3D ELASTOGRAPHY USING FREEHAND ULTRASOUND

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

CUED/F-INFENG/TR 531

July 2005

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

Email: jel35/gmt11/ahg/rwp @eng.cam.ac.uk

3D ELASTOGRAPHY USING FREEHAND ULTRASOUND

Joel E. Lindop, Graham M. Treece, Andrew H. Gee and Richard W. Prager

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

Abstract

We present an elastography system using freehand 3D ultrasound. A review is provided of the standard elastography methods that have been adapted for this purpose. The scanning protocol is simple and promising results are presented of 3D strain images from freehand scans. Robustness is a problem, however, and the main sources of error are explained. Measures have been developed to improve the quality of the freehand images by means of dropout correction and frame filtering. Results from the application of these techniques provide an indication of development strands which should lead to a system that is both easy-to-use and produces reliable, high quality images.

1 Introduction

Ultrasound elastography seems certain to become an important medical imaging tool. It has already been used in clinical trials for applications such as imaging diseased arteries [1], detecting prostate tumours [2] and categorising breast lesions [3, 4]. The most common approach is static compression elastography, where frames of data are recorded before and after a controlled compression has been applied to the tissue. *Quasistatic* compression elastography refers to variants in which the probe/compressor is still moving at the moment when data is acquired. Cross-correlation analysis is used to track tissue displacements from which the strain field can be estimated. The technique has been in development for over a decade [5], and can produce elastograms (strain images) with high SNR. Its most commonly cited application is the investigation of stiff tumours in soft tissues, where it has been shown that elasticity measurements could be useful for both detection and categorisation purposes [6].

Against this, static compression elastography presents several challenges which may partly explain why it has yet to be widely adopted into clinical practice. Firstly, it is usual during scanning that the probe or a separate compressing plate is controlled by a mechanical actuator to produce precise deformation of the tissue. This restricts the flexibility with which the technique may be applied. A second consideration with any imaging system is the delay after scanning before elastograms are ready for inspection. Increases in the speed of modern processors have helped to reduce computation times, but the processing time for an individual 2D elastogram may be anything from tens of milliseconds [2] up to minutes depending on the number of pixels required, the sampling rate, the choice of algorithm, and the numerous algorithmic parameters governing trade-offs between speed, accuracy and robustness. Fast processing is crucial for feasible 3D elastography, since to span a volume a much larger number of strain estimates must be produced. 3D imaging confers many advantages, such as improved accuracy of volume estimation and the possibility of viewing planes that are usually inaccessible; a comprehensive discussion is provided by Gee *et al.* [7].

There are two main alternatives to static compression elastography, each with advantages in certain situations. Firstly, the term sonoelasticity refers to motion detection when the tissue is excited by a vibrating actuator [8, 9]. It is straightforward to view sonoelastograms in real time because the signal processing is essentially equivalent to well established Doppler methods, and the images are useful since tissue's dynamic behaviour is a function of its stiffness field. The other alternative is radiation force elastography, which has arisen more recently: an ultrasound beam

of high intensity exerts a force at its focus inside the tissue; this is used either in a quasistatic mode [10] or as a means of inducing shear waves for a vibration analysis [11]. Radiation force elastography may confer advantages for inverse problem approaches to stiffness estimation, since boundary conditions are less important with the internally applied force. Nevertheless, for our system quasistatic compression was chosen ahead of sonoelasticity and radiation force because the elastograms are easier to interpret. Sonoelastograms often contain complicated vibration modes; quasistatic radiation force elastography is likely to require that the probe is held stationary for several milliseconds while each 2D elastogram is acquired, so it may be incompatible with freehand scanning. Meanwhile, it is often reasonable to interpret quasistatic compression elastograms as inverse stiffness images. The interpretation is strictly incorrect, because the applied stress field is generally inhomogeneous, especially in a freehand scan, but it will be seen from the results in this paper that the consequent artefacts are not severe.

The focus of the work presented is on developing quasistatic elastography techniques for use within an existing freehand 3D ultrasound system [12, 13]. The attraction of a freehand implementation is twofold. It confers superior ease of use and versatility, so that interactive scans can be carried out by a skilled practitioner to locate features of interest and investigate them in whatever way is most useful. Freehand technology also reduces the requirement for additional hardware, so elastography could be incorporated in a commercial system at minimal cost (at present 3D scanning necessitates the use of a position sensor, but in time this may be replaced by accurate sensorless position estimation [14]). These advantages are especially pronounced in 3D ultrasound, where the majority of systems scan volumes using either oscillating head or 2D phased array transducers, which increase the size of the probe footprint. Instead, freehand 3D ultrasound uses a normal 2D probe and the volume is populated with data from 2D scans where the locations of the planes have been recorded in 3D.

Freehand elastography has attracted a lot of interest in recent years. Doyley *et al.* [15] have shown that it is possible off-line to produce freehand elastograms of good quality by training an operator to move the freehand probe upon the region of interest in a compressive direction at a prescribed strain rate. Typically though, freehand systems require real-time imaging to guide the operator, which was a serious challenge for earlier research. For example, Bamber *et al.* [16] looked at a minimally intensive strain estimator based on speckle decorrelation, but increased speed came at the expense of reduced accuracy. Subsequent work by Pesavento *et al.* [17], however, produced a fast phase-based algorithm which was at least as accurate as conventional cross-correlation techniques. Hall *et al.* [4] have documented a system which exploits real-time operation by running 2D elastogram and B-scan displays side by side, providing feedback to the operator so that it is easier to practise a successful scanning technique.

It was anticipated that the main difficulty with extending freehand elastography to 3D would arise from the requirement that the probe be translated in the elevational direction to sweep out a volume. This means that pairs of consecutive ultrasound frames are usually non-coplanar, so the level of signal decorrelation is increased. Strain estimates in these circumstances are at best less accurate and at worst entirely erroneous. A more general problem is that whereas visual feedback may aid an operator in producing several successful freehand 2D elastograms, a 3D data set comprises a collection of many 2D images, where it is perhaps inevitable that some of the planes are dominated by strain estimation errors. For this reason, correction techniques are an important feature of the freehand 3D elastography system that is documented in this paper. Section 2 provides background information, with a review of tools from the literature that are reapplied here; Section 3 outlines the freehand 3D elastography system and the results that have been achieved; Section 4 documents the development of suitable correction techniques; finally some conclusions are drawn in Section 5.

2 Background

The elastography method incorporated in our 3D system is based on quasistatic techniques that have received a lot of attention since being proposed originally by Ophir *et al.* [5]. Two frames

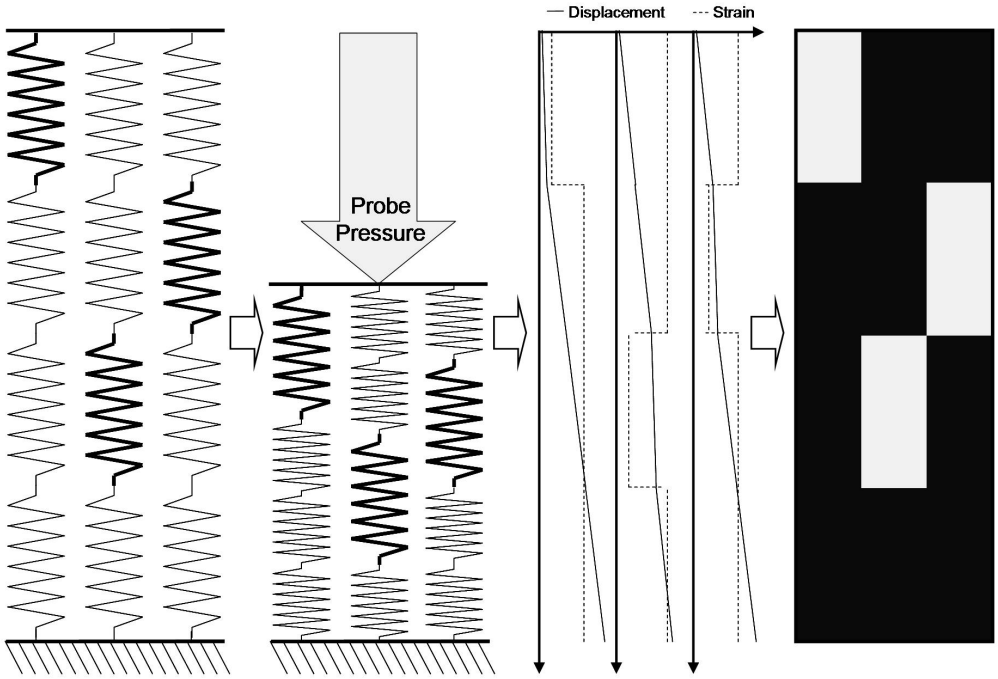

Figure 1: **Tissue modelled as a collection of elastic springs:** compressive pressure is applied and if the displacement field can be estimated, its derivative is strain, which may highlight salient tissue features.

of ultrasound data are recorded: one before and one after a section of tissue is compressed (see illustration in Figure 1). Tissue displacements give rise to shifts in the time-delays of corresponding sections in the recorded ultrasound signals. The shifts can be estimated by windowing data in the pre- and post-compression signals, and identifying the temporal displacement that produces the closest match between the windows — usually this is determined by locating the maximum in the normalised cross-correlation function. The temporal displacements correspond closely to mechanical tissue displacements, assuming that variations in the speed of sound are small, so the local gradient of the displacement estimates is used to estimate strain. By this method, a column of strain estimates is produced for each A-line in the recorded ultrasound signal. These are mapped to pixel intensities in the resultant 2D elastogram. We display low strains as bright regions (corresponding to stiff inclusions) and darker patches indicate higher strains (softer tissue).

Ultrasound is an excellent imaging modality for elastography. This is not only on grounds of low cost, but also because ultrasound signals exhibit variations due to microscopic tissue features, so even in fairly homogeneous tissue each piece has a unique speckle pattern. For this reason time-delay estimation can produce highly accurate estimates of the actual tissue motion. A distinguishing feature between the many different implementations is the method by which the tissue is compressed. Some systems have used a stationary probe, and relied on pressure variations from within the tissue caused by normal processes such as the cardiac cycle. However, a more common approach is the application of pressure through the probe, which is pressed into the tissue in precisely controlled increments using a stepper motor [3, 5]. The freehand scanning approach that will be described here falls into the separate quasistatic compression category because the probe is actually in continuous motion. The best interpretation of the resulting elastograms is still to assume that they show equilibrium strain states. However, it should be acknowledged that image interpretation may be more difficult in tissues where strain has a strong time dependency.

2.1 Axial strain estimation

The naming convention for a set of 3D axes relative to an ultrasound probe is illustrated in Figure 2. Displacement and strain estimation for elastography is often restricted to the axial direction [5, 17, 18]. The principal reason is that while ultrasound probes have excellent axial resolution, the lateral resolution is far poorer: the rate of decorrelation when tissue moves relative to an ultrasound A-line has been shown by Dickinson and Hill [19] to be roughly an order of magnitude lower for lateral movement, and while axial sample spacing depends on the RF sample rate (typically providing >3000 samples over a scan depth of 4cm) the lateral sampling rate is dictated by the spacing of the piezoelectric crystals on the probe (127 samples over 4cm for our probe).

A degree of lateral and elevational movement is nonetheless inevitable. Firstly, this is because uniaxial stress gives rise to secondary strains in the perpendicular directions, where Poisson's Ratio is the material property that measures this tendency. A second cause of non-axial motion is stress concentrations, a feature of inhomogeneous tissues which means that some regions bear non-axial stresses even if the applied pressure is entirely uniaxial. These sources of decorrelation are common to any quasistatic compression system, though freehand elastography is further compromised by the handheld probe: rotations and translations in five of the six degrees of freedom are sources of additional decorrelation. When the operator applies axial pressure, an unintended non-axial component is unavoidable. Furthermore, in freehand 3D elastography non-axial motion is a necessary requirement so that a volume may be swept out.

One way of mitigating decorrelation in general is rapid data acquisition. Reducing the temporal spacing of consecutive data frames lowers the between-frame motion in every direction: lateral, elevational and axial. The correlation can thus be improved, but it comes at the expense of a reduction in the axial strain that is to be estimated. If the pre- and post-compression frames are uncorrelated, the SNR will be tiny even if the axial strain is large. On the other hand, a correlation of 100% acquired with zero axial strain also yields zero SNR because there is no elastographic signal. Therefore, a suitable protocol must be devised to operate somewhere between these extremes.

It has become popular to regard elastography systems as strain filters, with passbands and stopbands as described by Varghese *et al.* [20]. Models developed by Varghese *et al.* incorporated the effects of electronic noise and signal decorrelation due to axial strain, and they predicted passbands at 1–10% strain. However, the same models do not apply to freehand elastography because of the additional decorrelation when axial strain is accompanied by lateral and elevational motion. This probably means that the freehand elastography passband is shifted to lower strains. Doyley *et al.* [15] acknowledged this problem in the design of their freehand 2D elastography system, but they also assumed that freehand palpation would necessarily involve strains of >2%, so they focused on restricting the extent of non-axial freehand motion. Since non-axial motion is unavoidable in a 3D scan, our approach is necessarily different. The focus is on identifying the low strain range at which the best freehand results are produced.

2.1.1 RF signal processing

The raw signals from the transducers on an ultrasound probe are in the radio frequency range. Normal ultrasound scanners produce B-scan amplitude images by detecting the signal envelope and applying a non-linear scaling to determine the intensities of the display pixels. Therefore, the envelope signal is the standard output from commercial ultrasound machines in the form of "video detected data". This is undesirable, however, because it provides only a few hundred samples per A-line. Even if the number of samples can be increased, there are advantages to using the full RF signal. The difference between the RF signal and its envelope is illustrated in Figure 3a. In the past it was common to perform cross-correlation analyses using the signal envelope [19, 21], but most recent elastography has used the full RF signal [5, 15, 17]. If it can be acquired, the RF signal offers much more accurate strain estimation, because the peak of the normalised cross-correlation function is sharper, so it is more robust in the presence of noise. Figure 3b demonstrates this by comparing the auto-correlation of the envelope and RF signals at small shifts of -2/+2 samples. RF data is especially important for achieving an acceptable SNR in freehand 3D elastography,

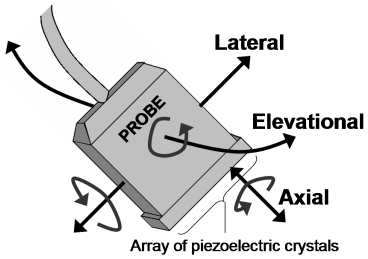

Figure 2: **Six degrees of freedom for the movement of the probe:** the probe is held in the hand and pressed lightly against the object to be scanned. Although the operator moves the probe in the axial (downward) direction, this is accompanied by small motions in the lateral and elevational directions which violate the assumptions of axial displacement tracking.

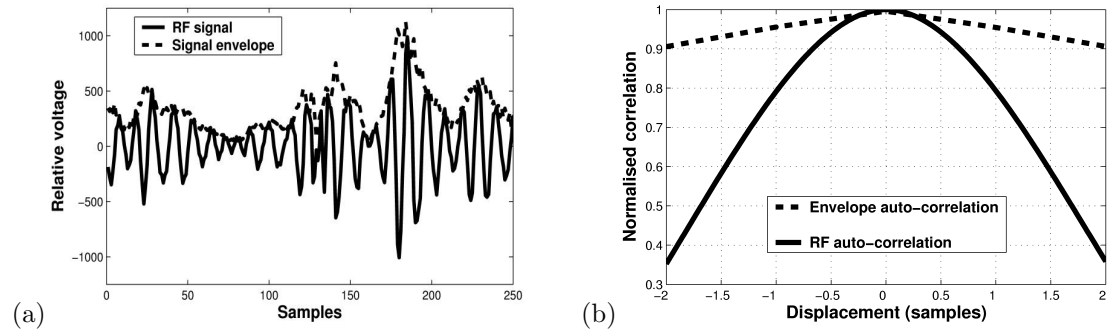

Figure 3: **RF signal versus envelope:** (1) Section from an RF ultrasound signal and the corresponding signal envelope. (2) Cross correlation values for this signal at temporal shifts in the range -2/+2 samples, calculated using both envelope and full RF data.

because the signal (tissue strain) is smaller than in most other systems.

2.1.2 Strain images from speckle tracking

Some form of search is required to estimate the displacements of corresponding data windows in the pre- and post-compression frames, achieving accurate sub-sample precision. This can be accomplished with good accuracy by interpolating the RF signals at sub-sample locations and recalculating the normalised cross-correlation, but an exhaustive search on this basis is extremely computationally expensive. Commonly the analysis is made faster by applying peak interpolation techniques [1, 22, 23] to estimate sub-sample displacements from values of the normalised cross-correlation to either side. Alam *et al.* [24] among others have noted that speckle tracking in this form is suboptimal for elastography, because tissue compression causes both displacements (which are estimated) as well as within-window deformations (which are ignored); Alam *et al.* have devised an alternative technique which addresses this issue, but at the low strains of freehand elastography it is unlikely that the performance improvements would be sufficiently large to warrant the significant increase in the computational overhead.

Speckle tracking is repeated at overlapping window positions spaced regularly along each A-line to produce vectors of displacement estimates. These are converted to strain vectors using the method of least squares (LSQ): a line is fitted through nearby displacements; its gradient is the strain estimate [18]. Note that in the limit when just two displacement estimates are used, the least squares filter is equivalent to taking the difference between consecutive displacements. Afterwards,

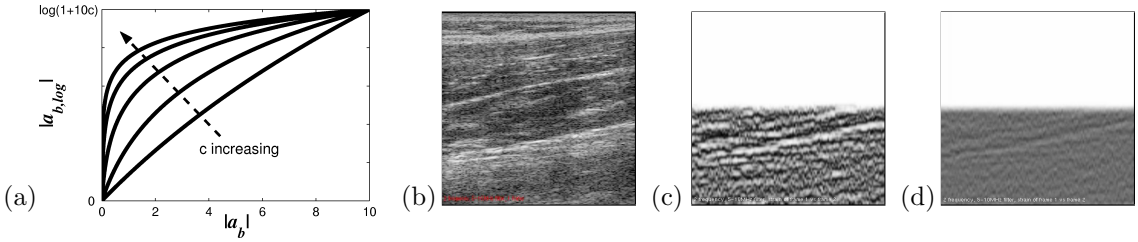

Figure 4: **Log compression:** (a) Input-output characteristic of the log compression of amplitude for compression factors $c = 0.1$, 1, 10, 100 and 1000. (b) Sonogram of human lower arm; 1% compressive strain simulated in the lower half. (c) Elastogram using basic EPZS. (d) Elastogram using EPZS with log compression ($c = 100$).

the strains are mapped to greyscale levels for display purposes.

The quality of the strain estimation depends on the properties of the scan subjects, the level of signal decorrelation, and also on three elastographic parameters. These are: window length (T), window spacing (Δt) and LSQ filter length (L_{lsq}). Together they govern the trade-off between axial resolution and estimation noise. An instructive discussion of this is provided by Righetti *et al.* [25].

Elastography in this framework produces unbiased strain estimates provided that the initial unbiased displacement estimates are performed at regularly spaced locations. In fact, this is rarely the case, because the locations of displacement estimates do not generally correspond to the centres of the analysis windows. Each displacement estimate is weighted towards locations where the RF signal has a large amplitude. High amplitude blips, such as the specular reflections at boundaries where the refractive index changes, distort the spacing of the estimates. For example, if neighbouring, overlapping windows contain a common bright boundary, their displacement estimates are effectively samples at the same position: therefore, the difference (strain estimate) is zero, regardless of the actual local strain. This artefact produces ghost images of B-scans superposed on the elastograms. However, it can largely be eliminated by log compression of the RF amplitude [26]. This is not an ideal solution, because it has been shown that log compressed signals introduce bias at the peak interpolation stage [23]. Less biased estimates can be made, however, using a search of the phase of the complex cross-correlation (CCC). Figure 4 shows the effect of log compression in this context, where the elastograms have been produced from a simulated uniform compression of RF data in the lower half of a real scan. The ghost image is almost entirely absent from Figure 4d where log compression is used. The search variation using CCC phase is preferable for 3D elastography anyway, because it has a low computational overhead.

2.1.3 Efficient phase zero searching

The speckle tracking algorithm in the freehand 3D ultrasound system is adapted from the original concept of Pesavento *et al.* [17], which is described in this section. They demonstrated superior accuracy and speed by working on analytic signals and using the CCC phase in an iterative search. It will be referred to as the efficient phase zero search (EPZS).

The window matching approach assumes that portions of the pre- and post-compression A-lines are time-shifted copies of the same signal. It makes sense, therefore, to consider the properties of the auto-correlation function: complex signals have pure-real auto-correlation values at zero lag. Similarly, the CCC of a pair of complex time-shifted signals has zero phase at the displacement where the signals match, as illustrated in Figure 5. The phase zero is easy to find when working with ultrasound signals because their average phase gradient is approximately equal to the probe centre frequency. A highly accurate estimate can usually be produced after a single iteration of gradient descent.

Consider a pair of complex signals, a and b. The unnormalised CCC is calculated as per Equation 1, where T is the window length, $n\Delta t$ is the position of the start of the window, and t'

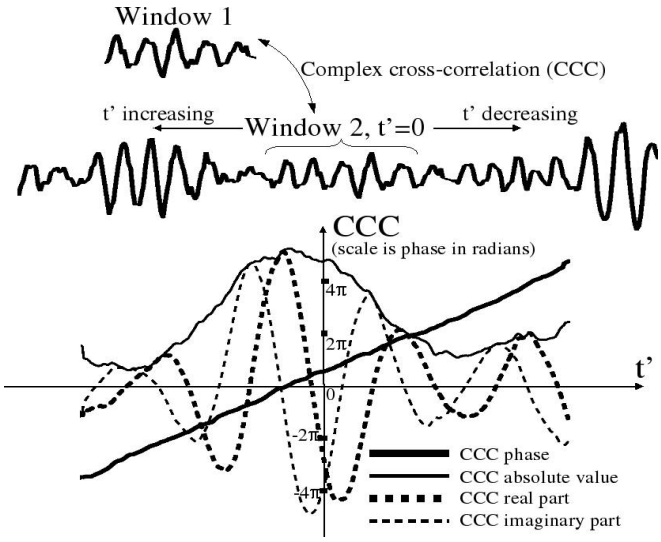

Figure 5: **CCC properties at the matching point:** the CCC phase is zero at the matching point and varies approximately linearly with displacement. Its average gradient is equal to the frequency centroid of the RF signal.

is the displacement of the post-compression window.

$$\langle a, b\rangle(n\Delta t, t') = \sum_{t=n\Delta t}^{n\Delta t+T} a^*(t)b(t+t') \tag{1}$$

Sub-sample precision is achieved by linear interpolation. This is most accurate at baseband frequencies, so a baseband conversion is calculated according to Equation 2, where a_b is the baseband analytic signal and ω_0 is a suitable modulation frequency.

$$a_b(t) = a(t)e^{-j\omega_0 t} \tag{2}$$

Using baseband analytic signals x_{b1} (pre-compression) and x_{b2} (post-compression), the iterative gradient descent motivated by Figure 5 was expressed in [17] in the form of Equation 3, where ω_c is the local frequency centroid. ${t'}_k$ and ${t'}_{k+1}$ are successive iterative estimates of the displacement (positive values indicate displacement towards the probe). The search is initialised with the final displacement estimate from the previous window, i.e. ${t'}_{0,n} = {t'}_{K,n-1}$.

$${t'}_{k+1} = {t'}_k + \frac{\arg\Big(e^{-j\omega_0 {t'}_k}\langle x_{b1}, x_{b2}\rangle(n\Delta t, {t'}_k)\Big)}{\omega_c} \tag{3}$$

In this form the arg function returns phase values in the range $-\pi$ to $+\pi$. This means that the search will fail if the change in displacement between one window and the next is $>\lambda/2$, where λ is the wavelength corresponding to the probe centre frequency. The upper limit this places on the strain rate is fortunately much higher than the strains that are encountered during scanning.

At the baseband conversion stage, signals are log compressed as mentioned in Section 2.1.2. This is performed according to Equation 4 following Pesavento *et al.* [17], where c is the compression factor. The phase is preserved, so the iterative search of Equation 3 can be applied without modification.

$$a_{b,\log} = \log\big(1 + c|a_b|\big)e^{j\arg a_b} \tag{4}$$

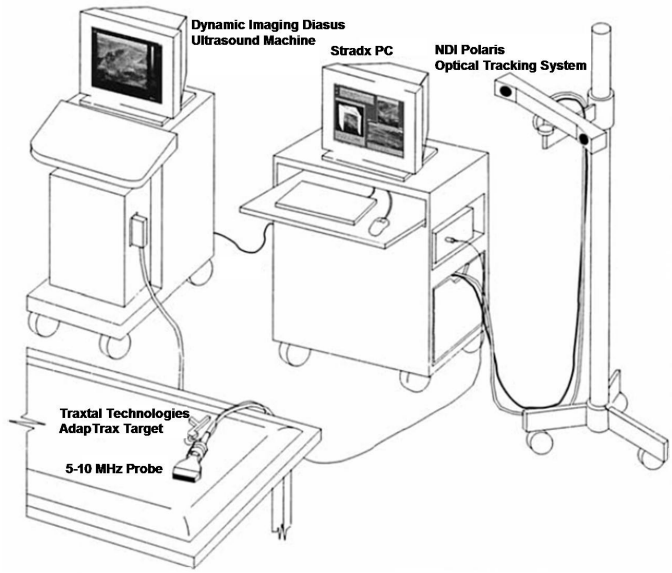

Figure 6: **Layout of the high definition freehand 3D ultrasound system.**

2.2 Freehand 3D ultrasound

The purpose of this work is to incorporate practical elastography within the freehand 3D ultrasound system of Treece *et al.* [13]. This is illustrated in Figure 6. It has an AdapTrax[1] target attached to the probe, tracked by a Northern Digital[2] Polaris optical position sensor. Custom calibration and imaging software are exploited, such that the system as a whole can register 3D point locations to an accuracy of 0.5 mm. Each ultrasound frame is treated as a rectangular plane in 3D space, and many planes are stacked together to produce 3D datasets. Typical applications include reslicing to view planes that are normally inaccessible, examining features that lie along curved surfaces, and estimating volumes.

3 Freehand 3D elastography

3.1 System outline

A modified Dynamic Imaging[3] Diasus ultrasound machine with a 5–10 MHz probe is used to acquire RF data at frame rates of 15–40 Hz. The probe centre frequency is 6.0 MHz, and samples are taken at 67 MHz using a Gage[4] CompuScope 14100 analogue-to-digital converter. A Hilbert filter with a 5–10 MHz passband is applied to produce analytic signals, so that axial strain estimation can be performed using EPZS[5] with parameters $T = 13.5\lambda$, $\Delta t = 3.5\lambda$ and $L_{lsq} = 2$, although experiments with longer LSQ filters are presented in Section 4. The criterion for stopping EPZS iterations is $|t'_k - t'_{k-1}| < 0.01\tau$, where τ is the RF sample spacing. This means that most searches consist of two iterations, although longer searches can occur when the signals are decorrelated.

[1] `http://www.traxtal.com`
[2] `http://www.ndigital.com`
[3] `http://www.dynamicimaging.co.uk`
[4] `http://www.gage-applied.com`
[5] In fact we also conducted experiments incorporating a lateral search. Neighbouring A-lines were searched in parallel. Then linear peak interpolation of the normalised cross-correlation was applied in the lateral direction to identify the best axial displacement estimate. The effect of this on the freehand elastograms was imperceptible, however, so lateral searching receives no further discussion in this paper.

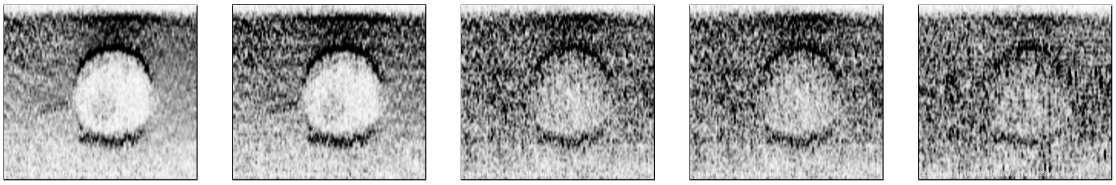

Figure 7: **Down-sampling levels:** from left to right: 1x (all data), 10x, 15x, 16x, 17x. The images degrade when the down-sampling level is increased, but the quality decays gracefully, with the SNR increasing only gradually while the sampling rate is still more than twice the probe bandwidth.

Some optimisation of the EPZS implementation has been applied in order to achieve a good frame rate. Three aspects are worth mentioning:

1. *Algorithm simplification.* We assume that $\omega_c \approx \omega_0$, where ω_0 was the modulation frequency used in the baseband conversion and ω_c is the frequency centroid. We choose ω_0 equal to the nominal probe centre frequency. This enables us to simplify the iteration formula in Equation 3 to obtain the more efficient expression of Equation 5. It should be noted that the phase of the baseband CCC may be $>2\pi$, so the arg function is potentially ambiguous. To handle this, our implementation assumes that the phase for each window remains within π radians of its initial value at ${t'}_{0,n} = {t'}_{K,n-1}$. This makes no difference to the upper limit on between-window displacements that was mentioned in Section 2.1.3.

$$ {t'}_{k+1} = \frac{\arg\Big(\langle x_{b1}, x_{b2}\rangle(n\Delta t, {t'}_k)\Big)}{\omega_0} \tag{5} $$

2. *Down-sampling.* The bandwidth of the probe is 2.1 MHz, so the Nyquist sampling frequency for baseband signals is 4.2 MHz. In fact all of the 67 MHz samples are used for the initial phase estimation with the Hilbert filter. However, thereafter up to 16x down-sampling can be performed assuming that the Nyquist sampling frequency is sufficient for accurate elastography. In practice, any down-sampling reduces the SNR, but it has been found that down-sampling up to 10x leads to only a minimal reduction in image quality. A demonstration using 2D elastograms of a phantom is shown in Figure 7. Down-sampling by a factor n reduces the computational overhead for Hilbert filtering, baseband conversion and log compression by $\mathcal{O}(n)$, and for the iterative search by $\mathcal{O}(n^2)$. With 10x down-sampling, a mean frame rate of 21.8 Hz has been achieved, which increases to 25.1 Hz if log compression is omitted (measurements were carried out on a machine with a 3 GHz CPU).

3. *Limit amplitude compression.* It was noted that log compression of the baseband signal, as in Figure 4, gave similar results for compression factors in the range 10^1–10^6. This motivated an experiment where the amplitude was set to unity at every sample position, so only phase information was retained. This is desirable, since calculating the CCC phase then requires only additions and subtractions. However, we were unable to produce sensible results by this method. It is assumed that after log compression the residual amplitude information provides a weighting, so that regions where the phase estimates are extremely uncertain are ignored. The log compression factor is 100 for all of the results that are presented here.

2D elastograms are produced by processing pairs of consecutive frames recorded during scanning. High speed elastography with 10x down-sampling provides a real time elastographic display for investigating suitable targets. For 3D elastography, however, the processing is performed off-line. This allows data to be acquired at a higher frame rate during the 3D sweep, averaging 30 Hz, which reduces between-frame decorrelation and has been found to give better elastographic results. Elastographic processing then proceeds off-line without down-sampling in order to achieve

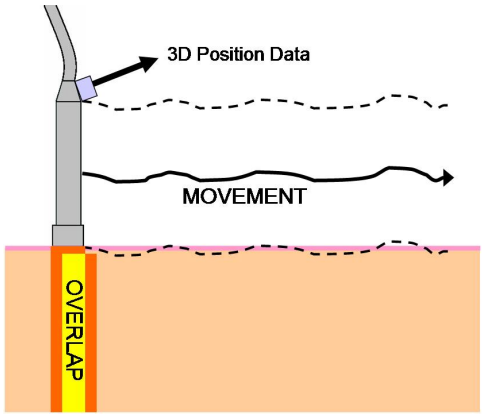

Figure 8: **A simple 3D scanning protocol:** axial and lateral motion are avoided while the freehand probe is translated slowly in the elevational direction.

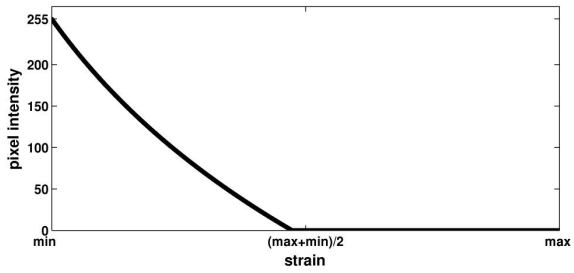

Figure 9: **Pixel intensity mapping:** a non-linear max-min strain scale focuses on variations in stiffer regions.

the maximum quality. The fast EPZS algorithm still offers a frame rate of several Hz, so a 3D elastogram consisting of several hundred 2D frames can be produced within 1–2 minutes.

The freehand 3D scanning protocol is surprisingly simple. It is illustrated in Figure 8. The best results are produced when no attempt is made at deliberately varying the pressure applied through the probe. The probe is swept slowly in the elevational direction, typically covering 3 cm in the course of 10 s. This gives an elevational spacing of ∼0.1 mm, which is well below the elevational width of the ultrasound resolution cell. The beam is narrowest at the focal depth, with a minimum elevational width of ∼2 mm for our probe. Therefore, consecutive frames overlap, and the level of decorrelation due to elevational translation is small. Between-frame strain is the result of small involuntary variations in the applied probe pressure, and *in vivo* scans include the additional effects of tissue-internal stresses. Mean strains between consecutive freehand data frames are in the range 0.03–0.50%.

Obviously the strain direction in this protocol is variable. Sometimes elastograms record an increase in compression and sometimes they record relaxation. Furthermore, since tissue is highly inhomogeneous, a single elastogram may contain some regions of compression, and others of relaxation. In any case, the interesting result for a qualitative interpretation is the magnitude of the strain, which is taken to indicate tissue stiffness. For this reason, the strain modulus is recorded, and the sign on each estimate is ignored.

The final processing stage maps absolute strain estimates to pixel intensities. The 2D elastograms must be normalised according to their strain distributions, so that pixels in the same type of tissue have similar intensities, regardless of the absolute level of strain in any particular frame. Another consideration is the property of interest: stiffness rather than strain. This is inversely

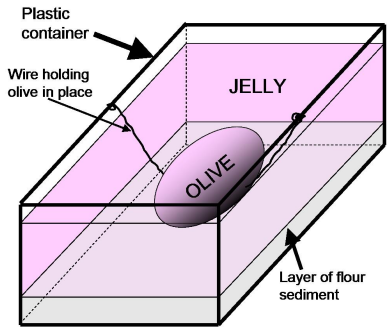

Figure 10: **Phantom:** stiff inclusion in soft tissue.

correlated with strain, and variations on a linear strain scale will tend to give the best contrast between different regions of soft tissue, which are perhaps of limited interest. Instead, a nonlinear scale linked to the minimum-maximum strain range is used, as shown in Figure 9. Stiff regions are bright and soft regions are dark.

To produce 3D elastograms, the pixel intensities of the 2D elastograms are written to file and stored together with the position sensor data. These files can be read by the Stradx[6] freehand 3D ultrasound software to produce a range of 3D visualisations.

3.2 First results

Phantoms were constructed to mimic stiff inclusions in soft tissue (olive in gelatin, see Figure 10). Note that wires prevented the olive floating out of the gelatin before it had set, but they were not intended to affect the mechanical properties thereafter. A suspension of flour in the gelatin provided suitable scattering properties, and a layer of flour sediment presented another relatively stiff region.

Figure 7 displays an example of a 2D elastogram from a phantom scan. Note that the image is uncalibrated — i.e. it is not intended for the purpose of taking quantitative strain readings — but the normalisation chosen gives a good contrast between the inclusion and the surrounding tissue. Another interesting point is that while the stiffness is fairly constant within the gelatin, the pixel intensities indicate higher strains towards the top of the image. This is because the stress spreads out away from the probe, so we witness the "target-hardening artefact", which has been mentioned in the literature [5]. Another typical artefact is the dark high-strain shadow around the edge of the olive, also mentioned in other studies [18, 27], which is caused by a stress concentration. We note that these artefacts are easy to interpret, and high-strain shadow can actually be *assistive* when it comes to identifying the boundaries between different stiffness regions.

The 2D elastogram was one among 389 combined in the 3D elastogram of Figure 11. Two reslices have been constructed, where intensities on these planes are assigned by nearest neighbour interpolation from the original 2D elastograms. The boundary of the olive can be seen easily thanks to the high-strain shadow, and there is an appreciable (though small) contrast between the regions of the resliced strain images within and outside the olive. The outline view in Figure 11c thus gives a clear impression of the 3D location of the inclusion.

Several artefacts are in evidence in Figure 11. Firstly, vertical streaks of estimation errors are present on the left hand edge of the 2D elastogram. We refer to these as "dropouts". They are caused by tracking errors, where the initial value in each iterative search is too far from the location of the correct match, so the search converges to an erroneous displacement. This limitation of the EPZS search is investigated in Section 4.1. Another region of poor strain estimates is present below the centre of the olive. From the B-scans it was apparent that a pocket of air had formed in this part of the phantom, so only noise was recorded. The air pocket is full of strain estimation errors,

[6]The software is available for free download from `http://mi.eng.cam.ac.uk/~rwp/stradx/`.

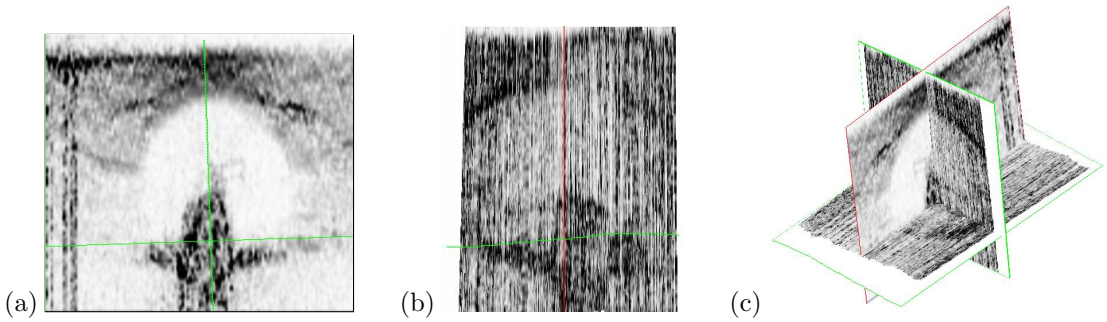

Figure 11: **3D elastography of olive/gelatin phantom:** (a) 2D elastogram. (b) Perpendicular reslice. (c) 2D elastogram and two such reslices.

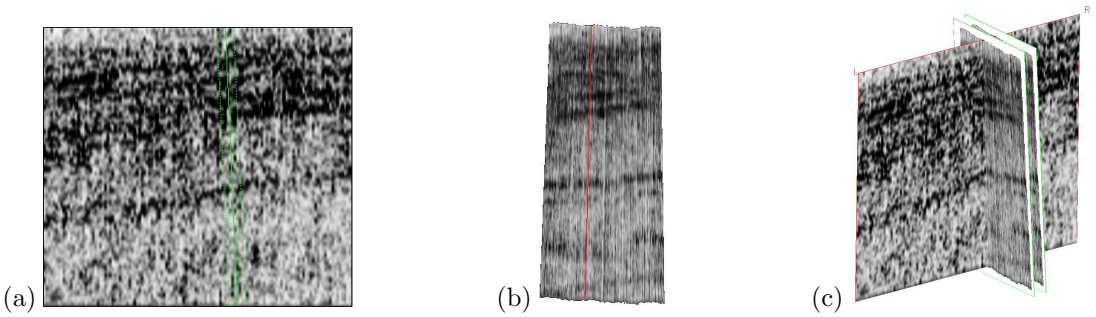

Figure 12: **3D elastography of the human calf:** (a) 2D elastogram. (b) Reslice. (c) 2D elastogram and reslice.

observable as a dark smear in both the 2D elastogram and in the reslice. A separate artefact of the reslices is the streaky appearance. This is the result of producing an image based on thin slices through many closely-spaced 2D elastograms of variable quality; it is not caused by variations in the applied stress — the normalisation accounts for different strain ranges. On the other hand, some elastograms exhibit more estimation noise than others, and outliers resulting from estimation errors can skew the normalisation. An approach to improving these 3D data sets is outlined in Section 4.2.

Freehand 3D elastography was also tested on a human subject. In this instance, the real-time display was used for locating suitable scan regions before performing the blind 3D sweep: speckle should be present throughout the tissue, and there should be few decorrelating phenomena (e.g. blood vessels). One such target within the human calf is displayed in Figure 12. Layers of soft fat (dark) and stiffer muscle (light) are easy to identify in both the 2D elastogram and in the perpendicular reslice.

4 Correction techniques for robust imaging

Our experience has shown that the difficulties in producing 2D elastograms during a freehand 3D scan are essentially the same as in normal freehand 2D elastography. This is because the level of degradation due to small elevational translations is not significant. One problem in both 2D and 3D cases is that quality can be compromised by the accumulated decorrelation due to involuntary movements in all five degrees of freedom that violate the assumptions of an axial search. This matters less in 2D elastography if elastograms with below average SNR can be ignored, but the 3D elastogram comprises all of the 2D frames, so any poor results detract from the overall quality of the data set.

(a) 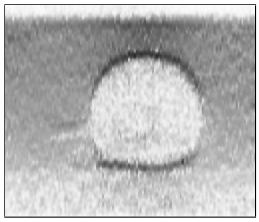 (b) 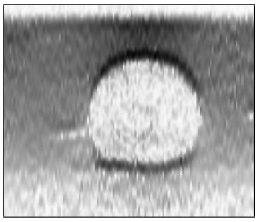 (c) 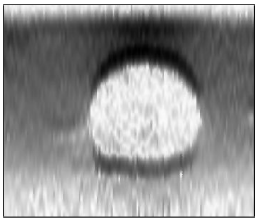

Figure 13: **LSQ strain estimation:** (a) Difference of neighbouring estimates ($L_{lsq} = 2$). (b) $L_{lsq} = 5$ (c) $L_{lsq} = 10$

One means of generally improving the SNR is by increasing the length of the LSQ filter that converts displacement estimates to strains [18]. Figure 13 presents a demonstration of this, where the SNR has been improved at the expense of the axial resolution. The optimal value for L_{lsq} depends on the data. A long LSQ filter may reduce the information content of good 2D elastograms, whereas for poor 2D elastograms the boost to the SNR is more important.

Two other methods have been devised for improving the 3D elastography results. Frame filtering automatically rejects 2D elastograms that are predicted to be of poor quality. A separate method corrects dropouts in 2D elastograms. From a visual inspection of the images, dropouts are the most severe artefact. They merit special attention because dropouts often mar images where the overall image quality would otherwise be good.

4.1 Dropout detection and correction

Dropouts are easy to spot by eye: they appear as incongruous vertical streaks usually extending to the bottom of the image, as in Figure 14a. A point is reached in such A-lines where the displacement tracking breaks down and subsequent displacement estimates are erroneous. This is perhaps made clear by inspection of the surface plots (produced using MATLAB[7]) of displacement estimates in two elastograms, one of which produced a good image (Figure 14c) the other of which contained dropouts (Figure 14d). An error-free elastogram from a volume of connected tissue must have a continuous displacement field, whereas it can be seen that dropouts give rise to discontinuities.

Dropouts propagate because at each window EPZS is initialised with the previous estimate, so a single large error can wipe out the remainder of an A-line. However, the cause of these large errors is not obvious. An initial theory was that mismatches occurred if the local strain went above a threshold, since the EPZS search cannot move further than $\lambda/2$ at a single window. As a test, EPZS was applied to synthetic data of the same form as used in Figure 4, where much higher strains were simulated. In fact, it was found that high strains *can* produce dropouts, but only at strains upward of 5.4%. This is almost two orders of magnitude higher than typical mean strains in the freehand data, so it is unlikely to be a frequent cause. A more plausible explanation for most dropouts was devised after careful inspection of the freehand elastograms where dropouts were present, together with the features evident from the corresponding B-scans. It was observed that dropouts begin at tissue features with local decorrelating properties. These include slip planes between the inclusion and the gelatin, pockets of fluid within the jelly, multiple specular reflections between the probe and metal wires, and all locations of fluid flow. A good example is provided by Figure 14b, showing a largely homogeneous section through the human calf. The dark patch is a blood vessel, where blood flow causes decorrelation between one frame and the next. Several estimation errors in the vessel are large enough to produce dropouts.

Dropouts can be detected automatically by identifying displacement outliers within each row of estimates. A-lines where the displacement is more than three standard deviations from the mean are marked as dropped-out. Starting from the top of the image, at each row the displacement mean and standard deviation are recalculated based only on the A-lines that have not dropped

[7]MATLAB is a registered trademark of The MathWorks Ltd.

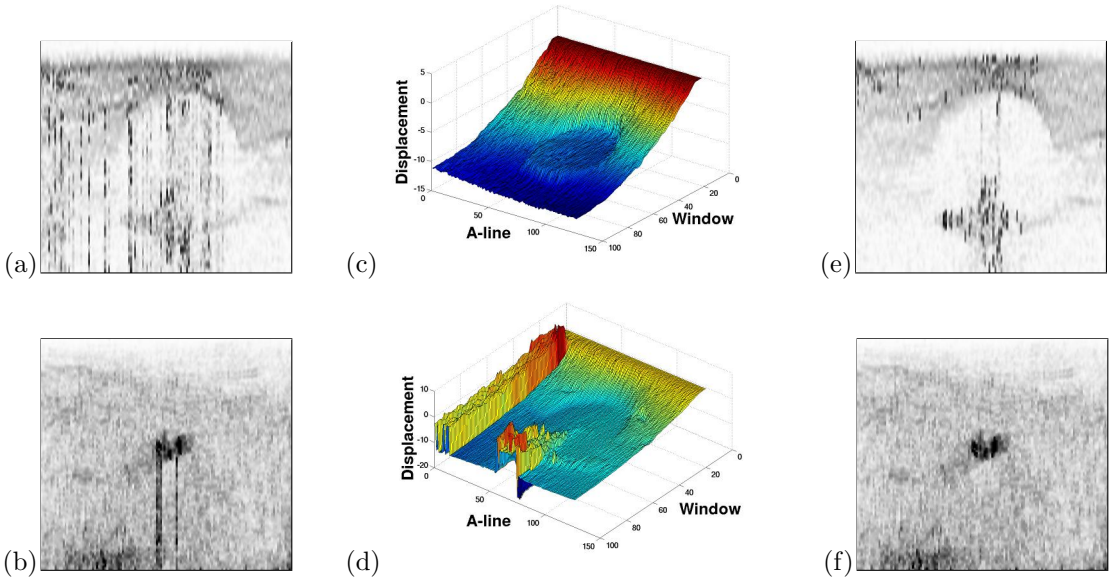

Figure 14: **Dropouts:** (a) 2D elastogram with dropouts. (b) Human calf elastogram including a blood vessel. (c) Surface plot of displacement estimates in a clean elastogram. (d) Displacement estimates in an elastogram with dropouts. (e) Corrected version of *a*. (f) Corrected version of *b*.

out, and new outliers are marked as dropouts. In a study of 80 freehand 2D elastograms, this method detected 100% of the dropouts that were spotted by eye, in addition to which 30% of the automatic dropout detections were either false positives or had been missed in the visual inspection.

Attempts have been made to devise data-driven dropout correction methods, since error propagation often masks sections of A-lines where the data is good. The general principle is to detect dropouts and reinitialise EPZS in the section below with the average of estimates in neighbouring clean A-lines. A challenge with this is finding the start of the dropout, so that errors are eliminated as far as possible.

The best method yet devised gets around this by working backwards up the image, so the dropout start position is not required. A first pass of EPZS displacement estimation proceeds as usual. Then an estimate is made of the mean axial displacement at the bottom of the elastogram — the bottom row of displacement estimates is averaged, with outliers excluded. This provides the initialisation for a second pass of EPZS beginning at the bottom of the image. Dropped-out A-lines are divided into two sections: above and below the decorrelation patch. Estimates from the first pass are reliable in the upper section, before breaking down at the patch. Second pass estimates are better in the lower section, but again they break down. In both passes the normalised cross-correlation is evaluated for every window match, and ultimately estimates are adopted from whichever pass had the higher correlation. The computational cost of this correction is not large, since the signal pre-processing stages that constitute most of the load are entirely unchanged. The iterative search must be carried out twice, but it is a small fraction of the overall processing. Normalised cross-correlation computations present a further small increase in the overhead.

The method corrects most dropouts, so decorrelation patches such as blood vessels cause displacement error blips, but errors no longer propagate. The most common exception to this occurs when more than one decorrelation patch is present within a single A-line, in which case estimates between the patches are likely to remain erroneous. However, the correction to many 2D elastograms is impressive, such as in Figures 14e–f (corrected versions of Figures 14a–b). An experiment was conducted using the data from a 3D data set with 309 2D elastograms, recording by eye the number of dropouts in each frame and their lengths, in order to produce an estimate of the image area fraction (IAF) lost to dropouts. The IAF before dropout correction was 0.65%, whereas

applying the correction method reduced it to 0.26%. The numbers sound small, but dropouts seem to have a disproportionate effect on the perceived image quality: for interpretation purposes, many images appear greatly improved.

4.2 Frame filtering

4.2.1 Observations

Aside from dropouts, the general image quality varies a great deal between 2D elastograms. Aspects of this are variations in the apparent level of estimation noise and in the contrast between different tissue regions (despite the min-max strain normalisation). The overall quality of 3D elastograms could be improved significantly by hand picking the best 2D results, so this section introduces automatic frame filtering.

A survey of 2D elastograms from a single 3D phantom scan is presented in Table 1. The images were manually selected, picking ten 2D elastograms perceived as being "good" and ten perceived as being "bad". Frames with large numbers of dropouts were avoided so as not to confuse general image quality determinants with the causes of the dropouts. Two metrics were evaluated *a posteriori* for each of the elastograms: (1) mean magnitude of the normalised cross-correlation; and (2) mean strain. The first of these entailed calculating the normalised cross-correlation between analysis windows where EPZS found a match, and averaging these values across the entire 2D elastogram. The mean strain was found by making a robust mean displacement estimate from the speckle tracking data, excluding outliers, and dividing this by the number of samples in the A-line. Note that the sign on the strain denotes either compression (positive) or relaxation (negative). This records the direction of the deformation that physically occurred during the scan. It does not affect the elastographic processing since the labelling of pre- and post-compression frames could just as well be reversed.

There are marked differences between the distributions of the metrics for the good and bad 2D elastograms in Table 1. It was anticipated that mean strain would be an important determinant of elastographic SNR, based on the notion that the mean squared strain is the elastographic signal power. Indeed, the strains of the good elastograms are generally larger. However, 4G has a lower strain than most of the bad elastograms, although the image is evidently superior. Conversely, 2B has a relatively high strain but the image is poor. The intuition that SNR is linked to mean strain may be correct, but this parameter alone does not separate the good results from the bad.

A more striking observation is that every one of the good elastograms was an instance of relaxation (negative strain), whereas more than half the bad elastograms were compressions. This suggests some asymmetry in the elastography scanning: either hysteresis behaviour in the phantom reducing decorrelation effects in the relaxation direction, or smoother movement of the freehand probe as pressure is reduced. This sample is not statistically significant, but it should not be ruled out that asymmetry might be a common feature of freehand 3D scanning.

Correlation values provide a better separation of the good and bad distributions. While all of the good elastograms were in the range 0.7719–0.8721, the values for seven of the bad elastograms were below this range. Of the remaining three, 8B and 9B had positive strains, which seems to have precluded successful elastography with this phantom. Closer inspection of the displacement estimates for elastogram 10B showed that in this frame the probe had twisted, with small relaxations in A-lines at one end and larger compressions in A-lines at the other. This unusual displacement field may explain why the axial strain display is ambiguous. Mean correlation appears more promising than mean strain as a metric for selecting good 2D elastograms, though qualitative observations alone provide insufficient evidence for any firm conclusions to be drawn.

4.2.2 Experimental verification

Contrast-to-noise ratio (CNR) is an image analogue of SNR. CNR measurements were used as the basis for a quantitative investigation of frame filtering. We adopt a definition of CNR from Chaturvedi *et al.* [28] as stated in Equation 6, where s_1 and s_2 are the mean pixel intensities in

Good elastograms				Bad elastograms			
A	**B**	**C**	**D**	**A**	**B**	**C**	**D**
1G	0.8370	-0.29		1B	0.6250	-0.12	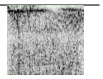
2G	0.8581	-0.12		2B	0.3641	+0.20	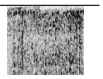
3G	0.8656	-0.17		3B	0.5693	+0.054	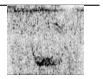
4G	0.8652	-0.051		4B	0.6031	+0.051	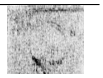
5G	0.8101	-0.31		5B	0.5336	-0.071	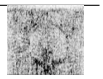
6G	0.8370	-0.16		6B	0.5490	0.0044	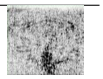
7G	0.7719	-0.36		7B	0.7092	+0.14	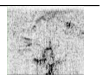
8G	0.8348	-0.34		8B	0.8675	+0.052	
9G	0.8721	-0.46		9B	0.8800	+0.079	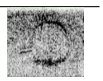
10G	0.8516	-0.40		10B	0.8374	-0.16	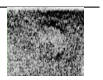

Table 1: **Quantitative survey of good and bad elastograms.** A: Elastogram number. B: Mean normalised cross-correlation magnitude. C: Mean strain (%). D: Thumbnail elastogram.

the inclusion and in the soft tissue respectively, σ_1 is the standard deviation of pixel intensities in the inclusion and σ_2 is the standard deviation in the soft tissue.

$$\mathrm{CNR} = \sqrt{\frac{2(s_1 - s_2)^2}{\sigma_1^2 + \sigma_2^2}} \tag{6}$$

Phantom scans were inspected off-line in B-scan mode to draw manual 3D segmentations of the inclusion and the soft tissue. In each data set the entire segmented volume was used as the basis for CNR evaluation. Usually the CNR would be calculated for a 2D resliced image, but the CNR values vary between reslices depending on the plane that is inspected: for example, reslices in elevational-lateral planes generally have lower CNR than elevational-axial reslices because they miss the stress concentrations, so the overall contrast is lower. Instead, 3D elastograms were resampled onto regular voxel arrays using nearest-neighbour interpolation, and all of the voxel intensities throughout the segmented regions were used for the CNR calculations. The effect of this was to estimate the expected CNR for a reslice through each data set on a plane of arbitrary orientation.

With CNR as the measured output variable, the primary input variables were the threshold levels in the minimum strain and minimum correlation filters. Mean strain and mean correlation were calculated as previously described, and 2D elastograms which fell below the threshold levels were omitted from the 3D elastograms. The elastography algorithm used for this experiment was largely identical to the setup described earlier, although only 60 windows were used in each A-line with $T = 40\lambda$ and $\Delta t = 4.5\lambda$ — these conservative values tend to result in smaller errors in the poor elastograms, although resolution is reduced in the best elastograms. Another change was

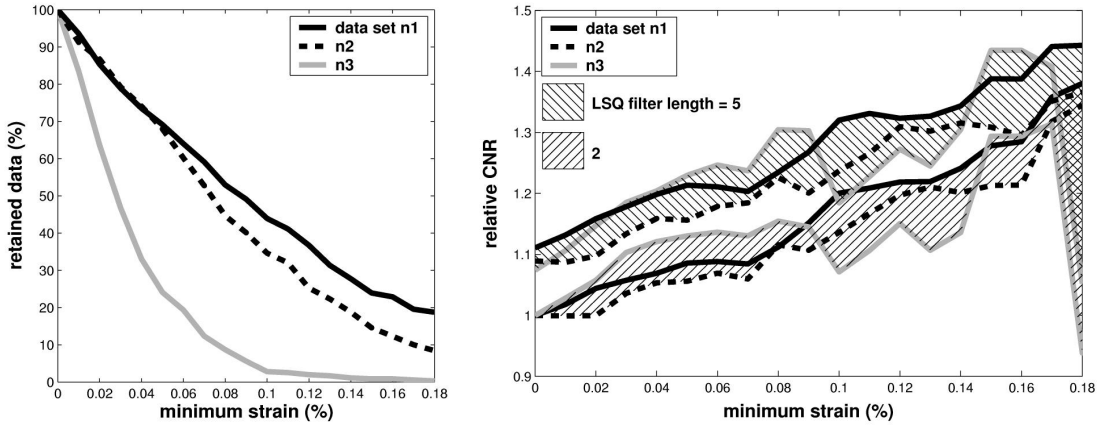

Figure 15: **Minimum strain filter:** Left: percentage of data retained versus minimum strain. Right: CNR performance versus minimum strain, with comparison of LSQ filters. The upper set of curves is for $L_{lsq} = 5$, while the lower set is for $L_{lsq} = 2$. Cross-hatching marks the range of results with each LSQ filter.

the use of different lengths of LSQ filters for strain estimation: they were applied for comparative purposes, and to see if frame filtering had different effects depending on L_{lsq}. Several values were tried for L_{lsq}, but for clarity only two are displayed in the graphs of results. The behaviour of the intermediate filters was unsurprising, so results are shown only for $L_{lsq} = 2$ (difference estimation) and $L_{lsq} = 5$ (a long LSQ filter).

Figure 15 presents results for the minimum strain filter and Figure 16 for the minimum correlation filter. The graphs on the left show the proportion of the data that was retained after the filter, which indicates the effect on elevational resolution (frame filtering reduces frame density in the elevational direction). On the right, two curves are plotted for each data set, one for each LSQ filter. Cross-hatching between the curves marks the sets corresponding to each LSQ filter, where the longer filter produces a higher CNR. Note that the vertical axis is relative (as opposed to absolute) CNR: this is the absolute CNR divided by the value of the CNR when no frame filter was applied and $L_{lsq} = 2$. The CNR varied significantly between the data sets, so performance comparison on an absolute scale would be difficult.

Figure 15 shows unambiguously that a filter on the minimum strain can improve the CNR, and the improvements may be larger than are yielded by the application of a long LSQ filter. LSQ filtering and frame filtering have a common feature in that the resolution is reduced. LSQ filtering reduces axial resolution, while frame filtering reduces elevational resolution. When the resolution becomes extremely poor, the CNR is eventually also affected, since nearest neighbour interpolation of sparsely distributed 2D elastograms distorts the apparent shape of the inclusion: high thresholds produce erratic results in data set n3 because few frames are retained, and the relatively poor CNR with a 0.18% strain threshold is based on extending just a single 2D elastogram to fill the whole elevational depth. However, a promising feature of these results is that data sets n1 and n2 (where 10–20% of the data was retained at the highest threshold) continued to exhibit increasing CNR. It might have been expected that the greater level of motion implicit in these elastograms would cause greater decorrelation and therefore a reduction in image quality, but this has not been recorded. It suggests that if sufficient frames at higher strains were available, then the strain filter would offer significantly improved 3D images.

The trend for minimum correlation filtering in Figure 16 again shows that a threshold can improve the CNR. In this case the data sets all had moderate or good coverage of the elevational dimension up to a correlation threshold of 0.93, at which point 15–20% of the data was retained. This is roughly where the relative CNR peaks, and the size of the improvement is similar to the maximum gain of the minimum strain filter. Note that the n3 data set "caught up" with the others

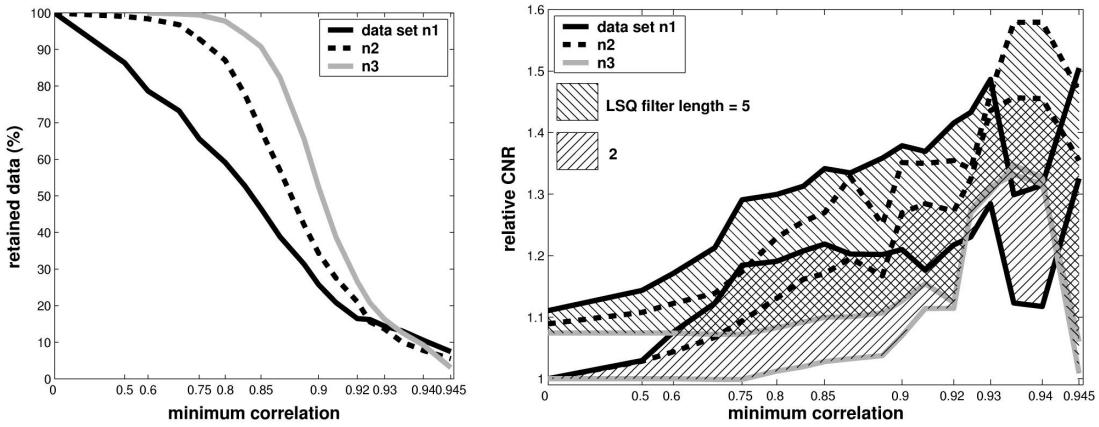

Figure 16: **Minimum correlation filter:** Left: percentage of data retained versus minimum correlation. Right: CNR performance versus minimum correlation, with comparison of LSQ filters. As in Figure 15, upper and lower curves are for $L_{lsq} = 5$ and 2 respectively, and cross-hatching marks the range of results in each case.

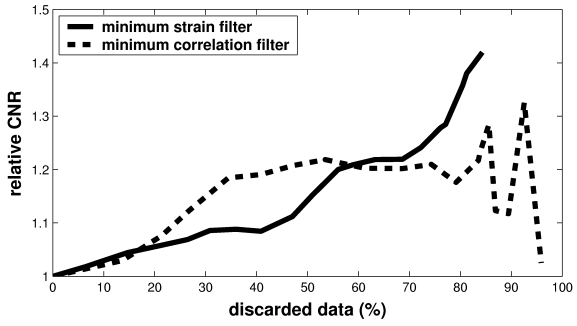

Figure 17: **Comparison of strain and correlation filters:** CNR versus the proportion of the data that has been discarded (data set n1, LSQ filter length = 2).

by 0.93, since less data was discarded for n3 at the lower thresholds. Up to 0.93 the curves appear to be converging towards similar performance improvements. Thereafter, the behaviour becomes erratic, so it is unclear whether there continues to be a correlation between higher thresholds and increased CNR. This cannot be explained simply as the effect of sparse data with inadequate resolution, because at least 10 frames were retained in each data set right until the 0.945 threshold. There may be an additional reason why high correlation thresholds eventually become unhelpful. Good correlation indicates a close match between data in the pre- and post-compression frames, so the estimation noise will be lower than average. However, the best correlation will be recorded for a stationary probe. While it is noted that surprisingly small strains can yield good elastograms, Figure 15 demonstrates that higher strains at the upper end of our low-strain range are to be preferred. A high correlation threshold disproportionately selects frames where the strain is low, and for this reason it will eventually yield a reduced SNR (and CNR) as the signal amplitude approaches zero.

An alternative way of comparing the merits of minimum strain and minimum correlation filters is to plot a characteristic curve of CNR against the proportion of the data that is retained. Figure 17 shows a pair of such curves for the n1 data set. It is clear that the filter initially offering the "best value" is minimum correlation, but as the data become more sparse (and the correlation filter begins to positively select low strain) the minimum strain filter eventually yields higher

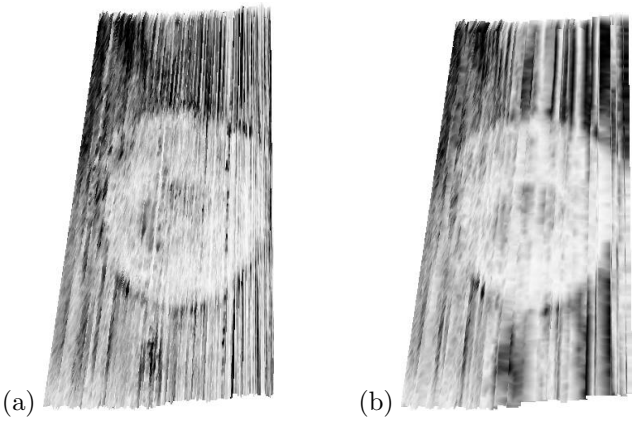

Figure 18: **Frame filtered reslice:** (a) Before and (b) after filtering.

performance. The same pattern is repeated in the n2 and n3 data sets. The implication is that neither minimum correlation nor minimum strain will alone yield an optimal filter — both are required.

As a visual example, Figure 18 shows a reslice through data set n2. Images with $L_{lsq} = 5$ are presented before and after frame filtering, where a strain threshold of 0.04% and a correlation threshold of 0.87 have been applied together. Unlike the procedure in the experiments, the display here is produced by linear interpolation. It is evident that the filter has improved both the contrast and the CNR.

A secondary implication of this investigation is that higher strains on average *do* produce better CNR. This is predictable given the strain filter analyses of researchers in static compression elastography. However, for the reasons explained it does not necessarily follow that high strains in excess of 1% are desirable in freehand elastography. Initial experience of freehand scanning showed prior to these results that low acquisition rates and coarser palpation through the probe produce *worse* elastograms. Despite this, the frame filtering results indicate that a somewhat rougher scanning technique might produce *superior* results if it could be combined with suitable filtering to remove error frames when they arise.

5 Conclusions

We have presented a novel system for freehand 3D elastography. The scanning protocol is simple and the elastograms are constructed in real time using a standard PC and unmodified freehand 3D ultrasound equipment.

Some conventional elastography techniques have been adopted for freehand 3D purposes, including the phase zero searching algorithm for displacement estimation. Down-sampling of the analytic signal is a useful way to increase the processing speed whenever high frame rates are required during scanning. It has been found that freehand 3D data acquired at a high frame rate is amenable to the same processing as 2D by exploiting the finite elevational width of the ultrasound beam. We find that at low strains planar elastograms with good contrast are often produced, and a simple normalisation scheme means that the display data can be combined to construct 3D data sets even though the applied compressions are uncontrolled. Reslices through the 3D elastograms are easy to interpret, especially for identifying stiffness boundaries, although the overall CNR is not high.

The challenge of freehand 3D elastography is in developing a suitably robust imaging system. Quality variations between the planar images contribute to a high level of noise in the 3D elastograms. This is mostly the result of variable CNR due to differences in the physical conditions of each frame, such as average strain and the level of decorrelation. A frame filtering method has been

shown to produce marked CNR improvements by omitting bad planes. Other errors are caused by dropouts: error propagation following decorrelation patches. Most of these are now removed by a second pass correction method.

Future work must investigate optimal criteria for frame filtering. It would be helpful if a unified measure of image quality could be devised, so that the CNR/resolution trade-off could be assessed within a clear framework. Frame filtering may ultimately need to apply variable thresholds across the volume to obtain an optimal balance between CNR and resolution at all locations. Furthermore, the improvements may be greater if similar filtering can be introduced at the sub-frame level, to construct 3D data sets from the best patches within the frames. Successful filtering methods will have knock-on implications for the freehand scanning protocol that delivers the best results. There is scope for both theoretical and experimental work to develop a solid understanding of this interaction.

References

[1] C. L. de Korte, A. F. W. van der Steen, E. I. Céspedes, and G. Pasterkamp. Intravascular ultrasound elastography in human arteries: experience *in vitro*. *Ultrasound in Medicine and Biology*, 24(3):401–408, 1998.

[2] A. Pesavento and A. Lorenz. Real time strain imaging — a new ultrasonic method for cancer detection: first study results. In *Proceedings of IEEE Ultrasonics Symposium 2001*, pages 1647–1652, October 2001.

[3] B. S. Garra, E. I. Céspedes, J. Ophir, S. R. Spratt, R. A. Zuurbier, C. M. Magnant, and M. F. Pennanen. Elastography of breast lesions: initial clinical results. *Radiology*, 202(1):79–86, January 1997.

[4] T. J. Hall, Y. Zhu, and C. S. Spalding. *In vivo* real-time freehand palpation imaging. *Ultrasound in Medicine and Biology*, 29(3):427–435, 2003.

[5] J. Ophir, I. Céspedes, H. Ponnekanti, Y. Yazdi, and X. Li. Elastography: a quantitative method for imaging the elastisticity of biological tissues. *Ultrasonic Imaging*, 13:111–134, 1991.

[6] A. R. Skovoroda, A. N. Klishko, D. A. Gusakyan, Y. I. Mayevshki, V. D. Yermilova, G. A. Oranskaya, and A. P. Sarvazyan. Quantitative analysis of the mechanical characteristics of pathologically changed soft biological tissues. *Biophysics*, 40(6):1359–1364, 1995.

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

[8] R. M. Lerner, S. R. Huang, and K. J. Parker. Sonoelasticity images derived from ultrasound signals in mechanically vibrated tissues. *Ultrasound in Medicine and Biology*, 16:231–239, 1990.

[9] K. J. Parker, S. R. Huang, R. A. Musulin, and R. M. Lerner. Tissue response to mechanical vibrations for sonoelasticity imaging. *Ultrasound in Medicine and Biology*, 16:241–246, 1990.

[10] R. K. Nightingale, M. L. Palmeri, R. W. Nightingale, and G. E. Trahey. On the feasibility of remote palpation using acoustic radiation force. *Journal of the Acoustical Society of America*, 110(1):625–634, 2001.

[11] A. P. Sarvazyan, O. V. Rudenko, S. D. Swanson, J. B. Fowlkes, and S. Y. Emelianov. Shear wave elasticity imaging: a new ultrasonic technology of medical diagnostics. *Ultrasound in Medicine and Biology*, 24(9):1419–1435, 1998.

[12] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. 3D elastography using freehand ultrasound. In *Medical Image Computing and Computer-Assisted Intervention — MICCAI'04*, pages 1103–1104, Saint-Malo, France, September 2004. LNCS 3217, Springer.

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

[14] A. H. Gee, R. J. Housden, P. Hassenpflug, G. M. Treece, and R. W. Prager. Sensorless freehand 3D ultrasound in real tissue: speckle decorrelation without fully developed speckle. *Medical Image Analysis*, in press.

[15] M. M. Doyley, J. C. Bamber, F. Fuechsel, and N. L. Bush. A freehand elastographic imaging approach for clinical breast imaging: system development and performance evaluation. *Ultrasound in Medicine and Biology*, 27(10):1347–1357, 2001.

[16] J. C. Bamber and N. L. Bush. Freehand elasticity imaging using speckle decorrelation rate. *Acoustical Imaging*, 22:285–292, 1996.

[17] A. Pesavento, C. Perrey, M. Krueger, and H. Ermert. 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, September 1999.

[18] F. Kallel and J. Ophir. A least-squares strain estimator for elastography. *Ultrasonic Imaging*, 19:195–208, 1997.

[19] R. J. Dickinson and C. R. Hill. Measurement of soft tissue motion using correlations between A-scans. *Ultrasound in Medicing and Biology*, 8(3):263–271, 1982.

[20] T. Varghese and J. Ophir. A theoretical framework for performance characterization of elastography: the strain filter. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(1):164–172, January 1997.

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

[22] I. Céspedes, Y. Huang, J. Ophir, and S. Spratt. Methods for estimation of subsample time delays of digitized echo signals. *Ultrasonic Imaging*, 17:142–171, 1995.

[23] S. K. Alam and J. Ophir. The effect of nonlinear signal transformations on bias errors in elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 47(1):297–303, January 2000.

[24] S. K. Alam, J. Ophir, and E. E. Konofagou. An adaptive strain estimator for elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 45(2):461–472, March 1998.

[25] R. Righetti, J. Ophir, and P. Ktonas. Axial resolution in elastography. *Ultrasound in Medicine and Biology*, 28(1):101–113, 2002.

[26] I. Céspedes and J. Ophir. Reduction of image noise in elastography. *Ultrasonic Imaging*, 15:89–102, 1993.

[27] I. Céspedes and J. Ophir. Elastography: elasticity imaging with application to muscle and breast *in vivo*. *Ultrasonic Imaging*, 15:73–88, 1993.

[28] P. Chaturvdi, M. F. Insana, and T. J. Hall. Testing the limitations of 2-D companding for strain imaging using phantoms. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 45(4):1022–1031, 1998.