
**DYNAMIC RESOLUTION SELECTION
IN ULTRASONIC STRAIN IMAGING**

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

CUED/F-INFENG/TR 566

September 2006

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

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

DYNAMIC RESOLUTION SELECTION IN ULTRASONIC 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

Ultrasonic strain imaging promises to be a valuable tool in medical diagnostics. Reliability and ease-of-use have become important considerations. These depend on selection of appropriate imaging parameters. Two tasks are undertaken here. The tradeoff between resolution and estimation precision is examined closely to establish models for the relationships with imaging parameters and data properties. These models are then applied in a system that automatically sets the imaging parameters responsive to the data quality and the required estimation precision, so as to produce more meaningful images under varying scan conditions. The new system is applied to simulation, *in vitro* and *in vivo* data for validation. It reduces the complexity of the sonographer's role in strain imaging, and produces images of reliable quality even when the level of signal decorrelation varies throughout the ultrasound data.

1 Introduction

There is considerable and growing interest in the measurement and imaging of tissue mechanical properties (primarily strain) by processing RF ultrasound signals [2]. Valuable applications seem likely in numerous areas of medical diagnostics, among them improved screening for soft tissue tumours including breast [1, 11] and prostate [9] cancer. With clinical adoption emerging as a realistic goal, practical requirements such as a high frame rate [10, 13], robustness [6, 14] and ease-of-use are becoming increasingly important considerations.

Strain imaging will be the first form of elasticity imaging in regular use. Not all ultrasound data can be used for this purpose, because data must to some degree be correlated between consecutive frames [5, 6]. In those frame pairs where strain imaging is possible, good performance depends on the selection of suitable processing parameter values. The optimal parameter set depends on numerous variables. Some of these are fixed for particular applications. Properties of the imaging hardware are usually well defined (transducer quality, focal depth, centre frequency, bandwidth, level of electrical noise). For image interpretation, clinical experience is likely to indicate the relative importance of clear shape definition versus accurate, quantitative measurement of strain contrast in different tasks. Some variables are less predictable, however, such as properties of the scanning target (echogenicity, scatterer density, mechanical feature scale) and motion of the ultrasound probe (size and direction of deformation, degree of lateral and elevational movement). In general, strain estimation precision can be improved by analysing at coarser resolution: mean strains can be estimated easily over large distances, but fine-scale displacement gradients are more problematic. An image with greater detail is more informative only if the detail is obtained at an acceptable estimation signal-to-noise ratio (SNR_e).

The construction of a strain image usually requires recording at least two RF ultrasound frames, with some intervening tissue deformation. There are two stages to the subsequent signal processing. (A) Displacement in the underlying tissue is estimated by finding the shifts between corresponding data windows in the pre- and post-deformation ultrasound frames. (B) The displacement field is spatially differentiated to estimate strain. Stage A typically involves shifting a candidate window over the post-deformation data until it maximises a similarity measure with the pre-deformation window. This is affected by parameter choices including window length, width and the spacing between successive windows. The gradient estimator in Stage B introduces further parameters,

such as the length and width of a least squares kernel. It follows that experienced researchers have numerous settings at their disposal to optimise the quality of each strain image.

The initial aim of this investigation is to examine the effects of strain imaging parameter selection on resolution and SNR_e . That information is applied in the form of a system for Dynamic Resolution Selection (DRS), whereby the parameters are steered automatically towards appropriate settings. This is motivated by two desirable outcomes. DRS will reduce the complexity presented to the sonographer, so the control options can be reduced to a single noise-rejection setting, which deals appropriately with inter-image variability. Additionally, the parameters may vary within each image, so as to produce reliable data despite intra-image variability in strain and decorrelation levels [4, 12].

This principle is pursued at a conceptual level, and is also implemented in the particular case of Weighted Phase Separation [8] with Amplitude Modulation Correction [7] and least squares strain estimation. Theory and empirical evidence from analysis of simulation and phantom data are examined to develop and validate DRS, which is also demonstrated under *in vivo* conditions.

2 Development

In this section we outline the development of DRS, employing a combination of theoretical and heuristic methods. In this investigation we are concerned only with axial strain estimation, although the concepts should be transferable to analysis of multidimensional estimation. Most approaches to strain image signal processing involve at least two stages, corresponding to displacement estimation followed by strain estimation. Both stages incur some reduction in resolving power, because deformation fields cannot be sampled with reasonable accuracy at the level of individual samples in the RF ultrasound signals. The dimensions of the analysis units for each processing stage are important parameters. We refer to the analysis units as “windows” in the case of displacement and “kernels” in the case of strain estimation. Parameter selection fixes the tradeoff between resolution and estimation precision. In order to assess the accuracy of strain imaging quantitatively, we invoke the concept of estimation signal-to-noise ratio (SNR_e). This is defined as the mean (or actual) strain, s , divided by the root mean squared estimation error, $\sigma_{\hat{s}}$.

$$\text{SNR}_e = \frac{s}{\sigma_{\hat{s}}} \quad (1)$$

SNR_e can always be improved by sacrificing resolving power in the strain images. We first consider the effect of parameter selection on resolution. We also consider how strain estimation parameter values affect the translation from displacement estimation errors into strain estimation error. Then we examine the noise source in displacement estimation, and we investigate the role of displacement estimation parameters in noise reduction. These concepts are combined practically in Section 2.5, forming the basis for the DRS procedure.

2.1 Resolution

The most obvious restriction on strain image resolution comes from the spacing between neighbouring estimation locations, but it is less important than the filtering effects of displacement estimation windows and strain estimation kernels. The only disadvantage of dense estimation spacing is increased computation time, while it improves both resolution and SNR_e . We use small spacing of estimation locations here, in order to investigate the significance of the other parameter choices.

The effects on strain estimates of changes to window and kernel size are similar (though not identical) to denoising by applying a moving average (MA) filter¹. The overall error decreases as an MA filter gets larger, because uncorrelated error components of measurements spanned by the filter average to zero. However, the output from large MA filters has coarse resolution.

¹The analogy becomes inaccurate for window size if within-window displacements are large enough to cause phase-wrapping, but this is not usually the case. This issue receives further attention in Section 2.4.




Figure 1: Effect of a MA filter on the contrast between low and high strain bands. (a) Our interest is the axial resolving limit. We determine the axial length scale at which features are just resolved, with a strain field that is uniform in the lateral direction. (b) Displacement and strain fields against distance. (c) A MA filter (length= L_r) easily resolves the different strain regions. (d) Resolution is still achieved with length= $1.5L_r$. (e) The contrast is zero when the filter length is $2L_r$. (f) Filter lengths $> 2L_r$ register negative contrast.

There is no universal definition of resolution that can be applied sensibly to all imaging tasks. We define the resolving limit as the feature scale, L_r , at which there is no longer any positive contrast between two or more low strain bands and regions of background material sandwiched between them (see Figure 1). The resolving limit is reached when strain estimates after the filter exhibit zero contrast between the low and high strain bands. For example, the resolving limit of a MA filter is half the filter length.

We will assume that the resolving limit is proportional similarly to window and kernel dimensions, although the constant of proportionality may not be the same in both cases. This is investigated empirically in Section 3.3. The greatest estimation accuracy at a given resolution will usually be achieved by windows and kernels of the maximum allowable size.

2.2 Strain estimation error

Strain estimation entails taking the derivative of displacement estimates. This amplifies any displacement estimation noise, particularly if the displacement estimates are closely spaced. Figure 2 illustrates this point for strain estimation in 1D. To be able to predict strain estimation error, we need to understand how displacement estimation error filters through.

There are various techniques available for gradient estimation to obtain strains from the displacement estimates. These include low-pass filtering and wavelet denoising. We will consider least squares strain estimation, which is more commonly applied. Analysis of the alternative approaches would involve similar considerations.

The simplest form of unweighted least squares gradient estimator is expressed in Equation 2. We follow a convention denoting lateral distance by x and axial distance by y . The strain estimate, \hat{s}_m , is produced using data from a set of displacement estimation windows, $\{\Omega_n\}$, comprising displacement estimates $\{\hat{d}_n\}$ at locations $\{\tilde{x}_n, \tilde{y}_n\}$ (measured relative to the centre of kernel \mathcal{K}_m).




Figure 2: 1D illustration of strain estimation. Errors in strain estimates depend on the resolution of the gradient estimator, but the original source of error is in the location and value of displacement estimates.

Note that although $\{\check{y}_n\}$ are axial location estimates, we mostly use 2D kernels, where displacement estimates are taken from multiple neighbouring columns. In order to ensure that the strain estimate has greatest validity at the centre of the kernel, the set of displacement estimate locations must be symmetric about the kernel centre.

$$\hat{s}_m = \frac{\sum_{\Omega_n \in \mathcal{K}_m} \check{y}_n \hat{d}_n}{\sum_{\Omega_n \in \mathcal{K}_m} \check{y}_n^2} \quad (2)$$

The scale of strain estimation errors is predicted by evaluating the variance of this estimator. In general there are errors in both displacement and location, all of which must be considered in the strain estimation variance. We will assume that errors in $\{\check{y}_n\}$ are negligible, however, because they are reduced substantially by our Amplitude Modulation Correction technique [7]. This leaves errors only in $\{\hat{d}_n\}$, resulting in the following strain estimation variance if covariances are negligible.

$$\sigma_{\hat{s}_m}^2 = \frac{\sum_{\Omega_i \in \mathcal{K}_m} \check{y}_i^2 \sigma_{\hat{d}_i}^2}{\left(\sum_{\Omega_i \in \mathcal{K}_m} \check{y}_i^2\right)^2} \quad (3)$$

It is also possible that the displacement estimator introduces significant covariances between nearby estimates, in which case a more complicated expression must be evaluated.

$$\sigma_{\hat{s}_m}^2 = \frac{\sum_{\Omega_i \in \mathcal{K}_m} \sum_{\Omega_j \in \mathcal{K}_m} \check{y}_i \check{y}_j \sigma_{\hat{d}_i \hat{d}_j}}{\left(\sum_{\Omega_i \in \mathcal{K}_m} \check{y}_i^2\right)^2} \quad (4)$$

Both approaches are tested in Section 3.4.

2.3 Displacement estimation error

2.3.1 Signal model

In order to estimate the errors and covariances for Equations 3 and 4, we need a suitable signal model with which to analyse the displacement estimator. It is helpful to separate signal components into two categories, which we call the common signal and the noise signal. This was introduced in our recent work [8]. (A) The common signal comprises those components that are transformed by an axial displacement field identical to the axial component of mechanical displacement in the underlying tissue. (B) The noise signal consists of components that are uncorrelated with the common signal; it arises because of changes to speckle interference patterns brought about by changed scatterer spacing, decorrelation due to off-axis motion and electrical noise.

Our signal model is expressed in Equation 5. Subscripts 1 and 2 denote pre- and post-deformation signals. We use a complex signal representation because we later employ a phase-based displacement estimator. a_1 and a_2 are the pre- and post-deformation analytic signals, which have envelopes s_1 and s_2 and phases ϕ_{s1} and ϕ_{s2} respectively. The axial displacement field in the tissue is $d(x, y)$. Common signal $f(x, y)e^{j\phi(x, y)}$ is warped precisely in accordance with the displacement field, so it is reproduced identically in Equations 5b and 5d, whereas the respective noise components, $n_1(x, y)e^{j\phi_{n1}(x, y)}$ and $n_2(x, y)e^{j\phi_{n2}(x, y)}$, are uncorrelated.

$$a_1(x, y) = s_1(x, y)e^{j\phi_{s1}(x, y)} \quad (5a)$$

$$\begin{aligned} &= \text{common signal} + \text{noise signal} \\ &= f(x, y)e^{j\phi(x, y)} + n_1(x, y)e^{j\phi_{n1}(x, y)} \end{aligned} \quad (5b)$$

$$a_2(x, y + d(x, y)) = s_2(x, y + d(x, y))e^{j\phi_{s2}(x, y + d(x, y))} \quad (5c)$$

$$= f(x, y)e^{j\phi(x, y)} + n_2(x, y)e^{j\phi_{n2}(x, y)} \quad (5d)$$

2.3.2 Displacement estimator

It is common in the literature to evaluate lower bounds, often Cramer-Rao lower bounds, on the displacement estimation error that could be achieved by maximum likelihood estimators. Instead in this work we look at the mechanism of a particular displacement estimator, to form an expression that can be used for predicting the actual error depending on properties of the recorded signals. The estimator that we analyse is Weighted Phase Separation [8]. This is a form of iterative displacement estimation where a scaled weighted sum of phase differences is taken to be the next displacement estimate. The formula for the k^{th} iteration at the n^{th} window, Ω_n , is as follows, where ω_0 is the probe centre frequency, or an estimate thereof.

$$\tilde{d}_{k+1} = \tilde{d}_k + \frac{\sum_{\{x, y\} \in \Omega_n} W(x, y, \tilde{d}_k)(\phi_{s1}(x, y) - \phi_{s2}(x, y + \tilde{d}_k))}{\omega_0 \sum_{\{x, y\} \in \Omega_n} W(x, y, \tilde{d}_k)} \quad (6a)$$

$$\text{where } W(x, y, \tilde{d}_k) = W_A(x, y, \tilde{d}_k)W_B(x, y, \tilde{d}_k) \quad (6b)$$

$$W_A(x, y, \tilde{d}_k) = s_1(x, y)s_2(x, y + \tilde{d}_k) \quad (6c)$$

$$W_B(x, y, \tilde{d}_k) = \frac{\pi - |\phi_{s1}(x, y) - \phi_{s2}(x, y + \tilde{d}_k)|}{\pi} \quad (6d)$$

The example weighting strategy of Equations 6c and 6d is employed throughout the investigation at hand. Iterations converge on displacement estimate \hat{d}_n . Additionally, Amplitude Modulation Correction [7] is applied in order to minimise location errors, as already discussed. This entails recording the centroid of the weightings from Equation 6, which corresponds to the location² at which the displacement estimate is most valid in the presence of intra-window strain. Both lateral and axial co-ordinates must be estimated when the windows are 2D (*i.e.*, for windows that span multiple A-lines).

$$\hat{x}_n = \frac{\sum_{\{x, y\} \in \Omega_n} W(x, y, \hat{d}_n)x}{\sum_{\{x, y\} \in \Omega_n} W(x, y, \hat{d}_n)} \quad (7a)$$

$$\hat{y}_n = \frac{\sum_{\{x, y\} \in \Omega_n} W(x, y, \hat{d}_n)y}{\sum_{\{x, y\} \in \Omega_n} W(x, y, \hat{d}_n)} \quad (7b)$$

It was previously observed that regular window positions result in an irregular distribution of estimation locations [7], which is inconvenient for subsequent processing and can introduce textural variation that correlates with amplitude features in B-mode images [7]. However, in the present

²These location estimates are recorded in image co-ordinates. For use of the least squares strain estimation formula as in Equation 2 the kernel centroid must be subtracted to obtain $\{\hat{x}_n, \hat{y}_n\}$ in kernel co-ordinates.

implementation this is avoided by an algorithm modification that positions windows irregularly in such a way as to return the estimation locations to a regular grid. With 1D windows, for example, this entails accumulating the location error while samples are added to the window, and taking each next sample from the end that opposes the current location error. The algorithm modification for 2D windows of a fixed aspect ratio is less simple, but it can be achieved at little computational cost by extending the same principle.

2.3.3 Error prediction

To predict the size of displacement estimation errors, we begin by considering point-wise displacement estimates (PWDEs). These are displacement estimates produced at a single sample of approximately aligned data by differencing pre- and post-deformation phase and scaling by the centre frequency. It was shown theoretically in [8] that PWDE variance is proportional to the square of the envelope of the noise signal divided by the square of the envelope of the common signal, provided that this value is large, or approximately proportional provided that the value is greater than unity. For this reason we introduce a modified signal-to-noise ratio, SNR_s .

$$\text{SNR}_s(x, y) = \frac{f(x, y)^2}{\frac{1}{2}(n_1(x, y)^2 + n_2(x, y)^2)} \quad (8)$$

Furthermore, if a set of uncorrelated PWDEs is averaged, constituting a window, the displacement estimation variance scales inversely with the number of samples.

In practice, PWDEs are not all independent, so displacement estimation variance is greater by a factor proportional to the autocorrelation length of the errors. The average degree of correlatedness between PWDEs depends on window length. However, once windows are longer than the error autocorrelation distance, the addition of each new sample adds an almost equal amount to the total information. This means that displacement estimation precision usually becomes a near-linear function of window length, as for example in Figure 3a. PWDEs in neighbouring A-lines are only weakly correlated, so estimation precision versus window width is usually more linear (see Figure 3b). It is therefore reasonable to predict displacement estimation precision with the product of SNR_s , window size and a constant of proportionality, K . Window dimensions are denoted by L_x and L_y respectively in the lateral and axial directions.

$$\sigma_d^2 \simeq (KL_xL_y\text{SNR}_s)^{-1} \quad (9)$$

This requires an estimate of SNR_s . If the underlying tissue deformation, $d(x, y)$, were known exactly, then SNR_s could be estimated with high accuracy as a function of the real part of the arithmetic correlation coefficient of correctly warped pre- and post-deformation analytic signals. The real part of the arithmetic correlation coefficient of windows aligned using displacement estimates is expressed in Equation 10.

$$\rho(\Omega_n, \hat{d}_n) = \frac{\sum_{\{x, y\} \in \Omega_n} \Re\{a_1^*(x, y)a_2(x, y + \hat{d}_n)\}}{\frac{1}{2} \sum_{\{x, y\} \in \Omega_n} |a_1(x, y)|^2 + |a_2(x, y + \hat{d}_n)|^2} \quad (10)$$

For perfectly aligned data, the terms of the product of pre- and post-deformation samples are listed in Equation 11.

$$\begin{aligned} a_1^*(x, y)a_2(x, y + d(x, y)) &= f(x, y)^2 \\ &+ f(x, y) \left(n_1(x, y)e^{j(\phi(x, y) - \phi_{n1}(x, y))} + n_2(x, y)e^{j((\phi_{n2}(x, y) - \phi(x, y)))} \right) \\ &+ n_1(x, y)n_2(x, y)e^{j((\phi_{n2}(x, y) - \phi_{n1}(x, y)))} \end{aligned} \quad (11)$$

All of the product terms are mutually uncorrelated, except for the square of the common envelope. Similarly, if the denominator of Equation 10 is expanded for perfectly aligned data, the only correlated terms are the sum of the squared common envelope and half the sum of squared noise




Figure 3: Displacement estimation precision versus window dimensions for simulation data with a mean SNR of 20 dB subject to a uniform compression of 0.01%. The horizontal scale is in RF samples of data with a 6.0 MHz centre frequency sampled at 66.7 MHz, so 500 samples correspond to 5.75 mm. 12 A-lines correspond to 3.75 mm. (a) 1D windows of lengths in the range 10–500 RF samples. The dotted line is included so that the curved section can be identified easily. (b) Analysis with windows of length 120 samples and widths in the range 1–12 A-lines.

envelopes. If they are integrated over a reasonable number of samples, the uncorrelated terms tend to zero, and a perfect estimate of SNR_s is produced by applying the following transformation.

$$\hat{\text{SNR}}_s(\Omega_n, \hat{d}_n) = \frac{\rho(\Omega_n, \hat{d}_n)}{1 - \rho(\Omega_n, \hat{d}_n)} \quad (12)$$

The problem in practice is that $d(x, y)$ is unknown, although SNR_s can be estimated roughly by assuming that displacement estimate \hat{d}_n is correct at all points in window Ω_n . Equation 12 still produces a useful estimate in this case.

Two types of alignment error limit the accuracy of $\hat{\text{SNR}}_s$. (A) Displacement estimation error (error in \hat{d}_n) leads to misalignment throughout the window, which can generally result in either over- or under-estimation of SNR_s . (B) Intra-window misalignment caused by intra-window strain is more problematic. It introduces severe bias into $\hat{\text{SNR}}_s$, because it always produces under-estimates. Figure 4 shows examples at different strain levels. Displacement estimation error is the main source of error at 0.01% strain, so values of $\hat{\text{SNR}}_s$ from longer windows are slightly more accurate. The situation is reversed at higher strain, because intra-window misalignment becomes the main source of error. Figure 4b shows the effect this has on $\hat{\text{SNR}}_s$ at different window lengths. The disadvantage of long windows is that the range of $\hat{\text{SNR}}_s$ values is much narrower than the true range. Estimates from shorter windows are more consistent across different strains. The estimates from longer windows could be improved by signal warping, but shorter windows are preferable if this is not pursued.

This is still not sufficient to predict the strain estimation error if the formula in Equation 4 is required. However, if the above analysis is used to estimate the size of mean squared displacement errors, the normalised covariance (covariance divided by mean squared error) can also be estimated reasonably, so all of the necessary information is available. Normalised covariance is estimated where necessary as the area overlap fraction between the windows. This is scaled to the full covariance by multiplying with the expected mean squared error. Examples in Figure 5 motivate this strategy.




Figure 4: SNR_s estimates at different 1D window lengths for simulation data similar to that used in the previous figure, with various levels of electrical noise. The vertical axis is image-wide $\hat{\text{SNR}}_s$, which has been evaluated by averaging the arithmetic correlation coefficients across all windows and substituting the result into Equation 12. Data are plotted for two levels of uniform strain. (a) 0.01%. (b) 0.5%.

2.4 Importance of the strain level

Displacement estimation error is sometimes also caused by phase-wrapping at the edges of long windows. This is mainly an issue at high strains. As window length increases, SNR_e performance flattens off when point-wise phase-wrapping errors become more common. Beyond a certain window length phase-wrapping errors occur with every new sample, so a severe performance reduction is observed. The length at which phase-wrapping occurs is inversely proportional to the strain. This effect is illustrated by the SNR_e contour plots in Figure 6. At low strains phase-wrapping errors never occur, so SNR_e goes up with every increase in both window length and kernel length. At higher strains, as in Figure 6b, SNR_e no longer rises so quickly beyond a certain window length, and it eventually drops. This is why the second contour plot has a different shape. The relationship between window length and estimation precision is also less linear when there is phase-wrapping, so error prediction as outlined so far would be inaccurate. As a rule-of-thumb, the window length should not be selected greater than $\lceil 5.0/\text{strain magnitude} \rceil$ samples to stay in the linear region, considering a 66.7 MHz sampling rate with a 6.0 MHz centre frequency. For example, if there is a minimum window length, say 50 samples, then the maximum strain that can reasonably be measured is 10%.

The strain level is also important because it determines the scale of significance of strain estimation errors. The analysis so far has considered predicting the absolute size of strain estimation errors, but performance in terms of SNR_e also depends on the size of the strain signal (see Equation 1). In a homogeneous image the noise level is obviously measured against the mean strain, but the appropriate benchmark for an inhomogeneous image is less obvious. Large regions of mostly low strain may be swamped by noise if the significance of errors is held to depend on the image-wide mean strain. However, if local strain sets the appropriate scale, then zero strain regions (which do arise in practice) always register zero SNR_e . We can alternatively pursue a compromise, mainly using local strains, but saturating if they deviate wildly from the mean. We adopt this final approach for DRS, but all three options are compared qualitatively in Section 3.6.




Figure 5: Examples of normalised covariance of displacement estimation between nearby windows from analysis of simulation data at 20 dB with a uniform compression of 0.01%. (a) Normalised covariance against axial separation. Displacement estimates separated in the axial direction become uncorrelated less quickly if they are from long windows. (b) Reciprocal of axial decorrelation rate against window length. The dotted line aids visual assessment of the divergence from linearity. (c) Normalised covariance against lateral separation. Displacement estimates separated in the lateral direction remain correlated longer when the windows are wide. (d) Reciprocal of lateral decorrelation rate against window width. Both axially and laterally, the separation of windows at which the covariance tends to zero is approximately equal to the window dimension.




Figure 6: Effect of phase wrapping on displacement and strain estimation performance with simulation data at an SNR of 20 dB in the presence of electrical noise, with further decorrelation caused by a 0.0625mm elevational translation of the scatterer field. For the horizontal axis, kernel length in samples comes from the length in windows multiplied by the spacing between windows. (a) SNR_e contours against 1D parameter settings at 0.5% strain. (b) SNR_e contours at 2.0% strain.

2.5 Dynamic Resolution Selection procedure

DRS brings together the strands of analysis above, and applies them so as to improve the value of real strain images. The tasks in DRS strain imaging divide essentially into three stages. These are illustrated with a flow chart in Figure 7. The aim is to set the imaging parameters automatically so that a threshold performance, minSNR_e (set by the user), is always exceeded.

Survey. An initial survey produces displacement and correlation data using Weighted Phase Separation with short windows throughout the image. The search algorithm is the standard combination of cross-seeding, multi-pass analysis and continuity checking, described in [8]. Short windows are used for two reasons. Firstly, the windows must be sufficiently short to operate without severe phase-wrapping throughout the stated dynamic range, so we use a length of 50, implying that strains should not exceed 10%. Short windows are also preferred because the resulting $\hat{\text{SNR}}_s$ values show relatively little variation in bias across different strains. 1D survey windows could be used for fast processing, but wide 2D windows (say 10 A-lines) may offer greater robustness. Both widths are tested in Section 3.4. The displacements are converted to strain estimates, \hat{s} , using a very large least squares kernel of 51 windows (510 samples) by 17 A-lines. Additionally, the arithmetic correlation coefficients are smoothed with a spatial Gaussian filter, and the smoothed values are transformed into signal-to-noise estimates, $\hat{\text{SNR}}_s$. The subsequent Analysis requires $\hat{\text{SNR}}_s$ and \hat{s} values throughout the image, so Survey data are extended to the edges by nearest-neighbour extrapolation.

Analysis. At each estimation location a search is undertaken to find the window and kernel dimensions of finest resolution that will produce performance exceeding minSNR_e . Trial values for the window and kernel dimensions can be combined with $\hat{\text{SNR}}_s$ and \hat{s} values from the survey to obtain performance estimates, SNR_e . Steps required include the prediction of displacement estimation error using $\hat{\text{SNR}}_s$ and window size (Equation 9), the estimation of normalised covariance between windows in a least squares kernel by evaluating their overlap fractions, and the substitution of both these sets of data into the formula for strain estimation error. This can be evaluated with Equation 3 if covariances have little impact, or (more rigorously) with Equation 4 when covariances are significant. This is converted to a prediction for SNR_e by substituting the strain estimation error along with an appropriate indicator of



Figure 8: Artefact caused by large window and kernel changes over a short distance. (a) Region from DRS strain image at boundary between soft background and stiff inclusion phantom. Boundary appears twice because of excessive window and kernel overlap. (b) Window lengths (white=500 samples). (c) Window lengths have been clipped, thus eliminating the double-boundary artefact.

strain level into Equation 1. The strain level is usually taken to be the local estimate, \hat{s} , but extreme values saturate at half and twice the mean strain.

Each time that $\hat{\text{SNR}}_e$ is calculated, the search terminates if it exceeds $\text{min}\hat{\text{SNR}}_e$. Otherwise, the window and kernel dimensions are increased for a new set of trial parameters. The kernel length is increased in increments of two windows to maintain symmetry about a particular estimation location. The kernel width also increases according to the aspect value. Window length and width are increased as well, in an appropriate ratio which is determined in Section 3.3. However, the window length is also subject to a limit of $[5.0/\text{local strain}]$ in order to avoid phase-wrapping, while window width continues to grow after the length has saturated.

Eventually adequate window and kernel dimensions are found to exceed $\text{min}\hat{\text{SNR}}_e$ at every estimation location. One final analysis step is required before proceeding with Refinement. The parameter selections are parsed once to eliminate an artefact that arises if parameter values change wildly over a small distance. If any window or kernel extends further in any direction than its neighbour on that side, this may result in the appearance of a higher level of information than is actually present. For example, a monotonic strain gradient could appear non-monotonic. This is eliminated trivially by clipping window and kernel lengths, so they never extend beyond 100% overlap with a neighbour. Figure 8 shows an example.

Refinement. Displacement and strain estimation is repeated using the DRS parameter values. Rather than the iteration seeding algorithm of [8], in this instance a single iterative search is employed at each window. Each initial displacement for a first Weighted Phase Separation iteration is taken from the Survey window with the highest arithmetic correlation coefficient in the same row, within five columns of the active estimation location.

3 Experimental validation

Experiments were undertaken to determine values for unknown constants required in DRS, and to assess the behaviour of the resulting system. All results were produced with estimation locations on a regular grid, with lateral spacing equal to the spacing between A-lines, and axial spacing equal to 10 samples (slightly less than one cycle at the centre frequency). The windows and kernels were 2D, with a fixed aspect ratio such that the length in samples was 30 times greater than the width in A-lines.

3.1 Image formats

Various types of images are presented: displacement estimates, strain estimates, DRS window lengths and local SNR_e . The images are presented in a consistent format. Data are obtained at every estimation location before being displayed, and the spacing is sufficiently dense that linear interpolation produces an acceptably smooth display.

A main motivation for this investigation is to examine performance variation at different resolutions, so a block pattern covers data at image edges, where estimates could not be evaluated using the stated parameter settings — *i.e.*, where the data were invalid at the stated resolution. Displacement and window length images are blocked-out at estimation locations where the windows did not fit inside the recorded data frame. Blocked-out borders on the strain images are roughly twice as thick, since strain data is displayed only from kernels that were filled entirely with valid windows. SNR_e images have the same borders as the corresponding strain images. For practical scanning the borders would be unnecessary, although it must be accepted that $\text{SNR}_e/\text{resolution}$ cannot be controlled arbitrarily at edges.

The images we present are normalised consistently. Values that determine the image scales are indicated in the figure captions. All valid displacement data are first processed in a single least squares filter to estimate the mean strain. This sets the mid-grey level in the strain images, with saturation at zero (black) and twice the mean strain (white). Similarly, displacement images have a mid-grey level at the mean strain multiplied by half the image height, with saturation at zero (black) and twice this value (white). SNR_e images have mid-grey set to the image-wide SNR_e , with saturation again at zero and twice this value. Finally, window length images are scaled such that the maximum window length is displayed white, mid-grey is half that value and black would be a window of zero length.

3.2 Simulations

Simulations were performed using Field II [3] to model the Dynamic Imaging Diasus ultrasound machine³ (*cf.*, Section 3.6). Frames were simulated before and after a compression of the scattering field in the axial direction. Each simulation had 128 A-lines spanning 40 mm, scanning to a depth of 40 mm with a single focus at 20 mm. (In fact, Field II is inaccurate in the near-field, so for reasonable accuracy data were recorded at 5–45 mm depth with a 25 mm focus.) Data were converted to the 16-bit Stradwin ultrasound format⁴ with the addition of uniform white noise, reducing the mean SNR to 20 dB where not otherwise stated. The centre frequency was 6.0 MHz, with a sampling rate of 66.7 MHz. Further details are in Table 1.

3.3 Resolution

We tested the effects of window and kernel size on resolution. Strain images were produced using a range of 1D windows (30–1000 samples) and 1D kernels (10–1000 samples), recording the contrast (difference of mean strains) between the low and high strain bands at the vertical centre of Simulation A. The thresholds at which the contrast goes from positive to negative indicate filter combinations that are at the resolving limit. Thresholds from four different feature scales are combined in Figure 9f, to form a contour map of resolving limit.

The contours are mostly rectangular, indicating that the resolution is limited essentially by either window or kernel length, whichever has the greater smoothing effect. It follows that window and kernel lengths should be set in a fixed ratio to achieve optimal estimation performance for any particular resolving limit, since increasing the window size or kernel size almost invariably reduces estimation noise (*cf.*, Figure 6). The extremes on the resolution contours are plotted in Figure 9f. These data are close to relationships of direct proportion between filter length and resolving limit. Least squares estimates of the constants of proportionality indicates that the kernel length should be 1.20 times greater than the window length for any particular resolution. This ratio is used for parameter selection in the Analysis stage of DRS, as described in Section 2.5.

3.4 Calibration

An implementation of DRS requires a scale factor, \sqrt{K} , which appeared in the formula for predicting displacement error (Equation 9). This section presents calibration results. Simulation data

³<http://www.dynamicimaging.co.uk>

⁴<http://mi.eng.cam.ac.uk/~rwp/stradwin/>




Figure 9: Strain estimation resolution assessed by analysis of Simulation A. (a–e) Example strain images (image mean strain 0.85%) of a 2.5 mm feature scale viewed with various parameter settings. Contrast is indicated as a fraction of the ideal contrast, ranging from high contrast in the first image to failure to resolve in the final image. The window/kernel lengths were 75/90, 175/210, 275/330, 375/450 and 525/630 respectively, using 2D windows and kernels. 500 samples corresponds to a physical length of 5.75 mm. (f–g) Quantitative results with 1D windows and kernels. (f) Contour plot of resolving limit over parameter setting space. (g) Maximum parameter values from the contours against resolving limit feature scale.

Purpose	Reference(s)	Scattering field	Other details
resolution experiment	Simulation A	2.0×10^5 scatterers, equal scattering strengths, $50 \times 50 \times 6$ mm volume (lateral \times axial \times elevational), uniform spatial distribution	background strain 1.0%, zero strain bands either side of focus, various feature scales, (<i>cf.</i> , Figure 1)
effect of electrical and uncoloured noise	Simulation B-20dB Simulation B-15dB <i>etc.</i>	as per Simulation A	range of levels of white noise, mean SNR is appended to simulation reference
effect of lateral motion decorrelation	Simulation C-0.0625mm Simulation C-0.125mm <i>etc.</i>	3.0×10^5 scatterers uniform dist. scat. strengths, $51 \times 50 \times 8$ mm volume, uniform spatial distribution	range of lateral motions, motion size appended to simulation reference
effect of elevational motion decorrelation	Simulation D-0.0625mm Simulation D-0.125mm	as per Simulation C	elev. motion appended to simulation reference
performance with non-uniform noise	Simulation E	1.0×10^6 scatterers equal scattering strengths, $50 \times 50 \times 6$ mm volume spatial distribution is non-uniform, gets more dense left-to-right	mean SNR 20dB, but SNR varies throughout the image, uniform strain 1.0%

Table 1: Details of simulations.

were processed with a range of uniform window and kernel settings (2D kernels in the range 70–610 samples; 2D windows 1.20 times shorter). These were performed instead of the Refinement stage in DRS, preceded by the Survey (windows of length 50 samples, either 1 or 10 A-lines wide). Rather than adjusting the window and kernel settings, Analysis produced a squared error prediction for each strain estimate. Strain divided by the root mean squared error prediction is image-wide $\hat{\text{SNR}}_e$ (before scaling by \sqrt{K}), which can be compared to the measured value, SNR_e . A suitable value for \sqrt{K} was found by comparing the predictions and measurements across a range of different strains, window/kernel dimensions and noise conditions.

By way of example, Figure 10a shows SNR_e against $\hat{\text{SNR}}_e$ for Simulation B-20dB at 0.5% strain. This is far from the expected straight line, because a baseline error is present in the strain estimates. This arises because the apparent depth in ultrasound images (time of receiving echo) does not vary perfectly linearly with scatterer depth (see Figure 10b). This occurs in real scans too, which are subject to further baseline error caused by fluctuations in the speed of sound. Baseline error sets a ceiling of $\text{SNR}_e \simeq 50$ in the simulation results. So as to avoid degrading the estimate of \sqrt{K} , we calibrate based on measured SNR_e values not greater than 30, from which the baseline error is removed following Equation 13, assuming a baseline SNR_e of 50.

$$\text{corrected } \text{SNR}_e = (\text{measured } \text{SNR}_e^{-2} - \text{baseline } \text{SNR}_e^{-2})^{-\frac{1}{2}} \quad (13)$$

In Figure 10c, corrected SNR_e measurements are plotted against SNR_e predictions for a particular scatterer field, subject to various strains. Covariances are ignored, so the calculation follows Equation 3. This is shown to be inapplicable, since the relation between predicted and measured SNR_e is neither linear within any single curve, nor do the curves at different strains come from a single function. Covariance estimates based on overlap have been used with Equation 4 to compute better SNR_e predictions plotted in Figure 10d. This shows a much higher degree of linearity, with very similar curves at the different strain levels. The least squares gradient (estimate of \sqrt{K}) is 0.186, with a correlation coefficient of 0.999 between the predictions and the measurements.

The type of noise may affect the value of \sqrt{K} , so calibration data from different types of simulations are shown in Figure 10e–f. All of these combine results across strains of 0.01%, 0.5%, 1.0% and 5.0%. The white noise in Simulation B can be thought of as the extreme opposite noise category compared to motion decorrelation in Simulations C and D. In the motion simulations SNR_s drops primarily because of the introduction of a noise ultrasound signal of similar statistical properties to the original signal. Despite this difference, within each simulation the linear relationship between predicted and measured SNR_e is good, as evidenced by the high correlation coefficients. The difference between the \sqrt{K} values of the different noise categories is discussed in Section 4. The variation means SNR_e cannot be controlled precisely by the DRS procedure outlined here, but since \sqrt{K} varies by much less than an order of magnitude, this approach may nonetheless be useful. The mean calibration value is $\sqrt{K} = 0.146$ for the 50×1 Survey or 0.131 for the 50×10 Survey. (The remainder of the experimental results are based on 50×10 Surveys.)

3.5 Variable noise experiment

Simulation E provides data with a variable noise level. Lateral variation in SNR is indicated by the mean SNR throughout different A-lines, shown in Figure 11a. Within each A-line, the SNR is highest at the focus, so the Survey produces a range of correlation coefficients as illustrated in Figure 11b. A normal strain image is noisier in the low correlation regions (Figure 11c), but the variation is reduced by applying DRS (Figure 11d).

DRS should maintain fixed SNR_e throughout the image, so for uniform strain fields we introduce a performance measure, denoted \mathcal{V} , recording spatial variation in SNR_e . This is evaluated by Gaussian filtering⁵ of strain squared errors to obtain local estimates of SNR_e . These are displayed in image form in Figure 11e–f. \mathcal{V} is calculated by dividing the overall SNR_e by the standard deviation of local SNR_e estimates. \mathcal{V} is improved substantially by applying DRS instead of fixed parameters.

3.6 *In vitro* and *in vivo* scanning

Data were recorded during scans using the Dynamic Imaging Diasus ultrasound machine with a 5–10 MHz probe (centre frequency of 6.0 MHz) from which RF data were sampled at 66.7 MHz using a Gage⁶ CompuScope 14200 analogue-to-digital converter attached to a PC running the Stradwin freehand 3D ultrasound software. Strain images were produced by analysis of consecutive frames of acquired data, and DRS was applied in off-line processing.

One subject is a homogeneous phantom. This is an agar cylinder (height 50mm, diameter 69mm) with 0.6wt% Al_2O_3 powder providing suitable scattering. A footprint extender was attached to the probe to apply a uniform compression. This is an interesting application for DRS, especially at high strains with a 1D estimator, since some regions decorrelate because of lateral expansion. A compression of 3.6% has been applied in Figure 12. Strain images in Figures 12c1–c2 demonstrate more uniform performance when DRS is applied. Figure 12b shows that larger windows and kernels are required at the sides, where lateral Poisson displacement is greatest. The \mathcal{V} values vary by only a small amount, which is perhaps because the baseline error in assuming uniform strain is quite high in this case.

To apply DRS to *in vitro* frames with varied SNR_s , the data in Figure 13 is from an elevational sweep performed for freehand 3D strain imaging, taken from [6]. The elevational movement is necessary to sweep a volume, but it also introduces noise, whilst the freehand scanning protocol means that the strain signal is sometimes very small. Figure 13 compares DRS (with suitable minSNR_e) against two fixed-parameter settings (one coarser than the other). DRS is shown to combine the advantages of both fixed settings, not smoothing unnecessarily when the data quality is good, but going to a coarse resolution in regions of low SNR_s to avoid presenting meaningless data.

⁵The Gaussian filter has the same dimensions as the filter for correlation that was mentioned in Section 2.5.

⁶<http://www.gage-applied.com>




Figure 10: DRS calibration. (a) Measured SNR_e against SNR_e predictions across a wide range of parameter settings, with no correction for baseline error. (b) Temporal shift rate of the PSF centroid against scatterer depth. This nonlinearity is a source of baseline error. (c–f) The effect of baseline error has been removed prior to evaluating the remaining results. (c) Calibration curves for Simulation B-20dB, if overlap covariances are ignored (Equation 3). The fit to any linear relation is poor. (d) Calibration curves for Simulation B-20dB with covariances considered (Equation 4). The correlation coefficient between predictions and measurements is 0.999, with a scale factor, $\sqrt{K} = 0.186$. (e) Calibration results with variances considered from various simulations for a 50×1 survey. (f) Calibration results for a 50×10 survey.




Figure 11: Variable noise simulation (1.0% strain). (a) Mean SNR values by A-line. (b) Smoothed arithmetic correlation coefficients from the Survey (black=0.7, white=1.0). (c) Strain image with fixed parameters, image-wide $\text{SNR}_e=7.70$ and $\mathcal{V}=2.08$. (d) Strain image with DRS, $\text{SNR}_e=7.51$ and $\mathcal{V}=4.63$. (e) Window length (max=91), strain and SNR_e images for fixed parameter strain imaging, image-wide $\text{SNR}_e=11.2$ and $\mathcal{V}=2.15$. (f) Window length (max=175), strain and SNR_e for DRS, $\text{SNR}_e=12.87$ and $\mathcal{V}=5.24$. (g) \mathcal{V} against SNR_e both for DRS and for fixed parameter values. Higher \mathcal{V} indicates that DRS reduces the variability of SNR_e .




Figure 12: Homogeneous phantom subject to 3.6% compression. (a) Survey strains. (b) DRS window lengths. (c–d) Strain and local SNR_e for 1: fixed parameters and 2: DRS parameters.

Further frames are presented from a freehand scan of human biceps. The images in Figure 14 illustrate the strain benchmark issue, which was introduced in Section 2.4. The extrapolated Survey strains are shown in Figure 14a. Figure 14b shows the DRS strain images when the SNR_e calculation is based on the mean strain. Basing $\hat{\text{SNR}}_e$ on the local strain, as in Figure 14c, gives a smoother image in the low strain region, but the window settings cover a very wide range. Figure 14d shows the image using the normal compromise strategy (local strains with saturation).

Finally, an *in vivo* image is presented in Figure 15, processed at a range of different $\min\hat{\text{SNR}}_e$ levels. Qualitative changes in smoothness and variability are obvious, although clinical experience may be required to determine appropriate settings for specific tasks.

4 Discussion

The calculation of $\hat{\text{SNR}}_e$, developed in Section 2, included numerous approximations, particularly the assumption of overlap fraction corresponding to the normalised covariance between neighbouring windows. However, the high degree of linearity in Figure 10d and the close fit to the linear model indicated by correlation values in Figure 10e–f serves as a partial validation of this approach.

The results also highlight a limitation: the calibration factors vary by as much as a factor of two, and they separate into two groups, based on whether the decorrelation was primarily due to white noise or to scatterer motion. In fact, this is to be expected. Equation 9 relies on the error autocorrelation distance having a constant value (*cf.*, justification of Equation 9 in Section 2.3.3), which is incorporated in \sqrt{K} . Errors with long-distance correlations cause greater loss of performance than short-distance errors with the same SNR_s . A more accurate prediction could be made if the error autocorrelation distance could be estimated. Otherwise, the approach followed here was to choose a suitable, typical value for \sqrt{K} , which offers useful if not precise control over SNR_e .

The variable noise experiment showed that, in addition to choosing sensible parameters for




Figure 13: Illustration of different strategies for consecutive frames from freehand 3D scanning of an olive-gelatin phantom. Image mean strains are noted on the right, where negative values indicate relaxation. (a) Fixed parameter settings for relatively fine resolution. (b) DRS with $\text{minSNR}_e=5.0$. (c) Fixed parameter settings for coarser resolution.




Figure 14: Demonstration of strain selection for the \hat{SNR}_e calculation, using scan of human biceps *in vivo* (image mean strain of -0.65% , $\min\hat{SNR}_e=6.0$). (a) Survey strains. (b) Window length and strain images when \hat{SNR}_e is based on the mean strain across the image. Low strain regions may be swamped by noise. (c) Images basing \hat{SNR}_e on the local survey strain. Low strain regions are analysed at extremely low resolution. (d) Compromise where \hat{SNR}_e is based on local strain with saturation.




Figure 15: Range of choices for $\min\hat{\text{SNR}}_e$ in another frame from the scan of human biceps (image mean strain of -0.64%).

any image, DRS is sensitive to within-image variation. This is evident from comparison of the SNR_e images in Figure 11e-f. It is also supported by the uniformity data in Figure 11g. At low SNR_e DRS is similar to fixed-parameter analysis, because it simply opts for the minimum window and kernel dimensions if $\min\hat{\text{SNR}}_e$ is already exceeded throughout the image. At moderate SNR_e levels, however, window and kernel dimensions are closely controlled so that SNR_e is roughly uniform. This pattern is illustrated in the plot of \mathcal{V} against SNR_e . \mathcal{V} becomes a less reliable measure at high SNR_e , however, since baseline errors eventually distort the results. Overall, this is strong evidence that DRS can lead to greater uniformity in SNR_e by means of varying the resolution.

However, it may not be intuitively obvious that uniform SNR_e is a sensible goal. As compared to the DRS image (Figure 11f), while the fixed parameter image (Figure 11e) is noisier in the corners, it clearly has higher SNR_e at the focal depth and on the right of the image. The important property of DRS is that $\min\hat{\text{SNR}}_e$ sets the degree of significance of contrast differences, so the information that is presented can be interpreted meaningfully. Sometimes this may blur over subtle features below the resolution setting, but the resolution level is at least discernable to the user, since a ripple of magnitude proportional to SNR_e is present with a wavelength that depends on the imaging parameters. By contrast, an image with variable noise at fixed resolution presents data where the significance of image features is less well defined, and interpretation is highly dependent on a process of cognitive filtering, where users must discriminate for themselves between those features that are meaningful, and others that are assumed to be noise.

The homogeneous phantom results in Figure 12 are further evidence that DRS can simplify image interpretation. The variable noise due to off-axis motion is a typical strain imaging phenomenon associated with high strains. No claim is made here that 1D tracking is optimal for high-strain work, where the addition of lateral tracking becomes advantageous. Against that, the outcome in Figure 12c2–d2 is encouraging, because it shows that DRS can adjust for the decorrelation caused by off-axis motion, so as to maintain the reliability of the image data. Indeed, this result is generally useful because lateral motion brings increased signal decorrelation, even if lateral tracking is employed.

Figure 14 illustrates a complication to this issue, in that the SNR_e is less well defined in the context of inhomogeneous scan subjects. Should noise be measured against the local strain, the local level of strain variation (perhaps a strain envelope) or simply compared to the same mean strain throughout the image? The mean strain option is unsatisfactory, because if a substantial region has relatively low strain, the local information can be completely lost under a noise level that is acceptable elsewhere in the image. On the other hand, it is obvious that the local strain cannot be used for $\hat{\text{SNR}}_e$, because strain zero crossings result in points where $\min\hat{\text{SNR}}_e$ can never be exceeded, even at the coarsest resolution — the long windows used in Figure 14c exemplify this scenario. Our compromise (saturated local strain) may be sensible, but there is undoubtedly scope for assessing more rigorously the effect on information content of estimation noise in different parts

of the image.

Nevertheless, in the present form DRS already behaves preferably compared to fixed-parameter settings when applied to real data. The sequence of images from a freehand 3D sweep in Figure 13 show that DRS moves between resolution settings in a manner that appears intuitively reasonable, so images are neither overly smoothed where the signal is strong, nor are misleading estimates displayed at inappropriately high resolution in regions of poor data. The question of exactly what level is suitable for the minSNR_e setting remains open, although various settings are illustrated in Figure 15. The optimal setting may depend on the degree to which users can discriminate for themselves between noise and signal features in poor image regions. For example, new users might benefit from a high minSNR_e to obtain reliable data (sometimes at low resolution). On the other hand, experienced users may be able to extract meaning from noisier images at finer resolution by comparing the presented data with their expectations. The significance of noise is probably application-dependent. This is equally likely to be the case regarding automated analysis procedures. Generally, clinical experience will be required to assess the most appropriate form and setting for future techniques similar to DRS.

5 Conclusions

We have examined resolution and estimation error in detail for a particular strain estimation algorithm. It was demonstrated that the best noise performance at any resolution is achieved if the window and kernel lengths are set to a particular ratio. A viable technique was outlined for predicting displacement estimation error, and translating this into strain estimation error.

These observations have been applied in the DRS system, whereby data are automatically processed with suitable parameters, so as to achieve a desired SNR_e whilst allowing the resolution to vary. This produces strain images in which the significance of image features can be interpreted with greater confidence.

Numerous tasks present themselves for ongoing development. The possibility of developing a superior DRS strategy based on a more rigorous analysis of image information content has already been discussed. There are also some performance improvements that could be achieved with relatively minor modifications to the current system. For example, estimation precision could perhaps be improved with negligible computational penalty by introducing covariance-weighted least squares strain estimation.

Ultimately, the key task that we have not addressed is adjusting DRS to reduce the computational load. The present system is time-consuming. For DRS (precisely as described here, with no particular attempt at optimisation) each strain image takes at least 30 seconds on a current standard PC with a 3 GHz processor, and extremely poor data takes up to five minutes if DRS opts for long, wide windows. The slow processing is despite an underlying strain estimation algorithm which is essentially fast (strain images from high quality data with reasonable SNR_e and acceptable resolution can be produced faster than 20 Hz). The algorithm is suited to a high level of parallel processing, operating simultaneously on each A-line, so real-time DRS processing could perhaps be achieved even with DRS in its current form. However, the algorithm will need substantial modification to achieve real-time imaging using standard, single-processor equipment. The Analysis is not the issue. Displacement estimation at both the Survey and Refinement stages can vary hugely in cost depending on numerous factors. Major costs are incurred by the unusually large number of estimation locations employed in the present version, and the occasional use of very large 2D windows. Several thousand samples are often covered by a single window. However, the same approach can perhaps be combined with intelligent downsampling at various stages, to reduce the computational load whilst avoiding excessive performance penalties. This demands further investigation.

References

- [1] 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.
- [2] J. F. Greenleaf, M. Fatemi, and M. Insana. Selected methods for imaging elastic properties of biological tissues. *Annual Review of Biomedical Engineering*, 5:57–78, April 2003.
- [3] J. A. Jensen. Field: a program for simulating ultrasound systems. In *Proceedings of the 10th Nordic-Baltic Conference on Biomedical Imaging*, volume 4, pages 351–353, 1996.
- [4] F. Kallel, T. Vargese, J. Ophir, and M. Bilgen. The nonstationary strain filter in elastography: Part II. lateral and elevational decorrelation. *Ultrasound in Medicine and Biology*, 23(9):1357–1369, 1997.
- [5] 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.
- [6] 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.
- [7] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Estimation of displacement location for enhanced strain imaging. Technical Report CUED/F-INFENG/TR 550, Cambridge University Department of Engineering, March 2006.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] T. Varghese and J. Ophir. The nonstationary strain filter in elastography: Part I. frequency dependent attenuation. *Ultrasound in Medicine and Biology*, 23(9):1343–1356, 1997.
- [13] M. Yamakawa and T. Shiina. Strain estimation using the extended combined autocorrelation method. *Japanese Journal of Applied Physics*, 40:3872–3876, May 2001.
- [14] Y. Zhu and T. J. Hall. A modified block matching method for real-time freehand strain imaging. *Ultrasonic Imaging*, 24:161–176, 2002.