ISSUES IN 3-D FREE-HAND MEDICAL ULTRASOUND IMAGING

R.N. Rohling and A.H. Gee

CUED/F-INFENG/TR 246

January 1996

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

E-mail: rnr20@eng.cam.ac.uk, ahg@eng.cam.ac.uk

Abstract

A drawback of conventional 2-D ultrasound imaging is the requirement that the physician mentally reconstruct 3-D anatomy given multiple 2-D image slices. This paper reviews attempts to overcome this problem using 3-D ultrasonic imaging. It is argued that *free-hand* imaging holds the most promise for effective and inexpensive 3-D ultrasound. In the free-hand paradigm, the physician is allowed to move the probe freely over the region of interest, while a sensing device records the position of each scan. The set of 2-D scans and position data are subsequently used to construct a 3-D data set which can be rendered on a computer monitor or used for volumetric data analysis. In this paper the strengths and weaknesses of free-hand imaging are identified, with reference to sources of error in the measurement, reconstruction, visualisation and volumetric analysis processes. The findings suggest several key research topics, aimed at overcoming some of the limitations of free-hand imaging.

1 Introduction and Outline

Ultrasound has proven its utility during more than 30 years of clinical use. High intensity ultrasound is used for surgery, medium intensity for physiotherapy and low intensity for ultrasound sonography and Doppler ultrasound [27]. This paper is concerned with only ultrasound sonography, hereafter referred to as simply ultrasound imaging.

2-D real-time ultrasound imaging (also called B-scan imaging) is commonly used for diagnosis in a variety of medical fields including obstetrics, gynaecology, opthalmology, cardiology, internal medicine, fetal neurology, as well as surgery and needle-guidance [27, 61]. A few of the reasons for its popularity are its safe non-invasive nature, its real-time capability, and relatively cheap cost of acquisition. There are several drawbacks of conventional 2-D ultrasound imaging compared to other imaging methods such as computed tomography and magnetic resonance imaging. These include the low signal-to-noise ratio of ultrasound images, the inability to image through bone or air, significant artifact content, non-uniform spatial resolutions and the obvious lack of 3-D images.

The last drawback is the main focus of this paper: the requirement that the physician mentally reconstruct 3-D objects given multiple 2-D image slices. The other drawbacks are hoped to be improved by the 3-D reconstruction process. Physicians have indicated that there is utility in 3-D medical ultrasound imaging of anatomical structures such as the urethra [38], vascular structure [20, 28, 29, 41, 44], fetus [11, 31, 37, 44], gall bladder [18], breast [35], kidney [22, 44], heart [32, 39, 47], uterine cavity [3] and others.

The 3-D ultrasound imaging process involves the acquisition of a set of 2-D images and subsequent reconstruction, rendering and volumetric data analysis. The purpose of this paper is to investigate the entire 3-D imaging process, discover the main limitations of current methods and identify several areas for further research. The main goal of 3-D ultrasound imaging is to provide the physician with consistent and accurate information of the anatomical structures of interest. A thorough investigation into the sources of errors at each step of the imaging process is required.

Section 2 begins with a statement of the ultimate goals of real-time 3-D ultrasound imaging, then describes the various methods used to achieve this goal. The relative advantages and disadvantages of the free-hand approach, compared to the other acquisition methods, are described. Section 3 describes the 3-D free-hand imaging process in more detail. It begins with the sources of measurement error in the acquisition process. The problems and issues in reconstruction are then discussed. Brief descriptions are also given of the popular visualisation and volumetric data analysis methods. Finally, Section 4 states conclusions

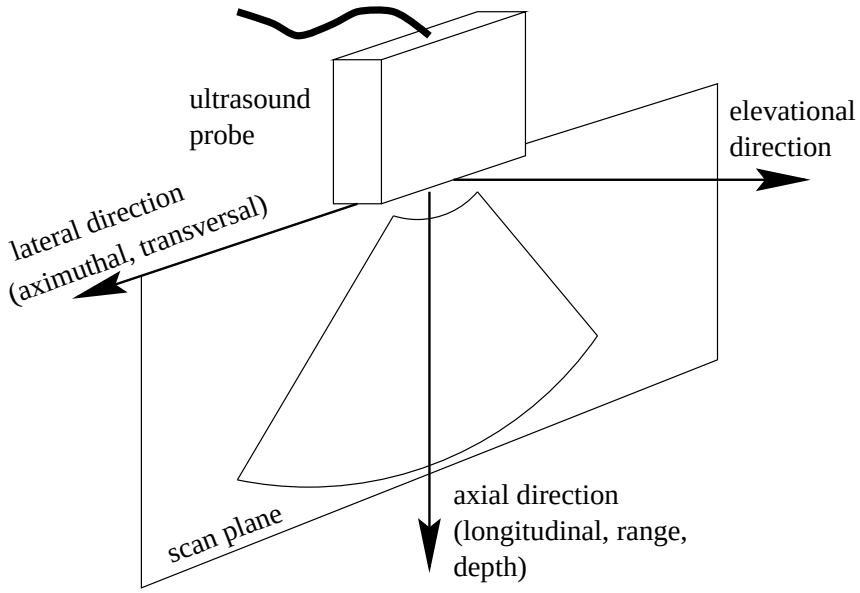

Figure 1: Nomenclature of imaging directions. Other common names of the directions are in parentheses.

about the strengths and weaknesses of free-hand imaging. Several relevant research areas, based on the investigation of the measurement errors, are then suggested .

2 3-D Ultrasound Acquisition Systems

Real-time 3-D ultrasound capable of high resolution imaging over a large volume is the ultimate goal. A 3-D voxel array with 128×128×128 8-bit elements and 0.5 vps (volumes per second) is a reasonable initial goal for imaging fetuses [11]. An array of 512×512×512 8-bit elements at 30 vps is the goal for echocardiography [61]. The finite propagation speed of the ultrasound pulses limits the acquisition rate. Parallel processing with high speed computers and large memory capacity are therefore required to reach these goals [61] . Real-time 3-D ultrasound research has not yet achieved a level suitable for clinical use and remains an active research area.

One unconventional approach — named Explososcan — proposes a two-dimensional transducer array. The array transmits a single wide-angle ultrasound pulse and uses multiple receive-mode array elements [61] . Parallel processing of the array elements for simultaneous multiple scan formations in 3-D is proposed. Currently the spatial resolution is too low for clinical use.

If the real-time requirement is relaxed, several other approaches are possible. Considerable research has been done on 3-D reconstructions from an accumulated set of 2-D image slices. A prerequisite for this method is the knowledge of the relative locations and orientations of the 2-D scan planes. Good overviews are found in papers of 3-D imaging in obstetrics [52] and 3-D echocardiography [39, 47]. Most of the issues discussed are applicable to other parts of the anatomy.

An unrestricted probe has six DOFs (degrees of freedom) of movement, three translational and three rotational. The names given to the three axes of the 2-D image slice are shown in Figure 1. The most common approach to scanning a volume with 2-D image slices

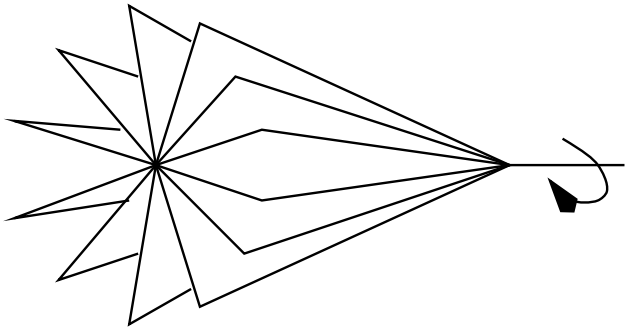

Figure 2: Cone-shaped swept volume. The volume is produced by rotating the probe about an axis in the axial direction.

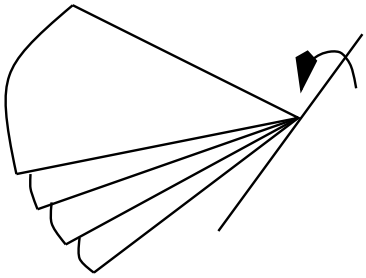

Figure 3: Fan-shaped swept volume. The volume is produced by rotating the probe about an axis in the lateral direction.

is to constrain the movement of the probe to one DOF. As the probe is moved with one DOF, the scan plane sweeps through a volume. The size of the swept volume is determined by the area of the image in the scan plane and the range of the DOF. Three different shapes of swept volumes can be produced by scanning with this method.

The first method is to rotate the probe about an axis in the axial direction. This method produces a cone-shaped swept volume (assuming that the 2-D image is sector shaped), as shown in Figure 2. The second method is to rotate the probe about an axis in the lateral direction. This method produces a fan-shaped swept volume, as shown in Figure 3. Finally, the probe can be translated along an axis in the elevational direction. A prismatic swept volume, shown in Figure 4, is produced.

There are advantages and disadvantages of each of these methods. The main advantage of the prismatic volume set is the equal intervals between the 2-D image slices that produce uniform inter-frame spacing. The fan-shaped data set has inter-frame spacing which varies

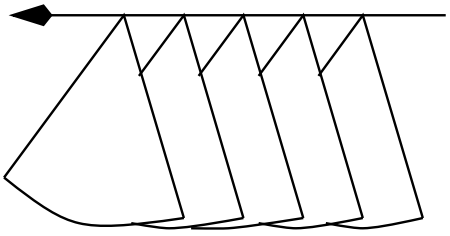

Figure 4: Prism-shaped swept volume. The volume is produced by translating the probe along an axis in the elevational direction.

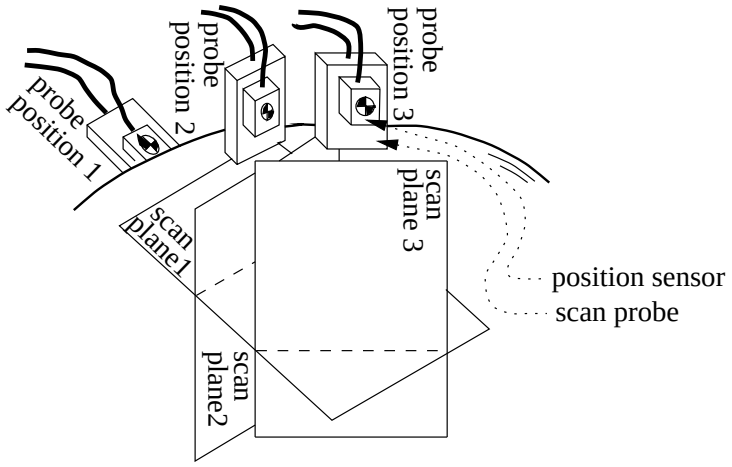

Figure 5: 3-D free-hand ultrasound imaging. The physician is allowed to move the probe as in a normal ultrasound examination while the position sensor measures the position and orientation of each scan plane.

axially. The cone-shaped data set has inter-frame spacing which varies laterally with the largest distance at the maximum axial and lateral positions. The main advantage of the cone-shaped and fan-shaped data sets is the use of a small acoustic window for the entire volume sweep. Acoustic windows on the human body are small and sometimes difficult to obtain. The translational motion limits the applicability of prismatic data acquisition. Prismatic sampling has been used successfully, however, in intravascular ultrasound [41]. Images perpendicular to the blood vessel are taken while moving a catheter-probe combination through the vessel. This method yields a prismatic set of stacked blood vessel cross-sections.

Limitations common to all three methods are the relatively low acquisition rate and the need for a specialised (and therefore expensive) probe. Furthermore, each position in the volume of interest can be sampled from only one direction. Therefore, compounding cannot be used to reduce noise and artifacts. Despite these limitations, most major manufacturers are developing prototype systems using these methods.

The simplest approach to 3-D reconstruction from multiple 2-D image slices is to allow free motion of the probe. This is called 3-D free-hand imaging and is depicted in Figure 5. The motion of the probe is unrestricted and the six DOFs are measured by a position sensor. With a free-hand approach, the scan planes through a region are chosen by the trained physician in a manner, for example, consistent with a normal examination. The scan planes may intersect each other so that the same region in space may be sampled from more than one probe direction (also called look direction). Gaps may also remain where no samples were taken.

There are several advantages of a free-hand imaging approach compared to swept-volume approaches. The first is that parts of the imaged region that are shadowed by strong echogenic regions (such as bone-tissue and air-tissue interfaces) will be filled-in by multiple look directions. Several artifacts — any part of the image that is not indicative of the structure imaged — can also be reduced by compounding multiple images.

Another benefit is the low cost, a welcome feature in an era of health-care cost reductions. 3-D imaging capability can be added to existing diagnostic ultrasound machines of most any type. Only a position sensor and a computer containing a video card are required.

Some parts of the anatomy are visible from only small acoustic windows [11]. The physician has undergone years of training to locate and image correctly the regions of interest, and some anatomical features may not be imaged if only automatic sweeping of the scan plane is used. The entire volume cross section does not need to appear in every image. The physician is allowed to focus his or her attention on the regions of interest or expand the scanned volume at will during the examination. The physician therefore makes decisions about the tradeoffs among volume, spatial resolution and acquisition time during the examination. This freedom is not possible with the other 3-D ultrasound imaging methods.

Of all the 3-D imaging techniques previously described, free-hand imaging appears to have the lowest capability for near real-time imaging. Fan, cone and prism volume scanning can all be implemented by using high speed motors to move the probe. Free-hand imaging scans a volume at a speed determined by the physician's hand motion. This is less of a problem for slow-moving or stationary objects. Errors from subject motion may be partly overcome by incremental reconstruction and visualisation [40], or by gating acquisition with ECG for echocardiography [47].

In summary, free-hand imaging is a feasible 3-D imaging method with several advantages over other methods. The prerequisite for improvement of free-hand imaging is a thorough investigation into each step of the imaging process. The rest of this paper is concerned with the sources of measurement errors, reconstruction, visualisation and volumetric data analysis of free-hand imaging.

3 3-D Free-hand Imaging

3.1 Overview

The imaging process, from 2-D image acquisition through to display of the rendered data and volumetric data analysis, is shown Figure 6. This diagram is meant not as an actual implementation but only as a graphical means of dividing the process into steps for the purpose of discussion.

In summary, the physician moves the probe over a region of interest. The location and orientation of the probe are measured by the position sensor as the received echoes are processed and displayed by the ultrasound system. The physician controls both the scanning and echo signal processing parameters. The video data from the ultrasound system display is converted to a 2-D image intensity array by the frame grabber. The frame grabber rate is chosen by the user up to a maximum dictated by hardware constraints. The 2-D intensity arrays representing image slices and the associated position data are stored in computer memory.

The reconstruction stage assembles a 3-D data set from the collected 2-D images and position data. The 3-D data set is most commonly a 3-D scalar array of intensity. In the visualisation stage, portions or all of the reconstructed 3-D data are displayed on the computer monitor. The volumetric data analysis stage is simply any algorithm used to extract diagnostically useful information from a reconstructed 3-D data set. An example of commonly requested data from physicians is the volume of an imaged object. Each stage of the imaging process in the diagram requires control from the user according to the type of tissue imaged and the type of final desired image and volumetric data analysis.

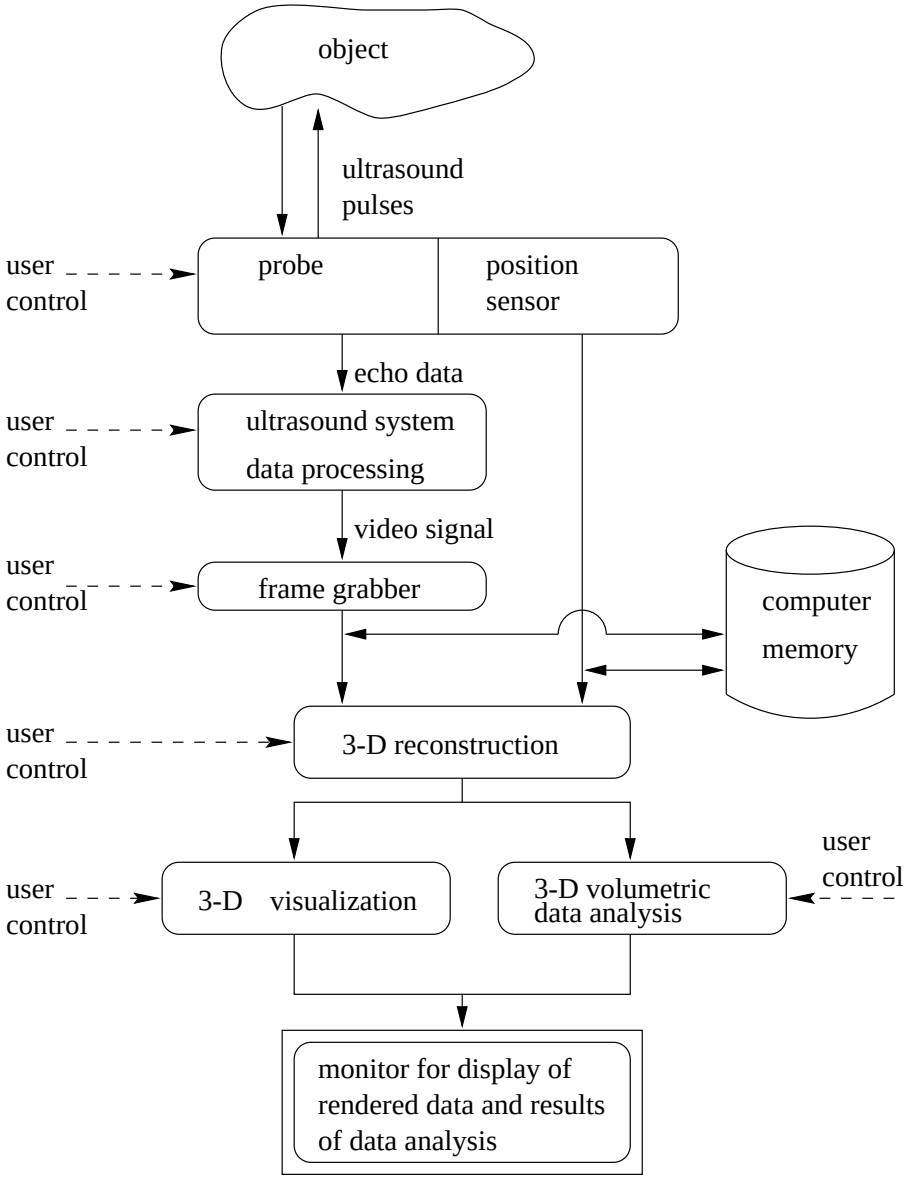

Figure 6: Overview of acquisition, reconstruction, visualisation and volumetric data analysis. The user is required to control several parameters at each stage.

3.2 Sources of Measurement Error

It is helpful to look at the physics of an idealised 2-D ultrasound imaging process: *A finite length ultrasound pulse is transmitted by the transducer in the probe. The pulse passes uniformly through the material. The scatterers that make up the material reflect part of the pulse and transmit the rest deeper into the material. The pulses are transmitted and reflected linearly. The amplitude and the timing of the pulses received by the transducer are converted into intensity-distance data. The higher attenuation of pulses returning from deeper locations is compensated by using the exact attenuation coefficient of the material. As the beam is swept through a plane, the set of collected intensity-distance data create a 2-D image.*

In reality, the interactions described in the ideal case are non-linear and include additional errors. The amplitude and frequency properties of the pulse are affected by the type of material it passes through in a non-linear manner. For example, the speed of propagation and amplitude attenuation depend on both the type of material the ultrasound pulse traverses and the pulse frequency. Both specular and diffuse reflection occur. Refraction diverts the beam direction and can cause de-focusing. Electronic noise and speckle also affect the image. Real transducers produce non-uniform beams including side lobes off of the main beam. Despite these phenomena and errors, diagnostically useful images are still obtainable by careful selection of imaging parameters.

The medical community has classified these ultrasound imaging issues into categories and sub-categories that are easy to conceptualise [11, 19, 30]. The following review of these categories is meant to give an intuitive feel for the types of problems which arise in 2-D image formation and position measurement. The problems that arise in 2-D imaging have a direct impact on the quality of the 3-D reconstructions. The categories are not independent nor comprehensive, but cover the most significant effects. The description of these categories demonstrates that the interaction of ultrasound in biological materials is not fully understood and is an active area of research [16, 17, 27].

3.2.1 Sonographic Artifacts: Propagation Group

Axial Resolution

Axial resolution is the minimum distance between objects positioned axially that are imaged individually. A first approximation is the spatial pulse length divided by two [24, 30, 61]. Typical values are shown in Table 1. For a given frequency, the axial resolution varies slightly throughout the image because of changes to the pulse shape as it travels through the medium [24, 42, 61]. Despite changes in pulse shape, the first approximation is close to the real axial resolution. A number of ultrasound systems have demonstrated axial resolutions similar to the Table 1 values [11].

Table 1 also shows that the physician is faced with the tradeoff between imaging depth and axial resolution. High frequency pulses give high axial resolution but low imaging depth because of higher attenuation [11, 30, 61].

Lateral Resolution

Lateral resolution is the minimum distance between objects positioned laterally that are imaged individually. 2-D images are produced by steering a beam through the scan plane. For example, older probes use a circular disk transducer which produces a narrow beam. The disk is rotated about an axis in the elevational direction by a stepper motor to sweep out a sector shaped image in the scan plane. More modern probes use an array

Frequency (MHz)	Maximum Imaging Depth (mm)	Axial Resolution (mm)
2	300	0.77
3.5	170	0.44
5	120	0.31
7.5	80	0.20
10	60	0.15

Table 1: Approximate values of maximum imaging depth and axial resolution for typical frequencies of probes scanning human tissue [30]. These values are for two cycle pulses.

of transducer elements. The beam is steered by selective activation of a subset of elements or by selective timing of the elements' activations. In all cases, the lateral resolution may be approximated as the beam width in the scan plane [30, 61]. The beam width is usually taken as the 6dB width of the amplitude distribution.

The beam shape depends on the type of probe used, the intensity of the ultrasound pulses and the properties of the material the pulse traverses. The beam width is a function of axial distance and has a minimal cross-sectional area at the focal point.

For fixed-focus circular disk transducers, shown on the left in Figure 7, as the effective aperture of the transducer increases the main beam becomes smaller and more cylindrical [34, 61]. Lateral resolution is therefore improved. The resolution for a focused circular transducer is given as the diameter of the focal spot and can be calculated [27] from the focal length (F), the wavelength (λ) and effective aperture (2a):

lateral resolution = w = 1.22 F λ/2a

For a probe with an array of transducers, shown on the right in Figure 7, it has been shown that diameter of the focal spot can be calculated [62] from:

lateral resolution = w = F λ/a

The lateral resolution calculated with this equation is only valid down to a minimum of twice the pitch (2p) of the array. Electronic focusing acts by delaying contributions from the central elements, which is equivalent to moving the centre elements back in the effective aperture. The focal length is therefore taken as the distance from the *edge* of the aperture.

For F = 50 mm, λ = 0.44 mm (3.5 MHz wave in average human tissue) and 2a = 20 mm (the transducer effective diameter) the beam widths are calculated as 1.3 mm and 2.2 mm for the disk and array transducers respectively. In vivo and in vitro experiments on circular disk beam widths verify the calculations [34]. Advanced techniques and higher frequency probes produce higher lateral resolutions.

For single focus transducers, the lateral resolution is worse away from the focal region. Lateral resolution is therefore dependent on the position of the objects in the image. Probes using multiple focuses help to maximise lateral resolution throughout the image.

Elevational Resolution

Resolution is also limited in the elevational direction because of the finite beam width in the elevation plane. Alternatively, artifacts are introduced onto the scan plane by off-plane (but still in beam path) objects [11, 19, 30]. The finite elevation thickness also blurs sharp material boundaries, as shown in Figure 8.

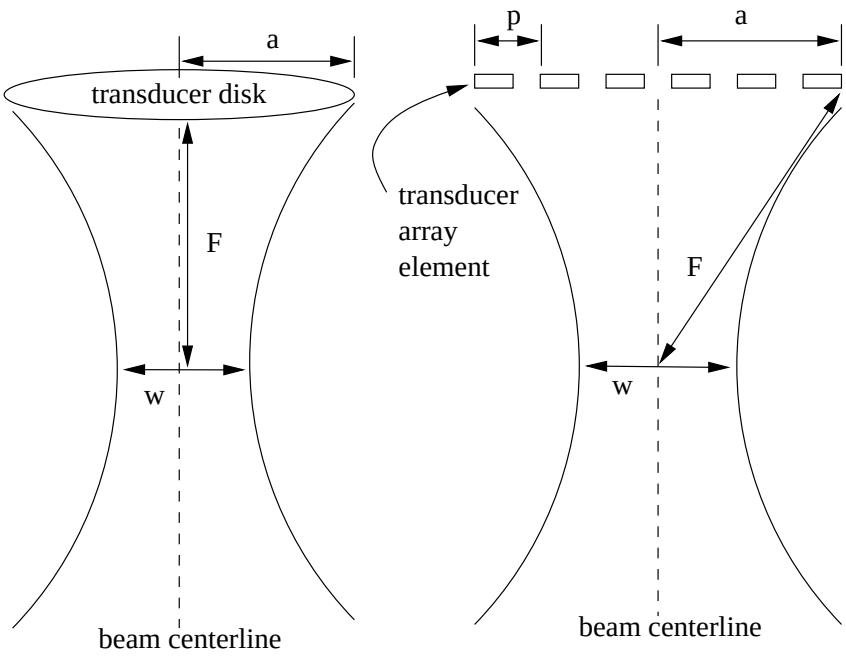

Figure 7: The lateral resolution of circular disk and linear array transducers is approximated as the width of the beam at the focus.

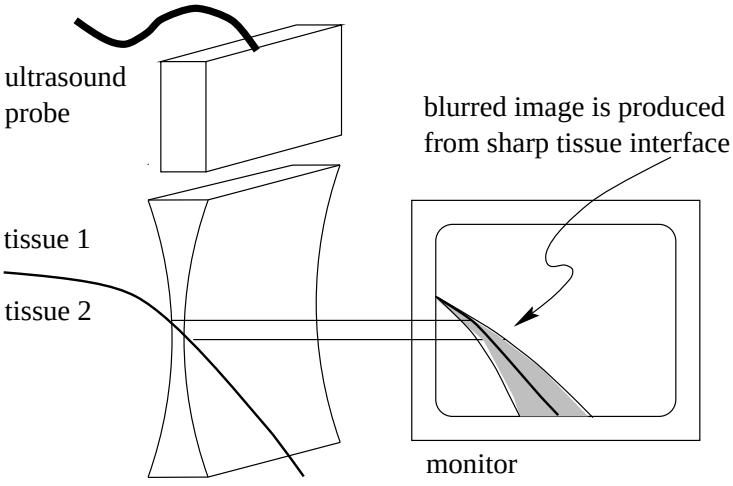

Figure 8: Reconstructed image of sharp interface becomes blurred because of finite width of beam.

The beam's elevational width varies with the type of probe, imaging depth and material properties. On current ultrasound machines it is typically focused by a fixed lens. A study of the elevational beam width of a 3.5 MHz phased array sector probe gives beam widths ranging from 10 mm at a 2 cm depth, to 6 mm at a 12 cm depth, with a minimum of 3 mm at the 8 cm focus [49].

Speckle

The grainy appearance of 2-D ultrasound images is due mainly to speckle. Speckle has a distinguishable onion-like structure in 3-D [6]. The speckle phenomenon results from the constructive-destructive interference of the coherent ultrasound pulses back-scattered from the tiny multiple reflectors that constitute biological materials. Speckle typically has the unfortunate aspect of falling into the high sensitivity region of human vision to spatial frequency. The frequency spectrum of speckle is also similar to the imaging system modulation transfer function [61]. Speckle can therefore obscure the diagnostically important information.

Speckle patterns become uncorrelated if a region is viewed from different look directions or with different frequencies [61]. The most common methods of speckle reduction involve compounding images taken from different look directions far enough apart to de-correlate the speckle pattern [35, 48, 59, 61]. The expected increase in signal to noise ratio is in proportion to the square root of the number of compounded images [48, 59, 61]. There is a fundamental trade-off between speckle reduction and spatial resolution: multiple compounding always results in a decrease in spatial resolution [61]. In other words, an object in one image will not be coincident with the same object in another image, taken from a different look direction, because of spatial positioning errors. Compounding images from different look directions reduces speckle but blurs the object.

The necessary condition for de-coherence between two interrogating beams is that the points of intersection of their axes with the transducer should be separated by a minimum distance. Translation of forty percent of one transducer aperture diameter was experimentally derived for de-correlation in the beam focus [59] of a phased array. If less than this distance is used then weighted averages of multiple images are required for optimal spatial compounding [48]. Using closely spaced look directions would minimise spatial resolution loss, so weighted compounding of only closely spaced images is possibly the best tradeoff between speckle reduction and spatial resolution.

Temporal compounding is a common alternative to spatial compounding. Temporal compounding is often implemented using a variable persistence of the images on the ultrasound system monitor [11, 61]. As the probe is slowly moved, the images are partly averaged and speckle is reduced. Spatial resolution is also reduced. Other approaches include frequency compounding and a variety of signal processing techniques. However, many of these produce images with insufficient spatial resolution [61].

The trade-off between spatial resolution and speckle reduction will likely be determined by the intended purpose of the final reconstructed 3-D data set. For example, removal of speckle for easier visualisation of object surfaces likely involves the loss of surface detail.

Reverberation

Reverberation is the phenomenon of multiple reflections within the material or from the sensor face. Acoustic boundaries transmit part of the pulse and reflect the remainder. The probe face can be a strong reflector, as shown in Figure 9, because of the large difference in acoustic impedance between the probe face and the tissue [11, 19, 30]. Matching layers, such as coupling gel, reduce this effect.

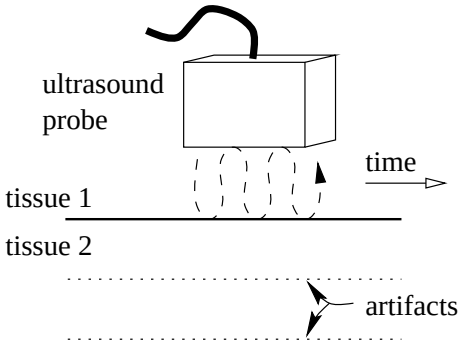

Figure 9: Reverberation off the sensor face produces multiple artifacts.

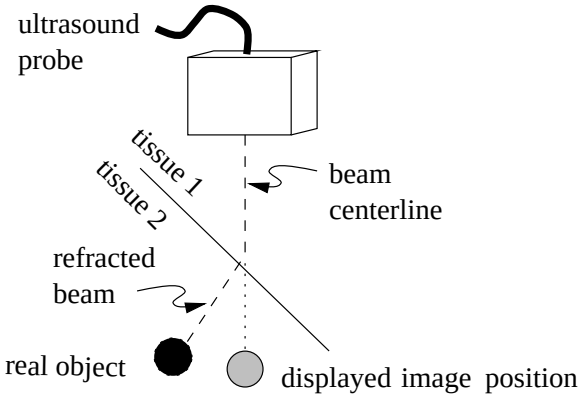

Figure 10: Incorrect placement of object due to refraction of beam direction.

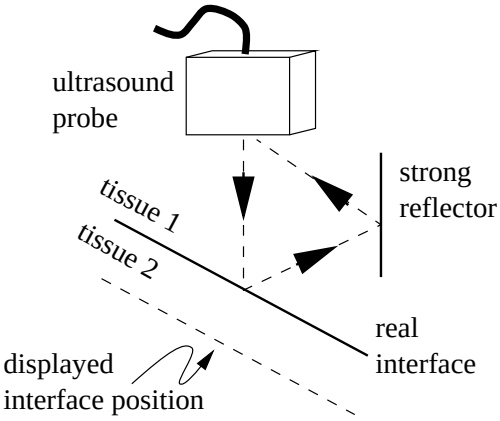

Figure 11: Incorrect placement due to lengthening of beam path through multiple reflections.

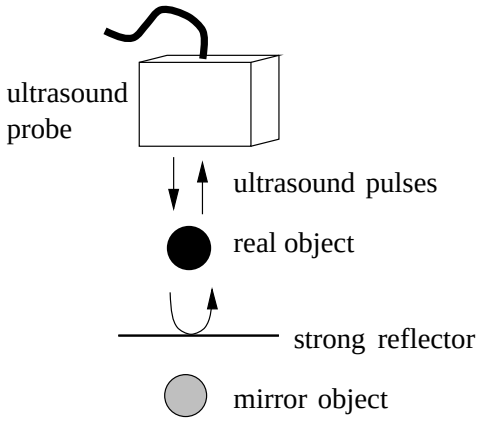

Figure 12: Mirror image of the real object is produced by a nearby strong reflector.

Refraction

Refraction of the beam, as shown in Figure 10, changes the beam direction and results in improper position placement of an object in the image [11, 19, 30]. This is one of the reasons why compounding loses spatial resolution. Different look directions have different position errors and a blurred object is produced in the compounded image [61].

Multi-path Reflections

Improper positioning of an object can result from a difference in the length of the return path from the incident ultrasound path. An example is shown in Figure 11. In general multi-path blurs the object boundaries in the image [19, 30].

Mirror Image

Strong reflectors near an object act as a mirror to produce an artifactual mirror image [11, 30]. The mirror image artifact appears below the reflector, as shown in Figure 12.

Side Lobes and Grating Lobes

Side lobes are minor beams of sound travelling in a direction different from the main beam [11, 19, 30]. Side lobes of sufficient intensity produce echoes from objects off the main beam path, as shown in Figure 13.

Grating lobes are lobes produced by array transducer probes [19, 30]. As the number of array elements increases and the gaps between the elements decrease, the size of the lobes decreases [62]. The lobes also change shape as the beam is steered. Apodization – unequal weighting of the transmit or receive elements' signals – reduces grating lobes but increases the width of the main beam [61]. Lobes off of the main beam are present to some degree in all ultrasound systems.

Comet Tail and Ring-down

Comet tails and ring-down appear as artifactual tails below certain objects [11, 30]. A comet tail is a series of closely spaced reverberation echoes. The reverberation echoes are produced by strong reflectors nearby an object, as shown in Figure 14.

Ring-down is related to a resonance phenomenon associated with gas bubbles [30]. Ring-down creates a fuzzy tail without distinguishable features.

Speed Errors

Ultrasound systems typically use a propagation speed of 1540 m/s to calculate the

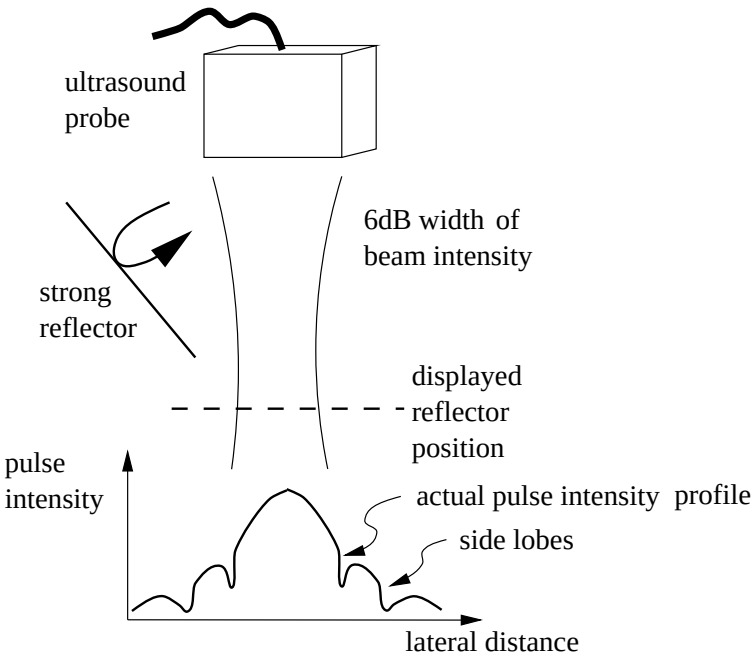

Figure 13: Side lobe of beam produces an artifact from strong reflectors off of the main beam path.

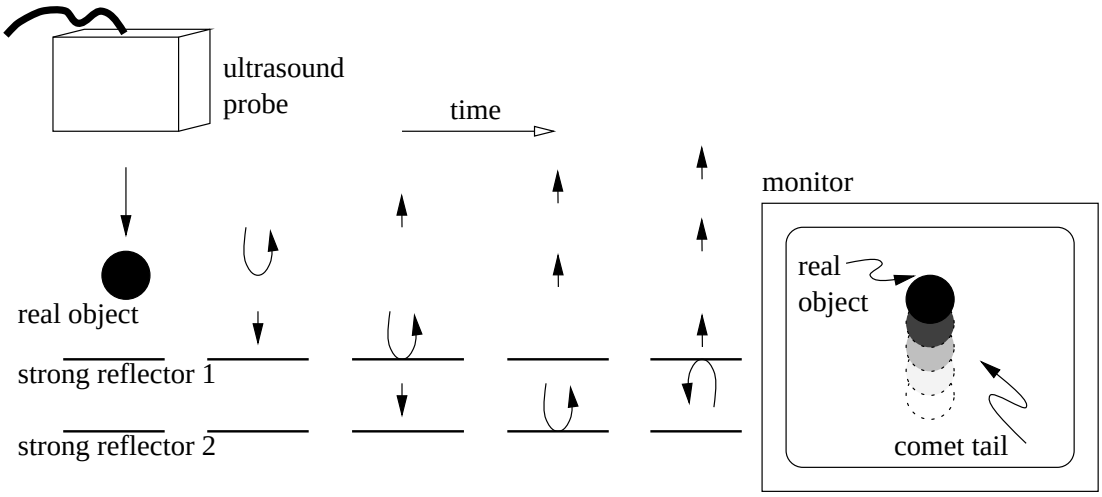

Figure 14: Comet tail is produced from strong reflectors nearby an object. The reflectors produce multiple reverberations that follow the original echo back to the probe.

distance to reflectors from the time it takes for the echoes to return [30]. The value of 1540 m/s is the average of the propagation speed of ultrasound in human tissue. The actual propagation speed varies with the type and temperature of the tissue. For example, body temperature fat has a propagation speed of 1440 m/s and muscle has a propagation speed of 1570 m/s [30]. Objects located in materials with a propagation speed greater than 1540 m/s will be incorrectly placed too near the probe and vice versa [19, 30]. The shape of objects within the image will also be distorted.

A less significant effect is the dependence of propagation speed on frequency for a given material. For example, the ultrasound propagation speed in body temperature human brain changes from 1561 m/s to 1566 m/s as frequency is increased from 1 to 10 MHz [27].

These combined effects cause misplacement of objects in the image. Speed errors contribute to the loss of spatial resolution when compounding.

Range Ambiguity

The maximum depth that can be imaged unambiguously is determined by the pulse repetition frequency [30]. In other words, a second pulse should be transmitted only after all echoes are received from the first pulse. Otherwise, echoes from the first pulse will appear as artifacts in the second image. In reality, this is not a problem if the physician chooses the correct probe and pulse repetition frequency for the desired imaging depth. The trade-off that occurs is between range ambiguity and frame rate.

3.2.2 Sonographic Artifacts: Attenuation Group

Contrast Resolution

Contrast resolution is the smallest detectable difference in echo intensity. The contrast resolution perceived by the physician depends on human perception, the computer monitor and the ultrasound system. The echo contrast dynamic range is usually compressed to a range most sensitive to the human eye. For example, the ratio of smallest to largest intensities of returning echoes may be 100 dB. The 256 grey level display has a dynamic range of 24 dB. The dynamic range of human perception of grey levels is approximately 20 dB [30].

A high contrast resolution benefits spatial resolution. Individual detection of two closely spaced objects depends on measuring the amplitude of two separate but partly overlapping distributions, as shown in Figure 15. Contrast and spatial resolution are also affected by noise and artifacts that obscure or distort objects in the image.

Furthermore, contrast resolution is also dependent on the processing of the echo data. Some ultrasound systems allow manual control of the post-processing curve which relates echo levels to grey levels on the display monitor[58]. A steep slope of the curve produces very different grey levels from echo levels of near equal amplitude. The physician can improve contrast resolution over a portion of the dynamic range at the expense of the remaining dynamic range. Figure 16 shows that a steep slope over one portion of the echo level dynamic range requires a shallow slope over the remaining portions. For example, consider two nearby objects which generate near equal echo amplitudes. The objects are more likely to be detected individually if the echoes fall in the dynamic range covered by the steep slope of the curve. The contrast resolution is reduced if the echoes fall into a range covered by the shallow slope of the curve. In general, image processing techniques that enhance visibility of features in 2-D images also enhance visibility in 3-D.

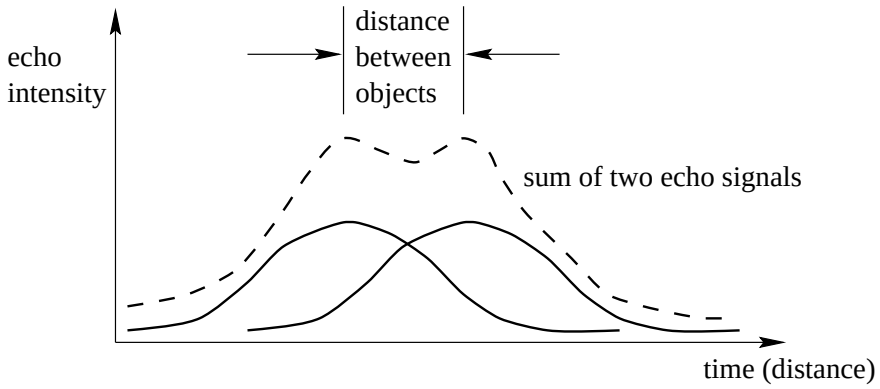

Figure 15: For two echoes to be distinguishable individually, the depression between the two peaks of the summed signals must be detectable — an issue of contrast resolution.

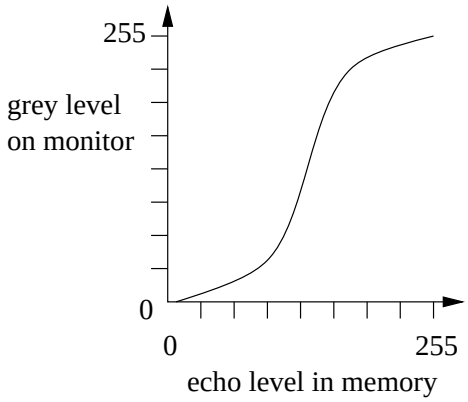

Figure 16: Post-processing curve between echo level in memory and grey level on the display monitor. A steep slope improves contrast resolution. A steep slope over one range of echo levels is accompanied by a shallow slope over the remaining ranges.

Shadowing

Shadowing occurs when ultrasound pulses are greatly attenuated behind strong reflectors such as bone or lung [11, 19, 30]. The only way to produce image data from the shadowed region is to view it from another look direction that is not blocked by the strong reflector.

Shadows can also be caused when the beam is refracted and de-focused by certain tissue interfaces. If the beam is sufficiently de-focused, then the returning echoes are too weak to be detectable. The image therefore exhibits a shadowed region below the interface [11, 30].

Enhancement

Enhancement is the opposite of shadowing. Some tissues or fluids (such as amniotic fluid) are weakly echogenic. Regions below weakly echogenic tissues or fluids return echoes of greater than expected amplitude [11, 30]. Enhancement can also occur because of the inherent higher intensities near the beam focus [30]. TGC (time-gain compensation) amplifies distant echoes more than proximal echoes to account for the greater attenuation of distant echoes. Manual control of the TGC curve is required because of the variation in attenuation coefficient among different types of materials. The physician typically controls the shape of the TGC curve to reduce the enhancement phenomenon.

Name of Sensor	Accuracy (mm)	resolution (mm)	notes
Polhemus AC magnetic sensor	0.8	0.38	manufacturer ratings for Fastrak (at 760mm)
Bird DC magnetic sensor	2.54	0.76	manufacturer ratings for Flock of Birds
FARO Mechanical Arm	0.4	0.25	manufacturer ratings for SpaceArm
OPTOTRAK LED system	0.15	0.01	independent tests for model 3020

Table 2: Position measurement ratings of several common sensors.

Electronic Noise

Electronic noise is unavoidably introduced during the analog processing of the echo pulses. The analog process may include amplification, demodulation and compression of the signals before digital conversion. Electronic noise dominates at low intensities and the noise level limits imaging depth [61].

3.2.3 Position Measurement Errors

Translation

A number of position sensors have been used to measure the location and orientation of the probe including mechanical arms [40, 45, 61] and transmitter-receiver systems. The transmitter-receiver systems use spark-gap acoustics [61], optical-video systems [45, 60] AC magnetic fields [40, 53] and DC magnetic fields [14] . A survey of some of the more accurate, recent devices is shown in Table 2. It does not include calibration errors which are typically as large as the accuracy ratings. Calibration is required to determine the offset of the position sensor to the scan plane.

Rotation

Rotational errors are separated from translational errors because of the lever-arm effect that rotational error has on the accuracy in the scan plane. The position sensor is located on or near the probe. Therefore, angular errors will result in larger positional errors at distant positions in the scan plane. The farther the sensor is from the top of the scan plane, the greater the errors will increase by the lever-arm effect. Mis-calibration will again cause systematic measurement errors. Some typical values of angular accuracy and resolution are given in Table 3.

The combination of position measurement errors, imaging system errors and human perception limitations were investigated [14] to determine the error in locating a crossed string target. A Bird DC magnetic sensor and a 7.5 MHz linear array scanner was used. The combined RMS errors in locating the point target in water were between 2.1 mm and 3.5 mm.

Name of Sensor	Accuracy (degrees)	resolution (degrees)	notes
Polhemus AC magnetic sensor	0.15	0.025	manufacturer ratings for Fastrak (at 760mm)
Bird DC magnetic sensor	0.5	0.1	manufacturer ratings for Flock of Birds (at 305 mm)
FARO Mechanical Arm	-	-	manufacturer ratings not available
OPTOTRAK LED system	-	-	dependent on number and placement of LEDs

Table 3: Orientation measurement ratings of several common sensors.

Latency and Acquisition Rate

A moving sensor will provide readings of past positions due to latency. Latency is the time between movement of the sensor and change in the recorded position. For example, a 4 ms latency (Polhemus manufacturer rating) for a probe moving at 0.1 m/s gives a position error of 0.4 mm. Others have estimated the latency of the Polhemus sensor as high as 100 ms [15].

A low acquisition rate requires the physician to move the probe slowly to avoid large gaps between images. For example, 1 mm spacing of images acquired at 120 Hz (rate of Polhemus sensor) requires the probe to move at 120 mm/s. The acquisition rate of the frame grabber may also be low. The smaller of the two acquisition rates determines the minimum time required for a complete scan.

3.2.4 Tissue-related Errors

Movement

Movement of the subject during a prolonged examination results in significant errors when constructing a 3D dataset from 2D images. A moving subject produces mis-registered slices in the 3-D data set. There are relatively few clinical applications where the imaged anatomy is near-stationary (one good example is the examination of a breast immobilised in a mammographic unit). Motion-related errors can be minimised by increasing the scan rate, so reducing the examination time. ECG gating and temporary suspension of breathing can also reduce motion artifacts. Nevertheless, the design of systems which are tolerant to tissue movement remains one of the major challenges in 3D ultrasound.

Anisotropy

It has been found that some tissues contain anisotropic scatterers [35] which only appear from certain look directions. Free-hand imaging has a clear advantage over other 3-D imaging methods by allowing the physician to choose the look directions. The physician has been trained to find the look directions that provide high visibility of anatomical features. The look directions with highest visibility are not necessarily used by the automatic volume sweeping methods discussed in Section 2.

Aliasing

Reconstructing 3-D data sets from 2-D image slices is analogous to sampling a time-varying signal at various times. Aliasing is a problem that occurs when the sampling rate is too low to reproduce the high frequency components of the signal. The Shannon-Nyquist sampling theorem states that the sampling must be carried out at a frequency which is at least twice that of the signal's highest frequency component. Anti-aliasing filters are typically low-pass filters used on the original signal before it is sampled and reconstructed. Because anti-aliasing filters cannot be used on human anatomy, reconstruction will only be free of aliasing errors if the object does not contain spatial variations at a frequency higher than the spatial sampling of the 2-D image slices. Fortunately, this is likely the case for many anatomical objects such as the liver and kidney. However, it may not be the case for the vascular system.

3.3 Problems and Issues in 3-D Reconstruction

The following discussion is based on the implicit assumption that the reconstructed 3-D data set is in the form of a 3-D scalar array of intensities, also called a cuberille. Each element of the array corresponds to a certain volume element called a voxel. Other 3-D data representations could include gradients or be simply a modified set of the original 2-D images.

Reconstruction of a 3-D array typically involves equating each voxel value to the value of the 2-D image pixel which intersects the voxel. Some voxels will be intersected and others remain empty. The values of voxels that are intersected more than once are calculated by compounding the pixels that intersect the voxel. Compounding is not restricted to simple averaging. Maximum value or newest value methods are two of several alternatives.

A number of issues must be resolved when combining 2-D images into a 3-D data set. The issues can be divided into two main categories: intensity non-uniformities and spatial non-uniformities among 2-D images.

3.3.1 Intensity Non-uniformity

The average intensity and intensity distributions for a given 2-D image depend on both the ultrasound system creating the images and the properties of the material. Single transducer probes and array transducer probes produce images of different intensities and window shapes. Image intensities also depend on the geometry, size, pulse frequency, pulse amplitude and the type of focusing of the probe. Artifacts, such as shadowing and enhancement, also affect the average image intensities.

The physician normally controls several parameters on the ultrasound system [11, 19, 30, 58, 61]. The following is a list of the important parameters controllable by the physician. Not all of these parameters are adjustable on all ultrasound systems, nor can they all be controlled independently.

- probe used
- time-gain compensation
- overall gain
- display persistence

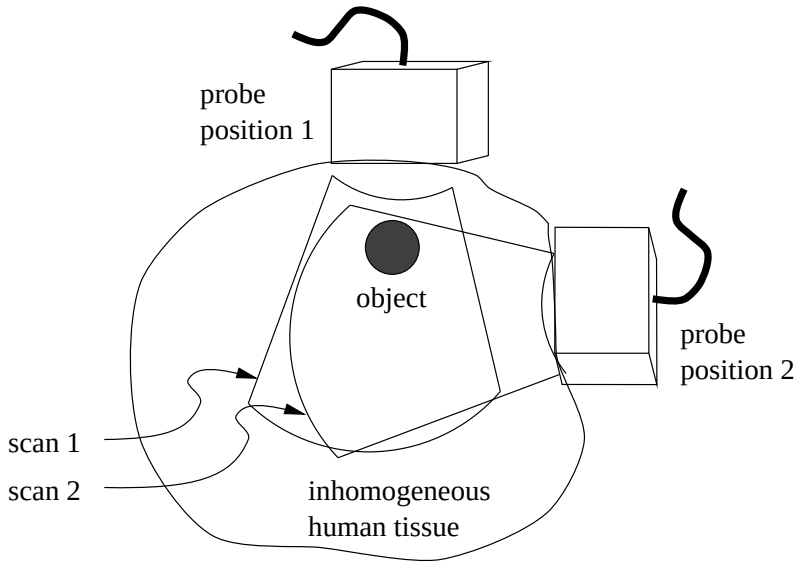

Figure 17: Two scans are taken of the same object from different look directions.

- dynamic range of display
- frame rate
- echo-enhance scanner option (image processing option)
- echo-filter scanner option (image processing option)
- post-processing curve
- focus parameters
- scan line density
- display reversal (display can be reversed upside down)
- pulse amplitude
- scan angular range
- enlargement of image (zoom)

The most obvious method to fill a consistent 3-D data set with 2-D image slices is to collect the slices with a single probe and constant user-controlled parameters. This method will improve, but not ensure, consistency of the intensities of a set of 2-D image slices. This is best demonstrated by the example depicted in Figure 17. The object is imaged from two look directions in the same plane. The object would ideally have the same intensity distribution in the two images. For examination of human patients, the objects in scans 1 and 2 are likely scanned through different tissue types and geometries. The TGC curve is usually adjusted for uniform average intensity in the portions of the image with the same tissue type. For example, if scan 2 images through a thicker layer of fat, the ultrasound pulse will be more attenuated before it encounters the object. The echo reflected from the object in scan 2 will return with a lower amplitude than if it came along a path in scan 1.

Constant TGC settings for the two scans will therefore yield two different average intensities of the object. Similarly, shadowing and enhancement artifacts will be different in the two scans.

3.3.2 Spatial Non-uniformity

The voxel dimensions should reflect the spatial resolution of the images. However, the axial, lateral and elevational resolutions are all different. Moreover, the resolutions vary within each image. A voxel size equal to the smallest resolution would minimise loss of image detail. In real free-hand systems, larger voxel sizes are used to reduce the computer memory requirements and rendering times.

Gap-filling may be required in regions of the volume which are under-sampled. The missing voxel values can be calculated by interpolation of the surrounding elements [60]. Alternatively, convolving each 2-D image with a 3-D Gaussian kernel also fills in gaps between images [40].

Using the same example in Figure 17, the two object positions, with respect to a global reference system, will not be the same in the two images because of errors in spatial positioning. The error analysis states that errors arise from refraction, propagation speed, multi-path, tissue motion and the position measurement device. The quality of the final reconstruction improves if some of the errors are removed before compounding. Transformation of an image to match the other images is called **registration**. The type of transformation is based on a model of the spatial positioning errors.

For example, registration has been attempted in previous research to account for misalignment of adjacent 2-D images acquired by a swept volume technique [7]. The transformation only allowed planar translations of the images. The parameters of the transformation were determined by a Fourier transform of the cross-correlation function between adjacent images. The error analysis in Section 3.2 indicates that planar translation will not account for all of the measurement errors in free-hand scanning. Ideally, registration accounts for all sources of spatial positioning errors to produce a consistent reconstruction. In reality, only the sources of error that can be conveniently modelled can be included in the registration. It is important, however, that registration not alter or remove diagnostically useful information.

3.4 Visualisation of 3-D Data

The important features of a good image are clear detail, obvious tissue texture, visualisation of detail at depth, and low noise [11]. A bad image has no visible diagnostic information. The goal of the visualisation method is to produce a good image that accurately represents the 3-D reconstruction data.

There are several common methods of displaying the 3-D intensity array data. The most conventional is to view slices of the data set in a manner similar to normal 2-D scanning [11, 23, 61]. The difference is that slices from any look direction may be displayed, whether or not any of the original 2-D images were taken in that particular look direction. This is called **any-plane slicing**. It may be of interest to a physician to view the anatomy from directions unobtainable by normal 2-D scanning. For example, if measurement of the size of an object is desired, it is helpful to be able to view the object from three perpendicular look directions. Another difference of any-plane slicing compared to normal real-time 2-D scanning is that the 3-D data set may have been substantially processed during reconstruction. 3-D image processing may improve certain aspects of the image (speckle reduction, for example) and

degrade others (such as spatial resolution). Therefore, even if a desired look direction is chosen as one of the original look directions, the displayed image may not be the same as the original 2-D image.

Sparsely sampled data presents a problem when the chosen look direction is perpendicular to the original scan directions. Only the edges of the original scan slices will be visible. Interpolation helps but introduces a large proportion of data not actually imaged. The displayed image is therefore mostly interpolated data and of diminished diagnostic use [60].

Other visualisation methods include surface fitting and volume rendering [1, 4, 50, 51, 63]. Surface fitting may simply involve selecting isovalue voxels and displaying by back-to-front traversal, ray tracing or surface reconstruction. Surface reconstruction typically involves fitting patches to constant value contour surfaces in the 3-D data set according to a certain threshold. The surface can then be displayed with a particular viewpoint and lighting source. When the physician is interested in viewing only one tissue type, the determination of thresholds becomes related to tissue classification, a difficult problem with ultrasound data [11, 51, 61].

The advantages of surface fitting include the relatively fast rendering and reduced storage requirements once the surfaces are determined [21, 36, 63]. The drawbacks are that surface fitting algorithms used on ultrasound data often produces artifactual surface pieces and incorrect handling of small features.

Volume rendering maps the 3-D data set directly into screen space. Ray casting is the most common method of volume rendering. Ray casting involves casting of rays from each screen pixel through the 3-D data set. The colour and opacity of each voxel that the ray passes through are blended until the opacities reach unity or the ray exits the rear of the volume of data. The blended opacities and colours are then displayed by the screen pixel. The procedure then repeats for all screen pixels. It has been suggested that summed voxel volume rendering reduces speckle in the rendered image [61].

The disadvantages of volume rendering is that the entire 3-D data set must be used for each new viewpoint, making the rendering slow [21, 36, 50, 63]. Additional visual cues, such as shadowing and reflections, may also be added. The advantages are that artifact generation is reduced compared to surface fitting, and that the entire 3-D data set is used for each 2-D display.

Standard surface-fitting and volume rendering methods used successfully in magnetic resonance and computed tomography imaging have limitations when used with ultrasound data. Ultrasound images have several unique properties [46]: significant amount of noise and speckle; lower dynamic range; blurred boundaries several voxels wide; boundaries with varying grey levels caused by variations in surface curvature and orientation; shadowed regions; intensity non-uniformities; spatial non-uniformities; and non-uniform spatial resolution. This list demonstrates that new visualisation techniques, possibly involving filtering and segmentation, are required [46]. Fetal imaging is an exception that is easier to render because the fetus is surrounded by a weakly echogenic amniotic fluid. It is therefore easier to isolate the fetus from surrounding tissue and noise [52].

3.5 Volumetric Data Analysis

The physician is often interested in obtaining quantitative information from the ultrasound data. On typical 2-D ultrasound systems cursors can be controlled by the physician to measure distances in the image. For example, fetal head diameter can be measured and

used to calculate the approximate gestational age. Other diagnostically useful information is available in the 3-D data set. Physicians have indicated a need for size and volume estimates of image objects. Examples where volume or size determination is useful include prostrate volume [43], ventricle volume [33, 39, 47], kidney volume [22], blood vessels [28] and fetus and placenta volumes [11].

Segmentation and modelling of the data [2, 5, 8, 9, 10, 12, 13, 25, 26, 53, 54, 55, 56, 57] are also important research areas. Texture classification — measurement of the statistical properties of material to aid diagnosis — is another active research area [11, 61].

4 Conclusions

3-D free-hand ultrasound imaging, compared to other 3-D ultrasound imaging techniques, has a number of advantages: the ability to fill in shadowed regions; reduction of artifacts by compounding; relatively low cost; adaptability to a wide range of ultrasound systems; flexibility in image volume size, number of image slices and acquisition time; use of existing examination techniques; and flexibility for the physician to image from any available look direction.

The disadvantages include acquisition times that are typically longer than other 3-D imaging techniques. The time required to scan a volume is determined by the speed at which the physician moves the probe. Free-hand imaging has the additional complications in reconstruction arising from the arbitrary positions and intersections of the image planes. Measurement errors also arise from sonographic artifacts, position measurement errors and tissue-related errors.

The study of the measurement errors reveals the type and size of errors that occur in free-hand imaging. The tradeoffs that exist in minimising some of the errors are also identified.

The design of a 3-D free-hand system should start with a high spatial and contrast resolution, low noise 2-D ultrasound machine. This includes choosing the highest frequency probe that can image to the desired depth. The highest frame rate that does not result in range ambiguity should be used. An accurate, calibrated position sensor with a high acquisition rate should be attached as close as possible to the probe face. Motion of the subject of the ultrasound examination should be minimised or measured for compensation during reconstruction. Finally, the acquired 2-D images should be closely spaced to reduce gaps and aliasing. If time permits, the volume of interest should be scanned with overlapping image slices taken from a variety of look directions.

The basic criteria for reconstruction are that it should be automatic and incremental. It should be automatic because it will make the system easier to use and therefore more likely to be accepted by physicians in a clinical setting. An incremental reconstruction means that the 3-D data set is updated as each new image is acquired. Incremental reconstruction is desired because it will reduce storage requirements and can make the system more interactive by allowing incremental rendering during the examination.

Most reconstruction methods in previous research have used solely the position data to fill a 3-D scalar array with 2-D image slices. The error analysis demonstrates that these methods may produce inconsistencies in the 3-D data set. Research should first focus on constructing a consistent and accurate 3-D data set by accounting for the measurement errors.

As discussed in Section 3.3, both the spatial and intensity non-uniformities must be addressed. Registration of the 2-D image slices with the 3-D data set must account for the

spatial non-uniformities and compounding the overlapping data must account for intensity non-uniformities. It is proposed that these two issues are separable: the reconstruction algorithm can perform registration of the 2-D image slices and then compound each image slice with the 3-D data set.

The following is a brief description of a prototype incremental, automatic reconstruction algorithm. A 3-D scalar array is chosen as the form of the reconstructed data set. All voxel values are set initially to zero. As each new 2-D image is acquired, each voxel that is intersected by the image scan plane is set to the value of the 2-D image pixel which intersects the voxel. When an image slice intersects previously filled voxels, registration aligns the image slice to the existing data in the 3-D array. It is proposed that the registration be based on matching features in the 2-D image with features in the 3-D array. The registered image is then added to the 3-D scalar array using a form of compounding. The original 2-D image data may now be discarded.

After a substantial portion of the 3-D data set is filled, it can be displayed on a computer monitor. Several different visualisation methods can be used including volume ray-tracing, surface rendering and any-plane slicing [21, 61]. Incremental rendering is also possible [40].

The visualisation and volumetric data analysis stages will benefit from improvements in the reconstruction stage. The overriding concern of each stage is the desire not to lose detail nor features and not to introduce artifacts. A beautifully smoothed, filtered and highly processed image is useless to a physician attempting to make a correct diagnosis of possible pathology.

General acceptance of 3-D free-hand imaging by physicians will occur if accurate and diagnostically useful 3-D renderings are available automatically, during or shortly after the examination. Solutions to some of the problems currently limiting 3-D free-hand imaging, such as registration and compounding, are feasible. Artifacts common to all ultrasound images, such as speckle, cannot likely be eliminated completely, yet 2-D ultrasound imaging is in general use with these artifacts. It is not unreasonable to suggest 3-D images with a small number of artifacts will still be clinically useful. Furthermore, as manufacturers continue to improve the quality of the 2-D images, the 3-D images will similarly improve. The inherent advantages of the free-hand imaging technique, with continued improvement, ensure that it will become a valuable tool for physicians in the future.

References

[1] N. Ayache. Medical computer vision, virtual reality and robotics — promising research tracks. In *Proceedings of the 6th British Machine Vision Conference*, pages 1–25, Birmingham, 1995.

[2] N. Ayache, I. Cohen, and I. Herlin. Medical image tracking. In A. Blake and A. Yuille, editors, *Active Vision*, pages 285–302. MIT Press, Cambridge, MA, 1992.

[3] F. G. Balen, C. M. Allen, J. E. Gardener, N. C. Siddle, and W. R. Lees. 3-dimensional reconstruction of ultrasound images of the uterine cavity. *The British Journal of Radiology*, 66(787):588–591, 1993.

[4] J. C. Bamber, R. J. Eckersley, P. Hubregtse, N. L. Bush, D. S. Bell, and D. C. Crawford. Data processing for 3-D ultrasound visualization of tumour anatomy and blood flow. *Proceedings of SPIE — The International Society for Optical Engineering*, 1808:651–663, 1992.

[5] H.H. Barrett and A.F. Gmitro, editors. *Lecture Notes in Computer Science 687: Proceedings of 13th Annual Conference of Information Processing in Medical Imaging.* Springer-Verlag, Germany, 1993.

[6] G. R. Bashford and O. T. von Ramm. Speckle structure in 3 dimensions. *Journal of the Acoustical Society of America*, 98(1):35–42, 1995.

[7] M. Belohlavek, D. A. Foley, J. B. Seward, and J. F. Greenleaf. 3D echocardiography: reconstruction algorithm and diagnostic performance of resulting images. *Proceedings of SPIE — The International Society for Optical Engineering*, 2359:680–692, 1994.

[8] A. Blake, R. Curwen, and A. Zisserman. A framework for spatiotemporal control in the tracking of visual contours. *International Journal of Computer Vision*, 11(2):127–145, October, 1993.

[9] A. Blake, M. Isard, and D. Reynard. Learning to track curves in motion. *Proceedings of the IEEE Conference on Decision Theory and Control*, 4:3788–3793, 1994.

[10] A. Blake and A. Yuille, editors. *Active Vision.* MIT Press, Cambridge, MA, USA, 1992.

[11] F.A. Chervenak, G.C. Isaacson, and S. Campbell. *Ultrasound in Obstetrics and Gynecology.* Little, Brown and Company, Boston, MA, USA, 1993.

[12] I. Cohen, N. Ayache, and P. Sulger. Tracking points on deformable objects using curvature information. In *Proceedings of the 2nd European Conference on Computer Vision*, pages 458–466, 1992.

[13] I. Cohen, L. D. Cohen, and N. Ayache. Using deformable surfaces to segment 3-D images and infer differential structures. In *Proceedings of the 2nd European Conference on Computer Vision*, pages 648–652, 1992.

[14] P. R. Detmer, G. Bashein, T. Hodges, K. W. Beach, E. P. Filer, D. H. Burns, and D.E. Strandness Jr. 3D ultrasonic image feature localization based on magnetic scanhead tracking: in vitro calibration and validation. *Ultrasound in Medicine and Biology*, 20(9):923–936, 1994.

[15] S. Emura and S. Tachi. Sensor fusion based measurement of human head motion. *Robot and Human Communication - Proceedings of the IEEE International Workshop*, pages 124–129, 1994.

[16] D.H. Evans and K. Martin, editors. *Physics in Medical Ultrasound II.* Institute of Physical Sciences in Medicine, London, England, 1988.

[17] J.A. Evans, editor. *Physics in Medical Ultrasound.* Institute of Physical Sciences in Medicine, London, England, 1986.

[18] D. Fine, S. Perring, J. Herbetko, C. N. Hacking, J. S. Fleming, and K. C. Dewbury. Three-dimensional (3D) ultrasound imaging of the gallbladder and dilated biliary tree: reconstruction from real-time B-scans. *The British Journal of Radiology*, 64:1056–1057, 1991.

[19] P. Fish. *Physics and Instrumentation of Diagnostic Medical Ultrasound.* John Wiley and Sons Ltd., Chichester, England, 1990.

[20] D. Franseschi, J. A. Bondi, and J. R. Rubin. A new approach for three-dimensional reconstruction of arterial ultrasonography. *Journal of Vascular Surgery*, 15(5):800–805, May 1992.

[21] H. Fuchs, M. Levoy, and S. M. Pizer. Interactive visualization of 3D medical data. *Computer*, 22(8):46–51, August, 1989.

[22] O. H. Gilja, A. I. Smievoll, N. Thune, K. Matre, T. Hausken, S. Odegaard, and A. Berstad. In vivo comparison of 3D ultrasonography and magnetic resonance imaging in volume estimation of human kidneys. *Ultrasound in Medicine and Biology*, 21(1):25–32, 1995.

[23] M. Halliwell, H. Key, D. Jenkins, P. C. Jackson, and P. N. T. Wells. New scans from old: digital reformatting of ultrasound images. *The British Journal of Radiology*, 62(741):824–829, 1989.

[24] R. A. Harris, D. H. Follett, M. Halliwell, and P. N. T. Wells. Ultimate limits in ultrasonic imaging resolution. *Ultrasound in Medicine and Biology*, 17(6):547–558, 1991.

[25] I. L. Herlin and N. Ayache. Features extraction and analysis methods for sequences of ultrasound images. In *Proceedings of the 2nd European Conference on Computer Vision*, pages 43–57, 1992.

[26] A. Hill, A. Thornham, and C. J. Taylor. Model-based interpretation of 3D medical images. In *Proceedings of the British Machine Vision Conference*, pages 339–348, Guildford, 1993.

[27] C.R. Hill, editor. *Physical Principles of Medical Ultrasonics.* John Wiley and Sons, New York, NY, USA, 1986.

[28] T. C. Hodges, P. R. Detmer, D. H. Burns, K. W. Beach, and D.E. Strandness Jr. Ultrasonic three-dimensional reconstruction: in vitro and in vivo volume and area measurement. *Ultrasound in Medicine and Biology*, 20(8):719–729, 1994.

[29] H-M. Klein, R. W. Gunther, M. Verlande, W. Schneider, D. Vorwerk, J. Kelch, and M. Hamm. 3D-surface reconstruction of intravascular ultrasound images using personal computer hardware and a motorized catheter control. *Cardiovascular and Interventional Radiology*, 15(2):97–101, 1992.

[30] F.W. Kremkau. *Diagnostic Ultrasound: Principles and Instruments.* W.B. Saunders Company, Philadelphia, PA, USA, 4th edition, 1993.

[31] A. Lee, J. Deutinger, and G. Bernaschek. Three dimensional ultrasound: abnormalities of the fetal face in surface and volume rendering mode. *British Journal of Obstectrics and Gynaecology*, 102:302–306, April 1995.

[32] R. Martin, E. Blood, F. Sheehan, G. Bashein, C. Otto, F. Derook, E. Filer, X. Li, and P. Detmer. A miniature position and orientation locator for three dimensional echocardiography. *Computers in Cardiology*, pages 25–28, 1993.

[33] R. W. Martin, G. Bashein, P. R. Detmer, and W. E. Moritz. Ventricular volume measurement from a multiplanar transesophageal ultrasonic imaging system: an in vitro study. *IEEE Transactions on Biomedical Engineering*, 37(5):442–449, May1990.

[34] M. Moshfeghi and R. C. Waag. In vivo and in vitro ultrasound beam distortion measurements of a large aperture and a conventional aperture focussed transducer. *Ultrasound in Medicine and Biology*, 14(5):415–428, 1988.

[35] A. Moskalik, P. L. Carson, C. R. Meyer, J. B. Fowlkes, J. M. Rubin, and M. A. Roubidoux. Registration of three-dimensional compound ultrasound scans of the breast for refraction and motion correction. *Ultrasound in Medicine and Biology*, 21(6):769–778, 1995.

[36] T. R. Nelson and T. T. Elvins. Visualization of 3D ultrasound data. *IEEE Computer Graphics and Applications*, 13(6):50–57, November, 1993.

[37] T. R. Nelson and D. H. Pretorius. Visualization of the fetal thoracic skeleton with three-dimensional sonography: a preliminary report. *American Journal of Roentgenology*, 164(6):1485–1488, 1995.

[38] K. J. Ng, J. E. Gardener, D. Rickards, W. R. Lees, and E. J. G. Milroy. Three-dimensional imaging of the prostatic urethra — an exciting new tool. *British Journal of Urology*, 74(5):604–608, 1994.

[39] E. O. Ofili and N. C. Nanda. Three-dimensional and four-dimensional echocardiography. *Ultrasound in Medicine and Biology*, 20(8):669–675, 1994.

[40] R. Ohbuchi, D. Chen, and H. Fuchs. Incremental volume reconstruction and rendering for 3D ultrasound imaging. *Proceedings of SPIE — The International Society for Optical Engineering*, 1808:312–323, 1992.

[41] G. Pasterkamp, C. Borst, A-F. S. R. Moulaert, C. J. Bouma, D. van Dijk, M. Kluytmans, and B. M. ter Haar Romeny. Intravascular ultrasound image subtraction: a contrast enhancing technique to facilitate automatic three-dimensional visualization of the arterial lumen. *Ultrasound in Medicine and Biology*, 21(7):913–918, 1995.

[42] R.C. Preston, editor. *Output Measurements for Medical Ultrasound*. Springer-Verlag, Berlin, Germany, 1991.

[43] A. Rahmouni, A. Yang, C. M. C. Tempany, T. Frenkel, J. Epstein, P. Walsh, P. K. Leichner, C. Ricci, and E. Zerhouni. Accuracy of in-vivo assessment of prostatic volume by MRI and transrectal ultrasonography. *Journal of Computer Assisted Tomography*, 16(6):935–940, 1992.

[44] R. N. Rankin, A. Fenster, D. B. Downey, P. L. Munk, M. F. Levin, and A. D. Vellet. Three-dimensional sonographic reconstruction: techniques and diagnostic applications. *American Journal of Roentgenology*, 161(4):695–702, 1993.

[45] R.N. Rohling, T.M. Peters, J.M. Hollerbach, and P.M. Munger. Comparison of relative accuracy between an optical and a mechinical position tracker for image guided neurosurgery. *Journal of Imaged Guided Surgery*, 1(1):30–34, 1994.

[46] G. Sakas, L-A. Schreyer, and M. Grimm. Preprocessing and volume rendering of 3D ultrasonic data. *IEEE Computer Graphics and Applications*, 15(4):47–54, July, 1995.

[47] A. Salustri and J. R. T. C. Roelandt. Ultrasonic three-dimensional reconstruction of the heart. *Ultrasound in Medicine and Biology*, 21(3):281–293, 1995.

[48] P. M. Shankar. Speckle reduction in ultrasound B-scans using weighted averaging in spatial compounding. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, UFFC-33(6):754–758, November, 1986.

[49] M. L. Skolnick. Estimation of ultrasound beam width in the elevation (section thickness) plane. *Radiology*, 180(1):286–288, 1991.

[50] C. Sohn. Challenges remain in 3-D ob/gyn ultrasound. *Diagnostic Imaging*, November 1991.

[51] E. Steen and B. Olstad. Volume rendering of 3D medical ultrasound data using direct feature mapping. *IEEE Transactions on Medical Imaging*, 13(3):517–525, September, 1994.

[52] H. Steiner, A. Staudach, D. Spitzer, and H. Schaffer. Three-dimensional ultrasound in obstetrics and gynaecology: technique, possibilities and limitations. *Human Reproduction*, 9(9):1773–1778, 1994.

[53] M. H-M. Syn and R. W. Prager. Mesh models for three-dimensional ultrasound imaging. Technical Report CUED/F-INFENG/TR 210, Cambridge University Department of Engineering, December 1994.

[54] M. H-M. Syn and R. W. Prager. Model based three-dimensional ultrasound imaging. Technical Report CUED/F-INFENG/TR 180, Cambridge University Department of Engineering, October 1994.

[55] M. H-M. Syn and R. W. Prager. Bayesian registration of models using FEM eigenmodes. Technical Report CUED/F-INFENG/TR 213, Cambridge University Department of Engineering, September 1995.

[56] M. H-M. Syn and R. W. Prager. A biological growth metric for 3D shape registration. Technical Report CUED/F-INFENG/TR 225, Cambridge University Department of Engineering, July 1995.

[57] M. H-M. Syn and R. W. Prager. FEM eigenmodes as shape features. Technical Report CUED/F-INFENG/TR 211, Cambridge University Department of Engineering, May 1995.

[58] Toshiba Corporation. *Operation Manual [Separate Volume (Application)] for Diagnostic Ultrasound System Model SSA-270A/HG.* Toshiba Corporation, Japan, 1989.

[59] G. E. Trahey, S. W. Smith, and O. T. von Ramm. Speckle pattern correlation with lateral aperture translation: experimental results and implications for spatial compounding. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, UFFC-33(3):257–264, 1986.

[60] J. W. Trobaugh, D. J. Trobaugh, and W. D. Richard. Three-dimensional imaging with stereotactic ultrasonography. *Computerized Medical Imaging and Graphics*, 18(5):315–323, 1994.

[61] P.N.T. Wells. *Advances in Ultrasound Techniques and Instrumentation.* Churchill Livingstone Inc., New York, NY, USA, 1993.

[62] T. A. Whittingham. Resolution and information limitations from transducer arrays. *Physics in Medicine and Biology*, 36(11):1503–1514, 1991.

[63] S. L. Wood. Visualization and modeling of 3-D structures. *IEEE Engineering in Medicine and Biology*, 11(2):72–79, June, 1992.