
**SPATIAL COMPOUNDING OF
3-D ULTRASOUND IMAGES**

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

CUED/F-INFENG/TR 270

October 1996

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, especially 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 reduce speckle. We develop a new statistical theory to predict the improvement in signal to noise ratio with increased levels of compounding, and verify the predictions empirically. We also investigate how registration errors can affect automatic volume estimation of structures within the compounded 3-D data set. Having established the need to correct these errors, we present a novel reconstruction algorithm which uses landmarks to accurately register each B-scan as it is inserted into the voxel array. In a series of *in-vitro* and *in-vivo* trials, we demonstrate that 3-D spatial compounding is very effective for improving the signal to noise ratio, but correction of registration errors is essential.

1 Introduction

Conventional diagnostic ultrasound 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. One of the limitations of conventional imaging is the requirement that the physician mentally reconstruct 3-D anatomy given multiple 2-D slices. Research is underway to overcome this limitation using **3-D free-hand ultrasound imaging**. In this 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 using CT or MRI, but with less expensive and invasive technology. Physicians have indicated that there is significant utility in 3-D ultrasound imaging of a variety of anatomical structures, including the fetus [5], vascular structure [10], gall bladder [8], breast [17], kidney [11], and heart [21].




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.




Figure 2: **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.

In a review article about 3-D ultrasound [14], the author suggests that 3-D visualisation and volume estimation are its most attractive capabilities. The main difficulty in performing these tasks automatically is that the quality of the ultrasound data can be severely degraded by artifacts, especially speckle. Speckle is an artifact common to all ultrasound images, a product of the constructive-destructive interference of the ultrasound echoes. While trained ultrasonographers can occasionally infer diagnostically useful information from speckle patterns, speckle is a nuisance when attempting automatic analysis of ultrasound data. For this reason there has been considerable research into speckle reduction.

There are three main techniques for reducing speckle in ultrasound images [7]: compounding (spatial [12, 15, 22, 24] and frequency [25, 26]), filtering [6, 7, 20], and phase-based methods [16]. While each technique has its advantages and disadvantages, spatial compounding is particularly attractive with 3-D ultrasound, since it can be performed with a standard free-hand system without the need for any modifications.

The principle behind spatial compounding is to image the region of interest repeatedly, from different look directions, and then average the values from the intersecting B-scans when constructing the 3-D data set — see Figure 2. 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.

There has been considerable research in the past on the spatial compounding of multiple B-scans for 2-D image quality improvement. Defining the signal to noise ratio (SNR) as the ratio of the mean grey level to the standard deviation for an image with no resolvable structures¹, it has been established that compounding multiple B-scans lying in the same plane improves the SNR by the square root of the number of B-scans used [3, 12, 15].

The main drawback of spatial compounding is that it requires accurate registration: registration errors will place the same anatomical feature seen from different look directions at different positions in the reconstructed volume. This phenomenon can be seen in Figure 12, which shows a reduction in speckle as more B-scans are averaged, accompanied by a blurring of the imaged boundaries. While the largest sources of registration error are

¹This is a commonly used figure of merit for imaging systems.

likely to be from inaccurate B-scan position measurement and tissue motion during the scan [19], refraction of the ultrasound beam and other imaging effects also contribute.

In this paper we take a fresh look at spatial compounding, this time from the perspective of 3-D free-hand ultrasound. Though our main focus is on speckle reduction, spatial compounding can also improve image quality by reducing several other kinds of artifacts, such as shadowing. Every free-hand system has to deal with compounding in some manner, since it is almost inevitable that the scan planes will intersect. Here we propose deliberate, extensive compounding, with the aim of producing high quality 3-D data sets which lend themselves to automatic segmentation for visualisation and volume measurement. The key to effective spatial compounding is to achieve a sufficiently high registration accuracy. Relying on the position sensor alone is usually not sufficient: there may be small errors in its calibration, and it does not take into account motion of the target or within-plane imaging artifacts. It is therefore necessary to improve the registration using image-based techniques.

Little work on the registration of 3-D ultrasound data sets is evident in the literature. In one exception [17], 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. Similarly, little work on 3-D ultrasound compounding has been performed. Only one brief article [18] has cited the improvements possible by 3-D compounding, but it simply stated the need for accurate registration without providing further detail.

In this paper we tackle three major objectives. The first is to demonstrate how accurate 3-D registration can be achieved. We describe a technique that takes sensor-based measurements of B-scans positions and applies small adjustments to align anatomical landmarks in the reconstructed volume. The second objective is to demonstrate how spatial compounding, coupled with accurate registration, can dramatically improve the SNR of the reconstructed data. Our final objective is to develop a statistical theory of 3-D spatial compounding and establish agreement between the observed and predicted improvements in SNR. In the course of the investigation we perform two empirical studies: an *in-vitro* phantom study and an *in-vivo* study. The phantom study allows the 3-D reconstruction process to be evaluated and verified before proceeding to the *in-vivo* study.

2 Acquisition system and test subjects

The acquisition system comprises a Toshiba model SSA-270A/HG ultrasound scanner, a standard 2-D probe, and a position sensor. The phantom study used a 7 MHz linear array probe and the *in-vivo* study used 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 video output of the scanner are recorded by an 8 bit frame grabber 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.

Laboratory tests were first performed on a phantom comprising a latex balloon filled with a combination of water, ultrasound coupling gel and talcum powder. This type of phantom was used because B-scans of its cross-section produce images of a uniformly speckled interior, a sharp boundary, and a uniformly speckled exterior with a lower mean grey level. Grey level statistics can be easily measured for the interior and exterior regions.

Figure 9 shows a typical B-scan of the phantom. The phantom was mounted in a bath of water at an elevated temperature so that the propagation speed of ultrasound in the water approximated the speed in human tissue. An *in-vivo* examination was also performed on the gall bladder of a healthy human subject. The examination was performed by an experienced ultrasonographer in a manner fairly consistent with a normal ultrasound examination.

3 3-D reconstruction

3.1 Reconstruction without registration

Most 3-D freehand systems use similar algorithms to construct a 3-D data set from the individual B-scans. An example of a typical reconstruction algorithm can be found in Figure 3, with a detailed description following below.

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}

$${}^C\mathbf{x} = {}^C\mathbf{T}_T {}^T\mathbf{T}_R {}^R\mathbf{T}_P {}^P\mathbf{x} \longleftrightarrow {}^C\mathbf{x} = \mathbf{T} {}^P\mathbf{x}$$
 - 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 3: Reconstruction algorithm *without* registration.

Each B-scan is represented as a 2-D array \mathbf{P} of intensity values 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 4 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 ${}^P\mathbf{x}$. 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 and remains constant throughout




Figure 4: Coordinate systems used for reconstruction.

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 to \mathbf{T} .

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} which 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 \mathbf{C} is filled, it can be displayed on a computer monitor by several different methods, including volume rendering, surface rendering and any-plane slicing. Examples of any-plane slicing can be found in Figures 11, 12, 13, 16, and 17, while surface rendering is used in Figure 14. Hereafter, the term **slice** is used to indicate an image produced by any-plane slicing.

3.2 Reconstruction with registration

3.2.1 Overview

While the standard reconstruction algorithm is adequate for many tasks, it is *not* ideal when acquiring heavily compounded data sets, where voxels are intersected many times by B-scans acquired from a variety of look directions in the course of a relatively lengthy examination. Such reconstructions tend to be plagued by registration errors, caused by inaccurate B-scan position measurement, motion of the target and other imaging artifacts: see, for example, Figure 12. To acquire high quality, spatially compounded data sets we have to correct these registration errors.

A variety of image registration techniques are available [2]. Correlation-based techniques are the simplest but do not work well with 3-D ultrasound, since they are sensitive to the average grey level which tends to vary through an ultrasound scan. Correlation techniques are also very inefficient when applied to 3-D (as opposed to 2-D) registration problems. For these reasons we prefer to attempt registration via *landmarks*.

Landmarks are anatomical features which are prominent in the B-scan images. When a newly acquired B-scan is compounded into a filled (or partially filled) voxel array, any landmarks in the B-scan should align with existing landmarks in the voxel array. Registration errors can be corrected by searching for corresponding landmarks and re-positioning the new B-scan plane so that the landmarks are brought closer together.

Registration has to be performed with respect to a reliable baseline. If *all* the B-scans are re-positioned, it is possible to construct a voxel array where all landmarks are in perfect alignment, but the reconstruction bears little resemblance to the underlying anatomy (for instance, the voxel array could contain a sheared image of the anatomy). For this reason, the ultrasound examination must commence with a quick pass over the region of interest, so that most if not all of the voxels are filled. No attempt is made to register these initial B-scans: they act as the baseline. At this stage the 3-D data set will be noisy (due to speckle) but relatively free from registration errors. Subsequent passes over the region of interest, from different look directions, are compounded into the voxel array to reduce the noise. The inevitable registration errors are automatically corrected by landmark alignment as each new B-scan is acquired.

Figure 5 describes a reconstruction algorithm with landmark-based registration. The main difference between this algorithm and the algorithm without registration (Figure 3) is that the image transformation \mathbf{T} is replaced with \mathbf{T}^* , the optimal transformation that registers the landmarks in the image \mathbf{P} with those already present in the voxel array \mathbf{C} . In this study, \mathbf{T}^* is constrained to a rigid 6 DOF transformation, consistent with the expected sources of registration error (motion of the target and inaccurate B-scan position measurements).

3.2.2 Detection of landmarks

The landmarks used in this study are edge elements (edgels) automatically extracted by the Canny edge detection algorithm [4] operating on a thresholded image. The edgels are produced at the resolution corresponding to the voxel size. 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. Figure 6 illustrates the landmark detection procedure.

1. acquire 2-D image \mathbf{P} and associated position data ${}^T\mathbf{T}_R$
2. find landmarks l_{mn} in image - see Figure 6
3. if the scan plane *does not intersect* existing data (previous B-scans) 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}


$${}^C\mathbf{x} = {}^C\mathbf{T}_T {}^T\mathbf{T}_R {}^R\mathbf{T}_P {}^P\mathbf{x} \longleftrightarrow {}^C\mathbf{x} = \mathbf{T} {}^P\mathbf{x}$$
 - 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 ${}^P\mathbf{x}$) with respect to \mathbf{C}

$${}^C\mathbf{x} = \mathbf{T} {}^P\mathbf{x}$$
 - 3.2.2 assign nearest vector in \mathbf{L} to l_{mn}
 - 3.3 repeat from 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}


$${}^C\mathbf{x} = \mathbf{T}^* {}^P\mathbf{x}$$
 - 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.

Figure 5: Reconstruction algorithm *with* registration.




(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 phantom.

(f) Chains of edgels of phantom.

Figure 6: **Landmark detection.** Image (a) is the original high resolution (490×380) image of a human gall bladder. Notice the speckling throughout the image, shadow-like artifacts and non-uniform intensity of the organ boundary. Image (b) is image (a) at the reduced resolution (99×77) required for the reconstruction volume. Image (c) depicts the 1409 edgels extracted from image (b). Image (d) depicts the chains that are formed from the edgels in (c). Notice the reduced number of edgels (854) in image (d) — the edgels forming small lines and circles are eliminated. As well as the organ boundary, the detector also finds edges which do not correspond to real physical structures: these must be tolerated by the registration algorithm. Images (e) and (f) show the same detection procedure applied to the phantom.

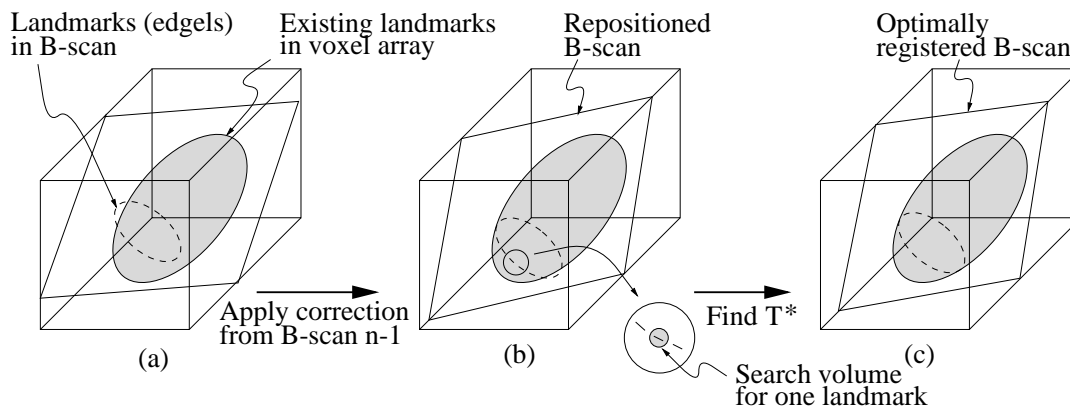


Figure 7: **Landmark-based registration.** In (a), B-scan n is inserted into the voxel array at the location indicated by the position sensor. Significant registration errors are evident (exaggerated here for clarity). After applying the rigid body correction found from registering B-scan $n - 1$, only small residual errors remain (b). These are corrected by landmark-based registration. A small ellipsoidal search volume is defined around each re-positioned landmark (edgel) in \mathbf{P} : any landmarks in \mathbf{L} found within this volume are marked as candidate correspondences for the landmark in \mathbf{P} . RANSAC regression is used to find the optimal rigid body transformation \mathbf{T}^* which aligns as many of the corresponding landmarks as possible — see Figure 8. Finally, the B-scan is inserted into its optimal position in the voxel array using \mathbf{T}^* (c).

3.2.3 Selection of correspondences

As the B-scans are acquired, landmarks are stored in a 3-D vector array \mathbf{L} which is aligned with \mathbf{C} . When a B-scan intersects non-empty voxels in \mathbf{C} , \mathbf{T}^* is determined by finding correspondences between landmarks in the new B-scan and landmarks in \mathbf{L} .

A minimum number of intersections between pixels in the scan plane and non-empty voxels in \mathbf{C} is needed for accurate registration. In the phantom study, landmark registration was attempted only when more than 25% of the pixels in the B-scan intersected filled voxels in \mathbf{C} . In practise, after the first sweep in the ultrasound examination, all subsequent B-scans have almost 100% levels of intersections.

Potential correspondences between landmarks in the B-scan and landmarks in \mathbf{L} are found by searching an ellipsoidal volume in \mathbf{L} for each landmark in the B-scan. Since the B-scans are acquired in rapid succession, any registration error will vary slowly from one B-scan to the next: we can use this observation to place the ellipsoidal search volume at an appropriate location in the voxel array, and also limit its size — see Figure 7. We are effectively *tracking* the registration error, which is far more efficient than performing an unconstrained search for \mathbf{T}^* for each B-scan.

3.2.4 Determination of transformation matrix and compounding

The set of correspondences produces more constraints than are required to determine \mathbf{T}^* . For example, each B-scan of the phantom generated approximately 1200 candidate correspondences, but only three are needed to determine \mathbf{T}^* . A least squares estimation of \mathbf{T}^* is inappropriate, since many of the candidate correspondences are outliers. Instead,

the RANSAC regression technique [9] is used to determine \mathbf{T}^* and reject the erroneous correspondences. Details of RANSAC regression can be found in Figure 8. RANSAC is sufficiently robust to tolerate a significant proportion of erroneous correspondences.

1. Randomly pick three 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}^* .

Figure 8: RANSAC regression for determining \mathbf{T}^* .

\mathbf{T}^* is used to transform the new B-scan 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 step 2.3 in Figure 3). Figure 13 shows slices through a volume reconstructed in this manner.

4 Results

4.1 Phantom study

4.1.1 3-D spatial compounding and its effect on the SNR

The main purpose of the phantom study is to investigate the effect of compounding on the SNR of the 3-D reconstructions. The B-scans of the phantom cross-sections contain two regions with almost homogeneous statistical features: the inside of the latex balloon and the outside. Figure 9 shows the two regions on a typical B-scan of the phantom.

The phantom was scanned in a continuous series of sweeps from one end of the balloon to the other, producing a large number of overlapping B-scans. Each sweep was carried out with the probe at a look direction slightly displaced from the previous sweep. About 4-5 sweeps were performed, giving 400 B-scans in total. The first 100 B-scans correspond to the first complete sweep, so they do not overlap each other.

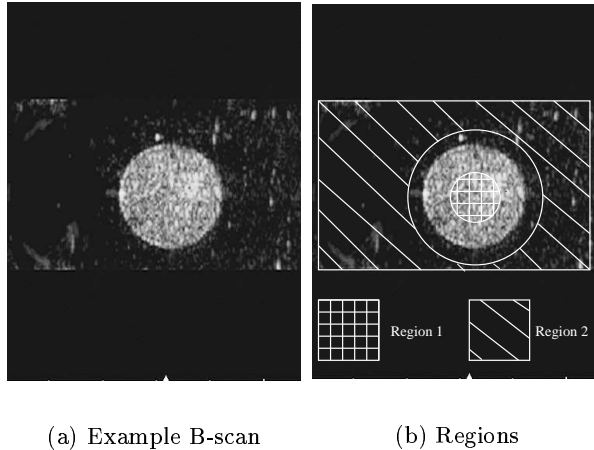


Figure 9: **Regions of the phantom B-scans.** To investigate the effect of compounding on the image statistics, two regions are defined: region 1 and region 2. Region 1 lies completely inside the latex balloon and region 2 on the outside. Similar volumetric regions are defined for the 3-D reconstruction volume.

Volumes were reconstructed from the 400 B-scans at both high and low resolution. The high resolution reconstructions featured cubic voxels with the same edge dimension as one pixel in the original B-scan (0.14mm). The low resolution reconstructions contained voxels with an edge dimension 7 times the size of a pixel (1mm). While the high resolution reconstructions preserve the full resolution of the B-scans, they require a large amount of memory (181 MBytes) and contain many unfilled voxels in regions between B-scans. Low resolution reconstructions (0.5 MBytes with 1 mm voxels) can be generated and visualised in significantly less time. For each resolution, two reconstruction volumes were created, one with registration (using the novel reconstruction algorithm in Figure 5), the other without (using the standard reconstruction algorithm in Figure 3). Four volumes were therefore reconstructed in total².

To investigate the improvement in SNR with spatial compounding, we analysed grey level statistics in each of the four reconstruction volumes at various stages of reconstruction: after 100 B-scans (initial sweep, no compounding), then 150, 200, 250, 300, 350 and finally 400 B-scans (heavy compounding). For each case we segmented volumetric regions 1 and 2 by hand, and then calculated the mean, standard deviation, and SNR for filled voxels in each region: these are tabulated in Tables 1, 2, 3, and 4.

It is immediately apparent that the SNR increases with the amount of compounding. Furthermore, the SNR improves almost identically for both the registered and unregistered cases. This is because regions 1 and 2 do not include the area where the phantom boundary is blurred by the registration errors.

The improvement in SNR can be predicted by statistical theory. Previous papers [3, 12, 15] have demonstrated a \sqrt{n} improvement in SNR for 2-D compounding of n uncorrelated B-scans. This theory is not directly applicable to the 3-D case, since the voxels are not

²In fact, registration was only performed at low resolution, with the results applied to the high resolution reconstructions. Registration at high resolution is infeasible because of the large number of unfilled voxels and the considerable memory requirements of the reconstruction algorithm.

Number of B-scans	Region 1			Region 2		
	mean	std. dev.	SNR	mean	std. dev.	SNR
100	103.76	18.37	5.65	33.68	5.76	5.85
150	102.45	18.02	5.68	33.60	5.49	6.12
200	103.09	17.65	5.84	33.39	5.14	6.50
250	103.33	17.03	6.07	32.23	4.77	6.97
300	102.80	16.55	6.21	33.17	4.66	7.12
350	102.30	16.00	6.39	33.07	4.45	7.43
400	101.41	15.74	6.44	33.01	4.34	7.61

Table 1: **High resolution reconstruction *without* registration.**

Number of B-scans	Region 1			Region 2		
	mean	std. dev.	SNR	mean	std. dev.	SNR
100	103.76	18.37	5.65	33.68	5.76	5.85
150	102.45	18.02	5.68	33.60	5.49	6.12
200	103.09	17.65	5.84	33.39	5.14	6.50
250	102.99	16.77	6.14	33.24	4.77	6.97
300	103.28	16.19	6.38	33.16	4.62	7.18
350	103.37	15.73	6.57	33.06	4.39	7.53
400	103.24	15.35	6.72	32.99	4.26	7.74

Table 2: **High resolution reconstruction *with* registration.**

Number of B-scans	Region 1			Region 2			Volume (ml)
	mean	std. dev.	SNR	mean	std. dev.	SNR	
100	100.39	11.07	9.07	33.22	3.78	8.79	7.20
150	100.09	10.63	9.42	33.15	3.41	9.72	7.17
200	99.94	9.77	10.23	33.07	3.07	10.77	7.19
250	100.05	9.27	10.79	33.00	2.88	11.46	7.27
300	99.93	8.67	11.53	32.97	2.86	11.53	7.31
350	99.71	8.16	12.22	32.93	2.82	11.68	7.43
400	99.44	7.81	12.73	32.90	2.80	11.75	7.56

Table 3: **Low resolution reconstruction *without* registration.** The volume of the phantom, as estimated by a semi-automatic segmentation technique, increases with the number of B-scans. This is because the registration errors blur the boundary of the phantom, affecting the semi-automatic segmentation.

Number of B-scans	Region 1			Region 2			Volume (ml)
	mean	std. dev.	SNR	mean	std. dev.	SNR	
100	100.39	11.07	9.07	33.22	3.78	8.79	7.20
150	100.09	10.63	9.42	33.15	3.41	9.72	7.16
200	99.94	9.77	10.23	33.07	3.07	10.77	7.19
250	99.91	9.09	10.99	33.00	2.84	11.62	7.20
300	100.02	8.13	12.30	32.96	2.78	11.86	7.14
350	100.06	7.76	12.89	32.91	2.68	12.28	7.20
400	100.17	7.27	13.78	32.86	2.57	12.79	7.18

Table 4: **Low resolution reconstruction *with* registration.** The volume of the phantom, as estimated by a semi-automatic segmentation technique, remains almost constant with higher levels of compounding.

all compounded the same number of times. The arbitrary positions and orientations of the B-scans result in some voxels being intersected more than others. Furthermore, the statistical theory for low resolution compounding must account for both the compounding due to the intersections of multiple B-scans as well as the reduction in resolution. For these reasons we have developed a full theory of 3-D spatial compounding, which can be found in the appendix.

The measured SNR of both the low and high resolution registered reconstructions is plotted against the theoretical increase in SNR in Figure 10. As the theory predicts, the SNR increases with an increasing level of compounding. Yet there are a number of differences between the idealised theory and the actual empirical study. The first is that no attempt was made to obtain completely uncorrelated B-scans for different sweeps. It has been shown that the look directions for different sweeps must be greater than 0.4 transducer widths apart to obtain completely uncorrelated speckle patterns [24]. To verify this, we deliberately performed the first few sweeps from similar look directions, saving large variations in the look direction for the last few sweeps. The results in Figure 10 confirm the theory: the SNR improves most dramatically for the last few sweeps, when the speckle patterns are less correlated. For this reason, we have highlighted the last four data points in Figure 10 to compare with the predicted values, which were derived assuming completely uncorrelated speckle (see the appendix).

The grey-level variations in the B-scans are also not due entirely to speckle. The echo amplitudes reflected by a homogeneous medium have an expected Raleigh distribution when there is a large number of scatterers per resolution cell. The SNR of a Raleigh distribution is constant and equal to 1.91 [3]. The logarithmic compression of echo amplitude in B-scan formation changes the Raleigh distribution to a Gaussian one, with an SNR of approximately 7.64 [23]. The measured SNR of a typical B-scan in this study is 5.6 for region 1 and 5.4 for region 2. The lower SNR means that the observed grey-level variations are greater than the variations due only to speckle. If some of the variation is due to real physical structure (such as the fine structure of the talcum powder suspension), then it will be correlated across different sweeps and cannot be reduced by compounding. Despite these differences from the assumptions underpinning the statistical theory, the slopes of the curves in Figure 10 are in remarkably close agreement with the theoretical predictions.

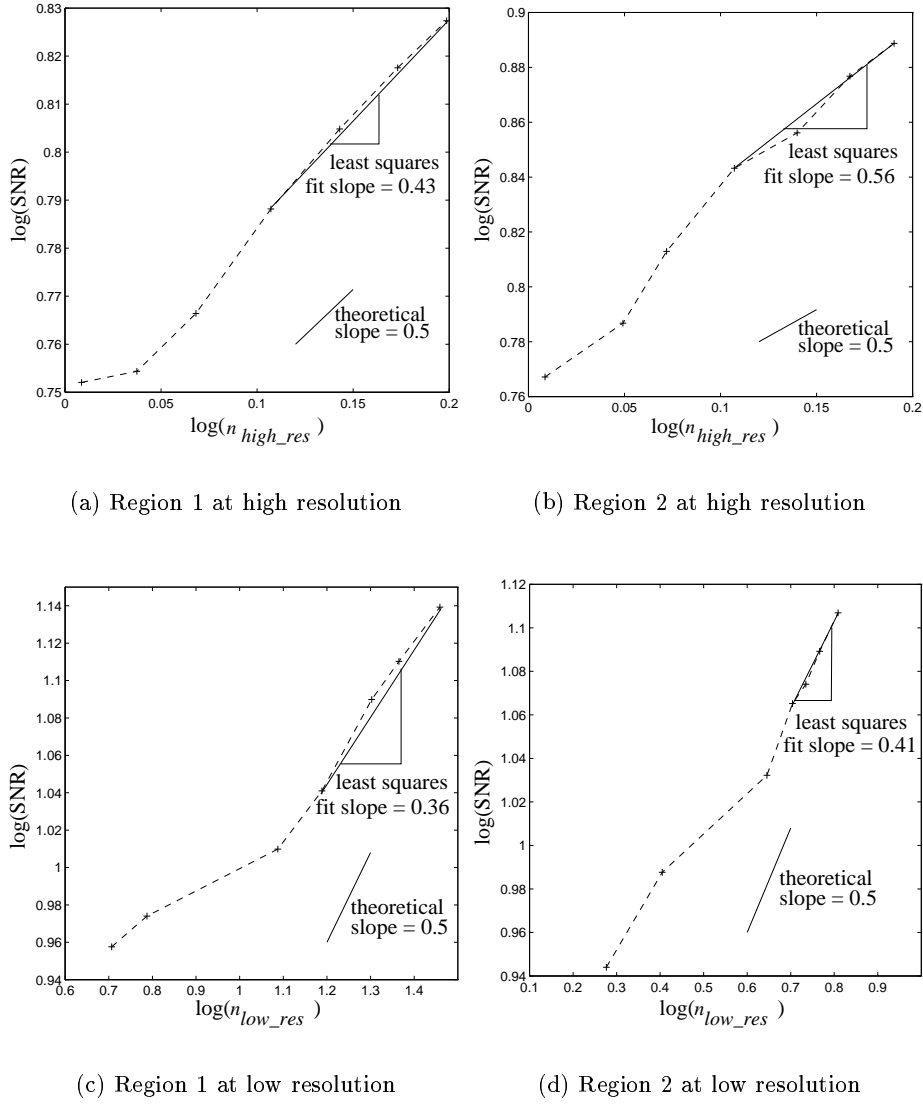


Figure 10: **Effect of compounding on SNR for registered reconstructions.** Both high resolution (a,b) and low resolution (c,d) show an increase in the SNR as the level of compounding increases. Each ‘+’ data point represents a SNR calculated at 100, 150, 200, 250, 300, 350, and 400 B-scans. The x-axis indices, n_{high_res} and n_{low_res} , are defined in the appendix. A log-log plot is used because a slope of 0.5 indicates agreement with the theory. In all cases, the final four data points are in remarkably good agreement with the theory.




Figure 11: **High resolution reconstruction.** A slice of the reconstruction volume is shown for both the unregistered and registered cases. All 400 B-scans are used in both reconstructions, but gaps still remain where voxels are not intersected by any of the B-scans. In (a), the registration errors substantially distort the circular cross-section of the phantom. The circular shape is restored in (b) by landmark registration.

4.1.2 Registration errors and their effect on volume estimation

Cross-sectional slices of the high resolution reconstructions are shown in Figure 11, both with and without registration. It is evident that the registration errors are large enough to significantly distort the reconstruction, but are dramatically reduced by landmark-based registration. The effect of compounding is not pronounced at high resolution, because each voxel is intersected only a small number of times.

Figures 12 and 13 show the effects of registration and compounding at low resolution. In both figures the speckle (and other artifacts) are greatly reduced by compounding. However, the shape of the phantom in the unregistered case departs significantly from the original shape. The registered reconstruction maintains the original shape. Figure 14 shows how compounding improves a surface rendering of the reconstruction volume.

A measure is required to quantify the effect of registration errors on the reconstruction. We chose to focus on the *volume* of the phantom, as estimated by “live-wire” segmentation [1] of slices through the reconstructed volume³. This is also a measure that is often sought after by physicians when scanning internal organs. Furthermore, it can be compared to the real volume of the phantom, which was measured with a graduated cylinder at $7.0 \text{ ml} \pm 0.2 \text{ ml}$.

³Live-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) that 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. For these tests, four sections of wire were sufficient to enclose the phantom cross-section accurately.






Figure 12: **Low resolution reconstruction *without* registration.** Figures (a) through (d) are of slices taken at the same location in the reconstruction volume at increasing levels of compounding. Speckle is reduced both inside and outside the object, but registration errors result in substantial blurring of the object boundary.



(a) 100 B-scans used

(b) 200 B-scans used



(c) 300 B-scans used

(d) 400 B-scans used

Figure 13: **Low resolution reconstruction *with* registration.** Figures (a) through (d) are of slices taken at the same location in the reconstruction volume at increasing levels of compounding. Speckle is reduced both inside and outside the object and blurring of the object boundary is minimal.




Figure 14: **Low resolution reconstruction *with* registration.** The 3-D surface renderings show how the speckle outside the object is reduced with increasing levels of compounding. (a) shows the reconstruction after a single sweep. The object is less obscured by speckle when the reconstruction volume is heavily compounded in (d).

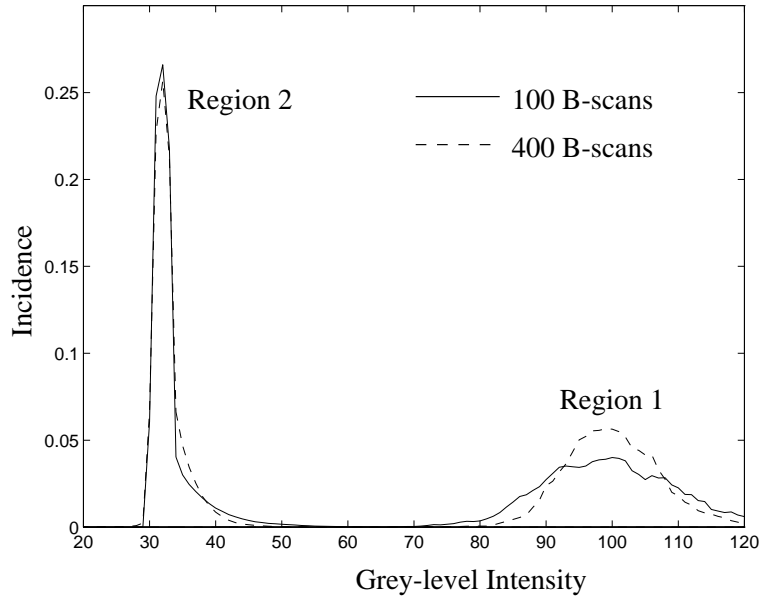



Figure 15: **Histogram of grey-levels in the reconstruction volume *with registration*.** The two regions are described by individual distributions. As compounding increases, the mean grey level remains constant but the standard deviation decreases. This effect is particularly visible for the distribution of region 1.


The volume of the phantom was calculated by live-wire segmentation at 7 levels of compounding: after 100, 150, 200, 250, 300, 350, and 400 B-scans. The resulting volume estimates are given in Tables 3 and 4 for the low resolution reconstructions. Live-wire segmentation is not feasible at high resolution, where there are too many gaps around the boundary.

Without registration, the blurring of the phantom boundary results in an increase in the segmented volume (range = [7.17–7.56]) that does not reflect the true volume. Conversely, the volume calculated from the registered reconstructions is stable (range = [7.14–7.20]). The increase in volume by blurring is less dramatic than would be expected from looking at slices through the reconstruction (Figure 12). This is because the live wire is attracted to the location with the highest intensity gradient, which is sometimes near the true boundary and not the edge of the blurred region. It is fair to say that considerably worse volume estimates can be expected under less favourable circumstances.


The grey level histograms of regions 1 and 2, shown in Figure 15, change as the level of compounding increases. As expected, the histograms become narrower with higher levels of compounding. If the histograms of the different regions are non-overlapping, then fully automatic segmentation can be achieved by thresholding the reconstruction at the appropriate grey level. Figure 16 shows several examples of fully automatic segmentation by grey level thresholding. One of the motivations for improving SNR is to allow more accurate automatic segmentation. It is evident that compounding with registration improves the accuracy of automatic grey level segmentation.



(a) Reconstruction using 100 B-scans. The area of the thresholded image on the right is 243.9 mm^2 . Live wire semi-automatic segmentation of the image produces an area estimate of 256.5 mm^2 . The difference arises mainly from speckle in the interior that falls outside the threshold range.



(b) Unregistered reconstruction using 400 B-scans. The area of the thresholded image on the right is 217.7 mm^2 . The area is lower than in (a) because the blurred regions near the phantom boundary fall outside the threshold range.



(c) Registered reconstruction using 400 B-scans. The area of the thresholded image on the right is 250.9 mm^2 . This is close to the area calculated by live-wire segmentation in (a).

Figure 16: **Low resolution reconstruction.** Corresponding slices through the reconstruction volumes are shown on the left of each figure. For each slice, the histogram of the grey levels is shown in the middle. The images on the right show the regions segmented by thresholding the grey levels falling into the range [78–122]. The compounding in (c) ensures that the entire object falls within the threshold range and the registration retains the original circular shape of the cross-section.

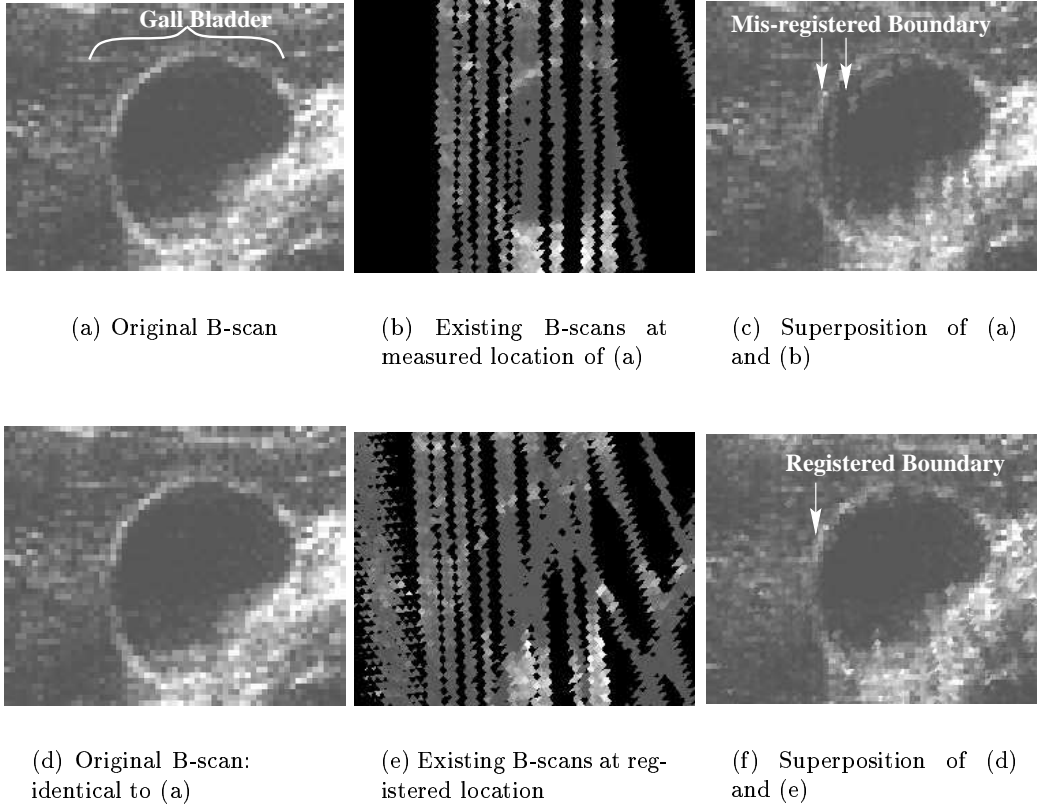


Figure 17: **Registration of an *in-vivo* transverse and longitudinal gall bladder examination.** The bottom row shows the elimination of a double boundary, evident in the top row, that is produced by mis-registered B-scans. This figure is produced with 1 mm^3 voxels to enhance clarity, but registration was performed with $2 \times 2 \times 2 \text{ mm}$ voxels. The level of compounding is not high: 35% of filled $2 \times 2 \times 2$ voxels are intersected by 1 B-scan, 33% are intersected by 2, 16% by 3, and 16% by 4 or more. The high intensity cloud-like artifact to the right and below the gall bladder is slightly suppressed by compounding. The effect is small because only a few B-scans are compounded in that region.

4.2 *In-vivo* study

To demonstrate that the registered reconstruction algorithm (Figure 5) can be applied to *in-vivo* images, we performed an ultrasound examination of the gall bladder of a healthy human subject. Reconstructions were performed at low resolution. Figure 17 shows that *in-vivo* registration errors are significant but can be minimised by landmark-based registration. A more subtle effect is that the high intensity cloud-like artifact is slightly reduced by compounding. Hardware limitations restricted this study to only 60 B-scans. With higher numbers of compounded B-scans, the artifacts should be further reduced, as in the phantom study.

5 Conclusions

We have shown how spatial compounding can improve the SNR of 3-D ultrasound images in agreement with theoretical predictions. The inevitable registration errors which come with higher levels of compounding can be corrected using an automatic, incremental landmark-based registration algorithm. The resulting high quality 3-D reconstructions are particularly well suited to automatic segmentation for visualisation and volume measurement.

Future work will investigate higher levels of compounding in *in-vivo* scans. We have recently installed a new data acquisition system which can acquire more than 1000 B-scans at 25 frames per second in a single examination. It is likely that the speed and robustness of the registration algorithm will have to be improved to handle the higher levels of compounding. We anticipate that the key to efficient registration will lie with more reliable landmark detection at larger scales, allowing registration of *contours* in the B-scans onto *surfaces* in the voxel array.

Acknowledgements

Ultrasound scanning was done with the help of Laurence Berman at the University of Cambridge Department of Radiology. The 3-D renderings were produced using 3DViewnix. Robert Rohling is supported by Churchill College and an ORS award.

A Statistical theory of 3-D spatial compounding

Simple statistical theory can be used to predict the increase in SNR with greater levels of 3-D spatial compounding. In this appendix we consider three distinct cases: 2-D compounding, 3-D high resolution compounding and 3-D low resolution compounding. First, however, we review some simple results from statistical theory.

Linear functions of random variables [13]: Consider the random variable Y as a linear function of *independent* random variables X_i :

$$Y = \sum_{i=1}^n a_i X_i$$

where $\{a_1 \dots a_n\}$ are constant coefficients. Then the expectation (or mean) μ_Y of Y , expressed in terms of the expectation of X_i (μ_{X_i}), is

$$\mu_Y = \sum_{i=1}^n a_i \mu_{X_i} \tag{1}$$

and the variance of Y (σ_Y^2), in terms of the variance of X_i ($\sigma_{X_i}^2$), is

$$\sigma_Y^2 = \sum_{i=1}^n a_i^2 \sigma_{X_i}^2 \tag{2}$$

Sets of samples of random variables: Consider Z as a set of samples Y_j drawn randomly from a number of independent distributions:

$$Z = \{Y_1, \dots, Y_j, \dots, Y_m\}$$

The expected mean of the samples in Z , expressed in terms of the expectation (μ_{Y_j}) of the distribution from which each Y_j is drawn, is

$$\mu_Z = \sum_{j=1}^m \frac{1}{m} \mu_{Y_j} \quad (3)$$

Provided all μ_{Y_j} are equal, the expected variance of the samples in Z , expressed in terms of the variances ($\sigma_{Y_j}^2$) of the distributions from which each Y_j is drawn, is

$$\sigma_Z^2 = \sum_{j=1}^m \frac{1}{m} \sigma_{Y_j}^2 \quad (4)$$

Equations (3) and (4) are readily derived from the definitions of expectancy and variance.

A.1 Case 1: 2-D compounding

The simplest form of compounding creates a 2-D image of m pixels by averaging n coplanar B-scans. To predict the improvement in SNR, we need to make two key assumptions:

Assumption 1: The subject of the ultrasound examination exhibits no resolvable structure, so any grey level variations in the B-scans are due entirely to speckle.

Assumption 2: Separate B-scans are taken from look directions spaced sufficiently far apart, so that the speckle across B-scans is uncorrelated.

Now consider an image Z , which is a set of pixels Y_j , each compounded from pixels X_i of the individual B-scans. Each X_i comes from a distribution with mean μ_0 and variance σ_0^2 . The X_i averaged for a particular Y_j are independent, because the speckle is uncorrelated across B-scans.

$$Z = \{Y_1, \dots, Y_j, \dots, Y_m\} \quad \text{where} \quad Y_j = \sum_{i=1}^n \frac{1}{n} X_i$$

From (1),

$$\mu_{Y_j} = \sum_{i=1}^n \frac{1}{n} \mu_0 = \mu_0$$

From(2),

$$\sigma_{Y_j}^2 = \sum_{i=1}^n \left(\frac{1}{n}\right)^2 \sigma_0^2 = \frac{\sigma_0^2}{n}$$

From (3),

$$\mu_Z = \sum_{j=1}^m \frac{1}{m} \mu_0 = \mu_0$$

From (4),

$$\sigma_Z^2 = \sum_{j=1}^m \frac{1}{m} \left(\frac{\sigma_0^2}{n}\right) = \frac{\sigma_0^2}{n}$$

Comparison of the SNR of Z to the SNR of an original B-scan Z_0 , with $\mu_{Z_0} = \mu_0$ and $\sigma_{Z_0}^2 = \sigma_0^2$, gives the familiar result quoted in several papers [3, 12, 15]:

$$\frac{\text{SNR}(Z)}{\text{SNR}(Z_0)} = \frac{\mu_Z / \sigma_Z}{\mu_{Z_0} / \sigma_{Z_0}} = \sqrt{n}$$

A.2 Case 2: 3-D high resolution compounding

The simplest 3-D compounding case occurs when the voxel size is the same as the pixel size. This means that each voxel is intersected no more than once per B-scan (discounting the voxels that can occasionally contain two pixels if intersected obliquely). Now consider Z to be the set of voxels in a 3-D voxel array. There are m voxels in the set Z , and voxel j is intersected n_j times. The maximum possible value of n_j is the total number of B-scans used in the reconstruction.

$$Z = \{Y_1, \dots, Y_j, \dots, Y_m\} \quad \text{where} \quad Y_j = \sum_{i=1}^{n_j} \frac{1}{n_j} X_i$$

From (1) and (2),

$$\mu_{Y_j} = \sum_{i=1}^{n_j} \frac{1}{n_j} \mu_0 = \mu_0 \quad \text{and} \quad \sigma_{Y_j}^2 = \sum_{i=1}^{n_j} \frac{\sigma_0^2}{n_j^2}$$

From (3) and (4),

$$\mu_Z = \sum_{j=1}^m \frac{1}{m} \mu_0 = \mu_0 \quad \text{and} \quad \sigma_Z^2 = \sum_{j=1}^m \frac{1}{m} \sum_{i=1}^{n_j} \frac{\sigma_0^2}{n_j^2}$$

Combining these results gives

$$\frac{\text{SNR}(Z)}{\text{SNR}(Z_0)} = \frac{\mu_Z / \sigma_Z}{\mu_{Z_0} / \sigma_{Z_0}} = \sqrt{\frac{1}{\sum_{j=1}^m \frac{1}{m} \sum_{i=1}^{n_j} \frac{1}{n_j^2}}}$$

To abbreviate references to this result in the main text, let us define the quantity n_{high_res} as follows:

$$n_{high_res} = \frac{1}{\sum_{j=1}^m \frac{1}{m} \sum_{i=1}^{n_j} \frac{1}{n_j^2}} \quad \Rightarrow \quad \frac{\text{SNR}(Z)}{\text{SNR}(Z_0)} = \sqrt{n_{high_res}}$$

Note that the result of case 1 can be derived by setting $n_j = n$ for all j .

A.3 Case 3: 3-D low resolution compounding

Most 3-D reconstructions involve cuberilles with voxels considerably larger than the pixels in the B-scans. Again, the set of voxels in the reconstruction can be written as

$$Z = \{Y_1, \dots, Y_j, \dots, Y_m\} \quad \text{where} \quad Y_j = \sum_{i=1}^{n_j} \frac{1}{n_j} X_i$$

but in this case the X_i are *not* all independent. This becomes clear if we consider each voxel Y_j as follows:

$$Y_j = \sum_{k=1}^{k_j} \sum_{p=1}^{p_{j,k}} \frac{1}{n_j} X_{p,k}$$

where k_j is the number of B-scans that intersect voxel j , and $p_{j,k}$ is the number of pixels that intersect voxel j for B-scan k : i.e. $\sum_{k=1}^{k_j} p_{j,k} = n_j$. We have labelled the individual pixels $X_{p,k}$ according to B-scan k and pixel p . Rearranging gives

$$Y_j = \sum_{k=1}^{k_j} \frac{p_{j,k}}{n_j} \Theta_{k,j} \quad \text{where} \quad \Theta_{k,j} = \sum_{p=1}^{p_{j,k}} \frac{1}{p_{j,k}} X_{p,k}$$

$\Theta_{k,j}$ is the mean of the pixels that intersect voxel j for B-scan k . The expected mean (μ_Θ) of $\Theta_{k,j}$ can be determined using lemma (1), which is still valid for non-independent $X_{p,k}$ [13]:

$$\mu_\Theta = \sum_{p=1}^{p_{j,k}} \frac{1}{p_{j,k}} \mu_0 = \mu_0$$

Since the size of the speckle is larger than one pixel, the grey level value of a pixel is not independent of its close neighbours. Thus the expected variance (σ_Θ^2) of $\Theta_{k,j}$ cannot be determined using (2), which is not valid if the $X_{p,k}$ are not independent. The variance of the sum of several neighbouring pixels *does* decrease as the number of summed pixels increases, but $\sigma_\Theta^2 \gg \sigma_0^2/p_{j,k}$. σ_Θ^2 depends on the spatial structure of the speckle and can be expressed as follows:

$$\sigma_\Theta^2 = \sigma_0^2 b^2(p_{j,k})$$

where $b(p_{j,k})$ is estimated empirically from the B-scans themselves. Returning now to the expression for Y_j , the summation over k *can* be considered as a sum of independent variables, since each element in the sum comes from a different B-scan. Hence

$$\mu_{Y_j} = \sum_{k=1}^{k_j} \frac{p_{j,k}}{n_j} \mu_0 = \mu_0 \quad \text{and} \quad \sigma_{Y_j}^2 = \sum_{k=1}^{k_j} \left(\frac{p_{j,k}}{n_j} \right)^2 \sigma_0^2 b^2(p_{j,k})$$

From (3) and (4),

$$\mu_Z = \sum_{j=1}^m \frac{1}{m} \mu_0 = \mu_0 \quad \text{and} \quad \sigma_Z^2 = \sum_{j=1}^m \frac{1}{m} \sum_{k=1}^{k_j} \left(\frac{p_{j,k}}{n_j} \right)^2 \sigma_0^2 b^2(p_{j,k})$$

Combining these results gives

$$\frac{\text{SNR}(Z)}{\text{SNR}(Z_0)} = \frac{\mu_Z/\sigma_Z}{\mu_{Z_0}/\sigma_{Z_0}} = \sqrt{\frac{1}{\sum_{j=1}^m \frac{1}{m} \sum_{k=1}^{k_j} \left(\frac{p_{j,k}}{n_j} \right)^2 b^2(p_{j,k})}}$$

Again, to abbreviate references to this result in the main text, let us define the quantity n_{low_res} as follows:

$$n_{low_res} = \frac{1}{\sum_{j=1}^m \frac{1}{m} \sum_{k=1}^{k_j} \left(\frac{p_{j,k}}{n_j} \right)^2 b^2(p_{j,k})} \quad \Rightarrow \quad \frac{\text{SNR}(Z)}{\text{SNR}(Z_0)} = \sqrt{n_{low_res}}$$

Note that the result of case 2 can be derived by setting $p_{j,k} = 1$, $b^2(p_{j,k}) = 1$ (because $\sigma_\Theta^2 = \sigma_0^2$), and $k_j = n_j$.

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] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, December 1992.
- [3] C.B. Burckhardt. Speckle in ultrasound B-mode scans. *IEEE Trans. Sonics and Ultrasonics*, 25(1):1–6, 1978.
- [4] J. Canny. A computational approach to edge detection. *IEEE Trans. PAMI*, 8(6):679–698, 1986.
- [5] F.A. Chervenak, G.C. Isaacson, and S. Campbell. *Ultrasound in Obstetrics and Gynecology*. Little, Brown and Company, Boston, MA, USA, 1993.
- [6] D. C. Crawford, D. S. Bell, and J. C. Bamber. Compensation for the signal processing characteristics of ultrasound B-mode scanners in adaptive speckle reduction. *Ultrasound Med. Biol.*, 19(6):469–485, 1993.
- [7] A. N. Evans and M. S. Nixon. Mode filtering to reduce ultrasound speckle for feature extraction. *IEE Proceedings: Vision, Image and Signal Processing*, 142(2):87–94, April 1995.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] A. Hernandez, O. Basset, P. Chirossel, and G. Gimenez. Spatial compounding in ultrasonic imaging using an articulated scan arm. *Ultrasound Med. Biol.*, 22(2):229–238, 1996.
- [13] R.V. Hogg and J. Ledolter. *Engineering Statistics*. Macmillan Publishing Company, New York, New York, USA, 1987.
- [14] 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.

- [15] A.T. Kerr, M.S. Patterson, F.S. Foster, and J.W. Hunt. Speckle reduction in pulse echo imaging using phase insensitive and phase sensitive signal processing techniques. *Ultrason. Imaging*, 8:11–28, 1986.
- [16] S. Leeman and D.A. Seggie. Speckle reduction via phase. In L.A. Ferrari, editor, *Intl. Symp. on Pattern Recognition and Acoustical Imaging*, pages 173–177. SPIE: 768, 1987.
- [17] 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.
- [18] T.R. Nelson and D.H. Pretorius. 3D ultrasound image quality improvement using spatial compounding and 3D filtering. *Med. Phys.*, 21(6):998, 1994.
- [19] 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.
- [20] 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.
- [21] A. Salustri and J. R. T. C. Roelandt. Ultrasonic three-dimensional reconstruction of the heart. *Ultrasound Med. Biol.*, 21(3):281–293, 1995.
- [22] P. M. Shankar. Speckle reduction in ultrasound B-scans using weighted averaging in spatial compounding. *IEEE Trans. Ultrasonics, Ferroelectrics, and Frequency Control*, 33(6):754–758, November 1986.
- [23] J.M. Thijssen, B.J. Oosterveld, and R.F. Wagner. Gray level transforms and lesion detectability in echographic images. *Ultrason. Imaging*, 10:171–195, 1988.
- [24] 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 Trans. Ultrasonics, Ferroelectrics, and Frequency Control*, 33(3):257–264, 1986.
- [25] G.E. Trahey, J.W. Allison, S.W. Smith, and O.T. von Ramm. A quantitative approach to speckle reduction via frequency compounding. *Ultrason. Imaging*, 8:151–164, 1986.
- [26] G.E. Trahey, J.W. Allison, S.W. Smith, and O.T. von Ramm. Speckle reduction achievable by spatial compounding and frequency compounding: experimental results and implications for target detectability. In L.A. Ferrari, editor, *Intl. Symp. on Pattern Recognition and Acoustical Imaging*, pages 185–192. SPIE: 768, 1987.