3D ULTRASONIC STRAIN IMAGING USING FREEHAND SCANNING AND A MECHANICALLY-SWEPT PROBE

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

CUED/F-INFENG/TR 623

January 2009

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

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

3D ultrasonic strain imaging using freehand scanning and a mechanically-swept probe

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

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

Abstract

This paper compares two approaches to 3D ultrasonic axial strain imaging. The first uses a tracked ultrasound probe swept manually over the region of interest to build up a co-registered sequence of 2D strain images, each obtained by comparing neighbouring B-scans in the sequence. The alternative uses a mechanically-swept 3D probe to record pre- and post-deformation volumes, which are then processed to yield a volume of strain data. The resulting strain images depend on the stress fields induced by the different probe footprints and also on the signal processing techniques used to produce the strain volumes. Both of these factors are considered in this paper, which presents a comparison of the two approaches based on finite element simulations validated through *in vitro* experimentation. The conclusion is that, for a given frame density, high quality axial strain data is more easily obtained using the 3D probe. However, the freehand approach might be preferable in situations where limited access to the scanning target restricts the use of a large footprint probe.

1 Introduction

Ultrasonic strain imaging is an emerging technique for visualising tissue stiffness. Several potential applications have been identified, including the detection of tumours in soft tissue [4, 9, 29], the monitoring of atherosclerosis [6, 27] and the measurement of skin elasticity [32]. This paper is concerned with quasi-static strain imaging [20], which is achieved in 2D by comparing radio-frequency (RF) echo signals from two or more B-scans with a small difference in the axial compression applied through the ultrasound probe. Axial strain is determined by calculating the gradient of axial displacement estimates. This 2D imaging technique has more recently been extended into 3D, offering the potential for more complete visualisation of the anatomy and more accurate measurement of the volume of stiff or soft regions.

There are currently two distinct approaches to quasi-static 3D strain imaging. The first approach will be referred to in this paper as volume-based 3D strain imaging. In this paradigm, whole volumes are recorded using specialised 3D probes, either 2D arrays or mechanically-swept 1D arrays. Strain data is obtained by comparing two overlaid volumes with a small compression applied between them [1, 2, 7, 8, 13, 19, 21, 31].

The alternative has received less interest but is also a feasible way to obtain 3D axial strain data. This approach uses a conventional ultrasound probe, which records a sequence of B-scans while being swept in the out-of-plane direction over the region of interest. Provided adjacent B-scans lie within each other's elevational decorrelation range, they can be compared, as in conventional 2D strain imaging, to produce a sequence of 2D frames of axial strain data. This sequence of strain images sweeps out a 3D volume. Each image is located in space either by a position sensing device attached to the probe [14] or, in the case of intravascular strain imaging, by a continuous pullback method [28]. We refer to this approach as freehand 3D strain imaging. The freehand in this case refers to the unconstrained motion of the transducer.

There are several fundamental differences between these two approaches that may affect the resulting strain images. Firstly, the volume-based approach uses probes that typically have a larger footprint than the conventional 1D-array probes used in the freehand technique. It is well known that the stress field, and therefore the strain, induced in the tissue by an applied compression

depends partly on the size and shape of the compressor [20, 23]. When interpreting stain data, it is often assumed that the stress field is uniform, so that the strain image represents the tissue stiffness. This assumption leads to artefacts caused by, for example, stress concentrations around stiff inclusions [12] and reduced stress with depth. The precise nature of these artefacts will depend, to some extent, on the probe footprint.

Secondly, there are differences in the way the two approaches record and process the data. The freehand approach is restricted to calculating strain between pairs of image frames, comparing RF signals in 2D windows and searching for displacements in the axial and lateral directions. In contrast, the volume-based approach can estimate displacements using 3D windows. This allows either an improvement in signal to noise ratio (SNR), because more data is used to estimate each displacement, or an improvement in axial resolution for the same SNR, by reducing the 3D window size axially. It is also possible to allow for displacements in the third, out-of-plane direction. This may further improve the SNR by matching pre- and post-deformation windows more accurately.

The objective of this paper is to compare the two approaches, looking at both the mechanical and signal processing differences, and establish their relative merits by considering resolution, contrast and SNR. The paper is structured as follows. Section 2 describes a practical implementation of the two techniques, followed by details of how they can be simulated using finite element methods. In Section 3, we present experiments based on these simulations, with the known ground truth facilitating the measurement of quantitative performance metrics such as contrast and SNR. This section also includes an *in vitro* validation of the simulation results, with additional commentary on the practical ease of use of the two scanning techniques. Finally, we present our main conclusions in Section 4.

2 Methods

2.1 Scanning protocol and processing

The two approaches considered in this paper have both been described previously, so we refer the reader to [14, 15, 17, 18, 31] for implementation details. Here, however, we review the main features and highlight particularly where the two approaches differ.

The freehand protocol [14] uses a linear-array probe, which is moved slowly over the region of interest while producing 2D strain images at a rate of about ten per second. At the same time, the axial contact pressure is varied gently by hand, to produce the necessary inter-frame deformation. For effective strain imaging, consecutive B-scans must not be separated by more than the elevational decorrelation distance. A position sensor is attached to the probe to record the location of each RF image frame.

Subsample axial displacements are estimated between adjacent frame pairs using the Weighted Phase Separation method [18] in 2D windows of equal aspect ratio. Subvector lateral displacements are also calculated, for the purpose of improving the axial estimates, by maximising correlation over a range of ± 1 A-line. The precision of each displacement value (estimated from the correlation [17]) is recorded. The location at which each displacement estimate applies is calculated using the Amplitude Modulation Correction technique [15] and the axial displacements are then interpolated onto a regular grid. Axial strain is calculated simply by differencing pairs of axial displacement estimates in the interpolated data.

Since each strain image in the sequence is produced with a different compression, the strain values are normalized to a consistent range so that each image displays the same relative stiffness between regions [17]. The normalization is done by fitting a function, by precision-weighted least squares, to the strain values. The strain values are then normalized by the local value of the function at each location in the image. The 2D normalization function used in this work is given by

$$f_{2D}(x, y) = a(1 + by)(1 + cx) \tag{1}$$

where x and y are the lateral and axial coordinates respectively and a, b and c are constants determined when the function is fitted to the data. The y variation compensates to some extent

for the expected reduction in stress with depth, while the x variation allows for non-uniform compression along the probe face.

Given the freehand nature of the scan, with varying compression and unconstrained motion, some pairs of images will produce better strain data than others. A final stage is therefore to spatially filter the 3D image sequence with a 3D Gaussian smoothing kernel, with each sample weighted by the precision of the strain estimate [17]. The Gaussian kernel is truncated beyond $\pm$ two standard deviations. Not only does this improve the appearance of the 3D data by emphasising the high quality data, it also smooths out axial variations in the positions of features caused by varying axial compression. This is the only stage in the freehand method's signal processing that considers the data as a volume rather than as a sequence of individual 2D strain images.

In the volume-based method, we use an RSP6-12 mechanically swept 3D probe (GE Healthcare, Chalfont St. Giles, UK). The probe is held stationary by hand while the internal mechanism sweeps the transducer over a known angle. This produces an approximately evenly spaced sequence of B-scans forming a volume of ultrasound data. A small compression is then applied by hand, roughly in the axial direction, before recording another volume.

The signal processing is similar to that used in the freehand method, but the volume is treated as a whole rather than as individual 2D frames. This means displacement estimation windows can be 3D with an elevational component, and elevational displacements can be tracked in the same way as lateral displacements are tracked in the freehand method. Since each volume is recorded with the probe held stationary, the normalization function can be fitted to the entire volume of data. The 3D normalization function used in this work is

$$f_{3D}(x, y, z) = a(1 + by)(1 + cx + dz) \tag{2}$$

where z is the elevational component and d is an additional constant allowing for varying compression in the elevational direction. The volume-based approach also applies a 3D Gaussian smoothing filter as a final stage. In this case, this simply improves the signal to noise ratio, since quality variation caused by variable compression is not an issue with volume-based scanning [17].

2.2 Simulations

Many of the experiments in this paper make use of simulations of the two methods. The principal benefit of such an approach is that quantitative error metrics can utilize the known, ground truth strain. Previous approaches to simulation have often adopted a two-step method, first determining the displacement field created by an applied compression and then generating RF data based on pre- and post-compression geometry [22, 24, 25, 26]. Following this approach, we used a finite element method (Abaqus 6.7-1, Simulia, Rhode Island, USA) to calculate the displacements produced in a model of the tissue by a probe pushing on the surface. To fully capture the pertinent properties of the resultant displacement field, the finite element model was three-dimensional [3, 22]. We then simulated 3D ultrasound data using Field II [10, 11], with the post-compression scatterer locations distorted by the displacements found in the finite element stage.

Figure 1 shows the geometry of the two finite element models. Each comprises a cylindrical volume (14 cm diameter, 6 cm high) of elastic material representing the tissue, and a rigid rectangular compressor representing the probe face. The tissue was modelled as an isotropic, linear elastic material with a Young's modulus of 10 kPa and a Poisson's ratio of 0.495. In Figure 1, the tissue contains a spherical inclusion with a different stiffness, although the experiments in this paper make use of various inclusion geometries. The compressor dimensions are different in the two models, reflecting the probe footprints in our two approaches. The freehand probe face was modelled as a 4 cm $\times$ 1 cm compressor, whereas the 3D probe was represented by a larger 5 cm $\times$ 5 cm compressor. In each case, there were frictionless slip conditions at the base of the model and the contact region between the probe and the top surface. All other surfaces were unconstrained.

The volume-based scanning protocol requires that the position of the probe is adjusted only once, to change the applied compression between the two volumes. This was straightforward to model as a single downward movement of the large compressor, with one volume of RF data

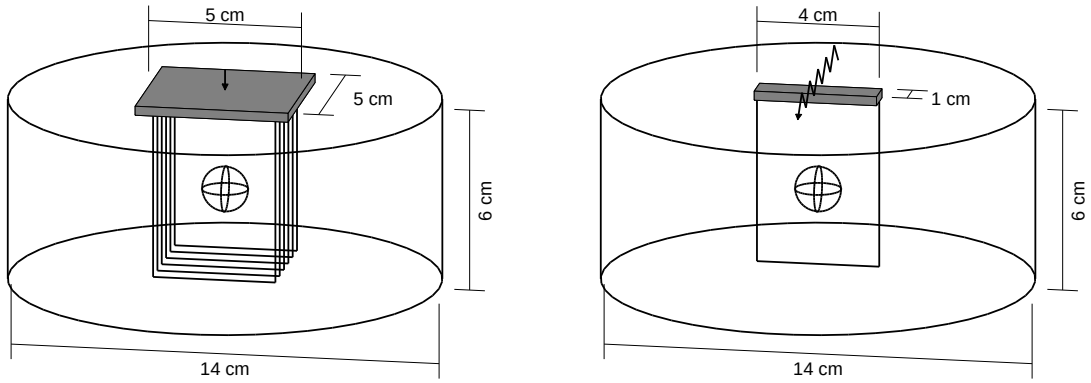

(a) **(b)**

Figure 1: **Geometry of the finite element simulations**. (a) shows the finite element model for the volume-based approach, which involves a large compressor applying a downward compression to a model of the tissue. (b) shows the model for the freehand approach. In this case, a smaller compressor applies a varying compression at different locations on the tissue surface, one for each image frame in the data set.

simulated on the uncompressed tissue and another on the compressed tissue. The downward movement in this case was 0.6 mm, producing a 1% axial compression of the tissue. Each simulated volume comprised 69 parallel frames of RF data centred on the compressor. The frame spacing was 0.3 mm, giving an elevational range of ± 10.2 mm. This is typical of a volume recorded by the RSP6-12 3D probe. However, it should be noted that the actual probe has a slightly curved surface and sweeps the transducer through an angle: these details are not captured by this simplified model.

In comparison, the freehand approach requires a larger set of finite element simulations, as each image frame is produced with the probe in a different location and with a different axial compression. We modelled this as a sequence of alternating compressed and uncompressed frames, with a fixed elevational spacing and no lateral probe movement between frame pairs. The axial displacement was again 0.6 mm, giving 1% compression. The frame spacing is one of the parameters under investigation and, as such, will be discussed in more detail in Section 3.

The Field II simulations were set up to model a typical ultrasound transducer with centre frequency 6.5 MHz. Each image comprised 127 A-lines at a pitch of 0.3 mm, giving a lateral B-scan width of 38.1 mm. Each A-line was sampled at 66.67 MHz to a depth of 44.0 mm. Both the lateral and elevational focal depths were set to 20 mm. The scattering density was constant throughout the tissue, so that any variations in stiffness in the finite element model were not visible in the simulated B-scans. For the strain SNR experiments described in Section 3.2.2, we also added Gaussian white noise, reducing the SNR of the RF signal to 20 dB. A typical example of a simulated B-scan and the associated strain image can be found in Figure 2.

3 Experiments and results

As explained in the introduction to this paper, we are interested in two comparisons. The first is how well the strain represents the underlying tissue stiffness, which depends to some extent on the stress field induced by the probe. The second is the trade-off between resolution and strain SNR: this is influenced by the scanning protocol and the signal processing inherent in the strain estimation algorithm. Experiments exploring both of these issues are presented in this section.

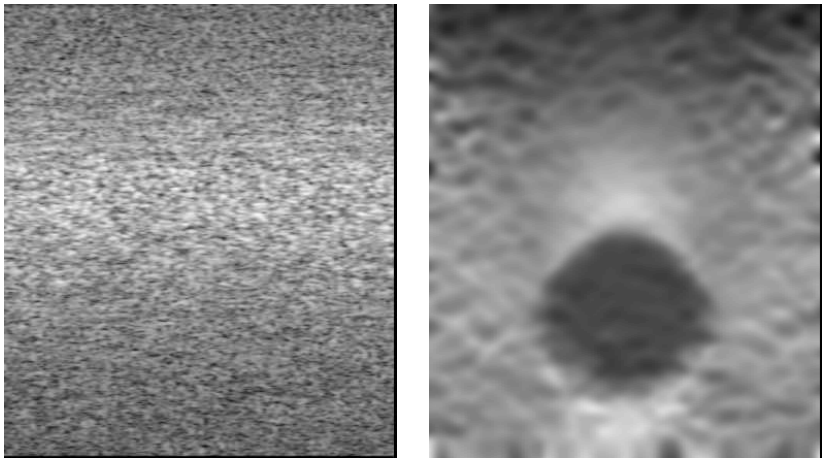

(a) **(b)**

Figure 2: **Simulated data**. The figure shows (a) a B-scan generated from simulated RF data and (b) a strain image derived from two frames of RF data. This example has a stiff, spherical inclusion (Young's modulus 40 kPa) in a uniform stiffness background (10 kPa). The inclusion is only visible in the strain image. The image size is 38.1 mm × 44.0 mm. The lateral and elevational focal depths are both 20 mm.

3.1 Mechanical stress artefacts

Mechanical stress artefacts were investigated in terms of the appearance of a spherical inclusion. For these experiments, the finite element model included a spherical region with a different stiffness, as in Figure 1. This region had a 1.5 cm diameter, was located at the centre of the model 3 cm below the tissue surface, and was bonded to the surrounding background material. The Poisson's ratio in this region was 0.495, the same as the background material, but the Young's modulus was set to a range of different values above and below the background value of 10 kPa.

Since we are interested in the fundamental mechanical properties of the two approaches, we generated the strain images directly from the displacements found in the finite element simulations. This avoids the inevitable errors introduced by displacement estimation from RF data. We initially looked at strain images without the location-varying normalization of Equations (1) and (2), thus displaying all strain images on the same scale. Figures 3(a–b) show the resulting images for the two approaches with an inclusion stiffness of 40 kPa (four times the background stiffness). Since we are also interested in the effect of normalization, Figures 3(c–d) show equivalent images using the location-varying normalization method. Finally, Figures 3(e–f) show normalized results for a soft inclusion (stiffness 2.5 kPa).

Considering first the unnormalized results, the images reveal several recognisable stress artefacts. In both cases, there are stress concentrations around the inclusion. Additionally, the freehand data exhibits a significant stress decay artefact, making it look softer at the top than at the bottom. This is less severe with the 3D probe, because the larger compressor area causes the stress to decay more slowly with depth [20]. When the data is normalized by the location-varying functions in Equations (1) and (2), the stress decay artefact is much reduced.

It is common, though of course technically incorrect, to interpret medical strain images by assuming a uniform stress field, since it is difficult to measure the actual stress field. Under this assumption, the axial strain is inversely proportional to stiffness. We are therefore interested in how well the strain images reflect the underlying tissue stiffness, and how the different stress fields induced by our two approaches affect this. In particular, while the inclusion has fairly uniform strain in all cases, the background does not. Also, while we cannot expect the strain ratio in the two regions to be quantitatively representative of the tissue stiffness, it is desirable that the difference in mean image values between the inclusion and the background is sufficiently high for

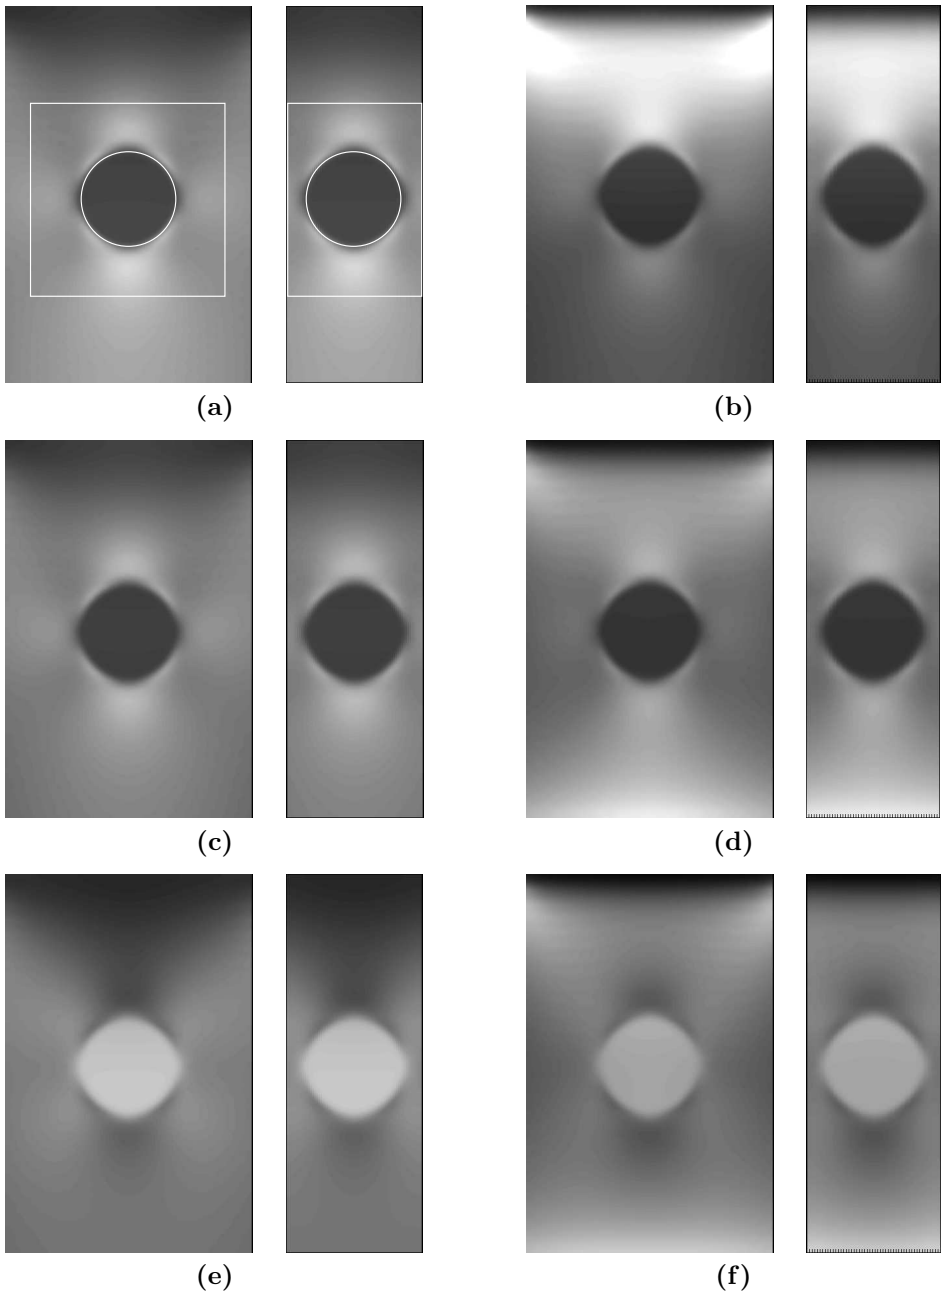

Figure 3: **Comparison of mechanical stress artefacts in strain images**. The images are derived from displacements taken directly from the finite element simulations. (a, c, e) show the data from the 3D probe model and (b, d, f) show similar results for the freehand model. (a–b) are unnormalized and therefore show strain data on the same scale (black is zero strain, white is 2.0% strain). (c–f) show results after normalization according to Equations (1) and (2). They are therefore not on equivalent scales, although each image spans the same range 0 (black) to 2 (white) of normalized values. The inclusion in (a–d) has a stiffness four times the background level. In (e–f) it is 0.25 times the background level. The white lines in (a) outline the 3D region of interest used to calculate contrast to noise ratio. In each subfigure, the left image is the central axial-lateral frame in the 3D sweep. The right image is an axial-elevational reslice along the length of the sweep. All images have a depth of 5.8 cm.

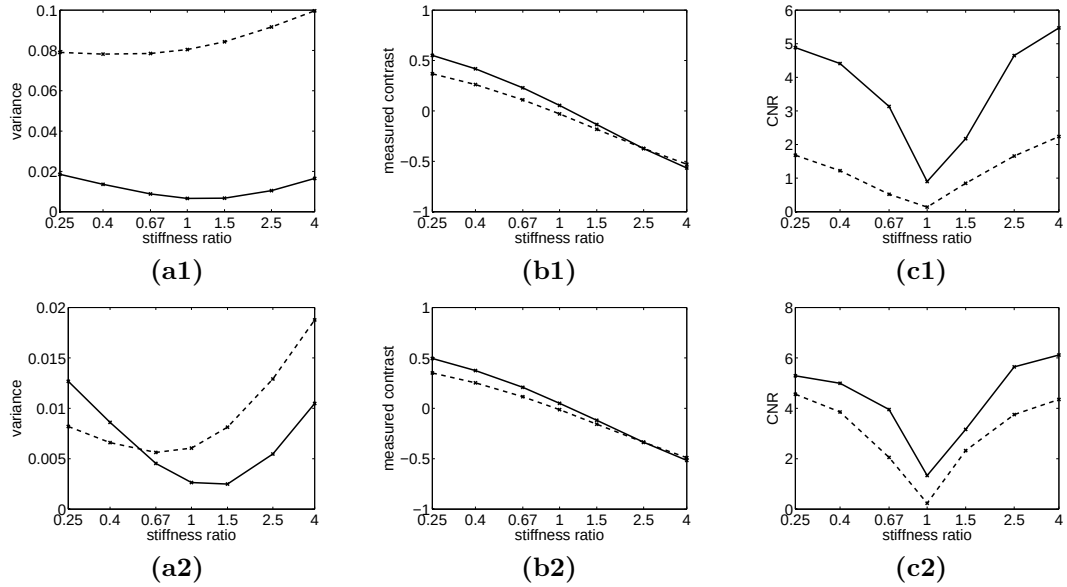

Figure 4: **Comparison of mechanical stress artefacts**. (a) shows the variance of the background region for the two techniques. (b) compares the difference in the mean strain value in each region. (c) shows the contrast to noise ratio between the two regions. In each graph, the solid line (—) gives the result for the volume-based approach and the dashed line (- -) is for the freehand approach. The upper row of graphs is without the location-varying normalization. Each strain value is instead scaled by 100 to give image values in the range 0 to 2. The lower row is with location-varying normalization, again giving image values in the range 0 to 2.

the inclusion to be easily detectable. For example, the soft inclusion in Figure 3 stands out more in the 3D probe image than in the freehand image.

Figure 4(a) shows the variance of the image values in the background region, as defined in Figure 3(a). Since we have no estimation noise in our images, the variance here is caused entirely by mechanical stress variations. Without normalization (Figure 3(a–b)), it is clear from the images that the 3D probe produces strain data with a more uniform background strain. This is due to the more significant stress decay with depth when using a smaller compressor. With the location-varying normalization (Figure 3(c–f)), both variances are reduced to some extent, bringing the freehand curve into a range comparable with that of the 3D probe.

Figure 4(b) shows the difference in the mean values of the two regions. For this measure, normalization does not make a significant difference. However, in both cases, the freehand method produces slightly lower contrast between the two regions for soft inclusions. With stiff inclusions, the freehand technique and 3D probe produce similar contrast.

A commonly used measure for comparing inclusion detectability is contrast to noise ratio (CNR) [5]. This is defined as

$$\mathrm{CNR} = \sqrt{\frac{2(\bar{s}_i - \bar{s}_b)^2}{\sigma_{si}^2 + \sigma_{sb}^2}}$$

where $\bar{s}$ is the mean value of strain and σ_s^2 is the variance of the strain. The subscripts i and b refer to the inclusion and the background regions respectively. This measure captures the combined effect of the variance and the contrast in Figures 4(a–b). Figure 4(c) shows the CNR for the two approaches, with and without normalization. The 3D probe produces images with superior CNR, over all stiffness ratios, both with and without normalization.

In summary, it is evident that without normalization, the smaller compressor is at a disadvantage because of the more significant stress decay with depth. Normalization improves this to some extent and, for this reason, we make use of normalization in all subsequent experiments in

this paper. However, it remains the case that the larger compressor produces images of inclusions that are more easily interpreted, even with location-dependent normalization.

3.2 Resolution and noise

The resolution and strain SNR of the two approaches were determined by making appropriate measurements from combined finite element/Field simulations. This approach provides realistic decorrelation due to scatterer displacement and allows for the fact that resolution depends partly on the parameters of the ultrasound system [25, 26]. It is well known that there is a trade-off between resolution and SNR [30]: the latter can always be improved at the expense of the former, for example by using larger windows for displacement estimation. In order to make a fair comparison, we must fix one of the two measures and compare the other. Here, we fix the resolution and compare the SNR of the two techniques.

3.2.1 Measuring resolution

To fix the resolution, we must determine parameters that result in both techniques resolving small features to the same degree. We can then ensure a fair comparison by using these parameters when assessing SNR. Strain image resolution is affected by several factors, including the imaging system's inherent resolution [25, 26], the elevational frame spacing and the signal processing parameters used to generate the strain image from the RF data (window size, window spacing, gradient estimation kernel size and filter size). The signal processing parameters provide a convenient route to directly *control* the resolution, though we still need some way to *measure* the resolution in order to know when the two techniques are fairly matched.

To this end, we produced finite element models with an array of small inclusions, as shown in Figure 5. There are eight stiff inclusions (Young's modulus 40 kPa) centred in the model at a depth of 2 cm (the focal depth). These inclusions are arranged in a $2 \times 2 \times 2$ array, so that there are four pairs of inclusions in each of the three directions. Each inclusion is an ellipsoid, with the lateral and elevational dimensions twice the axial dimension[1]. This size will be referred to as the *feature scale* in each direction. The inclusions are spaced so that there is a gap equal to the feature scale between them. This setup is an extension to 3D of the method used in [16], where bands of varying stiffness provided resolution estimates in one direction. We created five of these data sets, with axial feature scales of 0.75 mm, 0.875 mm, 1.0 mm, 1.125 mm and 1.25 mm, and lateral and elevational feature scales twice these values.

We then produced Field simulations at each feature scale, for both the 3D probe and a freehand scan. For the freehand data, the frames were distributed in pairs with zero separation within each pair and 0.3 mm separation between pairs. Although this scanning protocol is not practically feasible, it is useful here to give high quality strain estimates between adjacent pairs and therefore minimise the effect of decorrelation noise on the resolution measurements. Since the 0.3 mm frame pair spacing is well below the smallest elevational feature scale and therefore the smallest resolving limit considered here, the choice of frame spacing does not affect the resolution significantly. Having determined appropriate signal processing parameters on this ideal data set, we can then apply the same parameters to any freehand frame spacing, provided the spacing stays well below the elevational resolving limit.

Since we need to fix the resolution in all three directions, we must be able to adjust the elevational smoothing in both approaches. We therefore used the final filtering stage (see Section 2.1) to set the resolution, since this is applied elevationally even in the freehand approach. The other signal processing parameters that can affect resolution are the displacement estimation window size, window spacing and the gradient estimation kernel size. These were set to low values, so that they had a relatively small effect on the resolution. The window size was 6 RF cycles axially

[1] This finite element model was originally designed to assess B-mode ultrasound image resolution, accounting for the superior axial resolution relative to the other two directions: hence the ellipsoids. While a spherical inclusion would have been more natural for this work on strain imaging, and would have led to more uniform filter lengths in Figure 7, the use of ellipsoids does not materially alter any of this paper's results or conclusions.

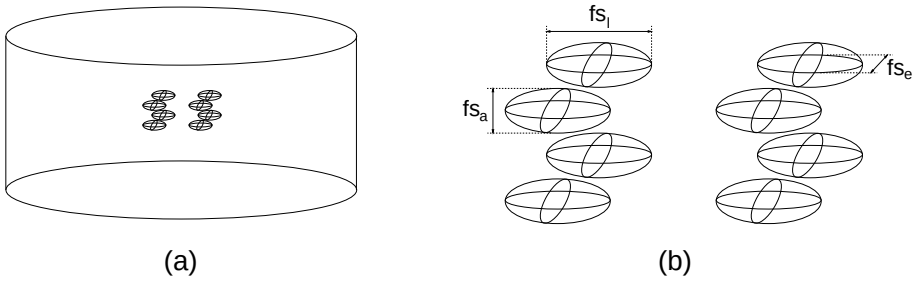

Figure 5: **Finite element model for measuring resolution**. (a) shows the geometry of the finite element model used for determining resolving limits. It has an array of small inclusions centred under the compressor at 2 cm depth. (b) is a magnified view of the inclusions, showing the arrangement of the eight inclusions in a $2 \times 2 \times 2$ array. The feature scales in the axial, lateral and elevational directions (fs_a, fs_l and fs_e respectively) are defined as the inclusion size in that direction. There are four pairs of inclusions in each direction, with a gap equal to the feature scale between them. The lateral and elevational feature scales are twice the axial feature scale.

(± 0.35 mm), ± 1 A-line laterally and, in the case of volume-based processing, ± 1 frame elevationally. While this is small, it does extend beyond a single frame elevationally, so there is a difference between the freehand and volume-based windows. The window spacing was set to 0.2 mm, well below the feature scale in any direction. Finally, the gradient was calculated by a simple differencing operation which again has only a small effect on the resolution.

We measured the resolution in a particular direction by considering the pairs of inclusions in that direction. We looked at resolvable contrast, defined as the difference between the mean strain value within a pair of inclusions and in the region in between, along a line through their centres. As the level of smoothing increases, this contrast decreases, until the inclusions are no longer individually distinguishable. Since we have an array of inclusions, we measured the average contrast across each of the four pairs. Figure 6(a–h) shows an example of this, for a lateral feature scale of 2 mm and four different Gaussian filter widths of 2.0 mm, 3.6 mm, 5.0 mm and 6.0 mm (filter width is defined as twice the standard deviation of the Gaussian). The filter sizes in the axial and elevational directions were set equal to the feature scale in that direction, which is well above the resolving limit for a Gaussian filter. The third of these filters (5.0 mm) gives zero contrast between the inclusions and the background: they are no longer individually distinguishable laterally in the image (Figure 6(g)). For larger filter sizes, the contrast becomes negative.

For the purposes of comparing SNR, it is not critical where we set the resolutions, provided that they are matched. We could, for example, compare SNR at the resolving limit where the contrast crosses zero and the inclusions become indistinguishable. However, in order to provide visual confirmation that the resolutions are indeed matched, we instead set the resolution at a level of smoothing just below that required to merge the two inclusions. This amount of smoothing is still sufficient to ensure that the filter size dominates the resolution.

Appropriate parameters were chosen by plotting the contrast in one direction against the Gaussian filter width in that direction, with the filter widths fixed at the feature scale in the other two directions. Figure 6(i) shows an example of this kind of curve. We then compared the measured contrast to the actual contrast (determined from the finite element output) and took our desired filter setting to be where the average measured contrast was 5% of the actual. Repeating this in all three directions for our two techniques, we determined matched filter settings at each feature scale. When these three settings are applied simultaneously to the data, the inclusions are just visible in the images (see Figure 6(j–k), preferably on a computer monitor, not in hard copy).

Figure 7 shows the Gaussian filter width settings in each direction. As expected, there is little difference in the required settings for the two techniques, because the resolution determining step (the Gaussian filter) works very similarly in both cases. Axial settings are around half the lateral and elevational settings. Again, this is what we would expect, given the ellipsoidal inclusions.

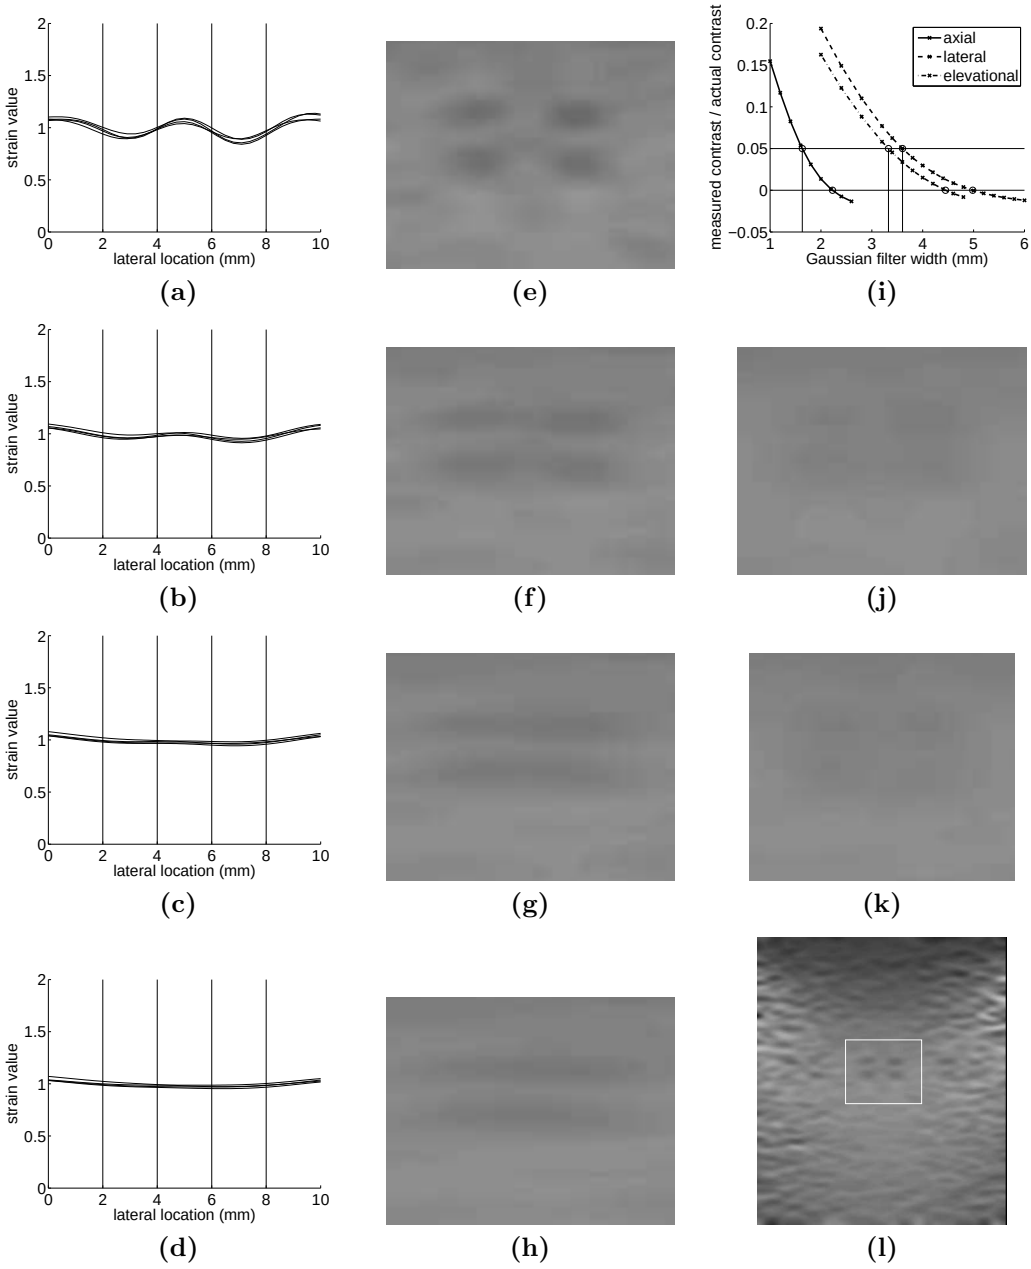

Figure 6: **Setting filter sizes**. (a–d) show how the strain varies along a line through the centre of the lateral pairs of 2.0 mm inclusions with four different Gaussian filter widths (2.0 mm, 3.6 mm, 5.0 mm and 6.0 mm). The vertical lines on the graphs indicate the true edges of the inclusions. (e–h) show axial-lateral views of the inclusions at each filter setting, magnified for the region marked in (l). (i) shows the measured contrast at each filter setting, for all three directions. From this graph, the 5% relative contrast level provides the filter settings for the subsequent SNR experiments. (j–k) are axial-lateral and axial-elevational views respectively for the 5% contrast filter settings, with the inclusions just resolvable (on a computer monitor) in each direction.

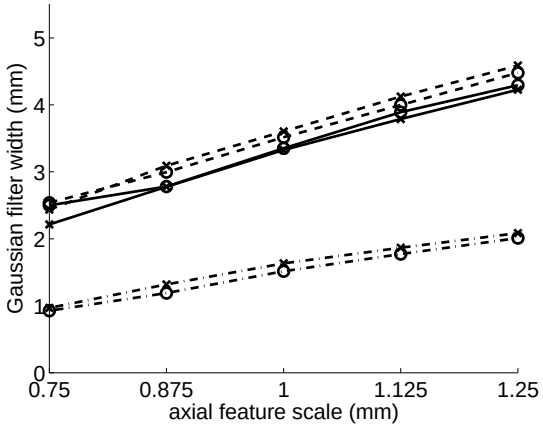

Figure 7: **Filter settings for matched resolution**. The graphs show the filter width, defined as twice the standard deviation of the Gaussian, required to give matched resolution just above the resolving limit for the two techniques. The dash-dot lines (-·-) are for the axial filter, dashed lines (- -) are for the lateral filter and solid lines (—) are for the elevational filter. In each pair, crosses (×) mark the line for the volume-based approach and circles (○) are for the freehand approach. The lateral and elevational feature scales are twice the axial feature scales.

3.2.2 Comparing SNR

Strain precision depends on the number of independent strain estimates captured by the Gaussian smoothing kernel (that is, the size of the filter relative to the data density) and on the precision of the individual estimates. The data density is therefore an important consideration, because, in the freehand technique, the elevational frame spacing can be freely varied by moving the probe at different speeds. The precision of the individual estimates depends on the window size and also on the correlation between the RF data within the windows. Correlation will be low when there is motion in an untracked direction, for example the elevational direction in freehand data, and also in regions of high deformation.

Strain signal to noise ratio is conventionally defined as

$$\mathrm{SNR}_e = \frac{\bar{s}}{\sigma_s}$$

where $\bar{s}$ is the mean strain level and, in data with uniform strain, σ_s is the standard deviation of the strain estimates. However, in order for our SNR measurements to capture the effects of deformation variations between our two approaches (and any consequent variations in estimation precision), we used Field-simulated data derived from finite element models with uniform stiffness. Then, to avoid biasing the σ_s value with true variations in strain, we set σ_s to the standard deviation of the strain *error*, with the finite element simulation providing the ground truth. For $\bar{s}$, we used the mean value of the estimated strain data. Our SNR_e metric thus captures strain estimation precision for a uniform stiffness data set.

Several estimates of SNR were obtained from different regions of each data set. We selected individual frames of data distributed through the sequence so that even at the largest filter size, they still provided almost independent estimates of SNR. The largest elevational filter size was 4.297 mm (see Figure 7), which allowed for up to five frames to be selected from the 2 cm elevational range[2]. Within each of these frames, we divided the image into two regions, as shown in Figure 8, giving a total of ten SNR estimates from each simulated data set.

[2]To be entirely independent, the spacing would need to be twice the filter width, because the Gaussian kernel was truncated two standard deviations from the centre. With the spacing used here, there is a little overlap at the kernel edges, though the small coefficients at the tails mean that the SNR estimates are almost independent.

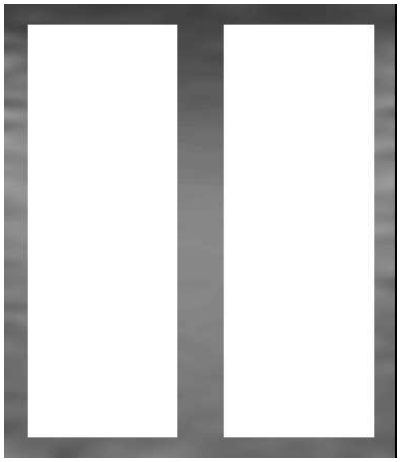

Figure 8: **Regions used to calculate SNR estimates**. Within a single frame, the two regions indicated by white boxes were used to calculate two estimates of SNR. They are separated laterally and leave a gap at the edges to allow for the filter capturing values outside these regions. Within each simulated data set, five frames like this were used to obtain ten estimates of SNR.

Since we need to consider the effect of varying frame density in the freehand data, we generated several freehand data sets with different frame spacings of 0.3 mm, 0.2 mm and 0.1 mm. We also produced one additional freehand data set with pairs of frames at zero separation and 0.3 mm spacing between pairs. While this scanning protocol is not feasible in practice, it offers a good comparison with the volume-based data set, as it has the same frame spacing and density. For this data set, SNR was calculated based only on the overlaid frame pairs, without including any strain estimates from the 0.3 mm-separated frames.

Figure 9 shows the SNR_e values measured in each data set at the five different resolution settings. There are several observations that can be made. First, the SNR of the freehand approach depends very much on the frame spacing and density, and can be greater than or less than that of the 3D probe. As expected, the SNR improves as the freehand spacing decreases, because there are then a larger number of better quality estimates within the filter range. The freehand data with zero-separation frame pairs every 0.3 mm actually has a lower SNR than the 0.1 mm evenly spaced data. This is because, even though the individual strain estimates are better with zero spacing, there are fewer of them within the filter range and this reduces the SNR.

Comparing with the volume-based result, the most similar freehand data set is the one with zero spacing, having the same frame spacing and density as the volume-based data set. The SNR is therefore very similar in these two results. There are, of course, many differences between the two, both mechanically, in the way the tissue is distorted by the probe and subsequently normalized, and also in the signal processing, where the volume-based approach uses 3D displacement estimation windows and normalizes over the entire data set. The combined effects of these differences give the volume-based approach a slightly lower SNR.

In practice, it would not be possible to obtain the zero-separation data set in a freehand scanning protocol. Considering the other freehand sequences, it is apparent that the data would need to be around twice as densely sampled as that recorded by the 3D probe in order to obtain a similar SNR. Moreover, considering that these simulations produced idealised freehand data, with no lateral probe motion and a good compression magnitude between each pair, it is likely that an even higher density would be required in practice.

3.3 *In vitro* validation

While the finite element simulations allowed quantitative assessment against a known ground truth, we must check that the results translate to practical strain imaging with a handheld probe.

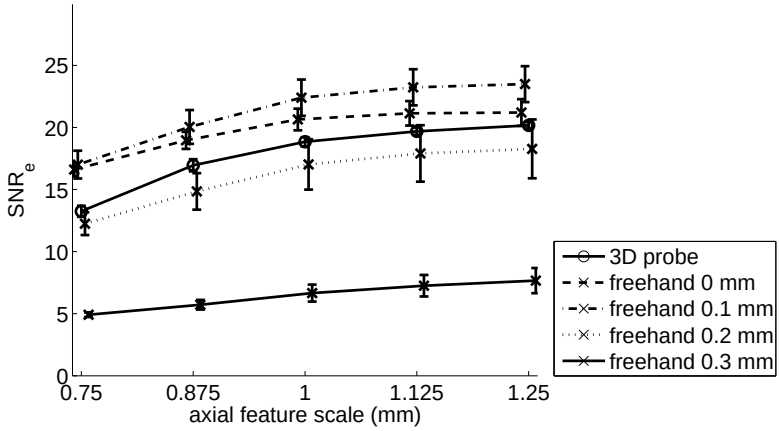

Figure 9: **Comparison of signal to noise ratio**. The graph shows the mean SNR_e values for the freehand approach at four different frame spacings, compared with the volume-based approach. The signal processing parameters were set to give matched resolutions in all three directions. The error bars show $\pm$ one standard deviation of the SNR_e estimates.

We therefore recorded some real data on an elasticity QA phantom (model 049, CIRS, Norfolk, Virginia, USA). The freehand scans were performed with a 8-16 MHz probe, which closely matches the transducer inside our 3D probe, interfaced to a Dynamic Imaging Diasus ultrasound machine. The location of each RF image was recorded using a Polaris optical tracking system (NDI, Waterloo, Ontario, Canada). Two scans were recorded, one with the probe moved deliberately slowly and carefully, the other with a faster and more practically viable scanning speed. The mean frame spacing in the slow scan was around 0.06 mm and in the fast scan was 0.13 mm. At a frame rate of around ten frames per second, this corresponds to a probe speed of 0.6 mm per second and 1.3 mm per second respectively[3]. For the volume-based scan, we used the RSP6-12 mechanically swept 3D probe interfaced to the Diasus scanner. Each volume comprised 50 frames and was swept over a range of 10 degrees, giving a spacing of 0.34 mm at the frame centre.

For all scans, the image frame dimensions were 33 mm axially by 26.7 mm laterally. The 3D probe had a footprint of approximately 5 cm laterally by 5.5 cm elevationally, and the freehand probe face was 3.5 cm by 0.9 cm. The filter widths were set to 1.4 mm in all directions: Figure 7 shows how this will achieve approximately matched, isotropic resolutions for the two techniques. The scanning target was a 10 mm diameter stiff inclusion, having a stiffness 3.2 times the background value and located at a depth of 15 mm.

Figure 10 shows the strain images produced from these three data sets. The slow freehand data clearly has the best SNR, while the 3D probe data appears to be slightly better than the fast freehand data. These results confirm that the freehand quality depends very much on the scanning speed. Furthermore, it would appear that in order to match the image quality of the 3D probe, the frame density must be higher than that suggested by the simulations. This is because the simulations do not account for many complications of the freehand method that would reduce the SNR, such as lateral motion. In our experience, it is easier to achieve the idealised conditions of the simulations using the 3D probe.

The disadvantages of the freehand protocol can be seen in the elevational reslices. The 3D probe always produces a well-sampled volume, whereas the freehand scanning protocol requires that the probe motion is carefully controlled over the length of the scan. The low quality regions in the fast freehand scan are caused by a lack of high quality strain estimates in that region, due to insufficient frame density, inappropriate axial compression, or a combination of the two.

[3]The frame rate is limited by the digitization pipeline in our research rig. Any commercial implementation would be able to achieve at least ten times this frame rate.

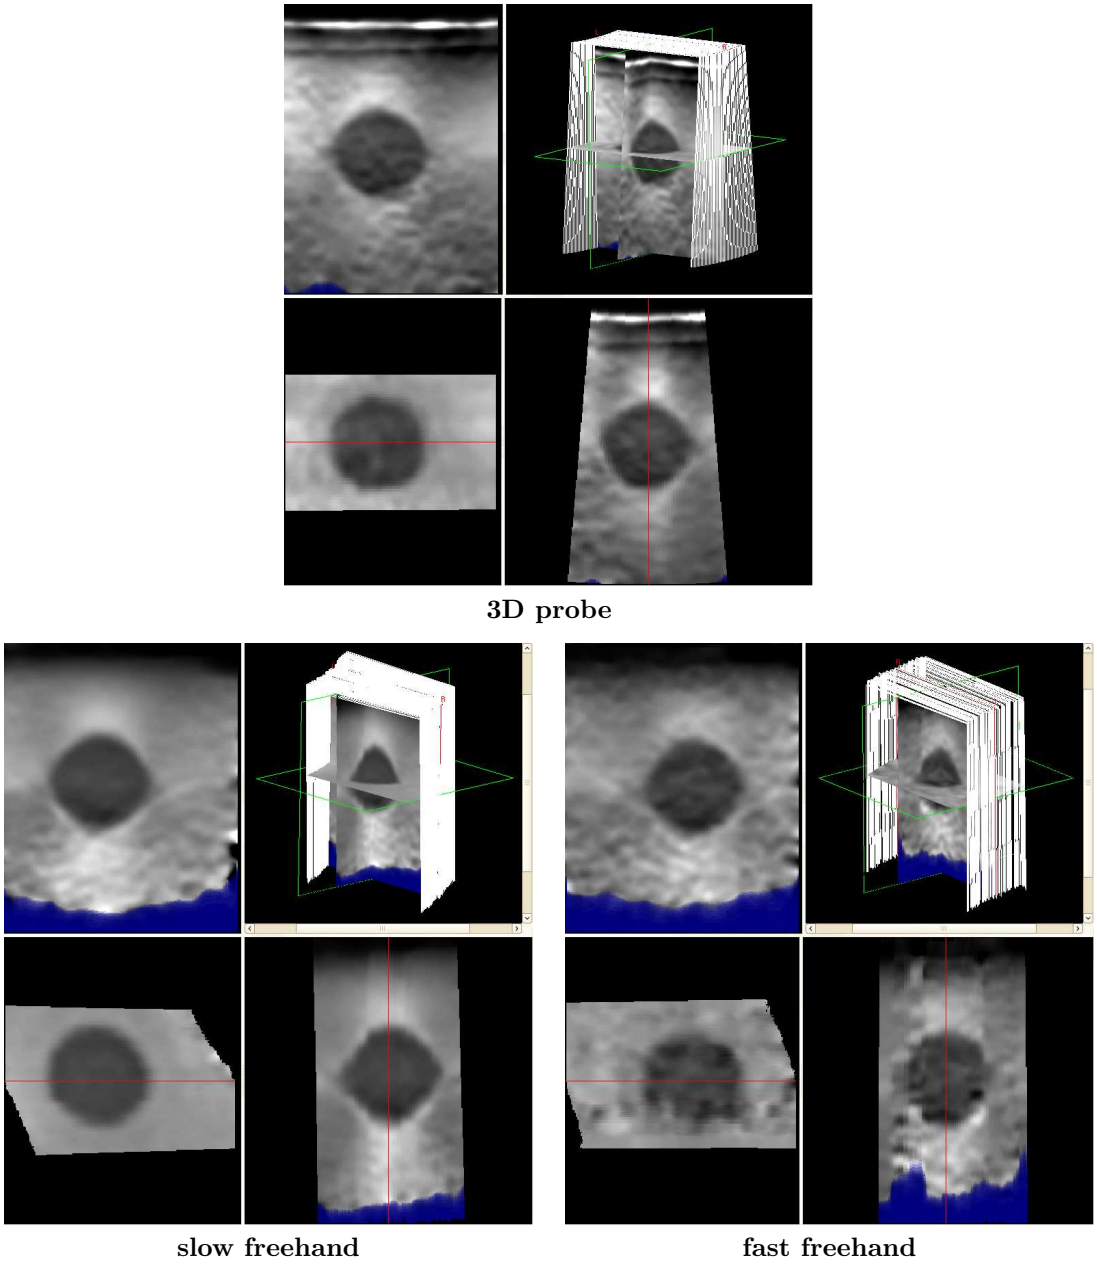

Figure 10: ***In vitro* comparison**. The figure shows examples of *in vitro* data sets for our two scanning techniques. For the freehand data, the target was scanned with the probe moving at two different speeds. In the slow scan, the average frame spacing is around 0.06 mm. In the fast scan, it is 0.13 mm. In each subfigure, the top left image is an axial-lateral view centred on the stiff inclusion, the bottom left is a lateral-elevational view, the bottom right is an axial-elevational view and the top right is a 3D view showing the outline of each B-scan and all three slices. The blue colour wash obscures low quality regions where the strain estimation precision falls below a certain threshold.

In general, these results validate the finite element simulations' findings that, for a given frame density, the volume-based approach yields superior results. Note, however, that the CIRS phantom is a relatively easy scanning target. In clinical practice, there are often situations where access to the target is limited, for example in neurosurgery, and a smaller probe face might offer an advantage. Even for a relatively accessible target like the breast, it can be difficult to maintain contact over the entire curved surface of the 3D probe, whereas this presents no problems for a smaller probe. So there may be some practical scenarios where a carefully controlled freehand protocol is preferable to the volume-based approach.

4 Conclusions

We have compared two practically feasible approaches to 3D strain imaging: freehand scanning with a conventional ultrasound probe, and volume-based acquisition with a mechanically-swept 3D probe. In terms of mechanical stress artefacts, we showed that the 3D probe, which has a larger footprint, produces strain images with more uniform strain in regions with uniform stiffness, and hence more easily identifiable inclusions. We also compared the two approaches in terms of resolution and noise. While it is possible to produce high quality strain images using the freehand method, this requires a considerably higher frame density than with the volume-based technique. The freehand scanning protocol is also more difficult to master in practice, but may nevertheless be preferable in clinical scenarios where access to the scanning site is difficult for a large 3D probe.

References

[1] S. I. Awad and J. T. Yen. 3-D strain imaging using a rectilinear 2D array. In *IEEE Ultrasonics Symposium, 2007*, pages 228–231, New York, USA, October 2007.

[2] S. Bharat, T. G. Fisher, T. Varghese, T. J. Hall, J. Jiang, J. A. Madsen, E. L. Zagzebski, and F. T. Lee Jr. Three-dimensional electrode displacement elastography using the Siemens C7F2 *four*Sight four-dimensional ultrasound transducer. *Ultrasound in Medicine and Biology*, 34(8):1307–1316, August 2008.

[3] M. Bilgen and M. F. Insana. Elastostatics of a spherical inclusion in homogeneous biological media. *Physics in Medicine and Biology*, 43(1):1–20, January 1998.

[4] E. S. Burnside, T. J. Hall, A. M. Sommer, G. K. Hesley, G. A. Sisney, W. E. Svensson, J. P. Fine, J. Jiang, and N. J. Hangiandreou. Differentiating benign from malignant solid breast masses with US strain imaging. *Radiology*, 245(2):401–410, November 2007.

[5] P. Chaturvedi, 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, July 1998.

[6] C. L. de Korte, A. H. W. van der Steen, E. I. Céspedes, G. Pasterkamp, S. G. Carlier, F. Mastik, A. H. Schoneveld, P. W. Serruys, and N. Bom. Characterization of plaque components and vulnerability with intravascular ultrasound elastography. *Physics in Medicine and Biology*, 45(6):1465–1475, June 2000.

[7] T. G. Fisher, T. J. Hall, S. Panda, J. Jiang, J. Resnick, S. Barnes, and E. L. Madsen. Volume elasticity imaging with a 2-D capacitive micro-machined ultrasound transducer (CMUT) array. In *Proceedings of the Fifth International Conference on the Ultrasonic Measurement and Imaging of Tissue Elasticity*, page 125, Snowbird, UT, USA, October 2006.

[8] T. G. Fisher, J. Jiang, and T. J. Hall. Volumetric strain imaging. In *IEEE Ultrasonics Symposium, 2007*, pages 355–358, New York, USA, October 2007.

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

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

[11] J. A. Jensen and N. B. Svendsen. Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 39(2):262–267, March 1992.

[12] F. Kallel, M. Bertrand, and J. Ophir. Fundamental limitations on the contrast-transfer efficiency in elastography: an analytic study. *Ultrasound in Medicine and Biology*, 22(4):463–470, 1996.

[13] M. Krueger, A. Pesavento, H. Ermert, K. M. Hiltawsky, L. Heuser, H. Rosenthal, and A. Jensen. Ultrasonic strain imaging of the female breast using phase root seeking and three-dimensional "optical flow". In *Proceedings of the IEEE Ultrasonics Symposium, 1998*, volume 2, pages 1757–1760, Sendai, Japan, October 1998.

[14] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. 3D elastography using freehand ultrasound. *Ultrasound in Medicine and Biology*, 32(4):529–545, April 2006.

[15] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Estimation of displacement location for enhanced strain imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 54(9):1751–1771, September 2007.

[16] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Dynamic resolution selection in ultrasonic strain imaging. *Ultrasound in Medicine and Biology*, 34(5):809–823, May 2008.

[17] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. An intelligent interface for freehand strain imaging. *Ultrasound in Medicine and Biology*, 34(7):1117–1128, July 2008.

[18] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Phase-based ultrasonic deformation estimation. *IEEE transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 55(1):94–111, January 2008.

[19] A. Lorenz, A. Pesavento, M. Pesavento, and H. Ermert. Three-dimensional strain imaging and related strain artifacts using an ultrasonic 3D abdominal probe. In *Proceedings of the IEEE Ultrasonics Symposium, 1999*, volume 2, pages 1657–1660, Caesars Tahoe, NV, USA, October 1999.

[20] J. Ophir, I. Céspedes, H. Ponnekanti, Y. Yazdi, and X. Li. Elastography: a quantitative method for imaging the elasticity of biological tissues. *Ultrasonic Imaging*, 13(2):111–134, April 1991.

[21] A. V. Patil, C. D. Garson, and J. A. Hossack. 3D prostate elastography: algorithm, simulations and experiments. *Physics in Medicine and Biology*, 52(12):3643–3663, June 2007.

[22] A. V. Patil, K. A. Krouskop, J. Ophir, and S. Srinivasan. On the differences between two-dimensional and three-dimensional simulations for assessing elastographic image quality: a simulation study. *Ultrasound in Medicine and Biology*, 34(7):1129–1138, July 2008.

[23] H. Ponnekanti, J. Ophir, and I. Cespedes. Ultrasonic imaging of the stress distribution in elastic media due to an external compressor. *Ultrasound in Medicine and Biology*, 20(1):27–33, 1994.

[24] H. Ponnekanti, J. Ophir, Y. Huang, and I. Céspedes. Fundamental mechanical limitations on the visualization of elasticity contrast in elastography. *Ultrasound in Medicine and Biology*, 21(4):533–543, 1995.

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

[26] R. Righetti, S. Srinivasan, and J. Ophir. Lateral resolution in elastography. *Ultrasound in Medicine and Biology*, 29(5):695–704, May 2003.

[27] J. A. Schaar, C. L. de Korte, F. Mastik, C. Strijder, G. Pasterkamp, E. Boersma, and P. W. Serruys. Characterizing vulnerable plaque features with intravascular elastography. *Circulation*, 108(21):2636–2641, November 2003.

[28] J. A. Schaar, C. L. de Korte, F. Mastik, L. C. A. van Damme, R. Krams, P. W. Serruys, and A. F. W. van der Steen. Three-dimensional palpography of human coronary arteries. Ex vivo validation and in-patient evaluation. *Herz*, 30(2):125–133, February 2005.

[29] R. Souchon, O. Rouvière, A. Gelet, V. Detti, S. Srinivasan, J. Ophir, and J-Y. Chapelon. Visualisation of HIFU lesions using elastography of the human prostate in vivo: preliminary results. *Ultrasound in Medicine and Biology*, 29(7):1007–1015, July 2003.

[30] S. Srinivasan, R. Righetti, and J. Ophir. Trade-offs between the axial resolution and the signal-to-noise ratio in elastography. *Ultrasound in Medicine and Biology*, 29(6):847–866, June 2003.

[31] G. M. Treece, J. E. Lindop, A. H. Gee, and R. W. Prager. Freehand ultrasound elastography with a 3-D probe. *Ultrasound in Medicine and Biology*, 34(3):463–474, March 2008.

[32] M. Vogt and H. Ermert. 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, March 2005.