AUTOMATIC REGISTRATION OF
3-D ULTRASOUND IMAGES

R.N. Rohling, A.H. Gee and L. Berman

CUED/F-INFENG/TR 290

May 1997

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

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

Abstract

One of the most promising applications of 3-D ultrasound lies in the visualisation and volume estimation of internal 3-D structures. Unfortunately, the quality of the ultrasound data can be severely degraded by artifacts and speckle, making automatic analysis of the 3-D data sets very difficult. In this paper we investigate the use of **3-D spatial compounding** to improve the quality of the ultrasound data. This involves imaging the region of interest repeatedly from a variety of insonation angles, and averaging the resulting data sets. For accurate compounding, it is important to register the multiple data sets precisely. We show how state-of-the-art techniques, developed elsewhere for multimodal CT to MRI registration, are equally successful with 3-D ultrasound. Results are based on *in-vivo* examinations of a human gall bladder, demonstrating clearly the superiority of the compounded data. In particular, it is possible to visualise and segment the compounded data using standard software packages developed for CT and MRI. We also demonstrate the accuracy of automatic volume estimates made from the compounded data, compared with labour-intensive, manual estimates from the uncompounded data. The registration algorithm also has applications in intra- and inter-patient comparitive studies.

1 Introduction

3-D ultrasound is a new imaging modality that has already been recognised as a valuable tool for a variety of clinical applications. Conventional 2-D diagnostic imaging is performed with a hand-held probe which transmits ultrasound pulses into the body and receives the echoes. The magnitude and timing of the echoes are used to create a 2-D grey-level image (B-scan) of a cross-section of the body in the scan plane. 3-D ultrasound extends this concept so that volumes of intensity data are created from pulse-echo information.

Unfortunately, high quality, instantaneous 3-D imaging remains a long term research goal. One promising approach centers around the development of a new type of phased array probe, which sends and receives echos from a 2-D array of elements (instead of the usual 1-D array) [19]. However, several technical challenges must be overcome before such probes receive clinical acceptance [19]. Alternative approaches, which make use of conventional 2-D ultrasound technology, include the free-hand and swept volume techniques [14, 20]. Instead of taking an instantaneous 3-D snapshot, these techniques construct a 3-D data set from a number of 2-D B-scans acquired in rapid succession.

In the free-hand paradigm, a 3-D position sensor is attached to the probe, so that each B-scan can be labelled with the position and orientation of the scan plane — see Figure 1. Subsequent processing can build up a 3-D description of the imaged anatomy, in much the same manner as is possible with CT or MRI, but with less expensive and invasive technology. Swept volume systems are similar but employ a specially constructed probe to mechanically sweep the B-scans through a volume of interest. In this paper we focus on the free-hand technique.

Both the swept-volume and free-hand techniques have been successfully applied to a number of clinical problems. These include imaging of the fetus [3], vascular structure [5], gall bladder [4], breast [11], kidney [6], and heart [18]. In review articles about 3-D ultrasound [7, 14], the authors suggest that 3-D visualisation and volume estimation are among its most attractive capabilities. Unfortunately, standard visualisation and analysis programs (designed primarily for magnetic resonance (MR) and computed tomography (CT) images) often yield unsatisfactory results with 3-D ultrasound. There are a number of properties of ultrasound data that make visualisation and volumetric data analysis

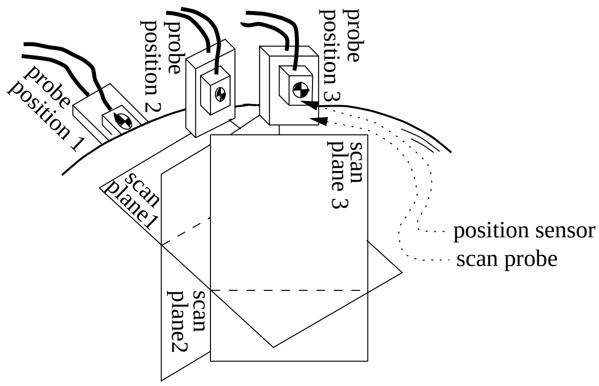

Figure 1: **3-D free-hand ultrasound imaging**. Free-hand imaging allows the physician to move the probe as in a normal ultrasound examination. The position sensor measures the position and orientation of each scan plane. Note that the planes intersect each other.

difficult [17]. Several of the items in the following list are illustrated in Figure 2.

- significant noise and speckle (speckle arises from the constructive-destructive interference of the coherent ultrasound pulses)
- lower dynamic range than MR and CT images
- blurred boundaries around anatomical features
- boundaries with varying intensities caused by changes in surface curvature and orientation
- partially or completely shadowed surfaces from objects nearer the probe
- variable resolution through the volume, dependent on both the spacing between the B-scans and the location within a B-scan

In response to these problems, considerable effort has gone into the development of new segmentation and visualisation techniques for ultrasound data. While some success has been achieved for special cases, such as fetal imaging [17], fast and fully automatic generic techniques remain elusive. Our aim is to exploit *spatial compounding* to improve the quality of the 3-D ultrasound data, so that existing visualisation and analysis techniques are more successful.

The principle behind spatial compounding is to image the region of interest repeatedly, from different look directions (also called isonation angles), and then average the values from the intersecting B-scans when constructing the 3-D data set — see Figure 3. The speckle signal, which de-correlates from different look directions, is suppressed by the averaging operation. Conversely, real anatomical features (tissue boundaries, for example) will be observed in the same location from all look directions. Provided the registration of the scan planes is accurate, the averaging operation will highlight the real anatomical features. Furthermore, the variety of look directions ensures that boundaries become more homogeneous and continuous, and shadowed regions are filled in.

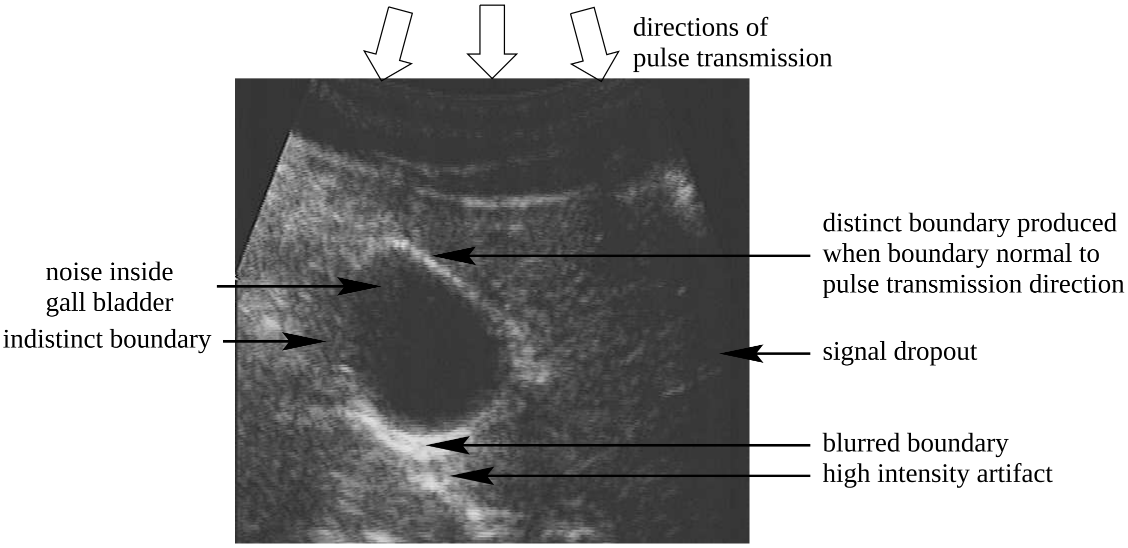

Figure 2: **B-scan of a human gall bladder**. The gall bladder is the dark circular region just left of center. Most of the properties of ultrasound that inhibit automatic analysis are present in this image. In particular, the speckle phenomenon is visible as a granular texture throughout the image. The intensity of the gall bladder boundary varies arounds its length. Sections of the boundary that are perpendicular to the pulse transmission direction produce high intensity echoes, and sections that are parallel produce low intensity echoes.

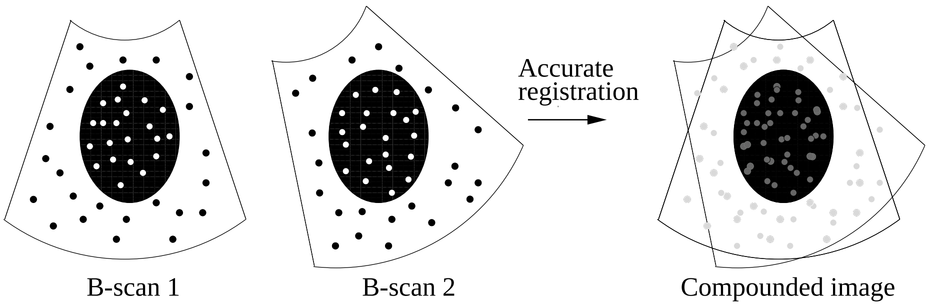

Figure 3: **Spatial compounding**. This is a simple illustration of 2-D spatial compounding. Two scans of the same plane are accurately registered and then averaged together to produce a compounded image with an improved SNR. The principle extends to 3-D, where compounding can be performed wherever scan planes intersect — see Figure 1.

Every free-hand system has to deal with compounding in some manner, since it is almost inevitable that the scan planes intersect. In this paper we propose deliberate, extensive compounding, with the aim of producing high quality 3-D data sets which lend themselves to clear visualisation and accurate volumetric analysis. Spatial compounding can also be performed on data gathered by 2-D array or swept volume systems. The only requirement is that the volume of data is reconstructed from pulse-echo information originating from different look directions. Two volumes of data acquired from similar look directions will not produce improved compounded images since no new information is available.

Although a considerable amount of research has been performed on spatial compounding of 2-D ultrasound images, we are aware of only two published articles on 3-D spatial compounding. In the first article [11], two 3-D data sets were registered using *manual* landmark matching. This constitutes a labour intensive solution to a specific registration problem. The second article [13] cited the improvements made possible by 3-D compounding, but simply stated the need for accurate registration without providing further detail.

Indeed, the key to effective spatial compounding is to achieve a sufficiently high registration accuracy. Registration errors will place the same anatomical feature seen from different look directions at different positions in the reconstructed volume. This phenomenon will result in a blurring of the imaged features. While the largest sources of registration error are likely to be from inaccurate B-scan position measurement and tissue motion during the scan, refraction of the ultrasound beam and other imaging effects also contribute [25]. The accumulated error is, in practice, significant, and it is therefore necessary to use image-based registration techniques to achieve the required level of accuracy.

We have already completed a pilot study into spatial compounding with image-based registration [15]. The study used an organ phantom (an artificial object that mimics the properties of an internal organ) to demonstrate that the registration errors can be corrected. We also showed that the SNR of the images increases with the square root of the average number of B-scans intersecting each voxel (in agreement with statistical theory). Furthermore, segmentation and volume estimates are improved by spatial compounding.

To continue this line of research it was necessary to redesign the free-hand acquisition system so that B-scans could be grabbed at the full video frame rate. This allows a heavily compounded data set to be acquired in the course of a short examination. This paper presents the first results of *in-vivo* examinations using the new acquisition system. In Section 2 we describe the algorithm to reconstruct a regular 3-D data set from a set of B-scans. The registration procedure is then described in Section 3. Results of spatial compounding of a human gall bladder are presented in Section 4, followed by conclusions and suggestions for future work in Section 5.

2 3-D reconstruction

The free-hand acquisition system comprises an ultrasound scanner, a standard probe and a position sensor. A Toshiba model SSA-270A/HG scanner was used with a 3.75 MHz convex curvilinear array probe. 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 scanner are recorded by an 8-bit frame grabber at a rate of 25 frames/s. The images and the position data are stored in the memory of a Silicon Graphics Indy workstation. The set of acquired B-scans is then used to

1. acquire 2-D image $\mathbf{P}$ and associated position data ${}^{T}\mathbf{T}_{R}$
2. insert image $\mathbf{P}$ into reconstruction volume $\mathbf{C}$
   - 2.1 determine location of pixel p_{mn} with respect to $\mathbf{C}$
     ${}^{C}\underline{\mathbf{x}} = {}^{C}\mathbf{T}_{T}\,{}^{T}\mathbf{T}_{R}\,{}^{R}\mathbf{T}_{P}\,{}^{P}\underline{\mathbf{x}} \longleftrightarrow {}^{C}\underline{\mathbf{x}} = \mathbf{T}\,{}^{P}\underline{\mathbf{x}}$
   - 2.2 if nearest voxel c_{ijk} in $\mathbf{C}$ 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 from step 1.

Figure 4: Reconstruction algorithm.

reconstruct a regular 3-D data set, as required by most visualisation and data analysis software packages. The reconstruction algorithm is illustrated in Figure 4, with a detailed description following below.

Each B-scan is represented as a 2-D array $\mathbf{P}$ of intensities p_{mn}. The reconstruction volume takes the form of a 3-D voxel array (or cuberille) $\mathbf{C}$. Each element c_{ijk} of $\mathbf{C}$ represents a voxel in space. The voxel size is chosen *a-priori*: small voxels (though no smaller than the pixel dimensions) produce high resolution reconstructions, larger voxels produce lower resolution reconstructions. While high resolution reconstructions reveal more detail, they also require considerable computational resources to generate and manipulate. There is a fundamental tradeoff between ease of data manipulation and resolution.

Figure 5 depicts the four coordinate systems used for reconstruction. The position sensor measures the relative position and orientation of the receiver with respect to the transmitter. These measurements are converted into a 4×4 homogeneous transformation matrix ${}^{T}\mathbf{T}_{R}$. A standard notation is used to describe ${}^{T}\mathbf{T}_{R}$ as the transformation *from* the coordinate system at the receiver (R) *to* the coordinate system at the transmitter (T).

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. ${}^{C}\mathbf{T}_{T}$ 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 remains constant throughout the reconstruction. ${}^{R}\mathbf{T}_{P}$ describes the transformation between the corner of the scan plane and the coordinate system of the receiver. It also remains constant throughout the reconstruction and is determined by calibration. The cumulative matrix multiplication of ${}^{C}\mathbf{T}_{T}\,{}^{T}\mathbf{T}_{R}\,{}^{R}\mathbf{T}_{P}$ is abbreviated to $\mathbf{T}$.

Determination of ${}^{R}\mathbf{T}_{P}$ is performed using a calibration rig comprising three orthogonal wires mounted in a water bath [2]. The three wires are scanned separately from a variety of look directions. The orthogonality of the wires can be used to write one equation in ${}^{P}\underline{\mathbf{x}}$, ${}^{T}\mathbf{T}_{R}$ and ${}^{R}\mathbf{T}_{P}$ for each B-scan. To improve the accuracy of the calibration, it is advantageous to acquire more B-scans than are necessary to determine the six parameters of ${}^{R}\mathbf{T}_{P}$, and use an iterative least squares algorithm to solve the over-determined

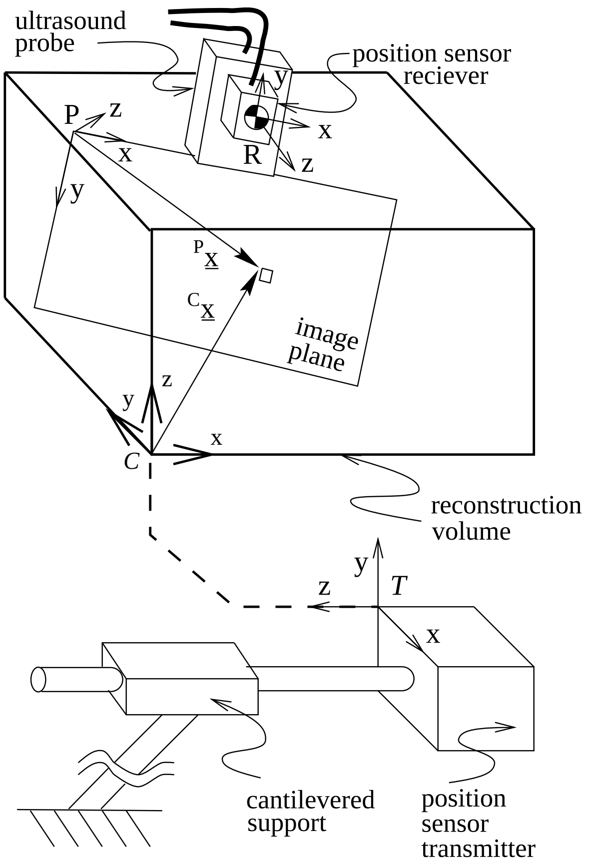

Figure 5: Coordinate systems used for reconstruction. The following notation is used to label coordinate systems: **T** = Transmitter, **C** = Cuberille, **P** = Plane and **R** = Receiver

problem [2]. Studies have shown this method to be accurate to 1 mm [2].

Before the start of the examination, the voxels in the reconstruction volume are all set to zero. As each B-scan is acquired, each voxel c_{ijk} is adjusted according to the pixels p_{mn} that intersect it. A single voxel will envelop many pixels if the voxel size is larger than the B-scan pixel size. Each voxel may also be intersected again by future B-scans. These possibilities are dealt with by step 2.3 of the reconstruction algorithm, which describes a compounding operation to average all pixels that intersect a voxel.

After a substantial portion of **C** is filled, it can be displayed on a computer monitor by several different methods, including surface rendering and any-plane slicing. Hereafter, the term **slice** is used to indicate an image produced by any-plane slicing.

3 3-D registration

In a typical *in-vivo* ultrasound examination, the physician scans smoothly and continuously a region of interest from one extent to the other. Each sweep of the region can be used to reconstruct a 3-D data set. For this study we examined the gall bladder of a healthy human subject. The subject was requested to lie motionless on a bed and to breathe lightly to minimise motion-induced errors during scanning. Approximately 100 B-scans

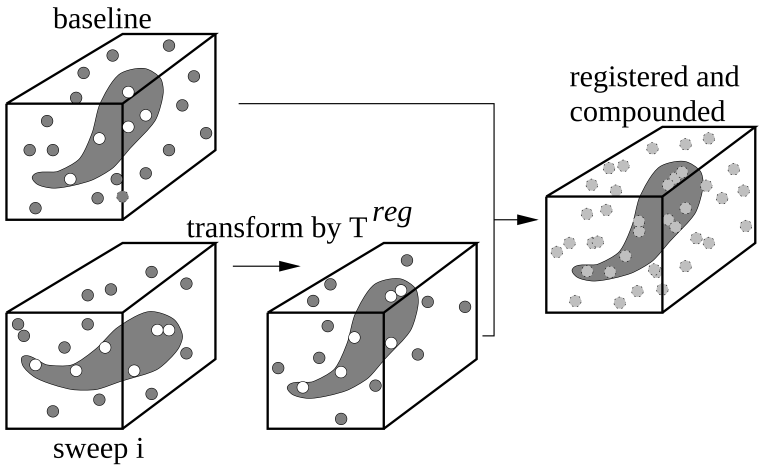

Figure 6: **Registration and spatial compounding of noisy 3-D data**. The first sweep is called the *baseline* since all other sweeps are registered to it. $\mathbf{T}^{reg}$ is the homogeneous transformation matrix describing the rigid body registration. Transformation of sweep i by $\mathbf{T}^{reg}$ aligns it with the baseline data so that compounding can be performed with minimal loss of spatial detail. Since sweep i is reconstructed from B-scans taken from different look directions to the baseline B-scans, the noise is uncorrelated. Compounding the two registered data sets therefore reduces the level of noise.

are required to cover the gall bladder from tip to tail with sufficient resolution to fill a data set with 1 mm voxels without gaps. Since our acquisition system records B-scans at 25 frames/s, 600 B-scans can be acquired in only 24 seconds. This allows six sweeps of the organ during that time. The look direction of the probe (i.e. window location and isonation angle) is changed slightly between each sweep so that different views of the gall bladder are obtained.

The basic idea of spatial compounding is to combine the data from each of the six sweeps into a single 3-D data set. If no errors are present, the six sweeps can be combined using only the position sensor readings. The manufacturer states that the RMS (root mean squared) accuracy of the FASTRAK position measurements is 0.8 mm with a resolution of 0.38 mm. The angle measurements are stated to have an RMS accuracy of 0.15° with a resolution of 0.025°. However, the accuracy of imaging a point is determined by the accumulated sources of error at each stage of the data acquisition process, including the speed of sound estimate, the FASTRAK readings and the calibration of the FASTRAK receiver to the plane of the B-scan. In practice, therefore, image-based registration is required to align the sweeps.

Spatial compounding is performed by accurately registering the last five sweeps to the first *baseline* sweep. Even though some nonrigid organ motion is undoubtedly present, we have found that the majority of the registration errors are caused by inaccurate position measurements and can be corrected by rigid body transformation [15]. The exception is when imaging pulsatile blood flows, when it is generally necessary to gate the B-scan acquisition to an ECG signal [18]. The registration process is illustrated in Figure 6.

There has been considerable research into the automatic registration of CT and MRI images, but very few attempts to apply these techniques to 3-D ultrasound. Two comprehensive surveys of the various registration techniques can be found in [21, 24]. We evaluated a number of the more popular approaches on the ultrasound data, with varying degrees of success.

The first class of techniques we tested were those based on *landmark matching*. The difficulty with these techniques lies in identifying suitable landmarks to match. For example, if it were possible to extract the organ boundaries from two data sets, it would be feasible to manually align the two surfaces (given suitable visual feedback). However, it is precisely because such features are difficult to extract from ultrasound data that we are attempting registration and compounding. Even with suitable surfaces to align, the manual approach is not sufficiently accurate to produce sharp, compounded data.

More accuracy is possible with assisted, point-based landmark matching. In this paradigm, the user identifies matching points in the two data sets, and then the transformation which brings one set of points into close alignment with the other is computed. The problem with this approach is the difficulty in identifying distinguished points on smooth organ boundaries. Other landmark-based approaches include automatic surface-surface matching [8] and automatic curve-curve matching [22]. However, all such techniques rely on the accurate extraction of suitable features to match, which is generally not feasible in ultrasound images.

The second broad class of registration techniques make use of some sort of correlation measure. Such techniques have been found to be robust for multimodal (MR to CT) registration [9, 23]. No user interaction, explicit segmentation or landmark identification is required. High accuracy can be achieved through a coarse-to-fine search, and it is possible to register 3-D data sets that are only partially overlapping. Registering two 3-D ultrasound data sets has much in common with registering an MR data set to a CT one: in both cases, the data sets to be aligned can appear substantially different, while sharing some common features. For this reason, we chose to investigate correlation-based registration of each sweep onto the baseline.

Two volumes of intensity data can be registered by searching the transformation parameter space for a peak in the correlation function. In the current application, we are only considering rigid body transformations, so the parameter space is six-dimensional. For two volumes L_1 and L_2 that only partially overlap, a normalised correlation coefficient is used:

$$corr(\underline{\mathbf{a}}) = \frac{\displaystyle\sum_{(x,y,z)\in L_1} L_1(x,y,z)L_2(\mathbf{T}_a(x,y,z))}{\displaystyle\sum_{(x,y,z)\in L_1} L_2(\mathbf{T}_a(x,y,z))}$$

where $\underline{\mathbf{a}}$ is the vector of the registration parameters, $\mathbf{T_a}$ is the rigid body transformation matrix corresponding to $\underline{\mathbf{a}}$, and the denominator is the portion of L_2 that overlaps L_1.

It remains to decide what features of the data sets to correlate. Since different ultrasound images of the same anatomy can appear substantially different (as explained in Section 1), simple correlation of the raw intensity data is not likely to succeed. Again, we look to multimodal MR-CT registration for inspiration. Raw MR and CT data sets cannot be compared directly, yet they share some common features which can be usefully correlated. For example, differential operators such as *edge* and *ridge* detectors can transform MR and CT scans of the skull into comparable data sets [9].

In a recent comparative study of multimodal 3-D image registration [9], correlation of the magnitude of the 3-D gradient was found to give the best performance. The 3-D gradient magnitude is invariant under the group of orthogonal transformations (translation, rotation, and reflection), and tends to produces better results than Laplacian and ridgeness operators [9]. For this reason, we attempted to register the ultrasound sweeps by correlating the 3-D gradient magnitudes of the data sets. When applied to ultrasound data, the gradient magnitude operator produces a transformed image with local maxima near organ boundaries. We calculated the gradients efficiently using a separable, recursive algorithm [10].

The main drawback of all correlation-based registration techniques is the significant computational expense of the search over the six-dimensional parameter space. However, considerable savings are possible using a multi-resolution approach [23]. A multi-resolution pyramid is constructed with the original data set at the base. Each higher level is produced at half the resolution of the previous level. The pyramid continues until the largest structures are no longer discernible. The basic idea is that wide range searches can be performed efficiently near the top of the pyramid and narrow but accurate searches at the bottom. Results from the searches at each level provide initial search locations at lower levels until the bottom of the pyramid is reached.

We found that the multi-resolution approach works well with the ultrasound registration problem. Since the position sensor readings provide a fairly good initial guess at the correct registration, only two pyramid levels are required. Our final reconstructions use 1 mm voxels, so searches started with a resolution of 2 mm.

The correlation coefficient is sufficiently well behaved to allow two searches, with different step sizes, at each level of the pyramid. So two searches were performed at the 2 mm resolution, the first with large steps to cover a large part of the parameter search space, the second with smaller steps around the first correlation peak. The maximum of the 2 mm search is then used as the center of the search at the 1 mm resolution, where again two searches are performed. Because the initial guess is close to the optimal registration, local minima are not encountered.

$\mathbf{T}^{reg}$ is constructed from the parameter values corresponding to the maximum correlation coefficient. To compound the registered sweep i with the baseline data, each B-scan of sweep i undergoes the additional transformation $\mathbf{T}^{reg}$ (c.f. step 2.1 of Figure 4):

$$ {}^{C}\underline{\mathbf{x}} \;=\; \mathbf{T}\,{}^{P}\underline{\mathbf{x}} \qquad \longrightarrow \qquad {}^{C}\underline{\mathbf{x}} \;=\; \mathbf{T}^{reg}\;\mathbf{T}\;{}^{P}\underline{\mathbf{x}} $$

An overview of compounded reconstruction, using the correlation-based registration method, is given in Figure 7.

4 Results

The gall bladder can be scanned from a number of different look directions, each giving a slightly different view. The variety of look directions used in this study is apparent in Figure 8.

Since all six sweeps (comprising, in total, 600 B-scans) are acquired in a single examination, the scans must be grouped into coherent sweeps. The gall bladder examination was easily partitioned by manually reviewing the B-scans in the order of acquisition. Software is currently being developed to allow the physician to start and stop acquisition using a touch sensitive display. This will allow the physician the flexibility of marking the sweeps during the examination.

1. Establish the baseline from the first sweep
    1.1 create 1 mm and 2 mm resolution data sets
    1.2 calculate the gradients at both resolutions
    1.3 the 1 mm data set is used for compounding the baseline with the following sweeps. (The baseline gradient used for correlation remains unchanged for registration of all sweeps)
2. Register each following sweep to the baseline
    2.1 create 1 mm and 2 mm resolution data sets
    2.1 calculate the gradients at both resolutions
    2.2 correlate the gradient magnitudes with the baseline gradient, first at 2 mm, then at 1 mm resolution
    2.3 $\mathbf{T}^{reg}$ is determined from the maximum of the correlation coefficient
3. Compound sweep data
    3.1 transform each B-scan in the sweep by $\mathbf{T}^{reg}$ and compound with the baseline and previous sweeps using average intensity — see step 2.3 of Figure 4
4. Repeat steps 2 and 3 for each sweep

Figure 7: Registration and compounding of multiple sweeps.

After grouping the B-scans into sweeps, a voxel array can be reconstructed for each sweep using the algorithm in Figure 4. At this stage it is possible to appreciate the registration errors by viewing corresponding slices through the reconstructions — see Figure 9. Compounding the sweeps using the position measurements alone gives unsatisfactory results, as can be seen in Figure 10. Clearly, image-based registration is essential for accurate spatial compounding.

The first stage of registration is the calculation of the gradients. Gradients calculated at both 2 mm and 1 mm resolutions are shown in Figure 11. The scale of the gradient operator was chosen to minimise the effect of speckle, yet retain the features of the anatomical structures. The correlation calculations were performed on a Silicon Graphics Indigo 2 Impact 10000 workstation, taking 3.6 hours per sweep. This includes calculation of the gradients and the two searches at both resolutions. While this is undoubtedly expensive, the procedure is readily parallelised and the times are comparable with those accepted for CT and MRI registration [23]. The registration results for each sweep are listed in Table 1.

Slices of the registered and compounded data set are shown in Figure 12. The borders of the gall bladder are all well aligned, compared to the unregistered case in Figure 10. The level of speckle noise is also reduced compared to the single sweep reconstructions of Figure 9. Furthermore, the organ boundaries of the registered and compounded reconstruction have better contrast and are more uniform than in the single sweep reconstructions. While all of these improvements are benefits of spatial compounding, it has to be admitted that, on close inspection, a slight loss of spatial detail is apparent. The residual

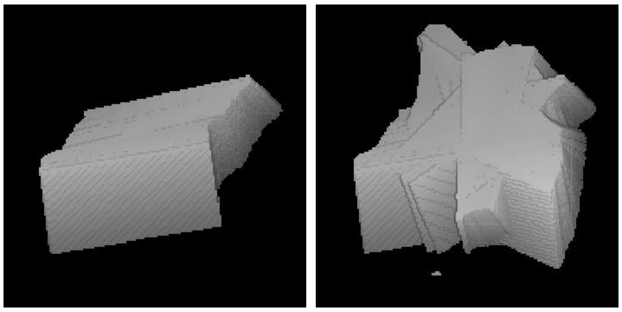

(a) single sweep (b) compounded

Figure 8: **Volume of reconstruction filled by sweeps**. Figure (a) shows the volume of the reconstruction filled by the baseline sweep 0. Figure (b) shows the volume filled by sweeps 0 to 5. Note how the volume is enlarged by compounding, since the individual sweeps only partially overlap each other.

Sweep	Registration Error					
	x (mm)	y (mm)	z (mm)	α (degrees)	β (degrees)	γ (degrees)
1	-2.0	9.0	0.0	-6.3	-1.7	0.6
2	-3.0	16.0	-4.8	-10.9	-3.4	2.3
3	1.0	-2.0	-9.0	-5.2	-8.6	-2.3
4	0.0	8.7	-16.0	-2.8	-6.9	4.6
5	10.0	-12.0	-10.0	6.9	-8.0	-3.4

Table 1: **Registration errors**. The six parameters of $\mathbf{T}^{reg}$ are listed. Translation is expressed as x,y,z and rotation about the z, y and x axes as α, β and γ respectively.

registration errors that are not corrected (such as non-rigid motion of the organ and local refraction of the ultrasound beam) result in a slight blurring of small features. The main benefits of spatial compounding are therefore improved segmentation and visualisation of larger structures.

Segmentation, the process of separating a particular anatomical region of interest from the surrounding data, is a key step in visualisation and volume estimation. Perfect segmentation of the gall bladder requires that all parts of the organ are accurately identified and no regions outside the organ are included.

In order to objectively compare the quality of the compounded reconstruction with the single sweep reconstructions, automatic segmentation by intensity thresholding was attempted. Figures 13 and 14 shows how spatial compounding improves the ability to segment by thresholding. Essentially, spatial compounding has the ability to fill in regions that are not well defined in some sweeps, but are clearer in others. Note that this improvement is not achievable by filtering, which introduces no new information.

The reduction of speckle noise by spatial compounding is difficult to see in the reconstruction slices. It is more clearly evident in 3-D surface renderings of original (unprocessed) 3-D data sets. Figure 15 shows a significant amount of speckle reduction by spatial

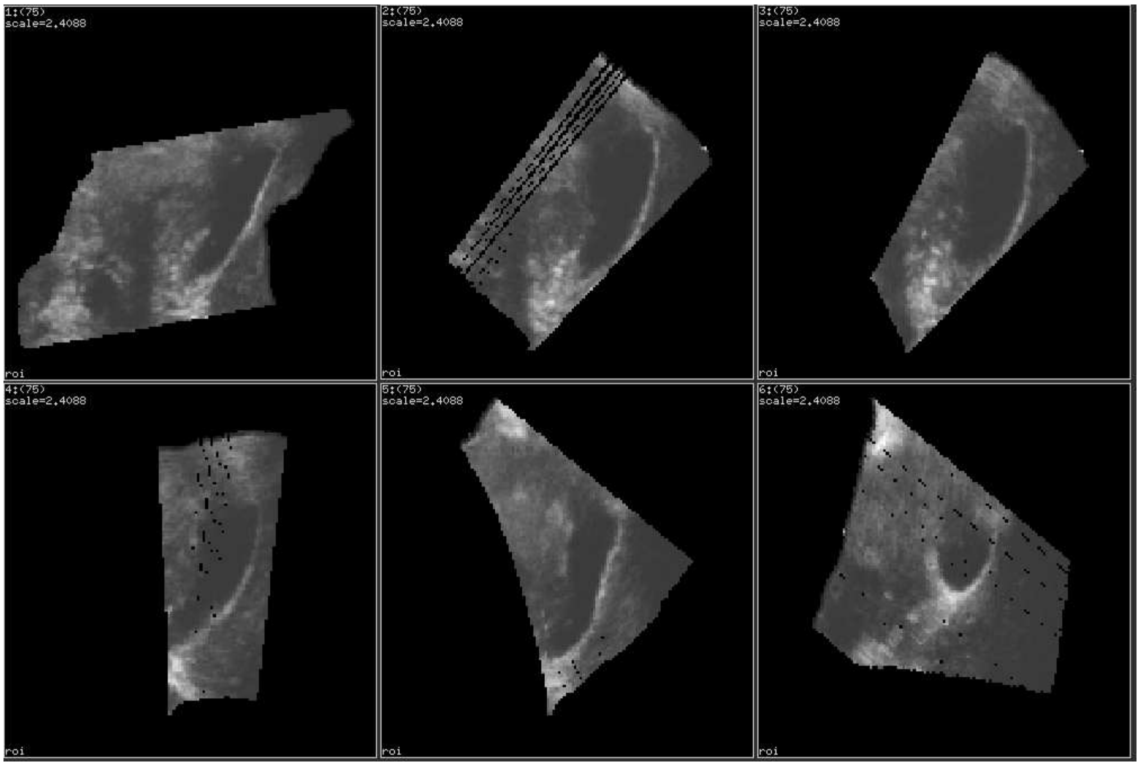

Figure 9: **Slices of reconstructions of individual sweeps *before* registration**. Sweeps 0 to 5 are shown in two rows, starting from top left. All slices are taken at the same location in the reconstruction, identical to the locations used in Figures 10 and 12. The data has not undergone any processing, so gaps remain in some reconstructions where the B-scans are widely spaced. The cross sections of the gall bladder appear differently in each sweep, indicating significant registration errors.

compounding. The compounding allows a significant fraction of the gall bladder surface to be viewed without any filtering.

Further improvements are possible with a small amount of data processing before segmentation and rendering. A typical ultrasound visualisation procedure involves interpolating and filtering the data before surface detection is attempted [12, 16]. Figure 16 shows that interpolation and filtering of single sweep reconstructions improves visualisation, but problems still remain. The main problem is that ill-defined boundaries (like the ones in Figure 13) are not improved by interpolation and filtering. The segmentation does not isolate the gall bladder from the surrounding noise and artifacts in single sweep reconstructions. Performing the same data processing on the registered and compounded data *does* isolate the gall bladder, as shown in Figure 17. Removal of the unwanted, disconnected regions becomes a trivial task, with the results shown in Figure 18.

Organ volume is a measure that is often sought after by physicians. Measuring changes in the volume of an organ over time is often used to monitor the progression of a disease or its response to treatment. For example, it is useful to measure the reduction in prostate size during reductase inhibitor therapy, and changes in splenic size during enzyme replacement therapy for Gaucher's disease. There is also a need to measure changes in the size of uterine

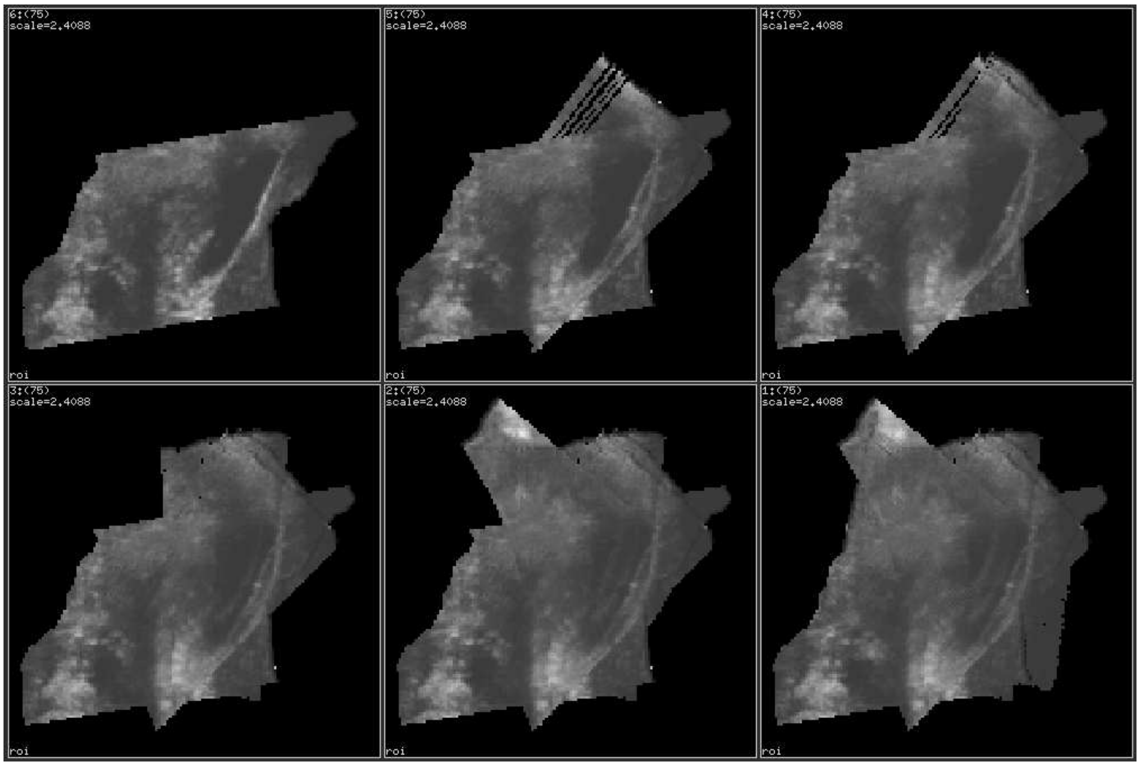

Figure 10: **Slices of compounded reconstructions *before* registration**. Slices of reconstructions compounded incrementally with sweeps 0 to 5 are shown in two rows starting from top left. All slices are taken at the same location in the reconstruction volume, identical to the locations used in Figures 9 and 12. As more and more sweeps are compounded together using the position sensor data alone, the gall bladder becomes more blurred and eventually indistinguishable. Image-based registration is therefore required for accurate spatial compounding.

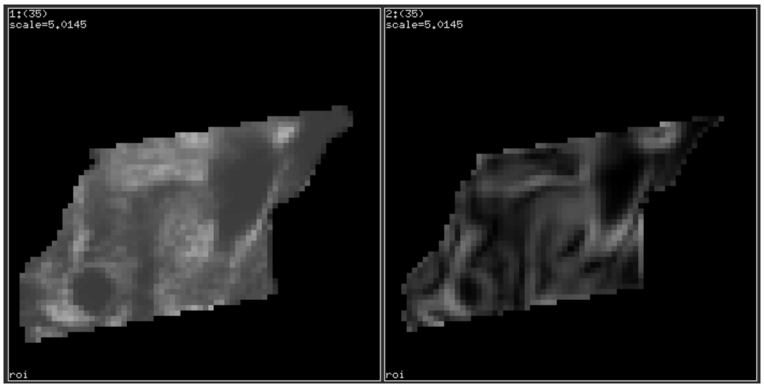

(a) 2 mm resolution

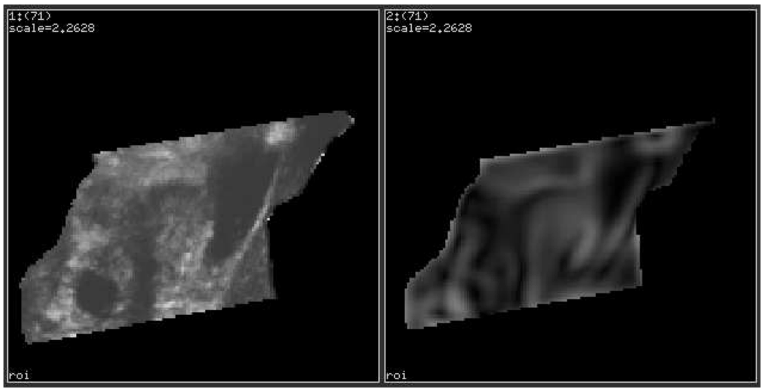

(b) 1 mm resolution

Figure 11: **Differential geometric features used for registration.** The left images are slices of the original data and the right are the magnitudes of the 3-D gradient. The features used for correlation at 2 mm resolution are shown in (a), and the features used for correlation at 1 mm resolution are shown in (b).

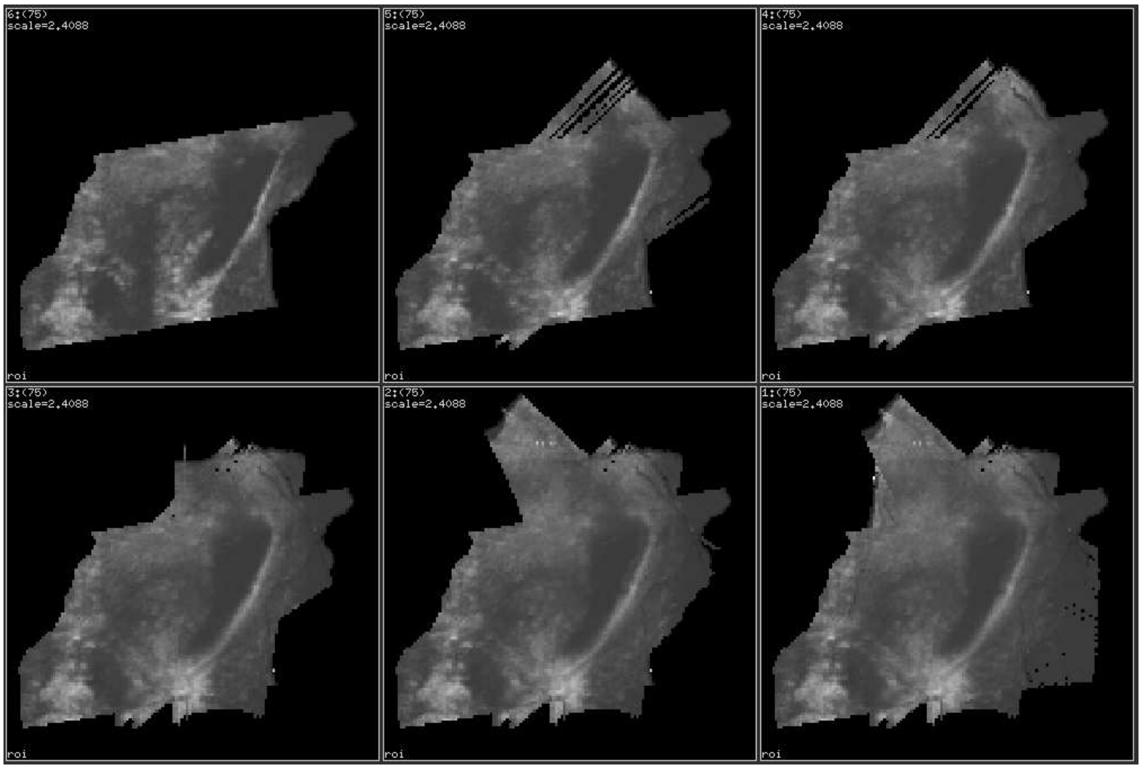

Figure 12: **Slices of compounded and registered reconstructions**. Slices of reconstructions compounded incrementally with sweeps 0 to 5 are shown in two rows, starting from top left. All slices are taken at the same location in the reconstruction volume, identical to the locations used in Figures 9 and 10. The registration has aligned the organ boundaries in the individual sweeps, so the fully compounded data set retains the original organ shape. Compounding has also reduced speckle noise and made the gall bladder more distinct.

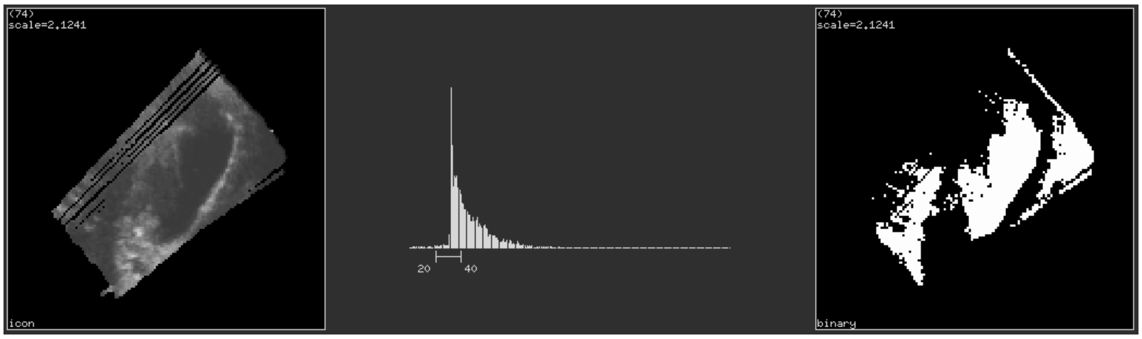

Figure 13: **Threshold segmentation of reconstructions from a single sweep.** A slice of the reconstruction from sweep 1 is shown on the left, identical to the location used in Figure 14. The slice's grey-level histogram is shown in the middle. The data has not undergone any processing. The segmentation by thresholding the intensities within the range [20, 40] is shown on the right. The problem with segmentation of ultrasound data by thresholding is that the organ is not easily isolated from the surrounding noise and artifacts. In this case, the image has a dark region from signal dropout near the gall bladder, which affects the segmentation. A narrower threshold range reduces the volume outside the gall bladder that falls within the range, but also creates gaps within the gall bladder. The range of [20, 40] represents the best tradeoff between the two throughout the reconstruction.

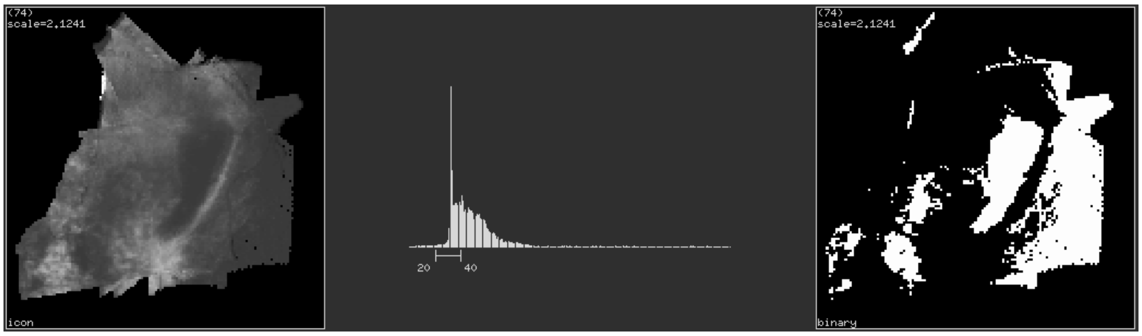

Figure 14: **Threshold segmentation of compounded reconstruction**. The segmentation has been performed in an identical manner to Figure 13, yet the segmentation is more accurate.

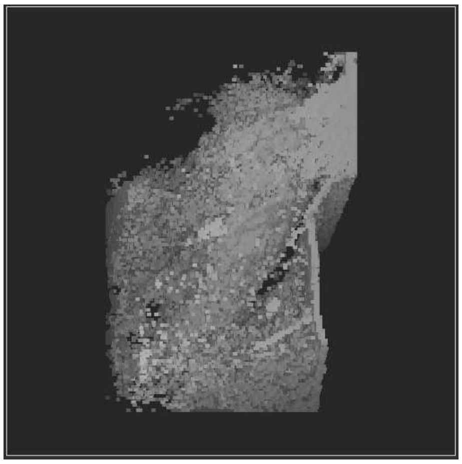

(a) single sweep

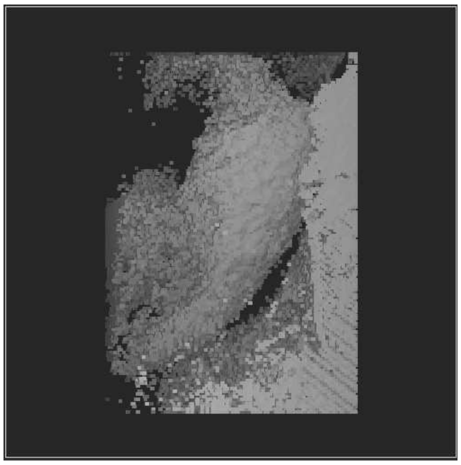

(b) compounded

Figure 15: **3-D segmentation of unprocessed data**. Surfaces are extracted by thresholding, similar to the method illustrated in Figures 13 and 14. No other data processing has been performed. Spatial compounding clearly suppresses the noise, making the surface more easy to detect.

leiomyomas (fibroids) in response to treatment by luteinising hormone releasing hormone antagonists.

Volume estimation follows a similar trend to visualisation since both rely upon accurate segmentation. Since the true volume of the gall bladder is unknown, we chose to use live-wire semi-automatic segmentation [1] of the baseline sweep as a "gold standard". Live-wire segmentation is a powerful tool for extracting boundaries in noisy images. It offers a good compromise between accuracy and amount of user intervention. The technique involves laying an active wire around the object on a slice by slice basis. The wire is attracted automatically to the object's boundary. The operator assists the live-wire by depositing small sections at a time near the boundary, so that the wire does not enclose nearby speckle.

In addition to the live-wire estimate, volumes for individual sweeps and the compounded reconstruction were calculated using the same process (interpolation, filtering and thresholding) used for the visualisation results in Figures 16 and 17. In order to calculate only the gall bladder volume, any segmented part of the data that was not connected to the gall bladder was removed manually. The volume was then calculated by summing the remaining filled voxels that contain the segmented gall bladder.

The calculated volume of the compounded data is closer to the live-wire estimate than any of the individual sweeps. The volumes for each of the individual sweeps are higher than the live-wire volume because the segmentation of the gall bladder typically exceeds the gall bladder boundary (as evident in Figure 13). This again shows that segmentation of spatially compounded data sets gives better results than single sweeps.

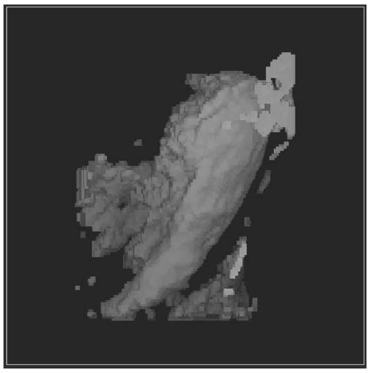

(a) sweep 0

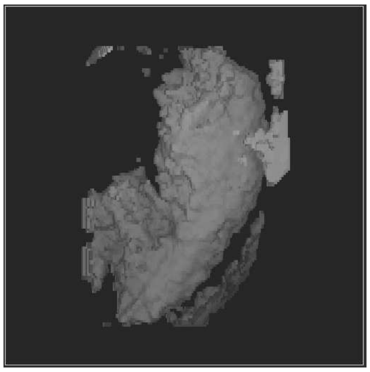

(b) sweep 1

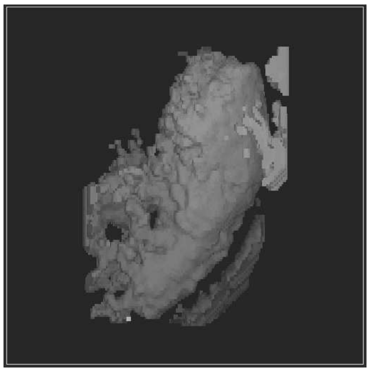

(c) sweep 2

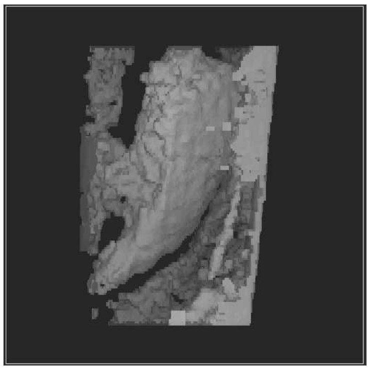

(d) sweep 3

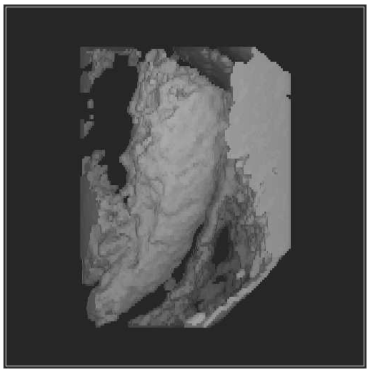

(e) sweep 4

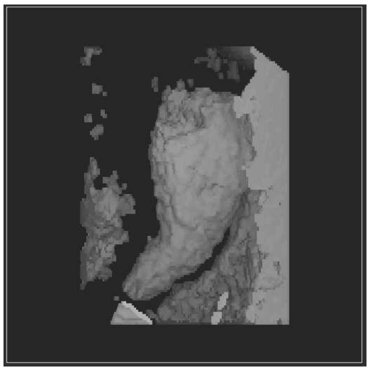

(f) sweep 5

Figure 16: **3-D segmentation of individual sweeps after data processing.** The reconstructed 3-D data sets are interpolated to fill gaps, and Gaussian filtered to remove speckle. Segmentation is then performed by thresholding the processed data. All images are produced with identical filter sizes, interpolation parameters and threshold ranges. Compared with Figure 15 (a), a considerable amount of noise has been removed by processing. In each sweep, however, the gall bladder is not disconnected from the surrounding noise and artifacts.

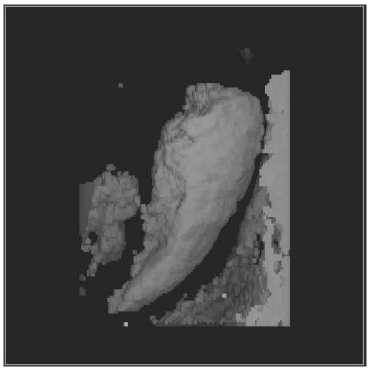

Figure 17: **3-D segmentation of compounded reconstruction after data processing.** This image is produced in the same manner as in Figure 16. Spatial compounding improves the organ segmentation. Compared to the segmentation of the individual sweeps, the compounded data set is the only one in which the gall bladder is completely disconnected from surrounding noise and artifacts. Complete isolation of the gall bladder is now easily performed — see Figure 18.

Figure 18: **3-D segmentation of high resolution compounded reconstruction.** This image is produced from a high resolution reconstruction compounded with sweeps 0 to 5. The data is first interpolated, median filtered, threshold segmented then manually trimmed of extraneous regions.

Sweep	Volume (mm^3)
0	13494
1	16662
2	15758
3	14945
4	15117
5	14166
compounded	13039
live-wire	12507

Table 2: **Volume of gall bladder**. The compounded data set provides the volume estimate closest to the "gold standard" live-wire segmentation.

5 Conclusions and future work

Ultrasound data is more difficult to register than data from other modalities such as MRI and CT. We have demonstrated, for the first time, the registration of 3-D ultrasound data using a technique developed for MR to CT registration. The registration technique was chosen according to the particular requirements of ultrasound data.

Registration is performed in a fully automatic manner, eliminating user subjectivity. The registration is based on the correlation of features obtained by calculation of the 3-D gradient, so neither explicit segmentation nor landmark identification is required. Implementation of the registration technique on *in-vivo* ultrasound data has shown it to be robust and accurate. The efficiency of the algorithm is also improved via a multi-resolution search for the correlation peak.

The main application we have considered in this paper is *spatial compounding*. We use a free-hand ultrasound acquisition system that can acquire many overlapping B-scans in a short period of time. The rapid acquisition allows spatial compounding of multiple 3-D data sets obtained from different views of the same organ in a single examination.

Although system calibration is performed prior to each examination, registration errors remain in the 3-D data sets. Accurate spatial compounding requires correction of the registration errors by image-based registration. We have shown that rigid body registration is sufficient to correct a large proportion of the registration errors. Since the registration errors are known to be small, registration by correlation is feasible by searching a small portion of the six-dimensional parameter space.

The resulting registered and compounded 3-D data sets offer improved segmentation for better visualisation and volume estimation. Application of state-of-the-art visualisation techniques to the compounded data has produced dramatically clear and accurate images of the gall bladder. Other users of 3-D ultrasound can use the techniques described in this paper, since they can be applied to any regular 3-D ultrasound data sets, independent of the particular method of data acquisition.

Since correlation-based registration can register volumes that are only partially overlapping, it is possible to combine several small volumes of data to form a larger volume. An example is shown in Figure 19. Accurate registration of the overlapping volumes will be required to avoid discontinuities in the data. This concept is especially useful for data acquisition systems that can examine only small volumes at a time, such as the commercially available Kretztechnik system.

Many conventional ultrasound scanners are also capable of producing Doppler ultra-

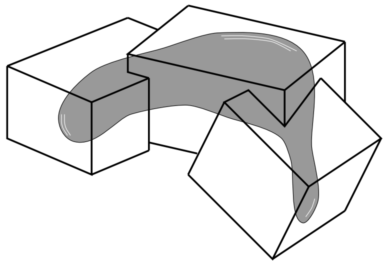

Figure 19: **Extended volume of interest**. Several small, partially overlapping volumes can be combined to encompass a larger organ.

sound images. In Doppler ultrasound, pulse-echo information can be used to calculate the speed of fluids passing through vessels, providing important vascular information. While Doppler ultrasound can provide information no other imaging modality can, it suffers from considerable signal dropout and noise. Spatial compounding of 3-D Doppler ultrasound data is another very promising direction for future research.

Acknowledgements

Calculations of the 3-D gradients were performed by a separable, recursive 3-D edge detector, kindly provided by the researchers at INRIA, Sophia Antipolis, France. Ultrasound scanning was performed with the help of Jonathon Carr. The free-hand acquisition system was developed by Richard Prager and Patrick Gosling, and the 3-D renderings were produced using the 3DViewnix visualisation package. Robert Rohling is supported by Churchill College and an ORS award.

References

[1] W.A. Barrett and E.N. Mortensen. Fast, accurate, and reproducible live-wire boundary extraction. In K.H. Hohne and R. Kikinis, editors, *Lecture Notes in Computer Science: Visualization in Biomedical Computing, Proc. VBC '96*, volume 1131, pages 183–192. Springer-Verlag, Heidelberg, Germany, 1996.

[2] J. Carr. *Surface Reconstruction in 3D Medical Imaging.* Ph.D. Thesis, University of Canterbury, Christchurch, New Zealand, 1996.

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

[4] 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. *Br. J. Radiology*, 64:1056–1057, 1991.

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

[6] 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 Med. Biol.*, 21(1):25–32, 1995.

[7] F. Hottier and A. Collet Billon. 3D echography: status and perspective. In K.H. Hohne, H. Fuchs, and S.M. Pizer, editors, *3D Imaging in Medicine: Algorithms, Systems, Applications*, pages 21–41. Springer-Verlag, Berlin, Germany, 1990.

[8] D. N. Levine, C. A. Pelizzari, G. T. Y. Chen, C. T. Chen, and M. D. Cooper. Retrospective geometric correlation of MR, CT and PET images. *Radiology*, 169(3):817–823, 1988.

[9] J. B. A. Maintz, P. A. van den Elsen, and M. A. Viergever. Comparison of edge-based and ridge-based registration of CT and MR brain images. *Medical Image Analysis*, 1(2):151–161, 1996.

[10] O. Monga, R. Deriche, G. Malandain, and J-P. Cocquerez. Recursive filtering and edge tracking: two primary tools for 3-D edge detection. *Image and Vision Computing*, 9(4):203–214, August 1991.

[11] 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 Med. Biol.*, 21(6):769–778, 1995.

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

[13] T.R. Nelson and D.H. Pretorius. 3D ultrasound image quality improvement using spatial compounding and 3D filtering. *Med. Phys.*, 21(6):998, 1994.

[14] 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.

[15] R. Rohling, A. Gee, and L. Berman. 3-D spatial compounding of ultrasound images. *Medical Image Analysis*, 1(3):177–193, April 1997.

[16] 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.

[17] G. Sakas and W. Stefan. Extracting surfaces from fuzzy 3D-ultrasound data. In Robert Cook, editor, *Proceedings of SIGRAPH*, pages 465–474. Addison Wesley, August 1995.

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

[19] S. W. Smith, R. E. Davidsen, C. D. Emery, R. L. Goldberg, and E. D. Light. Update on 2-D array transducers for medical ultrasound. In *Proc. of IEEE Ultrasonics Symposium, 1995*, pages 1273–1278, 1995.

[20] 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.

[21] R. H. Taylor, S. Lavallee, G. C. Burdea, and R. Mosges, editors. *Computer-Integrated Surgery*. MIT Press, Cambridge, Massachusetts, USA, 1996.

[22] J. P. Thirion, O. Monga, S. Benayoun, A. Gueziec, and N. Ayache. Automatic registration of 3D images using surface curvature. In *Proc. IEEE Intl. Symp. on Optical Applied Science and Engineering*, San Diego, USA, 1992.

[23] P. A. van den Elsen, J. B. A. Maintz, E. J. D. Pol, and M. A. Viergever. Automatic registration of CT and MR brain images using correlation of geometrical features. *IEEE Trans. Medical Imaging*, 14(2):384–396, 1995.

[24] P.A. van den Elsen, E. J. D. Pol, and M.A. Viergever. Medical image matching: a review with classification. *IEEE Engineering in Medicine and Biology*, 12(1):26–39, 1993.

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