RF and amplitude-based probe pressure correction for 3D ultrasound

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

CUED/F-INFENG/TR 491

August 2004

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

Corresponding e-mail: gmt11@eng.cam.ac.uk

Abstract

Anatomical deformation caused by variable probe contact pressure is a significant problem in freehand 3D ultrasound, particularly for high resolution musculoskeletal and breast scans. We have previously published an amplitude-based algorithm for correcting such errors. In this paper, we compare this approach with a novel, elastography-inspired algorithm which works with the higher resolution radio-frequency (RF) signal. The results show that, although the RF-based algorithm is more precise in certain circumstances, both algorithms are able to compensate for probe pressure in 3D ultrasound data. Consequently, freehand 3D ultrasound users who do not have access to the RF signal are still in a position to perform effective probe pressure correction using the readily available video output, as long as this signal is not affected by significant amounts of frame averaging (persistence).

Contents

1 Introduction 1

2 System 2
2.1 Estimation of compression due to probe pressure 2
2.2 Combination of image-based location with position sensor readings 4

3 Results 5
3.1 Static scan of an ultrasound phantom 8
3.2 3D scan in a water bath 9
3.3 3D scans in multiple directions 11
3.4 Panoramic scan 13

4 Conclusions 15

5 Acknowledgements 16

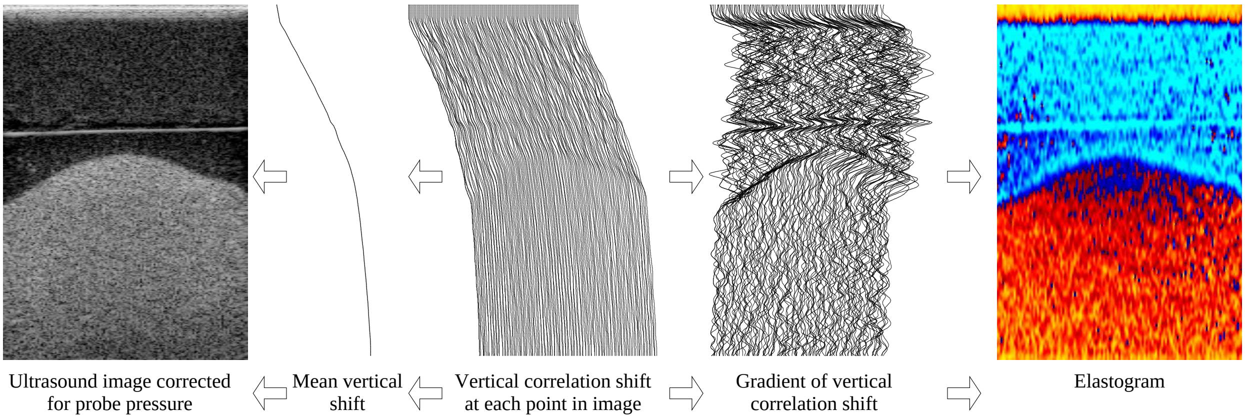

Figure 1: The relative vertical shift to give maximum correlation between small sections of RF data from two neighbouring ultrasound images (centre) can be used as the starting point both for probe pressure correction (left) and elastography (right).

1 Introduction

Routine ultrasound examination relies on contact pressure between the probe and the subject in order to provide a good acoustic match. Even with ultrasound gel, the skin is often substantially depressed: this is often necessary to achieve the desired view. Clearly this deformation affects the anatomy being scanned and hence also the ultrasound images. Probe contact pressure effects are particularly apparent in high resolution (> 10MHz) ultrasound data where the image is relatively shallow and the anatomical deformation can easily exceed 10% of its depth. This is exacerbated in areas which are relatively soft, for instance the breast (Xiao et al., 2002) or prostate (Artignan et al., 2004), and deformations of up to 1cm are fairly typical during such examinations. This leads to incorrect estimates of both the depth and the size of small tumours, which can be particularly problematic for image guidance techniques.

Typical high resolution ultrasound images have a limited field of view: in order to visualise a larger volume, they must be combined into a composite data set. The most appropriate technique for this is freehand 3D ultrasound, where the probe is moved by hand, and the resulting sequence of images is registered by either intrinsic (image-based) or extrinsic (position sensing) means. Current position sensing techniques alone are not able to accurately register high resolution ultrasound data, since sensor errors are at a significantly larger scale than the pixel size (< 0.1mm) in a high resolution image. Image-based registration, by contrast, has been used successfully to generate extended-field-of-view, or *panoramic* images (Weng et al., 1997). 3D data sets can be constructed by combining image-based registration with speckle decorrelation (Smith and Fenster, 2000; Tuthill et al., 1998), the latter providing an estimate of the out-of-plane probe movement. These techniques can achieve accurate local registration, are fairly robust to noise in the images, and require no user interaction.

However, the effect of varying probe contact pressure on such data sets is to generate distortions in 3D visualisations of the data. In a previous paper (Treece et al., 2002), we have demonstrated that it is possible to use image-based registration in combination with a position sensor to reduce the effects of probe pressure on freehand 3D ultrasound data. Unlike the method of Xiao et al. (2002) which requires multiple sweeps of the anatomy, this technique can be used to reduce probe pressure in a single 3D data set.

It has been noted that probe pressure correction, which has had very little attention in the literature, is closely related to ultrasound elastography, which is becoming increasingly popular (Greenleaf et al., 2003; Pesavento et al., 2000). In ultrasound elastography, an image is displayed of the visco-elastic properties of the scanned anatomy. The similarity is demonstrated in Fig. 1. In one particular technique (Pesavento et al., 1999), ultrasound images are acquired with the probe held over the same anatomy, but with very slightly varying contact pressure. The raw radio frequency (RF) data, from which these images are derived, is analysed to give a very accurate estimate of the relative vertical shift between the data at every point in these two images. An image of the gradient of this data, called an *elastogram*, is used to highlight anatomy of varying stiffness. However, the *average* relative vertical shift at depth gives an estimate of how much one image needs to be decompressed in order to match the other — and this is the raw information on which probe pressure correction is based.

This leads us to an obvious question: can adapted RF-based techniques from elastography be used to correct for probe pressure more accurately than amplitude-based (i.e. derived from an ultrasound image) techniques? The answer to this question is not obvious, due to a subtle but important difference between the two techniques: in elastography the probe is held still, over the same anatomy, while the pressure is varied. In freehand 3D ultrasound, the probe is moved during the scan in addition to the contact pressure varying. This means that changes in the ultrasound images cannot all be associated with pressure: some of them are due to differences in the anatomy being scanned.

In re-examining probe pressure correction, we also take the opportunity to present a more rigorous way of combining image-based positions with those from a position sensor. This step is necessary to maintain the long-range reliability from the position sensor whilst correcting the short-range registration and probe pressure errors. The original algorithm (Treece et al., 2002) contained an heuristic solution to this problem. In addition, it will be seen that the original results were adversely affected by small levels of ultrasound image persistence: this is analysed and persistence-free results included in this paper for the original algorithm as well as the novel RF-based algorithm.

2 System

2.1 Estimation of compression due to probe pressure

Amplitude-based probe pressure correction for a single freehand 3D ultrasound data set is described in detail in Treece et al. (2002). Ultrasound images are acquired at 20-30Hz, along with their location and orientation in space. Since the acquisition is relatively fast, only small changes in location, anatomy, and contact pressure are expected between neighbouring images. Probe pressure correction is hence based on accumulating sequential estimates between each pair of frames. The inevitable drift which this sequential accumulation induces on the location of each frame can itself be corrected by the re-introduction of the position sensor locations in the long range.

For each pair of frames, a rigid in-plane shift is first estimated which maximises the correlation between the two ultrasound images. Following this, each horizontal line in one image is correlated to groups of neighbouring lines in the other image, and this information is used to derive a sub-pixel estimate of the vertical shift which generates the best match. Estimates from neighbouring lines are averaged to a derive a smaller number of less noisy values. The size of the group from which each average is gathered is determined by the variance of the initial estimate: smaller groups are used where the values are tightly clustered, and larger groups where they are more widely distributed. Monotonicity is then imposed on these smoother estimates, in whichever sense best matches the data. Finally, they are linearly interpolated, to give the relative vertical shift at any depth due to the change in probe contact pressure between the two images. This is demonstrated diagrammatically

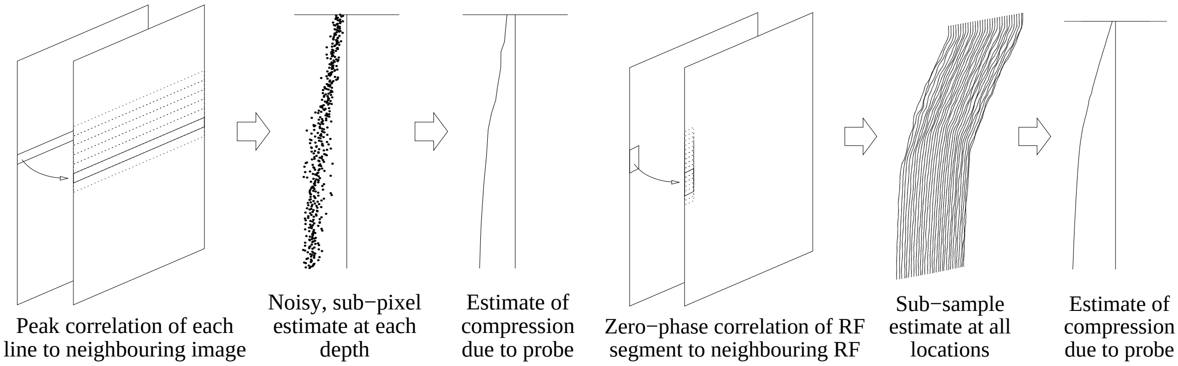

(a) Amplitude-based estimate (b) RF-based estimate

Figure 2: Compression due to probe pressure can be estimated both from (a) the amplitude of the ultrasound data, and (b) the raw RF signals.

in Fig. 2(a).

Clearly, the actual result of variations in probe contact pressure on real anatomy is much more complex than a simple vertical compression in the image, even one which varies in depth. However, the adoption of this simple model is vital in order to decouple the variations in contact pressure from the natural variations in the two images due to changes in anatomy: recall that, in a freehand 3D ultrasound scan, the probe is also moving as the contact pressure varies.

It is this estimate of compression between neighbouring images which can potentially be replaced by the RF-based estimates used in elastography. There is good reason for doing this, since the RF data generally has much better axial (depth) resolution than the ultrasound image, and contains phase information which is lost in the amplitude-based image data. To do this, we have developed a freehand 3D RF ultrasound system which is capable of recording accurate real time RF ultrasound data at over 30 frames per second, along with the 3D location of the data (Treece et al., 2004).

A good candidate for providing the RF-based probe pressure estimate is the algorithm developed by Pesavento et al. (1999). One of the motivating factors in the amplitude-based algorithm was the desire for real-time operation (it can run at about 14 frames per second on a 3GHz PC), since ultrasound is itself a real-time modality. This motivation is shared in the design of this elastography algorithm. The basis of the algorithm is to take small windows of the analytic RF signal and search for the zero-phase of the complex correlation of this signal with the corresponding signal in the neighbouring image. The phase of the complex correlation varies slowly at approximately the ultrasound centre frequency, and this allows a very fast iterative update to be used to improve the estimation of zero-phase shift. This zero-phase shift is tracked from shallow to deep tissue by repeating the process with subsequent overlapping windows of RF data. The complex correlation is based on a log-compressed RF signal in order to reduce the effects of speckle on the accuracy of the zero-phase search.

Fig. 2(b) demonstrates how this algorithm can be adapted to provide estimates of compression due to probe pressure. Multiple overlapping windows of RF data are used, each 8 periods long at the centre frequency of the ultrasound probe, leading to approximately 100 estimates of vertical shift in each RF vector. An estimate of vertical shift at depth due to probe pressure is derived simply

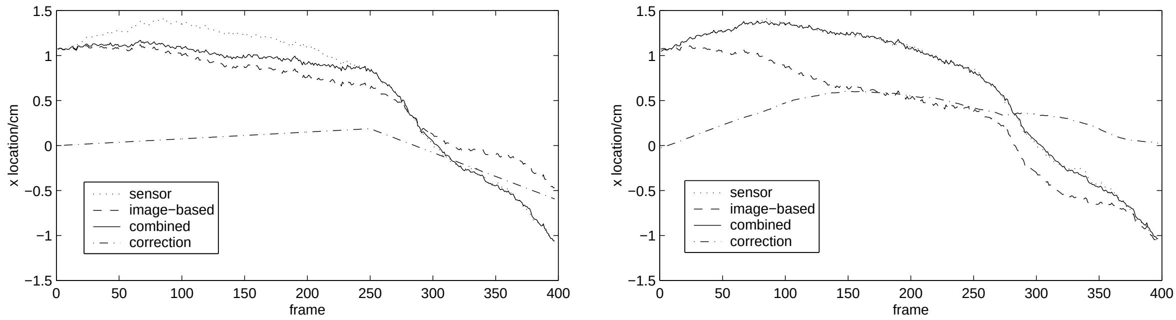

(a) Previous method

(b) New method

Figure 3: Combination of image-based and sensor positions. In the previous method (a), the correction was linear between the first and last frames, and other frames with significant differences between sensor and image-based values. With the new method (b), a polynomial-based correction is globally optimised, giving a much smoother function which more closely matches the actual differences. The combined positions still correctly track operator-induced tremor, but not noise in the readings from the position sensor.

from an average of these estimates over all RF vectors. The average is weighted by the square of the maximum normalised correlation at the location corresponding to the zero phase of the complex correlation. If this is high, the implication is that the anatomy has not changed other than being shifted in depth, which will lead to a good estimate. If this is low, then the anatomy has changed as well as (possibly) the probe contact pressure, which will lead to a degraded estimate.

2.2 Combination of image-based location with position sensor readings

Once image-based positions have been calculated, they are combined with the position sensor readings. The aim in this step is to maintain the smoothness and short range precision of the image-based locations, whilst re-incorporating the long range accuracy, but not the noisiness, of the sensor-based locations. In the algorithm described in (Treece et al., 2002), this was done by assuming the position sensor reading was correct for the first and last frames, and calculating a correction between image-based and sensor readings which was simply a linear interpolation of the errors at these locations. This correction was then added to the image-based positions to give the combined results. The process was repeated between each previous end-point and the frame with the maximum translational difference between combined and sensor-based positions, until the errors reached a sufficiently small scale, given the expected noise in the sensor-based readings.

Typical results of this simple and practically-motivated scheme are shown in Fig. 3(a). This suffered from two drawbacks: firstly it assumed that there was no error in the sensor-based positions of the first and last frame, and secondly the linear interpolation was *not* smooth if it had to be iterated from several end-points. Both of these can be overcome by performing a global optimisation, and using a polynomial to generate the shift in the image-based positions to make them match the sensor-based positions. The optimisation minimises both the difference in in-plane location and rotation between the two sets of readings, for all frames. Quartic polynomials of the form $c_4t^4 + c_3t^3 + c_2t^2 + c_1t + c_0$ are used to generate the three in-plane parameters (two translations and one rotation) which provide a shift of one frame relative to the last frame, additional to that predicted by the image-based positions. The coefficients $c_0 \ldots c_4$ of the three polynomials are found by the optimisation process. The parameter t is derived from the distance of the centre of a frame along a path through the centres of all the previous frames. The optimisation ceases when there is no further change in polynomial coefficients.

Fig. 3(b) shows the correction for the new algorithm. The global optimisation has ensured that the combined positions are much closer to the original sensor-based positions than with the previous algorithm. In addition, the correction is indeed much smoother, preserving the local smoothness of the image-based positions in the combined data.

3 Results

Ultrasound data was acquired with a Diasus ultrasound machine[1], using a 5-10MHz linear array probe, on a 4cm depth setting. Analogue RF ultrasound data after focusing and time-gain compensation was synchronously converted to 14-bit digital data using a Gage CompuScope 14100 PCI analogue to digital converter[2], and transferred at 30 frames per second to a 3GHz PC running Linux. The probe position was sensed by a Polaris[3] optical system tracking an AdapTrax[4] tracker mounted on the probe, and the system was calibrated to an accuracy of ± 0.6mm (Treece et al., 2003). Calibration, acquisition, processing and display of the data were performed by Stradx (Prager et al., 1999)[5]. Pressure corrections, for all the tested algorithms, were calculated at 14 frames per second for amplitude-based and 12 frames per second for RF-based techniques.

For the RF-based techniques, the RF data was converted to an analytic signal using matching filters with a unity gain 5-10MHz pass band, one with even and one with odd symmetry in the coefficients (i.e. a Hilbert filter). Each frame contained 127 vectors, each with 3827 samples at the sampling frequency of 66.67MHz. For amplitude-based techniques, the amplitude of the analytic RF signal was log-compressed suitable for an 8-bit range, and linearly interpolated to create an ultrasound image of 421×487 pixels, with equal scale in each dimension.

Creating the amplitude image from the RF data ensured that exactly the same data was used for all probe pressure correction techniques examined. However, there was also a subsidiary benefit of guaranteeing the images were entirely free from persistence, or frame averaging. This is a technique (originating from the construction of ultrasound machines on oscilloscope displays) for increasing the perceived signal-to-noise level in the ultrasound images by time-averaging several images, often implemented by adding a certain percentage of the new image to that already shown. Unfortunately, even small levels of persistence in ultrasound images have the undesirable result of corrupting image-based estimates of the relative positions of the frames. Each frame contains some of the previous data as well as the new data: the former will correlate maximally with the previous frame at zero offset; the latter at an offset dependent on the movement of the probe between frames. The result is to give a maximal correlation at a slightly smaller offset than was actually the case given probe movement.

This foreshortening effect is demonstrated by the panoramic scans shown in Fig. 4. A persistence-free panoramic data set was recorded, following which various data sets were created from this by simulating varying levels of persistence. Panoramas were constructed from these data sets using

[1] Dynamic Imaging Ltd., `http://www.dynamicimaging.co.uk/`
[2] Gage Applied Technologies Inc., `http://www.gage-applied.com/`
[3] Northern Digital Inc., `http://www.ndigital.com/`
[4] Traxtal Technologies, `http://www.traxtal.com/`
[5] `http://mi.eng.cam.ac.uk/~rwp/stradx/`

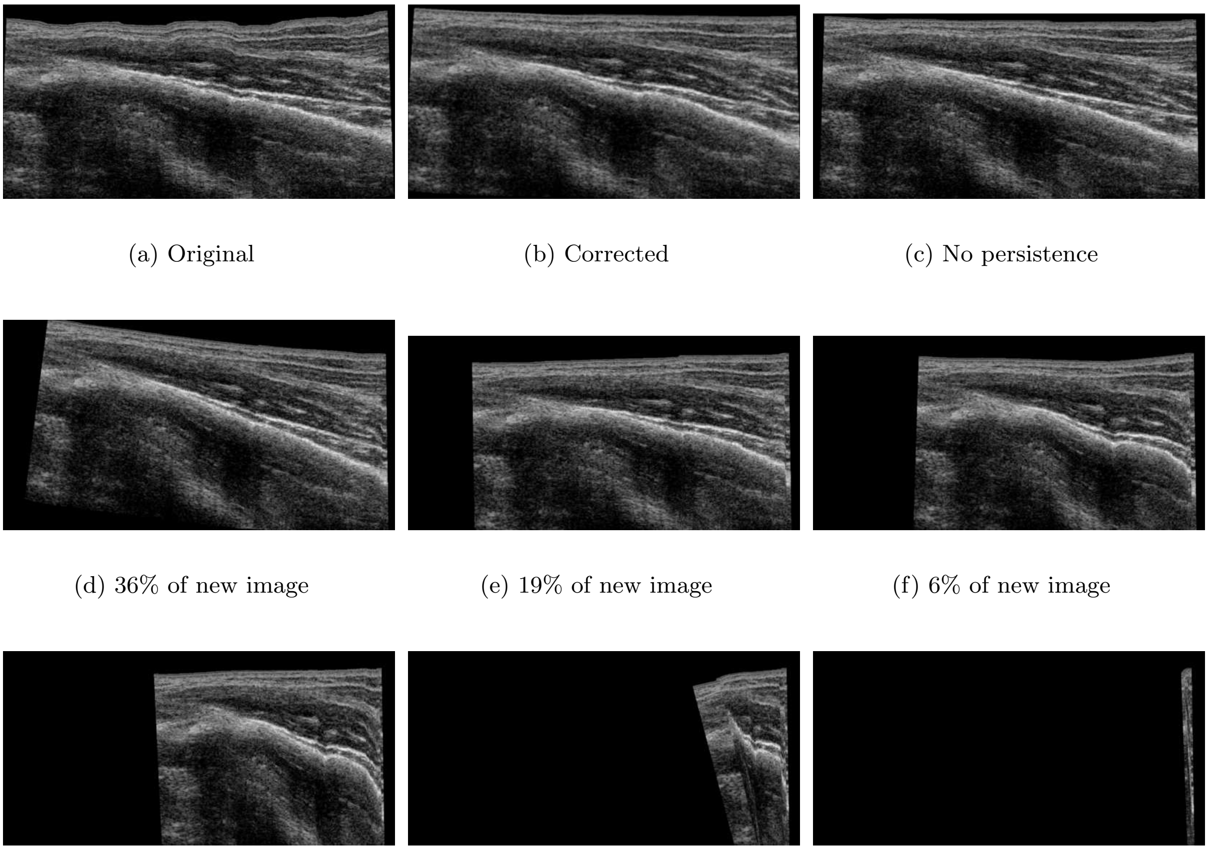

(a) Original (b) Corrected (c) No persistence

(d) 36% of new image (e) 19% of new image (f) 6% of new image

(g) 3% of new image (h) 2% of new image (i) 1.5% of new image

Figure 4: Projected panoramas constructed with varying levels of image persistence. (a) is based on the recorded positions from the position sensor. (b) is after probe pressure correction, combined with the original sensor readings. (c) to (h) are constructed from image-based positions only (i.e. the position sensor is not used), with varying persistence.

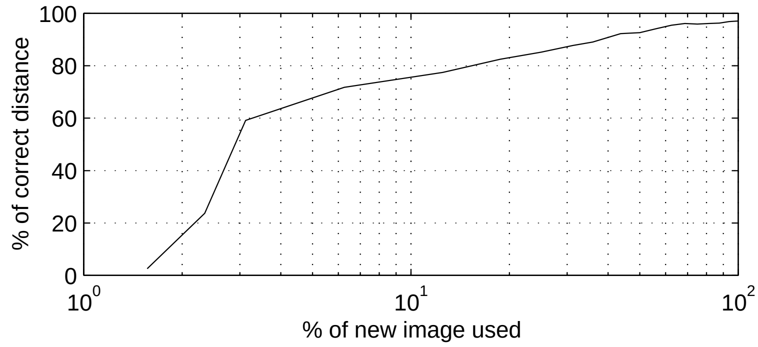

Figure 5: Amplitude-based correlation is affected by persistence, which tends to shorten the correlation distance. The graph demonstrates the percentage change in distance as the percentage of the new image used is lowered (100% implies no persistence).

image-based positions alone. It is indeed apparent from these images that high levels of persistence cause considerable foreshortening in the panoramic data. Fig. 5 is a graph of the apparent distance between two points at extreme ends of the panorama compared to the real distance given by the position sensor. This shows that there is some degradation even for relatively low levels of persistence. The data in (Treece et al., 2002) was recorded directly from ultrasound images, on an ultrasound machine on which it was impossible to turn off persistence completely. This may well explain the inconclusive results for probe pressure correction in this paper.

Confirming the accuracy of ultrasound probe pressure correction techniques is very hard. It is unfortunately not straightforward to compare the (corrected) ultrasound data with other non-contact modalities, e.g. Computed Tomography or Magnetic Resonance Imaging — ultrasound does not measure a material property, let alone one which is compatible with these modalities. The data would also need cross-modality registration before a comparison could be made, and this in itself is a difficult task with its own errors. It is also hard to make a realistic phantom for such experiments: not only must the phantom scatter and attenuate ultrasound in the same way as typical tissue, it must also contain complex structure which deforms appropriately under contact pressure. A variety of experiments have therefore been conducted, each revealing different properties of the algorithms, but also with different limitations.

Firstly, contact pressure was varied with the probe held otherwise still over the same location of an ultrasound phantom. The difference between each frame of corrected data and the first frame should then reveal the success of the probe pressure correction algorithm. However, this type of static data does not represent the situation in a true 3D scan. Secondly, several 3D scans were acquired of an arm in a water bath, and compared to a non-contact scan of the same arm. The experiment allowed the direct assessment of corrected data to non-contact data, nevertheless the scanning protocol in a water bath does not generate the same patterns of probe contact pressure as when using ultrasound gel. Thirdly, several 3D scans were acquired of the same area, using normal scanning methods. None of these scans were non-contact, but it is conjectured that if the correction is physically appropriate, then the sets of corrected 3D data would be more similar to each other than the sets of uncorrected data. Finally, several panoramic scans were compared to each other in a similar manner to the volume data.

In each case, for each acquired data set, five derived data sets were created, referred to as follows:

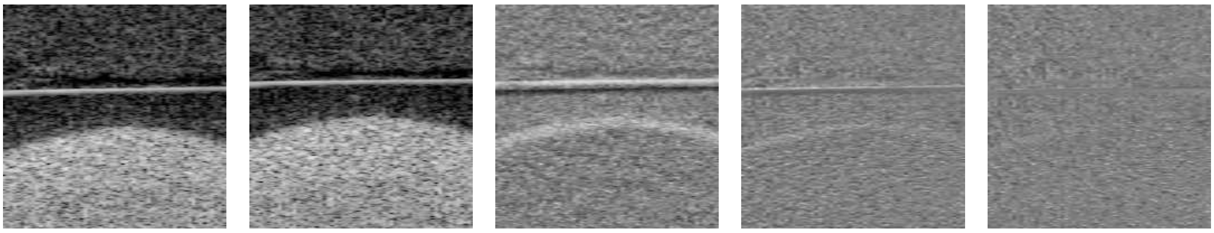

(a) (b) (c) (d) (e)

Figure 6: How a static ultrasound image is affected by probe contact pressure. A sequence of 150 frames was acquired with the probe and subject stationary except for three cycles of a manually applied contact pressure. (a) shows a section from the first image, and (b) from the 75th. (c) Shows the difference between these images before, (d) after amplitude-based and (e) after RF-based correction of the entire sequence.

ORIG The original uncorrected data.

RIGID The data corrected with rigid image-based correction only, which removes noise in the position sensor readings, but does not correct for probe contact pressure.

OLDAMP rigid plus amplitude-based probe pressure correction, using the previous method for combining image-based and sensor positions.

NEWAMP rigid plus amplitude-based probe pressure correction, using the new method for combining image-based and sensor positions.

NEWRF rigid plus RF-based probe pressure correction, again using the new method for combining image-based and sensor positions.

3.1 Static scan of an ultrasound phantom

An ultrasound phantom containing two distinct materials was scanned with the probe and phantom stationary except for three cycles of a manually applied contact pressure. 150 frames were acquired during the scan. Fig. 6 shows some sample data from this sequence. It is apparent from the difference images in Fig. 6(c) to (e) that in both the NEWAMP and NEWRF data sets, each frame is more similar to the first frame after correction than before correction. It is equally apparent that the RF-based algorithm performs better than the amplitude-based algorithm in this case. This similarity can be quantified by calculating the normalised correlation coefficient c_i between each frame i and the first (zero pressure) frame:

$$c_i = \frac{\sum_{x,y} u_0(x,y) u_i(x,y)}{\sqrt{\sum_{x,y} u_0(x,y)^2 \sum_{x,y} u_i(x,y)^2}} \tag{1}$$

where $u_i(x,y)$ is the 8-bit log-compressed amplitude of the data in frame i at pixel location (x,y).

Fig. 7 shows a plot of c_i for each frame, before and after amplitude-based and RF-based correction. Firstly it is clear that, as in Fig. 6, both amplitude-based and RF-based correction work well, with

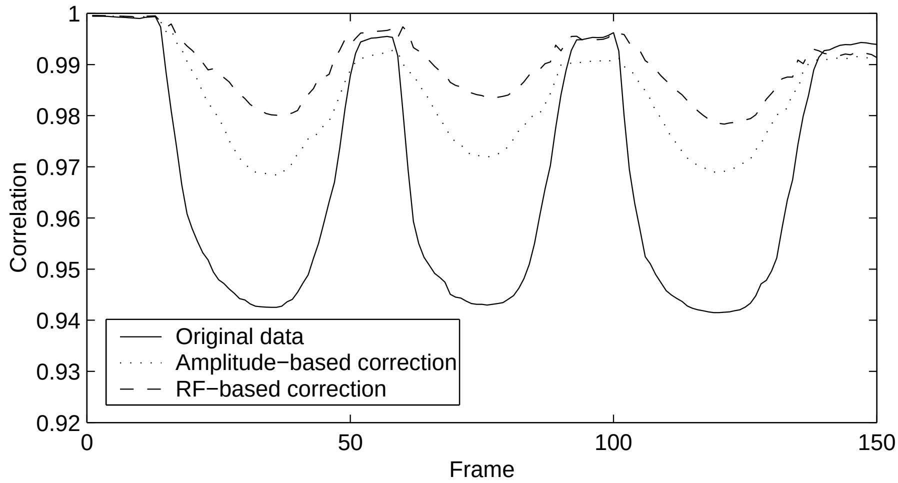

Figure 7: The correlation of each frame with the first in the sequence of Fig. 6 is shown, before and after amplitude-based and RF-based probe pressure correction.

RF-based correction out-performing amplitude-based correction on *all* frames. Secondly, it is clear that the drift in the accumulated corrections is only very slight — the correlation for the corrected 150th frame to the first is nearly the same as for the uncorrected data, despite summing sequential corrections across all 150 frames. Thirdly, there is a residual reduction in correlation for frames which had higher contact pressures. This is to be expected, since the pressure correction algorithm assumes that the subject compresses entirely in a vertical direction, whereas in practice vertical compression induces some horizontal expansion. The algorithm can correct the former, but not the latter.

Fig. 8 shows a comparison between amplitude-based and RF-based correction for three pairs of images from this sequence. The RF-based relative compression estimate is much smoother and in each case clearly reveals a kink in the compression curve: the lower region of the phantom is stiffer than the upper region as well as having a higher backscatter coefficient. Unsurprisingly, the raw amplitude-based estimate is considerably more noisy. What is perhaps more surprising is that the derived amplitude-based estimate is actually very close in each case to the RF-based estimate. This implies that the noisy amplitude-based estimate is nevertheless relatively free of bias, and explains why this performs so well.

3.2 3D scan in a water bath

Ten 3D ultrasound scans were acquired of a forearm held in a water bath, keeping the arm as still as possible during the scans, and moving the probe in the same manner each time, over the same anatomy. Sample frames after each correction was applied to one of the acquired data sets are shown in Fig. 9. In the first scan, the probe was held slightly above the arm, but beneath the surface of the water; the remaining scans were acquired with realistic levels of varying contact pressure. Each of the contact data sets were then corrected by the methods outlined previously, and compared to the non-contact data using a similar metric as in eq. (1):

$$k_i = \frac{\sum_{x,y,z} v_w(x,y,z) v_i(x,y,z)}{\sqrt{\sum_{x,y,z} v_w(x,y,z)^2 \sum_{x,y,z} v_i(x,y,z)^2}} \quad (2)$$

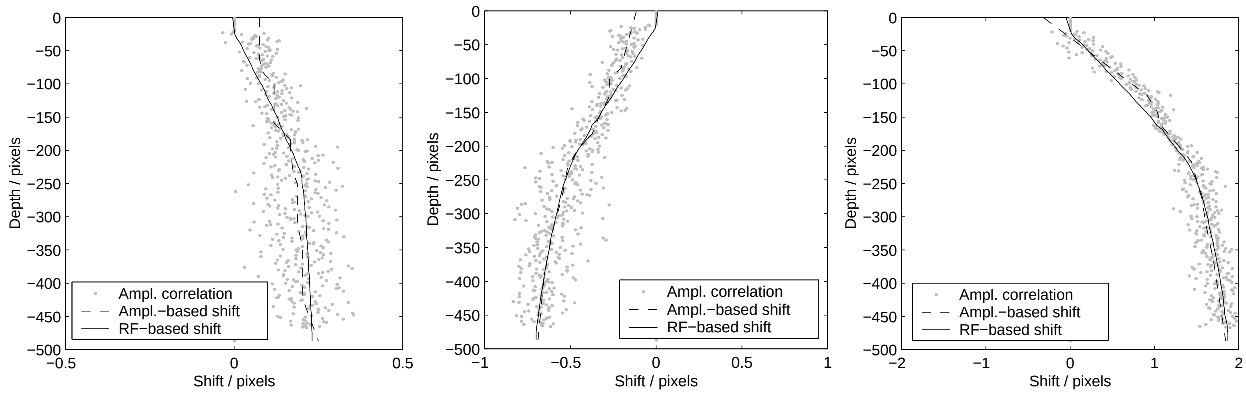

Figure 8: Amplitude-based and RF-based probe pressure estimates, for three pairs of images in the data of Fig. 7. The dots show the raw amplitude-based correlation, the dashes the derived probe pressure estimate, and the solid line the RF-based probe pressure estimate.

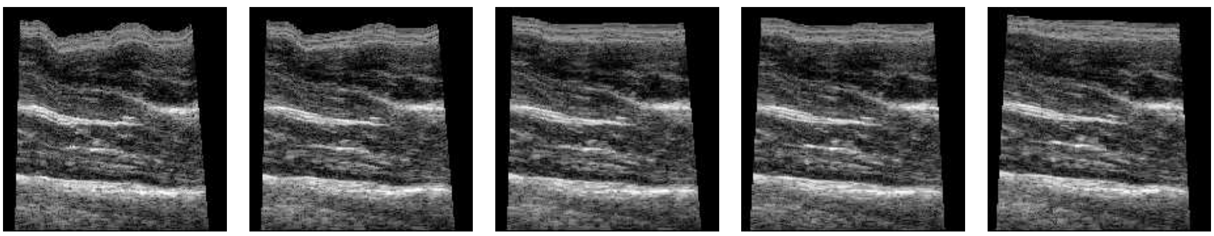

(a) ORIG (b) RIGID (c) OLDAMP (d) NEWAMP (e) NEWRF

Figure 9: The same frame from re-sampled 3D corrected data of an arm in a water bath. (a) shows the original data, (b) with the original rigid and (c) non-rigid correction. (d) is after new amplitude-based correction, and (e) after RF-based correction.

Table 1: Improvement in correlation coefficient x for 9 data sets of Fig. 9, compared to a non-contact data set, after correction. The mean $\bar{x}$, standard deviation σ, estimated probability $\Pr(x > 0)$ that the correlation technique improved the similarity of the data sets to the non-contact data, and the number $n(x < 0)$ for which the correlation was *degraded*, are shown. These figures are calculated assuming no gross movement between data sets, and allowing for a gross rigid shift to maximise the correlation.

	No movement				**Gross rigid shift**			
	$\bar{x}$	σ	$\Pr(x > 0)$	$n(x < 0)$	$\bar{x}$	σ	$\Pr(x > 0)$	$n(x < 0)$
RIGID	0.0023	0.0048	0.687	1	0.0029	0.0022	0.910	0
OLDAMP	0.0070	0.0110	0.739	3	0.0099	0.0047	0.983	0
NEWAMP	0.0058	0.0137	0.665	2	0.0121	0.0033	0.999	0
NEWRF	0.0086	0.0117	0.770	3	0.0124	0.0043	0.998	0

where the subscript $_w$ denotes the non-contact data, and $_i$ the remaining 9 data sets. In order to compare the data, all data sets were re-sampled to the same regular cubic 3D array, each side of length 128 voxels. This re-sampling was performed *after* correction for probe pressure. $v_i(x, y, z)$ is therefore the re-sampled 8-bit log-compressed amplitude of the data in volume i at voxel location (x, y, z).

k_i was first calculated for all data sets in the re-sampled coordinate frame, assuming that there was no movement in the (live) subject between scans. Since it was impossible to completely eliminate movement, k_i was also calculated across a small range of relative gross rigid shifts of $\pm$2mm in each direction, and the maximum k_i recorded. This successfully accounted for subject movement, increasing k_i and also reducing the range of k_i across scans, but unfortunately also masked any gross movement errors which may have been induced by the correction process. Hence both coefficients have been included in the following tables. In order to analyse whether each correction technique improved the similarity of the data with the non-contact scan, the difference x between k_i with correction, and k_i with no correction was calculated and analysed. The probability that the correction technique improved the data was also assessed, by using the mean $\bar{x}$ and standard deviation σ of x, and assuming that x is normally distributed. The results are contained in Table 1.

It is apparent from Table 1 that either there was some movement between scans, or the correction techniques caused incorrect gross shifts of the data volumes, since $\bar{x}$ is much higher and σ much lower for the data allowing for movement. The fact that the NEWAMP data, using global optimisation of image-based and sensed position recombination, generated poorer results than the OLDAMP data, indicates that these shifts were probably due to subject movement. The results allowing for them show that there is significant improvement in all data sets involving probe pressure correction, with the correction improved for the new position sensor optimisation strategy but the same for both the amplitude-based and RF-based techniques.

3.3 3D scans in multiple directions

Twenty 3D data sets were acquired, again of an arm, ten using one scanning pattern, and ten an approximately orthogonal scanning pattern, with the arm resting on a table, and using normal amounts of ultrasound gel. An illustration of the scanning pattern in each case is given in Fig. 10. All the data sets were re-sampled to the same 3D cubic array; a sample frame from one of the data sets from each scanning pattern is shown in Fig. 11. Since none of these data sets were non-contact,

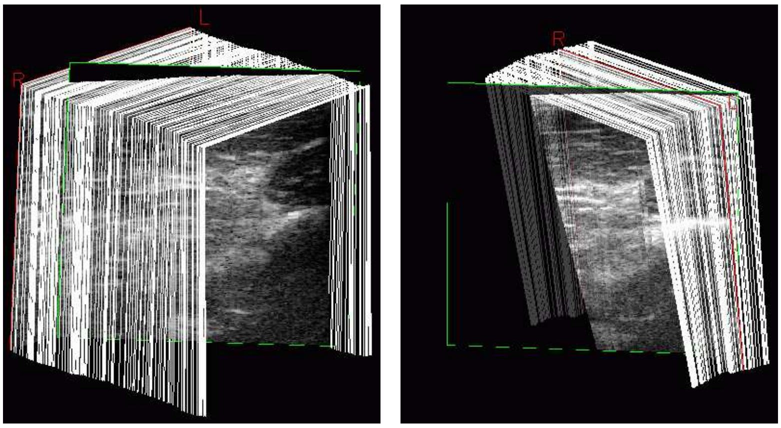

(a) Tangential scanning (b) Longitudinal scanning

Figure 10: The same section of arm was scanned ten times in direction (a) and ten times in direction (b), keeping both the arm and the position sensor coordinate reference steady throughout.

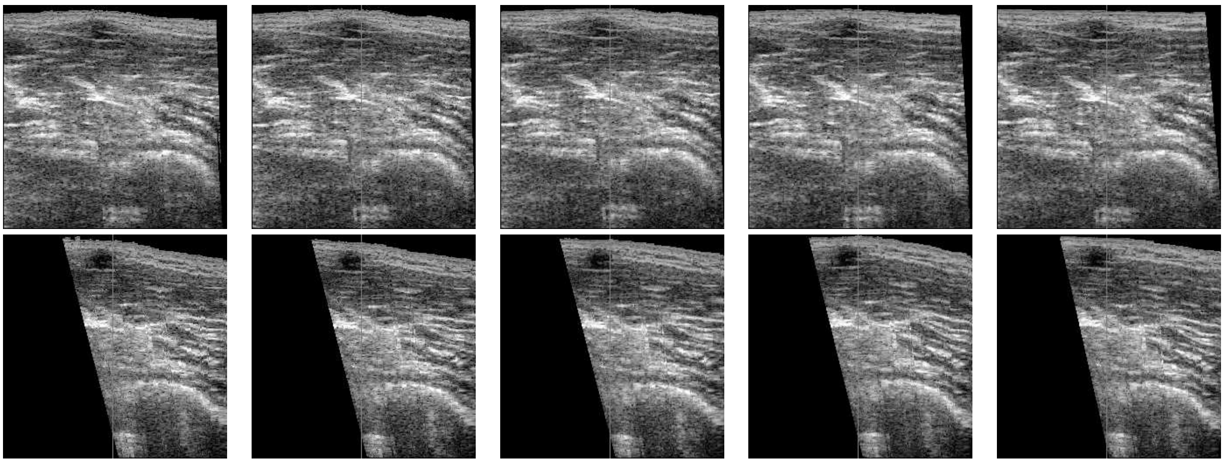

(a) ORIG (b) RIGID (c) OLDAMP (d) NEWAMP (e) NEWRF

Figure 11: The same frame from re-sampled 3D corrected data of an arm scanned as in Fig. 10(a) (top row) and Fig. 10(b) (bottom row). (a) shows the original data, (b) with the original rigid and (c) non-rigid correction. (d) is after new amplitude-based correction, and (e) after RF-based correction.

Table 2: Improvement in correlation coefficient x for all 190 pairs of the 20 data sets of Fig. 11, after correction. The mean $\bar{x}$, standard deviation σ, estimated probability $\Pr(x > 0)$ that the correlation technique improved the consistency of the data sets, and the number of pairs $n(x < 0)$ for which the correlation was *degraded*, are shown. These figures are calculated assuming no gross movement between data sets, and allowing for a gross rigid shift to maximise the correlation.

	No movement				**Gross rigid shift**			
	$\bar{x}$	σ	$\Pr(x > 0)$	$n(x < 0)$	$\bar{x}$	σ	$\Pr(x > 0)$	$n(x < 0)$
RIGID	-0.0062	0.0122	0.307	129	0.0001	0.0054	0.511	72
OLDAMP	0.0021	0.0152	0.556	66	0.0103	0.0069	0.933	18
NEWAMP	0.0088	0.0104	0.803	27	0.0088	0.0079	0.869	26
NEWRF	0.0129	0.0104	0.894	12	0.0146	0.0087	0.954	4

k_i from eq. (2) was calculated for all 190 pairs of the 20 scans. The analysis was otherwise identical to the scan in a water bath

Table 2 shows the change in k_i, x, for all sets of corrected data. k_i for no movement shows an improvement for both the NEWAMP and NEWRF data, where the global optimisation was used, but with fairly low confidence levels of 80% and 89% respectively. Nevertheless, the RF-based estimate outperforms all the other estimates, and significantly so when allowing for a gross movement in the data. The fact that the OLDAMP data improves so much after allowing for movement shows that the old correction algorithm applied an incorrect rigid shift to the data.

3.4 Panoramic scan

Panoramic, or extended-field-of-view, scans are not strictly 3D ultrasound data, since the probe is moved at least approximately *in its own plane*. Nevertheless, they are based on a sequential series of gradually varying images, and hence the probe pressure algorithms are equally applicable. Panoramas are usually constructed by matching each image to the previous image, and hence they *follow* the location and orientation of the ultrasound probe as it moves (Prager et al., 1999). In order to calculate how probe pressure correction affects panoramic data, we need a type of panorama which can be compared across data sets. Most panoramas are highly sensitive to not only the location of the probe but also the exact way in which it is moved, and hence the same anatomy will not in general appear in the same place in a repeated panorama.

With freehand 3D ultrasound data, it is possible to generate a panoramic image in a slightly different way which is much less dependent on the exact orientation of the ultrasound probe as it is moved. Each image is projected on to a plane, which is approximately parallel to all the images, and the data used at each point on this plane is from the image which had the nearest projected central column. Such a 'projected' panorama, an example of which is shown in Fig. 12, is much more suitable for comparison across multiple data sets.

Eleven panoramic data sets where recorded, once again keeping the subject still between each scan, but with varying probe pressure. A line was drawn on the subject's skin to ensure that the probe followed approximately the same trajectory during each scan. Projected panoramas were created on the same plane for each data set, corrected by each technique. Typical panoramas are shown in Fig. 13. Each pair of panoramas could then be compared using the correlation coefficient in eq. (1) and the differences in correlation analysed as before. In this case subject movement could only be allowed for within the plane of the panorama, by searching for the maximum correlation in

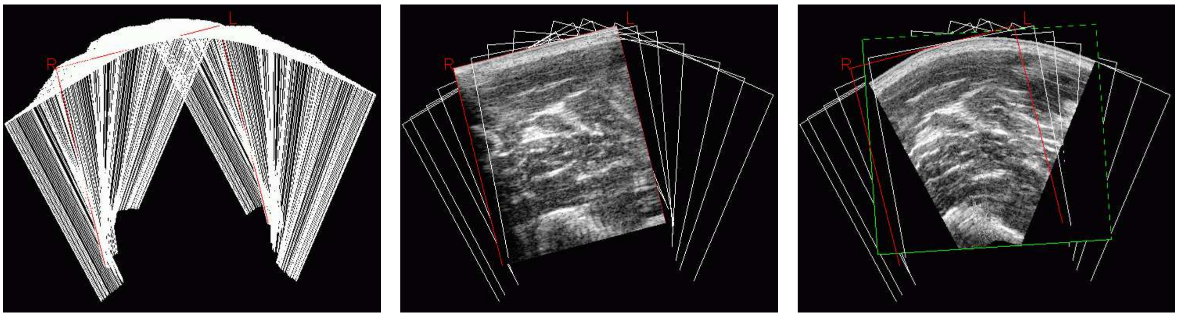

(a) All outlines (b) Sample image (c) Projected panorama

Figure 12: The panorama used for comparison between data sets is actually the projection of the centre columns of each B-scan onto a plane roughly parallel with the original images. (a) shows the entire data set — each white ‘goal post’ represents one ultrasound image. (b) shows a sample ultrasound image from a thinned-out data set, and (c) the projected panorama of this data.

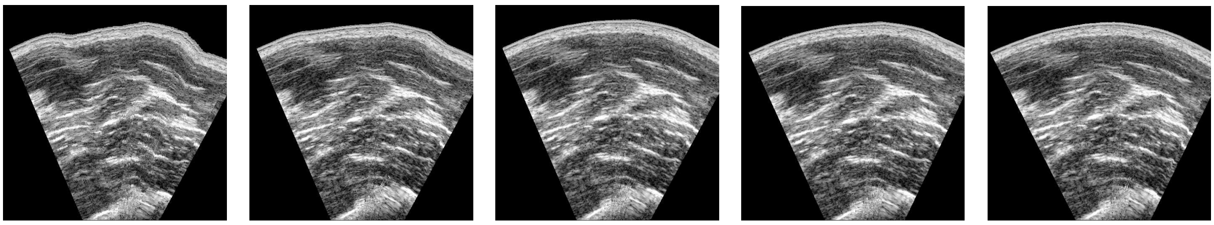

(a) ORIG (b) RIGID (c) OLDAMP (d) NEWAMP (e) NEWRF

Figure 13: Corrected projected panoramas from one data set. (a) shows the original data, (b) with the original rigid and (c) non-rigid correction. (d) is after new amplitude-based correction, and (e) after RF-based correction.

Table 3: Improvement in correlation coefficient x for all 55 pairs of the 11 data sets of Fig. 13, after correction. The mean $\bar{x}$, standard deviation σ, estimated probability $\Pr(x > 0)$ that the correlation technique improved the consistency of the data sets, and the number of pairs $n(x < 0)$ for which the correlation was *degraded*, are shown. These figures are calculated assuming no gross movement between data sets, and allowing for a gross rigid shift to maximise the correlation.

	No movement				**Gross rigid shift**			
	$\bar{x}$	σ	$\Pr(x > 0)$	$n(x < 0)$	$\bar{x}$	σ	$\Pr(x > 0)$	$n(x < 0)$
RIGID	0.0022	0.0134	0.567	20	0.0123	0.0059	0.982	0
OLDAMP	0.0142	0.0133	0.858	8	0.0203	0.0105	0.974	2
NEWAMP	0.0175	0.0102	0.957	3	0.0209	0.0095	0.986	0
NEWRF	0.0096	0.0153	0.736	12	0.0173	0.0101	0.957	1

each of the in-plane directions only.

Table 3 contains the results of this analysis. Although there is an improvement in correlation when allowing for movement, it is less striking than in the volume data, which is probably due to the reduced search range in this case. In both cases, the amplitude-based correction performed better than the RF-based correction, with the NEWAMP data giving the best results, although all improvements were significant once movement was compensated for.

4 Conclusions

Anatomical deformation due to probe contact pressure generates changes in high resolution ultrasound data which are significant and apparent from visualisations of the data. 3D ultrasound data acquired using the freehand scanning approach suffers from variation in this pressure, but this variation can be reduced by using correction techniques. All the experiments show that using either amplitude-based or RF-based probe pressure correction generates data which is more self-consistent, and (at least in the case of the water bath experiment) more similar to data which would have been acquired if non-contact ultrasound was possible. This is equally the case for volumetric and panoramic data.

In ideal scenarios like the static phantom experiment, RF-based correction has been shown to be slightly more accurate than amplitude-based correction, and results in less drift when estimates are accumulated across a large sequence of images. This slight improvement carries over to the volume data sets, although the difference between this and the amplitude-based correction is not significant. For panoramic data, the amplitude-based technique performs better, though once again, the difference to the RF-based technique is fairly small.

It is perhaps surprising that the amplitude-based technique performs so well compared to the RF-based technique, which has a much better axial resolution, and gives apparently much less noisy estimates. The fact that this is the case demonstrates that the assumptions made in deriving the amplitude-based probe pressure estimates are indeed valid. Both the smoothness and monotonicity which are enforced on these estimates turn out to already exist in the RF-based estimates.

In terms of convenience rather than performance, both amplitude-based and RF-based techniques are relatively efficient, running at approximately half the typical frame rate on a single processor 3GHz PC. This means correction can be performed after acquisition in typically 20 seconds on a 3D data sequence containing 300 frames. With advances in processor speed, or the use of several processors, it will no doubt soon be possible to run such algorithms in real time, as the data is acquired. Clearly the RF-based technique requires access to RF data, which is not always available. On the other hand, the amplitude-based technique can be performed on ultrasound images acquired via a standard frame grabber, so long as the ultrasound persistence or frame averaging is minimised.

In conclusion, amplitude-based probe pressure correction works well in most practical situations, and is comparable with more sophisticated RF-based techniques. Consequentially, freehand 3D ultrasound users who do not have access to RF data are still in a position to perform state-of-the art probe pressure correction.

5 Acknowledgements

Graham Treece is supported by an EPSRC/RAEng Postdoctoral Fellowship. Dynamic Imaging Ltd. provided a modified ultrasound machine to enable the acquisition of analogue RF and digital amplitude ultrasound data in real time.

References

Artignan, X., Smitsmans, M. H. P., Lebesque, J. V., Jaffray, D. A., van Her, M., Bartelink, H., 2004. Online ultrasound image guidance for radiotherapy of prostate cancer: Impact of image acquisition on prostate displacement. International Journal of Radiation Oncology, Biology and Physics 59 (2), 595–601.

Greenleaf, J. F., Fatemi, M., Insana, M., 2003. Selected methods for imaging elastic properties of biological tissues. Annual Review of Biomedical Engineering 5, 57–78.

Pesavento, A., Lorenz, A., Siebers, S., Emmert, H., 2000. New real-time strain imaging concepts using diagnostic ultrasound. Physics in Medicine and Biology 45 (6), 1423–1435.

Pesavento, A., Perrey, C., Krueger, M., Ermert, H., Sep. 1999. A time-efficient and accurate strain estimation concept for ultrasonic elastography using iterative phase zero estimation. IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control 46 (5), 1057–1067.

Prager, R. W., Gee, A. H., Berman, L., 1999. Stradx: real-time acquisition and visualisation of freehand 3D ultrasound. Medical Image Analysis 3 (2), 129–140.

Smith, W. L., Fenster, A., 2000. Optimum scan spacing for three-dimensional ultrasound by speckle statistics. Ultrasound in Medicine and Biology 26 (4), 551–562.

Treece, G. M., Gee, A. H., Prager, R. W., Cash, C. J. C., Berman, L. H., Apr. 2003. High definition freehand 3D ultrasound. Ultrasound in Medicine and Biology 29 (4), 529–546.

Treece, G. M., Prager, R. W., Gee, A. H., Sep. 2004. Freely available software for 3D RF ultrasound. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI'04). LNCS. Springer, Saint-Malo, France.

Treece, G. M., Prager, R. W., Gee, A. H., Berman, L., Oct. 2002. Correction of probe pressure artifacts in freehand 3D ultrasound. Medical Image Analysis 6 (3), 199–214.

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

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

Xiao, G., Brady, J. M., Noble, J. A., Apr. 2002. Nonrigid registration of 3-D free-hand ultrasound images of the breast. IEEE Transactions on Medical Imaging 21 (4), 405–412.