
**Correction of Probe Pressure Artifacts
in Freehand 3D Ultrasound
— Further Results and Convex Probes**

G.M. Treece, R.W. Prager, A.H. Gee and L.Berman

CUED/F-INFENG/TR 421

October 2001

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

E-mail: gmt11,rwp,ahg@eng.cam.ac.uk, lb@radiol.cam.ac.uk

Abstract

An algorithm has been presented in [5] which combines non-rigid image-based registration and conventional position sensing to correct probe-pressure-induced registration errors in free-hand three-dimensional (3D) ultrasound volumes. Visual results indicated that the algorithm increased both the clarity and accuracy of reslices through *in vivo* volumetric data sets. These initial findings are verified in this report by more rigorous experiments leading to numerical results. In addition, it is shown how the original algorithm can be extended to apply to convex probes which generate compressions varying in the lateral direction as well as in depth.

Contents

1	Introduction	2
2	Further results	2
2.1	Panoramic scans	2
2.2	Repeated volume scans	5
2.2.1	Repeated scans in the same direction	5
2.2.2	Repeated scans in varying directions	10
3	Extension to convex probes	11
3.1	Incorporating a model for convex probe deformation	15
3.2	Calculating pressure correction with convex probe deformation	15
3.3	Automatically detecting the type of probe	18
3.4	Examples of convex probe data correction	20
4	Conclusions	20

1 Introduction

The motivation and background to this work, a description of the algorithm, and initial results can all be found in the companion technical report [5].

2 Further results

B-scans were acquired with a Diasus ultrasound machine¹, using a 10-22MHz linear array probe, on a 3cm depth setting. 8-bit digital log-compressed data was transferred via ethernet at 25 B-scans per second to an 800MHz PC running Linux. The probe position was sensed by a Polaris² optical system tracking an AdapTrax³ tracker mounted on the probe, and the system was calibrated to an accuracy of $\pm 0.20\text{mm}$ RMS. Calibration, acquisition, processing and display of the data were performed by Stradx [3]⁴. Pressure corrections were calculated at between 3 and 4 B-scans per second.

There have been several improvements to the system over that described in [5]. Most notably:

- The AdapTrax tracker is more accurate and visible from a wider range of angles than the flat tracker previously used.
- The accuracy with which the B-scan image data was time-stamped has increased.
- The accuracy with which the time offset between probe positions and B-scans is calculated has been increased from $\pm 30\text{ms}$ to $\pm 5\text{ms}$.

These improvements combine to increase the spatial accuracy from $\pm 0.35\text{mm}$ to $\pm 0.20\text{mm}$, after temporal and spatial calibration.

2.1 Panoramic scans

The panoramic scan of a human thyroid in [5] was repeated with the more accurate system. Figure 1(a) shows how the scan is performed: the probe is moved within the plane of the B-scans such that many overlapping images are acquired. Data from the centre of each of these images is pasted together to form the panoramic image in Figure 1(b). The resulting images after rigid and non-rigid (pressure) correction are shown in Figures 1(c) and (d).

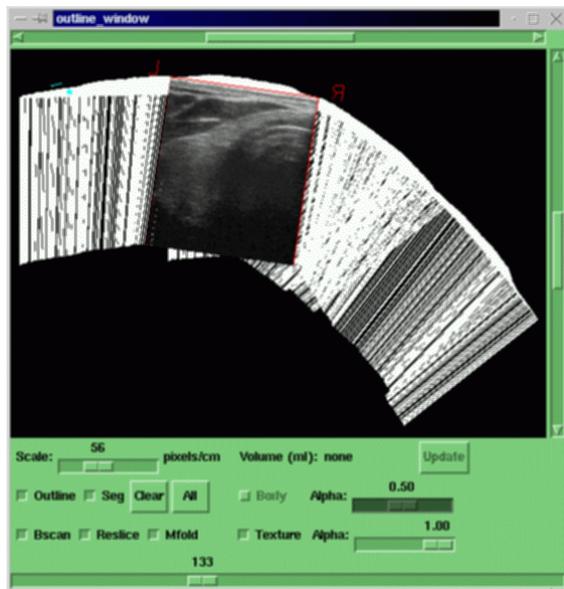
Figure 2 shows detail from two regions of the panoramas in Figure 1. If the non-rigid correction has been successful, Figures 2(d) and (h) should be similar to the B-scans in Figures 2(a) and (e) respectively. Clearly the non-rigid correction generates more accurate panoramas than the original data, or with only rigid correction applied. In addition, the panorama can accurately model the curve of the skin surface at the neck, even though each individual B-scan is flat at the top where it is compressed against the surface of the linear probe.

¹Dynamic Imaging Ltd., <http://www.dynamicimaging.co.uk/>

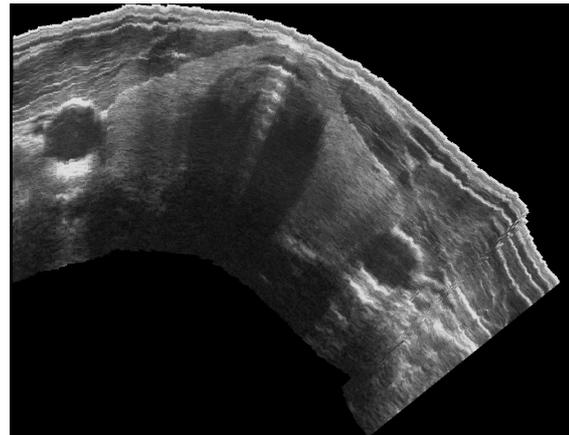
²Northern Digital Inc., <http://www.ndigital.com/>

³Traxtal Technologies, <http://www.traxtal.com/>

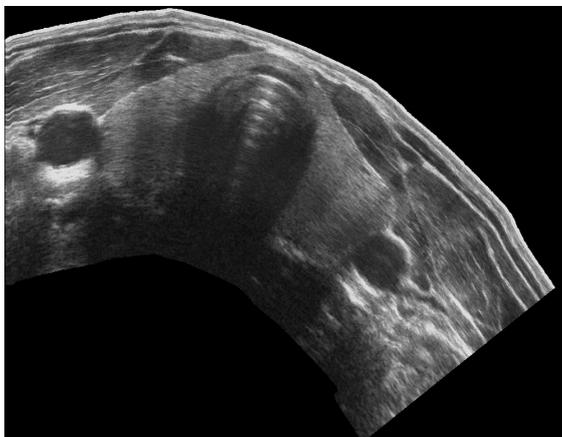
⁴<http://svr-www.eng.cam.ac.uk/~rwp/stradx/>



(a) B-scan outlines



(b) original panorama



(c) rigid correction



(d) non-rigid correction

Figure 1: Correction of probe pressure for a panorama of the human thyroid. (a) shows how the panorama in (b) is constructed from a sequence of individual B-scans (each ‘goal post’ in (a) represents one B-scan). (c) and (d) are the same panorama as in (b) after rigid and non-rigid correction respectively.

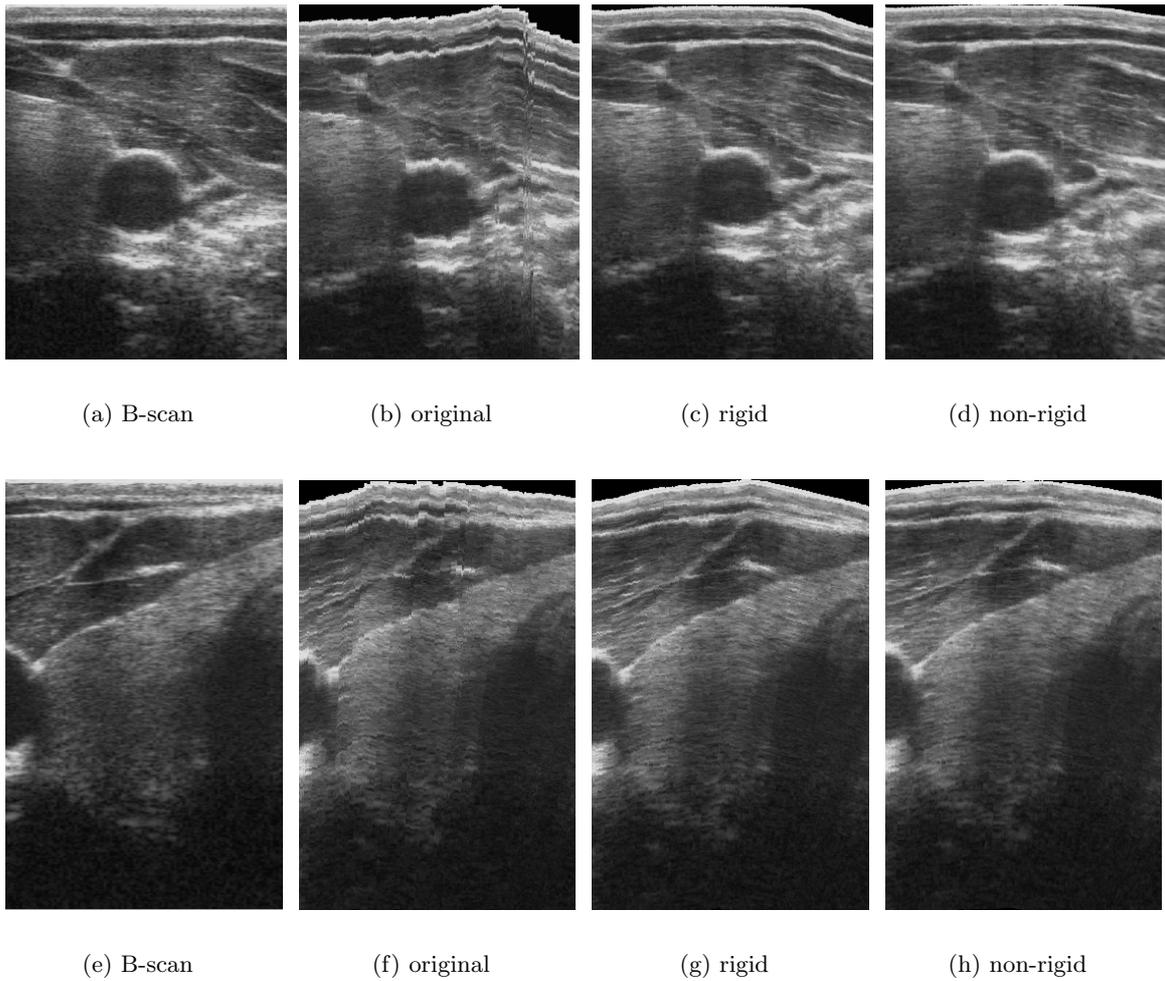
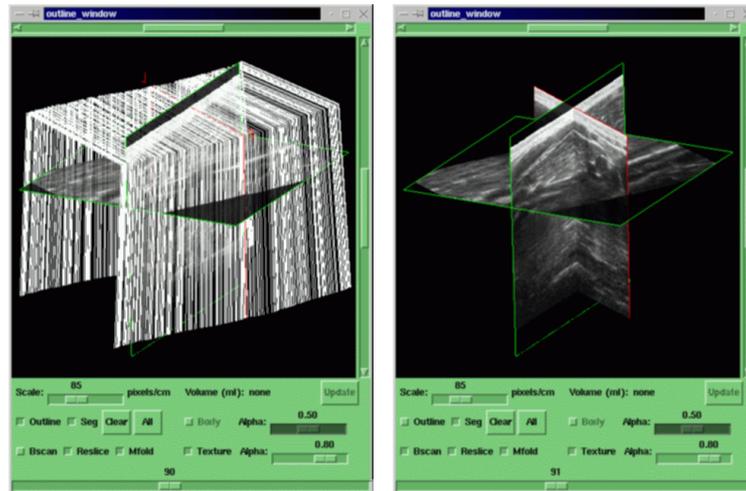


Figure 2: Selected regions from the panorama of Figure 1. (a) and (e) are individual B-scans, (b) to (d) and (f) to (h) are corresponding regions from the panoramas.



(a) sweep pattern

(b) reslices

Figure 3: 3D ultrasound examination of part of the forearm. (a) shows the sweep pattern which was used for ten examinations of the same location. (b) shows one of the original B-scans, and the two reslices in Figures 4 and 5.

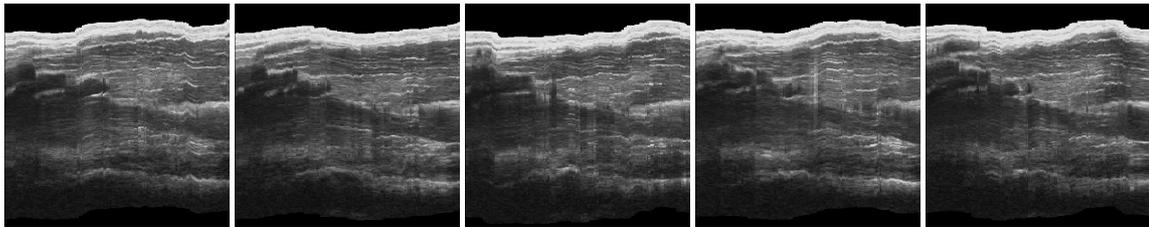
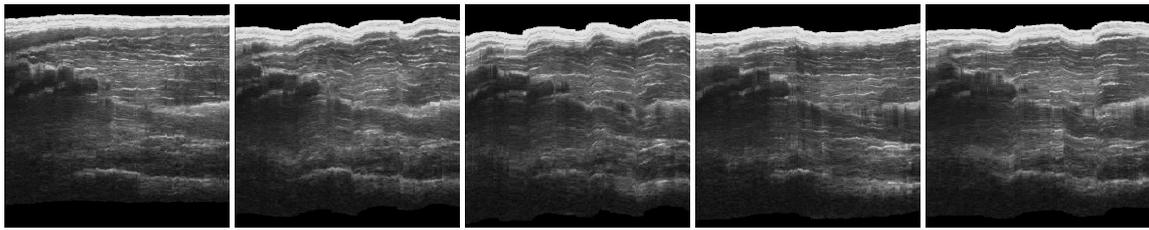
2.2 Repeated volume scans

Volume data sets are acquired by moving the probe in a direction perpendicular to the plane of each B-scan, thus ‘sweeping out’ a volume of data. As such, they are much harder to correct for pressure effects, since unlike a panoramic scan, in general no part of each B-scan overlaps with previous scans.

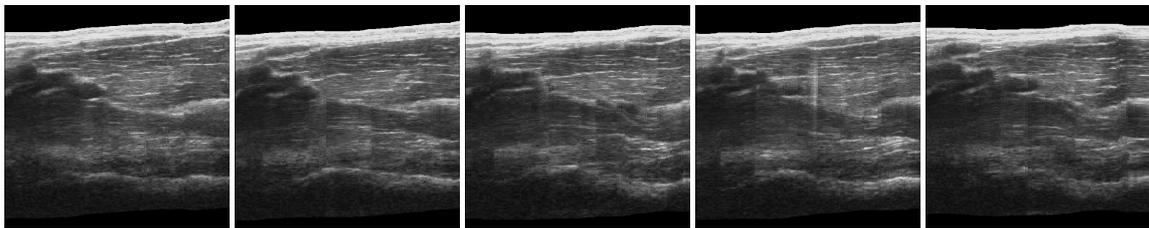
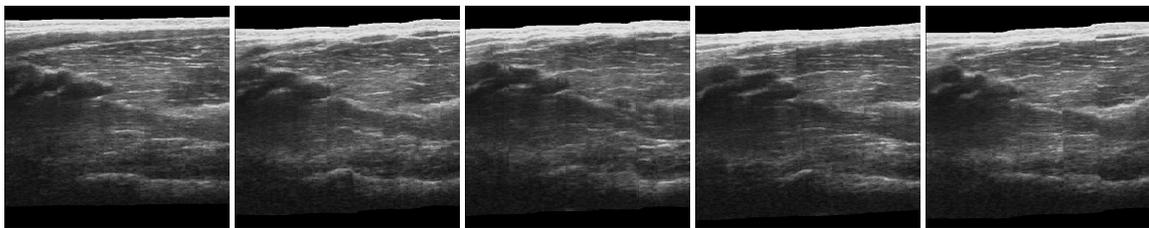
It is very hard to confirm the accuracy of such 3D ultrasound volume data. Comparing with other modalities, e.g. Computed Tomography or Magnetic Resonance Imaging is not straightforward — ultrasound does not measure a material property, let alone one which is compatible with other imaging modalities. Registering such data with ultrasound is a difficult task, and itself introduces errors which make this far from a ‘gold standard’ comparison. The following experiments are designed to assess repeatability by comparing multiple sets of 3D ultrasound data. It is conjectured that if the non-rigid correction is physically appropriate, then the sets of corrected 3D data will be more similar to each other than the sets of uncorrected data.

2.2.1 Repeated scans in the same direction

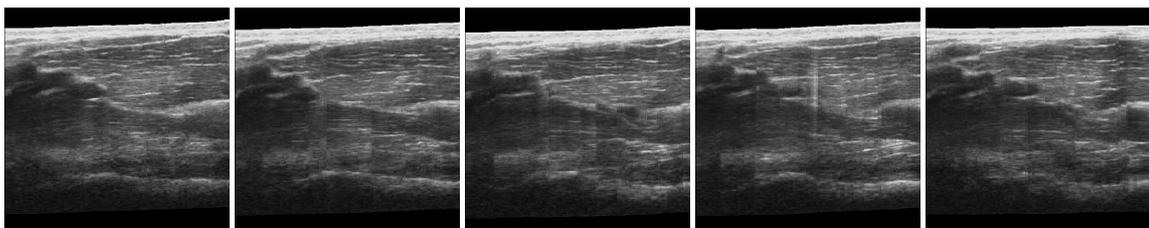
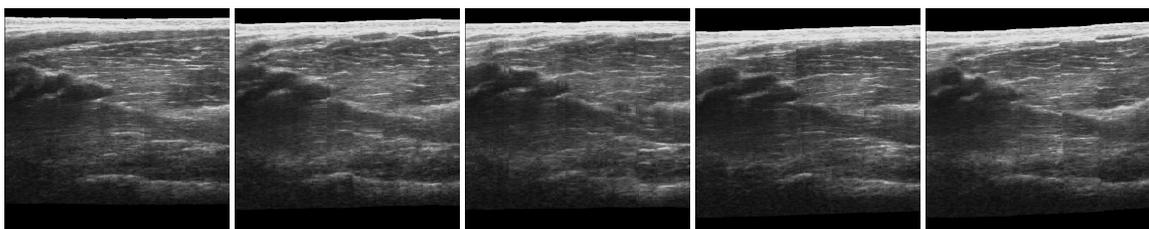
In the first experiment, part of a forearm was scanned ten times using the same linear sweep pattern in each case. This pattern is shown in Figure 3(a). Both the arm and the reference for the position sensor were kept as stationary as possible throughout the entire process. Figures 4 and 5 show reslices through each of these data sets (perpendicular and parallel to the skin surface, respectively). The relative location of these reslices is also shown in Figure 3(b). In each case, reslices are shown of the original data, and of that with rigid and



(a) original

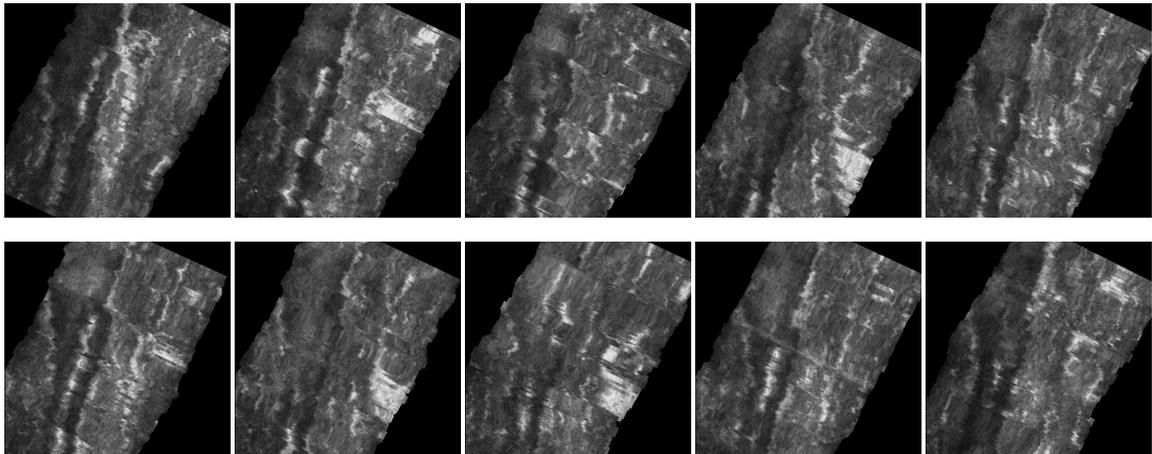


(b) rigid correction

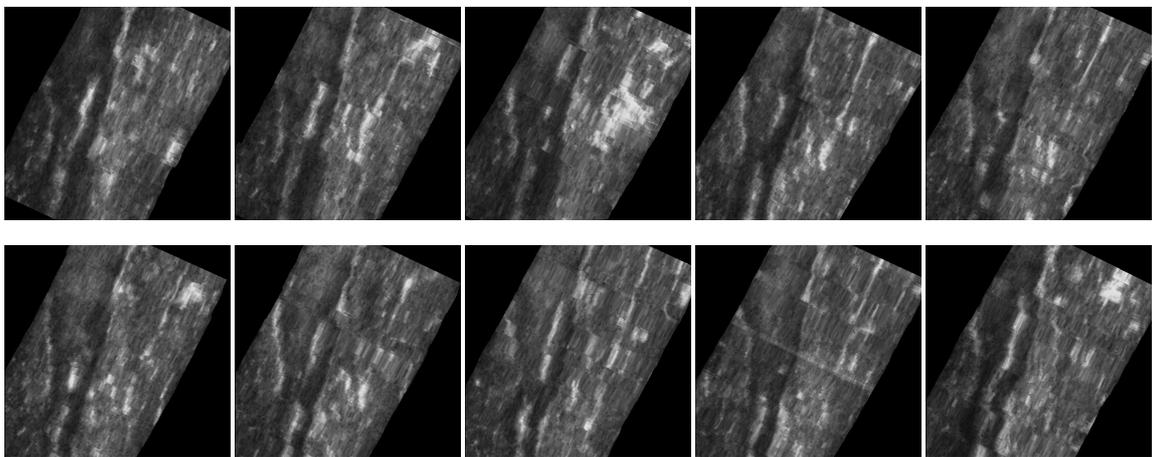


(c) non-rigid correction

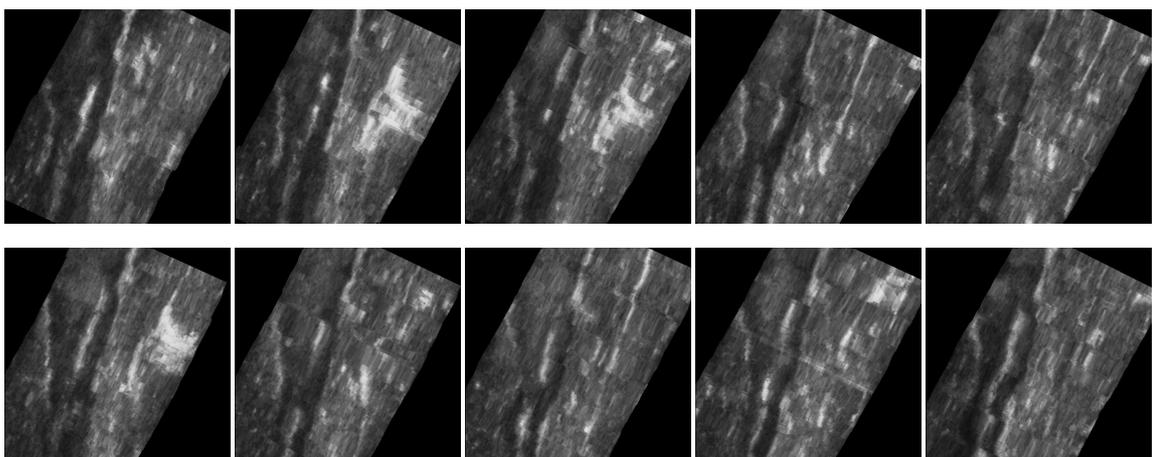
Figure 4: Reslices perpendicular to the skin surface as in Figure 3.



(a) original

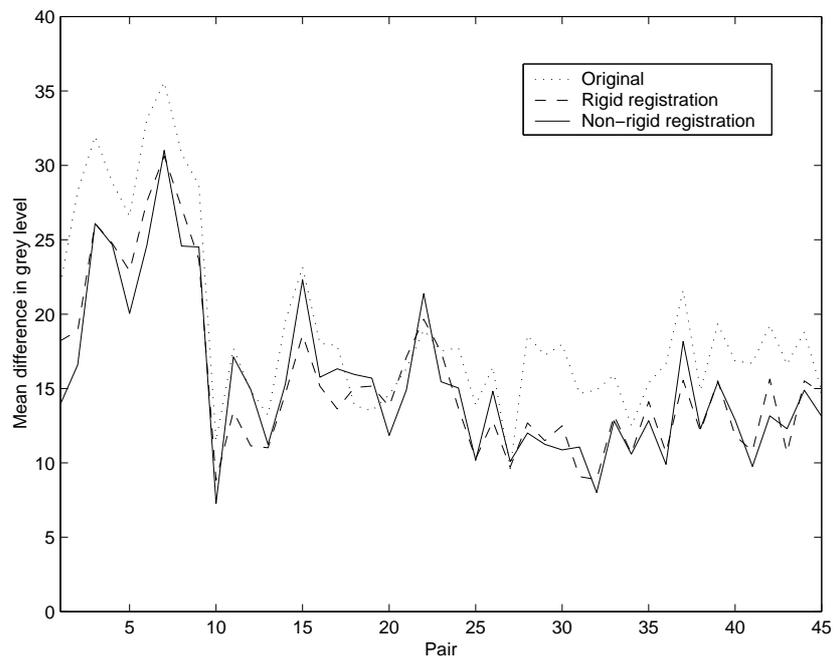
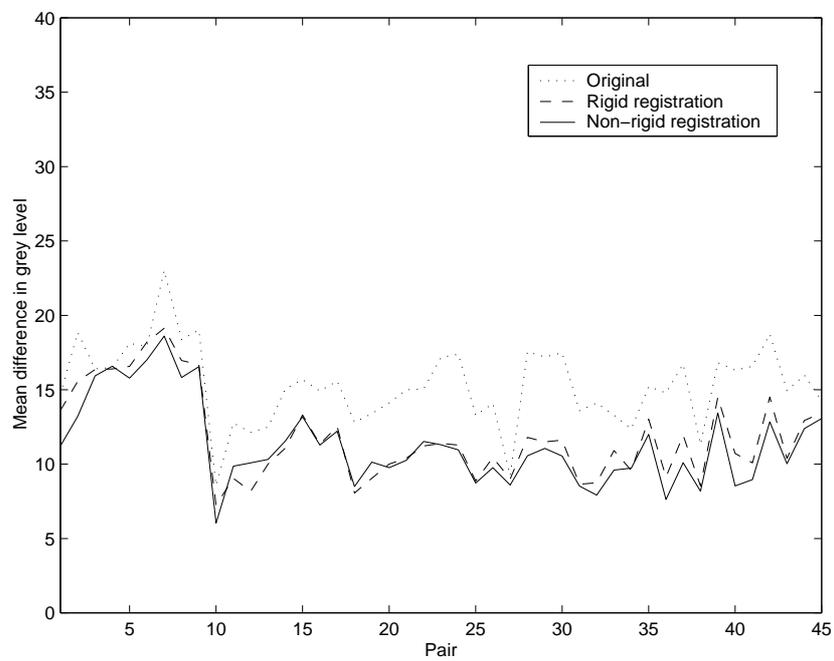


(b) rigid correction



(c) non-rigid correction

Figure 5: Reslices parallel to the skin surface as in Figure 3.

Figure 6: v_p for all pairs of data sets in Figures 4 and 5.Figure 7: Minimum v_p for all pairs of data sets in Figures 4 and 5.

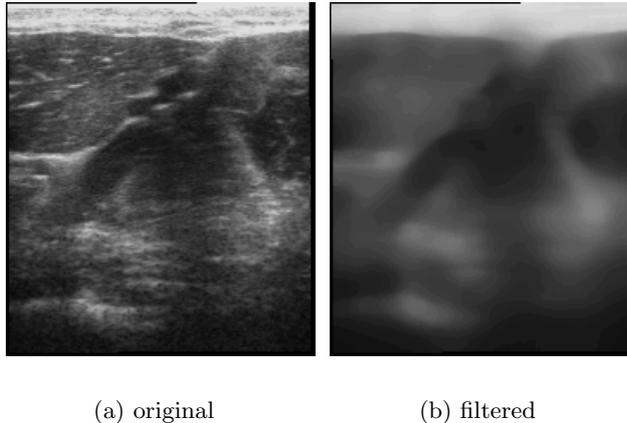


Figure 8: Filtering of 3D data sets. (a) shows a slice through one of the unfiltered data sets, and (b) shows the same slice after applying a 3D median filter.

non-rigid correction. The sets of reslices with rigid and with non-rigid correction are clearly smoother than their original uncorrected counterparts.

In order to investigate the variation in the original and corrected data, each of the ten data sets was resampled onto a cubic 3D array, using pixel nearest neighbour interpolation [4]. The spatial location and extent of this array was fixed over all the data sets, to enable subsequent comparison. Since it is not reasonable to expect the speckle pattern to be exactly the same over multiple scans, even if the location and orientation of the probe only differs by a small amount, the data set was then filtered to remove speckle using a 3D median (edge-preserving) filter. An example of the effect of this filter is shown in Figure 8. Having filtered each data set, it was then downsampled to an array of approximately $128 \times 128 \times 128$ voxels, in order to speed up subsequent processing.

A measure of data variation can be obtained by finding the correlation between the grey-levels of all possible pairs of the ten 3D data sets. For each pair of data sets p , the variation v_p is given by:

$$v_p = \sqrt{\frac{\sum_{x,y,z} (g_a(x,y,z) - g_b(x,y,z))^2}{\sum_{x,y,z} 1}} \quad (1)$$

where a and b are the data sets in pair p , and $g_a(x,y,z)$ and $g_b(x,y,z)$ are the grey levels at voxel (x,y,z) from each of a and b respectively. v_p can therefore be interpreted as the root mean square difference in grey level between the two data sets in pair p .

The mean and standard deviation of v_p was calculated across all pairs of the original data sets, and those with rigid and non-rigid correction. In addition, in order to allow for small movements in the arm between each of the ten sweeps, the local minimum v_p was calculated for each pair by allowing a relative movement of up to 10 voxels (approximately 2mm) in each direction, prior to the calculation. The mean and standard deviation was also calculated of the minimum v_p for all pairs. These results are contained in Table 1; black regions in the data sets containing no ultrasound data were not allowed to contribute to the correlation.

It can readily be seen from Table 1 that the similarity between the data sets was improved in all cases by rigid and non-rigid correction of the data. This improvement is masked by

		Same direction		Varying directions	
		v_p	Minimum v_p	v_p	Minimum v_p
Original	mean	19.04	15.31	22.04	19.36
	standard deviation	6.08	2.66	4.97	4.53
Rigid	mean	15.48	11.86	22.01	18.79
	standard deviation	5.42	2.98	5.01	4.36
	$p > \text{original}$	5	1	14	10
Non-rigid	mean	15.40	11.33	21.36	18.50
	standard deviation	5.28	2.82	4.75	4.21
	$p > \text{original}$	5	1	13	8

Table 1: Mean and standard deviation of v_p over all pairs of data sets in Section 2.2.1 (Figures 4 and 5) and Section 2.2.2 (Figures 10 and 11). ‘ $p > \text{original}$ ’ shows the number of pairs which were *less* similar after correction, out of 45 pairs (scans in the same direction) and 28 pairs (scans in varying directions) in total.

the inherently unrepeatable nature of ultrasound (the data is dependent on the direction of scanning as well as the material properties), and small movements of the arm between sweeps, however it is still significant. Note that the values for minimum v_p (which goes some way to account for movement of the anatomy) are much better. It is already clear that reslices through the corrected data are smoother, which makes them much easier to interpret, and this indicates that they are also more, rather than less, accurate.

Figures 6 and 7 show graphical results for v_p and the minimum v_p (i.e. taking into account possible movement of the arm) for all pairs. In Figure 7 there is only one case where non-rigid correction decreases the similarity.

2.2.2 Repeated scans in varying directions

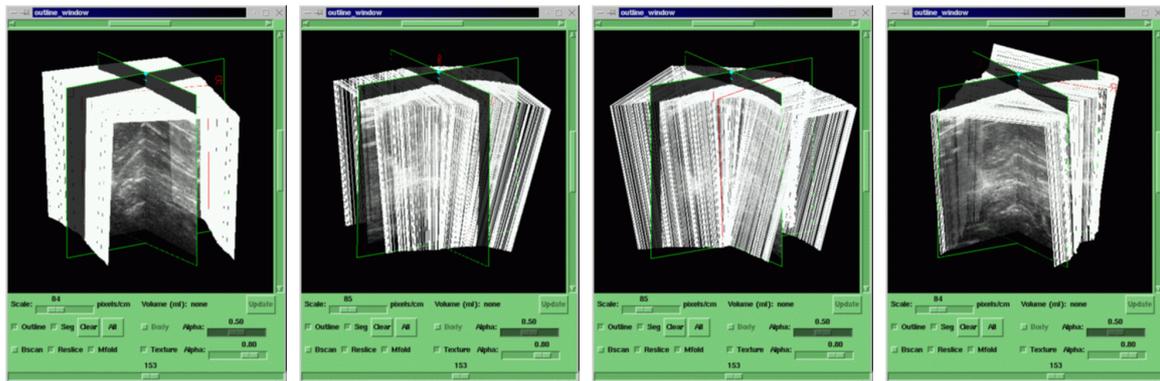
The second experiment was similar to the first, except that in this case the repeated sweeps were from different directions. Such data sets will be inherently less similar, even without probe pressure artifacts, due to the irregularity of the ultrasound resolution cell and the dependence of the ultrasound signal on the incidence angle. However, correction for pressure should still generate more similar data.

Four sets of two sweeps were acquired, at approximately 45° angles. As before, these were resampled to a cubic 3D array, filtered and downsampled to approximately $128 \times 128 \times 128$ voxels. The arrangement of the sweeps and the two reslices shown in Figures 10 and 11 can be seen from Figure 9. The reslices show the original and corrected data.

Once again, v_p and the minimum value of v_p were calculated for all pairs of scans (in this case 28). These results are also contained in Table 1, and the graphical results for v_p and the minimum v_p in Figures 12 and 13.

Although the corrected reslices through the data sets are clearly smoother and also appear to be more similar to each other, the values of v_p and the minimum v_p show little improvement. The graphs in Figures 12 and 13 indicate that if anything, the similarity is improved after correction, but unlike the previous experiment this improvement is not significant⁵.

⁵It is clear, however, that the corrected data sets are certainly not *less* similar than the uncorrected data sets.



(a) scans 1 and 2

(b) scans 3 and 4

(c) scans 5 and 6

(d) scans 7 and 8

Figure 9: 3D ultrasound examination of part of the forearm. The outlines of each B-scan, and the reslices used as examples of the data sets are shown in the same format as Figure 3.

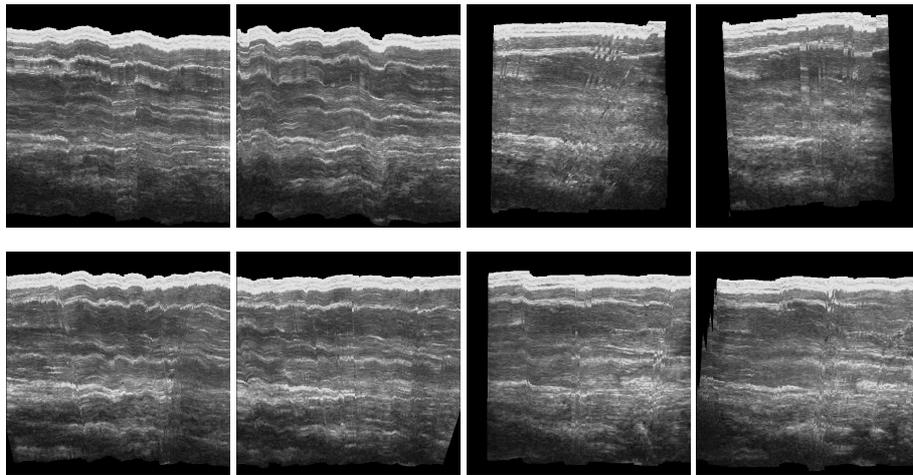
The reslices in Figure 11 highlight the probable cause of this inconclusive result. The curvature of the skin surface in each image varies considerably. This is due to a limitation in the ability of the algorithm to correct for all probe pressure effects. In reality, the skin surface *was* curved, but only in one direction. This curvature could be followed by the linear probe surface when scanning in one direction, but was flattened out when scanning in the other direction. The resulting tissue compression varied in the x (horizontal) as well as y (depth) direction, however we only attempt to correct for the y component of the compression. Differences in the location of the skin surface have a large effect on the value of v_p due to the higher signal (and hence whiter data) at these points.

3 Extension to convex probes

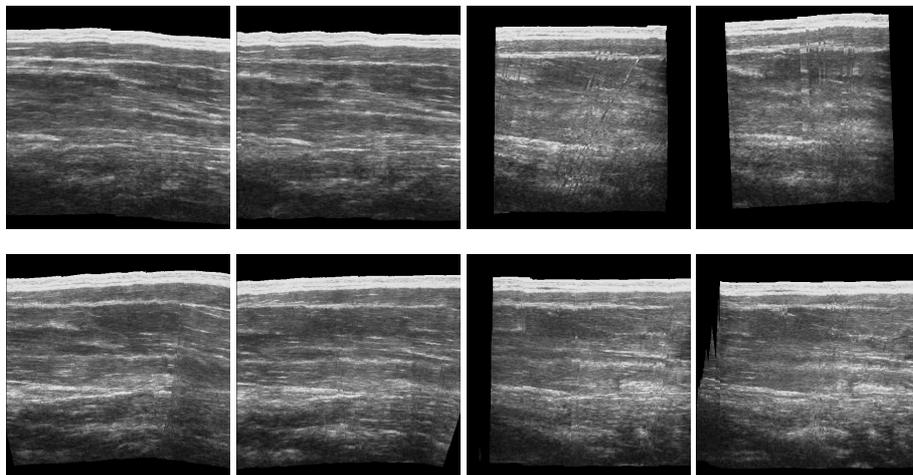
One of the assumptions in the probe pressure correction algorithm outlined in [5] is that probe pressure generates deformations which vary in the y direction only. We have already seen that this is not always the case, even with linear probes. The assumption becomes even more limiting when correcting data acquired with convex probes.

Consider for instance the B-scan of Figure 14(a). This is part of a sequence of scans of the abdomen, acquired whilst holding the probe in the same place, but applying a gradually increased pressure to the probe against the skin surface. Figure 14(b) shows another of the original B-scans with greater probe pressure, and Figure 14(c) is the result of correcting the whole data set using the linear deformation assumption.

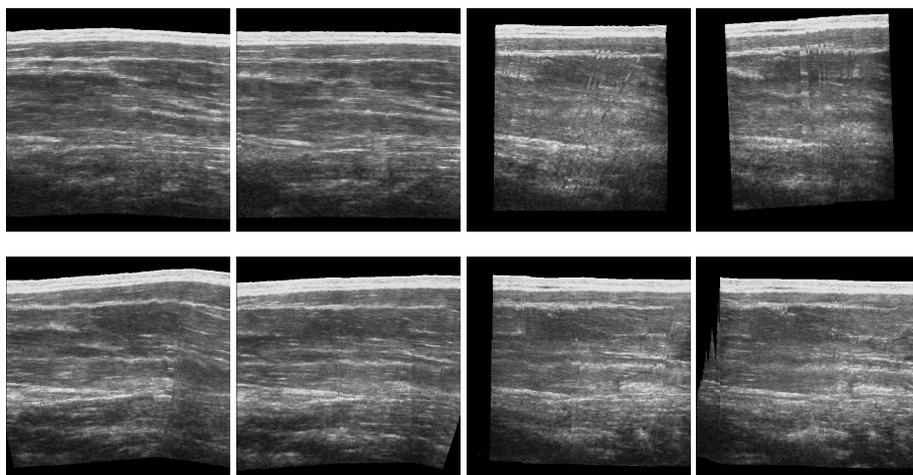
While the features at the centerline of the B-scan have been corrected to approximately the right locations, the horizontal surfaces have been deformed even further in the process. The deformations at the upper corners of the scan are particularly extreme, since it is at these points that the assumption of constant deformation in the x direction are most violated.



(a) original

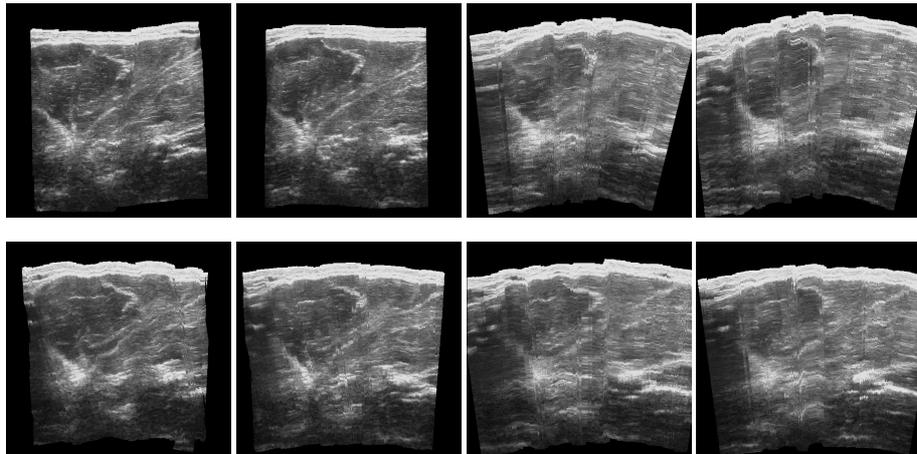


(b) rigid correction

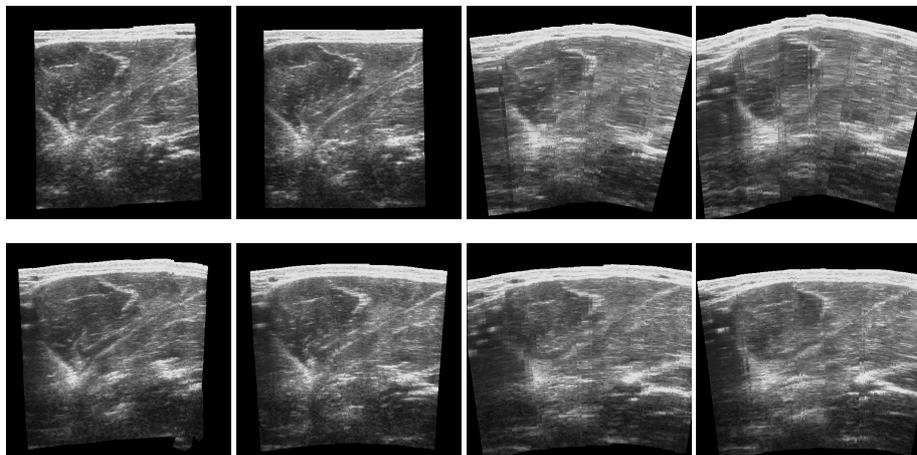


(c) non-rigid correction

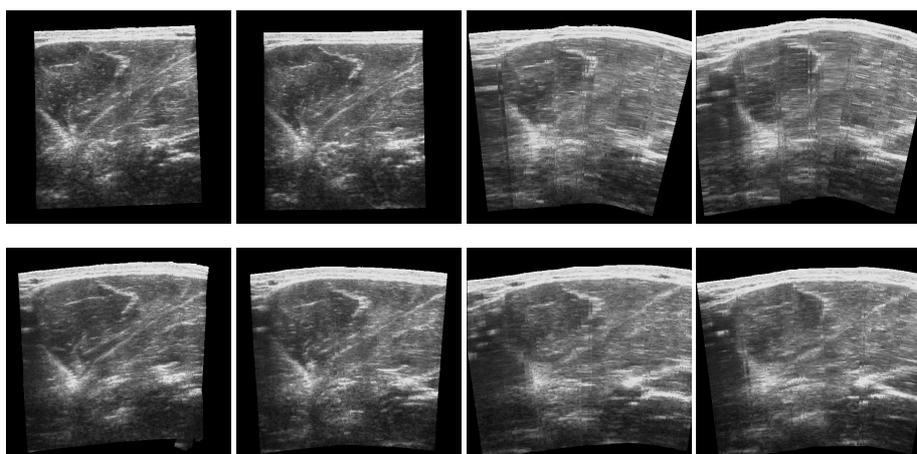
Figure 10: First reslice as in Figure 9.



(a) original

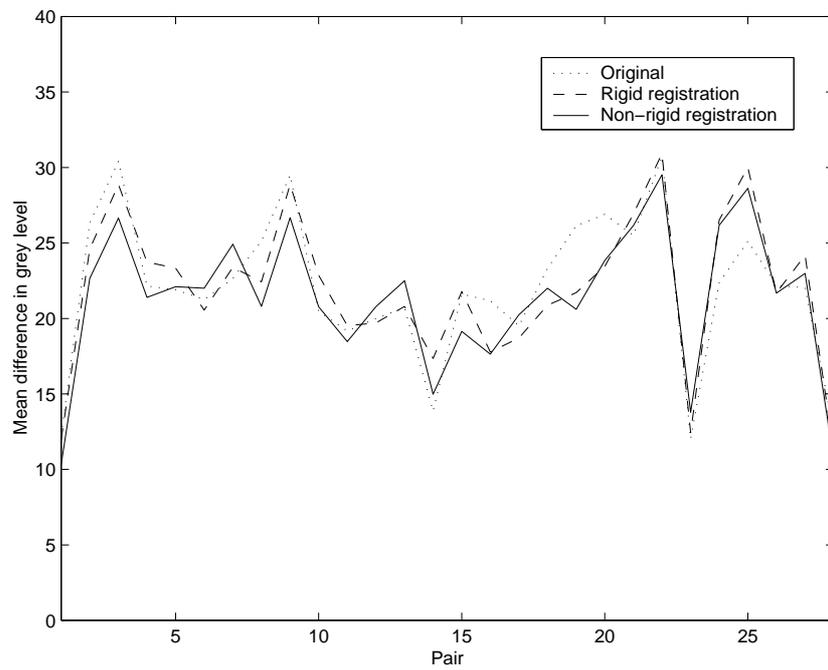
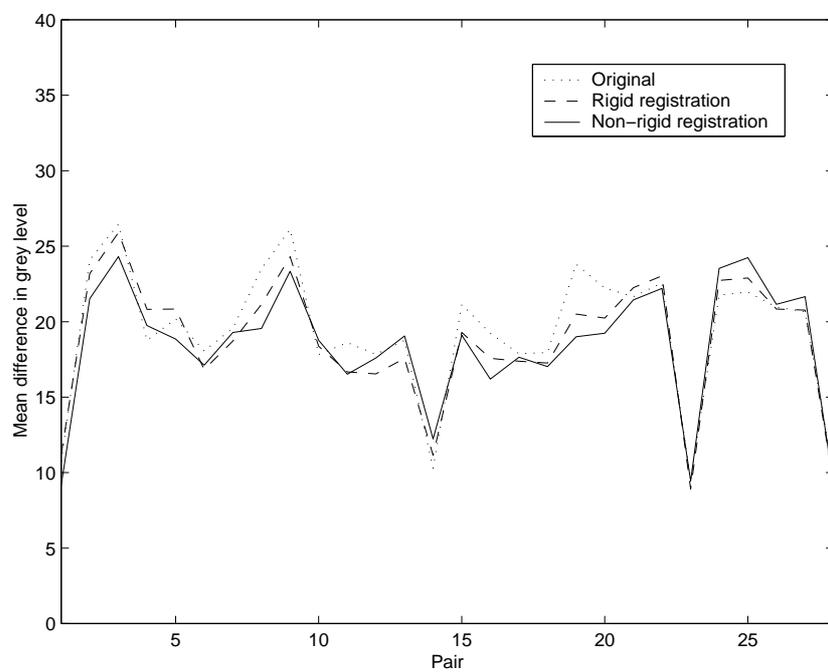


(b) rigid correction



(c) non-rigid correction

Figure 11: Second reslice as in Figure 9.

Figure 12: v_p for all pairs of data sets in Figures 10 and 11.Figure 13: Minimum v_p for all pairs of data sets in Figures 10 and 11.



(a) B-scan

(b) increased pressure

(c) non-rigid correction

Figure 14: Linear pressure correction for a convex probe. (a) and (b) are two B-scans from a stationary sequence acquired with varying pressure. (c) shows (b) after non-rigid correction of the entire sequence.

3.1 Incorporating a model for convex probe deformation

Non-rigid deformation is modelled by a vector $P(y)$ giving the shift in depth of each row at each value of y . If a model $S(x)$ can be assumed for the shape of the deformation in x at a given depth y and depth shift $P(y)$ for a particular probe, then this can be incorporated into the calculation without affecting the number of parameters to be estimated by the algorithm. Rather than applying the same $P(y)$ across all values of x , $S(x)$ is used, together with $P(y)$ and y , to provide a new estimate of pressure shift $P_{\text{convex}}(x, y)$ which varies in both x and y , and it is this which controls the deformation of the image pixels.

The actual deformation caused by a probe pressed against the skin surface is dependent on the geometry and properties of the part of the body under the probe, as well as the shape of the probe itself. The deformation is also not linear with increasing pressure — for instance a vein near the skin surface will deform until it is completely flat and will then stay relatively rigid. Hence $S(x)$ and the use of this to derive $P_{\text{convex}}(x, y)$ will necessarily be approximate. The shape used in the following examples was derived from a Gaussian distribution:

$$S(x) = 1 - e^{-\frac{1}{2\sigma^2}(x-x_c)^2} \quad (2)$$

$$P_{\text{convex}}(x, y) = P(y) \left[1 - \frac{d-y}{d} S(x) \right] \quad (3)$$

where σ and x_c in equation (2) and d in equation (3) are constants derived from the shape of the ultrasound probe. The application of $P_{\text{convex}}(x, y)$ for a typical convex probe shape and a fairly extreme deformation represented by $P(y)$ is shown in Figure 15. This function has the useful properties that it tends smoothly to zero at large values of x , and tends towards the linear probe case for large values of y (depth).

3.2 Calculating pressure correction with convex probe deformation

The key component in the calculation of the non-rigid pressure shift $P(y)$ for linear deformation is correlation of each horizontal line in one B-scan with a range of horizontal lines in

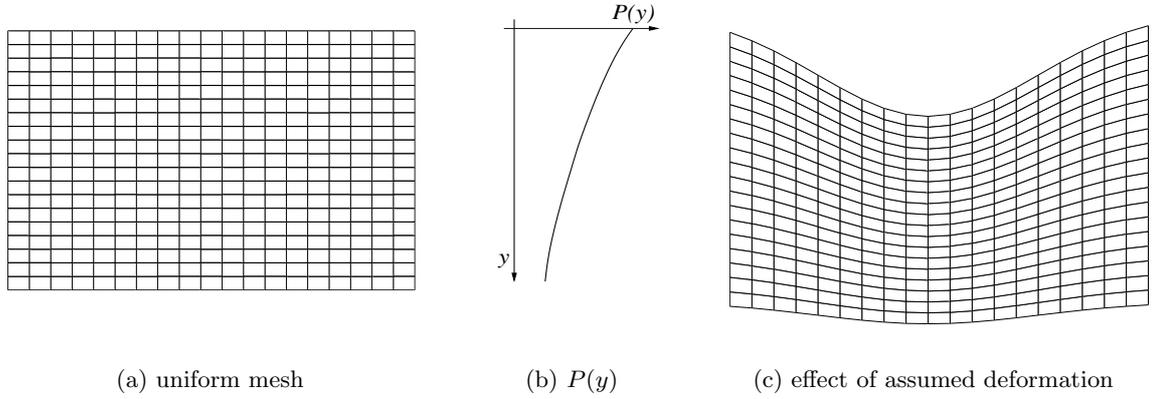


Figure 15: $P_{\text{convex}}(x, y)$ for a given probe shape and non-rigid deformation $P(y)$. (c) shows the effect of the deformation (b) on the mesh in (a).

the neighbouring B-scan. The relative y location of each line within this range directly gives the potential value of $P(y)$. For convex probes, rather than comparing the line in one B-scan with a range of lines, it is compared with a range of curves. The shape of the curve is given by equation (3), where in this case y is the depth of the line in the B-scan, and $P(y)$ is the y offset of the mid-point of the curve from this line.

An example of this search is given in Figure 16(a). Other than this first step, the remainder of the algorithm described in [5] remains the same.

There is one additional complication, which is introduced by the initial *rigid* correction which gives the starting point for the non-rigid correction search. The rigid correction is important in order to improve the non-rigid correlation values (due to optimising the alignment in x) and decrease the search range over y . However, the y offset which is introduced affects the set of curves over which the non-rigid search is made. Figure 16(a) shows an example of a correct non-rigid search for a pair of B-scans which have *not* already undergone rigid correction. Figure 16(b) shows non-rigid correction naïvely applied after rigid correction: the line is matched with a different shape of curve, and this will change the subsequent value of $P(y)$ slightly. Figure 16(c) shows the correct solution. The search range is the same, but equation (3) used to calculate the shape of the curves becomes:

$$P_{\text{convex}}(x, y) = P(y) - (P(y) + y_{\text{rigid}}) \frac{d - y}{d} S(x) \quad (4)$$

where y_{rigid} is the shift due to the rigid correction, and $P(y)$ and y are set the same as previously.

Figure 17 shows the result of applying this adjusted algorithm to the same sequence of B-scans as in Figure 14. Note that the non-rigid correction in Figure 17(c) has now positioned the main features in the correct locations whilst also correctly re-adjusting the shape of the top of the B-scan. More information is visible than in Figure 17(a), since the probe was in better contact with the skin surface: a direct result of increased contact pressure.

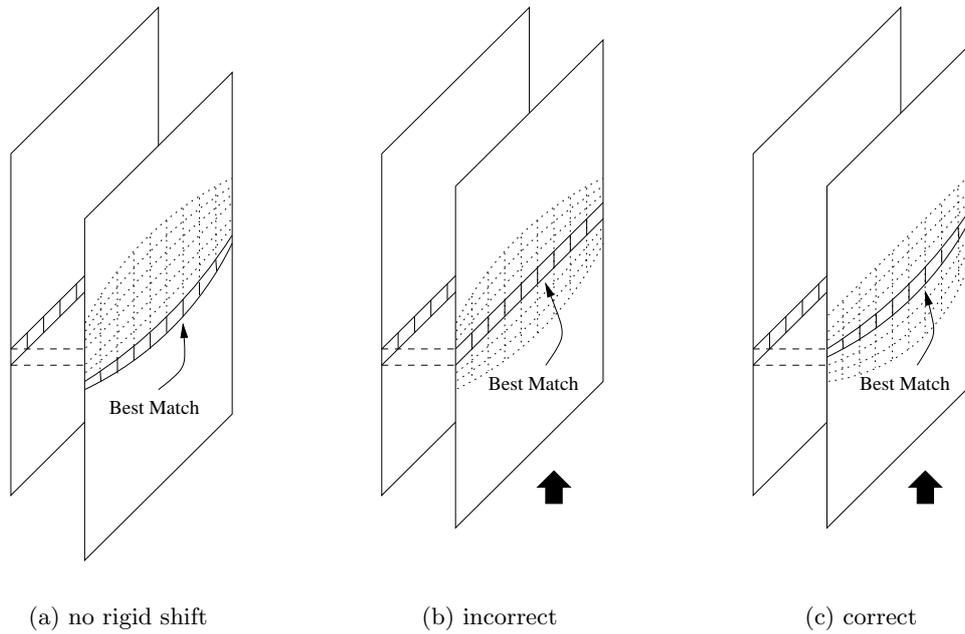


Figure 16: Range and curve shape used to estimate non-rigid convex deformations. (a) Curve shapes without a previous rigid correction. (b) The same curve shapes incorrectly used after a rigid correction has been applied. (c) The correct set of curves for (b).

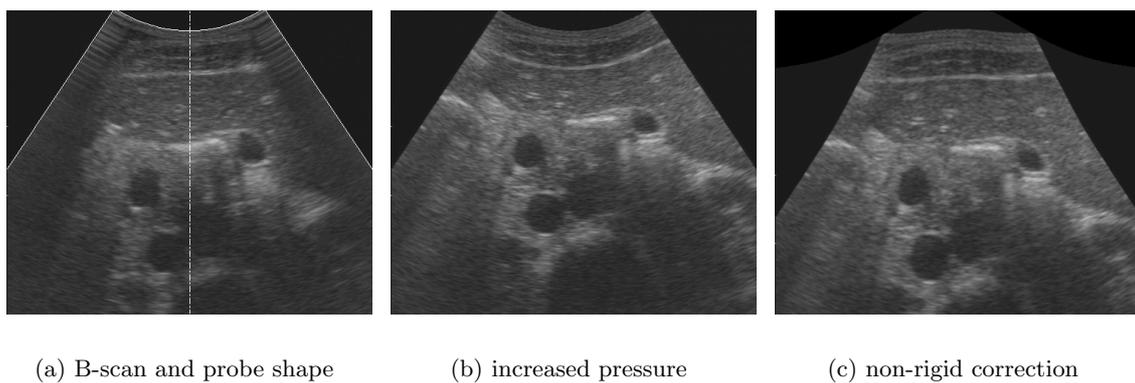


Figure 17: Convex pressure correction for a convex probe. (a) and (b) are two B-scans from a stationary sequence acquired with varying pressure. (c) shows (b) after non-rigid correction of the entire sequence.

3.3 Automatically detecting the type of probe

Given a particular ultrasound probe, equations (2) and (3) could be replaced by ones which matched the probe more exactly, across the expected range of anatomy to be scanned. However, our aim is to implement a more general solution which, although not ideal for a given probe, will nevertheless improve the non-rigid correction results over those with linear probe assumptions. Equations (2) and (3) are reasonable approximations, which leave sufficient parameters σ , x_c and d to tailor the pressure function according to the shape of the convex probe in use.

The freehand 3D ultrasound system we use is designed to run on a separate computer from the ultrasound machine, and the only link to the ultrasound machine is via a stream of images of the B-scans, either from an analogue video signal or digital data sent over the ethernet [3]. It is of considerable advantage, therefore, if the shape of the probe can be deduced from this stream of images, and the deformation parameters estimated from the shape of the probe.

A typical full-frame image from a convex probe is shown in Figure 18(a). Here the probe is in free air, but the gain is fairly high such that some noise is visible within the area of coverage of the B-scan itself. It is possible to exclude most of the background, and some of the text, by selecting all pixels in the image between a low and a high threshold, as in Figure 18(b). These thresholds are set interactively.

Morphological filtering [2] is then performed on the segmented image, by opening and then closing the image, with a square template a few pixels across. Opening the image enlarges then reduces the black regions, thus removing small white regions. Closing the image has the same effect on small black regions. Performing them in this order ensures that small segmented regions close to the B-scan tend to be removed, rather than subsumed into the B-scan data. Importantly, this type of filtering preserves both the location and the sharpness of strong edges in the original image. Figure 18(c) shows the result of applying this process.

Connected component analysis is now performed on the filtered, segmented image, by a grass-fire technique [1](page 458). The object with the largest number of connected pixels is retained and others are removed. This leaves a binary image of the B-scan as in Figure 18(d).

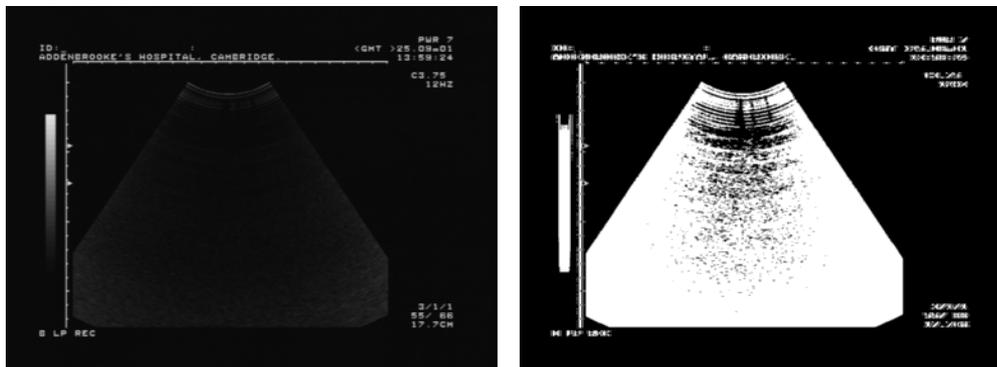
The centre of the B-scan and the slope of the edges are simultaneously extracted from this image by examining the locations of the left-most $x_l(y)$ and right-most $x_r(y)$ pixels on each horizontal line. For each pair of lines $y = a$ and $y = b$, the estimated centre x_c and slope s is given by:

$$x_c = \frac{1}{2} [x_r(a) + x_r(b) - x_l(a) - x_l(b)] \quad (5)$$

$$s = \frac{1}{2(a-b)} [x_l(a) - x_l(b) + x_r(b) - x_r(a)] \quad (6)$$

This estimate of the (symmetrical) edge location is compared against $x_l(y)$ and $x_r(y)$ for all values of y . The number of lines for which the error is less than one pixel, and the cumulative absolute error for these lines, is recorded. The process continues for all possible pairs a and b , and the final values for x_c and s kept for which most lines are in agreement. If there are several pairs with the same consensus, the pair with the lowest error is retained.

During the calculation, a separate set of results is recorded for linear and convex probe shapes, by testing for slopes s which are larger than a predefined threshold. A decision as to whether the probe is convex or linear is made at the end by testing whether the maximum consensus set for the convex case is above a predefined minimum. This separation of results



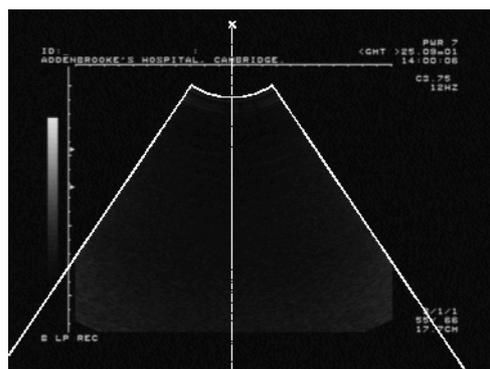
(a) input video frame

(b) after segmentation



(c) after morphological filtering

(d) after connected component analysis



(e) probe shape estimate

Figure 18: Automatic detection of probe shape. (a) to (e) are the stages in the estimation of probe shape.

is vital, since in general convex probes will contain regions which have vertical edges, as in the example of Figure 18, whereas linear probes will contain no non-vertical edges — we want to make the convex/linear decision on whether a non-vertical edge exists at all, *not* on which edge is the most dominant and hence has the largest consensus set.

Having found values for x_c and s , both the centre of curvature of the probe and the enclosed angle are therefore known. All that remains is to locate the top edge of the B-scan. For each x , the uppermost pixel is located, and the distance of this from the centre of curvature computed; the minimum distance gives the radius of the top of the B-scan, r . Figure 18(e) shows the calculated shape estimate in this case. σ and d used in equations (2) and (3) are calculated directly from s and r :

$$\sigma = \frac{3r}{5} \quad (7)$$

$$d = \frac{4r}{\sqrt{1+s^2}} \quad (8)$$

3.4 Examples of convex probe data correction

Figure 19 shows an example of convex probe pressure correction of two freehand 3D data sets. In both cases, one B-scan and a reslice through the entire data set are shown. These 2D images have been texture-mapped onto planes in 3D so that their relative locations can be seen. The corrected data sets in Figures 19(b) and (d) are improved in two ways. B-scans with a high contact pressure are ‘unsquashed’ such that anatomy near the surface is no longer deformed to the shape of the probe. In addition, the reslices show much less variation in the position of the anatomy, particularly that near to the skin.

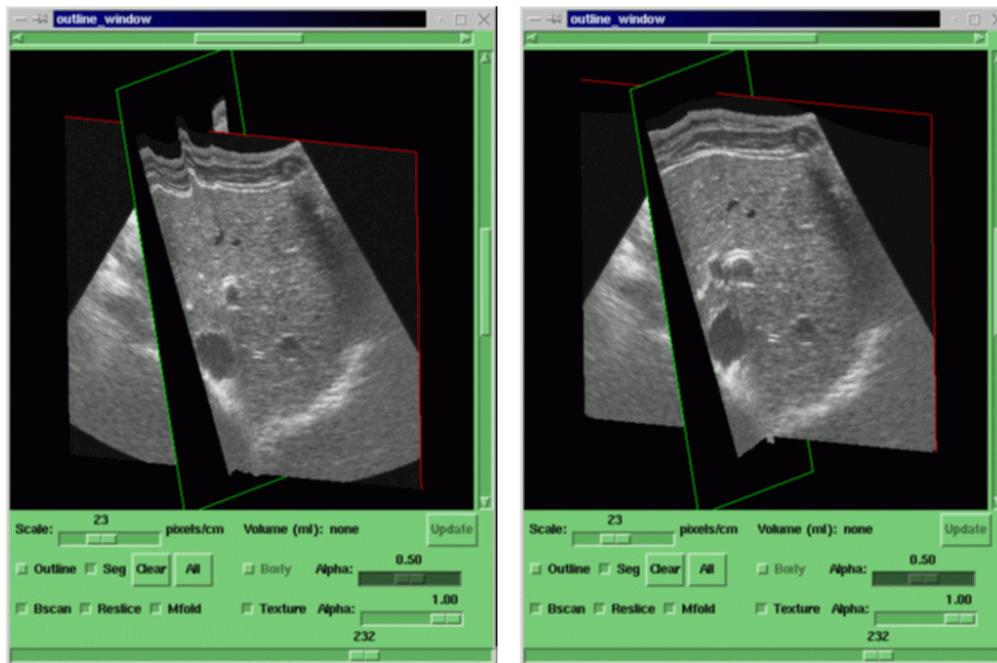
Note that both these results are achieved whilst maintaining the deeper anatomy at approximately the same location as in the initial data set. The location of each B-scan, on the other hand, is changed in order to ensure that the skin surface can still be seen at the top of the B-scan. Since the size of the B-scans is preserved, this inevitably means that some data is lost at the base of any B-scans which have been uncompressed.

4 Conclusions

A novel algorithm was presented in [5], combining image-based and position sensing techniques, to correct the most significant effects of probe pressure in freehand 3D ultrasound data. Here we have shown how this algorithm can be extended to apply to convex as well as linear probes. In addition, we have demonstrated that the algorithm increases the repeatability of multiple freehand 3D ultrasound examinations as well as increasing the clarity of reslices through such data. This increase in repeatability is however limited by the pressure-effect assumptions and the variability inherent in the ultrasound data itself.

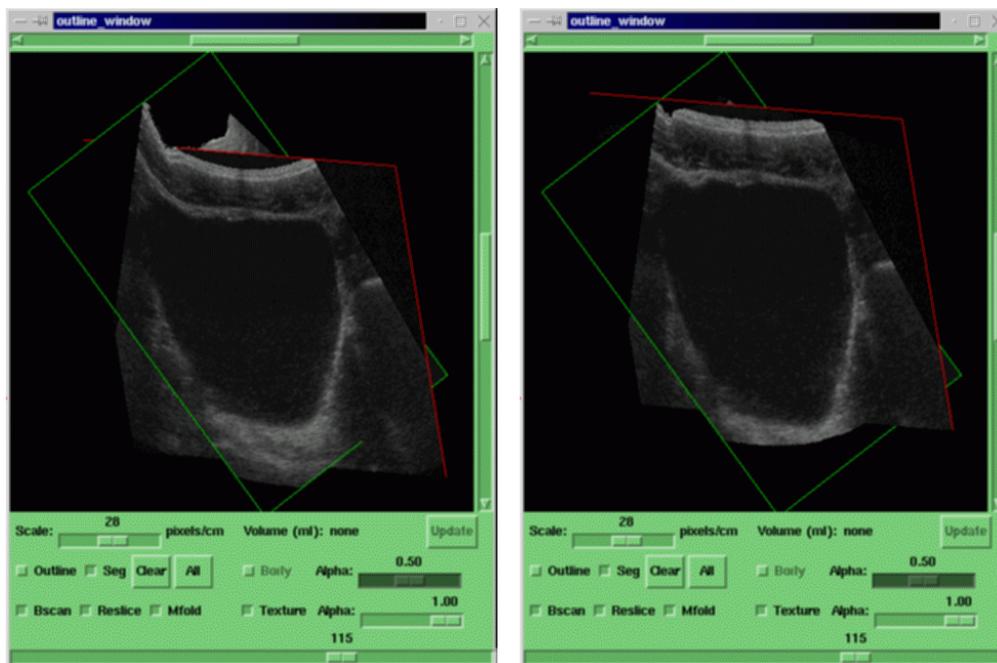
Acknowledgements

This work was carried out under EPSRC grant GR/N21062. Dynamic Imaging Ltd. provided a modified ultrasound machine to enable digital data acquisition.



(a) original liver

(b) corrected liver



(c) original bladder

(d) corrected bladder

Figure 19: Freehand 3D scans with and without convex probe pressure correction. In each case, a B-scan and a reslice through the whole sequence are shown in the correct 3D location.

References

- [1] R. Gonzales and R. Woods. *Digital Image Processing*. Addison-Wesley, 1993.
- [2] P. F. M. Nacken. Chamfer metrics, the medial axis and mathematical morphology. *Journal of Mathematical Imaging and Vision*, 6:235–248, 1996.
- [3] R. W. Prager, A. H. Gee, and L. Berman. Stradx: real-time acquisition and visualisation of freehand 3D ultrasound. *Medical Image Analysis*, 3(2):129–140, 1999.
- [4] R. N. Rohling, A. H. Gee, and L. Berman. A comparison of freehand three-dimensional ultrasound reconstruction techniques. *Medical Image Analysis*, 3(4):339–359, 1999.
- [5] G. M. Treece, R. W. Prager, A. H. Gee, and L. Berman. Correction of probe pressure artifacts in freehand 3D ultrasound — initial results. Technical Report CUED/F-INFENG/TR 411, Cambridge University Engineering Dept, July 2001.