First Year Summary Report:

3-D ULTRASOUND IMAGING: *OPTIMAL VOLUMETRIC RECONSTRUCTION*

R.N. Rohling

May 15, 1996

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

E-mail: rnr20@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 report 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 2-D image slice. The set of 2-D images and position data are subsequently used to reconstruct a 3-D data set which can be rendered on a computer monitor or used for volumetric data analysis. Initial tests show that measurement errors, including image artifacts, are significant enough to obscure and distort diagnostically important information. An investigation into the sources of measurement error is used as a basis for improving the quality of the 3-D reconstruction. Reduction of measurement errors by image registration is feasible during the reconstruction stage. A novel reconstruction algorithm which performs registration of image slices is described. Several experiments are then performed and the resulting reconstructions are much clearer and of greatly enhanced diagnostic utility. The registration is achieved by applying visual tracking techniques outside their traditional area of application. Since the algorithm is automatic and incremental, it can be easily incorporated into existing 3-D ultrasound systems. Future research will attempt to improve the registration method and apply it to a variety of clinical cases.

Contents

Chapter 1

Introduction to 3-D Ultrasound Imaging

Real-time 3-D ultrasound capable of high resolution imaging over a large volume is the ultimate goal. A 3-D voxel array with $128\times128\times128$ 8-bit elements and 0.5 vps (volumes per second) is a reasonable initial goal for imaging fetuses [12]. An array of $512\times512\times512$ 8-bit elements at 30 vps is the goal for echocardiography [48]. 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 [48] . Real-time 3-D ultrasound research has not yet achieved a level suitable for clinical use and remains an active research area.

Nevertheless, physicians have indicated that there is significant utility in off-line 3-D ultrasound imaging of a variety of anatomical structures, including the urethra [32], vascular structure [20], fetus [12], gall bladder [17], breast [31], kidney [22], heart [39] and uterine cavity [3]. 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 scan planes. Good overviews are found in papers of 3-D imaging in obstetrics [42] and 3-D echocardiography [33, 39]. Most of the issues discussed are applicable to other parts of the anatomy.

A hand-held ultrasound 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 (also called B-scan) are shown in Figure 1.1. The most common approach to scanning a volume with 2-D image slices 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 1.2. 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 1.3. Finally, the probe can be translated along an axis in the elevational direction. A prismatic swept volume, shown in Figure 1.4, is produced.

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

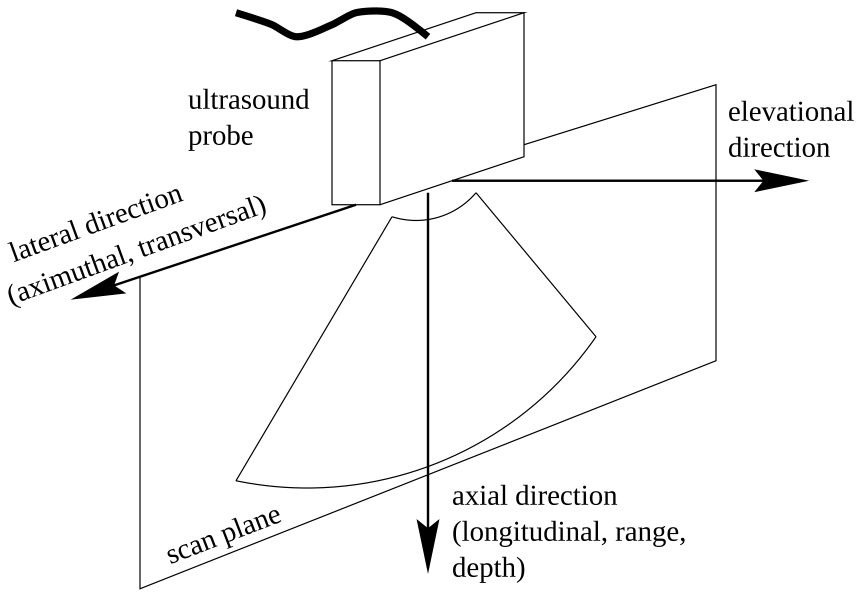

Figure 1.1: Nomenclature of imaging directions. Other common names of the directions are in parentheses.

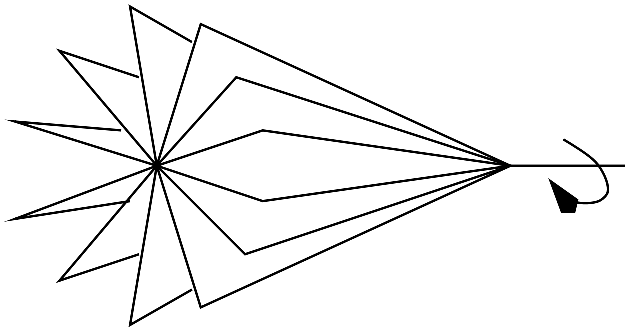

Figure 1.2: Cone-shaped swept volume. The volume is produced by rotating the probe about an axis in the axial direction.

varies axially. The cone-shaped data set has inter-slice 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 [35]. 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. 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 1.5. The motion of the probe is unrestricted and the six DOFs are measured by a position sensor. With

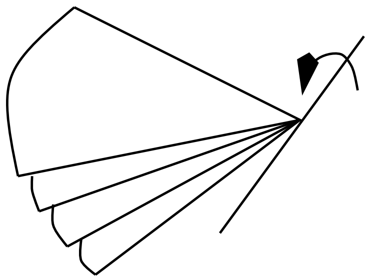

Figure 1.3: Fan-shaped swept volume. The volume is produced by rotating the probe about an axis in the lateral direction.

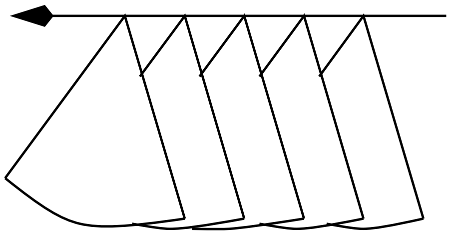

Figure 1.4: Prism-shaped swept volume. The volume is produced by translating the probe along an axis in the elevational direction.

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 are 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. For example, speckle typically obscures anatomical features in 3-D reconstructions of internal organs but is well-known to be reducible by compounding [46, 48]. Other researchers [44] have reported dramatic improvements in the quality of reconstructions from sets of many overlapping 2-D image slices acquired from a variety of look directions. These types of data sets cannot be acquired by swept-volume approaches.

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 [12]. 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, compounding and acquisition time during the examination. This freedom is not possible with the other 3-D ultrasound imaging methods.

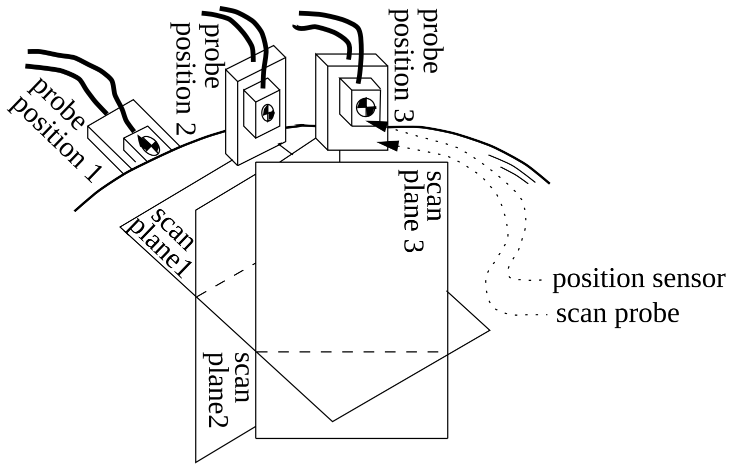

Figure 1.5: 3-D free-hand ultrasound imaging. The physician is allowed to move the probe as in a normal ultrasound examination while the position sensor measures the position and orientation of each scan plane.

Of all the 3-D imaging techniques previously described, free-hand imaging typically has the longest acquisition times. 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. Long acquisition times are not a problem for stationary subjects but subject motion introduces inconsistencies in the acquired data. Errors from subject motion may be partly overcome by incremental reconstruction and visualisation [34], image registration, or by gating acquisition with ECG for echocardiography [39].

In summary, free-hand imaging is a feasible 3-D imaging method with several advantages over other methods: 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 scanned volume size, number of image slices and acquisition time; use of conventional examination techniques; and flexibility for the physician to image from any available look direction. The main purpose of this report is to investigate how to improve 3-D free-hand imaging. The prerequisite for improvement is a thorough investigation into each step of the imaging process. Chapter 2 describes free-hand imaging with a focus on sources of measurement error. Chapter 3 describes a novel reconstruction algorithm which compensates for several of the sources of error. Several ultrasound examinations are undertaken and the results of applying the algorithm are then presented. Chapter 4 contains conclusions about the investigation and the direction in which future research will be taken.

Chapter 2

3-D Free-hand Imaging

2.1 Overview

The imaging process, from 2-D image acquisition through to display of the rendered data and volumetric data analysis, is shown Figure 2.1. 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. 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.

2.2 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

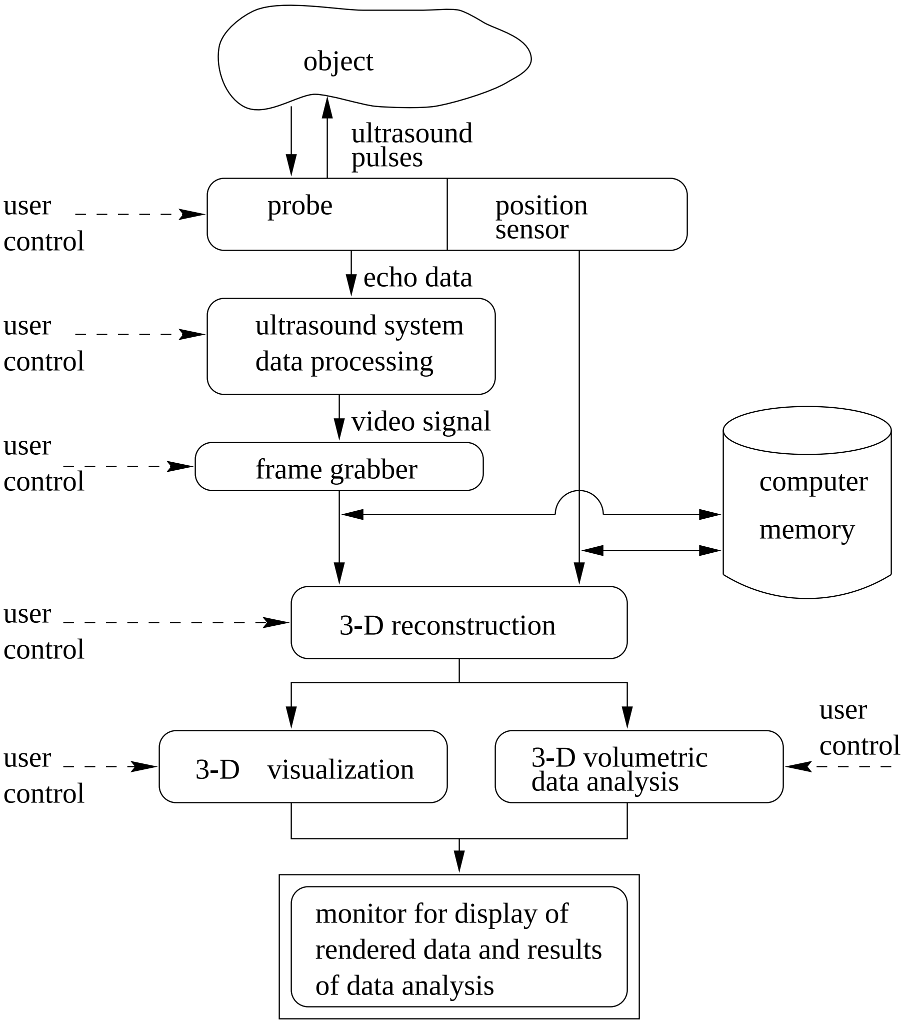

Figure 2.1: Overview of acquisition, reconstruction, visualisation and volumetric data analysis. The user is required to control several parameters at each stage.

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 [12, 19, 29]. 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. A full description of the errors is found in [37]. 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 interaction of ultrasound in biological materials is not fully understood and is an active area of research [15, 16, 26].

2.2.1 Sonographic Artifacts: Propagation Group

These artifacts are related to the ultrasound beam size, direction, speed and spatial wave properties.

axial resolution: the minimum distance between objects positioned axially that are imaged individually. Axial resolution improves with higher frequencies (shorter pulse lengths).

lateral resolution: the minimum distance between objects positioned laterally that are imaged individually. Lateral resolution improves with higher probe frequencies, shorter focal lengths and larger probe apertures. Highest resolution (smallest beam size) occurs at the focal spot.

elevational resolution: the minimum distance between objects, separated in the elevational direction, that are imaged individually. Elevational resolution improves with higher probe frequencies, shorter focal lengths and larger probe apertures. Highest resolution (smallest beam size) occurs at the focal spot.

speckle: the grainy appearance of ultrasound images that arises from constructive-destructive interference of the coherent pulses. Speckle can significantly obscure image features.

reverberation: multiple reflections within the material or from the sensor face. Artifacts are produced by the multiple reflections.

refraction: a change in beam direction at oblique interfaces that results in object position errors.

multi-path reflection: a difference in the length of the return path from the incident path that results in object position errors.

mirror-image: mirror images of objects that are produced by nearby strong reflectors.

side lobes: minor beams of sound travelling off of the beam direction that create echoes from off-beam objects. Artifacts are produced from the off-beam objects.

comet tail and ring-down: artifactual tails exhibited below certain objects due to reverberation or resonance.

speed errors: an incorrect estimate of propagation speed that results in object position errors. Scanning through a thick fat layer, for example, increases speed errors.

range ambiguity: a second pulse transmitted before all echoes from the first pulse return to the probe. Artifacts are produced in the second image by echoes produced from the first pulse.

2.2.2 Sonographic Artifacts: Attenuation Group

These artifacts are related to the amplitude properties of the ultrasound pulses, echoes and electrical signals.

contrast resolution: the smallest detectable difference in echo intensity. Resolution depends on human perception, signal processing and the display monitor.

shadowing: the loss of anatomical information due to high attenuation of pulses behind strong reflectors such as bone or lung.

enhancement: higher than expected echo intensity behind weakly echogenic tissue or fluid.

electronic noise: the noise introduced by analog processing of echoes. Electronic noise ultimately limits imaging depth.

2.2.3 Position Measurement Errors

These errors are related to the position measurement of each 2-D image slice.

translation: accuracy and resolution limitations of position sensor measurement of distance.

rotation: accuracy and resolution limitations of position sensor measurement of orientation. Errors are amplified by a lever arm effect caused by sensor offset from the image plane.

latency and acquisition rate: measurement errors that arise due to latency and low acquisition rate of sensing probe movement.

2.2.4 Tissue-related Errors

These errors are related to assumptions of non-ideal properties of the tissue in the images.

movement: blurring or multiple copies of objects caused by tissue movement. The 3-D reconstruction assumes the subject is stationary during examination.

anisotropy: scatterers in tissue that appear only from certain look directions. Anisotropic scatterers may be overlooked by certain look directions.

aliasing: the high spatial frequency components of anatomy that are not reproduced by widely spaced 2-D images produced by a finite width beam.

2.3 Volumetric 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 necessarily 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.

2.3.1 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. The physician normally controls several parameters on the ultrasound system [12, 19, 29, 45, 48]. Table 2.1 contains a list of the important parameters controllable by the physician. Not all of these

probe used	time-gain compensation	overall gain
display persistence	display dynamic range	frame rate
echo-enhancement	echo-filter	image post-processing
focus parameters	scan line density	display reversal
pulse amplitude	scan angular range	image enlargement

Table 2.1: Parameters than can be controlled by the physician during an examination.

parameters are adjustable on all ultrasound systems, nor can they all be controlled independently.

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 2.2. 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 images 1 and 2 are likely scanned through different tissue types and geometries. The TGC curve (manual control of time-gain compensation for differences in tissue attenuation coefficients) is usually adjusted for uniform average intensity in the portions of the image with the same tissue type. For example, if image 2 scans 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 image 2 will return with a lower amplitude than if it came along a path in image 1. Constant TGC settings for the two images will therefore yield two different average intensities of the object. Similarly, shadowing and enhancement artifacts will be different in the two scans. The physician is expected to minimise intensity non-uniformities by adjusting the TGC settings when changes in average intensity occur.

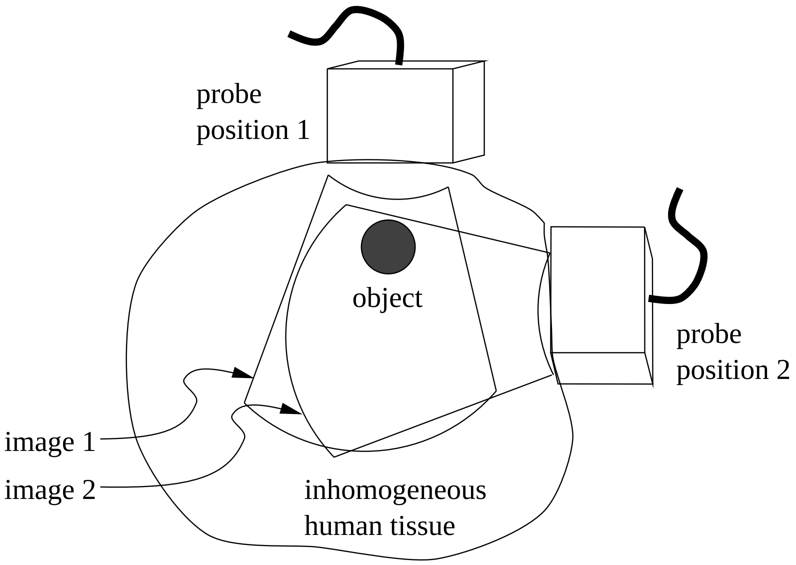

Figure 2.2: Two images of the same object are scanned from different look directions.

2.3.2 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, working on graphics workstations, 256×256×256 voxel reconstructions are manageable and the resolution limit is not reached.

Interpolation may be required in regions of the volume which are under-sampled [34, 47], but ideally, enough images are acquired to fill in all gaps. In fact, multiple compounding of each voxel from several look directions is not an unreasonable goal for some examinations. Compounding is a technique that can be used to reduce speckle artifacts in the reconstruction. Speckle patterns become uncorrelated if a region is viewed from different look directions. Compounding images taken from different look directions will theoretically increase the signal-to-noise ratio in proportion to the square root of the number of compounded images [46, 48]. Figure 2.3 shows an example of speckle reduction by compounding.

Compounding, however, has an inherent tradeoff with loss of boundary sharpness due to measurement errors. Using the same example in Figure 2.2, 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 reveals 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. Figure 2.4 shows an example of registration errors in an examination of a human gall bladder.

Figures 2.3 and 2.4 are derived from examinations performed by an experienced ultrasonographer on human volunteers. Subjects are requested to maintain a breath-hold and minimise body motion for the duration of the examination. The examinations last between 10 and 30 seconds. In clinical tests, strict subject stationarity and breath-holding may not always be possible, so errors may be larger compared to these examinations.

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.

2.4 Visualisation of Volumetric Data

The usual next step, after a 3-D data set has been reconstructed, is visualisation. The important features of a good image are clear detail, obvious tissue texture, visualisation of detail at depth, and low noise [12]. 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. A brief description of several visualisation techniques follows below, but a more thorough review is found in [37].

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 [12, 23, 48]. 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**.

Other visualisation methods include **surface fitting** and **volume rendering** [4, 21, 40,

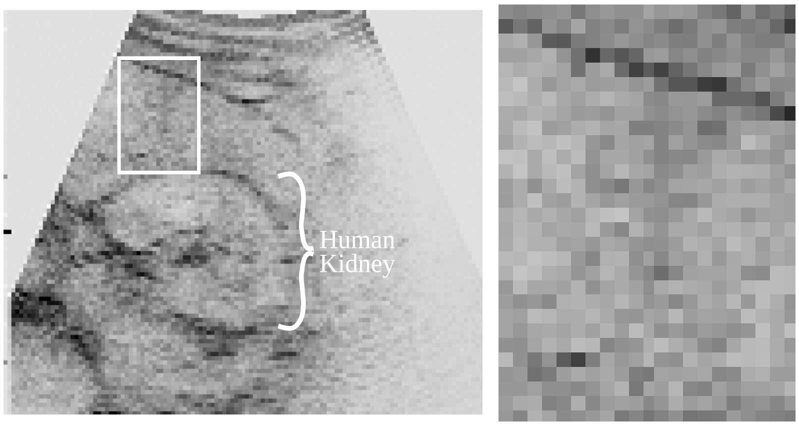

(a) Current Image. (b) Close-up of Selected Region.

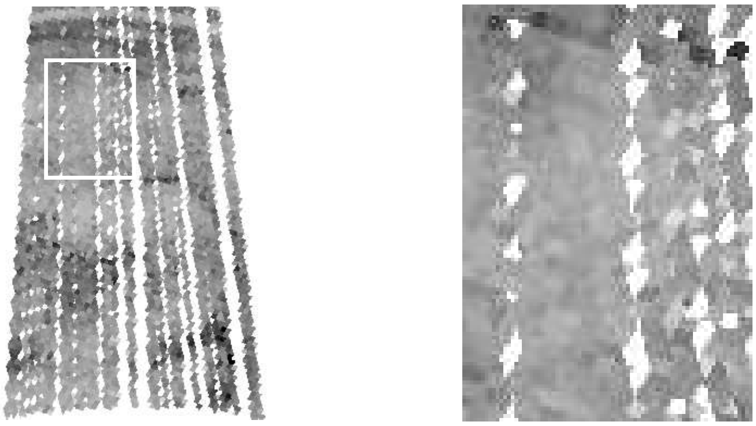

(c) Existing Images. (d) Close-up of Selected Region.

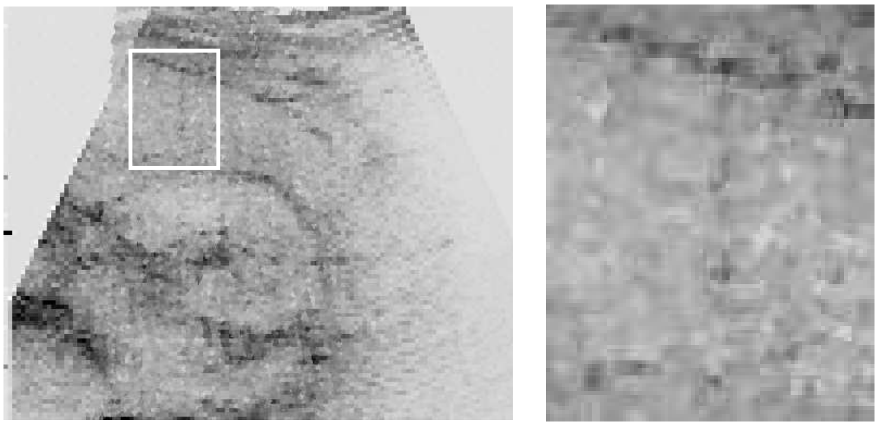

(e) Compounded Current and Existing Images. (f) Close-up of Selected Region.

Figure 2.3: Speckle reduction via compounding of multiple images of a human kidney. Image (a) shows the current image to be compounded with existing images. Image (b) shows the cross-section of the reconstruction cuberille where the current image (a) is to be inserted. The edges of 40 existing images taken from different look directions are visible. Image (c) shows the superposition of the current image on the existing images. Image (f) shows the reduction of speckle — compare to images (b) and (d). Object boundaries are also less distinct.

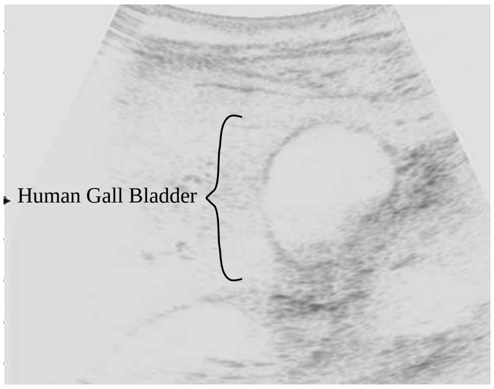

(a) Current Image.

(b) Existing Images (darkened for clarity).

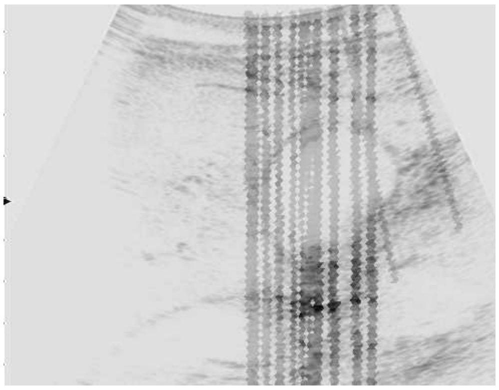

(c) Superposition of Current and Existing Images.

Figure 2.4: Mis-registration of 3-D free-hand ultrasound imaging of a human gall bladder. Image (a) shows the current image to be registered with existing images. Image (b) shows the cross-section of the reconstruction volume where the current image (a) is to be inserted. The edges of 20 existing images taken from different look directions are visible. Image (c) shows the superposition of the current image on the existing images. The position of the gall bladder in the current image does not exactly match the position of the gall bladder in the existing images. This mis-registration is evident as a double boundary on the left side of the gall bladder in image (c). The elapsed time to acquire the images is less than 10 seconds.

41, 49]. Surface fitting 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 [12, 41, 48].

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 [48].

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 [38]: 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; 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 [38]. 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 [42].

2.5 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 imaged objects. Examples where volume or size determination is useful include prostrate volume [36], ventricle volume [30, 33, 39], kidney volume [22], blood vessels [27] and fetus and placenta volumes [12].

Segmentation and modelling of the data [1, 2, 5, 6, 7, 9, 13, 14, 24, 25, 43] are also important research areas. Texture classification — measurement of the statistical properties of material to aid diagnosis — is another active research area [12, 48].

2.6 Summary

The investigation into the sources of error suggests how to best design a 3-D free-hand ultrasound system. The design 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.

Initial experiments with a system of this design reveals that the measurement errors limit its diagnostic utility. Image artifacts such as speckle, tissue motion and position sensor errors can severely limit visualisation and volumetric data analysis of the 3-D data sets. Creating a 3-D data set from a set of *overlapping* image slices is suggested as a way to reduce image artifacts, especially speckle.

Most reconstruction methods in previous research have used solely the data from the position sensor to fill a 3-D scalar array with 2-D image slices. However, the experiments reveal that these methods produce inconsistencies in the 3-D data set. Research should first focus on reconstructing a consistent and accurate 3-D data set from overlapping images by compensating for the measurement errors.

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 reduces storage requirements and makes the system more interactive by allowing incremental rendering during the examination.

The following chapter contains 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 is displayed on a computer monitor. Several different visualisation methods can be used including volume ray-tracing, surface rendering and any-plane slicing. Incremental rendering is also possible [34].

The visualisation and volumetric data analysis stages 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.

Chapter 3

Volumetric Reconstruction

3.1 Overview

Chapter 2 has demonstrated the need to correct for measurement errors during reconstruction. This chapter describes a registration technique to correct for several of the measurement errors, the largest of which are expected to be tissue motion, position sensor and calibration errors. Several ultrasound examinations are performed and reconstructed using the new technique. The resulting 3-D data sets are much clearer and of greatly enhanced diagnostic utility. The registration is achieved by applying established visual tracking techniques outside their traditional area of application.

Little work on the registration of 3-D ultrasound data sets is evident in the literature. In one exception [31], two separate data sets were *retrospectively* registered using *manual* landmark matching. This constitutes a one-off, labour intensive solution to a specific registration problem. Here we propose an *automatic*, *incremental* registration algorithm for use with generic free-hand ultrasound imaging.

3.2 Acquisition System and Test Subjects

The acquisition system comprises a Toshiba model SSA-270A/HG ultrasound scanner with a convex curvilinear array probe operating at a frequency of 3.75 MHz. The manufacturer states the spatial resolution (beam dimensions at the focal spot) as 2.3 mm, 1.8 mm, and 0.8 mm for lateral, elevational and axial resolutions respectively (measured at the 6 dB cutoff and assuming a 2 cycle pulse). These values are similar to the calculated theoretical values [37].

The position and orientation of each scan plane relative to a fixed transmitter are measured by an AC magnetic field receiver (Polhemus FASTRAK) mounted on the probe. Images from the video output of the scanner are recorded by an 8 bit frame grabber board at a rate of 5-7 frames per second. The images and the position data are stored in the memory of a Sun SparcStation 10 workstation.

Simulations are first performed on a computer generated octahedron. The octahedron shape is chosen because the boundary outline of the octahedron in each cross-section is unique within a local neighbourhood. In other words, each cross-section has a unique location where re-inserting the slice into the volume with the octahedron makes a perfect fit.

Laboratory tests are then performed on an organ phantom which consists of a latex balloon filled with water. The phantom is mounted in a bath of water at an elevated temperature so that the propagation speed of ultrasound in the water equals the speed in human tissue.

3.3 Reconstruction Without Registration

The following description is pseudo-code for the algorithm that creates, without any error correction, a volumetric reconstruction of the 2-D image slices.

Reconstruction Algorithm *Without* Registration

1. acquire 2-D image **P** and associated position data ${}^{T}\mathbf{T}_{R}$
2. insert image **P** into reconstruction volume **C**
   - 2.1 determine location of pixel p_{mn} with respect to **C**
     ${}^{C}\underline{\mathbf{x}} = {}^{C}\mathbf{T}_{T}\,{}^{T}\mathbf{T}_{R}\,{}^{R}\mathbf{T}_{P}\,{}^{P}\underline{\mathbf{x}} \longleftrightarrow {}^{C}\underline{\mathbf{x}} = \mathbf{T}\,{}^{P}\underline{\mathbf{x}}$
   - 2.2 if nearest voxel c_{ijk} in **C** is empty, set to p_{mn}
   - 2.3 else set c_{ijk} to weighted average of existing c_{ijk} and p_{mn}
     $c_{ijk} = \frac{n \times c_{ijk}}{n+1} + \frac{p_{mn}}{n+1}$
     where n is incremented after each calculation of c_{ijk}
3. repeat at step 1.

Each image is represented as a 2-D array **P** of intensity values p_{mn}. The 3-D representation of the set of images is chosen to be a cuberille **C**. The term **reconstruction volume** is also used to describe the 3-D data set. Each element c_{ijk} of **C** represents a voxel in space. The voxel size is chosen according to the tradeoff between reconstruction time and resolution but is never set lower than the system's spatial resolution. In these tests, the voxel sizes are approximately five times larger than the 2-D image pixels.

The x, y, z displacement and *roll, pitch, yaw* orientation data from the position sensor is converted into a 4×4 homogeneous transformation matrix ${}^{T}\mathbf{T}_{R}$. A standard notation is used which describes ${}^{T}\mathbf{T}_{R}$ as the transformation *from* the coordinate system at the receiver (R) *to* the coordinate system at the transmitter (T). Figure 3.1 depicts the four coordinate systems used for reconstruction.

The position of a pixel p_{mn} with respect to its plane (P) is expressed as a homogeneous vector ${}^{P}\underline{\mathbf{x}}$. The pixel position, with respect to the cuberille coordinate system (C), can be determined by transformation to the receiver coordinate system, then to the transmitter and finally to the reconstruction volume via ${}^{R}\mathbf{T}_{P}$, ${}^{T}\mathbf{T}_{R}$, and ${}^{C}\mathbf{T}_{T}$ respectively. ${}^{R}\mathbf{T}_{P}$ describes the transformation between the corner of the scan plane and the coordinate system of the receiver. It is determined by calibration on a phantom and remains constant throughout the reconstruction. ${}^{C}\mathbf{T}_{T}$ describes the transformation from the transmitter to the corner of the cuberille. It is set to the limits of the reconstruction volume that the physician scans and also remains constant throughout the reconstruction. The cumulative matrix multiplication of ${}^{C}\mathbf{T}_{T}\,{}^{T}\mathbf{T}_{R}\,{}^{R}\mathbf{T}_{P}$ is abbreviated as simply **T**.

Before the start of the examination, the c_{ijk} elements all have values of zero. As each image is acquired, each voxel element c_{ijk} that is intersected by the scan plane is calculated according to the p_{mn} elements which intersect the voxel. A single voxel will encompass many pixels p_{mn} from the higher resolution image and may be intersected again by future image slices. Step 2 above describes how the algorithm copes with both the reduction in resolution between the image **P** and the reconstruction volume **C**, as well as intersections of multiple scan planes.

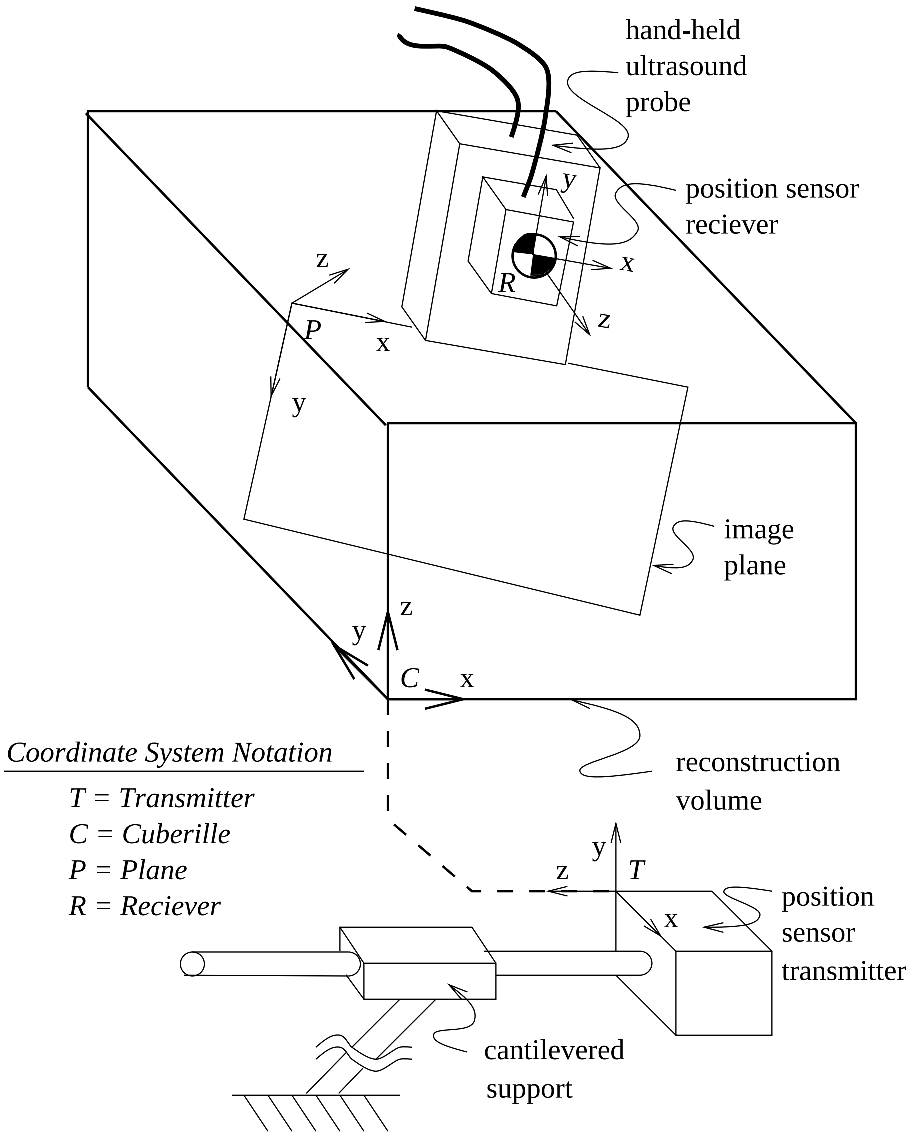

Figure 3.1: Coordinate systems used for reconstruction.

When an image intersects voxels whose values have previously been set, a registration procedure can be carried out to align the image to the existing data in the 3-D array — see Section 3.4.

After a substantial portion of **C** is filled it can be displayed on a computer monitor by several different methods. These methods include volume ray-tracing, surface rendering and any-plane slicing. Surface rendering is used in Figures 3.2, 3.4 and 3.5, and any-plane slicing in Figures 3.6 and 3.7. Figure 3.2 shows the result of applying the above algorithm to the organ phantom, i.e. *without* registration.

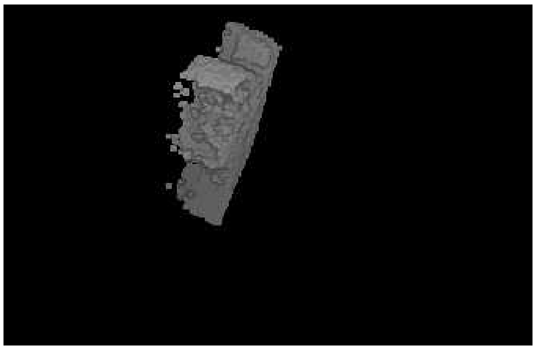

(a) 10 Images Reconstructed

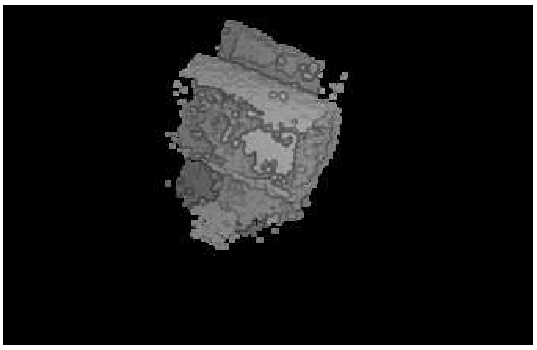

(b) 30 Images Reconstructed

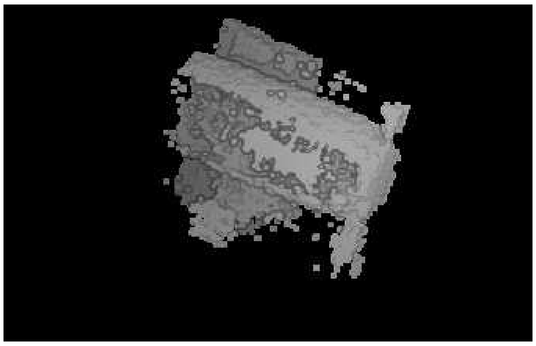

(c) 50 Images Reconstructed

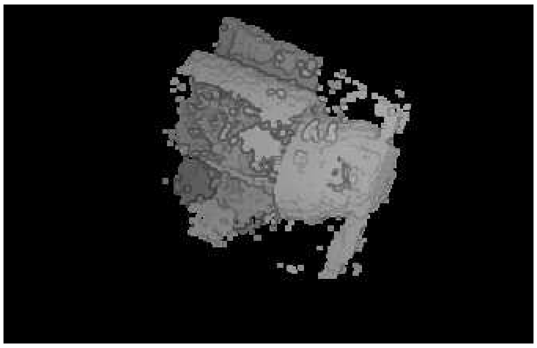

(d) 90 Images Reconstructed

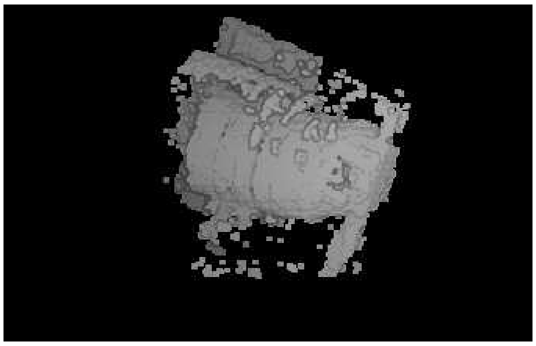

(e) 110 Images Reconstructed

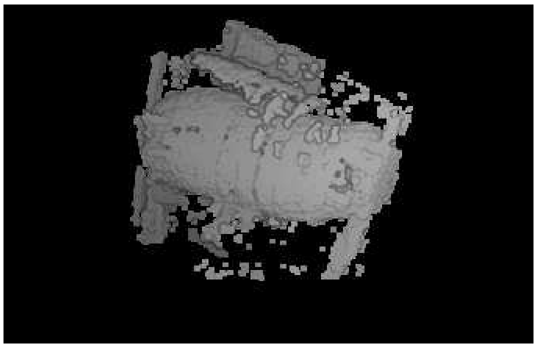

(f) 125 Images Reconstructed

Figure 3.2: Reconstruction of an organ phantom without correction for organ motion. A set of 125 images is obtained by free-hand scanning. The images are acquired by scanning first from left to right, then right to left. Motion induced registration errors are evident: the surface of the phantom constructed after the left-to-right scan (c) does not match the surface constructed by including the following right-to-left scan (f). In particular, the top of the reconstruction in (f) shows multiple surfaces created from the overlapping, mis-registered images. The rendering program attempts to generate surfaces from voxel intensities above a certain threshold. Portions of the surface of the phantom are missing in (c) because the threshold is chosen to minimise speckle and improve clarity of the image. The surfaces in all images are calculated with the same threshold. The floor of the water basin is also evident as well as the mounts at the ends of the phantom.

3.4 Reconstruction With Registration

3.4.1 Overview

Traditionally, registration problems in medical imaging are viewed as unconstrained searches for the transformation which brings one image (or set of images) into optimal alignment with another [10]. The search space is typically very large, resulting in a computationally expensive registration algorithm. With 3-D free-hand ultrasound imaging, we can take advantage of the fact that individual images are acquired in rapid succession, and any registration error should vary smoothly from one image to the next. This observation raises the possibility of performing efficient registration within a visual tracking framework.

Reconstruction Algorithm *With* Registration

1. acquire 2-D image $\mathbf{P}$ and associated position data ${}^{T}\mathbf{T}_{R}$
2. find landmarks l_{mn} in image - see Figure 3.3
3. if the scan plane *does not intersect* existing data (previous slices) in $\mathbf{C}$
   - 3.1 insert image $\mathbf{P}$ into reconstruction volume $\mathbf{C}$
     - 3.1.1 determine location of pixel p_{mn} with respect to $\mathbf{C}$
       ${}^{C}\underline{\mathbf{x}} = {}^{C}\mathbf{T}_{T}\,{}^{T}\mathbf{T}_{R}\,{}^{R}\mathbf{T}_{P}\,{}^{P}\underline{\mathbf{x}} \;\longleftrightarrow\; {}^{C}\underline{\mathbf{x}} = \mathbf{T}\,{}^{P}\underline{\mathbf{x}}$
     - 3.1.2 if nearest voxel c_{ijk} in $\mathbf{C}$ is empty, set to p_{mn}
     - 3.1.3 else set c_{ijk} to weighted average of existing c_{ijk} and p_{mn}
   - 3.2 insert l_{mn} into volume of landmarks $\mathbf{L}$
     - 3.2.1 determine location of l_{mn} (at ${}^{P}\underline{\mathbf{x}}$) with respect to $\mathbf{C}$
       ${}^{C}\underline{\mathbf{x}} = \mathbf{T}\,{}^{P}\underline{\mathbf{x}}$
     - 3.2.2 assign nearest vector in $\mathbf{L}$ to l_{mn}
   - 3.3 repeat at step 1.
4. if scan plane *intersects* existing data in $\mathbf{C}$, registration is performed
   - 4.1 optimise $\mathbf{T}$ to align l_{mn} with existing landmarks in $\mathbf{V}$
     $\mathbf{T} \longrightarrow \mathbf{T}^{*}$
   - 4.2 insert image $\mathbf{P}$ into reconstruction volume $\mathbf{C}$ via $\mathbf{T}^{*}$
     - 4.2.1 determine location of pixel p_{mn} with respect to $\mathbf{C}$
       ${}^{C}\underline{\mathbf{x}} = \mathbf{T}^{*}\,{}^{P}\underline{\mathbf{x}}$
     - 4.2.2 if nearest voxel c_{ijk} in $\mathbf{C}$ is empty, set to p_{mn}
     - 4.2.3 else set c_{ijk} to weighted average of existing c_{ijk} and p_{mn}
   - 4.3 repeat at step 1.

A variety of image registration techniques are available [10]. Correlation based techniques are the simplest but do not work well for 3-D ultrasound image registration. The first reason is the technique's sensitivity to average grey level which may not be consistent throughout the scan (Section 2.3.1). Secondly, the 6 DOF registration problem cannot be converted into more efficient calculations in the frequency domain, unlike 2 DOF problems. Registration times are therefore prohibitively long. Furthermore, correlation based techniques are also sensitive to the presence of residual non-rigid-body registration errors that are present in ultrasound images (Section 2.2). Registration via *landmarks* on objects in the image avoids these problems. Essentially, it is proposed to find $\mathbf{T}^*$ using landmarks [10]: anatomical features which can be reliably detected in all images.

The main difference between this algorithm and the algorithm without registration (Section 3.3) is that the image slice transformation $\mathbf{T}$ is replaced with $\mathbf{T}^*$. $\mathbf{T}^*$ is the optimal transformation that registers, via landmarks, the image $\mathbf{P}$ with the existing data in $\mathbf{C}$. It is constrained to compensate only for physically realistic registration errors. In general, it may be necessary to compensate for imaging errors, tissue motion and inaccurate measurement of the scan planes' positions. In this study, however, the main source of error is subject motion during the scan. $\mathbf{T}^*$ is therefore constrained to a rigid 6 DOF transformation consistent with motion-related errors. Smoothly varying registration errors will result in smoothly varying parameters of $\mathbf{T}^*$, hence the suitability of a tracking framework.

3.4.2 Detection of Landmarks

Landmarks generally lie on surface boundaries which are distributed smoothly in 3-D space. Thus, when a newly acquired image $\mathbf{P}$ is inserted into $\mathbf{C}$, its landmarks should lie close to landmarks detected in other images. Registration errors can be corrected by searching for corresponding landmarks and applying a transformation to the new plane so that the landmarks are brought closer together. Figure 3.3 depicts the landmark detection procedure.

The landmarks used in this study are edge elements (edgels) automatically extracted by the Canny edge detection algorithm [11] operating on a thresholded image. The edgels are produced at a lower resolution corresponding to the voxel size of $\mathbf{C}$. The edgel set is further pruned by chaining neighbouring edgels together and eliminating chains with fewer than three edgels. This reduces the number of edgels produced by speckle. The set of edgels produced in this manner is sufficient for the organ phantom reconstructions, though more sophisticated landmark extraction techniques [1] could be used for *in-vivo* studies.

Landmarks are stored in a 3-D vector array $\mathbf{L}$ as the 2-D images are acquired. The scan plane of the first image, and likely the next several images, will not intersect any existing data in the initially empty cuberille $\mathbf{C}$. When an scan plane intersects non-empty voxels in $\mathbf{C}$, $\mathbf{T}^*$ is determined by finding correspondences between landmarks in the new image and landmarks in $\mathbf{L}$.

A minimum number of intersections between pixels p_{mn} in the scan plane and non-empty voxels in $\mathbf{C}$ is needed for an accurate registration solution. Registration of the organ phantom images is performed when greater than 25% of the p_{mn} intersect data in $\mathbf{C}$. Fewer intersections than 25% produces a registration problem with too few landmark correspondences.

3.4.3 Selection of Correspondences

Potential correspondences between landmarks in the 2-D image and landmarks in $\mathbf{L}$ are found by searching an ellipsoidal volume in $\mathbf{L}$ around each landmark in the image. The search for correspondences must be restricted if the registration is to proceed rapidly. To this end, a *motion model* is used to describe the expected change in registration error from one 2-D image slice to the next. Then, given a registration error of one image, the search for correspondences in the next can be greatly constrained. A simple zero-velocity motion model, with a small uncertainty in position based on expected maximum motion, is found to be sufficient to account for the motion-induced errors in the organ phantom tests.

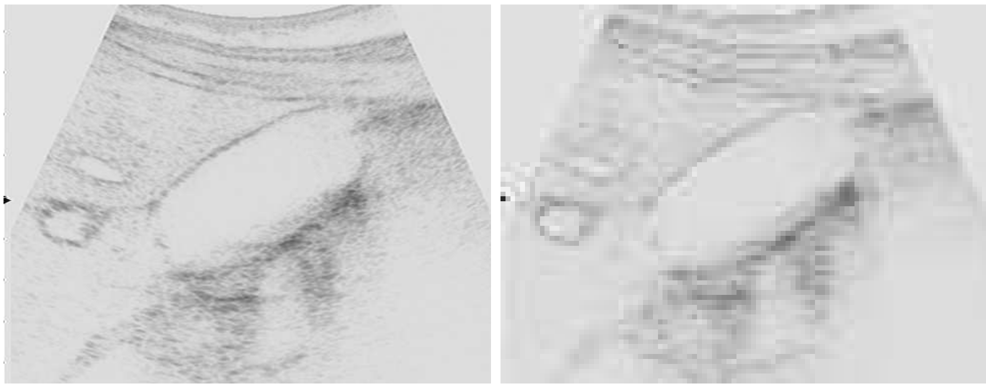

(a) High Resolution Image of Gall Bladder. (b) Low Resolution Image of Gall Bladder.

(c) Edgels Extracted from (b). (d) Chains of Edgels from (c).

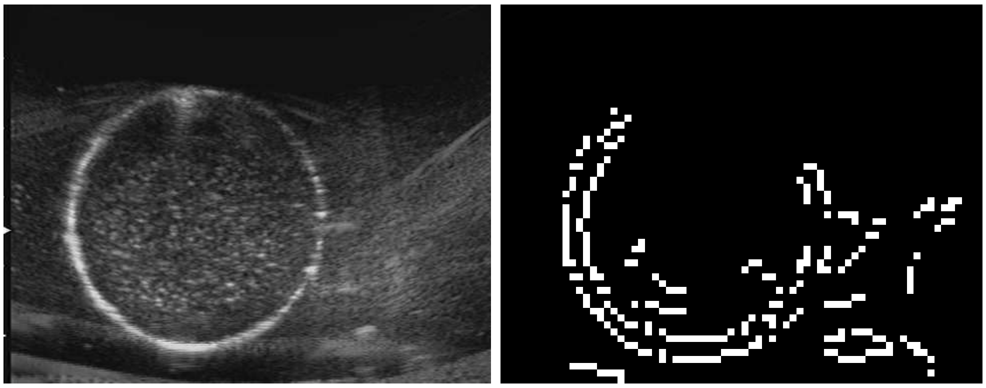

(e) High Resolution Image of Organ Phantom. (f) Chains of Edgels Extracted from (e).

Figure 3.3: Sequence of Landmark Detection. Image (a) is the original high resolution (490x380) image of the human gall bladder. Notice the speckling throughout the image, shadow-like artifacts amd non-uniform intensity of the organ boundary. Image (b) is image (a) at the reduced resolution (99x77) required for the reconstruction volume. Image (c) depicts the 1409 edgels extracted from image (b). Image (d) depicts the chains that are made from the edgels in (c). Notice the reduced number of edgels (854) in image (d) — the edgels forming small lines and circles are eliminated. Images (e) and (f) show the same detection procedure applied to the organ phantom. As well as the organ boundary, the detector also finds edges which do not correspond to real physical structures. These can generate false correspondences which must be tolerated by the registration algorithm.

3.4.4 Determination of Transformation Matrix and Compounding

The set of correspondences produces more constraints than are required to determine $\mathbf{T}^*$. For example, registration of the organ phantom images produces approximately 1200 correspondences, but only 3 are needed to determine $\mathbf{T}^*$. A least squares estimation of $\mathbf{T}^*$ is inappropriate for this problem, because many of the candidate correspondences are outliers. Instead, the RANSAC regression technique [18] is used to determine $\mathbf{T}^*$ and reject the erroneous correspondences:

1. Randomly pick 3 pairs of corresponding landmarks from the full set of correspondences.
2. Reject these correspondences if they are not consistent with a rigid body transformation.
3. Reject these correspondences if the landmarks in $\mathbf{P}$ are too close together or colinear (otherwise the calculation of $\mathbf{T}$ in step 4 is ill-conditioned).
4. Calculate a linear affine transformation $\mathbf{T}$ which brings the three landmarks in $\mathbf{P}$ into precise registration with the corresponding landmarks in $\mathbf{L}$. The pruning in step 2 ensures that $\mathbf{T}$ represents a rigid body transformation.
5. Transform all the remaining landmarks in $\mathbf{P}$ by $\mathbf{T}$.
6. Count how many of the transformed landmarks in $\mathbf{P}$ register with their counterparts in $\mathbf{L}$. Those landmarks that do register contribute to the *consensus set* for $\mathbf{T}$. The remaining landmarks are deemed outliers.
7. Repeat from 1 until a $\mathbf{T}$ is found with a consensus set larger than a preset threshold: this $\mathbf{T}$ becomes the estimate of the optimal transformation $\mathbf{T}^*$.

The technique described above is sufficiently robust to tolerate a significant proportion of erroneous correspondences. It is occasionally possible for more than one solution to exist, i.e., when a large consensus set can be produced by more than one transformation $\mathbf{T}$. This scenario occurs when there is degeneracy in the shape of the imaged object. For example, cross-sections of a perfect cylinder are ellipses which can register with the original cylinder at several positions. Experiments on the organ phantom show an occasional tendency for the registration to follow this degeneracy. The tendency is reduced, however, by the non-symmetrical shape of the phantom and the motion model constraints. This scenario is even less likely to occur for registration of more complicated objects in *in-vivo* examinations.

$\mathbf{T}^*$ is used to transform the new image $\mathbf{P}$ into the coordinates of $\mathbf{C}$. The pixels of the registered image can now be added to $\mathbf{C}$. Where the registered scan plane intersects already-filled voxels, weighted averaging is used to compound the new image with the existing data (see algorithm step 2.3 in Section 3.3).

3.5 Results

The registration algorithm yields a more consistent reconstruction of a 3-D data set from multiple 2-D image slices than reconstructions based on position sensor data alone. Figure 3.4 shows that the algorithm can compensate for different types of motion errors and track the motion accurately. Figure 3.5 shows the reconstruction, *with* registration, of the organ phantom which is allowed to move during the scan. Figures 3.6 and 3.7 show cross-sections of the reconstruction, with and without registration. Figure 3.8 shows the displacement parameters of the motion correction. These figures clearly demonstrate that motion induced errors must be corrected to achieve clear 3-D renderings.

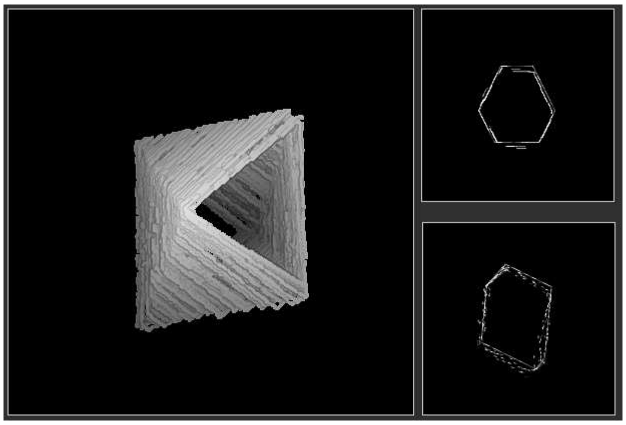

(a) Octahedron with Small Random Motion Errors.

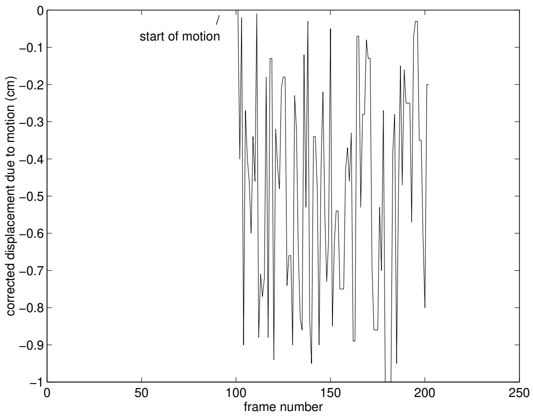

(b) Motion Errors Corrected by Registration.

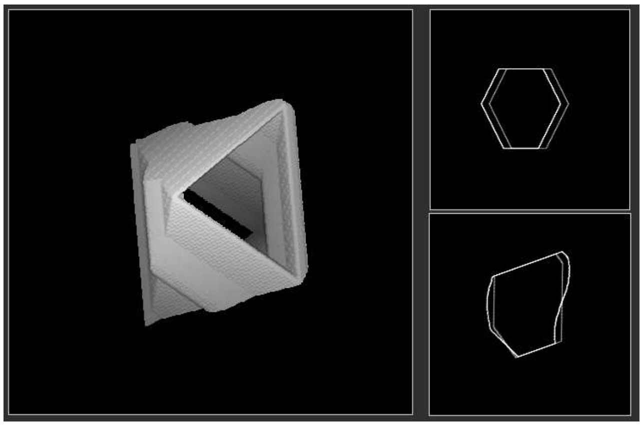

(c) Octahedron with Sinusoidal Motion Errors.

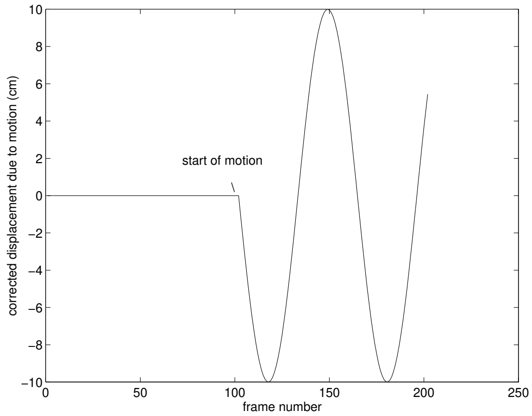

(d) Motion Errors Corrected by Registration.

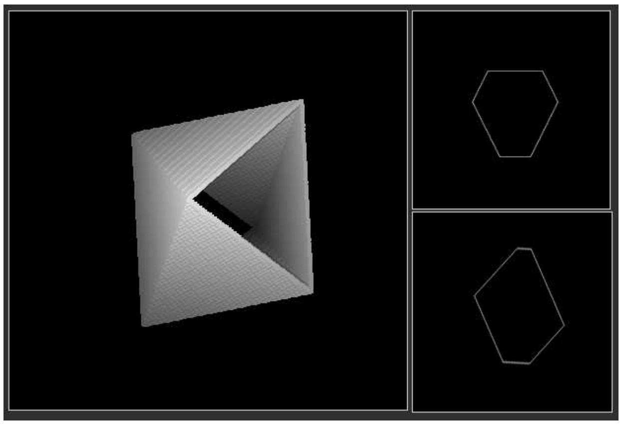

(e) Octahedron Reconstructed with Motion Correction.

Figure 3.4: Registration of an octahedron test object. A simulated scan of the test object from left-to-right, then right-to-left, is performed. The right-to-left portion of the scan contained errors due to simulated motion of the object. Figures (a) and (c) show the reconstruction (without registration) and two typical cross-sections of scans with a small amount of random motion errors and sinusoidal motion errors respectively. Figures (b) and (d) show the motion errors in (a) and (c) that are corrected by the registration algorithm. Figure (e) shows the result of reconstruction of (a) and (c) with registration.

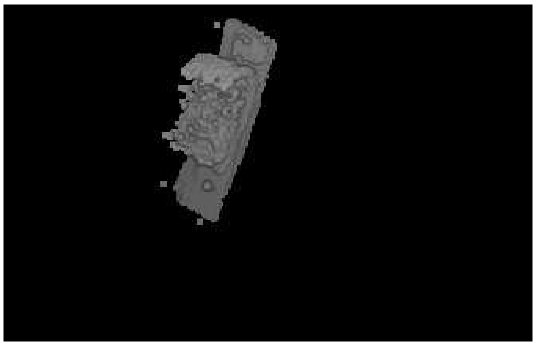

(a) 10 Images Reconstructed

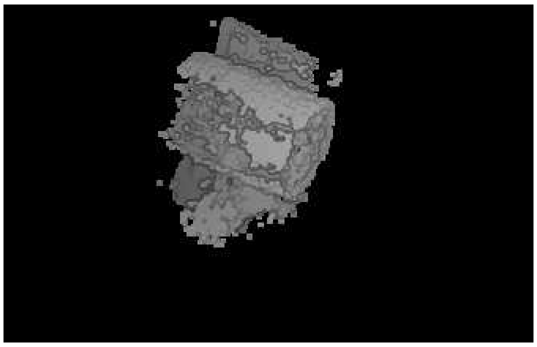

(b) 30 Images Reconstructed

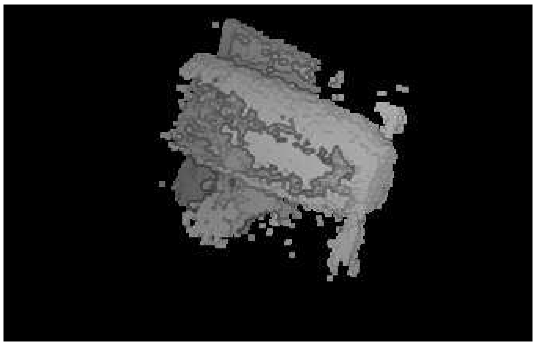

(c) 50 Images Reconstructed

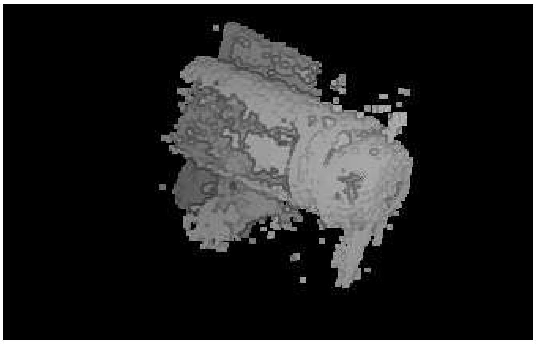

(d) 90 Images Reconstructed

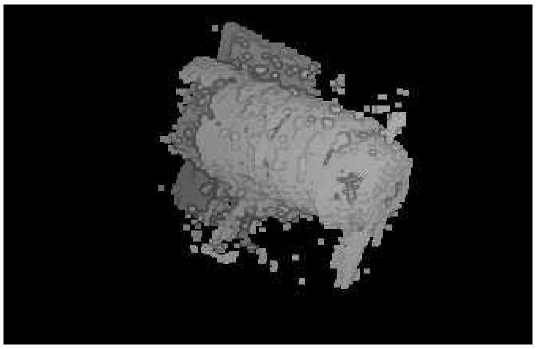

(e) 110 Images Reconstructed

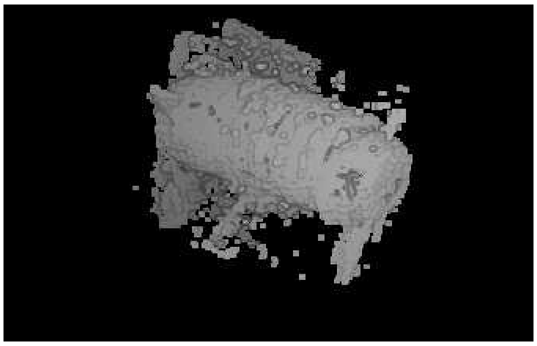

(f) 125 Images Reconstructed

Figure 3.5: Reconstruction of an organ phantom with correction for organ motion. The use of the registration algorithm during reconstruction results in a better alignment of boundary contours — compare with Figure 3.2. The compounding of both the left-to-right and right-to-left scanned images is particularly useful for reducing the uncorrelated speckle noise and filling in surface patches missed by the single left-to-right sweep. The surfaces in all images are calculated at the same threshold, equal to the value used in Figure 3.2.

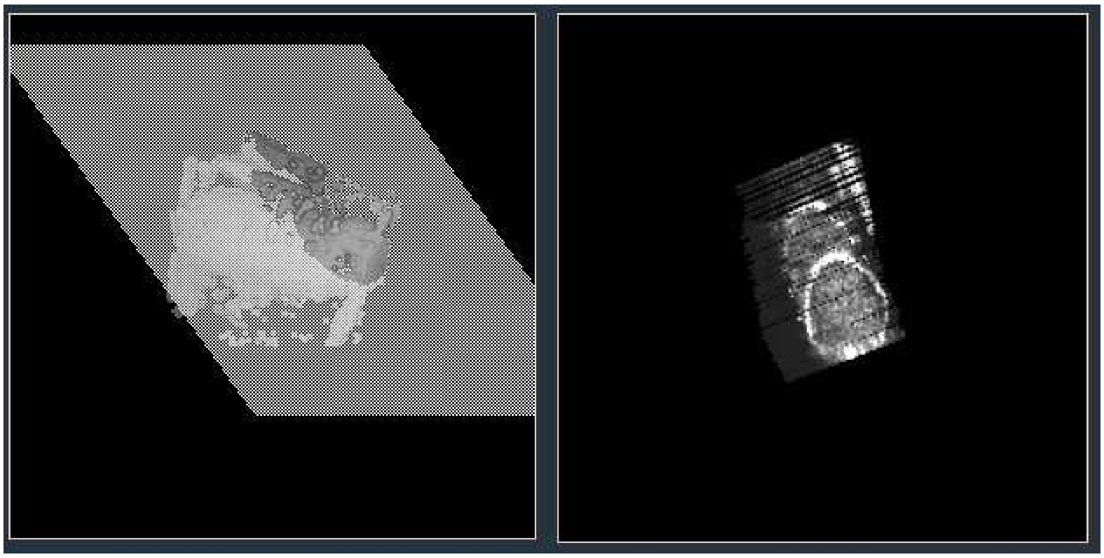

Figure 3.6: Cross-section of a reconstruction, *without* registration, of an organ phantom. The surface reconstruction (left image) is the same as in Figure 3.2(f). The cross-section clearly shows the two mis-aligned ellipses which correspond to the boundary of the phantom. The top ellipse corresponds to images taken during the left-to-right sweep of the scan and the bottom ellipse to the right-to-left sweep, after small motion of the phantom. The stratified look of the cross-section results from selecting a slice at an angle to the original images.

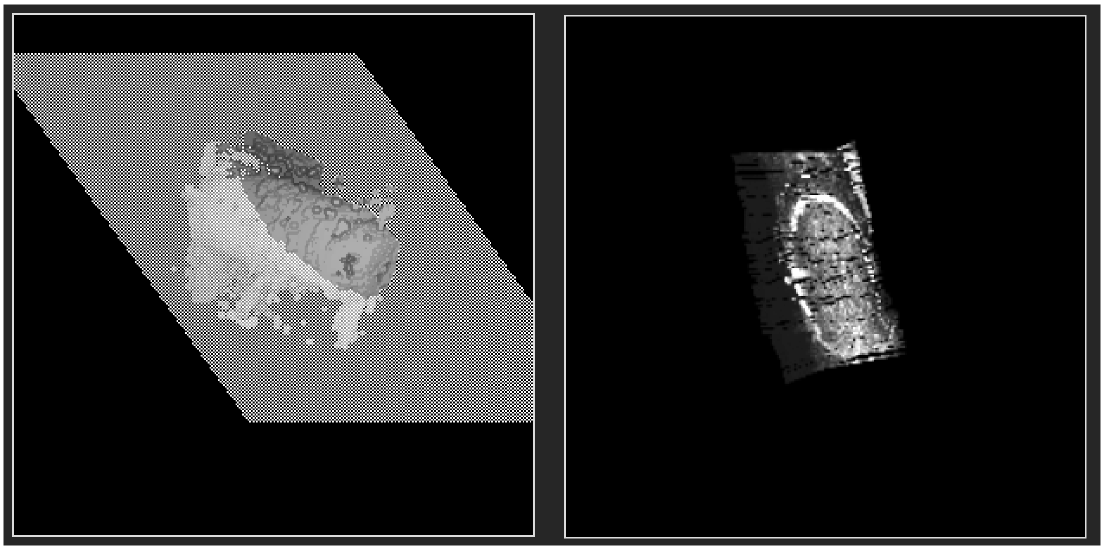

Figure 3.7: Cross-section of a reconstruction, *with* registration, of an organ phantom. The surface reconstruction (left image) is the same as Figure 3.5(f). The cross-section is taken at the same location as in Figure 3.6. The two ellipses have been aligned by the registration algorithm.

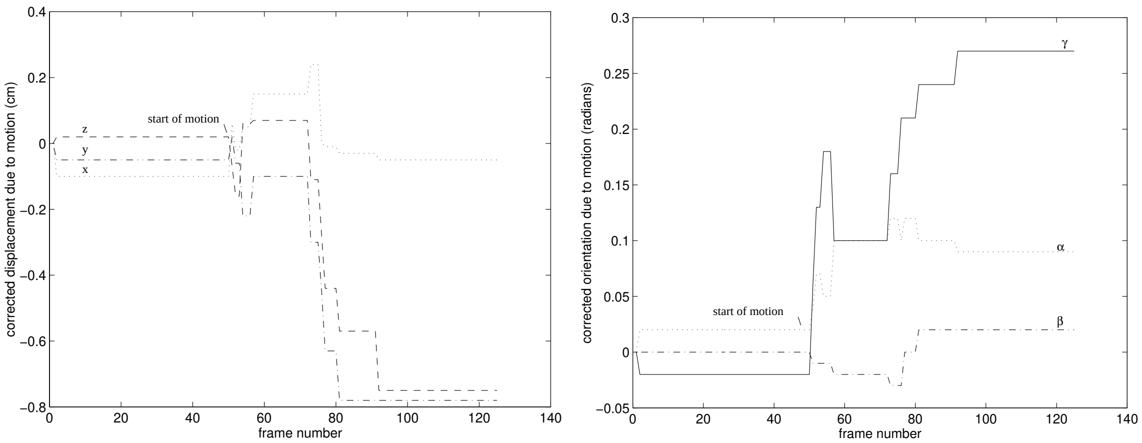

(a) Correction of the x, y, and z errors.

(b) Correction of the α, β, and γ errors.

Figure 3.8: Amount of the displacement (a) and orientation (b) motion errors in organ phantom that are corrected by the registration algorithm. These corrections are consistent with the type of motion observed during the examination of the phantom.

3.6 Discussion

The main purpose of development of the algorithm is to prove the feasibility of automatic error correction by landmark matching. The results prove its feasibility, yet the algorithm has several limitations. The first is the time required to register all image slices. Several hours of computation on a workstation are typically required for reconstruction of 100 image slices. The speed of each registration increases roughly cubicly with the size of the search space for correspondences. The search space must encompass the largest expected error, so larger errors means much longer reconstruction times. Another limitation is the algorithm's sensitivity to false landmarks. The parameters of the algorithm, such as consensus threshold, are set so that landmark outliers are rejected and measurement errors are tracked correctly. These parameters must be carefully adjusted for each type of examination or the algorithm will lose track of the errors and mis-register subsequent image slices.

These findings have implications in 3-D free-hand ultrasound imaging of human subjects. Patient and organ motion are two of the major problems which must be overcome when reconstructing 3-D anatomy from a series of scans. Application of the registration algorithm to the scanning of human subjects is currently being investigated.

The robustness of the algorithm could be improved by employing more sophisticated motion models tuned to specific motion patterns (eg. breathing). The motion models can be learned from examples using established system identification techniques [8]. For difficult data sets, where the detection of landmarks is unreliable, it may be necessary to employ a tracking algorithm which is capable of maintaining several motion hypotheses simultaneously, such as the CONDENSATION algorithm [28]. More reliable landmark detection will also improve tracking.

Chapter 4

Conclusions

4.1 Motivation for Future Research

Ultrasound is a safe and inexpensive imaging technology used in a variety of medical fields including obstetrics, gynaecology, internal medicine, opthalmology and surgery. Physicians in all of these fields have indicated a need for 3-D representations of ultrasound data. Among the different techniques, free-hand imaging holds the most promise for effective and inexpensive 3-D ultrasound. The flexibility of the free-hand paradigm allows the physician to use existing examination techniques to fill in shadowed regions and reduce artifacts by compounding. The size of the examined volume can also be expanded at will during an examination. 3-D free-hand imaging capability can easily be added to existing ultrasound machines, retaining the low-cost nature of ultrasound imaging.

Volumetric reconstruction of the image data acquired by free-hand imaging can be used for visualisation and volumetric data analysis. Both initial experimental results and an investigation into the sources of measurement error reveal significant inconsistencies in volumetric reconstructions unless the errors are corrected.

The largest sources of error for examinations of internal organs are expected to be organ motion, position sensor errors and position sensor calibration. Smaller, but still significant, errors include refraction, multi-path, propagation speed, position sensor latency and low acquisition rate errors. These errors are also likely to be significant in other types of examinations such as fetal imaging [12].

Reduction of the errors is possible by registration of the image slices during reconstruction. When an image slice intersects previously filled voxels in the reconstruction volume, registration aligns the image slice to the existing data. Registration is best performed by matching landmarks in the 2-D image with landmarks in the reconstruction volume.

The main goal of accurate registration is to allow compounding of multiple 2-D image slices to reduce image artifacts, especially speckle. Speckle can significantly obscure features in an image but can be reduced by compounding multiple images taken from a variety of look directions. The measurement errors, however, cause a loss of sharpness of image features when compounding. Accuracy of the registration is required to minimise the blurring of image features. Some researchers [44] have suggested that dramatic improvements in visualisation and volumetric data analysis are possible by compounding. Nevertheless, a tradeoff exists between feature sharpness and speckle reduction because not all measurement errors can likely be corrected. This tradeoff must be recognised when deciding how many image slices to compound together.

An automatic and incremental reconstruction algorithm was developed which proves the feasibility of accurate registration by landmark matching. The algorithm has limitations in speed and sensitivity to false landmark detection. Inaccurate registration of 2-D images at the beginning of the examination can cause failure of subsequent images. Use of more tolerant motion models makes the algorithm more robust to inaccurate registrations but greatly increases reconstruction times. Further research should focus on improving the algorithm then applying it to a variety of clinical cases.

4.2 Future Research

The next step of improvement of 3-D free-hand imaging is to perform a number of *in-vivo* examinations. The first purpose of the examinations is to determine the size of the measurement errors and the amount which can be removed by rigid-body transformations of each 2-D image slice. The investigation will ideally encompass a variety of human physiques, types of scanned tissue (organs) and breath-holding techniques. The second purpose of the examinations is to investigate the effects of compounding on the quality of reconstructions for each of the different types of examination. The initial tests show a small improvement of speckle with a small amount of compounding so tests with a higher number of intersecting scan planes need to be performed. The effects of compounding also need to be quantitatively measured.

In-vivo examinations with many intersecting images require an accurate and robust registration algorithm. Techniques developed for visual tracking in computer vision provide initial starting points for further improvement of landmark registration in the presence of noise and outliers.

The main improvement of registration accuracy and robustness, however, will be derived from more reliable landmark detection. In other words, what are the best features to use for matching, and how can they be reliably extracted from a typically noisy 2-D ultrasound image? Techniques developed for multi-modal image registration, model-based feature extraction, and texture analysis of ultrasound images are good starting points. The use of grey-level (correlation-based) matching techniques will also be researched.

In summary, the series of *in-vivo* examinations will provide the basis for investigations into error modeling, compounding, and landmark detection. Thorough investigations into each of these areas have not yet been attempted for 3-D free-hand ultrasound imaging. Furthermore, the quality of the registration depends on the efficacy of the registration method, so specialised registration techniques will also be investigated. Time permitting, clinical tests of the reconstruction algorithm will be performed to evaluate the utility of the optimised volumetric reconstructions for diagnosis. The time-table for future research is shown in Figure 4.1 and should lead to a comprehensive thesis on the optimal volumetric reconstruction for 3-D free-hand ultrasound imaging.

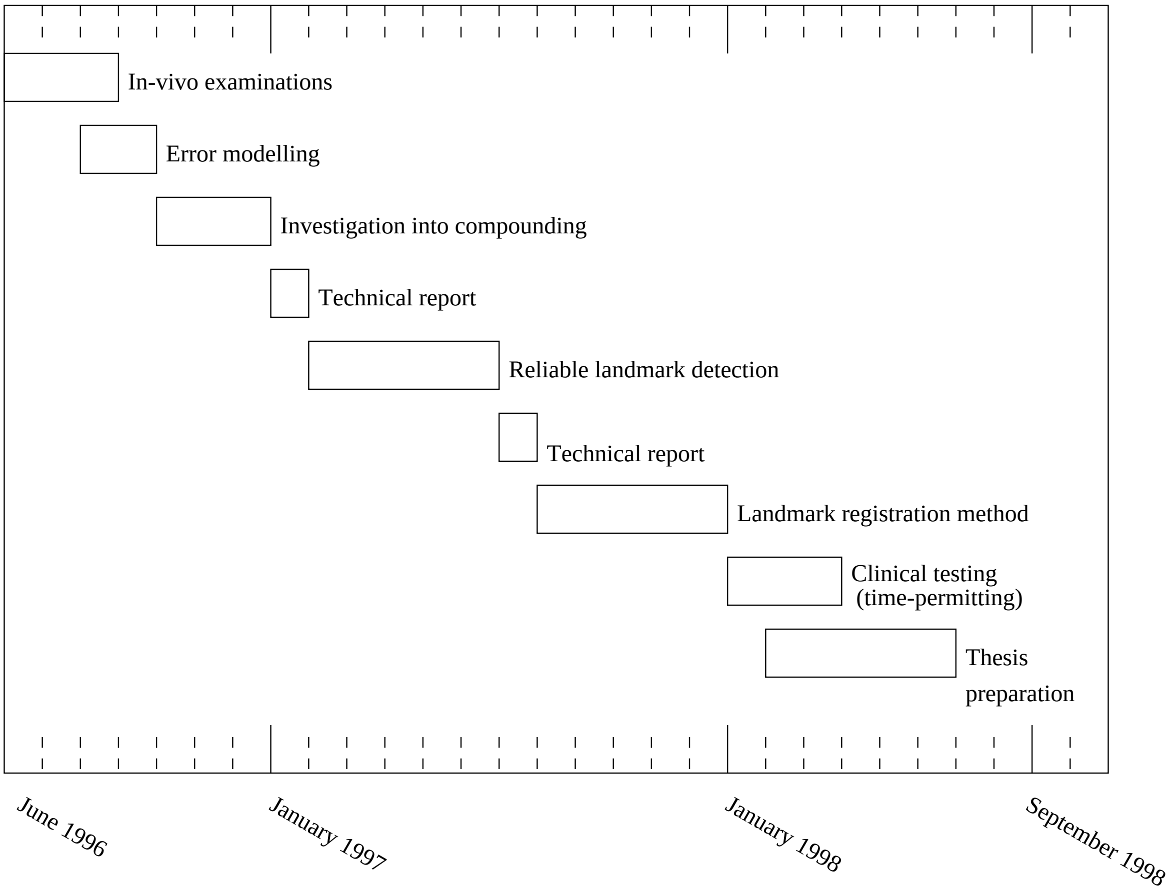

Figure 4.1: Timetable for thesis research.

Bibliography

[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] 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.

[7] 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.

[8] A. Blake, M Isard, and D. Reynard. Learning to track the visual motion of contours. *Artificial Intelligence*, 78:101–133, 1995.

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

[10] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, December, 1992.

[11] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.

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

[13] 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.

[14] 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.

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

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

[17] 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.

[18] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, June, 1981.

[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] 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.

[25] 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.

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

[27] 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.

[28] M. Isard and A. Blake. Contour tracking by stochastic propagation of conditional density. In *Proceedings of the 4th European Conference on Computer Vision*, volume 1, pages 343–356, 1996.

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

[30] 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.

[31] 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.

[32] 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.

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

[34] 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.

[35] 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.

[36] 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.

[37] R. N. Rohling and A.H. Gee. Issues in 3-D free-hand medical ultrasound imaging. Technical Report CUED/F-INFENG/TR 246, Cambridge University Department of Engineering, January 1996.

[38] 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.

[39] 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.

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

[41] 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.

[42] 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.

[43] 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.

[44] D.S. Thompson, D. Barry, and C. Allot. Zeneca pharmaceuticals. *Personal Communication*, 1996.

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

[46] 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.

[47] 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.

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

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