
Freehand ultrasound elastography with a 3D probe

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

CUED/F-INFENG/TR 577

May 2007

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

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

Abstract

This paper presents the first near-real-time freehand ultrasound elastography system using a three-dimensional (3D) mechanical probe. Acquisition is complete within 2 seconds, and only an additional 20 seconds are required for generation of a full 3D strain volume. The strain is *axial*, with estimates of lateral and elevational tissue movement used to increase the accuracy of the axial strain measurement. This is the first time all system components have been extended to 3D, i.e. 3D windows are used to track displacement, which is tracked in all directions, and 3D kernels are used for least squares gradient estimates. Normalisation of the freehand 3D strain data is also applied across the whole volume. The system is tested using a novel research 3D RF system with real time control over the stepper motor driving the ultrasound probe, and real-time streaming of RF ultrasound data. The paper proves the concept rather than making significant comments on the achievable accuracy in 3D, though we demonstrate that the high performance of the 2D techniques which we extend appears to carry through to *in vitro* and *in vivo* 3D data. The result is a fast and high resolution 3D image of normalised axial strain.

Contents

1	Introduction	1
2	Method	2
2.1	Calculating axial, lateral and elevational offsets	4
2.2	Tracking offsets over the 3D volume	5
2.3	AMC and interpolation to a regular grid	8
2.4	Gradient, normalisation and display	8
3	Experiments and results	12
3.1	<i>In vitro</i> strain images	13
3.2	<i>In vivo</i> strain images	13
4	Discussion	13
4.1	Processing and acquisition time	13
4.2	Lateral and elevational tracking	17
4.3	System performance	17
5	Conclusions	18

1 Introduction

Ultrasound strain imaging, or elastography, is a way of visualising tissue stiffness: it can be seen as the imaging equivalent of manual palpation. It is currently receiving much attention from within both the research and commercial sectors. Although the clinical usefulness has yet to be proven beyond doubt, it seems very likely that this will become an important technique. Applications have already been indicated in imaging certain tumours (Garra et al., 1997), atherosclerosis (de Korte et al., 1998) and any other masses which are expected to be stiffer than the surrounding tissue. It may also have a rôle in studying the physiological and pathological mechanical properties of soft tissue (Genisson et al., 2004; Vogt and Ermert, 2005).

High quality quasi-static freehand 2D ultrasound strain imaging is currently possible in real time (Lindop et al., 2007a), by comparing radio-frequency (RF) data from sequential ultrasound images during a slight deformation of the tissue due to contact pressure from the ultrasound probe. It has recently been shown that it is possible to produce stable, high quality strain data over a wide range of freehand motions of the probe (Lindop et al., 2007b). The approach has considerable benefits over other techniques in that no additional equipment is required to induce the required stress field in the tissue: the clinician simply moves a normal ultrasound probe over the anatomy. This potentially makes the technique simpler to use as well as easier to implement.

In just the same way as for 3D ultrasound, the availability of 3D strain data would be beneficial for more detailed geometry and more accurate measurements, in particular of the volume of small or irregular masses. Since the real strain during applied pressure from an ultrasound probe has a pronounced 3D variation, it may also be beneficial to show images of true strain, rather than axial strain, which is the current convention. However, the lateral and elevational components of true strain can generally be estimated with much less precision. Hence, in common with most work in this area, we focus on 3D volumes of *axial* strain, making use of estimates of lateral and elevational tissue movement only to increase the accuracy of such axial strain measurements.

Previous work on 3D strain imaging has generally followed one of two directions. In one approach, 2D data is acquired as the probe is gradually moved in the elevational direction, forming a sequential 3D data set. Neighbouring frames are used to estimate 2D strain images, which when stacked together create a 3D data set of 2D strain estimates. This technique has been successfully implemented for freehand 3D ultrasound (Lindop et al., 2006), and for intra-vascular ultrasound (Schaar et al., 2005; Li et al., 2006). The alternative approach involves the use of mechanical swept 2D probes or 2D phased array probes. Here the probe is held still, a volume of data is acquired, then either a controlled compression is applied (Krueger et al., 1998; Lorenz et al., 1999; Fisher et al., 2006), or in the case of strain estimation as a predictor of temperature change, no compression at all (Anand et al., 2007). 3D acquisition has also been investigated for sonoelastography (Taylor et al., 2000) and for prostate mechanical imaging using a pressure pad (Egorov et al., 2006).

The system presented in this paper is novel in several respects. Firstly, to the authors knowledge, it is the first time that fully 3D windows have been used for displacement and gradient estimation. This is an important step since, if we make the windows more representative of the fundamental resolution of the ultrasound data, we can maximise the quality of the strain estimates for a given strain resolution. This is also the first attempt to acquire 3D strain data with a freehand quasi-static approach on a mechanical 3D probe, with subsequent normalisation of the data also applied to the entire 3D data set (Treece et al., 2007). We achieve this in a system which matches the fastest reported processing time of about 20 seconds (Anand et al., 2007), even though this technique only employed 1D windows.

We start by outlining the important stages in this system in Section 2, before demonstrating results over a range of *in vitro* and *in vivo* data in Section 3. These results are discussed in Section 4 before concluding in Section 5.

2 Method

3D RF data is acquired in real time from a 3D ultrasound probe, which has a geared internal stepper motor controlling a 2D linear phased-array probe head. Each volume of data is acquired with the probe held steady and the motor sweeping in the same direction. A very slight additional pressure (or relaxation) is then applied to the probe in an approximately axial direction before a further volume is acquired. This can conveniently be achieved during the slight pause of just less than one second while the motor returns to the starting position and the RF data is transferred to PC memory. The whole process, including the gap between each frame, takes just less than two seconds. Good results can be achieved with just two volumes of data, however it is also possible to acquire several volumes with slight pressure changes between each, and then combine them in subsequent processing. The results presented hereafter make use of only two volumes.

Having acquired the RF data, strain estimation follows the procedure shown in Figure 1. The axial displacement estimator we employ requires baseband analytic ultrasound signals. Often, such IQ (In-phase and Quadrature) signals are directly available from the probe interface, in which case the displacement estimation can begin directly. However, in our research platform the ultrasound signals are digitised as radio-frequency (RF) passband signals — hence these must first be converted to baseband analytic form. This can easily be done by the use of a pair of passband Hilbert filters, followed by de-modulation at the nominal centre frequency of the ultrasound probe. The baseband signal r is therefore:

$$r = (p \otimes h_r + p \otimes h_i j) e^{-j\omega_c t} \quad (1)$$

where p is the original RF passband signal, ω_c is the approximate probe centre frequency, and h_r, h_i

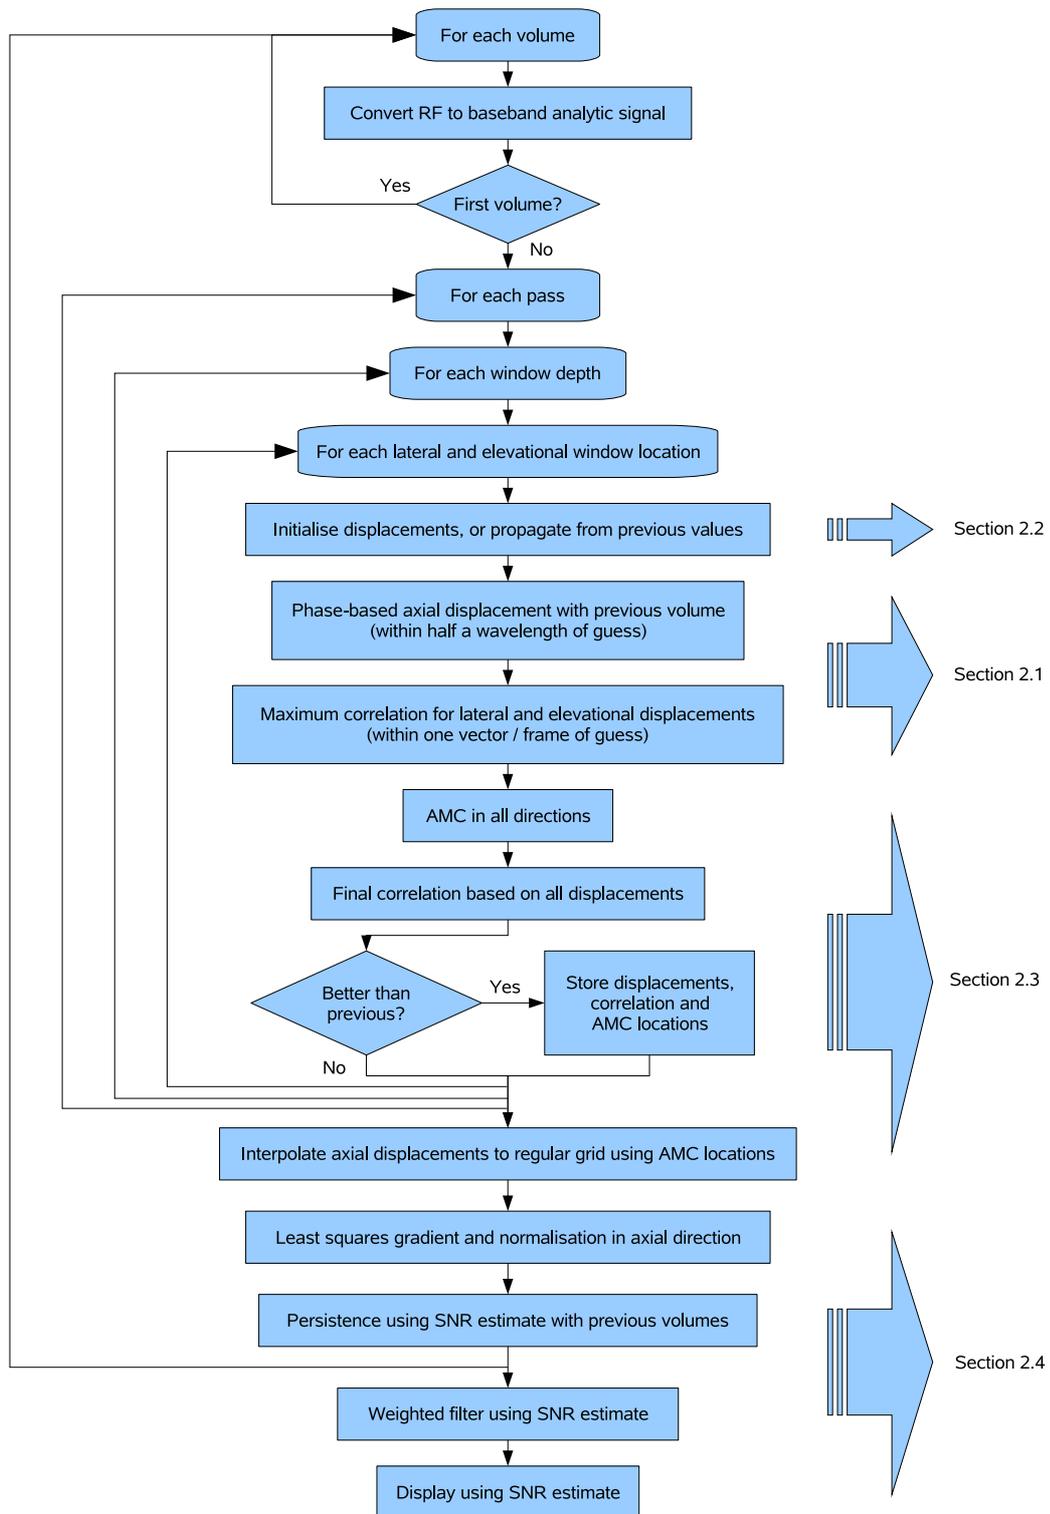


Figure 1: **Description of the 3D strain algorithm.** The flowchart shows the order of events in calculating a 3D axial strain image from recorded RF data. The data itself is recorded in two or more volumes, each with the probe held still. A very slight pressure is manually applied between volumes. The acquisition process takes only a second per volume.

are the filters. Best results are achieved when h_i is a Hilbert (anti-symmetric coefficients) filter with a passband covering only the expected ultrasound probe frequency range, and h_r is a symmetric filter with a frequency response carefully matched to h_i . The convolutions must be calculated at the sampling frequency of 66.67 MHz (synchronous to the ultrasound machine clock), to retain high accuracy phase information. All other operations can be performed after down-sampling, usually by a factor of 5, with substantial benefit in processing speed and negligible cost to the strain accuracy: a suitable choice of factor depends on the ratio of sampling frequency to signal bandwidth.

2.1 Calculating axial, lateral and elevational offsets

We estimate axial displacement using the Efficient Phase Zero Search (EPZS) of Lindop et al. (2006), which is a variant of phase root seeking. This is a fast and accurate method, providing that the tracking of displacements is handled appropriately: we explain how this is achieved in the following section. EPZS gives axial displacements which are inherently sub-sample resolution. It can be performed very efficiently, since only a handful of iterations are required to converge on a given displacement estimate. Indeed, since we do not require normalised correlations (only the phase is needed, and this remains unaffected by normalisation), it is generally only necessary to perform one correlation per displacement estimate. Assuming our initial displacement estimate is d , on the first iteration we require the complex cross-correlation c of the pre- and post-deformation signals, r_a and r_b at relative shift d :

$$c = \sum_{i,j,k} r_{a(i,j,k)} r_{b(i+d,j,k)}^* \quad (2)$$

where (i, j, k) is the extent of the 3D window over which each displacement estimate is calculated, in axial, lateral and elevational directions respectively. Since we only have samples of r_a and r_b , we calculate sub-sample shifts by linearly interpolating. Linear interpolation is sufficient since r is a baseband signal with low frequency content. Hence Equation (2) is actually calculated thus:

$$d_n = \lfloor d \rfloor \quad (3)$$

$$d_r = d - \lfloor d \rfloor \quad (4)$$

$$c = \sum_{i,j,k} \left\{ r_{a(i,j,k)} \left((1 - d_r) r_{b(i+d_n,j,k)}^* + d_r r_{b(i+d_n+1,j,k)}^* \right) \right\} \quad (5)$$

where $\lfloor d \rfloor$ denotes the largest integer less than d .

Equation (5) can be separated into four sums, which are independent of d_r :

$$\begin{aligned} c &= (1 - d_r) \sum_{i,j,k} \left\{ \Re(r_{a(i,j,k)}) \Re(r_{b(i+d_n,j,k)}) + \Im(r_{a(i,j,k)}) \Im(r_{b(i+d_n,j,k)}) \right\} \\ &+ d_r \sum_{i,j,k} \left\{ \Re(r_{a(i,j,k)}) \Re(r_{b(i+d_n+1,j,k)}) + \Im(r_{a(i,j,k)}) \Im(r_{b(i+d_n+1,j,k)}) \right\} \\ &+ (1 - d_r) \sum_{i,j,k} \left\{ \Im(r_{a(i,j,k)}) \Re(r_{b(i+d_n,j,k)}) - \Re(r_{a(i,j,k)}) \Im(r_{b(i+d_n,j,k)}) \right\} \\ &+ d_r \sum_{i,j,k} \left\{ \Im(r_{a(i,j,k)}) \Re(r_{b(i+d_n+1,j,k)}) - \Re(r_{a(i,j,k)}) \Im(r_{b(i+d_n+1,j,k)}) \right\} \end{aligned} \quad (6)$$

Hence, if the new estimate of d lies between the same samples as the previous estimate, i.e. d_n is constant between estimates, the following iteration can be performed simply by adjusting d_r in Equation (6): there is no need to re-calculate the summations. The iterative process therefore

converges with very few iterations actually requiring a summation, which is a significant cost in terms of processing time. In fact, with fairly low strains and fairly densely spaced estimates, the initial guess d is frequently between the same samples as the final converged d , with the result that only one summation is required for each new displacement estimate.

It is possible to use exactly the same zero-phase tracking technique to find the optimal lateral and elevational displacements, however this would require the calculation of an analytic signal in each of these directions. Since we do not require (nor are likely to attain) the same degree of accuracy in these dimensions, we adopt the simpler approach of looking for the maxima of the lateral and elevational correlations. We can speed up this process by noting that in any real scan, we do not expect the lateral or elevational displacements to change by more than one vector or frame between each axial window position. We therefore first search for the axial displacement, then use this displacement to look for the maximal correlation over the neighbouring vectors in the lateral direction, then over the neighbouring frames in the elevational direction.

We do not need to search for these three displacements simultaneously, since our tracking scheme involves a second pass over the data, at which point the searches will be repeated in all three dimensions using the previous best values. Equally, we do not need to search farther afield, since each search is seeded with the best values from the previous search (in the axial direction) so the cumulative shift can in general be much greater than the limited search region in each direction.

In contrast to the axial displacement, the resulting displacements in the lateral and elevational directions are only accurate to the nearest vector or frame¹. However, this accuracy is more than sufficient, since we are only using the lateral and elevational shifts to ensure that we are centred on the correct vector and frame for the much more accurate axial displacement calculation. In other words, the lateral and elevational searches keep us in the right place, but it is from the axial displacement that we calculate the strain for display.

2.2 Tracking offsets over the 3D volume

The searches for relative displacement in all three dimensions will only be successful if the initial guessed displacements are close to the real values: hence it is vital to have a tracking scheme which can ensure that this is the case. The scheme we use is based on the one described by [Treece et al. \(2006\)](#), where displacements at the probe face are initialised to zero (assuming there is no-slip contact between the probe surface and the anatomy being scanned). Each initial estimate is taken to be the estimate at the previous axial depth, over a small lateral range, which achieved the highest correlation value. We trivially extend this to 3D by searching all lateral and elevational locations for each axial location (depth) in turn. At each stage, we choose the displacements which had the highest correlation from a small lateral-elevational array of values centred about the previous axial location.

The resulting technique is better than an exhaustive search, in that it implicitly imposes a consistency on the displacement data which is physically sensible: i.e. we expect the anatomy we are scanning to displace in a piece-wise smooth sense, rather than as a random distribution of unrelated elements. In most cases candidate displacements which result in higher correlations represent the correct actual tissue displacement. There are occasional *specific* instances of displacements with higher correlations *not* being representative of the actual tissue displacement, but these do not propagate since such displacements rarely give rise to a higher correlation at any neighbouring point.

However, we cannot always assume that the displacements should start at zero at the probe

¹In fact, we can achieve slightly better accuracy by adopting the peak interpolation strategy of [Treece et al. \(2005\)](#).

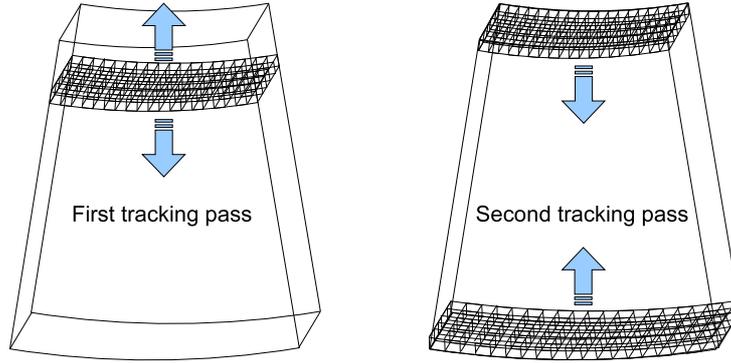


Figure 2: **Displacement processing order.** Axial, lateral and elevational displacements are first calculated at a shallow depth, across all frames (elevationally) and vectors (laterally). These results are used to initialise subsequent estimates at other depths. A first pass progresses outwards towards the top and bottom of the volume, followed by a second which progresses back towards the initial plane.

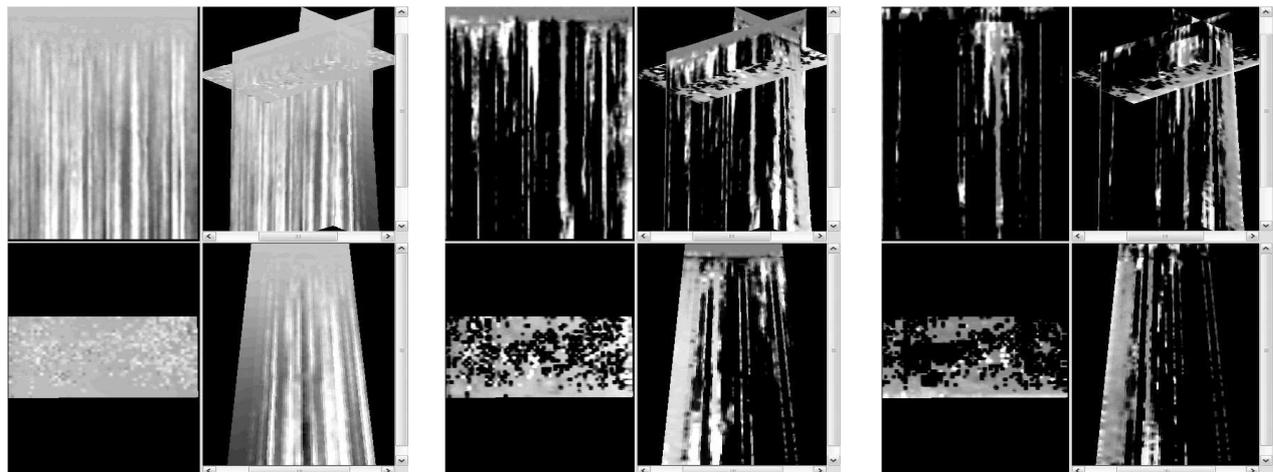
face — if, for instance, the mechanical interface at the probe face is not adequate, even though the ultrasound interface may be. This is likely to happen when scanning through gel if a part of the probe is slightly lifted off the skin. In this case we never get a good estimate to track from, and therefore subsequent estimates never reach the correct values. Equally, major slip near the probe face can be a problem. We can get around this by initialising away from the face of the probe, then progressively search for displacements up and down from this location, as shown in Figure 2. The problem now is that we do not know what the initial displacement at this depth should be, so we use a set of values distributed in a saw-tooth pattern between extreme positive and negative strains. We must ensure that there is at least one value per $\frac{1}{2}$ cycle axially, so that we always have one close to the real displacement. We then rely on the tracking algorithm to selectively propagate the correct values and ignore the incorrect values.

In the second pass shown in Figure 2, we repeat the displacement calculations in all three dimensions working back from the bottom and top of the volume towards the initialisation point, initialising the estimates in the same manner as the first pass, and keeping the new ones if they represent higher correlation values than in the first pass. This has two important effects. Firstly, it allows the propagation of good estimates in the other direction: this allows us to create good strain data both under and above blood vessels, for instance. Secondly, it allows the re-calculation of the displacements in all directions based on a much better estimate of all three.

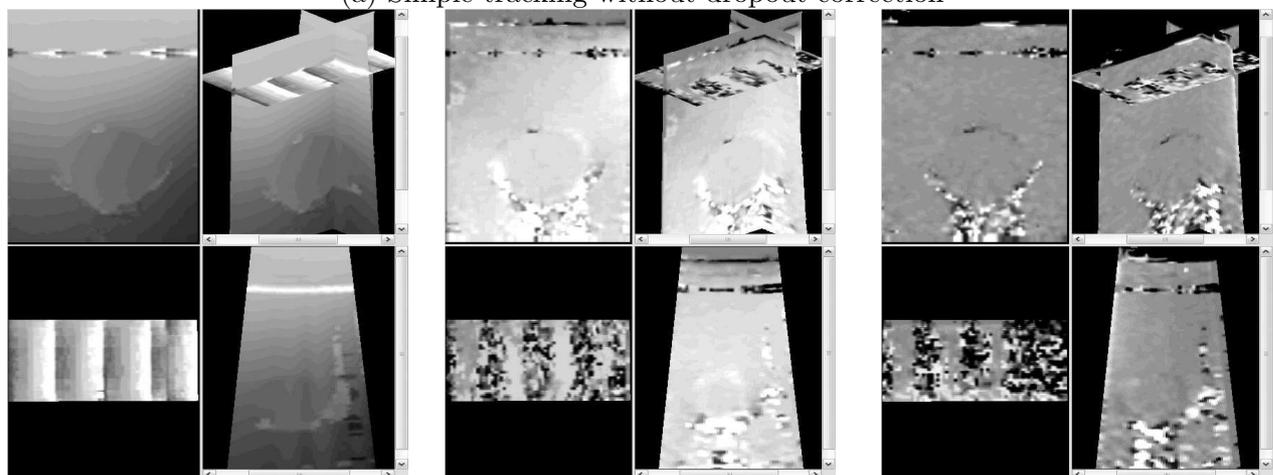
After the second pass, there are occasionally a few isolated incorrect displacements remaining, close to the initialisation point, which happen to have higher correlations away from the correct displacements. These are easily eliminated by comparing local values at the initialisation point with their neighbours and replacing any which are more than half a cycle away from the surrounding average with that average.

It would of course be possible, and in some cases perhaps beneficial, to complete any number of passes, each time alternating the direction of propagation, but in practice we have found that two passes is sufficient for practical scanning.

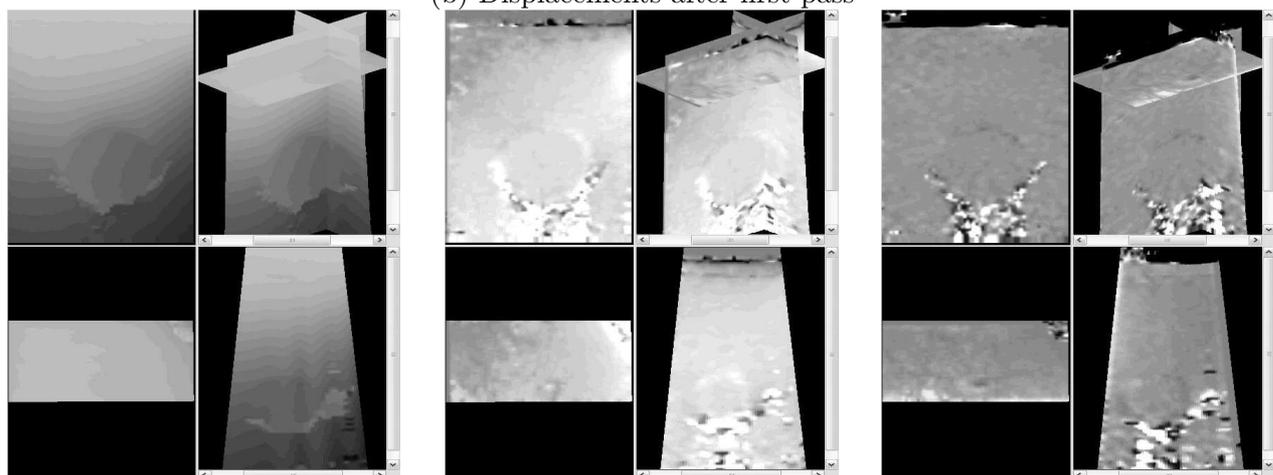
Figure 3 demonstrates simple tracking and the first and second pass of the tracking algorithm we actually use, showing in each case the axial, lateral and elevational displacements. In the simple



(a) Simple tracking without dropout correction



(b) Displacements after first pass



(c) Displacements after second pass

Figure 3: **Tracking of displacements.** Axial displacements are shown on the left, lateral in the centre, and elevational on the right. These are all to the same (normalised) grey scale, save that we take the absolute value for axial displacements, black representing zero, whereas lateral and elevational displacements are signed, with zero as mid-grey. Each set of data contains three orthogonal planes and a 3D composite image. Top left is aligned with the B-scan, bottom left with a C-scan (located at the point of tracking initialisation), and bottom right orthogonal to both.

scheme, the search starts at the top of each RF vector with a zero displacement, and subsequent estimates are initialised with the previous values. In this case, this scheme fails completely, due to the effect of the scanning window behind which the 2D probe is being swept. After the first pass of the more sophisticated scheme, the bottom left plane (aligned with a C-scan) shows the distribution of initial estimates. With a tracking window ± 6 vectors and frames each side of the current displacement, the correct estimates are selected and propagated within a few windows of the initialisation. The second pass entirely removes the residual incorrect starting estimates. In this case, there was considerable lateral movement (of the order of a few vectors) between volumes, but very little elevational movement.

2.3 AMC and interpolation to a regular grid

The estimates of displacement between volumes at each location, and any other value e.g. strain precision, are all based on matching 3D windows of data between the volumes. However, if the strain is non-zero, the displacement will not in fact be the same over these windows, and hence the actual recorded value will only match the real displacement at some particular point in the window. Often, this is assumed to be at the centre of the window, but due to the large amplitude variations inherent in ultrasound signals, this is often a poor assumption. Fortunately, it is possible to compute a good estimation of the correct location using Amplitude Modulation Correction, or AMC (Lindop et al., 2007a). In this context, AMC involves the calculation of a weighted sum of the signal envelopes over the 3D window — this can be done at very little computational expense, in all three directions, once the optimal displacements have been calculated.

Hence for each window, we not only have the axial, lateral and elevational displacements and the final precision values, we also have the axial, lateral and elevational location at which these displacements best match the underlying displacement field. This presents us with a problem, in that although the 3D windows are arranged on a regular grid, when we place our displacement estimates at the ideal locations, they are no longer arranged regularly (though they will maintain the same ordering in each direction — imagine an unevenly stretched fishing net). We need a regular grid in order to display the data, so we must interpolate the displacements.

Optimal interpolation could be achieved by a 3D spline interpolant, however the processing time would far exceed the improvement in accuracy this affords. Fortunately, it is possible to generate acceptable results very quickly using a much simpler linear interpolant. Figure 4 shows, by way of a 2D example, how we achieve this. Essentially, we interpolate in each direction in turn, starting by moving not just the displacement estimates and precisions, but also the other location estimates, to the correct locations. Then we interpolate all of these to a regular grid in one direction. We then discard the location estimates in this direction, and repeat the process in another direction. Finally, we are left with just the displacements and precisions arranged on a regular grid, the location estimates having been used in interpolating to this grid.

2.4 Gradient, normalisation and display

We use a weighted least-squares formulation to calculate axial strain, i.e. the axial gradient of the axial displacement. For completeness, we briefly record how this can be done very efficiently by noting that this form of regression essentially involves the calculation of a variety of summations over 3D kernels. Taking the sum of axial displacements d as an example, we therefore need *at each estimation location* terms similar to:

$$\sum_{(i,j,k) \in \mathcal{K}_n} d_{(i,j,k)} \quad (7)$$

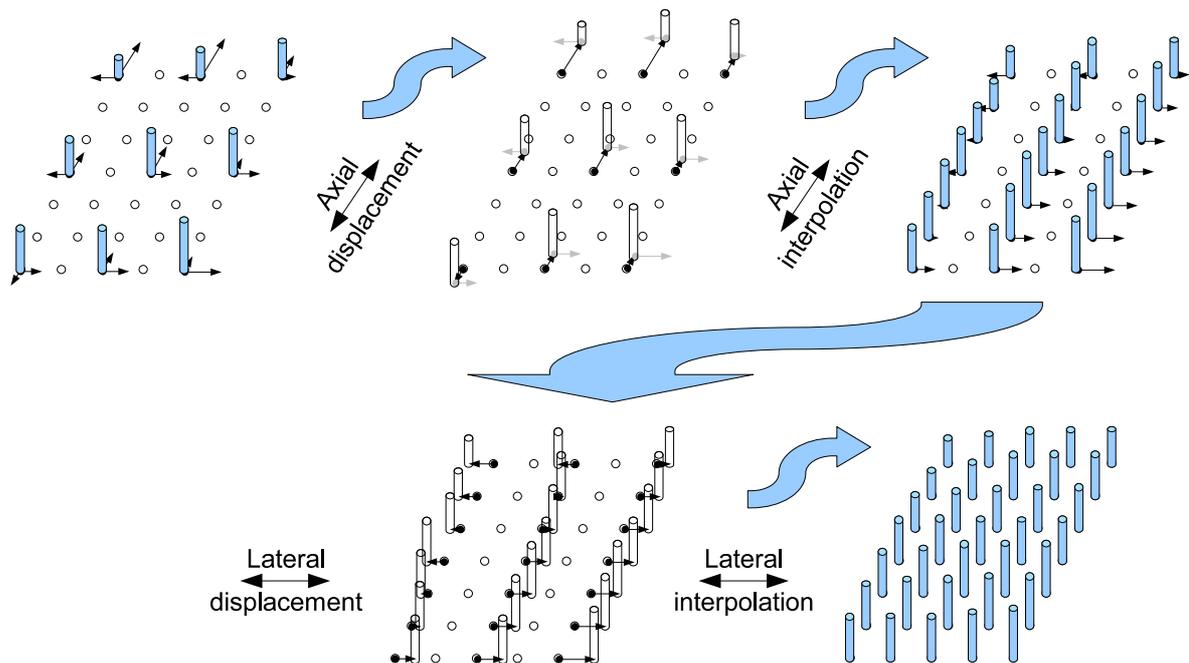


Figure 4: **Interpolation to a regular grid.** For clarity, the figure shows a 2D example: the same principle applies to the 3D data. Axial, lateral and elevational displacements and precision estimates are calculated at each initial window location. In addition, we can calculate where in the data these estimates are most applicable to subsample accuracy, using AMC. Each of the displacement and precision values (shown as vertical cylinders) must be interpolated to a regular grid using the AMC location estimated (shown as arrows). First the values *and* lateral AMC estimates are moved to the correct axial location using the axial AMC estimate, then these are linearly interpolated axially. Next these interpolated values are moved to the correct lateral location using the interpolated lateral AMC estimates, before finally interpolating laterally.

where \mathcal{K}_n is a 3D kernel centred on (x, y, z) with extent $\pm(n_x, n_y, n_z)$. This can be a very expensive calculation, particularly if the kernel is large and we require many independent estimates such that kernels overlap to a significant degree. Instead of calculating such sums directly, we first accumulate them in each direction, e.g.:

$$s_{d(x,y,z)} = \sum_{k=0}^z \left\{ \sum_{j=0}^y \left\{ \sum_{i=0}^x d_{(i,j,k)} \right\} \right\} \quad (8)$$

s_d , and the equivalent for whichever other sums are required in the regression, is stored at every location where we require an estimate of gradient. Having done so, the calculations *at each estimation location* are reduced to the sum of eight values at the corners of the 3D kernel:

$$\begin{aligned} \sum_{(i,j,k) \in \mathcal{K}_n} d_{(i,j,k)} &= s_{d(x+n_x, y-n_y, z-n_z)} + s_{d(x-n_x, y+n_y, z-n_z)} + s_{d(x-n_x, y-n_y, z+n_z)} \\ &+ s_{d(x+n_x, y+n_y, z+n_z)} - s_{d(x+n_x, y+n_y, z-n_z)} - s_{d(x+n_x, y-n_y, z+n_x)} \\ &- s_{d(x-n_x, y+n_y, z+n_z)} - s_{d(x-n_x, y-n_y, z-n_z)} \end{aligned} \quad (9)$$

We can then perform least-squares gradient estimation over a 3D kernel centred on every display voxel, with arbitrarily large 3D kernels, since each estimate has almost no processing overhead, and neither Equation (8) nor Equation (9) are dependent on the size of the kernel.

This technique is expensive in terms of memory consumption, since we require the storage of several sums at every display voxel, and they must be stored at double-precision so that we can detect small differences between very large values. However, it speeds up computation in any case where the 3D kernels at each estimation location overlap, approximately by the ratio of overlap to non-overlap volume.

Since we are moving the ultrasound probe manually, the applied stress will vary between volumes, and hence the absolute strain will also vary. However, we want a consistent image which shows the same relative stiffness independent of how the anatomy was scanned. For normalisation we fit some function to the entire 3D strain data set, and locally normalise the strain estimates using this function. The aim here is to approximate the stress field so we generate normalised strain images which are more related to tissue properties and less sensitive to the applied stress. The simplest function to fit is simply a constant. A better function for the normalisation value s_n which we use in this case is:

$$s_n = a(1 + by)(1 + cx + dz) \quad (10)$$

where a , b , c and d are all constants, which we estimate by a least-squares fit to the entire strain field. This allows for two important variations: firstly for a gradually decaying stress with depth, and secondly for a non-constant applied pressure across the surface of the probe face. Both of these are common features in strain imaging using freehand movement of the ultrasound probe.

Implicit in this approach is the assumption that likely variations in the applied stress field are sufficiently different from likely variations in stiffness due to anatomical variation. It is possible, though unlikely, that real tissue might have a stiffness which gradually increases with depth, where a stress field can nevertheless be applied which is constant. In this case normalisation by Equation (10) would remove real stiffness variation from the strain image, and some other normalisation function may be more appropriate.

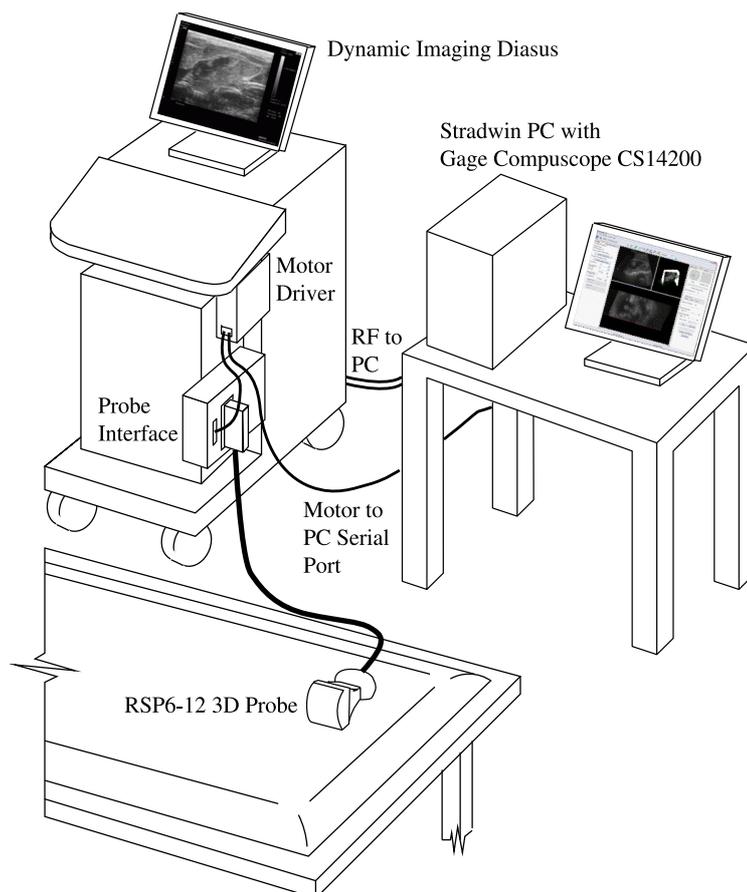


Figure 5: **Schematic of 3D RF system.** A 3D probe is interfaced to a Dynamic Imaging Diasus ultrasound machine. The analogue RF data from the Diasus is passed to the computer via a Gage Compuscope digitiser. The motor on the probe is driven by a controller, which responds to serial commands from the PC. The computer can therefore control and synchronise the motor and the RF ultrasound signals in real time.

3 Experiments and results

In order to acquire 3D RF data in real time, we used a Dynamic Imaging Diasus ultrasound machine ², with an adaptor allowing an interface to a 3D probe, GE RSP 6-12 ³. The ultrasound machine interfaced to the crystals on the probe, not the stepper motor. RF data was acquired using a Gage Compuscope 14200 ⁴ in a PC running Stradwin software ⁵. The serial port of this PC was connected to a motor controller which was interfaced to the stepper motor and top-dead-centre positioning switch for the sweeping probe. Hence both the RF acquisition and the stepper motor could be driven synchronously under software control.

Several modes of operation are available. In frame-at-once, the motor is moved to a specific location, one frame of RF data acquired and downloaded to the PC from the Gage card, and then the motor moved on to the next location. Roughly 15 frames per second can be recorded in this mode. In volume-at-once mode, both the motor and the RF acquisition are started synchronously, and a set number of frames covering the volume are acquired during motor movement. The position of these frames can be determined retrospectively by carefully timestamping each frame and comparing with the known probe motion. Then the frames can be downloaded from the Gage card while the motor is set back to its initial position, before the whole process is repeated. Roughly one volume per second can be recorded in this mode. Finally, in multiple-volume mode, a set number of volumes can be recorded with the probe sweeping to and fro and all positional information deduced at the end from the frame timestamps. In this way, roughly 7 volumes can be acquired per second. This speed is limited by only having one set of transmit and receive lines, whilst in practice 3D probes commonly operate with multiple receivers in parallel.

For the 3D strain data, the volume-at-once mode was found to be ideal, since each volume is acquired fairly quickly (less than 0.2 seconds), thus minimising anatomical movement. The pause between volumes of about 1 second was enough to apply a very slightly greater pressure to the probe, then hold it still again for the second volume acquisition. Hence two volumes of data (enough for the strain images included in this report) could be acquired in a little less than 2 seconds.

127 vectors of RF data with a centre frequency of approximately 6.5 MHz and a 4cm depth were sampled at 66.67 MHz, and converted to baseband analytic signals sub-sampled at 16.67 MHz. Equal aspect ratio 3D windows were used for displacement estimation, with an axial extent equivalent to 12 cycles of the ultrasound centre frequency and no overlap, giving roughly 80 windows over the entire axial depth. The least squares gradient operation was performed with an equal aspect 3D kernel of 16 pixels extent axially, centred on each of the 400 display pixels in the axial direction. Weighted mean filtering following (Lindop et al., 2007b) was also performed with an equal aspect ratio 3D kernel of extent 8 display pixels.

The software could display 2D strain images live, by comparing similar frames in the two volumes, which gave an immediate indication of data quality. 3D strain images were calculated in a post processing step, lasting typically between 20 and 30 seconds. On average at least every second acquisition was found to give a satisfactory strain image.

Strain images were calculated for a breast phantom, an olive in agar, and *in vivo* in a human forearm. The processing times (measured on a 2.1 GHz Core 2 Duo processor with the software running in a single thread) are contained in Table 1, in addition to the number of frames in each volume.

²Dynamic Imaging Ltd., Livingston, UK

³GE Healthcare, Chalfont St Giles, UK

⁴Gage Applied, Lockport, IL, USA

⁵<http://mi.eng.cam.ac.uk/~rwp/stradwin>

Table 1: **Processing of 3D strain data.** The table shows the number of frames in each of the two volumes, and the total processing time.

Data	Frames per volume	Processing time / secs
Breast phantom, 15 mm inclusion	60	17
Breast phantom, 6.5 mm inclusion	60	17
Breast phantom, mixed inclusions	60	21
Breast phantom, 1.6 mm and mixed inclusions	60	21
Olive in agar	120	32
Human forearm	50	8

3.1 In vitro strain images

3D strain data was acquired by scanning various targets in a CIRS breast biopsy phantom (Model 052)⁶. Strain images were produced with and without lateral and elevational tracking. Figure 6 shows a medium sized stiff inclusion and Figure 7 a small stiff inclusion. Figure 8 shows a conventional and strain image from a stiff inclusion next to a fluid filled cyst, demonstrating the performance of the system in regions where no strain data can be calculated (since there is very little ultrasound signal from the inside of the cyst). Figure 9 shows a conventional and strain image which contains a very small 1.5 mm diameter object near the centre which is not obvious in the B-scan, but appears as a black (and hence soft) region in the strain image.

Strain data was also calculated for an inhomogeneous agar phantom containing half an olive, which is slightly stiffer than the agar. This is clearly a more complex geometry than the spherical inclusions. Figure 10 shows the results for various tracking options.

3.2 In vivo strain images

Figure 11 demonstrates that the technique can also generate strain images for *in vivo* data, in this case of soft tissues in the forearm.

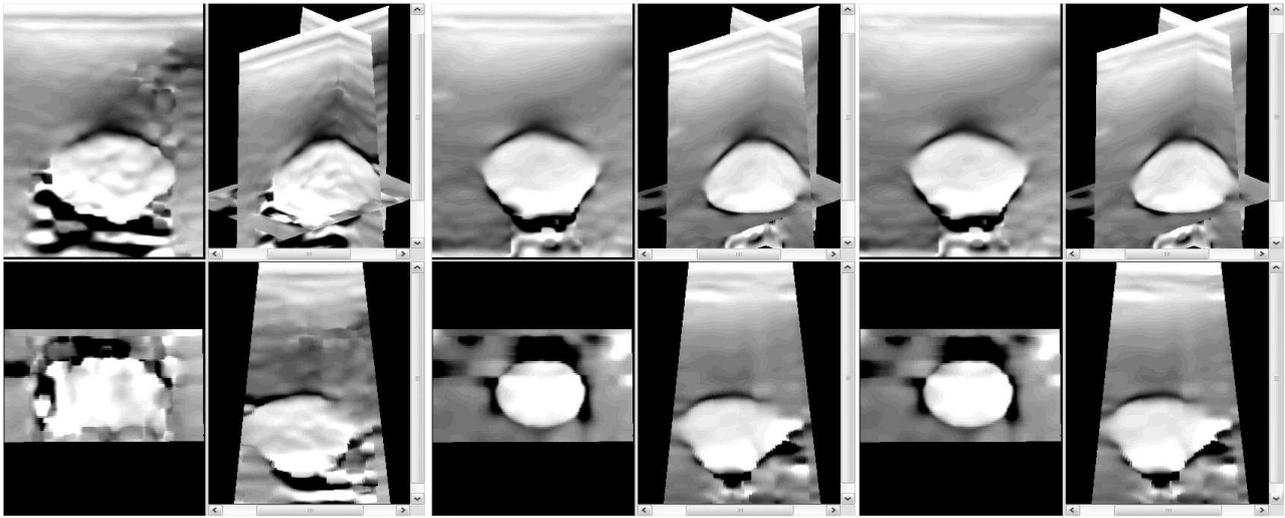
4 Discussion

4.1 Processing and acquisition time

It is evident from the processing times contained in Table 1 that 3D axial normalised strain imaging is possible with a very brief acquisition time and only a short wait for processing. This is despite employing an algorithm using very high accuracy phase tracking and AMC, which has been demonstrated to perform as well as any other strain algorithm, particularly at low strains (Lindop et al., 2007a). The two acquired volumes are sufficiently well correlated to generate a strain field which is well distributed over the volume and produces 3D strain data with uniform 3D properties.

This is partly due to the adoption of very fast iterative algorithms, and partly due to the use of a displacement tracking scheme which allows for short range searches which nevertheless converge on the correct displacements. Approximately two thirds of the processing time is devoted to the

⁶Computerized Imaging Reference Systems, Norfolk, Virginia, USA.

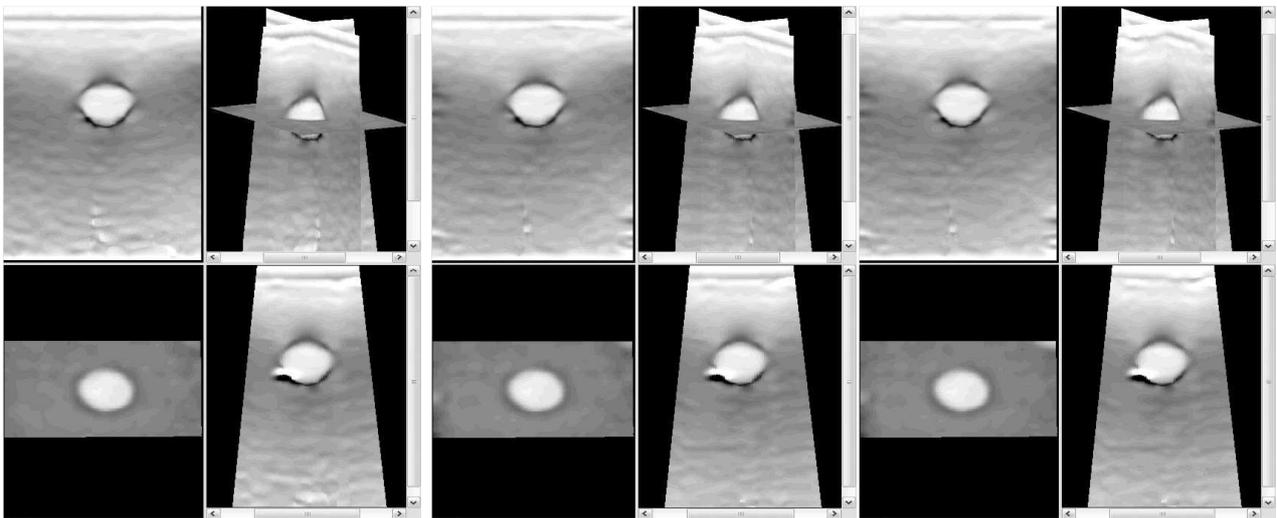


(a) No lateral or elevational

(b) Lateral

(c) Lateral and elevational

Figure 6: **Strain images for a 15 mm diameter stiff inclusion in a breast biopsy phantom.** (a) shows axial strain without lateral and elevational tracking, (b) shows axial strain with lateral tracking, and (c) also with elevational tracking.



(a) No lateral or elevational

(b) Lateral

(c) Lateral and elevational

Figure 7: **Strain images for a 6.5 mm diameter inclusion in a breast biopsy phantom.** (a) shows axial strain without lateral and elevational tracking, (b) shows axial strain with lateral tracking, and (c) also with elevational tracking.

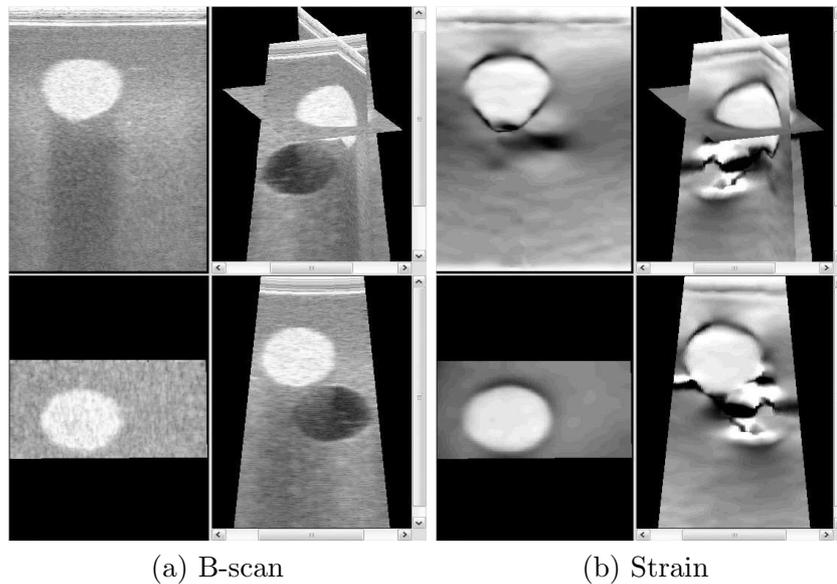


Figure 8: **B-scan and strain images for a hard inclusion near a cyst in a breast biopsy phantom.** (a) shows a conventional 3D ultrasound image, and (b) the equivalent strain image.

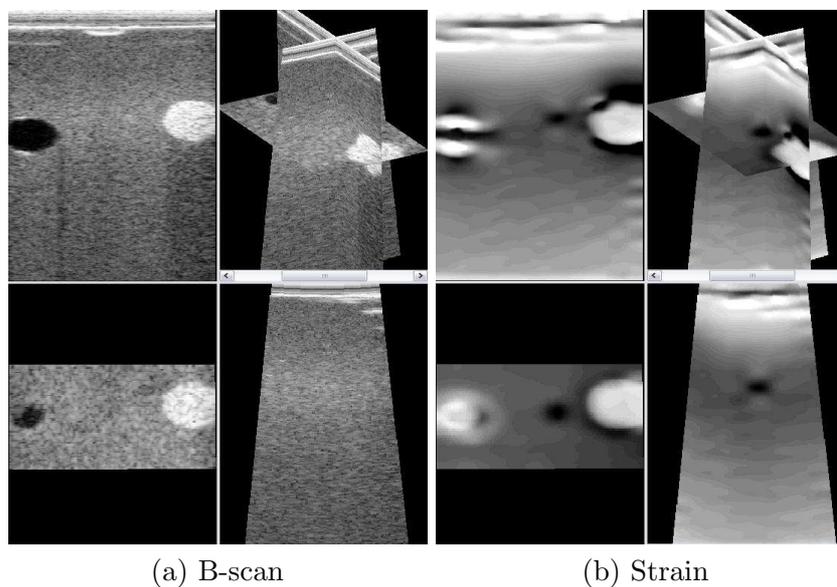


Figure 9: **B-scan and strain images for a small soft inclusion near a larger hard inclusion in a breast biopsy phantom.** (a) shows a conventional 3D ultrasound image, and (b) the equivalent strain image.

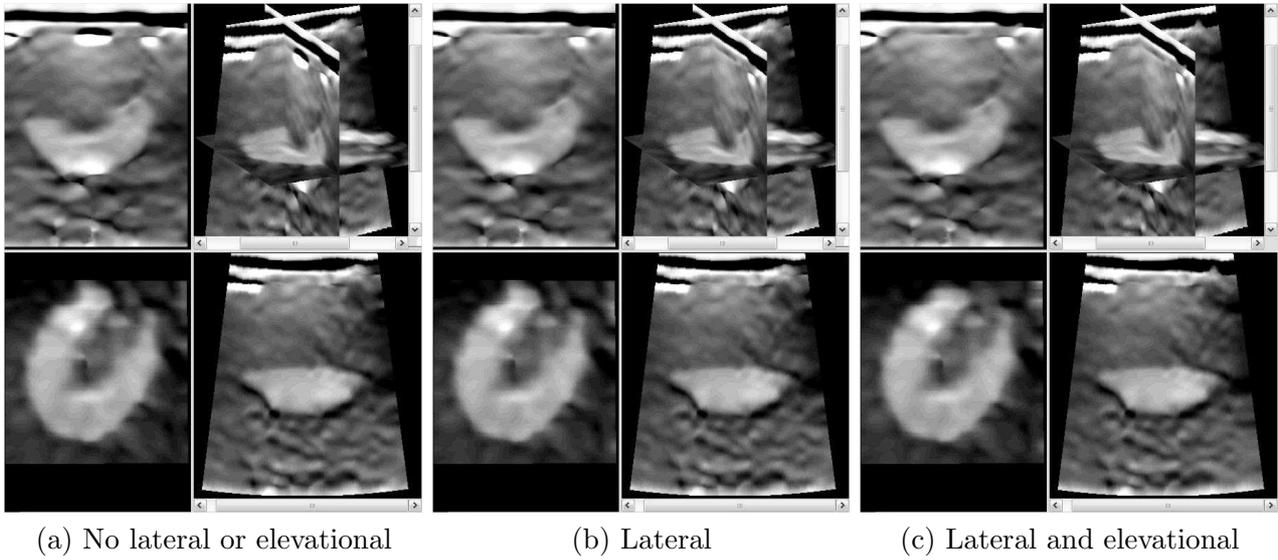


Figure 10: **Strain image of a half olive in agar.**

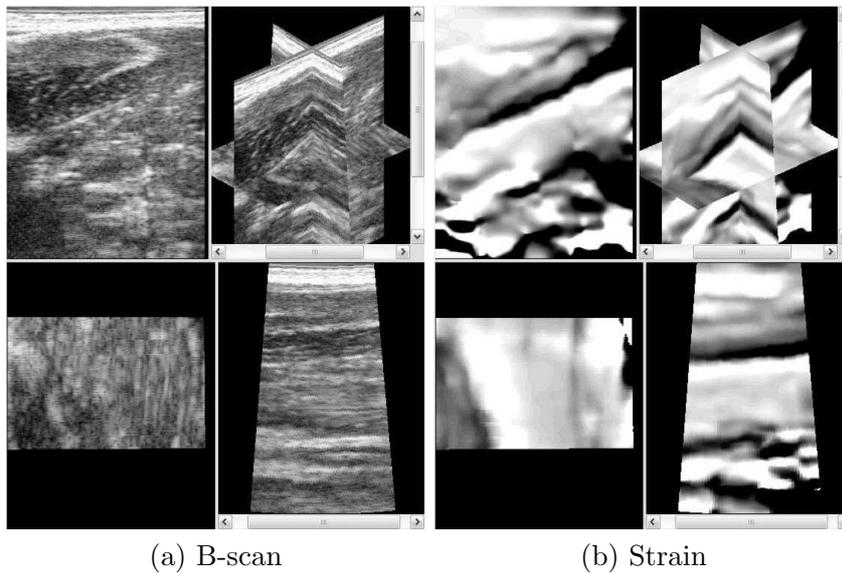


Figure 11: **B-scan and strain for forearm.**

displacement estimation algorithm, of which about 30% is for non-axial tracking. The next most demanding stage is conversion of the input RF to an analytic baseband signal: hence if a sampled IQ signal was available, the processing times would improve substantially.

4.2 Lateral and elevational tracking

In our experience with freehand 2D strain imaging, tracking lateral displacements is not critical when calculating axial strains. In this context, the frame rate can be very high (upwards of 20 Hz), the strains are very low (much less than 1%) and hence the displacements are correspondingly small. It is rare in this situation to have lateral movement of more than one RF vector due to the movement of the probe. More major anatomical movement, for instance imaging of a beating heart, does produce significant movement in lateral directions, but this corrupts the strain image, since the major stress direction is no longer axial.

With 3D data, the two volumes are acquired a second apart, and the probe face is much larger, both of which tend to lead to more movement in lateral and elevational directions. In this case lateral and elevational tracking are more important, though still not critical: in general the best strain image will be produced when there is little movement of this type, and when there is such movement the strain image will be degraded even if the movement is tracked.

Figure 3 shows a situation where lateral displacement (centre column) is significant. The lateral displacement is estimated with a lower accuracy than axial displacement, but a clear lateral strain field is apparent. Figure 6 shows the subsequent strain images from the same data, with and without tracking in the lateral and elevational directions. Tracking the lateral strain field (b) does indeed improve the accuracy of the strain images visibly, though the strain image (a) without lateral tracking is still meaningful.

Figures 7 and 10 show perhaps more typical situations where the non-axial tracking has a more limited, though still positive, effect. In the former, the lateral tracking makes the lower region of the images more homogeneous, in the latter elevational tracking improves the definition of the top left corner of the olive in the bottom left image. Nevertheless, the small overhead in processing time would still seem worthwhile, and it must be noted that all our images are at relatively shallow depths: deeper data is likely to suffer more from such displacements.

4.3 System performance

We make no attempt at a rigorous assessment of system performance here, since our aim is primarily to demonstrate that good quality 3D strain imaging is achievable at very acceptable rates. Nevertheless, a couple of important points can be highlighted from the images. Firstly, Figure 7 demonstrates that we can very easily resolve a stiff inclusion of only 6.5 mm diameter. Figure 9 contains an even smaller 1.6 mm diameter soft region which is clearly highlighted in the strain image despite the lack of contrast in the B-scan.

Secondly, Figure 8 demonstrates that the presence of cysts, which contain little RF signal and hence do not provide strain data, does not obstruct accurate strain estimates in all other parts of the image. This is obviously a key result, since we expect most anatomy to be full of features both with varying stiffness and echogenicity, as well as ultrasound artefacts such as shadowing and reverberations, which will locally corrupt the strain estimates. We demonstrate in Figure 11 that we can indeed provide meaningful strain images *in vivo* which are clearly correlated to the actual anatomical structure.

5 Conclusions

We have demonstrated high quality 3D strain images, using 3D windows for displacement estimation, least squares gradient calculation and filtering, with acceptable acquisition and processing times. It is quite possible to acquire two sequential volumes of data with a hand held 3D probe which are aligned sufficiently well to provide good signal correlation and have sufficient axial strain variation to produce good strain estimates. Lateral and elevational tracking can be included at a small additional processing cost and enable strain calculations over a larger range of probe movement.

Acknowledgements

Graham Treece is supported by an EPSRC/RAEng Postdoctoral Fellowship. Dynamic Imaging Ltd. provided a modified ultrasound machine with an interface to the 3D probe and to the analogue RF signals.

References

- Anand, A., Savéry, D., Hall, C., Jan. 2007. Three-dimensional spatial and temporal temperature imaging in gel phantoms using backscattered ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 54 (1), 23–31.
- de Korte, C. L., van der Steen, A. F. W., Céspedes, E. I., Pasterkamp, G., 1998. Intravascular ultrasound elastography in human arteries: experience *in vitro*. *Ultrasound in Medicine and Biology* 24 (3), 401–408.
- Egorov, V., Ayrapetyan, S., Sarvazyan, A. P., Oct. 2006. Prostate mechanical imaging: 3-D image composition and feature calculations. *IEEE Transactions on Medical Imaging* 25 (10), 1329–1340.
- Fisher, T. G., Hall, T. J., Panda, S., Jang, J., Resnick, J., Barnes, S., Madsen, E. L., Oct. 2006. Volume elasticity imaging with a 2-D capacitive micro-machined ultrasound transducer (CMUT) array. In: *Proceedings of the 5th International Conference on the Ultrasonic Measurement of Tissue Elasticity*. Snowbird, Utah, USA, p. 125.
- Garra, B. S., Céspedes, E. I., Ophir, J., Spratt, S. R., Zuurbier, R. A., Marnant, C. M., Pennanen, M. F., Jan. 1997. Elastography of breast lesions: initial clinical results. *Radiology* 202 (1), 79–86.
- Gennisson, J., Baldeweck, T., Tanter, M., Catherine, S., Fink, M., Sandrin, L., Cornillon, C., Querleux, B., Aug. 2004. Assessment of elastic parameters of human skin using dynamic elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 51 (8), 980–989.
- Krueger, M., Pesavento, A., Ermert, H., Hiltawsky, K. M., Heuser, L., Rosenthal, H., Jensen, A., 1998. Ultrasonic strain imaging of the female breast using phase root seeking and three-dimensional "optical flow". In: *IEEE Ultrasonics Symposium*. pp. 1757–1760.
- Li, Y., Patil, A., Hossack, J. A., Feb. 2006. High resolution three-dimensional prostate ultrasound imaging. In: *Medical Imaging 2006: Ultrasonic Imaging and Signal Processing*. Vol. 6147 of *Proceedings of SPIE*. San Diego, California, pp. 32–40.
- Lindop, J. E., Treece, G. M., Gee, A. H., Prager, R. W., Apr. 2006. 3D elastography using freehand ultrasound. *Ultrasound in Medicine and Biology* 37 (4), 529–545.

- Lindop, J. E., Treece, G. M., Gee, A. H., Prager, R. W., 2007a. Estimation of displacement location for enhanced strain imaging. In press for IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control .
- Lindop, J. E., Treece, G. M., Gee, A. H., Prager, R. W., May 2007b. An intelligent interface for free-hand strain imaging. Tech. Rep. CUED/F-INFENG/TR 578, Cambridge University Department of Engineering.
- Lorenz, A., Pesavento, A., Pesavento, M., Ermert, H., 1999. Three-dimensional strain imaging and related strain artifacts using an ultrasonic 3D abdominal probe. In: IEEE Ultrasonics Symposium. pp. 1657–1660.
- Schaar, J. A., de Korte, C. L., Mastik, F., van Damme, L. C. A., Krams, R., Serruys, P. W., van der Steen, A. F. W., 2005. Three-dimensional palpography of human coronary arteries. *Herz* 30 (2), 125–133.
- Taylor, L. S., Porter, B. C., Rubens, D. J., Parker, K. J., 2000. Three-dimensional sonoelastography: principles and practices. *Physics in Medicine and Biology* 45, 1477–1494.
- Treece, G. M., Gee, A. H., Prager, R. W., Apr. 2005. RF and amplitude-based probe pressure correction for 3D ultrasound. *Ultrasound in Medicine and Biology* 31 (4), 493–503.
- Treece, G. M., Lindop, J. E., Gee, A. H., Prager, R. W., Oct. 2006. Efficient elimination of dropouts in displacement tracking. In: Proceedings of Ultrasonic Measurement and Imaging of Tissue Elasticity. Snowbird, Utah, USA, p. 68.
- Treece, G. M., Lindop, J. E., Gee, A. H., Prager, R. W., 2007. Near-real-time 3D ultrasonic strain imaging. In: Akiyama, I. (Ed.), To appear in *Acoustical Imaging*. Vol. 29. Springer, pp. ??–??
- Vogt, M., Ermert, H., Mar. 2005. Development and evaluation of a high-frequency ultrasound-based system for in vivo strain imaging of the skin. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 52 (3), 375–385.