SENSORLESS FREEHAND 3D ULTRASOUND IN REAL TISSUE: SPECKLE DECORRELATION WITHOUT FULLY DEVELOPED SPECKLE

A. H. Gee, R. J. Housden, P. Hassenpflug,
G. M. Treece and R. W. Prager

CUED/F-INFENG/TR 510

January 2005

University of Cambridge
Department of Engineering
Trumpington Street
Cambridge CB2 1PZ
United Kingdom

Email: ahg/rjh80/ph305/gmt11/rwp @eng.cam.ac.uk

Sensorless freehand 3D ultrasound in real tissue: speckle decorrelation without fully developed speckle

Andrew H. Gee, R. James Housden, Peter Hassenpflug, Graham M. Treece
and Richard W. Prager

University of Cambridge
Department of Engineering
Trumpington Street
Cambridge CB2 1PZ

Abstract

It has previously been demonstrated that freehand 3D ultrasound can be acquired without a position sensor by measuring the elevational speckle decorrelation from frame to frame. However, this requires that the B-scans contain significant amounts of fully developed speckle. In this paper, we show that this condition is rarely satisfied in scans of real tissue, which instead exhibit fairly ubiquitous coherent scattering. By examining the axial and lateral correlation functions, we propose an heuristic technique to quantify the amount of coherency at each point in the B-scans. This leads to an adapted elevational decorrelation scheme which allows for the coherent scattering. Using the adapted scheme, we demonstrate markedly improved reconstructions of animal tissue in vitro.

1 Introduction

Freehand 3D ultrasound [4, 10] is an emerging medical imaging modality with many potential applications [5]. The data is normally acquired by attaching a position sensor to the ultrasound probe and recording a series of B-scans together with their positions and orientations. The B-scans thus form a 3D data set which can be visualised and processed in a number of ways to extract clinically useful information.

Perhaps the greatest barrier to widespread clinical uptake is the need for the add-on position sensor. This requires careful end-user calibration [9] and imposes cumbersome constraints on the scanning protocol. Typically, the operator must maintain a clear line of sight between the probe and the position sensor's fixed base station, and must be careful not to stray outside the sensor's operating region. It is therefore not surprising that researchers have been investigating alternative freehand acquisition techniques that dispense with the position sensor.

Hossack et al. [6] propose a modified probe with auxiliary arrays mounted at right angles to the main array. The probe thus records additional B-scans perpendicular to the main B-scan. As the probe is moved across the subject's skin, its motion can be deduced by tracking features in the main and perpendicular B-scans. Unfortunately, the auxiliary arrays give the probe a fairly large footprint, and it is difficult to keep the scanning window in full contact with the skin when going around corners. This causes drop-outs in the auxiliary B-scans and hence tracking failures. The extra probe complexity also adds considerable expense.

An alternative approach uses an unmodified probe and a technique called *speckle decorrelation* [3, 8, 15] to estimate elevational[1] motion. Consider two neighbouring B-scans A and B in a freehand sequence. Any in-plane motion between A and B (translation in the axial and lateral directions, roll around the elevational axis) is readily determined using standard 2D image registration techniques [11, 14]. That just leaves the three out-of-plane degrees of freedom: translation in the elevational direction, tilt around the lateral axis and yaw around the axial axis.

---

[1] The following standard nomenclature is used to refer to the principal directions in ultrasound imaging. The *axial* direction is in the direction of wave propagation, from the transducer into the skin. The *lateral* direction is the other principal direction in the plane of the B-scan. The *elevational* direction is perpendicular to the plane of the B-scan.

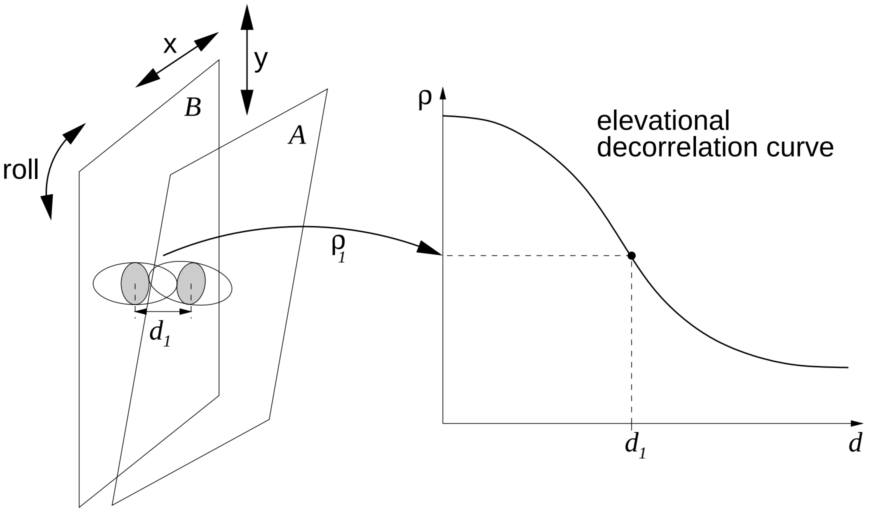

Figure 1: **Principle of elevational speckle decorrelation**. The in-plane motion between scans A and B (translation in the x and y directions, roll around the plane normal) is readily determined using conventional 2D image registration techniques. This leaves three degrees of freedom: translation in the elevational direction, tilt and yaw. Consider corresponding patches in scans A and B (the shaded ellipses). Because of the imperfect elevational focusing, the contents of the patches depend on scatterers within overlapping resolution cells (the hollow ellipsoids) and are therefore correlated. The correlation coefficient depends on the degree of overlap and hence the elevational separation. It follows that, given a suitable decorrelation curve, a measured correlation ρ_1 can be used to look up the corresponding separation d_1. Repeating this process for three (or more) non-collinear patches determines the elevational separation, tilt and yaw of A relative to B.

Perhaps surprisingly, these out-of-plane components can be estimated by analysing the images themselves. This is because the focusing of the ultrasound beam is far from perfect. Consequently, the backscattered signal at any point in a B-scan is a function of the scatterers in a certain *resolution cell* around that point. The resolution cells are particularly elongated in the elevational direction and there is considerable overlap between cells on neighbouring B-scans — see Figure 1. It follows that if we examine corresponding patches on A and B, we will observe a correlation between the two sets of backscattered signals. Moreover, the correlation will depend on the degree of overlap of the resolution cells, and hence the separation of the two patches. We can therefore measure the inter-patch correlation and look up the corresponding separation on a *decorrelation curve*, as illustrated in Figure 1. A minimum of three such lookups, for three (non-collinear) pairs of patches at different locations on the B-scans, allows us to estimate the elevational separation, tilt and yaw of A relative to B.

The precise form of the decorrelation curve can be derived using a simplified model of ultrasound imaging. For fully developed speckle, also known as Rayleigh scattering [16], it can be shown that the Pearson correlation coefficient between two patches of backscattered intensity I_1 and I_2 is given by

$$\rho(I_1, I_2) = \lambda^2 \tag{1}$$

and the corresponding formula for the backscattered amplitude is

$$\rho(A_1, A_2) = \frac{4E(\lambda^2) - 2\left(1 - \lambda^2\right) K(\lambda^2) - \pi}{4 - \pi} \tag{2}$$

where K and E are complete elliptic integrals of the first and second kinds respectively [1, 17, 18]. Assuming a Gaussian shaped resolution cell, the parameter λ depends on the standard deviation of the elevational cell width as follows [11]

$$\lambda^2 = \exp\left(\frac{-\delta^2}{2\sigma^2}\right)$$

where δ is the elevational separation of the patches. There is therefore a deterministic relationship between correlation ρ and elevational separation δ. The resolution cell width σ varies across and down the B-scan, especially with depth, and we must therefore use different decorrelation curves at different points in the scan plane. In theory, σ can be calculated using acoustic physics and a detailed specification of the ultrasound scanner and probe. A more practical way to calibrate the decorrelation curves is to scan a speckle phantom and measure ρ directly for different elevational offsets δ: this is the approach taken in this paper.

A serious difficulty with speckle decorrelation is that the theory only holds for fully developed speckle. Real tissue contains regions of coherent scattering, which decorrelate at a slower rate than regions of fully developed speckle. One way around this problem is to automatically detect those regions which contain fully developed speckle, and obtain elevational distance estimates in these regions only [15]: we shall return to this proposition later. Alternatively, we can accept that coherent scattering results in distortion of the 3D reconstruction and attempt to correct this in a post-processing stage. For example, we could correct any gross errors in length, tilt and yaw by registering the data to an additional B-scan acquired at right angles to the originals, as suggested in [2]. However, this complicates the scanning protocol and requires that the data contain sufficient distinguishable features for the registration to succeed. A fair assessment of the state-of-the-art is that speckle decorrelation is a good way to acquire *qualitative* 3D ultrasound data, but for accurate, *quantitative* reconstructions one has to fall back on a position sensor.

In this paper, we show how geometrically accurate reconstructions can be obtained using an adapted speckle decorrelation scheme. First, in Section 2, we describe the design and calibration of a conventional speckle decorrelation system. Through experiments conducted with this system, we establish the limitations of speckle detection as a means of overcoming the coherent scattering problem. In Section 3, we propose a model for adapting the calibration curves to account for local coherent scattering in the tissue being scanned. In a series of experiments, we demonstrate markedly improved reconstructions of animal tissue in vitro. Finally, in Section 4, we draw some conclusions and suggest some avenues for further work.

2 Conventional speckle decorrelation

The experimental framework

Experiments were conducted with a 5–10 MHz linear array probe connected to a Dynamic Imaging[2] Diasus ultrasound machine. The depth setting was 4 cm with a single focus at 2 cm. Analogue radio frequency (RF) ultrasound signals were digitised after receive focusing and time-gain compensation, but before log-compression and envelope detection, using a Gage Compuscope CS14100[3] 14-bit digitiser. Whole frames, each comprising 127 RF vectors by 3818 samples, were stored in on-board Gage memory before transferring to PC memory at 75 MB/s. The system operates in real time, with acquisition rates exceeding 30 frames per second. Sampling was at 66.67 MHz, synchronous with the ultrasound machine's internal clock: this synchronization minimises phase jitter between vectors. The acquired vectors were filtered with a 3–30 MHz broadband filter, then envelope-detected using the Hilbert transform. The resulting 127 × 3818 frames of backscatter amplitude data formed the basis of all further computation. Their resolution is approximately 0.01 mm per sample in the axial direction and 0.27 mm per vector in the lateral direction.

The real time RF acquisition platform constitutes a unique and powerful research facility which we are happy to make available to other users[4]. For this study, it offers two important advantages

---

[2] `http://www.dynamicimaging.co.uk`

[3] `http://www.gage-applied.com`

[4] `http://mi.eng.cam.ac.uk/~rwp/stradx`

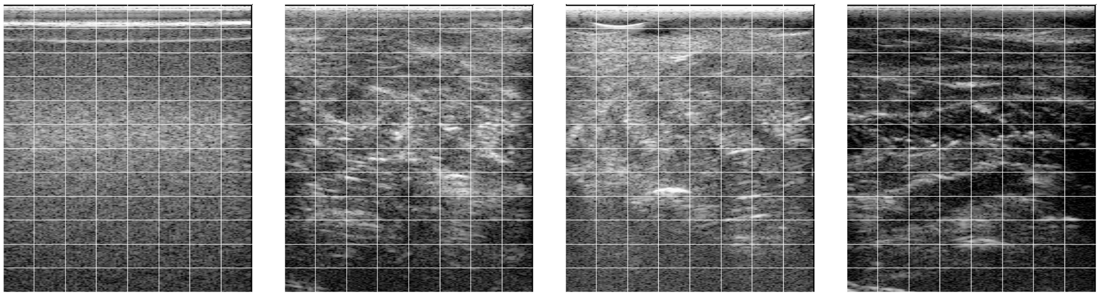

(a) phantom (b) beef1 (c) beef2 (d) turkey

Figure 2: **Calibration and test data**. The figure shows typical B-scans from the calibration (a) and test (b)–(d) data sets, with the 8×12 patch grid superimposed on each. Note that all processing was performed on the 127×3818 RF amplitude data. These 382×439 B-scans are geometrically correct (equal scales of 0.09 mm/pixel in the axial and lateral directions) and are derived from the RF data by log-compression, subsampling axially and interpolating laterally.

over conventional, frame-grabbed B-scans. Firstly, the data is acquired before log-compression inside the ultrasound machine: equations (1) and (2) are therefore directly applicable, without having to first undo the log-compression process [12]. Secondly, the RF data's axial resolution is approximately ten times better than that of a B-scan (around 4000 samples compared with 400 pixels). We will soon see the value of this when we come to compute axial decorrelation curves. Further details of the RF acquisition system can be found in [13].

The probe was mounted on a linear screw thread allowing accurate translation in the elevational direction. Its displacement was measured using a dial gauge. Four 3D data sets were acquired using this framework, each comprising 100 parallel B-scans at an elevational spacing of 0.04 mm or 0.02 mm. The first was of a speckle phantom with Rayleigh backscatter and uniform attenuation of 0.4 dB/cm/MHz [7]. This was used to calibrate the elevational decorrelation curves at different points in the scan plane. The elevational scan spacing was 0.04 mm. The next two were of a beef joint: one was acquired parallel to the muscle fibres, the other perpendicular to them. Both had an elevational scan spacing of 0.04 mm. The final data set was of a stack of turkey fillets lying in a water bath, with an elevational scan spacing of 0.02 mm. The three in vitro test sets allow us to carefully evaluate reconstruction accuracy without having to worry about the effects of physiological motion.

Each frame was divided into an 8 column × 12 row grid of patches. This gives more than 5000 samples per patch, sufficient to calculate meaningful second order statistics. Using the speckle phantom, separate calibration curves were obtained for each patch, thereby capturing any variation in the resolution cell across and down the B-scan. Subsequently, separate distance estimates were obtained for each patch in the three test sets. Typical B-scans from each data set can be found in Figure 2.

The calibration scan

For each patch in the phantom data set, we recorded not only elevational, but also axial and lateral decorrelation curves. Preliminary experiments indicated that the data decorrelated in the elevational direction over a distance not greater than 0.8 mm, and in the axial direction over a distance not greater than 0.2 mm. Wishing to represent each decorrelation curve with around 10 samples, we therefore computed elevational correlations at offsets of 0, 0.08, 0.16 . . . 0.72 mm (i.e. skipping every other frame) and axial correlations at 0, 0.02, 0.04 . . . 0.18 mm (i.e. skipping every other sample). Lateral correlations were calculated at the finest possible resolution of 0.27 mm (unit vector steps). Typical decorrelation curves can be found in Figure 3. They are for four patches from the same column near the centre of the B-scan, but at different depths.

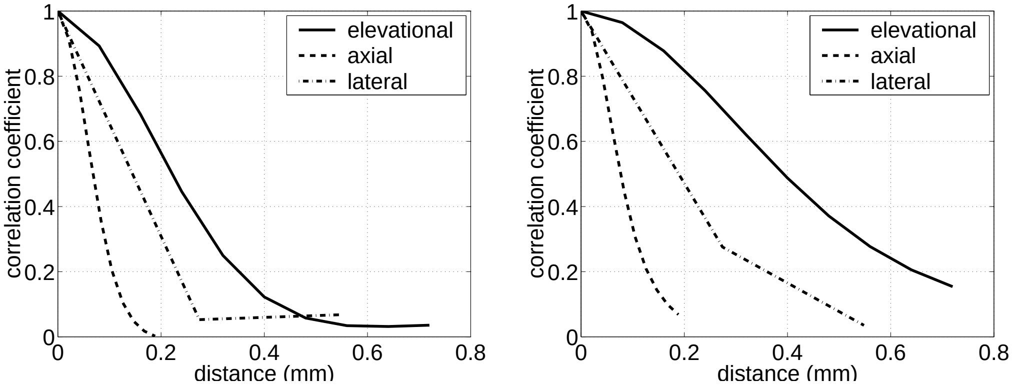

(a) row 3 column 5 (b) row 6 column 5

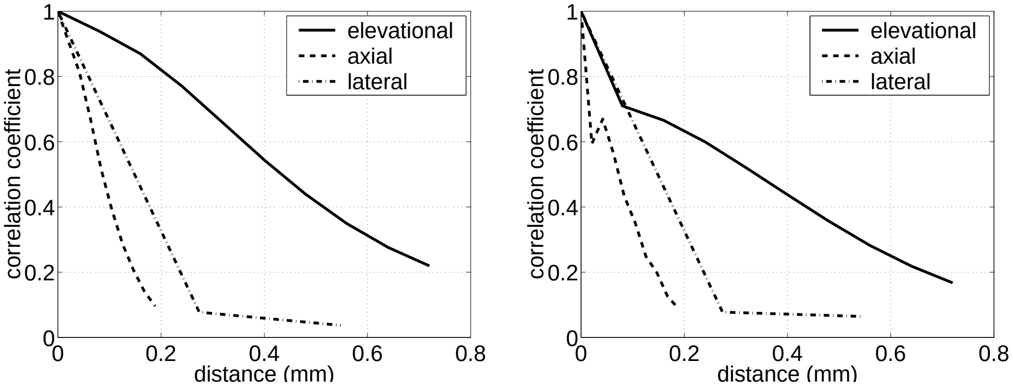

(c) row 9 column 5 (d) row 12 column 5

Figure 3: **Decorrelation curves for the speckle phantom**. The figure shows average curves in the axial, lateral and elevational directions for four patches from the same, central column of the grid. Each curve is represented by ten samples, though only the first three are shown for the lateral curves. Each data point on each curve is the average of a large number of observations. For example, the 0.08 mm sample of the elevational curve was based on the correlation of all pairs of frames 0.08 mm apart, i.e. the average of 98 separate correlation measurements.

	frame separation (mm)				
	0.08	**0.24**	**0.40**	**0.56**	**0.72**
phantom	0.077 ± 0.010	0.239 ± 0.011	0.404 ± 0.014	0.552 ± 0.017	0.612 ± 0.011
beef1	0.029 ± 0.005	0.127 ± 0.009	0.216 ± 0.011	—	—
beef2	0.044 ± 0.004	0.154 ± 0.009	0.246 ± 0.013	—	—
turkey	0.034 ± 0.003	0.125 ± 0.005	0.205 ± 0.006	—	—

Table 1: **Elevational distance estimation using conventional speckle decorrelation**. The table shows the mean distance estimate ± one standard deviation for all frames separated by 0.08, 0.24, 0.40, 0.56 and 0.72 mm. For separations greater than 0.40 mm, the estimates are biased for reasons explained in Figure 4. Results for the test data sets, and for all further experiments in this paper, are therefore given for separations of up to 0.40 mm only.

The approximately Gaussian shape of the curves is in agreement with theory [3]. The widths of the curves indicate significant variation in the resolution cell down the B-scan. However, there is little variation across the B-scan: different patches in the same row exhibit similar curves. Note how the elevational focusing is far worse at the bottom of the scan than at the top. Also note how the data decorrelates in the lateral direction after only one or two samples, which is to be expected given the relatively good lateral focusing but sparse lateral sampling. Since there is always complete decorrelation after two lateral shifts (0.54 mm), only the 0.27 mm sample is potentially meaningful, and we shall disregard all higher shifts in the rest of this work.

At the bottom of the B-scan, the curves lose their characteristic Gaussian shape, with rapid decorrelation in all directions for small offsets. Visual inspection of the B-scans — see Figure 2(a) — indicates severe attenuation at depth: the backscattered signal for rows 11–12 is dominated by noise which is essentially uncorrelated in all directions. This suggests a rather simple technique for automatic noise detection. If a_0, a_1 and a_2 are the first three samples of the axial decorrelation curve, then a patch is classified as containing significant noise if

$$a_1 - a_2 < c(a_0 - a_1) \tag{3}$$

The constant c was set to 0.9 for all the experiments in this study, though repeating some of them with $c = 1.0$ made very little difference to the results: this is not a critical parameter. Noisy patches are disregarded for the purpose of elevational distance estimation. This includes all patches in rows 11–12 of the B-scan, since these patches fail the noise test in the calibration scan of the speckle phantom. We also disregard all patches in rows 1–2, since we were unable to obtain valid calibration curves for these patches: the phantom's scanning window introduced strong specular features in these rows, as is evident in Figure 2(a).

The evaluation scans

The calibrated decorrelation curves were used to estimate the spacing between pairs of scans in the phantom, beef and turkey data sets. All pairs of scans i and j, separated by 0.72 mm or less, were considered. A separate distance estimate was obtained for each patch, by measuring the elevational correlation between the corresponding patches on i and j and looking up the distance on the patch's calibration curve. To facilitate the noise test (3), axial decorrelation curves were also calculated for each patch. The elevational distance lookup was deemed invalid if the patch failed the noise test, or the measured correlation was below the range recorded in the calibration curve. The results from valid lookups were averaged to obtain a single distance estimate for each pair of scans. The mean and standard deviation of these inter-scan spacings are tabulated in Table 1.

As expected, the results for the phantom are very good, since this was the very data used to derive the calibration curves. Small errors arise because the calibration curves are based on the average of many measurements, whereas individual correlations between a particular pair of

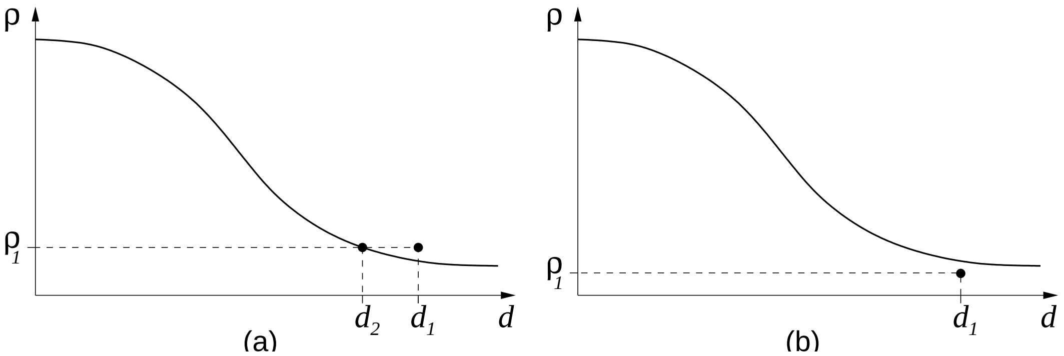

Figure 4: **Systematic errors for large inter-scan separation.** When two patches are widely separated in the elevational direction, their correlation will be relatively low. Suppose two patches, which are actually separated by a distance d_1, exhibit a slightly higher than average correlation ρ_1, as shown in (a). Lookup on the calibration curve leads to a slight underestimate of distance d_2. Now suppose ρ_1 is slightly lower than average, as in (b). In this case, the lookup fails and no distance estimate is available for this patch. The set of valid lookups is thus dominated by cases like (a), resulting in systematic underestimation of elevational separation.

	frame separation (mm)		
	0.08	**0.24**	**0.40**
beef1	0.057 ± 0.009	0.215 ± 0.020	0.358 ± 0.025
turkey	0.063 ± 0.004	0.193 ± 0.005	0.319 ± 0.005

Table 2: **Elevational distance estimation using tissue-based calibration**. The table shows the results of conventional speckle decorrelation, but this time using calibration curves derived from the beef2 data set.

patches may differ slightly from the average. The larger errors for separations above 0.40 mm are more systematic. They arise because patches exhibiting a lower than expected correlation may cause an invalid lookup. The set of valid lookups is thus dominated by patches exhibiting a higher than expected correlation, leading to systematic underestimation of elevational separation. This phenomenon is illustrated in Figure 4. All remaining results in this study are for separations not exceeding 0.40 mm, and are therefore assumed to be unaffected by this systematic error.

The results for the beef and turkey data sets are very poor, with severe underestimation of distance at all inter-scan separations. This is because these data sets contain very little in the way of fully developed speckle. They are dominated by coherent scattering, which correlates over a longer range than fully developed speckle. The resulting high correlation values are misinterpreted as close spacing. The traditional solution to this problem is to somehow detect those patches which contain fully developed speckle, and disregard all other patches when estimating elevational separation. Speckle detectors typically draw on low order statistics [12], such as the ratio of the mean backscattered signal intensity to the standard deviation [15]. Using this simple scheme, the authors of [15] found that only 5% of an in vivo breast scan qualified as fully developed speckle.

To establish to what extent the results in Table 1 could be improved by speckle detection, we implemented the “ideal” speckle detector by manually rejecting all patches except those whose actual elevational decorrelation matched the calibration curves. In the beef data, we found extremely little that even approached fully developed speckle. Figure 5 shows, for a typical frame, the only two patches which came close. Even if a practical speckle detector were able to single out these two patches, we would still have residual distance errors of around 20–30%, and, with just two offsets, we would be unable to estimate all degrees of freedom between the two scans.

As a last resort, we might try a less principled approach to improve the results in Table 1.

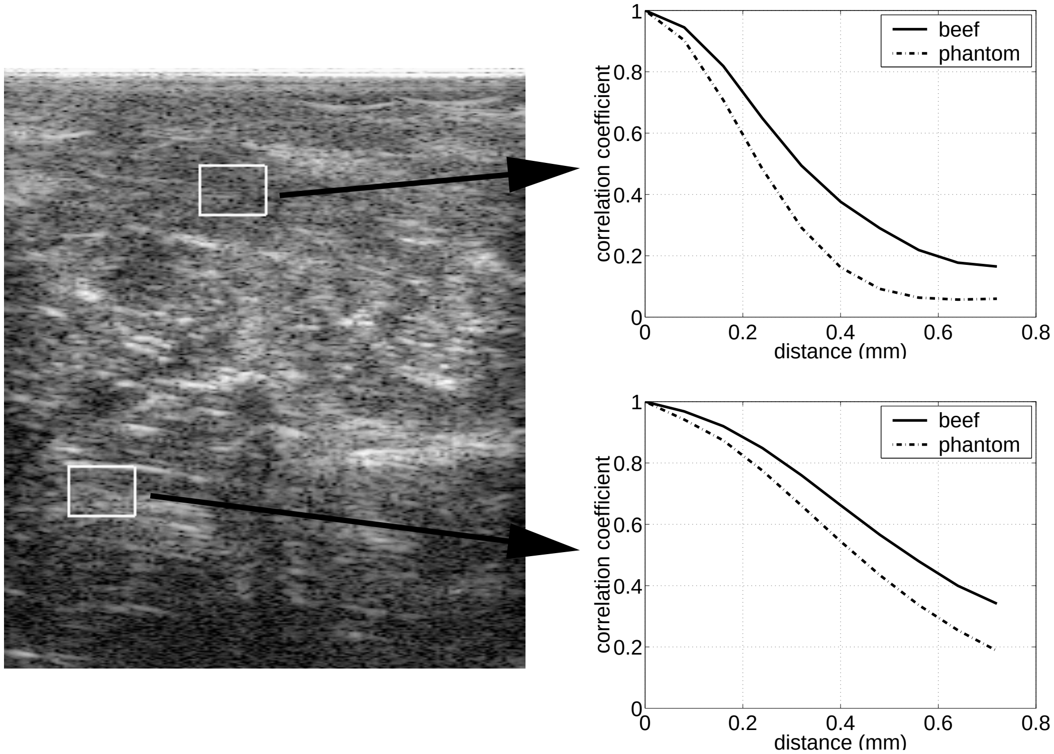

Figure 5: **Elevational decorrelation curves for beef**. These are the only two patches whose elevational decorrelation curves approach those of the speckle phantom. All other patches deviate even further from the calibration.

Even though the decorrelation theory holds only for fully developed speckle, we could consider an alternative calibration, this time based on scans of real tissue instead of the speckle phantom. Table 2 shows what happens if we repeat the beef1 and turkey experiments, but this time using decorrelation curves calibrated on the beef2 data set. Though much improved, the results still indicate errors of around 20–30%. This is because the average decorrelation rate in turkey is different to that in beef, and even the same piece of beef decorrelates at different rates when scanned in different directions. So what is required is not a refined, static calibration, but a dynamic scheme which can adapt to the local properties of the tissue being scanned.

3 Adaptive speckle decorrelation

The scatterer adaptation model

The principal reason for the poor results in Table 1 is that real data exhibits little to no Rayleigh scattering. Interestingly, this is observable in the lateral and axial decorrelation curves as well as in the elevational curves — see Figure 6(a). Compared with the speckle phantom, the coherent scattering in beef results in slower decorrelation in *all* directions. This begs the following question: can the observed discrepancies in the axial and lateral curves be used to adapt the elevational calibration curve to allow for coherent scattering? To answer this question, we first consider models for decorrelation in the presence of coherent scattering.

Wagner et al. [17] postulate a model for mixed coherent and Rayleigh scattering in which a fixed coherent component is added to the Rayleigh model in equations (1) and (2). This leads to

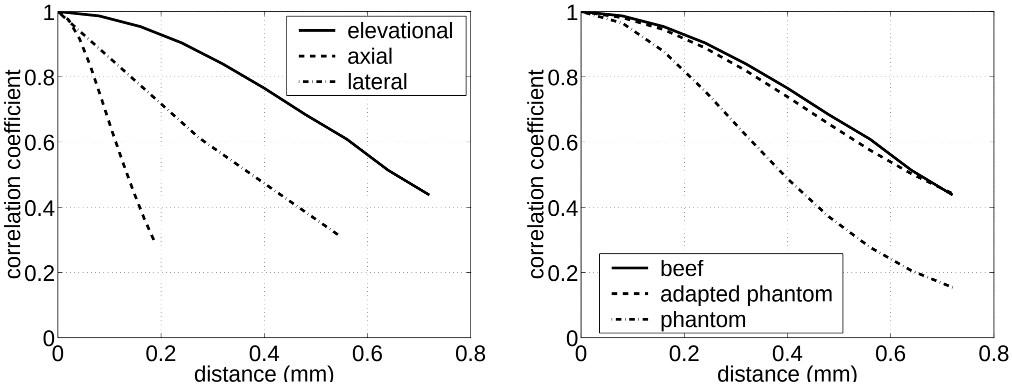

(a) row 6 column 5

(b) elevational curves

Figure 6: **Decorrelation curves for beef**. (a) shows decorrelation curves for a patch near the centre of one of the beef B-scans. Compared with the same patch in Figure 3(b), the coherent scattering leads to slower decorrelation in all directions. There is even significant correlation between neighbouring vectors laterally. (b) highlights the elevational curves, showing how, for each elevational offset, the correlation in beef is always greater than the corresponding correlation in the calibration phantom. Adapting the calibration curve to allow for the coherent scattering results in much closer agreement between the phantom and beef curves.

an expression for the intensity correlation as follows

$$\rho(I_1, I_2) = \frac{\lambda^2 I_r + 2\lambda I_c}{I_r + 2I_c} \tag{4}$$

where I_r is the intensity of Rayleigh scattering and I_c is the intensity of coherent scattering. The equivalent amplitude expression results in a graph of similar shape, but is complicated to express algebraically [17].

As Wagner et al. illustrate in Figure 8 of [18], the consequence of equation (4) is that as the scattering changes from Rayleigh to purely specular, the decorrelation lengths increase by a factor of $\sqrt{2}$. This is thus a poor model for some sorts of coherent scattering where the decorrelation lengths may be arbitrarily large, depending on the geometry of the objects being scanned. By comparing the decorrelation curves in Figures 6(a) and 3(b), it is clear that increases of more than a factor of $\sqrt{2}$ should not be ruled out. Hence, equation (4) does not provide a sufficiently flexible model for our empirical observations.

Instead, we propose a simple, heuristic model for decorrelation in the presence of coherent scattering. Consider a patch of RF amplitude data A_1 in B-scan i and a corresponding patch A_2 in B-scan j. Assuming both patches contain fully developed speckle, their correlation coefficient can be written $\rho_r = \rho(A_1, A_2)$, where the subscript r denotes Rayleigh scattering. Now add to each patch a component of coherent scattering. In our model, we assume that the patches now contain data $A_1 + kA_2$ and $A_2 + kA_1$, where $0 \leq k \leq 1$. The correlation coefficient is then $\rho = \rho(A_1 + kA_2, A_2 + kA_1)$. The parameter k accounts for coherent scattering correlating over the distance between the two patches: when $k = 0$ we have fully developed speckle and $\rho = \rho_r$, when $k = 1$ we have identical data $A_1 + A_2$ in the two patches and $\rho = 1$. Intermediate values of k give correlation coefficients between ρ_r and 1.

This model is entirely heuristic and certainly does not account for all the phenomena observed in ultrasound images. For example, when $k = 1$ the patches are perfectly correlated but the data $A_1 + A_2$ would still have the appearance of speckle: so this is not a convincing model for specular reflection. But this simple model does account for the sort of decorrelation curves observed in real tissue, with ρ varying between ρ_r and 1 depending on the amount k of coherent scattering present in both patches.

The precise relationship between ρ and ρ_r is readily derived by expanding the expression for the correlation coefficient:

$$\begin{aligned}
\rho &= \rho(A_1 + kA_2, A_2 + kA_1) \\
&= \frac{\mathrm{cov}(A_1 + kA_2, A_2 + kA_1)}{\sqrt{\mathrm{var}(A_1 + kA_2)\mathrm{var}(A_2 + kA_1)}} \\
&= \frac{\mathrm{cov}(A_1, A_2) + \mathrm{cov}(A_1, kA_1) + \mathrm{cov}(kA_2, A_2) + \mathrm{cov}(kA_2, kA_1)}{\sqrt{[\mathrm{var}(A_1) + \mathrm{var}(kA_2) + 2\mathrm{cov}(A_1, kA_2)]\,[\mathrm{var}(A_2) + \mathrm{var}(kA_1) + 2\mathrm{cov}(A_2, kA_1)]}} \\
&= \frac{\mathrm{cov}(A_1, A_2) + k\mathrm{var}(A_1) + k\mathrm{var}(A_2) + k^2\mathrm{cov}(A_2, A_1)}{\sqrt{[\mathrm{var}(A_1) + k^2\mathrm{var}(A_2) + 2k\mathrm{cov}(A_1, A_2)]\,[\mathrm{var}(A_2) + k^2\mathrm{var}(A_1) + 2k\mathrm{cov}(A_2, A_1)]}}
\end{aligned}$$

Since A_1 and A_2 both represent Rayleigh scattering, $\mathrm{var}(A_1) = \mathrm{var}(A_2)$, and this simplifies to

$$\rho = \frac{(1 + k^2)\mathrm{cov}(A_1, A_2) + 2k\mathrm{var}(A_1)}{(1 + k^2)\mathrm{var}(A_1) + 2k\mathrm{cov}(A_1, A_2)}$$

Dividing through by $\mathrm{var}(A_1)$, and noting that $\rho_r = \mathrm{cov}(A_1, A_2)/\mathrm{var}(A_1)$, we arrive at

$$\rho = \frac{(1 + k^2)\rho_r + 2k}{(1 + k^2) + 2k\rho_r} \tag{5}$$

Equation (5) allows us to predict the correlation coefficient between two patches given an amount k of coherent scattering. There is also the inverse problem to address: given observed correlation ρ, how much coherent scattering k is there? By simple rearrangement of (5), we find:

$$k^2(\rho_r - \rho) + 2k(1 - \rho\rho_r) + \rho_r - \rho = 0$$

This is a quadratic equation in k. Solving, we find

$$k = \frac{(1 - \rho\rho_r)}{(\rho - \rho_r)} \pm \frac{\sqrt{(1 - \rho^2)(1 - \rho_r^2)}}{(\rho - \rho_r)} \tag{6}$$

For physically plausible scenarios $\rho > \rho_r$, and the negative solution then gives k in the range $0 < k \leq 1$, as required.

Equations (5) and (6) suggest a possible scheme for coping with coherent scattering in elevational distance estimation. Consider a patch of data which exhibits axial correlation ρ_a at axial displacement d_a. Given a suitable calibration for this patch, we know the expected correlation ρ_r for fully developed speckle. We can then use equation (6) to calculate k_a, the amount of coherent scattering correlating over a distance d_a in the axial direction. Assuming physical isotropy of the coherent scatterers in the axial and elevational directions, we would expect to observe the same amount of coherent scattering correlating over a distance d_e in the elevational direction. The correspondence between d_a and d_e should account for the resolution of the imaging system: we assume $d_e = 4d_a$, since the elevational resolution is, on average, around four times worse than the axial resolution (see Figure 3). Hence, by calculating k_a for the full set of axial displacements $d_a \in \{0 \ldots 0.18\}$ mm, we obtain a matching set of values k_e for elevational offsets $d_e \in \{0 \ldots 0.72\}$ mm.

We can repeat the same analysis in the lateral direction, though in this case we have only one meaningful measurement of ρ_l, and hence k_l, this being for a unit vector lateral displacement of $d_l = 0.27$ mm. We assume that this value of k_l is valid for the full set of elevational displacements $d_e \in \{0 \ldots 0.72\}$ mm. There is no reason to favour either the axial or lateral k_e estimates, so we simply average the two to obtain the final set of k_e values. We use these in equation (5) to calculate an adapted elevational decorrelation curve which accounts for the apparent coherent scattering. Figure 6(b) shows one such adapted curve. The entire adaptation process is illustrated in Figure 7.

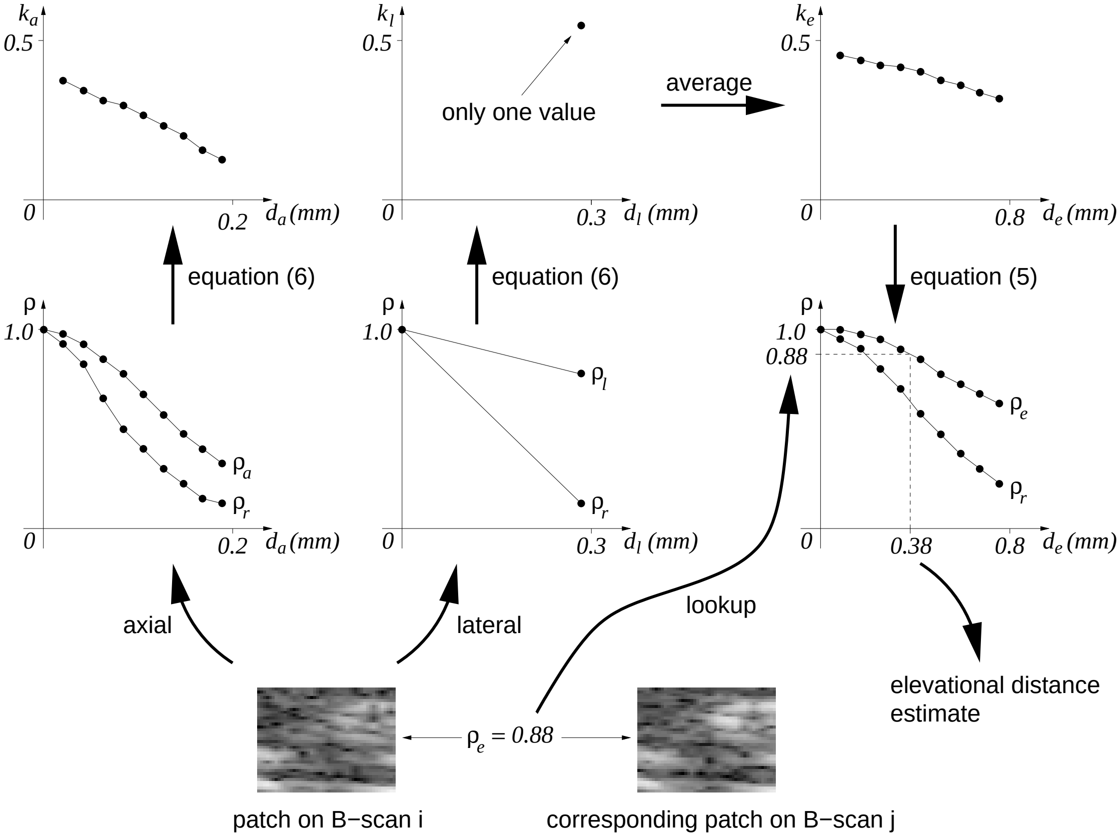

Figure 7: **The adaptation procedure**. Consider a patch of RF amplitude data on B-scan i and the corresponding patch on B-scan j. The starting point for elevational distance estimation is a set of decorrelation curves ρ_r in the axial, lateral and elevational directions, obtained by one-off scanning of a speckle phantom. We can also calculate decorrelation curves ρ_a and ρ_l in the axial and lateral directions for the patch on B-scan i (since decorrelation curves are slowly varying, it makes little difference whether we use B-scan i or B-scan j). By comparing ρ_a with ρ_r using equation (6), we can calculate $k_a(d_a)$, the apparent amount of coherent scattering correlating over a distance d_a in the axial direction. Likewise, we can calculate $k_l(d_l)$ in the lateral direction, though in this case we have only one meaningful distance d_l to work with. In this particular example, the patch exhibits pronounced streaks in the lateral direction and k_l is higher than k_a. By averaging k_a and k_l we can estimate $k_e(d_e)$, the amount of coherent scattering correlating over a distance d_e in the elevational direction. Since the resolution cell is approximately four times wider elevationally than axially, we assume $k_e(4d) = 0.5(k_l + k_a(d))$: in other words, coherent scattering correlates four times further in the elevational direction than in the axial direction. The k_e values are then used in equation (5) to adapt the elevational decorrelation curve ρ_r to obtain ρ_e. ρ_r was measured for fully developed speckle, whereas ρ_e accounts for the observed coherent scattering. The ρ_e curve allows lookup of elevational distance given elevational correlation: in this case, the patches are estimated to be 0.38 mm apart. Note how using the unadapted curve ρ_r would have resulted in a gross underestimate. This is because these patches exhibit significant coherent scattering and therefore decorrelate at a slower rate than fully developed speckle.

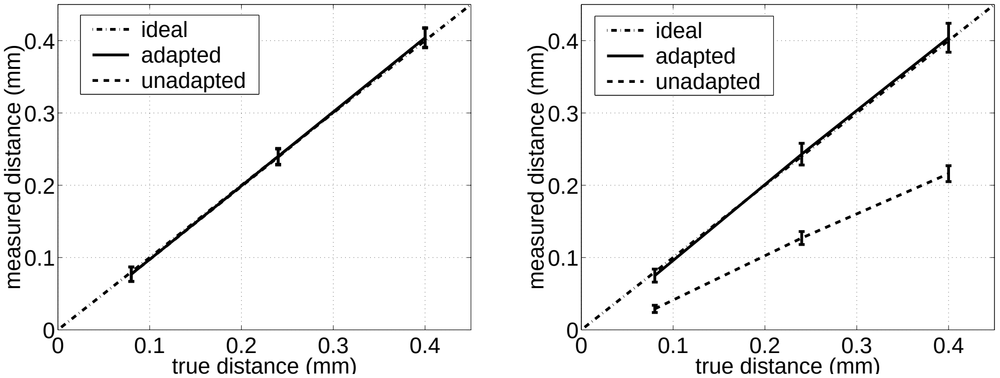

(a) phantom

(b) beef1

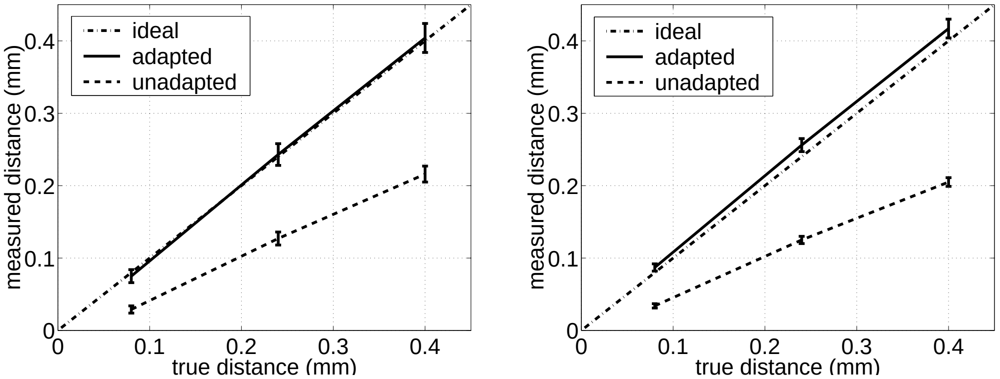

(c) beef2

(d) turkey

Figure 8: **Distance estimation using adaptive speckle decorrelation**. The distance between any pair of frames was estimated by averaging all valid patch lookups on the elevational decorrelation curves. The graphs show the mean of these inter-frame distances, with the error bars representing ± one standard deviation.

		frame separation (mm)		
		0.08	**0.24**	**0.40**
phantom	distance (mm)	0.077 ± 0.010	0.240 ± 0.011	0.404 ± 0.013
	tilt (deg)	0.005 ± 0.017	-0.008 ± 0.021	-0.025 ± 0.044
	yaw (deg)	0.000 ± 0.011	-0.000 ± 0.013	0.002 ± 0.034
beef1	distance (mm)	0.075 ± 0.009	0.243 ± 0.015	0.404 ± 0.020
	tilt (deg)	0.083 ± 0.034	0.188 ± 0.104	0.305 ± 0.169
	yaw (deg)	-0.001 ± 0.011	0.013 ± 0.026	0.033 ± 0.055
beef2	distance (mm)	0.085 ± 0.006	0.246 ± 0.012	0.403 ± 0.020
	tilt (deg)	0.152 ± 0.055	0.164 ± 0.148	0.200 ± 0.268
	yaw (deg)	-0.045 ± 0.024	-0.080 ± 0.041	-0.099 ± 0.074
turkey	distance (mm)	0.087 ± 0.005	0.256 ± 0.009	0.417 ± 0.013
	tilt (deg)	0.055 ± 0.019	0.146 ± 0.038	0.285 ± 0.066
	yaw (deg)	-0.022 ± 0.014	-0.084 ± 0.013	-0.131 ± 0.035

Table 3: **Distance and angle estimation using adaptive speckle decorrelation**. This table includes the distances plotted in Figure 8, and also inter-frame tilt and yaw estimates as described in Figure 9. For the parallel frames used in this study, both angles should be zero. Although the average distances are very good, there is a bias in the tilt. The yaw angles are generally acceptable.

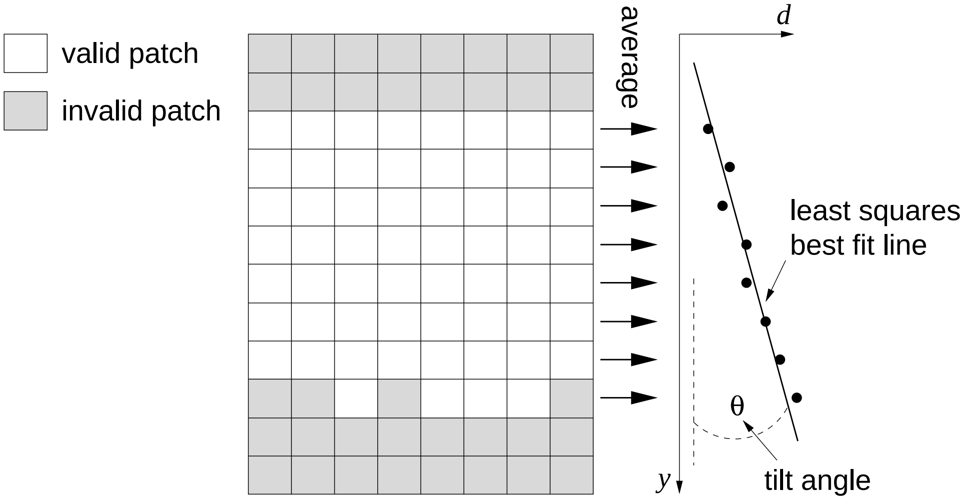

Figure 9: **Estimation of yaw and tilt**. For any pair of frames, the valid elevational distance estimates are averaged across rows to produce a single estimate per row. Invalid patches (the top and bottom two rows of the grid, and any additional patches which fail the noise test or cause an out-of-range lookup) are disregarded. A best fit line is found for the averaged distances: the orientation of this line provides an estimate of the tilt between the two frames. A similar analysis along columns produces an estimate of the yaw between the two frames.

The adaptive scheme was applied to the calibration and test data sets, yielding the results in Figure 8. At first sight, the performance appears to be excellent. The beef data is almost perfectly corrected, the turkey results are slightly biased but much better than before. The adaptation makes negligible difference to the phantom results, since $k_e \approx 0$ for all patches.

While the graphs in Figure 8 are a common way of presenting results of this kind, they can obscure significant deficiencies. It transpires that, while the average distance estimate per frame is very good, patches near the top of the B-scan tend to underestimate, while those near the bottom tend to overestimate. If these distances were used for three degree of freedom position estimation, the reconstruction would exhibit an incorrect tilt, with the frames bunched tightly at the top and more widely spaced at the bottom. To quantify this error, we calculated yaw and tilt estimates for each pair of frames, as explained in Figure 9. The complete results, including all three degrees of freedom, are presented in Table 3. While by no means perfect, the performance of the adaptive scheme is quite remarkable considering its heuristic nature and total lack of adjustable parameters.

Scatterer adaptation with axial correction

By introducing an ad hoc axial correction term, with two adjustable parameters, we can significantly improve the tilt estimates in Table 3. The correction takes the simple, linear form

$$k_e^* = k_e\,[1 - a(r - b)] \tag{7}$$

where k_e^* is the corrected k_e term, r is the row in which the patch resides (in the range 1–12), and a and b are adjustable parameters. We tuned a and b using a fifth data set, an additional scan of the beef made several days later. By this time the beef had toughened, producing scans with a different appearance to those in beef1 and beef2. We found that $a = 0.048$ and $b = 5.2$ corrected the tilt errors in this additional data set. Using the same ad hoc correction with the original four data sets produced the results in Table 4. Figure 10 shows the marginal effect of the axial correction on decorrelation curves from the top and bottom of a B-scan.

It would appear that the axial correction generalises well from the independent data set to the original four data sets. Interestingly, and perhaps significantly, little to no correction is needed

		frame separation (mm)		
		0.08	**0.24**	**0.40**
phantom	distance (mm)	0.077 ± 0.010	0.240 ± 0.011	0.404 ± 0.013
	tilt (deg)	0.005 ± 0.017	-0.008 ± 0.022	-0.026 ± 0.045
	yaw (deg)	0.000 ± 0.011	-0.000 ± 0.013	0.001 ± 0.033
beef1	distance (mm)	0.071 ± 0.009	0.235 ± 0.014	0.394 ± 0.019
	tilt (deg)	0.000 ± 0.034	0.016 ± 0.101	0.045 ± 0.188
	yaw (deg)	0.004 ± 0.010	0.022 ± 0.027	0.044 ± 0.053
beef2	distance (mm)	0.084 ± 0.006	0.244 ± 0.012	0.402 ± 0.021
	tilt (deg)	0.075 ± 0.062	0.007 ± 0.157	-0.028 ± 0.255
	yaw (deg)	-0.044 ± 0.023	-0.079 ± 0.042	-0.096 ± 0.073
turkey	distance (mm)	0.083 ± 0.005	0.248 ± 0.009	0.409 ± 0.011
	tilt (deg)	-0.042 ± 0.014	-0.056 ± 0.032	-0.044 ± 0.057
	yaw (deg)	-0.011 ± 0.012	-0.057 ± 0.026	-0.077 ± 0.059

Table 4: **Distance and angle estimation using adaptive speckle decorrelation with axial correction**. These results are for the corrected elevational coherency terms k_e^*, calculated using equation (7) with parameters a and b tuned using an additional, independent data set. Compared with Table 3, the tilt errors are much improved.

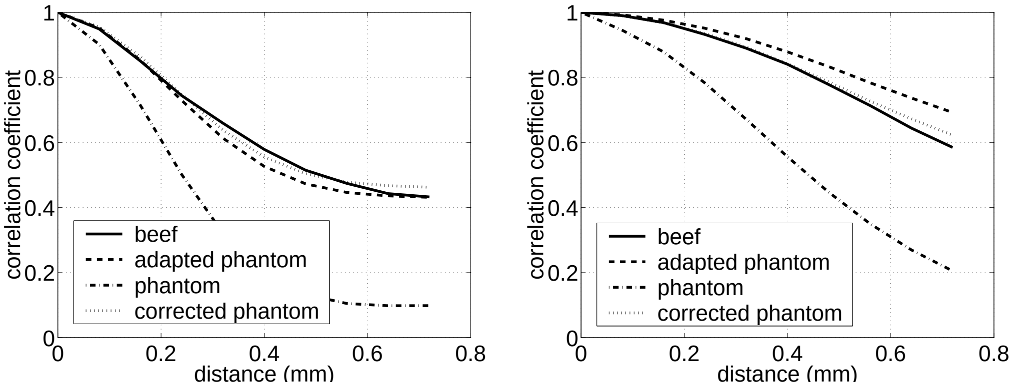

(a) top (b) bottom

Figure 10: **Scatterer adaptation with axial correction**. The axial correction makes a small but significant difference to the decorrelation curves at the top (a) and bottom (b) of the B-scan. While the adapted phantom curves are close to the true decorrelation curves for beef, the additional axial correction improves the match even further and largely removes the top-to-bottom bias responsible for erroneous tilt estimates.

	width (mm)	tilt (deg)	yaw (deg)
phantom	3.95	-0.06	0.05
beef1	3.84	1.13	0.54
beef2	3.98	0.55	-0.85
turkey	2.04	0.10	-0.07

Table 5: **Accumulated errors in full 3D reconstructions**. The table shows the width of the reconstruction, measured from the centre of B-scan 1 to the centre of B-scan 100. The correct width is 3.96 mm for the first three data sets and 1.98 mm for the turkey data set. Also tabulated is the tilt of B-scan 100 relative to B-scan 1 (positive if the tops of the B-scans are closer than the bottoms), and the yaw of B-scan 100 relative to B-scan 1 (positive if the left sides of the B-scans are closer than the right sides). Both angles should be zero.

around rows 5–6, which is approximately where the ultrasound beam is focused (2 cm focus on a 4 cm depth setting). The original, parameter-free adaptation scheme would appear to work well at the focus, but produce slightly biased results away from it. Further work is required to establish the connection between the heuristic adaptation scheme and ultrasound physics, and explain the scheme's apparent sensitivity to beam focusing.

The axial correction has improved not only the tilt, but also the distance estimates for the turkey data set. This is because this data had a fair number of invalid lookups in the top half of the B-scans, caused by the echogenic water between the fillets producing noise-like axial decorrelation. Hence, compared with the other data sets, the turkey distances were based on a greater proportion of patches from the lower half of the B-scans, leading to systematic overestimation. Axial correction improves the situation by reducing the dependence of distance on depth.

We used adaptive speckle decorrelation, with axial correction, to compute full 3D reconstructions of the four data sets. For the data sets recorded at a pitch of 0.04 mm, the reconstruction algorithm operated as follows.

1. Compare B-scans 2–11 with B-scan 1 (i.e. up to 0.4 mm separation).
2. For B-scans i and 1, $i \in \{2 \ldots 11\}$, use the corrected coherency terms k_e^* to obtain an elevational distance estimate for each patch.
3. Apply a 3×3 median filter to the grid of elevational offsets (to remove outlying estimates).
4. Find the least squares best fit plane through the filtered offsets: this locates B-scan i with respect to B-scan 1.
5. Now position B-scans 12-21 relative to B-scan 11, and so on up to B-scan 100.

For the turkey data set, recorded at a pitch of 0.02 mm, the procedure was the same except that the reference frame was advanced every 20 frames, not every 10 frames: this maintains the maximum 0.4 mm separation for decorrelation-based distance estimation.

The reconstruction results can be found in Table 5 and Figure 11. In the design of the reconstruction algorithm, we have to address an inherent trade-off regarding error accumulation. When concatenating n rigid body transformations to find the position of B-scan $i+n$ relative to B-scan i, we would expect angular and distance errors of the order ne, where e is the mean per-frame error given in Table 4. It therefore makes sense to minimise the amount of concatenation, by moving the reference B-scan as infrequently as possible within the limits of elevational distance estimation (i.e. up to around 0.4 mm). This is why, for the first three data sets, the algorithm locates B-scans 2–11 relative to B-scan 1, then B-scans 12–21 relative to B-scan 11, and so on. However, there are also potential problems with too little concatenation. The first three data sets are reconstructed in ten steps, the turkey in only five. This makes the reconstructions susceptible to random perturbations in e (the standard deviations in Table 4). We have been particularly unlucky with beef1, accumulating a tilt error of 1.13° over only ten steps. The expected tilt error

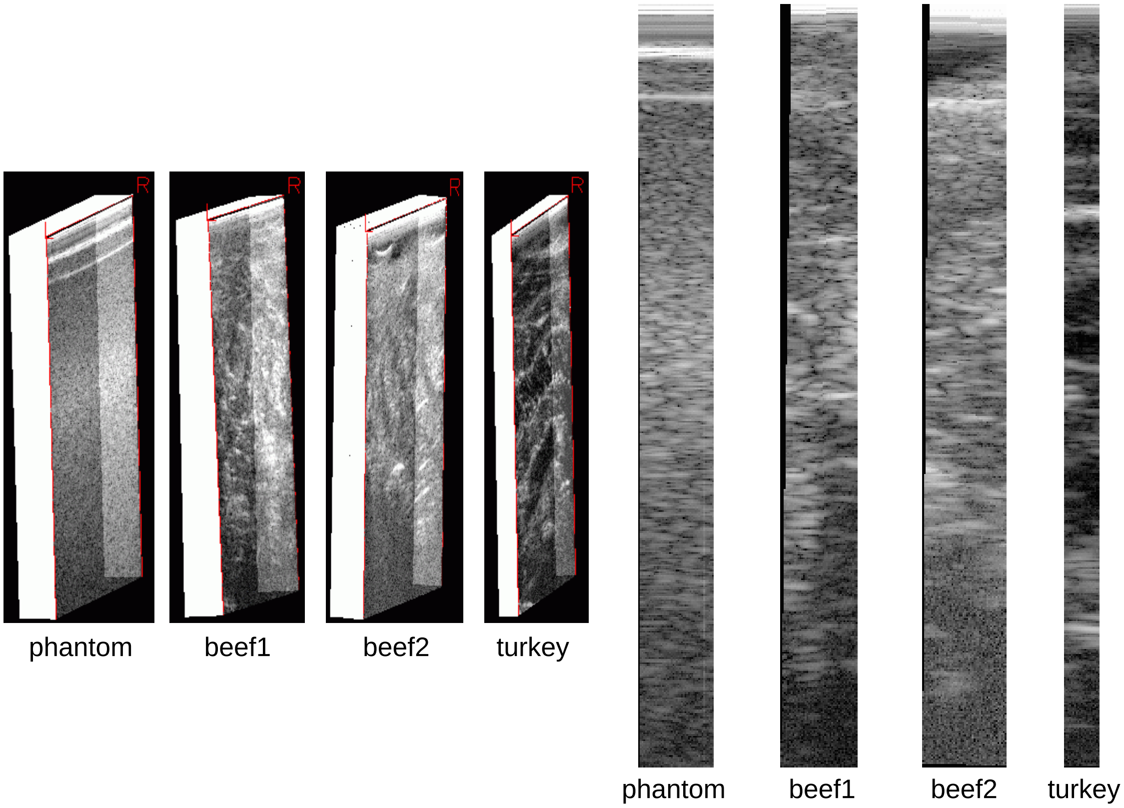

phantom beef1 beef2 turkey

phantom beef1 beef2 turkey

Figure 11: **Full 3D reconstructions**. The figure shows the outlines of the 3D reconstructions (left) and reslices down the centre of each reconstruction (right). The reslices are perpendicular to the plane of the B-scans. The turkey data set was acquired at a pitch of 0.02 mm, compared with 0.04 mm for the others, and is therefore narrower.

is around 0.045° per 0.4 mm step — see Table 4. For larger data sets, we would expect the accumulated errors to approach ne.

4 Conclusions and further work

We have demonstrated how conventional speckle decorrelation is not a feasible technique for B-scan positioning in real tissue. Much improved results can be obtained by adapting the decorrelation curves to account for apparent coherent scattering. Our heuristic adaptation model is conceptually simple and entirely free of adjustable parameters. An additional ad hoc correction term can improve estimates of tilt: this requires two adjustable parameters, though these do not appear to be overly sensitive to tissue type.

We can identify three major avenues for further work. The first is to analyse the adaptation scheme from the perspective of ultrasound physics. This may lead to a more principled way to correct the tilt bias without introducing adjustable parameters. The second is to develop a more powerful reconstruction algorithm. There is much redundancy of information to exploit: each frame-pair provides many more elevational offsets than are required to estimate three degrees of freedom, and we can correlate each frame with any number of frames within the decorrelation range, not just one. Finally, we need to apply these algorithms to true freehand scans of in vivo data. The difficulty here is that decorrelation will not depend solely on elevational motion, but also on in-plane and physiological motion. In-plane probe motion can be tracked using established image registration techniques [14]: we have already demonstrated their efficacy for sensorless freehand 3D ultrasound [11]. Physiological motion will cause localised, poor estimates of elevational displacement. However, these should not compromise a robust reconstruction algorithm which takes full advantage of the aforementioned redundancy.

Acknowledgments

This work was carried out under EPSRC grant number GR/S34366/01. Dynamic Imaging provided a modified Diasus ultrasound machine with direct access to the analogue RF signals.

References

[1] M. Abramowitz and I. A. Stegun. *Handbook of mathematical functions.* Dover, 1970.

[2] R-F. Chang, W-J. Wu, D-R. Chen, W-M. Chen, W. Shu, J-H. Lee, and L-B. Jeng. 3-D US frame positioning using speckle decorrelation and image registration. *Ultrasound in Medicine and Biology*, 29(6):801–812, 2003.

[3] J-F. Chen, J. B. Fowlkes, P. L. Carson, and J. M. Rubin. Determination of scan-plane motion using speckle decorrelation: Theoretical considerations and initial test. *International Journal of Imaging Systems Technology*, 8:38–44, 1997.

[4] A. Fenster, D. B. Downey, and H. N. Cardinal. Three-dimensional ultrasound imaging. *Physics in Medicine and Biology*, 46:R67–R99, 2001.

[5] A. Gee, R. Prager, G. Treece, and L. Berman. Engineering a freehand 3D ultrasound system. *Pattern Recognition Letters*, 24(4–5):757–777, February 2003.

[6] J. A. Hossack, T. S. Sumanaweera, S. Napel, and J. S Ha. Quantitative 3-D diagnostic ultrasound imaging using a modified transducer array and an automated tracking technique. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 49(8):1029–1038, August 2002.

[7] J. M. Kofler and E. L. Madsen. Improved method for determining resolution zones in ultrasound phantoms with spherical simulated lesions. *Ultrasound in Medicine and Biology*, 27(12):1667–1676, 2001.

[8] M. Li. System and method for 3D medical imaging using 2D scan data. United States patent 5,582,173, application number 529778, September 1995.

[9] F. Lindseth, G. A. Tangen, T. Langø, and J. Bang. Probe calibration for freehand 3-D ultrasound. *Ultrasound in Medicine and Biology*, 29(11):1607–1623, November 2003.

[10] T. R. Nelson and D. H. Pretorius. Three-dimensional ultrasound imaging. *Ultrasound in Medicine and Biology*, 24(9):1243–1270, 1998.

[11] R. Prager, A. Gee, G. Treece, C. Cash, and L. Berman. Sensorless freehand 3D ultrasound using regression of the echo intensity. *Ultrasound in Medicine and Biology*, 29(3):437–446, March 2003.

[12] R. W. Prager, A. H. Gee, G. M. Treece, and L. Berman. Decompression and speckle detection for ultrasound images using a homodyned k-distribution. *Pattern Recognition Letters*, 24(4–5):705–713, February 2003.

[13] G. Treece, R. Prager, and Gee. A. Freely available software for 3D RF ultrasound. In *Medical Image Computing and Computer-Assisted Intervention — MICCAI'04*, pages 1099–1100, Saint-Malo, France, September 2004. LNCS 3217, Springer.

[14] G. M. Treece, R. W. Prager, A. H. Gee, and L. Berman. Correction of probe pressure artifacts in freehand 3D ultrasound. *Medical Image Analysis*, 6(3):199–215, 2002.

[15] T. A. Tuthill, J. F. Krücker, J. B. Fowlkes, and P. L. Carson. Automated three-dimensional US frame positioning computed from elevational speckle decorrelation. *Radiology*, 209(2):575–582, 1998.

[16] R. F. Wagner, S. W. Smith, J. M. Sandrik, and H. Lopez. Statistics of speckle in ultrasound B-scans. *IEEE Transactions on Sonics and Ultrasonics*, 30(3):156–163, May 1983.

[17] R.F. Wagner, M. F. Insana, and D. G. Brown. Statistical properties of radio-frequency and envelope-detected signals with applications to medical ultrasound. *Journal of the Optical Society of America A*, 4(5):910–922, May 1987.

[18] R.F. Wagner, M. F. Insana, and S. W. Smith. Fundamental correlation lengths of coherent speckle in medical ultrasonic images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 35(1):34–44, January 1988.