper is to use an incorrect *psf* during restorations to investigate the robustness of the method under this condition.

The bright spot (4; Fig. 4) becomes larger in the restoration. Also, the restorations of the reflectance image seem a bit more diffuse, the edge definitions are poorer than the situation

was in Fig. 5 and Fig. 6. The black areas 1 and 2 (Fig. 4) have become slightly smaller. The spot 5 from Fig. 4 is almost completely lost. As can be seen in Fig. 9, the effect of a change in frequency seems to be slightly more severe for the restoration of the image degraded with the measured *psf*.
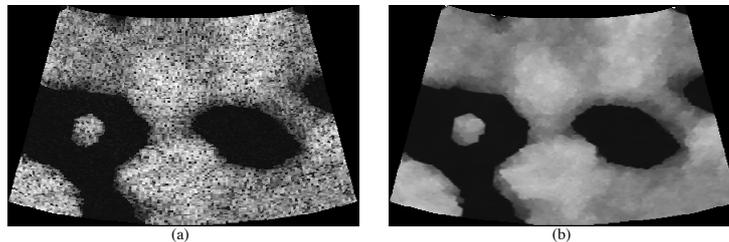

Fig. 9. Results from restoration of simulated ultrasound image degraded with measured *psf* (see Fig. 3d). a) Estimated reflectance image using a restoring *psf* similar to the one in Fig. 1a (same frequency, but 50% shorter in the radial and lateral direction) and b) estimated variance image using a restoring *psf* similar to the one in Fig. 1a (same frequency, but 50% shorter in the radial and lateral direction).

D. Long iteration

In Fig. 10, we show samples from an extended iteration. We observed that brighter areas seemed to grow into the darker regions after about 1000 iterations, but with a lower intensity. However, after a while, this effect reversed, *i.e.,* the brighter areas (or edges) that grew into the darker ones disappeared again. In fact, this was an oscillating effect. Depending on the choice of parameters, this effect started after approximately 1000 iterations and was fully reversed after roughly 2000 iterations. The fact that our method is based on a probability distribution and that we sample pixels randomly from the current image make such effects possible. We can see in Fig. 10b, that after about 1500 iterations, the dark circular region (area 2 in Fig. 4) has gained an edge of gray level between the surrounding bright area and the black. This edge is completely gone after 2000 iterations (Fig. 10c). This effect repeats again, and we observed it at least until the total number of iterations was 100 000 (results not shown).
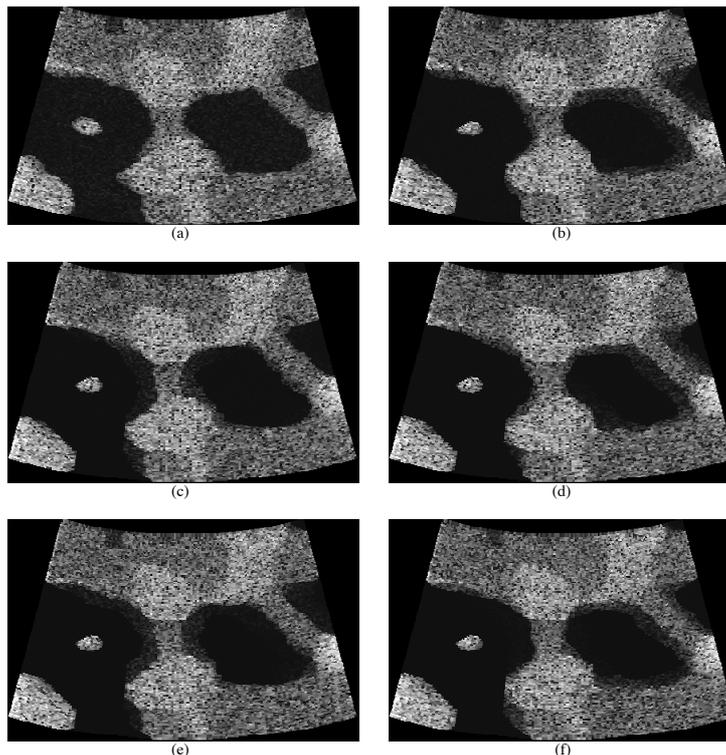
Fig. 10. Reflectance image samples from a long iteration (totally 100000). Samples after a) 500, b) 1500, c) 2000, d) 2500, e) 3100, and f) 3700 iterations. It is easy to see that the brighter regions grow a bit into the darker regions, particularly in the dark regions (1 and 2, Fig. 4). The parameter settings in the restoration were $\beta = 7$ and $\delta = 2$. The starting image was the simulated ultrasound image in Fig. 3c.

V. DISCUSSION

This paper investigated the effects and robustness of a model described in section III and in a companion paper [1] when using spatially invariant *psf*s with different frequency, width, and length during restoration of simulated ultrasound images. We found that fairly large variations, within ±25% of the frequency, width, and length gave satisfactory results. When altering the restoring *psf* parameters more than this, the results were not acceptable. Nevertheless, as can be seen from the results, too small estimates of the length, width, or frequency of the *psf* yielded better results than too large estimates. The worst restorations occur with too high frequencies for the restoring *psf*s (see Fig. 7).

Although testing the algorithm on an image degraded with a measured *psf*, we used a symmetric modeled *psf* for our restoration experiments. The *psf* in ultrasound imaging, particularly in vivo, varies considerably over the imaged area, both radially and laterally [7, 25]. Its shape in one area of the image is strongly dependent on the acoustical properties of the tissue it has traversed on its path from the ultrasound transducer. This means that the *psf* is different in different parts of the image and also changes from image to image when the ultrasound probe is moved. Consequently, any method that attempts to restore ultrasound images should be robust with respect to inaccurate estimates of the *psf*. As mentioned, a few methods exist for estimating 2-D *psf*s [4, 7, 8, 14-19], but this is a field of research that requires further study. Effort should be invested in finding methods for estimation of *psf*s to achieve good restorations.

In our initial tests with the method, we ran restorations using parameter values for β and δ within a certain range. The best results from these test runs resulted in the values used in the restorations in this paper. However, the choice of value for the δ parameter is more or less straightforward, guided by physical considerations (see section III in this paper or the companion paper [1]). The parameter β should ideally be estimated by some kind of data-

driven technique, e.g., the approaches given in [26-28] for parameter estimation in Markov random field models. Nevertheless, we determined appropriate values for β by trial and error from the intervals used in the test runs. The values for the model parameters are indeed subject for further studies.

Our prior model favors homogenous areas with sharp transitions between regions. This is a reasonable characteristic for a model representing ultrasound images because tissues are mainly homogenous with sharp transitions between the different tissue types. A result of this, and a disadvantage to the model, is that we may lose smaller structures in the image. However, these structures are, depending on their size and the size of the degrading *psf*, lost anyway in the blurring process. For instance, the small black spot (3 in Fig. 4) is not visible after degrading the reflectance image with the *psf*. Therefore, such structures, *i.e.*, structures smaller than the extent of the main lobe of the *psf*, will be difficult to recover in the restoration process with methods based on the observed image.

Furthermore, our model allows smooth transitions [1], and hence the errors in Fig. 7 represents mainly classification errors caused by restoration errors along the edges. This is understandable by looking at the initial simulated ultrasound image (Fig. 3c and d), *i.e.*, the edges are very diffuse and poorly defined.

The oscillating effect described in section IV D of the results was discovered early during test runs. This algorithmic issue may be the result of the fact that we use a probability-based model for sampling, hence revealing such effects along edges caused by the jump in intensity, or it may be caused by a lack in mixing of the underlying Monte Carlo Markov chain. However, by running all iterations past at least one oscillation, we confirmed the reversal of the effect and picked iteration number 500 as the result, which was approximately 500 iterations prior to the start of the effect. Nevertheless, this effect needs further investigation.

In addition to the shape of the *psf*, the noise level is a significant characteristic of the observed ultrasound image. In the simulated ultrasound images, the noise level is assumed to be constant, contrary to real ultrasound images, where the noise level can vary greatly depending on the probe and on which parts of the body one wishes to investigate. Furthermore, the noise increases as a function of depth in ultrasound images because of the time gain compensation, a method used to increase the intensity of echoes from the deepest parts of the region being imaged. This fact is not included in our model, but could easily be implemented, and this is of interest in a further investigation of our method.

In conclusion, we have shown that this method for restoration [1] of ultrasound images has potential and should be investigated further, especially with respect to the prior model, real ultrasound images, and spatially invariant pulse estimation techniques. Furthermore, the results show that too small estimates of the length, width, and frequency of the *psf* are preferable to too large estimates. Nevertheless, effort should be invested in finding methods to estimate the *psf* directly from ultrasound images and information about the transmitting probe frequency and depth settings, etc., which might be necessary to obtain full advantage of the possibilities of this method for cases in which the *psf* changes significantly over the tissue volume. However, 2-D Bayesian restoration using our method and a fixed *psf* may yield acceptable results as long as the size of the true variant *psf* does not vary more than ±25% (particularly for the frequency) over the imaged area. This is within normal variation in many cases, as indicated in [3, 25].

VII. APPENDIX

A simplified pseudo code for the implementation of our sampling algorithm.

Initialize the images r and σ^2 as random distributed noise images.
Make visiting scheme table, *i.e.,* table for sequence of pixel positions to visit.
For all pixel positions *i,j* from visiting scheme table,

> calculate $\Upsilon_{m,n}$ according to Eq. 8
> calculate $\mu_{i,j}$ according to Eq. 7
> calculate $\kappa^2_{i,j}$ according to Eq. 9
> generate u_1 and u_2 randomly from (0,1) and set (Box-Müller [21])
> $$norm = \sqrt{-2\ln(u_2)} \cdot \cos(2\pi u_1)$$
> $r_{i,j} = \mu_{i,j} + norm \cdot \kappa^2_{i,j}$ then generates the new sample of *i,j* for the reflectance image.

For all pixel positions *i,j* from visiting scheme table,

> generate a candidate value for σ^2 at the next location
> set $\alpha = \min\left\{1, p\left(\sigma^2_{candidate}|...\right)/p\left(\sigma^2_{current}|...\right)\right\}$ according to Eq. 5.
> If random(0-1)$<\alpha$, set σ^2 = candidate value; if not, keep the current value (Metropolis [22]).

For theoretical background on this algorithm, and an exploration of more sophisticated sampling algorithms, see our companion paper [1].

VIII. REFERENCES

[1]. O. K. Husby, T. Lie, T. Langø, J. Hokland, and H. Rue, "Bayesian 2-D restoration: A model for diffuse ultrasound scattering," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.,* vol. 48, no. 1, pp. Jan 2000.

[2]. E. Steen and B. Olstad, "Volume rendering of 3D medical ultrasound data using direct feature mapping," *IEEE Trans. Med. Imag.,* vol. 13, no. 3, pp. 517-525, 1994.

[3]. L. Ødegård, E. Halvorsen, B. Ystad, H. G. Torp, and B. Angelsen, "Delay and amplitude focusing through the body wall; A simulation study," in *Proc. IEEE Ultrason. Symp.,* San Antonio, Texas, USA, 1996, vol. 2, pp. 1411-1414.

[4]. T. Taxt, "Restoration of medical ultrasound images using two-dimensional homomorphic deconvolution," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.,* vol. 42, no. 4, pp. 543-554, 1995.

[5]. D. Iraca, L. Landini, and L. Verrazzani, "Power spectrum equalization for ultrasonic image restoration," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.,* vol. 36, no. 2, pp. 216-222, 1989.

[6]. T. Loupas, S. D. Pye, and W. N. McDicken, "Deconvolution in medical ultrasonics: Practical considerations," *Phys. Med. Biol.,* vol. 34, pp. 1691-1700, 1989.

[7]. J. A. Jensen and S. Leeman, "Nonparametric estimation of ultrasound pulses," *IEEE Trans. Biomed. Eng.,* vol. 41, no. 10, pp. 929-936, 1994.

[8]. T. Taxt and G. V. Frolova, "Noise robust one-dimensional blind deconvolution of medical ultrasound images," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.,* vol. 46, no. 2, pp. 291-299, 1999.

[9]. S. F. Burch, "Comparison of SAFT and two-dimensional deconvolution methods for the improvement of resolution in ultrasonic B-scan images," *Ultrasonics,* vol. 25, no. 5, pp. 259-266, 1987.

[10]. A. K. Jain, "Fundamentals of Digital Image Processing", Prentice-Hall, Englewood Cliffs, N.J., 1989.

[11]. G. Winkler, Image analysis, random fields and dynamic Monte Carlo methods, Berlin, Springer-Verlag, 1995.

[12]. S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Machine Intell.,* vol. PAMI-6, no. 6, pp. 721-741, 1984.

[13]. R. F. Wagner, M. F. Insana, and D. G. Brown, "Statistical properties of radio-frequency and envelope-detected signals with application to medical ultrasound," *J. Opt. Soc. Amer.,* vol. 4, no. 5, pp. 910-922, 1987.

[14]. T. Taxt, "Comparison of cepstrum-based methods for radial blind deconvolution of ultrasound images," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.,* vol. 44, no. 3, pp. 666-674, 1997.

[15]. J. A. Jensen, "Estimation of pulses in ultrasound B-scan images," *IEEE Trans. Biomed. Eng.,* vol. 10, no. 2, pp. 164-172, 1991.

[16]. U. R. Abeyratne, A. P. Petropulu, and J. M. Reid, "Higher order spectra based deconvolution of ultrasound images," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.,* vol. 42, no. 6, pp. 1064-1075, 1995.

[17]. K. F. Kåresen and E. Bolviken, "Blind deconvolution of ultrasonic traces accounting for pulse variance," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.,* vol. 46, no. 3, pp. 564-573, 1999.

[18]. A. P. Petropulu and U. R. Abeyratne, "System reconstruction from higher order spectra slices," *IEEE Trans. Signal Processing,* vol. 45, no. 9, pp. 2241-2251, 1997.

[19]. H. Pozidis and A. P. Petropulu, "System reconstruction based on selected regions of discretized higher order spectra," *IEEE Trans. Signal Processing,* vol. 46, no. 12, pp. 3360-3377, 1998.

[20]. J. H. Hokland and P. A. Kelly, "Markov models of specular and diffuse scattering in restoration of medical ultrasound images," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.,* vol. 43, no. 4, pp. 660-669, 1996.

[21]. B. D. Ripley, Stochastic Simulation, New York, Wiley, 1987.

[22]. M. A. Tanner, Tools for Statistical Inference. Methods for the exploration of posterior distributions and likelihood functions, New York, Springer-Verlag, 1993.

[23]. W. K. Hastings, "Monte Carlo simulation methods using Markov chains and their applications," *Biometrika,* vol. 57, no. 1, pp. 97-109, 1970.

[24]. P. J. Green, "Reversible jump MCMC computation and Bayesian model determination," *Biometrika,* vol. 82, no. 4, pp. 711-732, 1995.

[25]. L. Ødegård, "Phase aberration correction in medical ultrasound imaging," Ph.D. dissertation, Norwegian University of Science and Technology, Trondheim, 1995.

[26]. C. Geyer and E. A. Thompson, "Constrained Monte Carlo maximum likelihood for dependent data," *J. R. Stat. Soc., Ser. B,* vol. 54, no. 3, pp. 657-699, 1992.

[27]. A. M. Thompson, J. C. Brown, W. Kay, and D. M. Titterington, "A study of methods of choosing the smoothing parameter in image restoration by regularization," *IEEE Trans. Pattern Anal. Machine Intell.,* vol. 13, no. 4, pp. 326-339, 1991.

[28]. D. M. Higdon, J. E. Bowsher, V. E. Johnson, T. G. Turkington, D. R. Gilland, and R. J. Jaszczak, "Fully Bayesian estimation of Gibbs hyperparameters for emission computed tomography data," *IEEE Trans. Med. Imag.,* vol. 14, pp. 516-526, 1997.

Appendix: Paper G

SonoWand®, an Ultrasound-Based Neuronavigation System

Aage Gronningsaeter[1], Atle Kleven[1], Steinar Ommedal[2], Tore Erling Aarseth[1], Torgrim Lie[1], Frank Lindseth[2], Thomas Langø[2], Geirmund Unsgård[3]

[1] SINTEF Unimed, Ultrasound, 7465 Trondheim, Norway, presently with MISON AS, 7489 Trondheim, Norway
[2] SINTEF Unimed, Ultrasound, 7465 Trondheim, Norway
[3] Department of Neurosurgery, University Hospital and Medical Faculty, Norwegian University of Science and Technology (NTNU), 7006 Trondheim, Norway.

ABSTRACT

OBJECTIVES: We have integrated a neuronavigation system into an ultrasound scanner and developed a single-rack system that enables the surgeon to perform frameless and armless stereotactic neuronavigation by means of intraoperative 3D ultrasound data as well as preoperative MRI- or CT images.

TECHNICAL DEVELOPMENTS: The system consists of a high-end ultrasound scanner, a modest cost computer and an optical positioning/digitizer system. Special technical and clinical efforts have been made in order to achieve high image quality. A special interface between the ultrasound instrument and the navigation computer ensures that there is a rapid transfer of digital 3D data with no loss of image quality. The positioning system tracks the position and orientation of the patient, ultrasound probe, pointer and various surgical instruments.

RESULTS: The image quality improvements have enabled us, in most cases, to extract information from ultrasound that has a similar clinical value as preoperative MRI. The overall clinical accuracy of the ultrasound-based navigation system is expected to be comparable to, or better than that of an MRI-based system.

CONCLUSIONS: SonoWand enables neuronavigation by the direct use of intraoperative 3D ultrasound. However, further research will be necessary to explore the potential clinical value and the limitations of this technology.

INTRODUCTION

Some years of experience with computer-based navigation systems have clearly demonstrated the need for an intraoperative imaging modality that can cope with the anatomical changes that normally occur during cranial surgery. Three concepts seem to represent the options for the foreseeable future: I) Open MRI in which the patient has to be transported in and out of a sterile draped magnet for updating 3D data set, a process that typically takes 30-70 minutes (18), II) interventional MRI where the surgeon is operating inside a sterile draped magnet and thus can perform 2D imaging in close to real time and slower 3D scans, the typical 3D acquisition time is 5 minutes, and III) ultrasound which only requires a small sterile draped scanhead in the field, making it possible to perform real-time 2D scanning whenever needed and rapid 3D acquisition several times during the procedure. The typical time for updating the 3D ultrasound data set is 2-3 minutes, but new technology will enable 3D imaging in real time in the foreseeable future.

There seems to be significant expectations about intraoperative MRI-systems on the market, but these systems require very high investments, high running costs, and a special operating room and surgical equipment. Ultrasound has not been given very much attention until recently, probably due to limited image quality, lack of dedicated equipment and limited skills to interpret such images. Some groups and companies have lately connected an ultrasound scanner to a conventional navigation system, digitized the analog video signal from the scanner and displayed a real-time 2D image on the navigation computer side-by-side with the corresponding MRI-slice. This modality has proven beneficial for simplifying the interpretation of ultrasound and for the bulky identification and quantification of brain shift (3, 9, 10, 12, 16). However, these solutions require two space-consuming racks in the operating room, and a number of compromises are normally made on image quality and data transfer capabilities.

We have developed a different system that integrates a neuronavigation system into a high-end ultrasound scanner. The navigation computer and the optical digitizer or camera system are built into the ultrasound instrument, meaning that only one rack is required in the operating room. There is a digital interface between the ultrasound scanner and the navigation computer which enables the rapid communication of digital data with no loss of image quality. A phased array probe is normally selected, this is a convenient probe with a relatively small footprint. The image quality is optimized for imaging the brain by selecting suitable scanning parameters. Furthermore, special actions are taken to optimize image quality during patient preparation, surgical planning as well as during the clinical procedure.

The result is a single-rack neuronavigation system called SonoWand® that can work as a stand-alone ultrasound scanner, a conventional MRI- or CT based neuronavigation system and as a neuronavigation system with rapid access to intraoperative 3D ultrasound data. The system enables the surgeon to *plan* the surgery by conventional MRI navigation, acquire a 3D ultrasound scan before opening the dura mater in order to compare the two modalities, and then *guide* the resection and surgery using repeated 3D ultrasound acquisitions.

A two-rack prototype has been used for clinical testing since November 1997 at the University Hospital in Trondheim. A single-rack prototype has been recently built, see Figure 1, and this version will be subject to clinical trials and further R&D at the University of Heidelberg as well as in Trondheim. Commercialization of the system is being done by MISON AS, the company that has the rights the brand name SonoWand®. The purpose of this paper is to describe the two-rack prototype and present the results of our image quality enhancement work. The resolution and accuracy of the system is described in a companion paper (7), while the clinical experience from more than 45 cases will be described in a paper which is in preparation.

FIGURE 1.   The SonoWand prototype represents a single-rack system that can operate as a stand-alone ultrasound scanner, a conventional MRI/CT-based neuronavigation system and a neuronavigation system with rapid access to intraoperative 3D ultrasound.

TECHNICAL DEVELOPMENTS

The navigation equipment

The prototype system consists of a high-end ultrasound scanner (System-fiVe, GE Vingmed Ultrasound, Horten, Norway), a medium cost computer (Daystar Genesis MP900, Daystar, USA) for image processing and navigation, and an optical positioning system (3D digitizer or tracking system). We have used the camera unit of the VectorVision system (BrainLAB, München, Germany) as the positioning system in the two-rack prototype. However, an interface has lately been made for FlashPoint 5000 (Image Guided Technologies Inc., Boulder, USA) and the Polaris (Northern Digital Inc., Ontario, Canada) systems, a prerequisite for integration into a commercial single-rack solution. Two adjustable arms have been attached to the top of the scanner, one for the camera system and one for two flat-screen monitors. One monitor is used for the ultrasound scanner, while the other is used for the navigation computer. A direct Ethernet link has been established between the ultrasound scanner and the navigation computer in order to provide rapid transfer of high-quality digital

3D ultrasound data. MRI or CT data can be imported into the navigation computer through an Ethernet connection using a DICOM standard.

The position and orientation of the Mayfield frame, pointer devices, ultrasound probes and surgical instruments can be measured by the camera system when a positioning frame has been attached to the device. All positioning frames are equipped with three reflecting spheres. Figure 2 shows the three ultrasound probes that have been subject to clinical trials. A special adapter is glued to each probe in order to enable repetitive and precise placement of the positioning frame, even through the sterile drape.
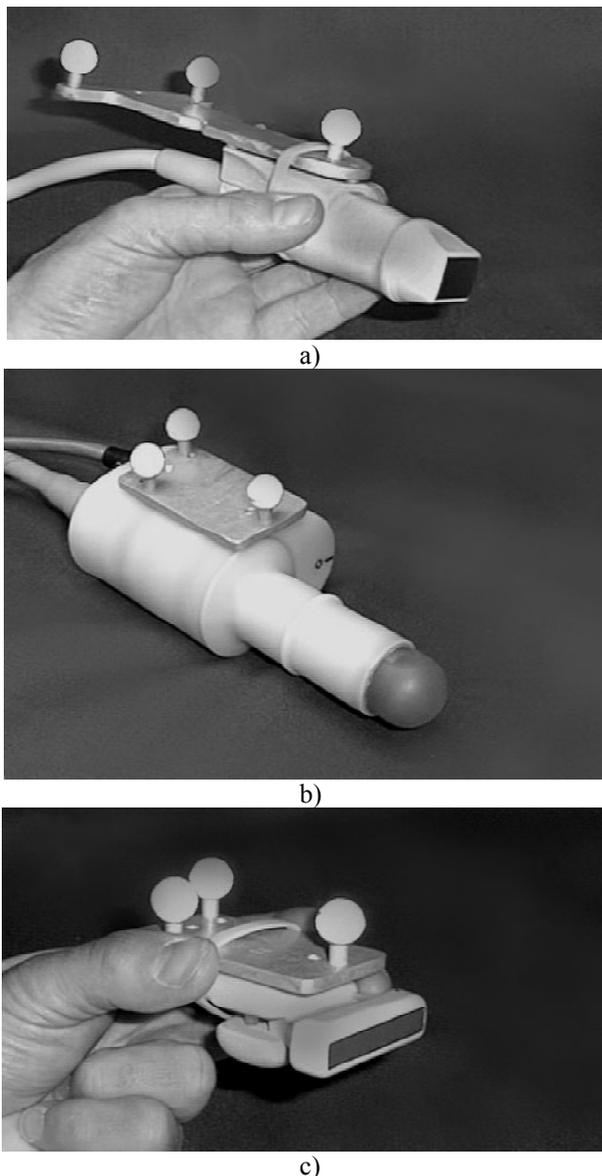

a)

b)

c)

FIGURE 2.   Ultrasound probes with passive optical positioning adapters attached: a) 4-8 MHz phased array for overview imaging and free-hand 3D scans, b) Motorized 3D-probe with a 5 MHz annular array transducer and c) a 5-9 MHz linear array for superficial imaging.

Two different surgical instruments have also been equipped with positioning frames as illustrated in Figure 3. One frame is permanently fixed to a biopsy forceps, while another frame can be attached to a CUSA System 200 ultrasound aspirator (Valleylab, Boulder, Colorado, USA) with a simple hand grip. These instruments can be used as pointer devices, and their distal tips can thus be navigated into the brain by means of the image information on the monitor. A pointer device with three reflecting spheres (two is normal) has also been

designed which makes it possible to not only measure the position and direction of the pointer, but also the degree of rotation.
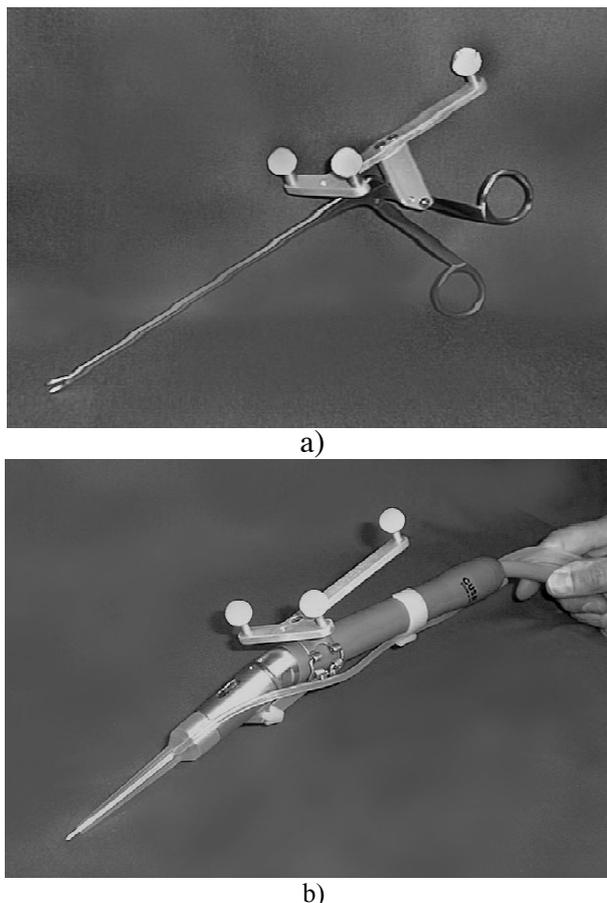

a)

b)

FIGURE 3.   Surgical instruments such as (a) biopsy forceps and (b) CUSA ultrasound aspirator can be equipped with positioning adapters. The instruments can be navigated deep into the brain by means of image-guidance.

Registration and calibration algorithms

A standard patient registration algorithm has been implemented in the system in order to allow conventional navigation by means of preoperative MRI-images. The algorithm is based on pointing at skin fiducials or anatomical landmarks using a pre-calibrated pointer device. Furthermore, we have developed a method for calibrating the ultrasound probe so that the computer can calculate the exact position and orientation of the ultrasound scan plane based on the measured position and orientation of the attached positioning frame. Accurate probe calibration takes place for each individual probe in the laboratory using a special designed water-tank. Finally, a special calibration algorithm has been developed for the ultrasound aspirator since the positioning frame has to be mounted to the aspirator under sterile conditions. Three-dimensional ultrasound data is acquired by tilting the probe approximately 90 degrees by hand for approximately 15 seconds.

The navigation software

The navigation software contains the means for importing MRI or CT data, performing patient registration and displaying navigation images on the monitor. In addition, the system can acquire 3D ultrasound data by tracking the position and orientation of the ultrasound probe during a free-hand probe movement. A pyramid-shaped volume of the brain is acquired by tilting the probe typically 90 degrees for 15 seconds. The digital images are reconstructed

into a regular volume, and the navigation software treats this ultrasound volume in a similar way as the MRI and CT volumes. The system supports three different navigation features:

Ultrasoundprobe-driven anyplane slicing: The first navigation feature is based on using the ultrasound probe for real-time 2D scanning where the ultrasound image is displayed on the ultrasound monitor. The position and orientation of the probe is tracked by the positioning system and the corresponding cross-sectional slice through the preoperative MRI or CT data set is displayed on the navigation monitor. This technique is similar to that reported previously by other groups, however with the difference that we display the two images on two different monitors.

Instrument-driven orthogonal slicing: The second navigation feature is similar to the conventional orthogonal display technique which is available in most neuronavigation systems. The surgeon can use a pointer to obtain axial, coronal and sagittal views from a preoperative MRI or CT data set. However, an additional set of axial, coronal and sagittal views, taken from an intraoperative 3D ultrasound data set can also be provided, see Figure 4. The visualization is controlled by the active tool which can be a pointer, a biopsy forceps or an ultrasound aspirator. The instrument axis is indicated on the images with a dashed line and the tip is indicated by crosshairs. This display mode facilitates easy identification and localization of residual tumor fractions using the instrument tip.
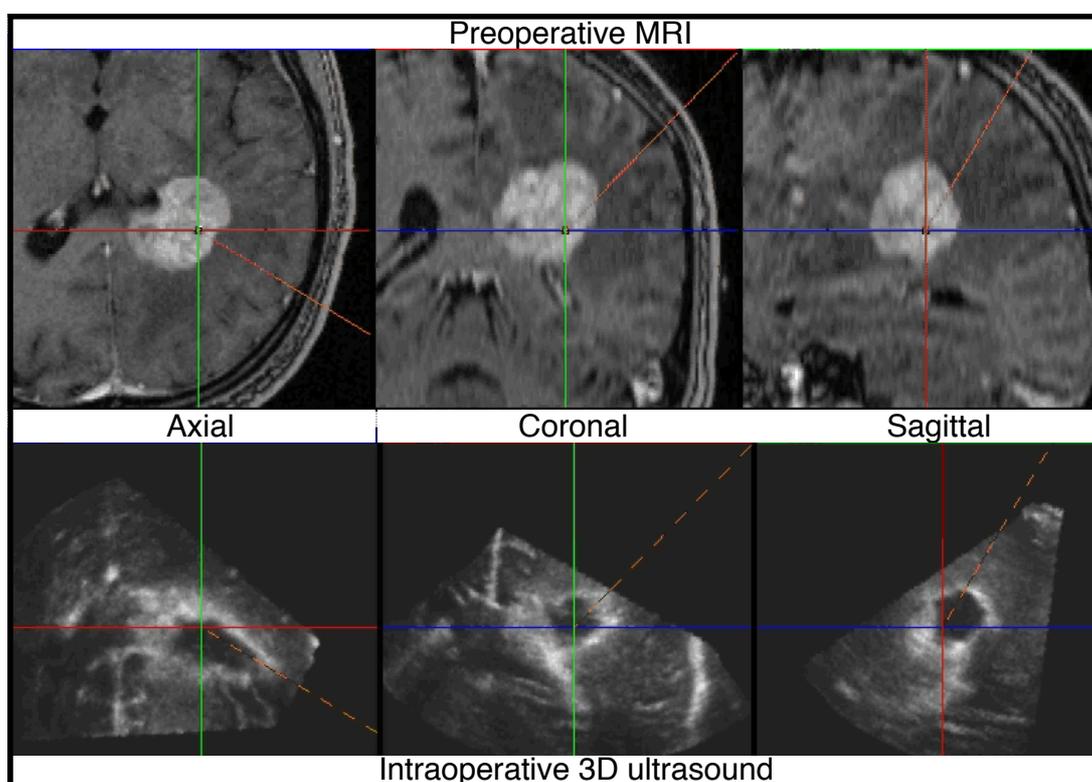

FIGURE 4.   The *instrument-driven orthogonal-slicing technique* facilitates comparison between preoperative MRI and intraoperative 3D ultrasound using pairs of corresponding axial, coronal and sagittal images. The ultrasound aspirator tip has been navigated into the resection cavity in order to identify and localize residual tumor tissue in a glioblastoma.

Instrument-driven anyplane slicing: The third navigation feature represents a combination and simplification of the two methods that have been mentioned. Only two images are displayed in order to reduce the amount of information, one MRI slice and the corresponding ultrasound slice as described in the first navigation features. However, the navigation device is no longer the ultrasound probe, but a pointer or surgical instrument. Furthermore, both slices are taken from 3D volumes as in the second method, but the slices are not restricted to coronal, axial or sagittal views. This technique is illustrated in Figure 5 and gives the surgeon

more flexibility to display arbitrary cross sections, and the reduced information content on the monitor simplifies the interpretation for the surgeon.
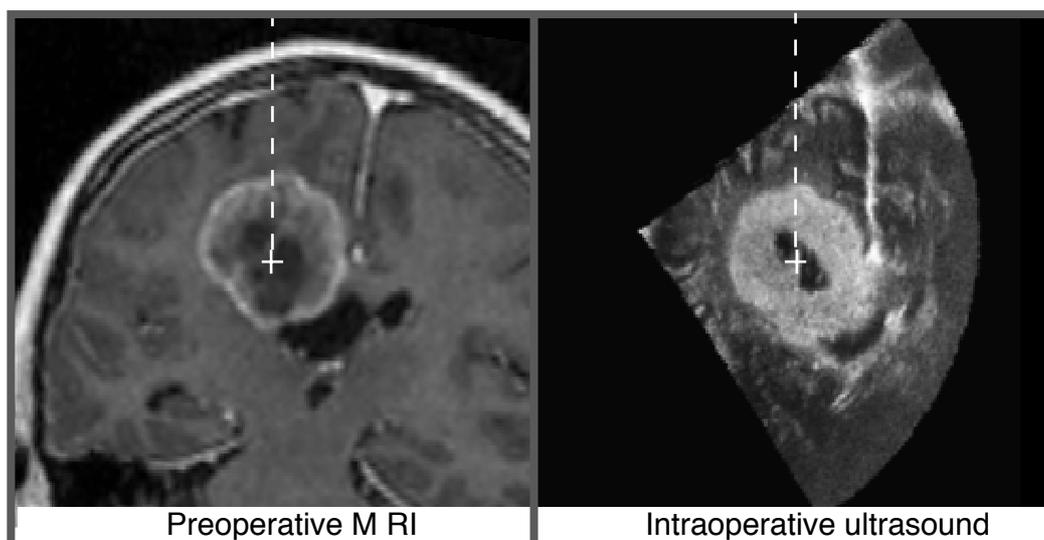

FIGURE 5. The *instrument-driven anyplane-slicing technique* displays only one set of corresponding images. The two images are taken from their corresponding 3D volumes in a cross section which is located in front of the instrument. This technique simplifies the perception, and gives the operator more flexibility to display any desired cross section in the region of interest.

Image quality improvement efforts

Phased array probes (Figure 2a) have been primarily developed for cardiac imaging. The probe footprint is relatively small in order for the beam to pass between the ribs, and the scanning parameters are optimized for imaging deep-seated moving structures such as the valves and the myocardium. The small footprint fits well with the needs in brain imaging from a small craniotomy, but the high frame rate setup does not. We have optimized the scanning setup for stationary structures by building up one image from up to 9 beam shots per beam direction in order to optimize the transmit- and receiving focus in each depth zone. The result is very high resolution in both the radial direction (along the beam) as well as in the lateral direction (perpendicular to the beam). Unfortunately, the resolution in the elevation direction, i.e. the "thickness" of the scan plane, cannot be adjusted electronically with this probe type. The resolution along the direction of 3D acquisition will therefore be limited, especially outside the focal region (very close to, and far from the transducer). Furthermore, the scanning sector is narrow in the near field, and broad in the deeper regions. This probe is therefore best suited for lesions that are located some centimeters from the surface.

An alternative to a phased array transducer is an annular array. These transducers have a circular symmetric beam, which means that the resolution in the lateral and elevation directions are equal and can be adjusted electronically. Two-dimensional sector scanning is realized by a rapid motorized tilting movement. We expected this probe to yield better spatial resolution, especially in the direction of 3D acquisition and far from the transducer. A motorized 3D prototype probe was therefore built (Figure 2b) in order to compare the two different arrays and acquisition techniques. The 3D acquisition is achieved by rotating the internal 2D-probe assembly using a second motor.

Linear array probes (Figure 2c) have been primarily designed for transcutaneous vascular imaging. The transducer frequency is higher than the phased array probe, yielding a better spatial resolution at the expense of lower penetration. The footprint is larger ($\approx$10x40 mm) and the field of view is rectangular. These properties make this probe best suited for imaging

superficial lesions from a larger craniotomy. The scanning parameters have been optimized for this probe in order to improve the quality of brain imaging.

RESULTS

The technical developments and clinical work that have been taken to improve image quality have been very successful. The quality of our ultrasound images enables the surgeon to use 3D ultrasound in a similar manner as other groups use 3D MRI data for navigation and surgical guidance. An example of the image quality from the 4-8 MHz phased array probe is shown in Figure 6a. Fine details can be seen with high spatial resolution even at depths of 7-8 cm. Since the system is capable of updating the 3D map several times during surgery, brain shift is no longer a severe problem. Furthermore, we have found that the accuracy of the system is comparable to, or even better than conventional navigation systems that are based on preoperative MRI.

The motorized 3D probe with an annular array transducer did not give improvements in 3D resolution as expected compared to the phased array transducer. One explanation is the fact that mechanically scanning probes have limited ability to be optimized for 2D image quality compared to electronically scanned probes. Furthermore, the transducer array was also slightly smaller yielding a lower resolution in general. Consequently, the probe was rarely used in the clinic.

The 5-9 MHz linear array probe has occasionally been used for superficial lesions. An example from a glioblastoma with multiple cysts is shown in Figure 6b where the spatial resolution seems to be smaller than one millimeter. Although this is not documented yet, we have an impression that this probe yields images with better resolution than the corresponding preoperative whole-head MR-images (Gyroscan S15 HQ, Philips Medical Systems). Regardless of probe type, ultrasound often provides details and structures than cannot be seen in the corresponding MRI data set.
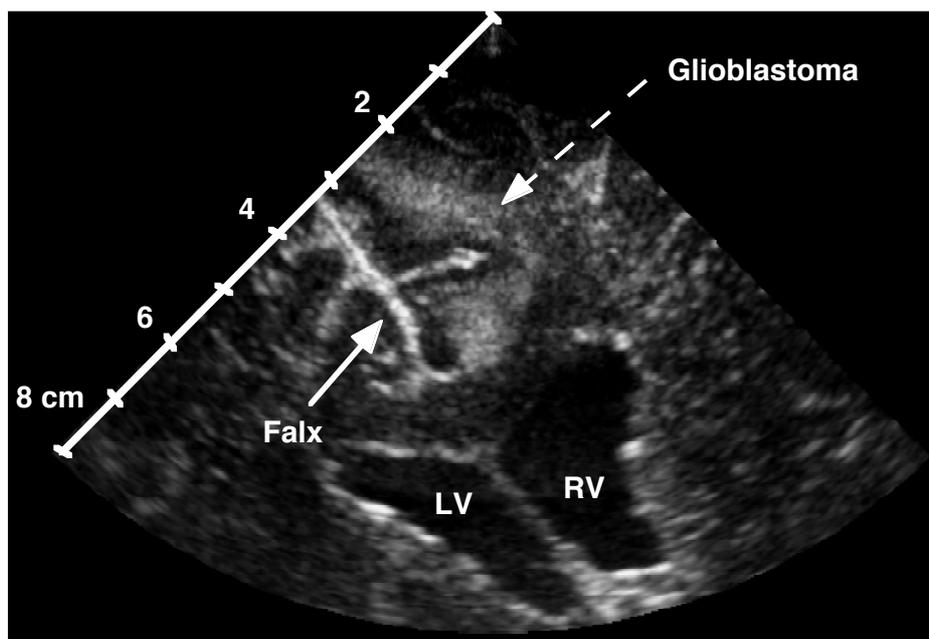
a)
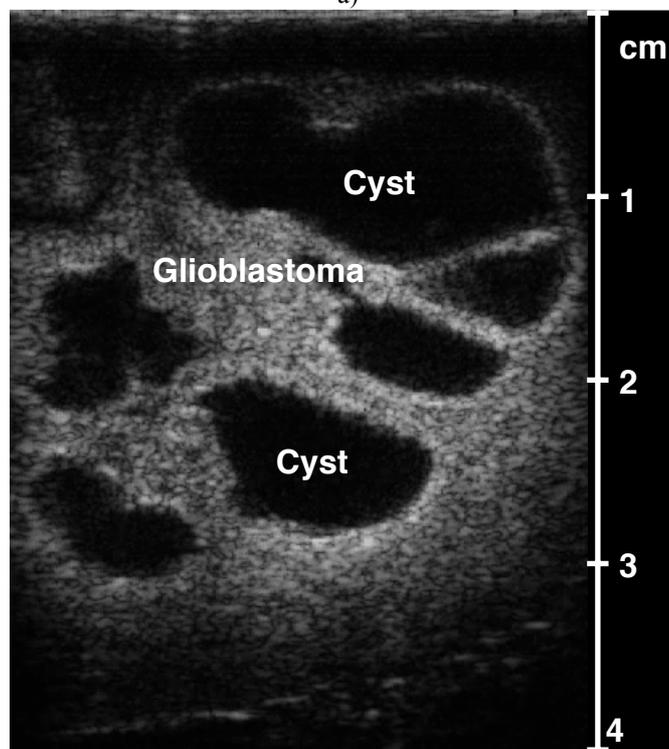
b)

FIGURE 6.   High quality ultrasound is possible to achieve by optimizing a high-performance ultrasound scanner with a convenient probe to brain imaging. a) Coronal overview image using the 4-8 MHz phased array. b) Superficial image of a glioblastoma with multiple cysts using the 5-9 MHz linear array.

DISCUSSION

The importance of high image quality: Based on our experience with ultrasound in neurosurgery (8, 11, 17), we are surprise that ultrasound has not gained more acceptance in such surgery. The first experiments with ultrasound pulse-echo in the human brain was performed in the early fifties (14). However, real-time 2D imaging was not introduced until the late sixties and the beginning of the seventies. Researchers and surgeons believed that this technique would become an important tool for brain surgery as commercial instruments became available in the eighties. However, the first period of enthusiasm was followed by disappointment. The situation today seems to be that neurosonography has gained a certain acceptance for localizing various tumors, hematomas, cysts, blood vessels, aneurysms, and necrotic areas (2, 4, 6). Some groups have also applied ultrasound for localizing residual tumor tissue in order to improve gross total resection (15, 19). Ultrasound-guided biopsy is quite common, and some research groups are working with more sophisticated projects such as ultrasound-guided endoscopy (1) and 3D ultrasound (3, 5, 13). However, the current use of ultrasound is sparse compared to the use of CT and MRI.

Some possible explanations are addressed by D.W. Roperts, P.J. Kelly, R.J. Maciunas and G.H. Barnett in their comments on a paper by N. Hata et al. (9):   I) Ultrasound has demonstrated limited quality due to poor spatial and contrast resolution and due to artifacts or drop-outs from blood, air and instruments. II) Commercial ultrasonography only provides 2D cross-sectional images, that are normally of oblique orientation, making it difficult to relate the structures seen on the monitor to the anatomy of the patient. All these limitations have been addressed during the design and clinical application of SonoWand®, and we consider that these problems have now been minimized.

Direct or indirect use of 3D ultrasound: Like other research groups, we also find the comparison between real-time 2D ultrasound and the corresponding MRI-slice useful for the interpretation of ultrasound as well as for identification of brain shift (3, 9, 10, 12, 16). However, a conventional MRI-based neuronavigation system benefits only slightly from a 2D ultrasound system that only *identifies* brain shift. The system should ideally be able to *modify* the images so that the 3D map corresponds to the anatomy at any time. Some groups have tried to solve this problem by using intraoperative 3D ultrasound in an *indirect* way by measuring brain shift and transferring this information to the navigation system (5). The movement of some anatomical landmarks can be registered in the ultrasound volumes and transferred to an elastic model that manipulates the preoperative MRI volume correspondingly (3). The surgeon is then supposed to trust and navigate according to this *manipulated* preoperative MRI volume.

We have solved this problem in a more *direct* way by simply navigating by means of 3D ultrasound. A direct comparison between preoperative MRI and intraoperative 3D ultrasound is useful before the resections start since MRI and ultrasound represent the tissue characteristics differently. However, when brain shift occurs, we navigate purely by means of 3D ultrasound. This is possible due to the high image quality in SonoWand, high navigation accuracy and an acceptable ability to differentiate tumor tissue from normal brain structures, even on low grade astrocytomas. This fact is not yet fully understood, and consequently this topic is subject to further research.

Real-time 3D ultrasound: A modern high-end ultrasound scanner is capable of scanning approximately 20 high quality, wide sector images per second. By reducing the sector width and the spatial resolution, several hundred images can be scanned per second. This means that a limited 3D sector, for example 30 degrees in both directions, can be scanned with several 3D volumes per second. This technique will require an advanced scanner with electronic beam-steering in both directions, or a probe with a motorized movement of the 2D scan plane. Both techniques are expected to be commercially available within a few years. Real-time 3D imaging will make it possible to *see* the moving surgical instrument *directly* in the image, relative to the surrounding structures. Today, only a *computer model* of the instruments, such as a colored line with crosshairs, is overlaid the images. Thus, we expect this future

technology improvement to increase the surgeon's confidence in the navigation system, and also make it possible to develop a simpler and more intuitive user interface for computer aided neurosurgery. We have prepared the SonoWand® system for this future upgrade by planning a rapid custom-designed communication protocol between the ultrasound scanner and the navigation computer. This protocol will enable real-time transfer of digital raw data with no loss of image quality.

Video signal or digital data transfer: Vendors of conventional neuronavigation systems are now putting a lot of efforts into integrating ultrasound into their products. Ideally they want to interface to any ultrasound scanner in order to give their customers the ability to chose their ultrasound system more or less freely. Since composite video is the only image standard all ultrasound vendors provide today, most companies design a video cable between the two systems, using a video-grabbing board in their navigation system. However, this strategy has a number of limitations including: quality loss due to digital-to-analog-to-digital conversions, limited control with the exact geometry of the ultrasound image as well as a constant and limited frame rate (typically 25-30 interlaced, resulting in 12.5-15.0 full frame images per second). Many ultrasound vendors now provide a DICOM standard which means that the images can be transferred in digital form. However, this standard does not provide real-time data transfer capabilities, a prerequisite for real-time 2D and future 3D ultrasound.

SonoWand is currently based on a custom designed direct Ethernet link between the ultrasound scanner and the navigation computer, thus there is no quality loss associated with data transfer because of digital representation. The system also has full control of the geometry of the ultrasound images since such information is transferred with the 3D data set. Finally, this interface is well suited for further upgrading to real-time digital transfer.

CONCLUSIONS

We have developed an ultrasound-based neuronavigation system that enables the surgeon to perform free-hand and arm-less stereotactic neuronavigation by means of intraoperative 3D ultrasound images as well as preoperative MRI or CT images. The navigation system is physically integrated into an ultrasound scanner, and this single-rack solution occupies less space in the operating room than more common two-rack solutions. A high-performance ultrasound scanner has been selected, and a number of technical and clinical actions were taken to improve image quality. The quality of the images and the accuracy of the system, make it possible to navigate directly by 3D ultrasound in a similar manner as with conventional MRI- or CT-based systems. Tight electronic integration between the ultrasound scanner and the navigation computer maintains the rapid transfer of high-quality digital ultrasound data. The system is thus well prepared for a future upgrade to real-time 3D operation.

Acknowledgments

.

References

1.   Auer LM: Ultrasound stereotaxic endoscopy in neurosurgery. Acta Neurochir Suppl (Wien) 54:34-41, 1992.

2.   Auer LM, Velthoven Vv: Intraoperative Ultrasound (US) Imaging. Comparison of Pathomorphological Findings in US and CT. Acta Neurochir (Wien) 104:84-95, 1990.

3.   Bucholz RD, Yeh DD, Trobaugh J, McDurmont LL, Sturm CD, Baumann C, Henderson JM, Levy A, Kessman P: The Correction of Stereotactic Inaccuracy Caused by Brain Shift Using an Intraoperative Ultrasound Device. Lecture Notes in Computer Science Proceedings of the 1997 1st Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery, CVRMed-MRCAS'97, Grenoble, 1997.

4.   Chandler WF, Knake JE, McGillicuddy JE, Lillehei KO, Silver TM: Intraoperative use of real-time ultrasonography in neurosurgery. J Neurosurg 57:157-163, 1982.

5.   Erbe H, Kriete A, Jödicke A, Deinsburger W, Böker DK: 3D-Ultrasonography And Image Matching For Detection Of Brain Shift During Intracranial Surgery. Computer Assisted Radiology, CAR '96, Paris, France, 1996.

6.   Gaab MR: Intraoperative ultrasound imaging in neurosurgery. Ultraschall Med 11:62-71, 1990.

7.   Gronningsaeter A, Lindseth F, Langø T, Unsgård G: Accuracy Assessment of a Prototype Ultrasound-Based Neuronavigation System. Neurosurgery Vol. ?, No. ?:??-??, 2000.

8.   Gronningsaeter A, Unsgård G, Ommedal S, Angelsen BAJ: Ultrasound-guided neurosurgery: a feasibility study in the 3-30 MHz frequency range. Br J Neurosurg 10, no 2:161-168, 1996.

9.   Hata N, Dohi T, Iseki H, Takakura K: Development of a Frameless and Armless Stereotactic Neuronavigation System with Ultrasonographic Registration. Neurosurgery 41:608-614, 1997.

10.   Hirschberg H, Unsgaard G: Incorporation of Ultrasonic Imaging in an Optically Coupled Frameless Stereotactic System. Acta Neurochir Suppl (Wien) 68:75-80, 1997.

11.   Kleven A, Lie T, Kaspersen JH, Aarseth TE, Lindseth F, Gronningsaeter A, Unsgård G, Myhre HO, Mårvik R: An Ultrasound Based Navigation System for Minimally Invasive Surgery. The 6th World Congress of Endoscopic Surgery, Rome, Italy, 1998.

12.   Koivukangas J, Louhisalmi Y, Alakuijala J, Oikarinen J: Ultrasound-controlled neuronavigator-guided brain surgery. J Neurosurg 79:36-42, 1993.

13.   Koivukangas J, Ylitalo E, Alasaarela E, Tauriainen A: Three-dimensional ultrasound imaging of brain for neurosurgery. Ann Clin Res 18 Suppl. 47:65-72, 1986.

14.   Leksell L: "Echo-Encephalography.I. Detection of Intracranial Complications Following Head Injury". Acta Chir Scand 110:301-308, 1955.

15.   LeRoux PD, Winter TC, Berger MS, Mack LA, Wang K, Elliott JP: A Comparison Between Preoperative Magnetic Resonance and Intraoperative Ultrasound Tumor Volumes and Margins. J Clin Ultrasound 22, No. 1:29-36, 1994.

16.   Trobaugh JW, Richard WD, Smith KR, Bucholz RD: Frameless stereotactic ultrasonography: Method and applications. Comput Med Imaging Graph Vol. 18, No.4:235-246, 1994.

17.   Unsgård G, Kleven A, Ommedal S, Gronningsaeter A: Intraoperative ultrasound and neuronavigation. European Congress of Neurosurgery, Copenhagen, Denmark, 1999.

18.   Wirtz CR, Bonsanto MM, Knauth M, Tronnier VM, Albert FK, Staubert A, Kunze S: Intraoperative Magnetic Resonance Imaging To Update Interactive Navigation in Neurosurgery: Method and Preliminary Experience. Comput Aided Surg 2:172-179, 1997.

.

19. Woydt M, Krone A, Becker G, Schmidt K, Roggendorf W, Roosen K: Correlation of Intra-Operative Ultrasound with Histopathologic Findings After Tumour Resection in Supratentorial Gliomas. Acta Neurochir (Wien) 138:1391-1398, 1996.

.