ROTATIONAL MOTION IN SENSORLESS FREEHAND 3D ULTRASOUND

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

CUED/F-INFENG/TR 587

October 2007

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

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

Rotational motion in sensorless freehand 3D ultrasound

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

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

Abstract

Freehand 3D ultrasound can be acquired without a position sensor by using image-based positioning methods. In-plane motion can be tracked using image registration between nearby images. Elevational probe motion can be determined from the decorrelation between images. However, a freehand scan involves rotational as well as translational motion, and this affects the decorrelation. If this effect is ignored, it leads to errors in the image-based positions. In this paper, we present a technique to compensate for out-of-plane rotations, which we test using simulations and in vitro experiments. We show that by using this technique, the accuracy of image-based positioning is improved.

1 Introduction

3D ultrasound [5, 17] is a medical imaging modality with many potential applications [7]. The data can be acquired using special 3D probes, with either a 2D array of crystals or an internal mechanism to sweep the scan head over a 3D volume. Another possibility is the freehand approach, which uses a conventional 2D probe. The clinician manually sweeps the probe over a volume of interest while a sequence of images is recorded. At the same time, a position sensor attached to the probe labels each image frame with its position and orientation, resulting in a 3D volume of data. The 3D data can then be visualised in various ways to extract clinically useful information.

The freehand approach has certain advantages over 3D probes: it allows arbitrary volumes, is relatively low-cost and works with any standard 2D probe. One of the limitations of this approach is the requirement for an add-on position sensor. Sensors have only a limited operating range and they often need a direct line of sight between the device attached to the probe and the base unit. Also, careful calibration [14] is required each time the sensor is reattached to the probe.

In previous work [8], we have developed an alternative to the position sensor, which determines the relative position of two image frames using information in the images themselves. Consider the two image frames A and B in Figure 1. We can determine the in-plane motion between them (axial and lateral translation, and roll in the plane of the images) by using standard image registration and motion tracking techniques [1, 24]. For the out-of-plane motion, we use speckle decorrelation [3, 12, 21] to track the transducer motion. This makes use of the fact that the ultrasound beam has a finite width in the elevational (out-of-plane) direction, even at the focus, due to imperfect focusing. The backscattered signal at a particular point therefore depends on the scatterers in a resolution cell around that point. For nearby image frames, the resolution cells at any particular point in the two images overlap and there is some correlation between the two signals. The correlation depends on the amount of overlap, and therefore on the separation at that point. By comparing corresponding regions (patches) of data in the two frames, we can calculate a correlation value. We can then use a precalibrated decorrelation curve to convert this to a distance, as illustrated in the figure. By obtaining distance estimates for at least three non-collinear locations in the image, we can determine the elevational separation, tilt and yaw[1] of B relative to A.

[1] We use the convention that roll is the rotation in the plane of the image, about an axis in the elevational direction, yaw is about an axis in the axial direction and tilt is about an axis in the lateral direction.

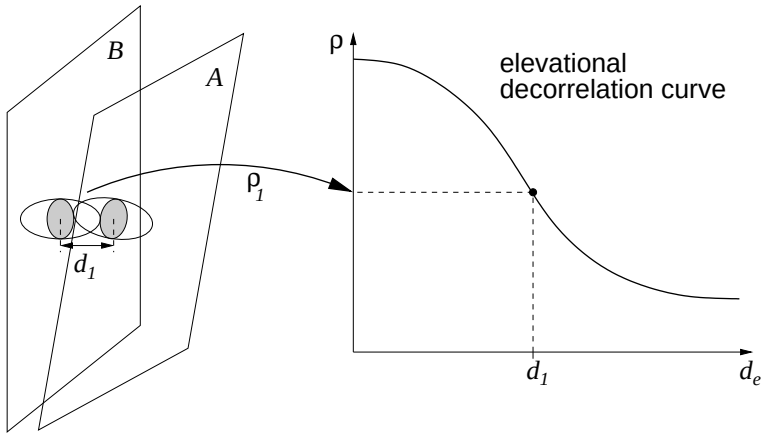

Figure 1: **Speckle decorrelation**. Consider two patches of data on nearby image frames A and B. The correlation between the data varies with the resolution cell separation. The precalibrated decorrelation curve can be used to look up a distance d_1, given the correlation ρ_1.

The elevational decorrelation curve can either be determined from ultrasound physics, or more practically by scanning a speckle phantom and measuring the correlation at different elevational separations. Since the elevational beam width, and therefore the resolution cell, varies over the image, it is necessary to calibrate several decorrelation curves for different locations in the image frame.

There are several difficulties with this speckle decorrelation technique that limit its accuracy when applied to freehand scans of in vivo data. First, the decorrelation curves are only valid for data consisting of fully developed speckle, as the calibrations are derived from a scan of a speckle phantom. Scans of real tissue include other types of scattering, which decorrelate at a different rate. In [6], we show how to correct for this effect by adjusting the decorrelation curves to a form better suited to the tissue being scanned. A further difficulty is that elevational probe motion is not the only cause of decorrelation. Other causes include physiological motion, tissue compression, transducer rotation and noise. The effect of these other sources of decorrelation is that the distance estimates will be biased. This will typically result in an overestimate of the total length of a sensorlessly reconstructed sequence. Also, depending on how the bias varies over the image frame, it can lead to over- or underestimated yaw or tilt in the reconstruction. Consequently, any measurements of anatomical features in the data will also be biased.

In recent work, we have shown that it is possible to correct for this bias by warping the reconstruction so that its overall size matches measurements from a position sensor [9]. Since one of the objectives of image-based positioning is to avoid the inconvenience of a sensor, we must be careful in our choice of sensor. We found that while it is possible to almost completely remove the bias using any six degree of freedom sensor, the sensor required is too obtrusive in the scanning process. A compromise is to use an Xsens MT9-B sensor (`http://www.xsens.com`), which uses Micro-Electro-Mechanical Systems (MEMS) magnetometers, accelerometers and rate gyros to determine its orientation. While this provides only orientation information, it does so accurately while being relatively unobtrusive. This sensor allows bias in the tilt and yaw to be corrected, but gives no real improvement in the length. In this paper, we are interested in correcting for the unwanted decorrelation at a more fundamental level, rather than as a post-processing stage. While most of the decorrelating effects mentioned above are beyond the scope of this paper, we consider here the problem of rotation. Our objective is to use measurements of the inter-frame rotation (from the orientation sensor) to compensate for the decorrelation due to this rotation, and hence obtain more accurate reconstructions of the data.

The decorrelation caused by rotation is not an unknown effect. One paper that has considered how this affects image-based positioning is [13]. It is shown that significant decorrelation can occur

with surprisingly small angles (only a few degrees) and therefore that this effect should not be ignored. Rotation decorrelation has also been used to some advantage in angular compounding [2]. This relies on the fact that a scatterer field viewed from different directions will have a different appearance in the images. In other words, the images are decorrelated. Some attempts have been made to either measure [19] or theoretically predict [2, 4, 18, 22] the rate of decorrelation with probe motion. The theoretical work has been mostly limited to analysis at the focal depth. In this paper, we are interested in the effect of rotation over the whole frame, so we take the approach of experimentally measuring the rate of decorrelation.

There are two distinct reasons why rotation would be expected to affect correlation. The first is that correlation values are not calculated at a single point in the image frame: we use the correlation between finite patches of data in the images. Typically, the image frame is divided up into small rectangular patches, as shown in Figure 7(a). For a parallel frame pair with sufficiently small patch sizes, it can be assumed that the overlap between individual resolution cells is the same over the entire patch. However, when there is some rotation, the overlaps will vary across the patch. For patches that would be coincident if not for the rotation, there will be an overall reduction in overlap over the patch and therefore a reduction in the calculated correlation value. This effect has been studied for individual A-lines in [4]. It is important to realise that this is not a significant effect at the angles considered in our application. In a typical high frame rate freehand sequence with deliberate probe rotation, the rotation between nearby frames is almost certainly going to be less than 1 degree. Over the small area of an individual patch of data, the change in overlap across the patch will make almost no difference to the correlation.

The second cause of decorrelation is the repositioning of scatterers in the resolution cell. The signal at a particular point can be modelled as the sum of backscattered signals from a large number of point scatterers within the resolution cell [23]. Each of these individual backscattered signals has a phase depending on its location relative to the transducer. The overall signal depends on whether these phases tend towards constructive or destructive interference between the individual signals. By viewing the scatterers from a different direction, their relative positions, and therefore their relative phases, will be changed and the overall signal will be decorrelated from that in the original view.

Considering rotations in terms of changes in viewing direction along an individual A-line, it can be seen that rotations about an axis parallel to the A-line will not change the relative scatterer positions. Therefore, the only decorrelation due to rotations about this axis will be due to a change in overlap over a patch, and we have explained how this is not significant. For that reason, there is no need to account for the yaw angle when considering the effect of rotations on correlation[2]. This explains why, in our previous work on sensorless reconstruction [8], we have observed considerably more error in the tilt angles than the yaw angles. That leaves just the roll and tilt rotations as a significant source of decorrelation.

In general, the effect of roll or tilt is more complicated than a simple reduction in correlation. Consider a pair of A-lines, A and B, as in Figure 2(a). The solid decorrelation curve in Figure 2(b) shows how the correlation would vary at the centre of the A-lines as B is moved past A in the direction shown, for the case where θ is zero. This is a standard, parallel-frame decorrelation curve of the type assumed for elevational speckle decorrelation. The dashed line shows a typical example of how the curve changes when θ is not zero.

Surprisingly, the peak of the adjusted decorrelation curve is offset from the zero position. This effect has been explained in [16] in terms of a preferred scanning direction for a group of scatterers. Figure 3 shows how this causes a peak offset. Due to various diffraction and focusing effects through the transducer aperture, the viewing direction can vary over the width of an ultrasound beam, as shown in Figure 3(a). Typically, the beam will be focused so that it is wide and converging above the focus, narrow at the focal depth, and wide and diverging below the focus, in both the lateral and elevational directions. It is this convergence and divergence that causes the varying viewing direction over the width of the beam. A group of scatterers will be viewed from a particular

[2]Note that this is only true for a linear array probe. For a convex probe, the effect of yaw angles would also need to be considered.

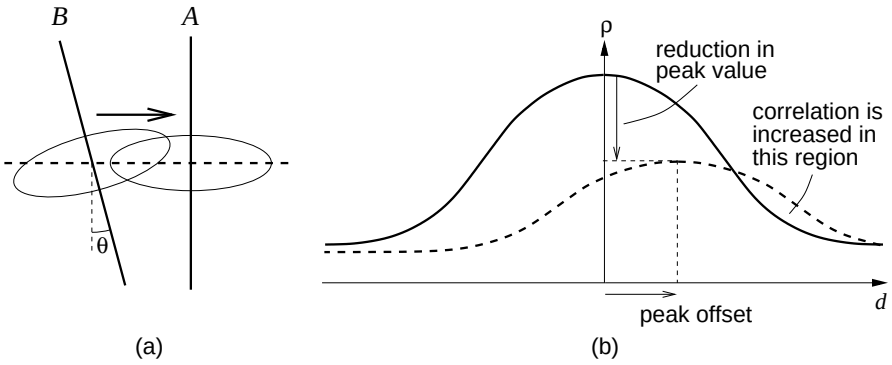

Figure 2: **Decorrelation curves**. (a) Resolution cells at the centre of a pair of A-lines, A and B, at an angle θ. (b) The decorrelation curve shows how the correlation of the signals from the two resolution cells varies as they move past each other. The solid curve is for A-lines at $\theta = 0$. The dashed curve is for $\theta > 0$.

direction, depending on their location in the converging or diverging beam. As we have discussed, these different viewing directions will result in different signals from a group of scatterers. There will be one direction, called the preferred direction in [16], that gives a high signal and therefore a bright speckle spot.

Consider now a situation where the transducer is rotated relative to a group of scatterers, as shown in Figure 3(b). The thick line indicates the location in the beam profile that looks in a particular direction, equivalent to the thick line in Figure 3(a). Note that this line is not vertical, since the preferred viewing direction for a group of scatterers is not necessarily from directly above. Comparing the rotated and unrotated beams, it can be seen that above the focus, the equivalent viewing direction is further to the right in the rotated beam than in the unrotated beam. Similarly, it is further to the left below the focus. The consequence is that in order to view a group of scatters from the same, preferred direction after rotation, and hence observe the maximum correlation, it is necessary to move the transducer left or right. This is illustrated in Figure 3(c). Above the focus, the same viewing direction is achieved after rotation by moving the transducer to the left. Below the focus, it must be moved to the right. The required offset depends on how quickly the direction varies across the beam profile and therefore on the depth-dependent width of the beam. At the focus, all directions pass through the same point, so no offset is required. As the distance from the focus increases, the direction varies over a larger beam width, so a larger offset is required to achieve the same viewing direction.

Note that for a significant peak offset, the beam intensity must be strong enough away from the beam centre to produce equivalent speckle features over a range of positions. This is more the case away from the focus, where the beam intensity decays slowly from the beam centre. Near the focus, the beam is relatively narrow, effectively meaning that a group of scatterers can only be viewed when the transducer is aimed directly at them. There is therefore very little peak offset at the focal depth.

Considering the converging beam in Figure 3(a), it can be seen that for a scatterer at a fixed depth below the transducer, the pulse reaches it sooner when it is at the edge that when it is in the centre of the beam, due to the focusing applied to the beam. The scatterer therefore appears to be at a shallower depth at the edge. This would cause the image of the scatterer (the point spread function (PSF)) to be curved upwards at the edges. Similarly, a divergent beam would cause a downward curving PSF. This relationship between focusing, PSF curvature and peak offset has been analysed in detail in [11] for in-plane rotation, where it was shown that the lateral peak offset can be determined from the curvature of the axial-lateral PSF. In fact, the same arguments apply equally well to out-of-plane focusing, hence the peak offset also occurs in the elevational decorrelation curve.

These effects of rotation on correlation have various implications for image-based positioning.

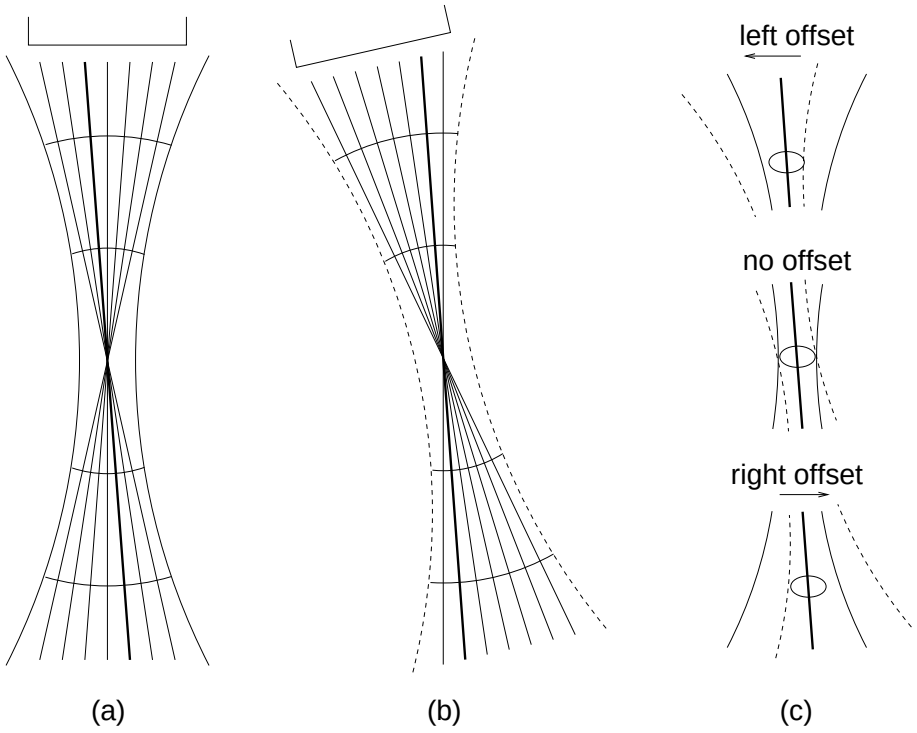

Figure 3: **Correlation peak offset**. (a) A transducer aperture produces a finite width ultrasound beam which varies in direction across its width. The curved lines across the beam indicate where the pulse will be at a particular time after it is produced. They therefore show contours of constant apparent depth. They are curved due to the diffraction and focusing of the beam. (b) As the transducer rotates, the location in the beam that points in a particular direction varies. (c) In order to view a group of scatterers from the same direction as in the unrotated position, the rotated beam must be offset to the left or right relative to the unrotated beam. The required offset depends on whether the beam is converging or diverging, and therefore on the depth in the image.

For in-plane speckle tracking, the lateral peak offset effect means that the standard approach of searching for the correlation peak will not give the correct in-plane alignment when there is in-plane rotation. However, this effect could be compensated for using the theoretical peak offset determined from the PSF curvature, derived in [11].

For out-of-plane motion tracking, the effect of rotation is more complicated and there have been no previous attempts to compensate for it. Since any effect that causes decorrelation will bias the out-of-plane motion estimates, it is necessary to consider both the roll and the tilt. Tilt causes the change in the elevational decorrelation curves shown in Figure 2 and this results in incorrect distance estimates. Notice that it is possible in some situations for the tilt to result in a higher correlation value, causing underestimated distances. Also, the in-plane roll can affect the out-of-plane distance estimates, because it changes the correlations between patches. In general, the peak correlation will be reduced by in-plane rotation and this will give incorrect out-of-plane distance estimates regardless of whether the elevational decorrelation curve is corrected for tilt. However, this paper is concerned only with the effect of out-of-plane rotation (specifically tilt) on the elevational distance estimates.

A further complication associated with speckle decorrelation is illustrated in Figure 2: for each correlation value, the decorrelation curve gives two possible distances. For parallel A-lines, the two distances have equal magnitude, but opposite signs, so it is only the direction that is uncertain. We have considered this direction ambiguity in detail in [8]. Each corresponding pair of data

patches provides a distance estimate, which could be in either the positive or negative elevational direction. In order to reconstruct the data from these distance estimates, we must determine a consistent set of directions throughout the data. The necessary details of this process are reviewed in Section 3, but for a complete description, we refer the reader to [8]. The ambiguity becomes more difficult to resolve in the current situation where the curve is adjusted for tilt. We must now consider two distinct distances, rather than a single distance with two possible directions.

There are therefore two objectives in this paper. First, we need to adjust the decorrelation curve to allow for tilt, which will give us two possible distance values, one of which will be correct. We then need to disambiguate the distance estimates to produce a correct, image-based reconstruction of the data. The paper is organised as follows. In Section 2, we describe the method we use to produce a tilt-corrected decorrelation curve. We then explain in Section 3 how we produce a correct reconstruction from the ambiguous distance data. Section 4 describes our experimental methodology for evaluating the various techniques, along with the results. Finally, we conclude in Section 5.

2 Rotation correction

In our previous work, we produced precalibrated decorrelation curves by scanning a speckle phantom with parallel image frames at known elevational separations and recording the correlation values at each separation. We allowed for the varying elevational beam width by dividing the image into a grid of patches and calibrating a separate curve for each patch. The full experimental details of this method are given in [6].

Here, we are interested in adjusting the standard decorrelation curves according to the measured tilt angle between frame pairs. We take the very simple approach of calibrating an additional decorrelation curve for each patch at a known tilt angle. We then have a pair of decorrelation curves like those in Figure 2: one for zero angle and one for a known tilt. The adjusted decorrelation curve is produced by interpolating between the two according to the actual tilt angle between the image frames, measured using the MT9-B. Note that no attempt is made to correct for the effect of in-plane rotation. If there were a roll rotation, this would have a detrimental effect on the reconstruction in both the in-plane and out-of-plane directions, as discussed in Section 1.

The interpolation is performed by fitting a model to the two calibrated decorrelation curves and linearly interpolating the parameters of the model. We have observed that a four parameter offset scaled Gaussian gives a good fit to the range of decorrelation curves in a typical calibration. The curve model over elevational distance d is defined as

$$\rho(d) = Ae^{(-\frac{(d-c)^2}{2\sigma^2})} + B$$

with parameters A, B, c and σ. We find the best least squares fit for the curves using an iterative optimisation algorithm (Levenberg-Marquardt [15]). Figure 4 shows a typical example of the model fitted to a pair of decorrelation curves, along with an interpolated decorrelation curve at a tilt angle of 30% of the calibrated tilt. Given the parameters of the interpolated curve, it is then simple to show that the pair of distance estimates for a correlation ρ is given by

$$d = c \pm \sqrt{-2\sigma^2 \ln\left(\frac{\rho - B}{A}\right)}.$$

For tilt values above the calibrated value, we simply extrapolate the curve parameters. However, in the experiments in this paper, we have been careful to calibrate for a tilt value that is at the upper end of the inter-frame tilt angle expected in a deliberately tilted freehand sequence, so we do not expect to encounter large extrapolation errors. When the measured tilt is negative, the appropriate decorrelation curve is the same as for positive tilt but reflected in the ρ axis. In this case, we interpolate as usual using the absolute tilt value, then multiply the resulting distance estimates by -1 to get appropriate values for negative tilt.

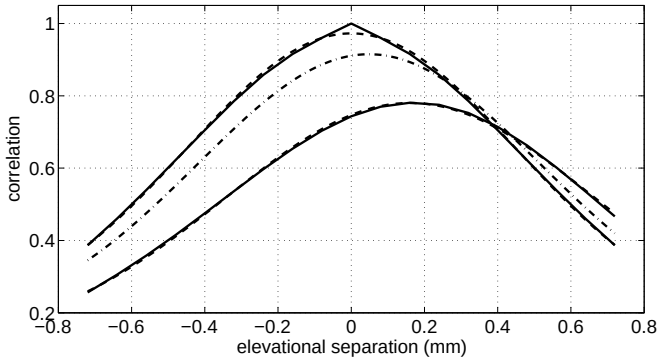

Figure 4: **Interpolation of decorrelation curves**. The solid lines show two calibrated decorrelation curves, one with and one without tilt. The dashed lines show the curve model fitted to the calibrated curves. The dash-dot line is an interpolated curve at a tilt of 30% of the calibrated tilt, obtained by linearly interpolating the parameters of the fitted curves.

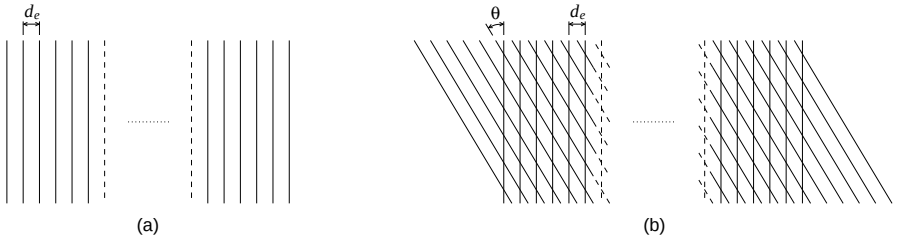

Figure 5: **Calibration sequences**. (a) A linear sequence of frames with elevational spacing d_e, used for calibrating a zero tilt decorrelation curve. (b) An elevationally spaced sequence, tilted by θ and overlayed on the linear sequence of (a), used for calibrating a tilt-compensated decorrelation curve. Patch centres in the tilted sequence coincide approximately with patch centres in the linear sequence. In our experiments, we divide the image frame into 8 patches across by 12 down the image. For clarity, the figure shows the situation for only 6 patches down the image.

Calibration

In order to calibrate the decorrelation curves, we must measure the correlations between patches of data at various elevational separations, both with and without a tilt rotation. Figure 5(b) illustrates the frames we recorded for this purpose, comprising two overlayed sequences. The first is a sequence of 101 parallel frames, with an elevational separation of 0.08 mm between each adjacent frame pair, as shown in Figure 5(a). We refer to this as the linear sequence. From this, we determine the decorrelation curves for no tilt, by comparing nearby frames at various elevational separations. The second sequence consists of 112 parallel frames, at a tilt of 1.275° relative to the first sequence. By comparing frames from this tilted sequence with frames from the linear sequence, we can build up a tilted decorrelation curve.

The image frame is divided into a grid of patches — 8 across and 12 down the image. The tilt angle of 1.275° was chosen so that each patch centre in the linear sequence coincides approximately with an equivalent patch centre in the tilted sequence. By selecting appropriate linear and tilted frames from the sequences, individual patches can be compared for elevational separations at multiples of 0.08 mm, in the same way as for the zero tilt decorrelation curves. Note that because of the rotation, the vertical positions of the patch centres towards the top and bottom of the image do not line up exactly. However, the misalignment is very small and we can assume that it is negligible.

For this calibration, and for all subsequent experiments, we used a 5–10 MHz linear array probe

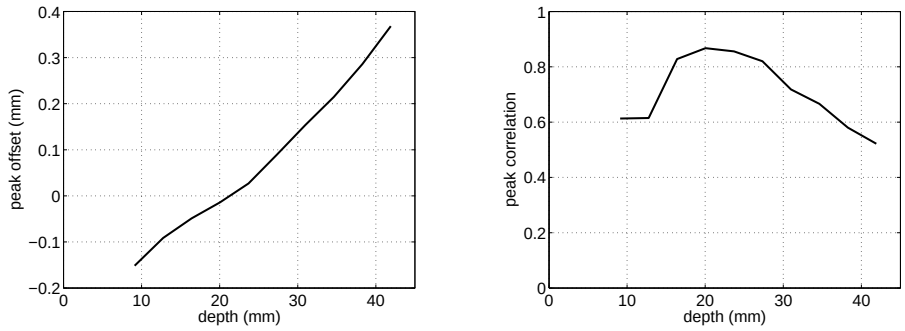

(a) **(b)**

Figure 6: **Variation of tilt calibration with depth**. Variation of (a) the peak offset and (b) the peak value with depth for a tilted decorrelation curve. For comparison, a zero tilt curve always has a peak offset of 0 and a peak value of 1.

connected to a Dynamic Imaging Diasus ultrasound machine (`http://www.dynamicimaging.co.uk`). The depth setting was 4 cm with a single lateral focus at 2 cm. Analogue RF ultrasound signals were digitised after receive focusing and time-gain compensation, but before log-compression and envelope detection, using a Gage Compuscope CS14200 14-bit digitiser (`http://www.gage-applied.com`). The system operates in real time, with acquisition rates of about 30 frames per second. Sampling was at 66.67 MHz, synchronous with the ultrasound machine's internal clock: this synchronisation minimises phase jitter between vectors. The acquired vectors were filtered with a 5–10 MHz passband 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 0.3 mm per vector in the lateral direction. The scanning subject was a speckle phantom consisting of an agar cylinder with a uniform distribution of aluminium oxide powder providing scattering.

Figure 6(a) shows how the peak offset effect varies with depth in the image for the calibrated tilt angle. This is for a positive tilt, which is defined as having a more positive elevational offset at the bottom of the image than at the top. The peak offset depends on the focusing of the ultrasound beam. The offset is therefore smallest around the elevational focal depth, which is about 20 mm for our probe. Similarly, Figure 6(b) shows that the reduction in peak value is smallest at the focus. We would therefore expect the distance estimation errors caused by ignoring the tilt effect to be more severe at the top and bottom of the image than around the focal depth.

3 Multiple distance ambiguity

Each pair of distance estimates consists of one correct distance and one other that has no physical significance. The challenge is to find the correct distance in each pair and use these to reconstruct the frame sequence.

The first stage in this process is to resolve the ambiguity over individual frame pairs. This is not necessarily a simple task. For each patch pair, the correlation value gives an upper and a lower distance estimate, as shown in Figure 7. Depending on the relative angle and position of the two frames, the correct distance could be either of these and it will not necessarily be the same one for each patch over the frame pair.

In order to resolve this ambiguity, we make the assumption that the image frames are planar and therefore that the correct distances will define a flat plane. We can then look for the correct distance in each pair by looking for those that best define a plane. This process is simplified considerably by the fact that we have a sensor to measure orientations. We proceed by considering each individual distance estimate in turn and using the measured orientation to define a plane from that single estimate. For each of these candidate planes, we then consider the other patches

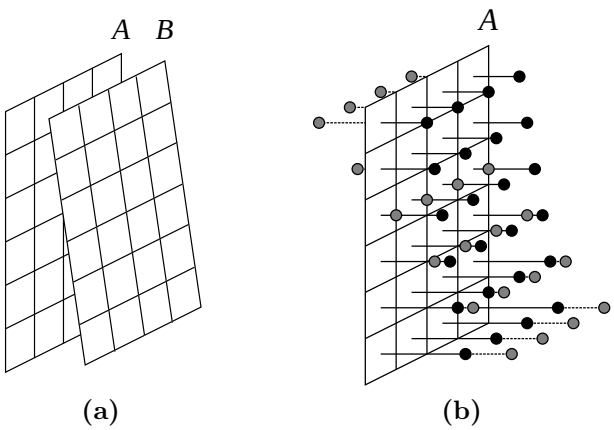

Figure 7: **Frame pair ambiguity**. (a) Each corresponding patch pair on frames A and B provides a correlation value. (b) This is converted to two distance values, one of which is the correct patch separation. The correct distances, shown by the black circles, define a flat plane, indicating where frame B is relative to A. The other distance in each pair, indicated by the grey circles, is an artefact of the decorrelation curve and has no physical significance.

and for each, we take the distance estimate that is closest to the defined plane. We then define a new plane by performing a least squares fit to these distances and noting the least squares error. The distances contributing to the plane with the minimum error are taken as the correct distances in each pair.

It may appear to be sensible at this stage to continue with just the 'correct' distances, distinguished by the sensor measurements, and ignore the rest, which are simply an artefact of the decorrelation curve and do not have any physical significance. However, as the peak offset in the decorrelation curves varies approximately linearly with depth (see Figure 6(a)), the two sets of distances often look similarly planar. In addition, since the peak offset varies in a way consistent with the tilt direction, it is often the case that the two sets of distances give planes at very similar angles. It is therefore not possible to use the sensor measurements to distinguish the correct solution. The result of this process is therefore not to produce a single set of planar distance estimates, but rather to sort the pairs of individually ambiguous estimates into two consistent solutions.

Having obtained consistent solutions for each individual frame pair, the next stage of the disambiguation process is to form them into a complete reconstruction of the frame sequence. In our previous work on reconstruction with a direction ambiguity [8], we have shown how to produce a consistent reconstruction by considering three frames at a time, rather that just frame pairs. This is illustrated in Figure 8. Given two frame pairs, AB and BC, each will have its own ambiguous pair of solutions, which we refer to as a forward and backward solution. Considering the forward solution of AB, we cannot tell whether the correct sequence is obtained by concatenating the forward or backward solution of BC. We overcome this ambiguity by also considering the two solutions of the AC frame pair. By considering the three together, we can see which of the two directions of BC results in the most consistent set of solutions for the three frame pairs. We then continue to build up the sequence by considering frames B, C and a new frame, D, in the same way. A more detailed description of this reconstruction process is included in [8].

The final reconstruction of the whole sequence depends on whether it is initialised with the forward or backward solution of frame pair AB. If we consider both possibilities, we arrive at two different, but self-consistent, reconstructions of the sequence, one of which is correct. Without tilt compensation, the two solutions are mirror images of each other and measurements taken from either reconstruction would be equally valid. With tilt compensation, the two solutions are completely different, so it is important to choose the correct one. Unfortunately, there is no reliable way to determine automatically which one is correct. For the same reason that the sensor cannot be used to determine the correct offset for each frame pair, it also cannot reliably determine the

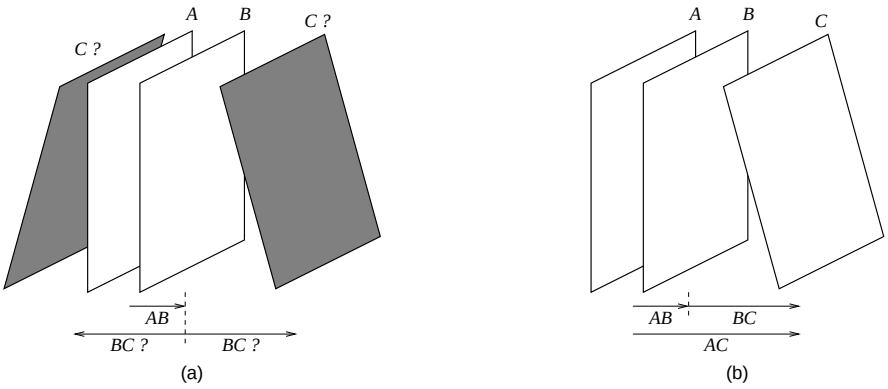

Figure 8: **Reconstruction ambiguity between frame pairs**. (a) Each frame pair has two possible solutions. It is not possible to position frame C relative to frame A by considering only the offsets between adjacent frame pairs. (b) This ambiguity can be resolved by considering three frames at a time.

correct final reconstruction from the two possibilities. This must be determined manually, either from the scanning protocol or from features in the images.

4 Experiments and results

In order to demonstrate the improvement in reconstruction accuracy achieved by correcting the decorrelation curve, we recorded several test sequences of frames. The scanning subject was the same phantom used for the decorrelation calibration. Four sequences of 15 frames were recorded with known elevational separation and tilt, as shown in Figure 9. The elevational offset at the frame centre was 0.2 mm between each frame, and the tilt angle was set to either 0.30° or 0.45°. We also produced one additional test data set and calibration using simulated data (generated using Field II [10]) modelling the 5–10 MHz probe used for the real experiments.

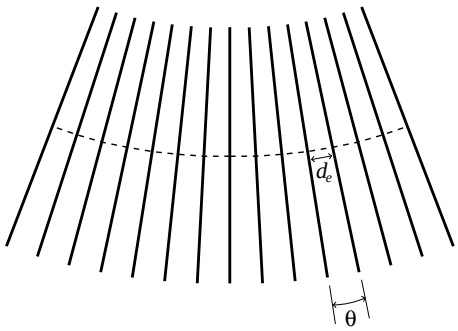

Figure 9: **Frame sequence for test data**. The test sequence consists of evenly spaced frames with an elevational centre offset (along the dashed line) of $d_e = 0.2$ mm and a fixed tilt angle, θ, between each adjacent frame pair. We produced sequences with tilt angles of 0.30° and 0.45°.

For the initial experiments, we avoided the distance ambiguity issue by making use of the known frame positions. With the correct distance known for each patch pair, we can determine which side of the decorrelation curve should be used and therefore avoid having multiple distance results. In this way, we are able to demonstrate the advantage of using a corrected decorrelation curve independently of the other factors that may affect the accuracy of the sensorless reconstruction. For comparison, we reconstructed the sequences both with and without tilt compensation, by

		actual	uncorrected	corrected
simulated data	**length**	2.80	2.41	2.57
tilt = 0.30°	**tilt**	4.20	0.94	3.99
data set 1	**length**	2.80	4.42	4.15
tilt = 0.30°	**tilt**	4.19	1.04	3.08
data set 2	**length**	2.80	3.37	2.93
tilt = 0.30°	**tilt**	4.19	0.23	1.96
data set 3	**length**	2.79	3.71	3.15
tilt = 0.45°	**tilt**	6.28	1.73	4.92
data set 4	**length**	2.79	3.43	2.85
tilt = 0.45°	**tilt**	6.28	0.73	3.84

Table 1: **Accumulated elevational length and tilt for test sequences**. The lengths are in millimetres and the tilt angles are in degrees. The values were obtained by concatenating the out-of-plane transformations between adjacent image frames in the reconstructed sequences. Length is defined as the accumulated value of elevational offset at the centre of each frame. Adjacent frame pairs had a tilt angle of either 0.30° or 0.45°, as indicated. The accumulated values are the result of concatenating 14 transformations over the 15-frame sequences.

using either the known inter-frame tilt or assuming a tilt of zero. Table 1 shows the accumulated elevational offset at the frame centre, along with the accumulated tilt angle, for each of the reconstructed sequences.

The same five test sequences were also reconstructed by automatically resolving the multiple distance ambiguity. In each case, the automatic methods were able to correctly resolve the pairs of distance estimates into two consistent solutions. However, it is not the purpose of this paper to thoroughly assess the robustness of the automatic disambiguation techniques. This has already been done for the more basic direction ambiguity in [8]. Here, we have simply explained the modifications that are necessary in order to apply the automatic methods to tilt-compensated distance estimates. The success of these disambiguation techniques depends very much on the quality of the distance data and since we are currently scanning a speckle phantom, the success is not surprising.

As a final test, we recorded two freehand sequences, again on the same speckle phantom used for the calibration. The sequences had deliberate elevational and tilt motion, but no intentional in-plane motion or yaw rotation. Frame positions were recorded using a Polaris optical tracking system (`http://www.ndigital.com`) from which we used only the orientation measurements for tilt compensation. The multiple distance ambiguity was resolved using the automatic techniques. Figure 10 shows reslice images along the length of the sequences, reconstructed both with and without tilt compensation, compared to the reconstruction obtained using the positions recorded by the position sensor. For each image-based method, the figure shows both of the final reconstructions produced by the automatic techniques. Clearly, it is much more important to choose the correct one when performing a tilt-compensated reconstruction.

The effect of tilt compensation is shown clearly in Table 1. By using a corrected decorrelation curve, both the lengths and tilts of the reconstructed sequences are improved to some extent. This improvement can also be seen in the freehand reconstructions of Figure 10, although the most noticeable difference is in the tilt angle. Despite this clear improvement, the tilt compensation does not result in a complete removal of the sensorless bias, even for the simulated data. The only source of bias in the simulated data is the use of linear interpolation of the decorrelation curve parameters. The linear variation is not necessarily a valid assumption: in particular, the peak value of the decorrelation curve does not decrease linearly with rotation. Figure 11 shows how this affects the distance estimates. The interpolated decorrelation curve will have a peak value slightly less than the correct value. As a result, the pair of distances derived from the curve for a particular correlation value will be closer to the peak position. When the offset in the peak position is reasonably small, as in our experiments, this bias tends to give a slightly underestimated length

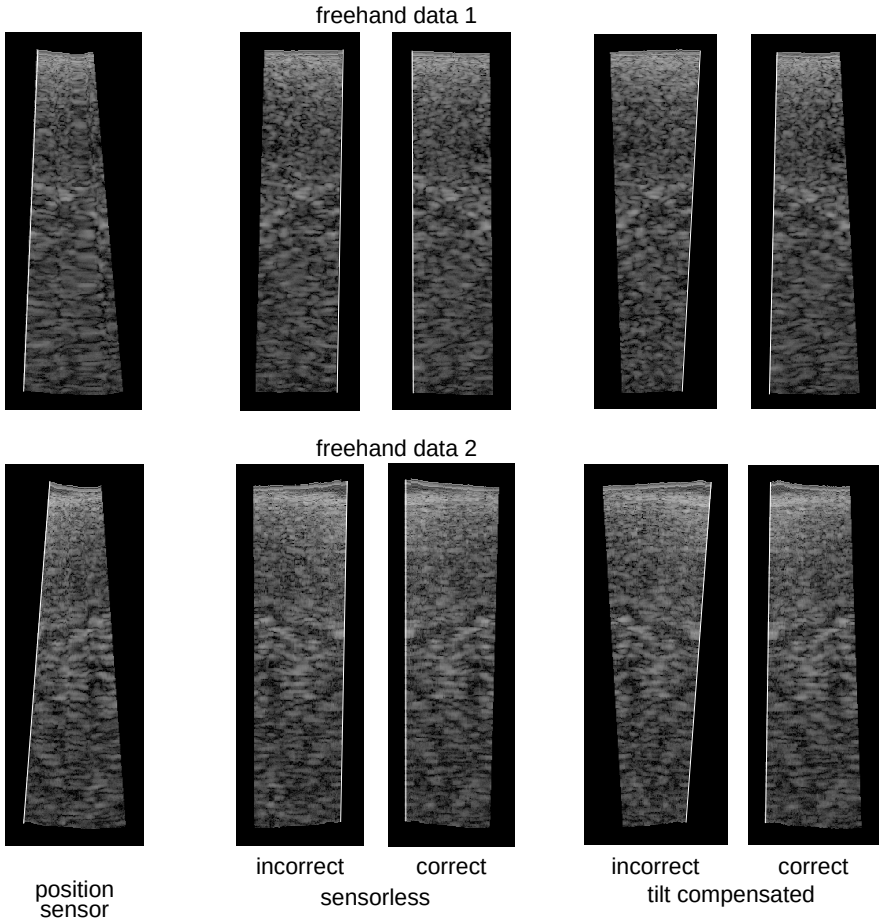

Figure 10: **Reslice images of freehand data sets**. The reslices are in the axial-elevational direction, showing the length and tilt of the reconstructed sequences. For the image-based reconstructions, both of the two possible reconstructions are shown. The solid white line indicates the first frame in the sequence. Given the elevational direction of the sensor data, it is obvious that the right of the two images is the correct one, but this would not be apparent in practice where we have only orientation measurements.

in the tilt-compensated reconstruction.

For the data recorded on the speckle phantom, the distances tend to be overestimated. This difference can be attributed to the other sources of decorrelation that we mentioned earlier. Probe pressure distortion was a particular difficulty in our experimental setup and there is always some amount of noise in the recorded data. This reduction in correlation leads to overestimated distances. However, it is still clear that tilt compensation improves the reconstruction.

In terms of practical usefulness, the reduction in the length bias is much more important. In recent work, we have shown that it is possible to almost completely correct the tilt bias in a sensorless reconstruction by warping it to match orientations measured by a sensor [9]. Given that we require an orientation sensor to correct the decorrelation curves, there is no reason not to correct the tilt completely using the sensor. The real advantage of tilt compensation is that we achieve a reduction in the length bias, which we cannot correct directly using an orientation sensor.

A limitation of this work is that the evaluations were performed on the same speckle phantom, with the same scattering behaviour, as was used to calibrate the various decorrelation curves. In contrast, real tissue would exhibit regions of more coherent scattering that decorrelate at a slower

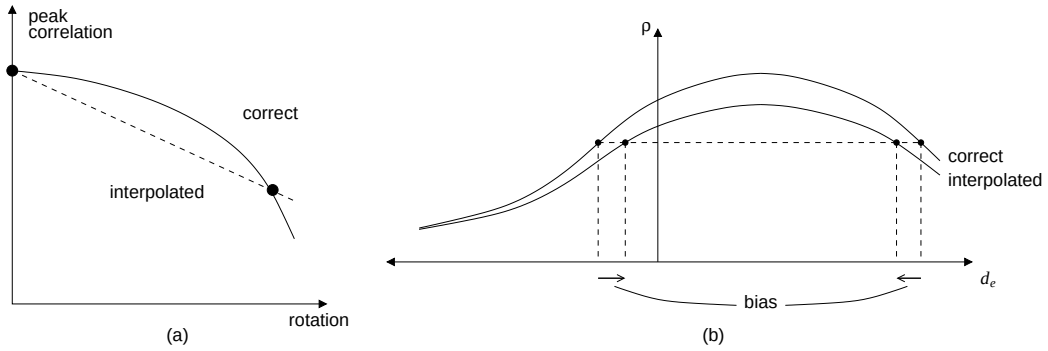

Figure 11: **The effect of linear peak interpolation**. (a) The linear interpolation of the decorrelation curves means that the peak value will be lower than that of the correct curve. (b) This introduces a bias in the distance estimates.

rate than fully developed speckle. In [6], we describe how the calibrated decorrelation curves can automatically be adapted to allow for such coherent scattering, with significant benefit to the sensorless reconstruction accuracy. However, simultaneously adapting the curves to allow for both tilt and coherent scattering is a challenge beyond the scope of this paper. The two phenomena are not independent. For example, the extent of the peak offset depends on the type of scattering, with resolvable features offset by less than the overlayed speckle pattern [11].

5 Conclusions

We have demonstrated that an out-of-plane tilt rotation between a pair of frames causes a change in the elevational correlation value and, as a result, image-based reconstructions of the data are inaccurate. By measuring the inter-frame tilt, it is possible to obtain tilt-compensated decorrelation curves. We have shown that by using these modified decorrelation curves, we obtain improved reconstructions and, in particular, the overestimate in the total length of a reconstruction is reduced.

There is scope for future work to further improve the accuracy of image-based positioning. In order to correctly reconstruct entirely unconstrained freehand data, there is a need to compensate for in-plane roll. Not only does this affect the accuracy of in-plane alignment estimates, through the position of the peak correlation, but also out-of-plane estimates through the value of the peak. Also, the limitation of requiring fully developed speckle needs to be addressed, in order to apply the technique to scans of real tissue. Finally, there are other sources of decorrelation, including the effect of probe pressure distortion, that need to be allowed for. While it is possible to correct for in-plane distortion by motion tracking [20], the effect this has on the correlation value has not yet been investigated.

References

[1] L. N. Bohs, B. J. Geiman, M. E. Anderson, S. C. Gebhart, and G. E. Trahey. Speckle tracking for multi-dimensional flow estimation. *Ultrasonics*, 38(1–8):369–375, March 2000.

[2] C. B. Burckhardt. Speckle in ultrasound B-mode scans. *IEEE Transactions on Sonics and Ultrasonics*, 25(1):1–6, January 1978.

[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(1):38–44, 1997.

[4] Q. Chen, A. L. Gerig, U. Techavipoo, J. A. Zagzebski, and T. Varghese. Correlation of RF signals during angular compounding. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 52(6):961–970, June 2005.

[5] A. Fenster, D. B. Downey, and H. N. Cardinal. Three-dimensional ultrasound imaging. *Physics in Medicine and Biology*, 46(5):R67–R99, May 2001.

[6] A. H. Gee, R. J. Housden, P. Hassenpflug, G. M. Treece, and R. W. Prager. Sensorless freehand 3D ultrasound in real tissue: speckle decorrelation without fully developed speckle. *Medical Image Analysis*, 10(2):137–149, April 2006.

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

[8] R. J. Housden, A. H. Gee, G. M. Treece, and R. W. Prager. Sensorless reconstruction of unconstrained freehand 3D ultrasound data. *Ultrasound in Medicine and Biology*, 33(3):408–419, March 2007.

[9] R. J. Housden, G. M. Treece, A. H. Gee, and R. W. Prager. Hybrid systems for reconstruction of freehand 3D ultrasound data. Technical Report CUED/F-INFENG/TR 574, Cambridge University Department of Engineering, March 2007.

[10] J. A. Jensen. Field: a program for simulating ultrasound systems. In *Proceedings of the 10th Nordic-Baltic conference on Biomedical Imaging*, pages 351–353, Tampere, Finland, June 1996.

[11] F. Kallel, M. Bertrand, and J. Meunier. Speckle motion artifact under tissue rotation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 41(1):105–122, January 1994.

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

[13] P-C. Li, C-Y. Li, and W-C. Yeh. Tissue motion and elevational speckle decorrelation in freehand 3D ultrasound. *Ultrasonic Imaging*, 24(1):1–12, January 2002.

[14] L. Mercier, T. Langø, F. Lindseth, and L. D. Collins. A review of calibration techniques for freehand 3-D ultrasound systems. *Ultrasound in Medicine and Biology*, 31(2):143–165, February 2005.

[15] J. J. Moré. The Levenberg-Marquardt algorithm: implementation and theory. In *Numerical Analysis*, pages 105–116. Lecture Notes in Mathematics 630, Springer-Verlag, 1977.

[16] D. C. Morrison, W. N. McDicken, and D. S. A. Smith. A motion artefact in real-time ultrasound scanners. *Ultrasound in Medicine and Biology*, 9(2):201–203, March–April 1983.

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

[18] M. O'Donnell and S. D. Silverstein. Optimum displacement for compound image generation in medical ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 35(4):470–476, July 1988.

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

[20] 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):199–214, September 2002.

[21] 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, November 1998.

[22] 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.

[23] P. N. T. Wells and M. Halliwell. Speckle in ultrasonic imaging. *Ultrasonics*, 19(5):225–229, September 1981.

[24] L. Weng, A. P. Tirumalai, C. M. Lowery, L. F. Nock, D. E. Gustafson, P. L. von Behren, and J. H. Kim. US extended-field-of-view imaging technology. *Radiology*, 203(3):877–880, June 1997.