

---

**AN INTELLIGENT INTERFACE  
FOR FREEHAND STRAIN IMAGING**

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

**CUED/F-INFENG/TR 578**

**May 2007**

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

**Email: [jel35/gmt11/ahg/rwp@eng.cam.ac.uk](mailto:jel35/gmt11/ahg/rwp@eng.cam.ac.uk)**

---

# AN INTELLIGENT INTERFACE FOR FREEHAND STRAIN IMAGING

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 a new, intelligent interface for freehand strain imaging, which has been designed to support clinical trials investigating the potential of ultrasonic strain imaging for diagnostic purposes across a broad range of target pathologies. The aim with this interface is to make scanning easier, and to help clinicians learn the necessary scanning technique quickly, by providing real time feedback indicating the quality of the strain data as they are produced. The images are also easier to interpret, because data at unacceptably low signal-to-noise ratios do not reach the display. Overall, the interface also considerably reduces the difficulty in producing volumes of strain data from freehand 3D scans. Its main components are novel normalisation, persistence and display methods. These not only present data in a more meaningful format, but the level of noise in the displayed images is actually reduced compared to other methods that use the same strain estimates with the same level of persistence.

## 1 Introduction

Ultrasonic strain imaging is an emerging technique, which is likely to have numerous applications in the clinical examination of soft tissues. In this paper we are primarily interested in the subset of elasticity imaging techniques that are categorised as “static” or “quasistatic” strain imaging [19]. In this paradigm, small tissue deformations are caused by varying pressure between the ultrasound probe and the tissue surface; two or more ultrasound frames are recorded during this deformation, and some form of tracking is applied to the recorded ultrasound data to estimate tissue deformations, amounting to displacement fields that vary with position. Spatial derivatives of such a displacement field are tissue strain, which indicates stiffness; there are sometimes further stages of analysis to estimate quantitative tissue properties directly, such as elastic moduli [10]. Quasistatic strain imaging was first tested clinically for breast scanning [4] and breast screening has ever since been a key driver for research [7, 23, 24]. Numerous studies have been motivated by prostate screening [18, 21]. Detection and staging of deep vein thrombosis also seems particularly promising [3], and there are many other possible applications.

One of the engineering challenges in strain imaging is the development of a suitable clinical interface. Ultrasound clinicians have extensive experience with existing scanning modes including B-mode greyscale, colour Doppler and power Doppler. Given the highly interactive nature of ultrasound examinations, the established modes have advantages in that clinicians are already well practised in the required scanning techniques, understand the significance of typical images, and are generally familiar with the uses, benefits and disadvantages of each mode. The likelihood of an addition to the ultrasound tool-set gaining clinical favour may be boosted if it possesses an interface that is practically helpful: actively fostering the development of a successful scanning technique, by providing either visual or audio feedback; displaying data in an intuitively meaningful format; and automatically guarding against the presentation of misleading data.

The aforementioned issues concern *how* we present information. We may also consider *what* information to present. This raises at least two further issues. Qualitatively, what type of information can be provided (stiffness, strain, or an alternative compromise)? Quantitatively, how much data should be amalgamated to form each display image? This latter question is relevant to many types of imaging system, particularly those pertaining to time series data (where persistence




Figure 1: Flowchart illustrating aspects of the interface that will be discussed.

may help, whether to improve a real time display during acquisition or for post-processing) and to volumetric data (where spatial averaging can be applied to reduce the level of noise).

Regarding the type of information, we note that ultrasonic strain imaging falls within a broader set of emerging elasticity imaging techniques. These are all essentially concerned with mechanical properties such as tissue stiffness, of which strain is only an indicator. Strain measurements can be converted into stiffness estimates if the stress field is known, but it is highly unlikely that this can be inferred from either static or quasistatic deformation data without reducing the resolution and imposing certain limiting assumptions [1]. Furthermore, such assumptions are unlikely to hold even approximately under *in vivo* scanning conditions, especially not with freehand scanning. On the other hand, strain images can be misleading, because an interpretation of low strain as indicating relatively high stiffness may be erroneous if the stress field varies substantially throughout the tissue [13, 19]. Some types of stress field variation occur repeatedly, and can hence be adjusted for. We will discuss the use of strain normalisation that varies both between images and within every individual image, so as to reduce the ambiguity of strain. The modified data after non-uniform normalisation are referred to as “pseudo-strain”.

In practice, an often more severe obstacle in freehand strain imaging is the basic challenge of achieving an acceptable strain estimation signal-to-noise ratio. Although many frames individually produce good images, typically a substantial fraction (sometimes a majority) of frames may be difficult to interpret because of high estimation noise. One of the common approaches to noise reduction amounts to averaging a sequence of strain images [27]. Rather than crude frame averaging, we present a more sophisticated weighting approach, which we use for persistence in the real time display, and for spatial averaging in the display of volumetric data [26].

The goal of this report is to describe aspects of a novel interface that we have developed to support a wide-ranging clinical trial of ultrasonic strain imaging.<sup>†</sup> The new interface tackles all of the issues mentioned above, to improve the quality of data that clinicians can acquire, and to improve the interpretability of the display. We present results based on example images that demonstrate the effects of all aspects of the interface, using recorded RF ultrasound data from freehand scans of *in vitro* and *in vivo* targets.

## 2 Method

The interface that we outline here is applicable to any static or quasistatic strain imaging system, almost regardless of the approach taken in the earlier stages of signal processing. It is likely to be particularly valuable in conjunction with freehand scanning. We provide illustrations based on an example, in which displacement tracking is by Weighted Phase Separation [14] with Amplitude Modulation Correction [16], and axial strain estimation is performed by piecewise-linear least squares regression [11]. This offers a good demonstration, not primarily because of its competitive estimation accuracy, but more importantly because it has already been analysed and tested rigorously, resulting in a promising method for predicting the strain estimation variance [15, 17]. Nonetheless, the aim of this paper is to describe our interface concept in general; the reader may envisage numerous specific applications. We now provide an overview of the interface as a whole. This is followed by a brief discussion of predicting estimation precision, and descriptions of each of the three subsequent stages of processing in the interface — normalisation, persistence or spatial averaging, and display (see Figure 1).

<sup>†</sup>Several aspects of this interface are covered by a recent UK patent application.




Figure 2: Typical images during freehand scanning without the new interface. (a) The strain display is filled with noise before the probe comes into contact with tissue, although (b) the B-mode image shows that there is obviously nothing to be seen. However, the images are actually more difficult to interpret when the probe is in contact with the tissue. For example, while scanning a thyroid (c) the strain display contains a mixture of good estimates alongside regions of noise, without an obvious boundary. This is unsurprising given that (d) the B-scan has regions with a high signal-to-noise ratio alongside other regions where there is simply no signal, and a region of severe decorrelation around the artery caused by blood flow and pulsatile motion.

## 2.1 Interface concept

Strain image quality varies substantially depending on the sonographer’s scanning technique, physiological motion in the tissue, and changes in the analytical parameters for converting RF ultrasound data into strain data. In order to produce consistently meaningful images, these parameters need to be controlled locally so as to adjust for different conditions during the scan. [15] and [17] describe such a system. However, adjustment of parameters cannot alone overcome all of the difficulties associated with practical strain imaging. For a start, at some stage it becomes impossible to produce meaningful deformation data from frames that are extremely weakly correlated. An adequate minimum level of correlation may not always arise, depending on the scanning technique, and with a very poor technique it may not even occur often. Even in the majority of frames where a uniform estimation signal-to-noise ratio can be achieved by adjusting the resolution settings, it is desirable to improve the quality of the recorded ultrasound data, so as to achieve the maximum resolving power. The best data may arise from relatively substantial deformations (*i.e.*, typically a large fraction of 1%, sometimes lower or higher depending on the target) accompanied by relatively low decorrelation. The acquisition of good data therefore depends on the combined properties of the scanning technique and the tissue.

Strain imaging with a typical interface requires a high level of expertise, in terms both of scanning technique and of image interpretation. The examples in Figure 2 illustrate some common difficulties. The use of a side-by-side display has been suggested [6], with the B-scan next to the strain image, because it is then easier to match strain data with features of interest that have already been identified in the B-scan. It also means that in cases where there may be little or no data, owing to an absence of coupling to the tissue, as for example in Figure 2a–b, then the sonographer knows to ignore the strain display. Image interpretation may nonetheless be rather difficult when the coupling is good, as for example in Figure 2c–d, because some scan targets do not offer signals suitable for strain imaging throughout the entire image.

One approach that to some extent handles this problem is a display in which each strain image is overlaid on the B-scan as a “colour wash”, where colour indicates strain and brightness is partly determined by the ultrasound signal amplitude [7, 18]. In so far as ultrasound signal amplitude correlates with the accuracy of strain estimates, this goes some way to indicating the quality of the strain data. It only helps to a limited extent, however, since signal amplitude is a very weak indicator of overall decorrelation. While a complete absence of signal would certainly mean that strain estimates were dominated by noise, it is often the case that strain estimates

from regions with medium signal amplitude are less noisy than other estimates where the signal is stronger. Furthermore, the blend of strain with B-mode data could actually make insightful image interpretation more difficult, by mixing strain data with fine features of B-scans such as the speckle pattern, that are in fact not generally related to tissue stiffness.

Therefore, the basis for our approach is the availability of more accurate indicators of the precision of each strain estimate, which influences our use of these estimates at every stage downstream including the display. An appropriate strain normalisation may be calculated by fitting a suitably constrained surface to the entire set of displacement data in each frame, possibly by the method of precision-weighted least squares. Normalisation can be applied both to the strain data and also to the associated precision data, producing a new array of pseudo-strain data with updated precision values. Having produced a single frame of pseudo-strain, the signal-to-noise ratio may be boosted by applying some form of persistence or spatial averaging, which may again be weighted according to precision, from which the output is a set of persisted pseudo-strain values and appropriately updated precisions. Finally, the display scheme may be tailored to indicate both strain and precision data on a two-dimensional (2D) scale represented by a 2D colour map.

## 2.2 Predicting estimation precision

Our system exploits the availability of useful predictions of strain (and/or displacement) estimation precision. The method for achieving these predictions is not critical, although the advantages of our concept are likely to be greatest if the precision predictions are highly accurate. The prediction method employed to produce results in this report is based on the work reported in [15].

To summarise, precision is the reciprocal of variance or mean squared error. Displacement precision can be predicted by evaluating  $Tc/(1-c)$ , where  $c$  is the correlation coefficient between pre- and post-deformation data in the displacement estimation window, and  $T$  is the window length. In each least squares kernel, the overall strain estimation variance can be predicted by evaluating an average of the displacement variances weighted by the square of the distance from the kernel centre, dividing by the sum of squared distances. A more accurate estimate can be produced by applying a more complicated formula that accounts for the correlations between nearby errors [15]. This was not applied in the results that we present here — although it might bring a slight improvement in some cases — because it is less important when we only consider fixed analytical parameters, and we only require predictions of relative (rather than absolute) precision. The results that we present use the following approximation for strain estimation precision,  $W_A$ :

$$W_A(x, y) = \frac{(\sum_i \check{y}_i^2)^2}{\sum_j \check{y}_j^2(1-c_j)/(T_j c_j)} \quad (1)$$

where the sums are over displacement estimation windows in the least squares regression kernel centred on pixel  $(x, y)$ , and  $\check{y}$  denotes distance from the centre of the kernel along the axial direction in which strain is being estimated.

## 2.3 Normalisation

Careful design of the normalisation strategy may contribute to valuable improvements in the quality of the strain images, particularly if real time images are required or the scanning procedure is freehand. Various approaches have been reported in the past [6, 12, 13]. The basic problem of finding an appropriate strain scale for each image can be solved robustly by fitting a plane to the entire set of displacement estimates,  $\{d(x, y)\}$ . This is performed in our examples by the method of precision-weighted least squares, thereby determining an “average” strain. The equation of the fitted plane is then as follows:

$$\hat{d}(x, y) = \alpha + \hat{s}y \quad (2)$$

The strain estimates can be scaled so that the dynamic range in the display is from zero up to a fixed multiple of the average strain,  $\hat{s}$ .

In the new interface, we also introduce extensions of this approach, by fitting other parametric surfaces to the set of displacement estimates. For instance, we can adjust for the reduced stress at greater depths away from the probe, as the stress spreads out into the surrounding tissue:

$$\hat{d}(x, y) = \alpha_0 + \alpha_1 x + \beta_1 y + \beta_2 y^2 \quad (3a)$$

$$\hat{s}(x, y) = \beta_1 + 2\beta_2 y \quad (3b)$$

It is therefore more appropriate to refer to  $\hat{s}(x, y)$  as a “normalisation” strain, rather than an average, since it is a function of image position. Having found the parameters  $\beta_1$  and  $\beta_2$  — which we again evaluate in our examples by precision-weighted least squares regression — we can divide through by the local value of  $\hat{s}(x, y)$  in order to normalise each strain estimate.

A further extension can be made to adjust for the possibility that the probe may rotate about the elevational axis during the scan, resulting in stress variation over the lateral direction.

$$\hat{d}(x, y) = \alpha_0 + \alpha_1 x + \beta_1 y(1 + \beta_2 y)(1 + \beta_3 x) \quad (4a)$$

$$\hat{s}(x, y) = \beta_1(1 + 2\beta_2 y)(1 + \beta_3 x) \quad (4b)$$

Again, the parameters  $\beta_1$ ,  $\beta_2$  and  $\beta_3$  can be found by precision-weighted least squares regression, or any suitable alternative, thereby defining the normalisation strain at every position throughout the image.

Scans using 3D probes (2D arrays or mechanically-swept 1D arrays) lead to volumes of displacement estimates,  $\{d(x, y, z)\}$ . In these cases we can extend our normalisation further, adjusting for linear variation in pressure over the elevational direction:

$$\hat{d}(x, y, z) = \alpha_0 + \alpha_1 x + \alpha_2 z + \beta_1 y(1 + \beta_2 y)(1 + \beta_3 x)(1 + \beta_4 z) \quad (5a)$$

$$\hat{s}(x, y, z) = \beta_1(1 + 2\beta_2 y)(1 + \beta_3 x)(1 + \beta_4 z) \quad (5b)$$

Strain estimates in volumetric frames are normalised simply by dividing through by the local value of  $\hat{s}(x, y, z)$ .

Many further extensions are possible. However, as the number of parameters defining the displacement surface increases, it may be more convenient to determine the normalisation strain by fitting a surface to the pre-normalisation strain estimates, rather than to the displacements. This reduces the computational complexity, although it is sometimes less accurate. Whether it is better to fit a displacement surface or a strain surface depends on the number of available estimates, their accuracy, and the complexity of the parametric hyper-surface that is being used. Normalisation in our later examples is by fitting displacement surfaces, except where otherwise specified.

It is worth noting that all of our normalisations can be applied both to the strain estimates and also to the associated precision values. Since normalisation applies a scaling of  $1/\hat{s}(x, y)$  both to good measurements and to errors, we correspondingly scale the precision (reciprocal of mean squared error) by  $\hat{s}(x, y)^2$ . If we denote pre-normalisation strain estimates and post-normalisation pseudo-strain by  $s_A$  and  $s_B$ , with  $W_A$  and  $W_B$  respectively denoting the pre- and post-normalisation precisions, then

$$s_B(x, y) = s_A(x, y)/\hat{s}(x, y) \quad (6a)$$

$$W_B(x, y) = W_A(x, y) \times \hat{s}(x, y)^2 \quad (6b)$$

The practical effect of the combined normalisation is to place each individual strain estimate on a broad scale of possible interpretations, depending on its relative properties in the context of the entire frame of scan data. Depending on the value of  $\hat{s}(x, y)$ , the normalisation of any single strain estimate and its precision value is to place it within a range spanning (1) relatively low pseudo-strain at relatively high precision, through to (2) relatively high pseudo-strain at relatively low precision. The form of the normalisation therefore potentially influences not only the type of image, but also its accuracy.

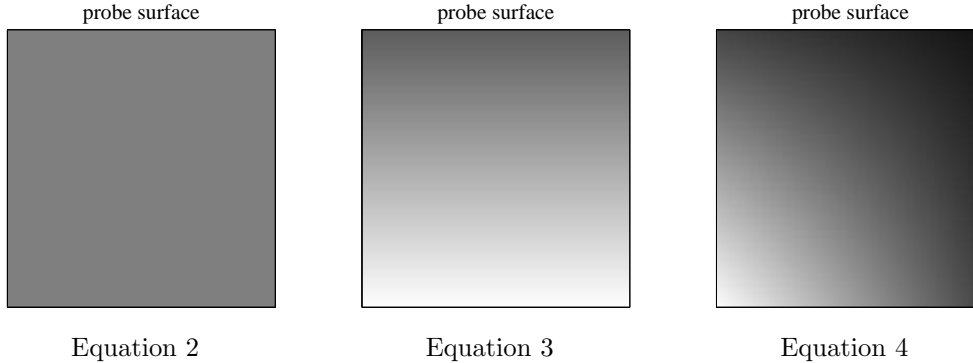


Figure 3: Illustration of the types of strain fields that may produce uniform pseudo-strain fields — indicating homogeneous stiffness — with each of our normalisation options. Here low strain is shown by white and high strain is black. The first option only adjusts for stress variation on the level of the whole image, while the second adjusts for lower stress away from the probe surface, and the third also adjusts for uneven probe pressure.

It bears noting that the main computational expense of normalisation comes from fitting a parametric displacement or strain surface, but this in itself is typically a negligible cost on widely available GHz processors in the context of 2D frame rates below 100 Hz. Computational efficiency is one of the main factors behind the selection of the particular parametric forms that we provide here as examples. However, Equations 3, 4 and 5 imply linear variation with depth in the normalisation strain. This may usually be a good approximation, but it leaves open the possibility that the sign of the normalisation strain could invert at some depth within the image. If this were a reasonable form of normalisation, it would imply that at some depth the direction of the stress field inverts, *i.e.*, that a compression at the surface causes extension at greater depth within the tissue. This is unrealistic, but it can be prevented, for example by constraining the fitted surface to avoid the strain crossing zero within the image depth (as in the images presented in this report) or alternatively data below the zero crossing of the normalisation strain can be treated as uninformative, by setting the precision to zero.

The normalisation surface might ideally reflect exponential variation with depth, but a least squares fit would then incur much greater computational cost. Our demonstration employs the normalisation surfaces outlined above for efficiency, not precluding the possibility that other parametric or constrained non-parametric forms may be found in the future, offering better performance at reasonable cost.

Figure 3 illustrates the strain fields that are implied by each of the normalisation scheme examples, or equivalently the stress fields that might produce such a field in homogeneous material. The key with the normalisation is to fit a suitably constrained surface, that with high probability corrects for artefacts associated with the uneven distribution of stress within the tissue, without removing information that has arisen owing to genuine differences in stiffness.

It is possible — but unlikely — that there may be tissue in which stiffness in fact varies with the reciprocal of depth, and the application of a uniform stress field may also be possible, in which case normalisation using Equation 3 or Equation 4 would remove real stiffness data from the display. The frequency with which this sort of ambiguity arises will depend on the scanning target, so it may be that different normalisation surfaces are required for different clinical applications.

## 2.4 Persistence or spatial averaging

Persistence refers to time-averaging of image data, while in some applications spatial averaging is more appropriate. The concept of weighted averaging of multiple frames is not new in ultrasonic strain imaging [8]. In general, our preferred approach is to perform averaging after normalisation

in the form of a precision-weighted average on a per-pixel basis. We also sum the precision values, since it can be shown that the overall precision of a correctly precision-weighted average of data with uncorrelated errors is equal to the sum of precisions.

In the context of producing a real time interface for freehand imaging with two spatial image dimensions, we perform this averaging as a form of persistence on the arrival of each new frame,  $f$ . The values that persist in pre-display buffers at pixel  $(x, y)$  are a precision-weighted sum,  $S(x, y, f)$ , and the sum of precisions,  $\Omega(x, y, f)$ . These buffers are updated as each new frame arrives, providing new pseudo-strain data,  $s_B(x, y, f)$ , and new precision data,  $W_B(x, y, f)$ .

$$S(x, y, f) = \gamma S(x, y, f - 1) + W_B(x, y, f) s_B(x, y, f) \quad (7a)$$

$$\Omega(x, y, f) = \gamma \Omega(x, y, f - 1) + W_B(x, y, f) \quad (7b)$$

Here  $\gamma$  is a number between 0 and 1 that determines the level of persistence. Each persisted pseudo-strain is given by  $S(x, y, f)/\Omega(x, y, f)$ , accompanied by a precision (quality) value for the display,  $\Omega(x, y, f)$ . Therefore, we can present sonographers with a meaningful display, representing data quality as well as strain — the appearance of each image is determined jointly by these two quantities.

Our method for spatial averaging is very similar. We apply it to reduce the noise in 3D data spanning a volume. In general, our spatial averaging involves convolving the data with a filter kernel, spanning time and the three spatial dimensions. Persistence as described above is usually suitable for implementing the filter over time — it amounts to a smoothing kernel, with an infinite impulse response. We treat the spatial dimensions differently, on the other hand, by explicitly expressing the kernel in the form  $K(|\Delta x|, |\Delta y|, |\Delta z|)$ . Spatially averaged data  $S(x, y, z)$  and  $\Omega(x, y, z)$  are similar to the persisted data: spatially averaged pseudo-strain values are again given by  $S(x, y, z)/\Omega(x, y, z)$ . These data are calculated as follows.

$$S(x, y, z) = \sum_i K(|x_i - x|, |y_i - y|, |z_i - z|) W_B(i) s_B(i) \quad (8a)$$

$$\Omega(x, y, z) = \sum_i K(|x_i - x|, |y_i - y|, |z_i - z|) W_B(i) \quad (8b)$$

This is symmetric smoothing, for which kernel values in the range 0 to 1 express the weighting of normalised pseudo-strain data at  $(x_i, y_i, z_i)$  when spatially averaged data are being calculated at  $(x, y, z)$ . Generally, sets of normalised pseudo-strain  $s_B(i)$  and precision  $W_B(i)$  data may be either regularly or irregularly distributed over 3D space. Examples of both situations are presented in Section 3. If the location data  $(x_i, y_i, z_i)$  are continuous, then  $K(|x_i - x|, |y_i - y|, |z_i - z|)$  must be defined in a functional form, rather than as a discrete kernel. Ideally, spatial averaging might be implemented with a smooth kernel, such as a Gaussian, but this is inconvenient to calculate, especially when the data are irregularly distributed. Instead, in our examples we revert to a rectangular moving average filter,

$$K(|\Delta x|, |\Delta y|, |\Delta z|) = \begin{cases} 1 & \text{if } |\Delta x| < L \cap |\Delta y| < L \cap |\Delta z| < L \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

where  $L$  sets the kernel size. This filter can be applied very efficiently to irregularly as well as to regularly distributed data. The central lobe of its spatial frequency response is essentially a low pass filter. Some high frequency noise remains owing to side lobes, but this can mostly be removed simply by explicitly low-pass filtering the pixel data before the final display.

## 2.5 Display

Returning to Figure 2, an advantage of traditional ultrasound imaging is that signal intensity displays automatically tend to show the most data where the signal is strong, and they show less data where the signal is weak (the image turns black). Similarly, one of our options when imaging



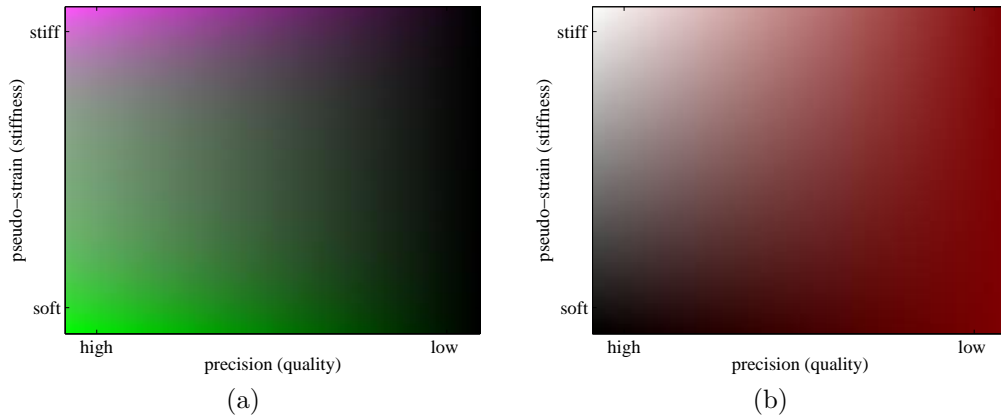


Figure 4: Two examples of 2D colour maps. (a) Green through to magenta provides the strain scale, while pixel intensity indicates the data quality within a range from a lower threshold — below which everything appears black — up to a maximum threshold — above which colours are displayed with the maximum intensity. (b) The strain scale is based on intensity variation between black and white, which blends with dull red when the precision is low.

pseudo-strain is to control image luminance or intensity based on the precision data, and to use changes in colour (preferably independently of luminance/intensity) to indicate strain. Regarding the colour scheme, our options include the use of a wide range of saturated colours, producing the effect of a contour display (as for example in the blue-cyan-green-yellow-red scale of [7] and [18]), but for the present demonstration we favour a dichromatic scale, which is qualitatively closer to traditional intensity-based displays, and which may avoid distorting the features that are perceived to appear in each image. Aiming for maximum colour variation across our example scale, we use green and magenta at the extremes, varying from strong green (high strain/soft) through grey (medium strain/medium stiffness) to magenta (low strain/stiff). Since we perceive different colours with different sensitivity, colour variation at a fixed intensity is achieved following the convention of holding constant the value of  $59 \times (\text{green pixel value}) + 30 \times (\text{red pixel value}) + 11 \times (\text{blue pixel value})$  [9]. The overall colour map, considering both strain and precision, is illustrated in Figure 4.<sup>†</sup>

There are likely to be both advantages and disadvantages associated with representing strain with colour instead of intensity, since image features encoded in these two alternative ways are processed with different accuracy and at different speed by the human visual system [2]. We therefore also test a 2D colour map in which strain is indicated by intensity, and a colour (red in our example) is introduced to indicate precision. We include this in Figure 4 and in our results to provide a comparison. In any event, our aim in relation to displays is simply to demonstrate that a 2D colour map can be used effectively to depict strain and precision data simultaneously.

Note that in the results section we demonstrate 2D colour maps that are encoded with eight bits per pixel. This is usually sufficient to produce good images, because distinctions within the dark or red regions of the colour maps are less perceptible, so these regions can be encoded at low precision. However, 16-bit encoding might be preferable if this type of interface came into widespread clinical use, since the appearance of the display images would then be marginally smoother.

<sup>†</sup>High quality colour printing is required in order to appreciate the full meaning of the remaining figures.

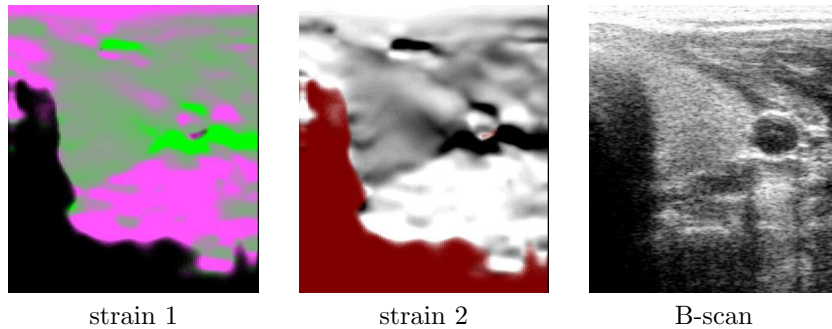


Figure 5: Thyroid images using the full interface.

### 3 Results

By applying our full interface scheme — with intelligent normalisation, persistence and display — to the scan shown earlier in Figure 2, we produce a substantially better image in Figure 5. Precision-weighted persistence gives rise to good strain estimates throughout most of the image, and the poor precision of estimates in the shadowed regions is clearly indicated by both of the 2D colour maps. This scan was undertaken freehand using the 5–8 MHz 7L3-V probe of the Terason<sup>†</sup> T3000 laptop-based ultrasound machine, running Stradwin<sup>‡</sup> freehand 3D ultrasound software. The remainder of this section provides further illustrations from scans of phantoms, highlighting the properties of each interface component, and the interactions between different components.

Figure 6 illustrates the effect of persistence, using images from an inhomogeneous gelatin phantom with stiff inclusions at a depth of 7 cm, scanned using the convex 2–5 MHz 4C2-A Terason probe. These images are normalised using Equation 4. Individual strain images usually produce some regions of good strain estimates, alongside other regions with lower precision. Unweighted frame averaging as in Figure 6b might eventually converge on a good image, but for short integration times it is usually less accurate than some of the best individual images. The advantage of persistence, as in Figure 6c–d, is that it makes efficient use of the data, so better strain images are produced easily, with larger regions of good data and generally less noise. In some scans, as in this example, the use of an image-wide weighting is sufficient to cut out most of the noise, although pixel-level weightings often give better results. The other advantage of pixel-level weightings is that a precision value is retained by each pixel in the persisted image, so it is still possible to indicate the data quality using a 2D colour map. However, we include an image produced using a less robust displacement tracking algorithm in Figure 6e, in which the best form of persistence has been employed. The image is poor, demonstrating that persistence is far more effective if the rate of severe outliers can be kept to a minimum, because estimates of the precision of gross outlier errors tend to be too high. Persistence is highly effective in conjunction with our phase-based algorithms, because of the implicit continuity constraint that was introduced by the advanced iteration seeding strategy described in [14] and [25]. This issue also affects exhaustive searching based on correlation coefficient or SAD, where in both cases the imposition of continuity constraints — whether explicitly [20] or implicitly [14] — substantially reduces the rate of outlier errors.

The next example is from freehand scanning of a breast biopsy phantom (Computerised Imaging Reference Systems, Inc.<sup>§</sup> Model 052) using the linear array 5–8 MHz Terason probe. The data quality in this case is less dependent on maintaining even probe pressure, because displacement tracking near to the surface is subject to less motion decorrelation, even if the probe does rotate substantially. However, this means that a wider range of motion types register high precision values, which actually makes correct normalisation more important than in the example of Figure 6. A

<sup>†</sup><http://www.terason.com>

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

<sup>§</sup><http://www.cirsinc.com>

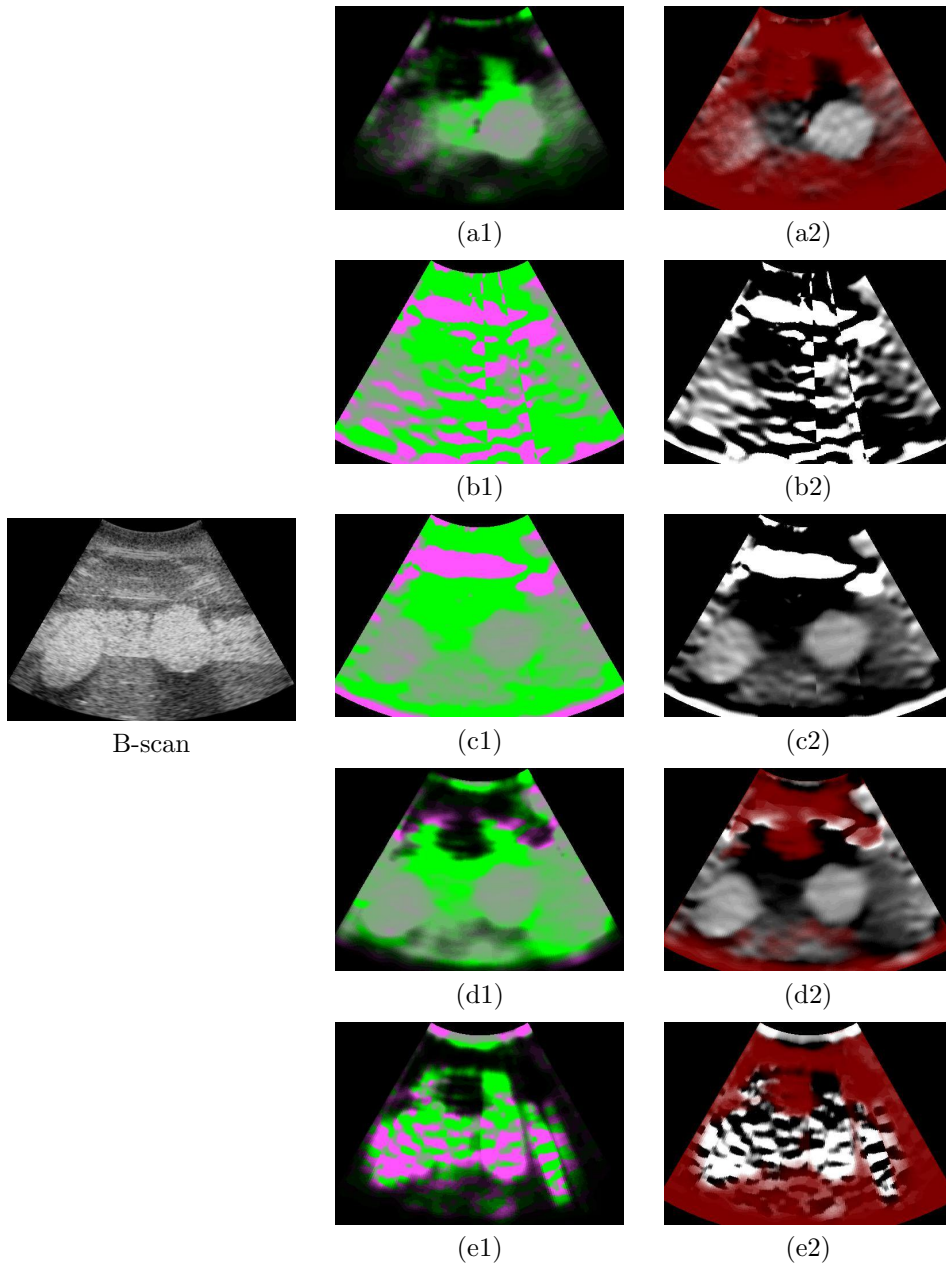


Figure 6: Example of alternative persistence methods in the strain images. (a) *Best individual image without persistence*: Individual frames produce a mixture of good and bad image regions, which register different levels of precision. (b) *Unweighted frame averaging*: The unweighted average of an image sequence is noisier than many of the individual frames. (c) *Precision-weighted frame averaging*: A sequential average weighted by each frame's mean precision significantly reduces the level of noise. (d) *Pixel-level precision-weighted persistence*: Performing the average with a different weight for every pixel further reduces the level of noise, but only slightly in this example. Its main advantage in this case is the retention of pixel-level persistence data, hence the remaining poor data can be hidden. (e) *Less robust displacement tracking*: This image has the same persistence, but displacement tracking is by the fragile method of [22]. Persistence is more effective in conjunction with robust displacement tracking algorithms.




Figure 7: Comparison of normalisation applied to a single frame (no persistence). The normalisation surface is based on (a) Equation 2, (b) Equation 3 and (c) Equation 4.

single frame of strain data with relatively even compression is illustrated in Figure 7, exhibiting two noteworthy features. The uniform normalisation in Figure 7a gives the impression of there being stiffer material towards the bottom of the image, where the stress disperses into the surrounding material. The images in Figure 7b–c are better because the region with lower stress registers instead as having similar pseudo-strain at a lower signal-to-noise ratio, resulting in larger hidden regions. It is also clear from the image that the motion of the probe was slightly rotational, so that greater pressure was applied on the right hand side. This gives an appearance of soft material on the right of the image in Figure 7a–b, including a particularly soft region with low precision data. The background material correctly appears more uniform when we apply the more sophisticated normalisation in Figure 7c, particularly in the top right of the image, where the data now register an acceptable level of precision.

Rotational movement of the probe often results in stark differences depending on the form of the normalisation. The image in Figure 8 is a relatively extreme example. This demonstrates the importance of appropriate normalisation for making best use of the recorded data. The inhomogeneity of pseudo-strain precision in these images means that they also highlight the value of correctly applying weighted persistence at the pixel-level. Figure 9 shows that in this instance precision-weighted frame averaging is no better than unweighted frame averaging, whereas an excellent pseudo-strain image is produced by applying precision-weighted persistence at the level of individual pixels. Our sophisticated normalisation with lateral stress correction is advantageous because it both reduces the level of noise and produces a pseudo-strain image that corresponds much more closely to the stiffness of the phantom material.

Final 2D results are presented in Figure 10, showing a typical image sequence indicative of the sonographer’s experience when beginning a freehand scan using the new interface. The scan target in this example is an inhomogeneous agar phantom containing half of an olive, which is slightly stiffer than the agar. The screen is initially black (or red) before acceptable data become available. It begins to colour almost immediately on contact with the scan target, although some parts of




Figure 8: This image records a frame in which the main motion was rotational, so one side extended while the other compressed. The label 0 denotes the use of a black-white colour map without precision data. (a) If a uniform normalisation is used the resulting pseudo-strain image has one half coloured white and the other black about a pivot. Fortunately the precision data correctly register an absence of useful data, so (a1) and (a2) are blank. However, the more sophisticated normalisation applied in (b) registers many useful measurements, with acceptable precision at the edges of the image, away from a central pivot, the position of which is clearly visible.

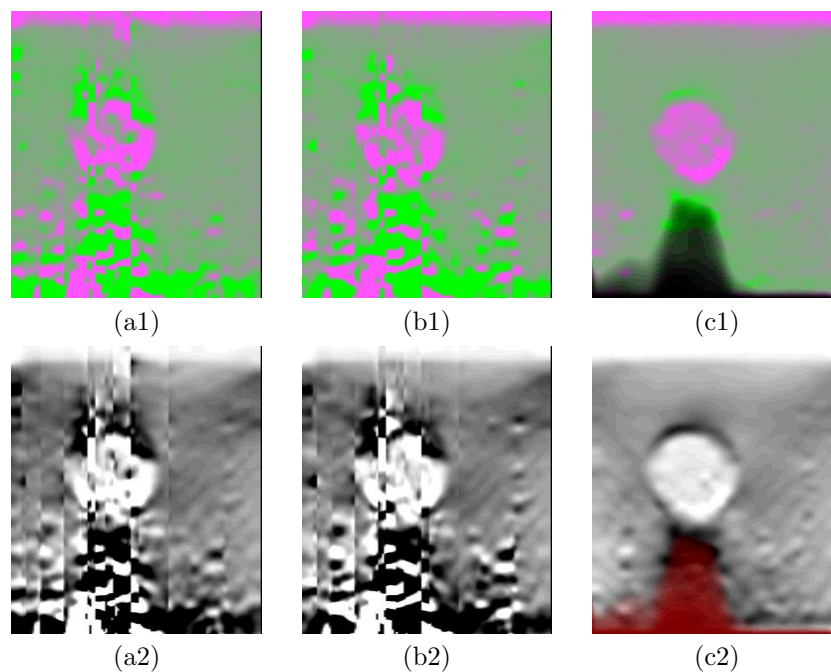


Figure 9: Examples of persistence alternatives applied to a sequence of strain images from the scan of the breast biopsy phantom, where the scan has been conducted inexpertly, frequently rolling the probe about the elevational axis. (a) Unweighted frame averaging still produces poor results. (b) In this instance, precision-weighted frame averaging is no better than unweighted averaging, because the precision of each individual estimate correlates poorly with the mean precision in each frame. (c) Precision-weighted persistence at the pixel level produces a far better image.




Figure 10: Sequence of images at the start of a freehand scan.

the image colour less quickly than others, while regions without data — such as the shadow on the right — may remain black/red. Stable images are achieved easily, and the development of a successful scanning technique is supported by visual feedback: good technique illuminates the display, whereas poor movements cause it to darken.

We have also undertaken two forms of freehand 3D strain imaging. Firstly, semi-freehand 3D scanning employs a mechanically-swept array that sweeps elevationally within its casing, while contact pressure is controlled manually. Secondly, we use a standard linear array tracked by a 6DOF position sensor for more traditional freehand 3D scanning [5, 13]. Each 3D example that we present naturally requires many images to illustrate different planes slicing through each volume of data, so to save space in these figures we only show the second of our 2D colour maps.

Our mechanically-swept system consists of a GE<sup>†</sup> RSP 6-12 probe interfaced to a Dynamic Imaging Diasus<sup>‡</sup> ultrasound machine, from which RF ultrasound data are sampled by a Gage<sup>§</sup> CompuScope 14200 analogue-to-digital converter, and the data are processed and displayed in real time on a PC running Stradwin software. The scanning protocol consists of pressing the probe against the scan target, and holding fairly still during the first mechanical sweep, then manually applying slightly more pressure, before again holding still for the second sweep. Full 3D displacement estimation is applied to the resulting volumes of pre- and post-deformation RF ultrasound data, and axial strain is estimated using 3D least squares kernels, as described in [26]. The volumes of strain data are normalised to pseudo-strain data by following Equation 5 — in this instance a strain hyper-surface is fitted directly to find the parameters in Equation 5b, which is a

<sup>†</sup><http://www.gehealthcare.com>

<sup>‡</sup><http://www.dynamicimaging.co.uk>

<sup>§</sup><http://www.gage-applied.com>




Figure 11: Using the mechanically-swept array, freehand 3D scan of half of an olive in agar. The images show slices through 3D data volumes. Top row from left to right: axial-lateral planar intensity image (*i.e.*, a conventional B-scan), axial-elevational intensity, lateral-elevational intensity, and these planes located in 3D space. Middle row: corresponding pseudo-strain images without spatial averaging. Bottom row: pseudo-strain images with spatial averaging. Fine red lines through the axial-elevational and lateral-elevational slices indicate their intersections with the displayed axial-lateral slice.

reliable method given the large quantity of data available in the volume. Example images from a scan of the olive-agar phantom are shown in Figure 11. The images show that spatial averaging is not necessarily required in this case, but the 3D results are nonetheless improved substantially in the bottom row of Figure 11, in which a weighted moving average has been applied following Equations 8 and 9 with  $L=0.8\text{mm}$ .

The spatial averaging incurs minimal loss of resolution, because much of the increase in signal-to-noise ratio arises from the weighted averaging of pseudo-strain data that are spaced more closely than the true resolution of the strain imaging system. This means that weighted spatial averaging is in fact more important in traditional freehand 3D strain imaging, where movement of the ultrasound array is determined entirely by manual scanning technique. In this case, strain estimates from successive 2D frames often overlap almost precisely in 3D space. This arises because of the relatively haphazard motion of the probe. The errors in overlapping strain estimates are in fact only weakly correlated (if at all), because each overlapping strain estimate arises from a different deformation of the tissue, making weighted averaging all the more effective. The form of normalisation is also very important in this application, because every frame has a different pressure distribution, so the effectiveness of spatial averaging depends critically on the use of data






Figure 12: Freehand scan of an olive-gelatin phantom. Top row: intensity. Middle row: strain. Bottom row: pseudo-strain with spatial averaging ( $L=1\text{mm}$ ). In every row, fine red lines through the middle two slices (axial-elevational and lateral-elevational) indicate their intersections with the axial-lateral slice on the left. These images may be compared with previous results based on the same data in Fig. 11 of [13].

that have been normalised suitably to adjust for pressure variations.

The examples in Figures 12 and 13 are based on data acquired again by sampling RF from the Dynamic Imaging Diasus machine, this time using the standard 5–10 MHz probe, tracked during freehand scanning by a Northern Digital<sup>†</sup> Polaris optical position sensor. Figure 12 shows an olive-gelatin phantom, which was scanned as part of our original investigation of freehand 3D strain imaging [13]. The middle row of Figure 12 shows strain images with a uniform normalisation of each frame, without spatial averaging, so the axial-elevational and lateral-elevational planes have been constructed by nearest-neighbour interpolation. This fails to produce a useful 3D image. Previously we discussed the use of frame-level quality measures for automatic rejection of frames with poor strain images [12, 13], but we now achieve far higher quality overall by applying weighted pixel-level filtering, *i.e.*, producing pseudo-strain images using our new interface with spatial averaging and normalisation following Equation 4. One part of the 3D image in Figure 12 still produces poor data, because the phantom contained a pocket of trapped air, causing the shadowing shown in the axial-elevational slice. Our 2D colour map is useful in this region, because it correctly marks the shadowed region red, where no accurate strain data are available.

We also show a similar freehand 3D scan in Figure 13 of the same olive-agar phantom that was

<sup>†</sup><http://www.ndigital.com>




Figure 13: Freehand scan of the olive-agar phantom. Top row: intensity. Middle row: strain. Bottom row: pseudo-strain with spatial averaging ( $L=1\text{mm}$ ). In every row, fine red lines through the middle two slices (axial-elevational and lateral-elevational) indicate their intersections with the axial-lateral slice on the left.

scanned with the semi-freehand method to produce Figure 11. Again, spatially averaged pseudo-strain in Figure 13 offers a great improvement over naive strain imaging. The reader may also note the benefits of experience: the olive-agar phantom contains half of an olive, upturned to avoid the formation of the air pockets that were problematic in our previous olive-gelatin phantoms [13].

## 4 Conclusions

We have presented an overview of a novel interface for real time freehand strain imaging, with brief explanations of the underlying theoretical principles. The preferred inputs for the interface are strain estimates from a robust strain estimator, together with accurate precision estimates. This means that the interface can be incorporated as the front end on a wide range of strain imaging systems, although the best results are likely to be achieved in systems that include robust displacement estimation that neither relies on exhaustive searching nor on tracking methods that exhibit excessive fragility.

Notable aspects of our interface include a normalisation stage, persistence or spatial averaging and a novel display using a 2D colour map. Normalisation reduces the ambiguity of strain imaging, and actually reduces the level of noise in persisted images. It follows that good, informative pseudo-strain images can be produced by a wide range of probe motions, rather than relying heavily on

careful, even compressions. In order to exploit these benefits fully, persistence should be weighted at the level of individual pixels, rather than at the level of sequential images.

We have also demonstrated the application of our interface to freehand 3D scanning. With the new interface, the general robustness of pseudo-strain imaging means that the acquisition of 3D volumes of pseudo-strain data is essentially no more challenging than the basic problem of reliably producing good 2D images. This is also made considerably easier by the presentation of persisted real time pseudo-strain images in 2D, while the sonographer acquires data across the volume.

In general, this system not only improves the quality of the results from particular data sets, but it also supports the acquisition of suitable data, by helping the sonographer to develop a successful scanning technique. Wide-ranging clinical trials will begin mid-2007 in Addenbrooke's Hospital (Cambridge, UK) to investigate suitable applications for this system. Although the interface has already been received enthusiastically by our clinical collaborators, it is possible that further modifications could be made to tailor it further for specific applications as clinical experience is accumulated. At the very least, it is probable that the best normalisation and 2D colour scheme will be application-dependent, so a range of application-specific settings is likely to be required within the framework of the existing interface.

## References

- [1] P. E. Barbone and J. C. Bamber. Quantitative elasticity imaging: what can and cannot be inferred from strain images. *Physics in Medicine and Biology*, 47:2147–2164, 2002.
- [2] B. Dresch and S. Grossberg. Spatial facilitation by color and luminance edges: boundary, surface and attentional factors. *Vision Research*, 39:3431–3443, 1999.
- [3] S. Y. Emelianov, X. Chen, B. Knipp, and D. Myers. Triplex ultrasound: elasticity imaging to age deep venous thrombosis. *Ultrasound in Medicine and Biology*, 28(6):757–767, 2002.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] A. Itoh, E. Ueno, E. Tohno, H. Kamma, H. Takahashi, T. Shiina, M. Yamakawa, and T. Matsuura. Breast disease: clinical application of US elastography for diagnosis. *Radiology*, 239(2):341–350, 2006.
- [8] J. Jiang, T. J. Hall, and A. M. Sommer. A novel image formation method for ultrasonic strain imaging. *Ultrasound in Medicine and Biology*, 33(4):643–652, 2007.
- [9] D. B. Judd and G. Wyszecki. *Color in Business, Science, and Industry*. Wiley, New York, 1975.
- [10] F. Kallel and M. Bertrand. Tissue elasticity reconstruction using linear perturbation method. *IEEE Transactions on Medical Imaging*, 15(3):299–313, June 1996.
- [11] F. Kallel and J. Ophir. A least-squares strain estimator for elastography. *Ultrasonic Imaging*, 19:195–208, 1997.
- [12] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Frame filtering for improved freehand 3D US elastography. In *Ultrasonic Measurement and Imaging of Tissue Elasticity*, page 74, Austin, Texas, October 2005.

- [13] 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.
- [14] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Phase-based ultrasonic deformation estimation. Technical Report CUED/F-INFENG/TR 555, Cambridge University Department of Engineering, May 2006.
- [15] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Dynamic resolution selection in ultrasonic strain imaging. Technical Report CUED/F-INFENG/TR 566, Cambridge University Department of Engineering, September 2006.
- [16] 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*, 2007. *In press*.
- [17] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Dynamic resolution selection in ultrasonic strain imaging. To appear in *Acoustical Imaging*, Volume 29, I. Akiyama (editor), Springer 2008.
- [18] N. Miyanaga, H. Akaza, M. Yamakawa, T. Oikawa, N. Sekido, S. Hinotsu, K. Kawai, T. Shimazui, and T. Shiina. Tissue elasticity imaging for diagnosis of prostate cancer: a preliminary report. *International Journal of Urology*, 13:1514–1518, 2006.
- [19] J. Ophir, I. Céspedes, H. Ponnekanti, Y. Yazdi, and X. Li. Elastography: a quantitative method for imaging the elasticity of biological tissues. *Ultrasonic Imaging*, 13:111–134, 1991.
- [20] C. Pellot-Barakat, F. Frouin, M. Insana, and A. Herment. Ultrasound elastography based on multiscale estimations of regularized displacement fields. *IEEE Transactions on Medical Imaging*, 23(2):153–163, February 2004.
- [21] 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.
- [22] 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.
- [23] D. M. Regner, G. K. Hesley, N. J. Hangiandreou, M. J. Morton, M. R. Nordland, D. D. Meixner, T. J. Hall, M. A. Farrell, J. N. Mandrekar, W. S. Harmsen, and J. W. Charboneau. Breast lesions: evaluation with US strain imaging — clinical experience of multiple observers. *Radiology*, 238(2):425–437, 2006.
- [24] W. E. Svensson and D. Amiras. Ultrasound elasticity imaging. *Breast Cancer Online*, 9(6):1–7, 2006.
- [25] G. M. Treece, J. E. Lindop, A. H. Gee, and R. W. Prager. Efficient elimination of dropouts in displacement tracking. In *Ultrasonic Measurement and Imaging of Tissue Elasticity*, page 68, Snowbird, Utah, October 2006.
- [26] G. M. Treece, J. E. Lindop, A. H. Gee, and R. W. Prager. Near-real-time 3D ultrasonic strain imaging. To appear in *Acoustical Imaging*, Volume 29, I. Akiyama (editor), Springer 2008.
- [27] T. Varghese and J. Ophir. Performance optimization in elastography: Multicompression with temporal stretching. *Ultrasonic Imaging*, 18(3):193–214, 1996.