
**NONPARAMETRIC SURFACE REGRESSION
FOR STRAIN ESTIMATION**

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

CUED/F-INFENG/TR 598

March 2008

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

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

NONPARAMETRIC SURFACE REGRESSION FOR STRAIN ESTIMATION

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 usually involves applying some form of smoothing filter to estimate gradients (corresponding to strains) after displacement measurements have been calculated from RF ultrasound data. All methods involve trade-offs between resolution and accuracy. Technical research often focuses on aspects of the RF signal processing, but differences in the smoothing method have a significant bearing on overall performance. We introduce a nonparametric regression method. Comparing with other linear-filtering techniques, this has advantages of higher image quality, greater versatility for practical applications, and (potentially) low computational cost. Its properties for strain imaging are examined through analysis and experiments. In addition to smoothing with uniform resolution, we demonstrate that nonparametric regression can be used to vary the resolution automatically based on indicators of data quality, thereby avoiding variation in the noise level within the image displayed. Interactions with RF signal processing parameters are also considered. This work indicates promising avenues for future research to improve the overall properties of practical strain imaging systems by taking a holistic approach to algorithm design.

1 Introduction

This paper is concerned with smoothing noisy displacement measurements to improve the quality of strain images. At the time of writing, it seems likely that ultrasonic strain imaging will be adopted in some form into routine clinical practice, within a decade, to support a still unestablished set of diagnostic tasks, primarily within the broad category of soft tissue examinations. Candidate applications discussed in the academic literature have included detection of soft tissue tumours [11, 31, 39], discrimination without biopsy between complex cysts and malignant breast lesions [3], monitoring of atherosclerosis [7, 8], detection and grading of deep vein thrombosis [10], assessment of skin pathologies [45] and evaluation of myocardial fitness [16].

Despite an increasingly large volume of literature, from a technical perspective it is not obvious how best to generate axial strain images from RF ultrasound data. Changes in details of the signal processing bring significant changes in the quality of strain images, and impact greatly on their suitability for any clinical application, yet the technical literature is far from converging on a consensus in terms of a preferred method.

Two essential tasks are required when computing a strain image from ultrasound data. Firstly, the data must be analysed to estimate tissue displacement between frames acquired before and after a deformation. Many methods have been proposed, *e.g.*, [1, 5, 4, 19, 21, 23, 25, 26, 29, 30, 37, 44, 49]. Secondly, strain must be estimated by differentiating the displacement field — this has less often been a focus for research. The two tasks can potentially be combined [28], but in this paper we are concerned with differentiation in the usual scenario, where displacement measurements are first produced by an independent stage of RF signal processing. The simplest differentiation method is differencing consecutive samples [27]. In practice, strain images are improved by incorporating some form of smoothing, since differentiation tends to amplify high-frequency components in the measurement noise.

A large number of differentiation techniques can be classified as linear filters: Every strain estimate is calculated as a linear combination of the displacement measurements, which can be inter-

preted as convolving the displacement data with a particular kernel. Examples include piecewise-linear least squares regression (PLLSR) [15], moving-average filtering [26] and staggered strain estimation [36]. In our preliminary analysis [22], we focused on identifying the resolution and accuracy (precision, *i.e.*, reciprocal of variance) of the resulting strain estimates based on typical assumptions. Our analysis showed that the precision of strain estimates at a given resolution is optimised by filters with kernels that are themselves smooth in shape. For instance, linear least squares regression is optimal for producing a unique gradient estimate over an entire set of data that genuinely consists of noisy measurements from a single linear trend, but PLLSR is suboptimal for estimating varying strains within a wider field of displacement measurements, because smoother kernels (*e.g.*, Gaussian) can include more displacement data when achieving the same resolution, and consequently achieve lower estimation noise [22].

With any particular level of smoothing, kernels with a smooth shape have the general advantage that they avoid leaving behind artefactual high-frequency noise. We argue that high-frequency artefacts after non-smooth filtering such as PLLSR not only detract from the aesthetic quality of strain images, but they can also lead to errors in image interpretation. There are various ways of generating and applying smooth kernels, so we acknowledge alternatives later in our discussion.

In our approach to smoothing, we focus on the roughness penalty method for nonparametric regression (NPR),[†] which is supported by a large body of literature [13]. Its advantages for strain imaging include the fact that the output solves an inference problem, which can be tuned based on any valid prior knowledge to produce strain imaging results that are in fact more probable than the output from other filters. More pragmatically, the behaviour of NPR depends on settings that may be thought of as a smoothing strength and data weights. NPR provides a convenient and versatile framework for applying any desired level of smoothing by choosing suitable settings.

The ratio between the smoothing strength and the data weight determines the frequency response, resolution and accuracy of NPR through simple, intuitive relationships. Furthermore, NPR is shown in Section 3.1 to have an appealing frequency response, with 80 dB/dec. attenuation after the cut-off frequency and no side-lobes. It is therefore easy to appreciate the effect of a change in the level of smoothing, so strain images can more easily be interpreted correctly.

There are two extreme modes of NPR smoothing. A uniform level of smoothing is applied throughout images where the data weights are constant. Alternatively, strain estimates can be produced with approximately uniform precision after the correct non-uniform smoothing, which is achieved by setting the weight of each displacement measurement equal to the reciprocal of its expected mean squared error. NPR can thus be harnessed to improve strain images when it is possible to estimate the accuracy of the displacement measurements.

Section 2 describes the principle behind NPR, followed by details of our implementation. Section 3 presents analysis characterising the resolution and accuracy of strain estimates as functions of the smoothing strength and the data weights. We present experimental results in Section 4, some of which illustrate the qualitative differences between NPR and PLLSR, as well as comparing the absolute accuracy in familiar quantitative terms, while other results verify the expected characteristics of NPR alone, for both uniform and non-uniform modes of smoothing. Finally, the discussion in Section 5 concerns the advantages of NPR compared to other strain estimation filters, the role of optimal filtering in strain imaging, and possible implications regarding the suitability of different displacement estimation algorithms.

While this paper focuses on strain image quality, computational issues are also important in the context of real-time applications. When reading the main text, it suffices to note that uniform smoothing by NPR can be implemented extremely efficiently with IIR filters, at a cost that is insignificant in most contexts. The cost of non-uniform smoothing is greater, but can be restricted to a similar level by optimisation techniques such as those discussed in Appendix A.

[†]In general, NPR refers to all forms of regression that are not based on fitting a parametric/functional model to measured data. The roughness penalty method is also known as spline smoothing. We prefer the general term, NPR, so as to avoid confusion with spline interpolation and B-spline approximation (techniques for solving different problems), and to acknowledge that practically identical results can be reached via differently motivated types of NPR such as variable-kernel smoothing [34].

2 Method

We apply NPR before differentiating displacement data, so as to replace the raw displacement measurements with a relatively smooth regression surface. In so doing, we aim to achieve a good balance between resolution and noise. The regression displacements are transferred to a grid of image pixels by a smooth interpolation method, and strain is estimated simply by differencing the values of neighbouring pixels.

2.1 Principle

NPR can be described most simply as an optimisation problem. A regression curve (in one dimension) or a surface (in two or more dimensions) is found by minimising a particular cost function [48]. A first group of terms, known as the data penalty, expresses the cost of a poor fit between measured displacements and values on the regression surface at the same positions, while a second group of terms, known as the roughness penalty, ensures that a relatively low cost is incurred by simple surfaces, such as planes, whereas a relatively high cost is incurred by complicated surfaces, such as those containing large numbers of sharp peaks and troughs. The mathematical form of the roughness penalty is an integral over the entire surface of the square of appropriate surface derivatives. For example, in this paper we consider the cost function

$$C_A = (\hat{\mathbf{d}} - \mathbf{d})^T \mathbf{W} (\hat{\mathbf{d}} - \mathbf{d}) + \int \int \left[r_x \left(\frac{\partial^2 \hat{d}}{\partial x^2} \right)^2 + r_y \left(\frac{\partial^2 \hat{d}}{\partial y^2} \right)^2 \right] dx dy, \quad (1)$$

where x and y denote lateral and axial distance (differentiation to estimate strain is in the axial direction); \hat{d} is the regression surface; $\hat{\mathbf{d}}$ and \mathbf{d} are vectors containing raw displacement measurements and corresponding points on the regression surface respectively, which are multiplied by the weight matrix, \mathbf{W} . The roughness penalty consists of squared partial second derivatives in the x and y directions, multiplied by the smoothing strengths, r_x and r_y .

The properties of the data penalty depend on the choice of \mathbf{W} . In general, \mathbf{W} must be a symmetric positive definite matrix. Broadly speaking, there are three possible modes:

- Uniform smoothing is enabled by setting \mathbf{W} equal to a constant data weight, w , multiplied by the identity matrix. In this case, the data penalty is a multiple of the sum of squared errors. Choosing large w implies that the displacement measurements are known to be accurate, so the regression surface should follow them closely (little smoothing), whereas small w implies that errors are probably large (resulting in coarse smoothing).
- Non-uniform smoothing can be applied if the level of measurement noise is known to vary, and the variation can be estimated. Non-uniform confidence in the data is then expressed by setting each element in the leading diagonal of \mathbf{W} to the expected precision of the corresponding displacement measurement.
- Whether smoothing is uniform or non-uniform, the filter can be adjusted to deal better with correlations between the errors in different displacement measurements, by including some non-zero values in the off-diagonal elements of \mathbf{W} , usually by setting \mathbf{W} equal to an estimate of the precision matrix (the inverse of the covariance matrix of measurement errors). We do not test this final extension, but the effects of correlated errors will be discussed.

As for the roughness penalty, we test the penalty set out in Equation 1. Other roughness penalties might be useful, but this is a good example. The aim is to associate a high cost with surface derivatives that are known *a priori* to be unlikely — *i.e.*, derivatives for which the measured data must present strong evidence if they are to be reflected in the shape of the regression surface. Penalties on second derivatives are a common choice [13], but a better penalty would reflect real-world properties of the strain estimation problem.

There is at least a strong argument for roughness penalties applying to derivatives no lower than second in the axial direction: A first derivative penalty would warp the regression surface

towards local estimates of the average displacement, biasing the strain estimates towards zero. In the limit of large r_y , each column of the image would converge on zero strain, with a single displacement measurement. The second derivative penalty instead has the effect of restricting the rate at which the strain changes in the axial direction, tending in the limit of large r_y to estimate a single strain in each column.

From a theoretical perspective, it can also be argued that a second derivative penalty in the axial direction leads to maximum *a posteriori* strain estimates if one assumes a certain prior model: the strain at any particular depth is most likely to be equal to the strain at the immediately preceding depth; changes in strain over depth are Gaussian-distributed; and errors in displacement measurements are also Gaussian-distributed. However, bearing in mind the simplicity of that theoretical argument, it may be that the roughness penalty could be improved with penalties on higher derivatives, possibly including mixed derivatives of the form $\partial^{n+m}/\partial x^n \partial y^m$. For example, a penalty on the third axial derivative would suppress the second derivative of strain, leading to strain varying linearly with depth in the limit of large r_y .

2.2 Implementation

Efficient NPR was first described by Reinsch [32], who converted the roughness penalty into a matrix expression that can be solved extremely efficiently by band-matrix methods for one-dimensional smoothing. Reinsch’s algorithm produces precisely the optimal continuous function, $\hat{\mathbf{d}}$, that minimises C_A in Equation 1 if we omit terms associated with the lateral direction. However, strain imaging is essentially an at-least-two-dimensional problem, so Reinsch’s algorithm cannot be applied directly.[†] The usual method for nonparametric surface regression from irregularly sampled data over two-dimensional space is to minimise C_A using a set of radial basis functions (the “thin plate spline” [9]), but this is far less convenient to compute. Fortunately, our data are regularly sampled, so we continue with a finite-difference approach, which to some extent retains the advantages of Reinsch’s algorithm.

We do not concern ourselves with finding the precise continuous function that minimises C_A . We find a similar solution by replacing second derivatives in the roughness penalty with second differences. A matrix \mathbf{M} is constructed such that $\mathbf{M}\hat{\mathbf{d}}$ is a vector listing every second difference (multiplied by r_x or r_y as necessary) throughout the grid of regularised displacement data. Details are provided in Appendix B. Our modified cost function is

$$C_B = (\hat{\mathbf{d}} - \hat{\mathbf{d}})^T \mathbf{W}(\hat{\mathbf{d}} - \hat{\mathbf{d}}) + \hat{\mathbf{d}}^T \mathbf{M}^T \mathbf{M} \hat{\mathbf{d}}. \quad (2)$$

A formula for the regression surface is found by setting $\nabla C_B = \mathbf{0}$, differentiating with respect to the value at every point on the regression surface, $\hat{\mathbf{d}}$, which gives

$$(\mathbf{W} + \mathbf{M}^T \mathbf{M}) \hat{\mathbf{d}} = \mathbf{W} \hat{\mathbf{d}}. \quad (3)$$

Efficient solution methods maximise and exploit sparsity in $(\mathbf{W} + \mathbf{M}^T \mathbf{M})$, which we discuss in Appendix A. For our main experiments we apply Cholesky factorisation, finding a lower triangular matrix, \mathbf{L} , with ones in its leading diagonal, and a diagonal matrix, \mathbf{D} , such that

$$\mathbf{W} + \mathbf{M}^T \mathbf{M} = \mathbf{L} \mathbf{D} \mathbf{L}^T. \quad (4)$$

The regression surface is found by a standard method: forward substitution solves $\mathbf{L} \mathbf{a} = \mathbf{W} \hat{\mathbf{d}}$, division solves $\mathbf{D} \mathbf{b} = \mathbf{a}$, and back substitution solves $\mathbf{L}^T \hat{\mathbf{d}} = \mathbf{b}$ [12].

3 Analysis

We now analyse NPR to provide simple models that predict the properties of the resulting strain images. We also discuss the situations in which these models lose accuracy; this occurs if NPR

[†]However, in Appendix A we discuss the use of one-dimensional algorithms to smooth each dimension independently in the fastest implementations of uniform smoothing.

is applied to displacement measurements produced with very large spacing between consecutive displacement samples, or using large RF-windows with very large overlaps, which introduces correlated measurement errors.

3.1 Resolution

Resolution is characterised by the spectral properties of the NPR filter, which is applied before differentiation. Differentiation changes the frequency response of the displacement filter, but it has unity gain with respect to strain data — it simply reveals the strain data [22].

It is helpful to note that NPR is similar to a convolution kernel with infinite extent. Although each point in the image takes a slightly different set of kernel coefficients, the effective kernel is substantially constant away from the edge of the image, so it is meaningful to refer to the shape and size of the effective kernel [34].

We consider continuous signals, so as to derive the ideal frequency response of NPR. This result is inaccurate where effective kernels would ideally be narrower than the sample spacing, but it is increasingly accurate for larger kernels. Consider a continuous displacement measurement, $\hat{d} = \cos(\omega_x x + \omega_y y + \phi)$, and the regularised displacement, $\hat{d} = H(\omega_x, \omega_y) \cos(\omega_x x + \omega_y y + \phi)$, with a uniform data weight, w . The frequency response is found by substituting these terms into Equation 1:

$$\begin{aligned}
 C_A &= \int \int (w(1 - H(\omega_x, \omega_y))^2 + (r_x \omega_x^4 + r_y \omega_y^4) H(\omega_x, \omega_y)^2) \cos^2(\omega_x x + \omega_y y + \phi) dx dy, \\
 \therefore \frac{\partial}{\partial H(\omega_x, \omega_y)} C_A &= 0 = w(H(\omega_x, \omega_y) - 1) + (r_x \omega_x^4 + r_y \omega_y^4) H(\omega_x, \omega_y), \\
 \therefore H(\omega_x, \omega_y) &= \frac{1}{1 + \frac{r_x}{w} \omega_x^4 + \frac{r_y}{w} \omega_y^4}. \tag{5}
 \end{aligned}$$

Unsurprisingly, NPR is a low-pass filter. The -6dB cut-off frequencies are $(w/r_x)^{1/4}$ in the lateral direction and $(w/r_y)^{1/4}$ in the axial direction. Similarly, for any resolution criterion we expect the resolving limit in each direction to be proportional to $(r_x/w)^{1/4}$ and $(r_y/w)^{1/4}$. This means that with a high data weight or low smoothing strength it is possible to resolve fine details in strain images, provided that the signal-to-noise ratio is adequate.

We empirically verify this relationship in Figure 1, with plots showing the equivalent kernel over the axial dimension for six examples of different smoothing strengths. The smoothing strengths are expressed in terms of a constant scale factor, a . The coefficients are found by calculating the response to an impulse in the centre of the image, employing the method from Section 2.2 with a uniform data weight. We also evaluate an indication of kernel width based on these kernel coefficients, which is plotted against $r_y^{1/4}$ to confirm the expected relationship.

3.2 Accuracy

It is important to be able to predict the accuracy of the strain estimates. This information can be exploited most simply by providing the user with an indication of strain image accuracy. It can also be used to improve the way that strain image data are subsequently processed, such as in weighted image persistence schemes [20], or to tune non-uniform smoothing, to even-out the level of accuracy throughout the strain image [18].

Predictions of strain accuracy are possible because precision indicators for displacement measurements can be calculated based on the RF ultrasound data [18]. Several different methods are applied in tests in Sections 4.3 and 4.4. Further analysis of subsequent filtering is required when predicting the precision of strain estimates.

The analysis must account for properties of the full kernel, including differentiation, which is found by differencing the equivalent kernels at consecutive sample positions — see examples in Figure 2a, where a again denotes a constant scale factor. However, for practical application we need

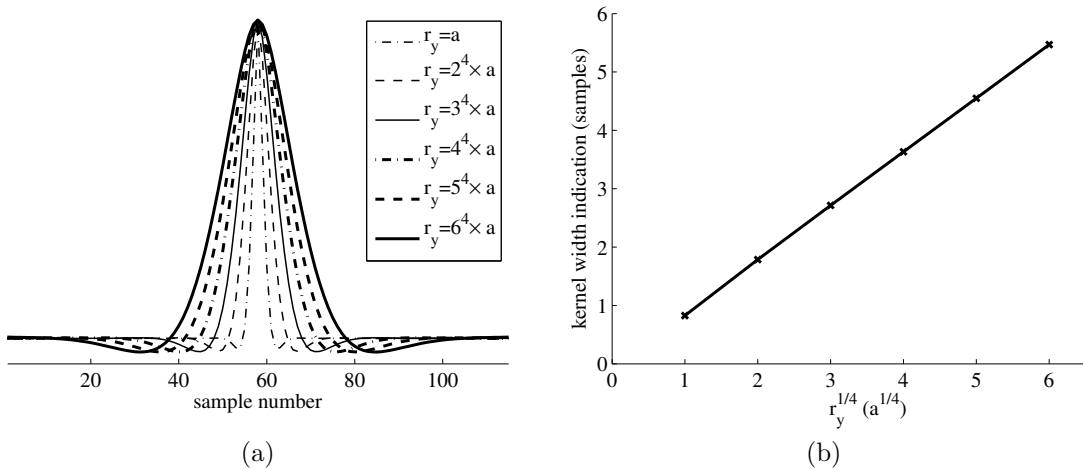


Figure 1: Illustration of (a) the coefficients of the equivalent kernel over the axial dimension for various smoothing strengths. The kernels have been scaled to show them at equal height. In (b) a plot of the kernel width indication (the standard deviation of sample number, weighted by kernel coefficients squared) against $r_y^{1/4}$ confirms the expected relationship.

to identify a simple relationship that can be calculated without evaluating the kernel coefficients, since they are not otherwise required for calculating the regression surface (*cf.*, Section 2.2).

A suitable relationship can be identified, bearing in mind the shape of the full kernel as illustrated in Figure 2a, with its length being proportional to $r_y^{1/4}$, and its width being proportional to $r_x^{1/4}$:

- Both positive and negative parts of the full kernel cover more data if its length increases. The resultant averaging confers a first-order increase in the precision of strain estimates.
- Lengthening of the full kernel also increases the distance between the centroids of its positive and negative parts. As with increasing the distance over which displacement measurements are differenced (*i.e.*, increasing Δy when calculating $\Delta \hat{d} / \Delta y$), this confers a further second-order increase in precision.
- Since there is no differentiation over the lateral direction, the lateral profile of the kernel resembles Figure 1a. Widening of the full kernel results in averaging, conferring a first-order increase in precision.

Overall, this indicates strain precision should be proportional to $r_x^{1/4} r_y^{3/4}$. If the smoothing strengths in both directions are linked to a single setting, $r = \alpha_x r_x = \alpha_y r_y$, then strain precision scales with r .

While the reasoning behind the r_x -dependency in this model is relatively straightforward, the r_y -dependency is more subtle. However, it can be tested by comparison with a brute-force method from Equation 6 of [22] for calculating the precision based on the coefficients of the full kernel. If the strain estimate, s , is given by

$$s = \sum_i f_i \hat{d}_i,$$

then its precision is the reciprocal of the estimation variance,

$$\sigma_s^2 = \sum_i \sum_j \left[f_i \times f_j \times \sigma_{\hat{d}}(i, j) \right], \quad (6)$$

where $\sigma_{\hat{d}}(i, j)$ is the covariance of displacement estimates. Errors in displacement estimates from non-overlapping RF-windows ($i \neq j$) are assumed to be uncorrelated, so $\sigma_{\hat{d}}(i, j)$ is zero. The




Figure 2: Illustration of (a) the full kernel for NPR strain estimation, formed by subtracting the equivalent smoothing kernels at successive sample positions. Three of the six examples are shown, scaled to display them with uniform height. In (b) the relative precision (calculated by the brute-force method) of strain estimates produced by applying six example kernels to uncorrelated displacement data is plotted against $r_y^{3/4}$ to test our model.

estimation variance, $\sigma_d(i, j)$ where $i = j$, may be assumed constant to assess the relative precision achieved by different filters. Based on Equation 6 the relative precision of six example NPR filters is plotted against $r_y^{3/4}$ in Figure 2b. This supports our model, because the relationship between these two quantities is represented by a line passing through the origin.

3.3 Effects of large sample spacing and large RF-windows

We have a model relating resolution and precision in strain images to the smoothing strength, but it becomes inaccurate if the sample spacing and/or RF-windows are large compared to NPR's equivalent kernel. If the sample spacing is sufficiently large, reasonably accurate estimates of the average strain between a pair of displacement measurements can be calculated just by differencing, without any smoothing. The sample spacing places lower limits on resolution and precision. If the strain estimates are already very coarse, smoothing that incurs further loss of resolution may be undesirable. However, aside from small savings on computation time, there is no strong argument in favour of sparse sampling.

Lower limits on resolution and precision are also raised by using RF-windows that are larger than the sample spacing. This is significant, because many past investigations have shown that large RF-windows reduce the level of noise in strain images.[†] We usually assume that (a) the precision of displacement measurements is proportional to the RF-window size, and (b) to a good approximation, displacement measurements from overlapping RF-windows are correlated, with error covariance equal to the variance of individual measurements multiplied by overlap as a fraction of RF-window size [18, 22, 43]. These assumptions imply the following consequences:

- Without smoothing, the precision of strain estimates produced by differencing consecutive displacement measurements is proportional to the square of the RF-window length for overlapping RF-windows, because not only are the displacement measurements more accurate, but differencing removes the correlated components of the measurement errors. This can set a high lower limit on precision.
- With light smoothing, when NPR's equivalent kernel is similar in size to the RF-windows, we

[†]However, the level of noise is only reliably reduced by large RF-windows if the displacement estimation algorithm incorporates amplitude compression, Amplitude Modulation Correction or adaptive stretching [19].




Figure 3: The effect of overlapping windows. (a) Plots of strain precision (brute-force method) against $r_y^{3/4}$ for various RF-window lengths, M , as multiples of the sample spacing. (b) Distorted kernels (the convolution of NPR kernels with 15-sample rectangular filters) showing the overall filtering effect of NPR with long RF-windows ($M = 15$).

have no simple formula for calculating the precision. It exceeds the lower limit imposed by the RF-window size, but smoothing brings a relatively small increase in performance because of the correlated errors.

- With heavy smoothing, when NPR’s equivalent kernel is much greater in size than the overlapping RF-windows, the likely outcome is that precision is approximately equal to the prediction of our model for smaller RF-windows (with dimensions matching the sample spacing). In other words, correlated components in the displacement data do not count.

Figure 3a illustrates these points for axial smoothing with various overlapping RF-window lengths. M indicates length as a multiple of the sample spacing. For example, $M = 2$ implies that the covariance between consecutive measurement errors is half the variance. The likely strain precision based on our assumptions is calculated following Equation 6 and plotted against $r_y^{3/4}$. In the absence of smoothing, strain estimation precision depends on the RF-window length. However, if the smoothing strength is large compared to the RF-window length, any difference in precision compared to $M = 1$ is an insignificant fraction of the total precision.

Numerically, the assumptions we have made are equivalent to saying that large RF-windows apply a moving-average filter of length M to the displacement data that would have been produced by RF-windows with length and width equal to the sample spacing. Deviation from the model of Section 3.2 can be considered as arising due to distortion of NPR’s full kernel when it is convolved with a moving-average filter, so we illustrate distorted kernels in Figure 3b, which may be compared with the original full kernels in Figure 2a.

4 Experiments

We test NPR with various smoothing levels, and PLLSR with various kernel sizes. In every test, the strain output is assessed after smoothing and differentiation of displacement measurements produced by the Weighted Phase Separation algorithm (WPS) [21],[†] applied to recorded and simulated RF ultrasound data.

[†]The displacement data are initially spread over a slightly irregular grid following Amplitude Modulation Correction [19]. We apply interpolation as described in [41] to return them to a regular grid.




Figure 4: Examples of simulations. (a) B-mode image. Strain images are calculated using RF-windows of 1 A-line by 31 RF-samples, followed by NPR with quite high smoothing strength. (b) Horizontal bands (2mm wide). (c) Diagonal bands.

The recorded data are *in vitro* and *in vivo* freehand scans from a Dynamic Imaging Diasus ultrasound machine (Dynamic Imaging Ltd., Livingston, UK) with a 5-10 MHz probe focused at 20 mm depth. Frames consist of 127 A-lines at 0.3 mm pitch, sampled at 66.7 MHz using a Gage[†] CompuScope 14200 analogue-to-digital converter. The spacing of RF-samples is 0.0115 mm, assuming a sound speed of 1540 ms⁻¹.

Simulated data were produced using Field II [14] with parameters to mimic the above ultrasound machine. Pre-deformation scatterer fields representing tissue with uniform echogenicity were constructed with 10⁶ scatterers distributed randomly with uniform probability density throughout a 50 × 50 × 6 mm volume, scanned over a 40 × 40 mm image region. Simulations offer precise control over the deformation parameters for quantitative tests of resolution and precision, while accurately representing the decorrelation due to changed interference patterns in strained scatterer fields. White noise was added, reducing the electrical SNR to 20 dB.

4.1 Resolution

The method for this test is similar to that in [18]. Measurements of the resolving limit due to the parameters of NPR or PLLSR enable us in later tests to compare different smoothing methods across appropriately matched settings.

We simulated narrow bands of zero axial strain separated by a narrow band at the 1% background strain. The smoothing methods were applied to displacement data from small RF-windows (31 RF-samples long, 1 A-line wide, sampled at intervals of 15 RF-samples axially and 1 A-line laterally). For example, Figure 4 shows a typical B-scan, in which no features are visible, alongside strain images with horizontal bands and diagonal bands, testing axial and lateral resolution respectively (vertical bands cannot be used for testing lateral resolution, because they do not represent a physically feasible deformation).

The feature scale in Figure 4 — *i.e.*, the width of each band — is 2.0 mm, and the stiff-soft-stiff pattern is repeated at three depths. Various feature scales were simulated, spanning 0.5–2.0 mm. Contrast is measured by differencing the mean strain estimate in the soft band and the mean strain estimate in the stiff bands on either side. For the diagonal bands, contrast is recorded based only on strain estimates at the centre of the image (1 mm around the focal depth), because a high shear strain causes extremely high decorrelation away from the focus, making strain data much less reliable.

The resolving limit is found at the feature scale where contrast crosses zero. This usually occurs between two of the simulated feature scales, where the zero crossing is estimated by linear

[†]<http://www.gage-applied.com>

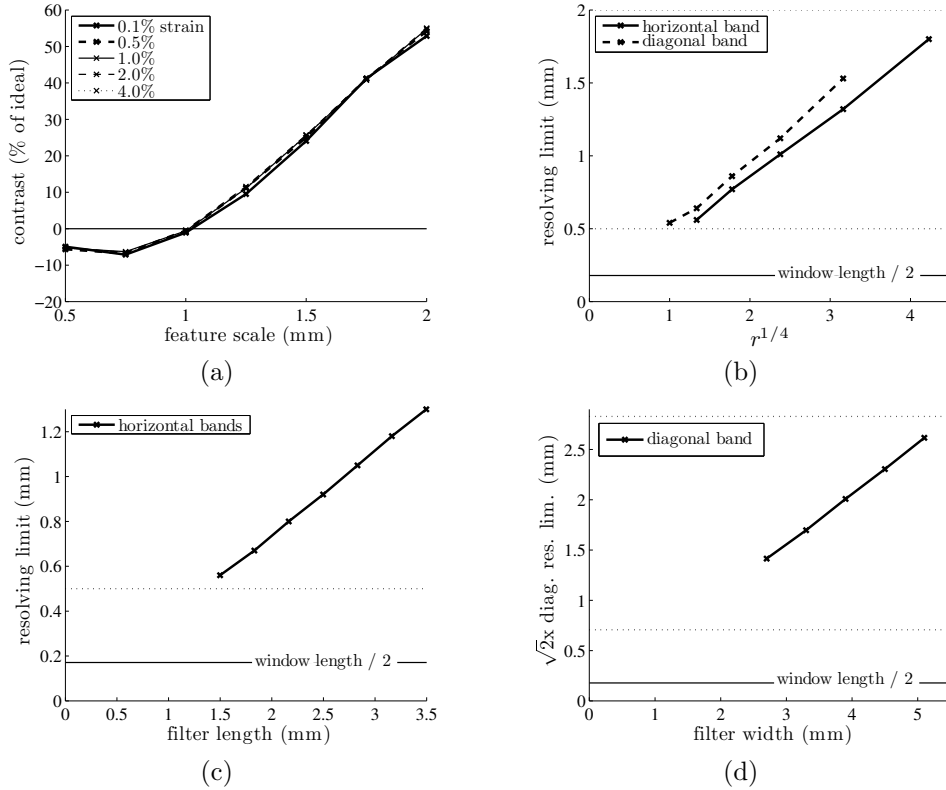


Figure 5: Resolution tests. (a) Contrast against feature scale for NPR applied to horizontal bands with $r^{1/4} = 2.38$. (b) Axial and lateral resolving limits against $r^{1/4}$. (c) Axial resolving limit against PLLSR filter length. (d) Lateral resolving limit based on contrast of diagonal bands against PLLSR filter width, using short, wide filters.

interpolation, and results are averaged across five strain levels, as in Figure 5a. No result is recorded if there is no zero crossing between 0.5–2.0 mm, so the bounds of estimable resolution limits are indicated on each plot of results in Figure 5b–d. We also indicate half the window length, to show that the window length did not affect these resolution data.

NPR is tested with r_x and r_y in a fixed ratio chosen by eye to give isotropic resolution ($r = \alpha_x r_x = \alpha_y r_y$). Results in Figure 5b indicate proportional relationships with $r^{1/4}$, as expected. Least squares estimates of the constants of proportionality are 0.424 mm for axial resolution and 0.478 mm for diagonal resolution (1.13 times worse). Therefore, resolution in the lateral direction is expected to be 1.53 times worse than the axial resolution, based on theory in Appendix C.

PLLSR is tested with various filter lengths, leading to the results in Figure 5c. A least squares fit indicates that the axial resolving limit is 0.371 times the filter length. The lateral resolving limit is tested using short kernels (0.832 mm over the axial direction) with much greater width (number of A-lines multiplied by pitch in mm). A least squares fit to the results in Figure 5 indicates that the lateral resolving limit is 0.532 times the filter width.

Based on these results, NPR and PLLSR are compared in later tests across settings that are matched in terms of lateral and axial resolution. This seems sensible, although it is only approximate, since matched resolution in the axial and lateral directions does not imply matched resolution of diagonally-oriented inhomogeneities.[†] Nevertheless, this approach is important so that we can make meaningful comparative observations, which would be impossible if NPR and PLLSR were tested across arbitrary settings.

[†]Judging by the filters' 2D frequency responses, the version of NPR that we tested is likely to achieve finer diagonal resolution than the "matched" PLLSR filter.




Figure 6: Scan of tensed forearm (no ground truth). (a) B-scan. The strain images should have the same resolution based on (b) heavy NPR smoothing, and (c) a large PLLSR kernel.

4.2 Comparison

For subjective assessment, several example images are presented with fairly heavy, uniform smoothing by NPR and PLLSR, using displacement data from RF-windows of 49 RF-samples by 1 A-line, sampled at intervals of 15 RF-samples axially and 1 A-line laterally. The approximately matched resolution settings are $r^{1/4} = 3.64$ for NPR, compared with PLLSR kernels of 13 A-lines by 25 displacement samples.

Figure 6 shows strain images from a transversal scan of the forearm. The unsmoothed strain image, which appeared to consist only of noise, is not shown. After smoothing, it is clear that NPR yields a strain image with greater aesthetic appeal. This is partly because the PLLSR image contains a perceptually significant level of high-frequency noise, even after smoothing. These images show typical qualitative differences. However, owing to the absence of a ground truth, it is difficult to judge the significance of differences in the structure that appears within each image.

Simple differences can be interpreted more confidently considering strain fields with known properties. Figure 7 shows the region of interest from a simulation along the lines of Figure 4b for testing axial resolution, with a feature scale of 1 mm. The detail is shown with light smoothing in Figure 7a1–2, but it is unresolved under heavy smoothing in Figure 7b1–2.

Inevitably, heavy smoothing loses resolution, but qualitative differences between the images in Figure 7b1–2 might nonetheless be significant. The NPR result is good, because the blurred feature appears fairly uniform, and must be interpreted either as a homogeneous stiff inclusion, or as an inhomogeneous stiff inclusion with unresolved detail. The disadvantage with the PLLSR image is that it contains a narrower band of very low strain, with bands of medium strain on either side. This is suggestive of an inhomogeneous inclusion with elevated stiffness towards the centre. In other words, PLLSR gives a misleading representation of the unresolved detail. Strain estimates are plotted against depth in Figure 7c–d, and compared with the unresolved true strain. The problem shown in Figure 7d is residual high-frequency information that remains after PLLSR filtering, giving an impression of details such as sharp strain boundaries, which do not match the details in the true strain.

Figure 8 provides another example of different behaviour between the two smoothing methods. Scans of this phantom are affected by an artefact associated with reflections from a wire, which is actually outside the image plane, but reflects sufficiently strongly from the edge of the resolution cell that it produces a bright streak. This is most clearly visible to the left of the olive. In the NPR strain image, the main effect is a high strain gradient at the location of the original artefact, marked B in Figure 8e. In the PLLSR strain image, the artefact appears as two artefactual stripes, above and below the original, marked A and C in Figure 8e. Although the artefact's appearance

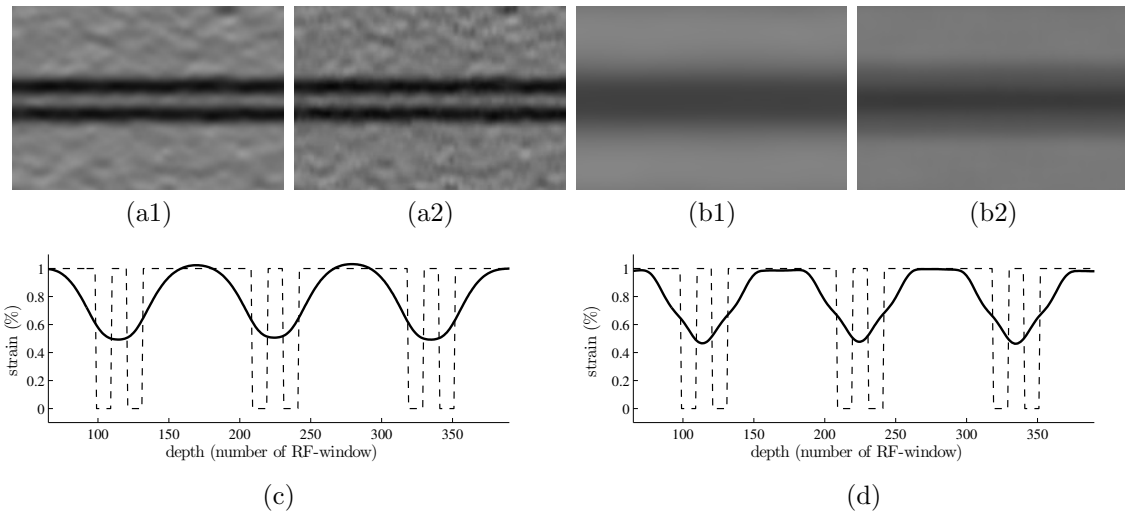


Figure 7: Detail from strain images of the axial resolution simulation with a 1 mm feature scale. Light smoothing resolves the detail, using (a1) NPR and (a2) PLLSR. Heavy smoothing using (b1) NPR and (b2) PLLSR produces strain images in which the detail is lost. PLLSR produces an image which is more likely to be misinterpreted. The differences are clearer to see in (c) and (d), where plots show the strain estimates produced after heavy smoothing (solid lines) against depth, compared with the unresolved true strain (dashed line). (c) Strain estimates by NPR vary smoothly over depth. (d) Strain estimates by PLLSR are suggestive of sharp boundaries that misrepresent the true strain.

in both strain images is essentially unrelated to any real change in stiffness, that does not make differences in its appearance inconsequential. The NPR strain image has the advantage of being affected most strongly by the artefact at the location where it appears in the B-mode image, so the correspondence between related image features is obvious. On the other hand, the response of PLLSR to a displacement impulse produces the most distinct strain responses at either end of the PLLSR kernel, which has the unfortunate effect of doubling the number of artefacts. This makes image interpretation more difficult for the sonographer.

It seems fairly clear that NPR is preferable to PLLSR, but the differences observed so far relate essentially to the way that inhomogeneous strain data are displayed. It is desirable to make quantitative comparisons. We record measurements of the estimation signal-to-noise ratio [43],

$$\text{SNR}_e = \frac{s}{\sigma_s}, \quad (7)$$

where s is the strain and σ_s is the root mean squared estimation error in a homogeneous region. σ_s may alternatively be thought of as the reciprocal of the square root of strain precision.

SNR_e can be measured easily by analysing uniform-strain simulations. Results are shown at two different strains in Figure 9, plotted against the axial resolving limit of each smoothing filter. These are average results over five frames of data, each with 128 A-lines and 240 strain estimates per A-line, ignoring data in the outer 10% of the image, where the level of smoothing by PLLSR cannot be changed arbitrarily. The raw displacement data were produced using RF-windows of 1 A-line by 31 RF-samples, sampled at intervals of 15 RF-samples axially and 1 A-line laterally.

The data in Figure 9 are matched in terms of both the axial and lateral resolving limits. However, owing to the discrete steps in which the PLLSR kernel size can be changed, the lateral resolution cannot be varied in perfect proportionality with the axial resolution. Instead, we provide a lower comparison (where the kernel is slightly too narrow) and an upper comparison (slightly too wide). The ideal comparison would be a smooth curve approximately bisecting these experimental results.

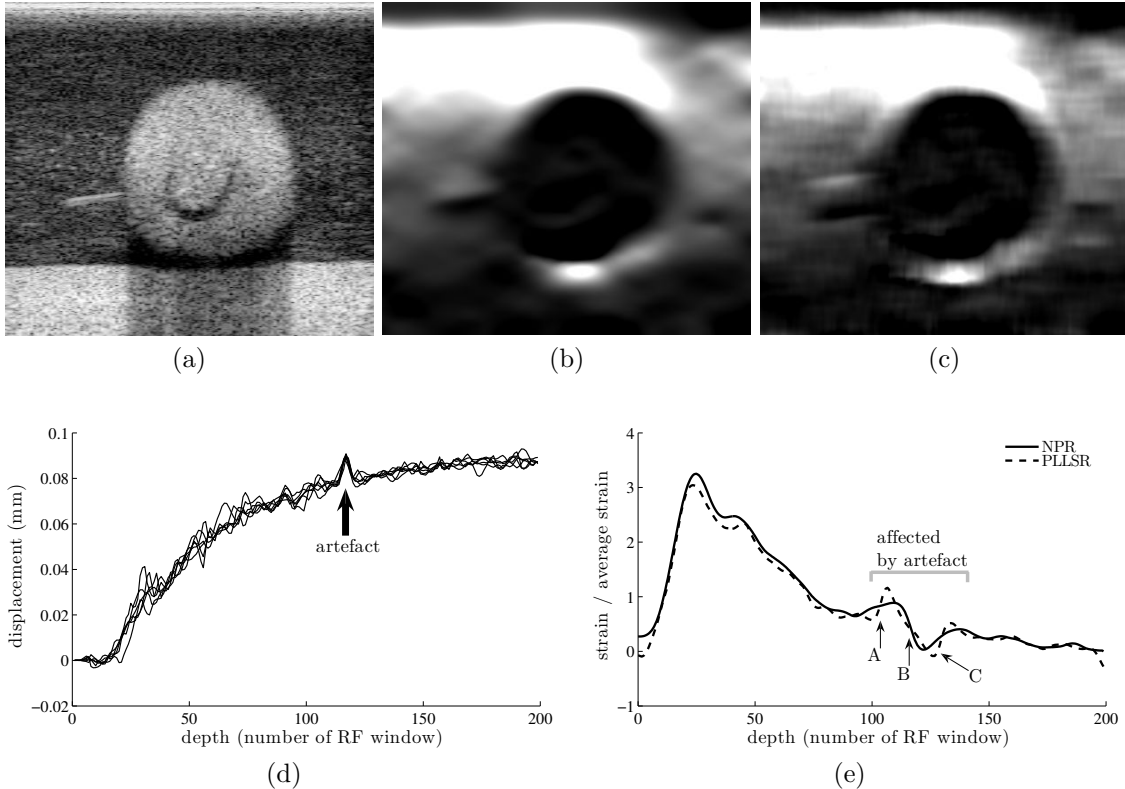


Figure 8: Scan of a phantom (olive in gelatin) showing an artefact. (a) B-mode. The artefact is a streak on the left, caused by a strong reflection from a metal wire at the edge of the elevational resolution cell. The wire is mostly outside the scan plane, so its displacement is not continuous with the rest of the image data. This results in a strain image artefact. Heavily smoothed strain images are shown using (b) NPR and (c) PLLSR. The data properties are highlighted by (d) showing a plot of displacement against depth in seven A-lines intersecting the artefact, and (e) showing a plot of strain estimates against depth from NPR and PLLSR in a single A-line.

Figure 9 shows that NPR achieves slightly higher SNR_e than PLLSR at each resolution. However, the differences are undramatic. SNR_e indicates the amount of data covered by the filter. Theoretically, we expect filters with a smooth profile to achieve slightly higher SNR_e than PLLSR, because they span more data at each resolution [22]. However, this improvement is of less significance than the obvious differences in Figures 6–8 in the way that inhomogeneous strain fields are reconstructed, owing to the different shapes of the kernels. In any event, the SNR_e comparison shows that if we opt for NPR because of obvious qualitative advantages, improvements in the images are not achieved at any cost in terms of absolute accuracy in a quantitative sense.

4.3 Accuracy

Improving the accuracy and interpretability of individual strain images is just one aspect of NPR’s appeal as a smoothing method. Images could probably be improved similarly by filters designed in a more heuristic manner, and heuristics might end up approximating the behaviour of NPR. We noted in Section 1 that behaviour identical to the roughness penalty version of NPR can be achieved by different approaches.

The advantage of implementing NPR in the framework presented is that the level of smoothing can be modified very conveniently by changing the smoothing strength, which can be set to any positive value. Kernel-based filtering methods are often limited by parameters that have to be changed in discrete steps. The continuously-variable smoothing strength should have a well-behaved impact

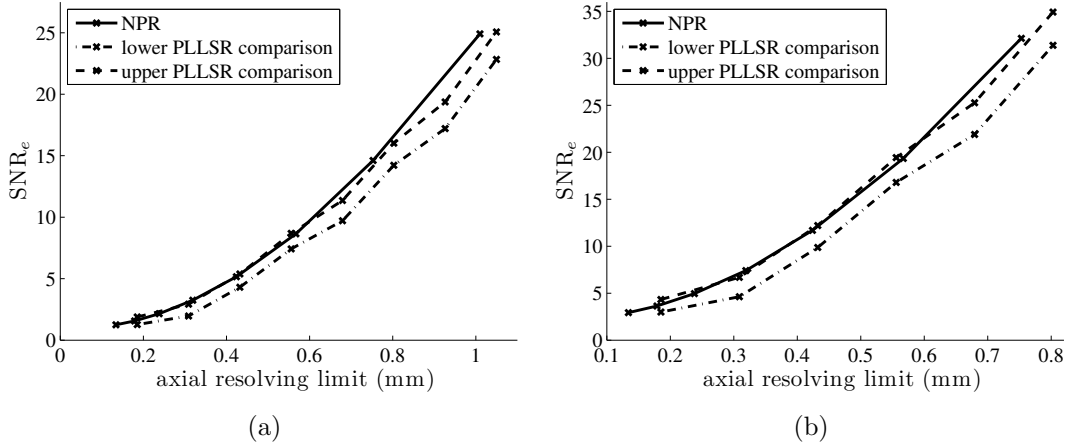


Figure 9: Comparisons of SNR_e between NPR and PLLSR using the same data with (a) 0.5% strain and (b) 2.0% strain.




Figure 10: Axial-only smoothing ($r_x = 0$, variable r_y). Plots of SNR_e against $r_y^{3/8} \times s \times \Delta t \times \sqrt{T \times X \times \bar{Q}}$. The plots on different axes differ in terms of the precision indicator, \bar{Q} , which is based on (a) correlation coefficient, and (b) phase residuals.

on NPR's properties, which we now test.

According to the model developed in Section 3.2, strain estimation precision should be proportional to $r_y^{3/4} \times \Delta t^2 \times T \times X \times \bar{Q}$, where Δt is the sample spacing, T is the RF-window length, X is the RF-window width and \bar{Q} is the mean value of a precision indicator based on the RF data. It follows that SNR_e in any particular image should be proportional to $r_y^{3/8} \times s \times \Delta t \times \sqrt{T \times X \times \bar{Q}}$.

Plots of SNR_e against this quantity are shown in Figure 10 for simulations of uniform-strain deformations with strains of 0.1%, 0.5%, 1%, 2% and 4%, using RF-windows at a spacing of 20 RF-samples down each A-line, with dimensions of 5 A-lines by 21 RF-samples (almost no overlap). The role of \bar{Q} is to adjust for different levels of decorrelation, which arise here owing to the different levels of strain. Results are shown for two alternative precision indicators. Firstly, $c/(1-c)$ is a relatively well-known precision indicator, where c is the correlation coefficient between aligned pre- and post-deformation RF-windows. Secondly, \bar{Q} is set equal to the reciprocal of the mean squared residual error when a weighted linear least squares regression line is fitted to the phase differences




Figure 11: Variable noise experiment: (a) B-mode image with variable SNR. Strain images are shown with $\text{SNR}_e=7.0$, based on displacement measurements from RF-windows of 1 A-line by 51 RF-samples at spacing of 1 A-line laterally and 50 RF-samples axially. (b) Strain image with uniform smoothing. (c) Strain image with non-uniform smoothing. The remaining images show variation in SNR_e , based on Gaussian-filtering of mean squared errors. (d) Variation in SNR_e with uniform smoothing. (e–g) show variation in SNR_e with non-uniform smoothing. The precision measures used for varying the data weights are based on (e) correlation coefficient, (f) phase variance and (g) phase residuals.

throughout aligned RF-windows.[†]

Our model of NPR’s behaviour is validated by the fact that both graphs show straight lines. The only departure from linearity is a small curved section on the left, where the finite sample spacing imposes a lower limit on precision. Notably, although the two methods for calculating precision indicators are quite different, both approaches are successful in adjusting for changes in the levels of decorrelation affecting the displacement data from short RF-windows. This is shown by the colinearity of results at different strains.

4.4 Non-uniform smoothing

We expect non-uniform smoothing to be a particularly good application for NPR, highlighting its versatility. Kernel smoothing methods are difficult to apply for non-uniform smoothing, because there is a risk of introducing false boundaries where filtering parameters change, which appear to represent changes in tissue properties [18]. Continuously-variable smoothing should be a definite advantage with NPR.

We aim to minimise the level of variation in the accuracy of strain estimates, by allowing the resolution to vary throughout every image. Figure 11 shows results for a simulated 1% strain in a scatterer field with non-uniform scatterer density, resulting in a non-uniform signal-to-noise

[†]This is particularly easy to compute when using WPS [21].




Figure 12: Two examples with and without a modified weighting matrix. Top line: (a) Strain image from diagonal band simulation with uniform smoothing. (b) Strain image with non-uniform smoothing, approximately matched in terms of resolution in the centre of the image. (c) B-scan from cyst phantom. (d) Strain image with uniform smoothing. (e) Strain image with non-uniform smoothing.

ratio.[†] Displacement measurements are calculated using RF-windows of 1 A-line by 51 RF-samples, sampled every 50 RF-samples axially and 1 A-line laterally. Strain estimates produced by uniform smoothing have non-uniform accuracy, as shown in Figure 11b where $\text{SNR}_e=7.0$. Non-uniform smoothing is applied to produce Figure 11c, by setting the diagonal elements in \mathbf{W} equal to the precision indicator of each displacement measurement, calculated based on correlation coefficient, again achieving $\text{SNR}_e=7.0$ for the image as a whole.

We measure local SNR_e by filtering the mean squared error with a Gaussian kernel, so as to observe the degree of variation. These data are illustrated in Figure 11d–g using a scale from black (0) through mid-grey (7.0) to white (14.0). Figure 11d is the result with uniform smoothing. Local SNR_e with non-uniform smoothing (as above) is shown in Figure 11e to be far more consistent. Figure 11f–g shows two further results for non-uniform smoothing, where the data weights are set according to different precision indicators. One alternative is to calculate the reciprocal of a weighted variance of phase differences between pre- and post-deformation RF-windows. The final alternative is the calculation involving phase residuals, as described in Section 4.3. Anyway, these results show remarkably similar variation in SNR_e . There are probably yet more alternatives for calculating precision indicators that could be devised to achieve similar results.

Figure 12 shows further qualitative examples of non-uniform smoothing, using the same dimensions and spacing of RF-windows, with the precision indicator based on correlation coefficient. Ultrasound data in the diagonal resolution simulations, such as the example with a 1.5 mm feature

[†]These data were previously used to test a parametric method for non-uniform smoothing. See Fig. 11 in [18].




Figure 13: Resolution with interaction between NPR and large RF-windows. (a) Contrast against feature scale for NPR applied to horizontal bands with $r^{1/4} = 2.38$. (b) Axial and lateral resolving limits (resolution of horizontal bands and diagonal bands respectively). Note the axial resolving limit that corresponds to half the window length, as indicated.

scale in Figure 12a–b, are subject to high decorrelation away from the focus. This gives rise to obvious variation in the level of accuracy when the smoothing is uniform, which is reduced greatly by non-uniform smoothing. The images in Figure 12c–e come from a scan of an anechoic cyst in a breast biopsy phantom (Computerised Imaging Reference Systems, Inc.[†] Model 052). Displacement estimation based on ultrasound data from the anechoic region is essentially impossible. The strain image produced by uniform smoothing is mostly convincing, but the cyst appears as a noisy region. Non-uniform smoothing in Figure 12e produces strain data with similar resolution throughout most of the image, but it blurs over the cyst, showing no detailed data, just indicating at low resolution that the cyst is stiff.

This shows that non-uniform smoothing is potentially a useful technique. It is worth pointing out that the mechanism for achieving uniform SNR_e — setting data weights equal to the precision indicator — is a special property of the roughness penalty that we are testing. To achieve the same effect with a different roughness penalty, or in 3D, the precision indicator would need to be raised to the correct power when setting data weights. Anyway, eliminating all variation in the noise level is probably not the ideal approach for exploiting precision indicators. Alternative types of non-uniform smoothing are discussed in Section 5.

4.5 Large RF-windows

We have demonstrated advantages of applying NPR to displacement measurements from small RF-windows. However, we must also consider the effect of applying NPR to displacement measurements from large RF-windows, because the strain imaging literature describes many displacement estimation algorithms based on overlapping RF-windows that may be long [18, 35] and/or wide [49]. In this case the effect of NPR filtering is less easy to predict, since filtering properties of the large RF-windows may significantly distort the equivalent kernel.

Figure 13 shows how large RF-windows change the resolution of strain images. RF-windows of 3 A-lines by 101 samples, sampled every 15 RF-samples axially and 1 A-line laterally, are used for repeating the resolution test (*cf.*, Figure 5a–b). In terms of physical distance, the likely axial resolving limit due to the window length is indicated in Figure 13b. This is in good agreement with the measured results for light NPR smoothing, when RF-window size dictates the resolving limit. Heavier NPR smoothing achieves more or less the same resolution as with shorter windows (*cf.*, Figure 5b).

[†]<http://www.cirsinc.com>




Figure 14: Strain images from scan of forearm. Five images using displacement data from large RF-windows of 151 RF-samples by 7 A-lines sampled every 1 A-line laterally and 15 RF-samples axially: (a) No smoothing. (b) NPR smoothing to match the resolution of the RF-windows. (c) PLLSR smoothing to match the resolution of the RF-windows. Strain images are also shown with the same smoothing parameters as in Figure 6 for (d) NPR and (e) PLLSR. (f) The final image is Figure 6c convolved with a moving-average filter of the same size as the large RF-windows.

We repeat some of the comparisons between NPR and PLLSR, now using larger RF-windows. Figure 14 shows further versions of the strain image from a transversal scan of a forearm, now using RF-windows of 151 RF-samples by 7 A-lines (*cf.*, Figure 6). These large RF-windows give an informative strain image even in Figure 14a where no smoothing is applied.

In the past, it has been suggested that the size of smoothing filters should be chosen to match the resolving limit imposed by the RF-windows [18, 36]. This is done for NPR and PLLSR to produce the images in Figure 14b–c. These NPR and PLLSR images are very similar, which may initially be surprising. The reason is that, if the smoothing kernels and RF-windows are more or less the same size, then the similarity of the overall kernels — effectively a convolution of two filters — reflects the large extent to which the properties are determined by the RF-windows.

The quantitative comparison is repeated, recording SNR_e based on displacement data from relatively large RF-windows of 3 A-lines by 101 RF-samples. The results are plotted in Figure 15, where the horizontal axes are the axial resolving limit for the smoothing filters when operating on data from small RF-windows. We expect the axial resolving limit associated with these RF-windows to be approximately 0.56 mm, so results to the left of 0.56 mm on the graph mostly reflect properties of the RF-windows, whereas differences between the smoothing filters become

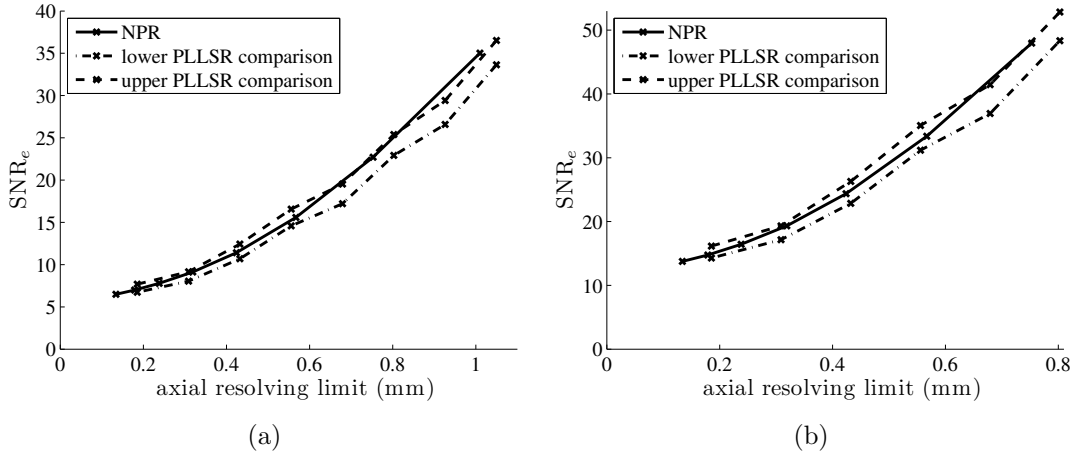


Figure 15: Comparisons of SNR_e between NPR and PLLSR using the same displacement data from large overlapping windows. (a) 0.5% strain. (b) 2.0% strain. The horizontal axis is the axial resolving limit imposed by the smoothing filter. The actual resolution will be different in this case, because the large RF-windows impose a lower limit of approximately 0.56 mm.

more important to the right of 0.56 mm. NPR and PLLSR achieve essentially the same SNR_e where the RF-windows are larger, but NPR does better than PLLSR when heavy smoothing is applied.

Continuing with this observation, Figure 14d–e shows further strain images with heavy smoothing, where the smoothing kernels are larger than the RF-windows. Comparing against Figure 6b–c, we see that heavy smoothing causes differences to re-emerge between strain images produced by NPR and PLLSR.

The smoothing parameters in Figure 14d–e are the same as in Figure 14b–c. Interestingly, the PLLSR strain image in Figure 14e appears greatly improved compared to the result for smaller RF-windows. It is no longer visually obvious which is more believable out of the NPR and PLLSR strain images. However, going by our earlier observations regarding structural differences between NPR and PLLSR strain images, we still expect that the NPR image is more accurate.

Any improvement in the NPR strain image is barely apparent between Figures 6b and 14d. NPR already achieves a good result using data from small RF-windows. The differences are much more obvious between the PLLSR images with different RF-window dimensions, because most of the high-frequency noise has been eliminated in Figure 14e. However, we could achieve a similar effect at much lower computational cost by applying a moving-average filter to Figure 6c. This is undertaken to produce the similar image in Figure 14f. We suggested in Section 3.3 that the main “benefit” of large RF-windows might be a filtering effect. When this is the case, the smoothing associated with large RF-windows is suboptimal in conjunction with NPR, because a more accurate result would be achieved by smaller RF-windows with an increased smoothing strength.

A further indication of the role of large windows is provided in Figure 16, showing precision plotted against $r_y^{3/4}$ for axial-only smoothing of a simulated 1% strain (*cf.*, Figure 3a). Each curve represents experimental results based on displacement data from RF-windows of a particular length, with the same width in all cases. Although the proportional difference in precision is significant when the level of smoothing is small (on the left of the graph) the proportional difference diminishes as the level of smoothing increases (on the right of the graph). This is consistent with the effect of different RF-window sizes being similar to applying another convolutional filter, which does increase precision, but does not achieve extra performance beyond that which can be achieved by modifying the smoothing filter.

Figure 17 shows results from repeating the accuracy test using displacement data from RF-windows of 1 A-line by 101 RF-samples with axial spacing of 20 RF-samples (*cf.*, Figure 10). SNR_e is plotted against $r_y^{3/8} \times s \times \Delta t \times \sqrt{T} \times X \times Q$, with plots showing different results for each of two different methods of calculating the precision indicator, Q . We show results for Q based

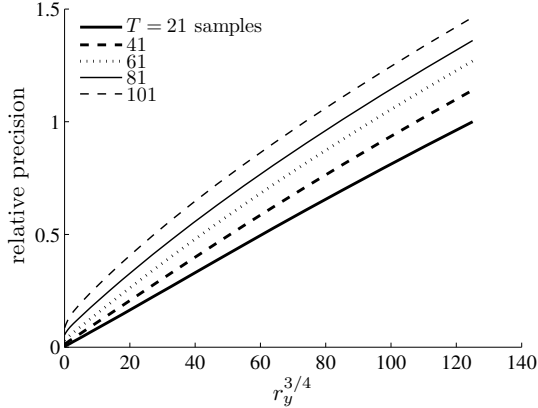


Figure 16: Plots of strain precision following axial smoothing of displacement data from RF-windows at an axial spacing of 20 samples, width of 5 A-lines and various lengths, based on simulations of 1% strain. The values are scaled relative to precision with the heaviest smoothing of displacement data from the smallest RF-windows. In terms of the covariance between neighbouring measurements, the RF-window lengths of 21, 41, 61, 81 and 101 samples correspond to $M = 1$ (uncorrelated data), 2, 3, 4 and 5 respectively (*cf.*, Section 3.3).




Figure 17: Axial-only smoothing ($r_x = 0$) applied to displacement data from long, overlapping RF-windows. Plots of SNR_e against $r_y^{3/8} \times s \times \Delta t \times \sqrt{T \times X \times \bar{Q}}$. The plots on different axes differ in terms of the precision indicator, \bar{Q} . (a) \bar{Q} based on correlation coefficient. (b) \bar{Q} based on phase-difference residuals.

on the correlation coefficient $[c/(1 - c)]$, because this measure has often been mentioned in the literature, although we know that it becomes inaccurate when using long RF-windows [18]; this is shown by the results in Figure 17a, because the precision indicator fails to adjust for the different levels of decorrelation in the data at different levels of strain.

When long RF-windows are used for measuring displacement, the precision indicator based on phase residuals is much more accurate. If we are correct in assuming that large RF-windows have a very similar effect to moving-average filters, then we expect the SNR_e lines to converge asymptotically on the same linear relationship as for non-overlapping RF-windows. Since the RF-windows used in Figure 10 are five times wider than those in Figure 17, we expect the asymptotic linear relationship in Figure 17b to be $\sqrt{5}$ times lower than the measured linear relationship in Figure 10b. This provides the predicted asymptote indicated by a dot-dash line in Figure 17b. In




Figure 18: Variable noise experiment: These images show variation in SNR_e within smoothed strain images, using displacement measurements from the same data as Figure 11, produced using RF-windows of 1 A-line by 151 RF-samples, sampled every 50 RF-samples axially and 1 A-line laterally. The three original strain images each had an average SNR_e of 15.0, so the scale in SNR_e images ranges from black (0) to white (28.0). The smoothing was (a) uniform, or non-uniform with Q based on (b) correlation coefficient, or (c) phase residuals.

view of the experimental results, it seems likely that this is close to the slope of the true asymptote.

We have validated our analysis regarding the role of large RF-windows in strain imaging, *i.e.*, that the main effect on strain images is similar to the application of a moving-average filter. Large RF-windows make our estimates of strain estimation precision less accurate than for non-overlapping RF-windows. However, SNR_e is nevertheless related to the predicted performance by a relationship that is approximately linear, so non-uniform smoothing may still be successful to a limited extent. This is demonstrated in Figure 18, showing images of the variation in SNR_e after the smoothing of displacement data from overlapping RF-windows, processing the same ultrasound data as in Figure 11. Uniform smoothing produces a highly non-uniform level of SNR_e . Non-uniform smoothing with Q based on correlation coefficients does only slightly better. Greater uniformity in SNR_e is achieved using Q based on phase residuals, although this is still less successful than the corresponding result for small RF-windows in Figure 11.

5 Discussion

NPR is a good smoothing method for improving the properties of strain images. In the absence of smoothing, strain images based on displacement measurements from small RF-windows are very noisy. We demonstrated a well-known fact: strain estimates can be made more accurate by smoothing displacement data with a linear filter. However, strain image quality varies significantly between different types of linear filter: in order to determine which filter is most efficient, different methods can be compared across settings that achieve the same resolution.

Subjectively, in our comparison between NPR and PLLSR, the resulting images differ in two main respects. NPR strain images are free from high-frequency noise, the presence of which is potentially misleading in PLLSR images, because it may erroneously be interpreted as high-resolution information. NPR is also better than PLLSR at avoiding distorting the appearance of image features wherever the strain field is inhomogeneous. We have not devised metrics that quantify these differences. However, we have recorded performance in terms of SNR_e , which is based on the accuracy of strain data in homogeneous regions. Although this bears little relation to the two differences mentioned above, it does represent another meaningful aspect of performance. The fact that NPR also does slightly better than PLLSR in terms of SNR_e reassures us in saying that it is unambiguously a preferable smoothing method.

Plenty of other linear filters might also be preferable to PLLSR, owing in part to the smoother

shapes of their kernels. Examples include higher-order Savitzky-Golay filters (*i.e.*, polynomial fitting) [24, 33] and suitably-designed low-pass filters [38]. We have not compared these directly with NPR, because it would be difficult to devise suitable tests of inhomogeneous strain fields with which to establish conclusively the preferable method for uniform smoothing. From a Bayesian theoretical perspective, there are many potential justifications for NPR, such as the argument that it can be used to find the optimal polynomial fit to data where the polynomial degree is unknown [17, 46]. We cannot make statements regarding absolute optimality for strain image smoothing, but we can at least conclude that NPR is a fairly good method. The more straightforward reason for our preference is the versatility of NPR for practical applications. It achieves uniform smoothing with greater computational efficiency than even short FIR filters. We showed both analytically and experimentally how the resolution and accuracy of NPR are determined by the smoothing strength, balanced against the data weight. The absolute precision of strain estimates can be predicted when indications of displacement estimation precision are available.

The versatility of NPR is all the more striking when it is applied to non-uniform smoothing. The approach that we demonstrated tends to produce uniform accuracy throughout the image, blurring where necessary when the data weights are modified according to expected precision. In this approach, the smoothing strength is a very useful parameter, allowing simple, continuous control over the level of smoothing. It means that the smoothing is varied both within individual images as well as between different images.

Given past developments in the strain imaging literature, it is important that we also consider the effect when smoothing is applied to displacement data from large overlapping RF-windows. In fact, NPR seems to achieve similar — perhaps even better — results when applied to displacement data from smaller RF-windows, since large RF-windows behave as filters, distorting NPR’s equivalent kernel. NPR offers greater advantages without this confounding effect, because alone NPR already has good properties. Strain images computed by PLLSR, on the other hand, are improved by the extra filtering from large RF-windows, because PLLSR applied in isolation is a very suboptimal filter.

Our NPR results indicate that there may be no benefit at all from using large overlapping RF-windows. The potential risk with small RF-windows is that substantial outlier errors could arise, owing to gross ambiguity in finding correct displacements among multiple local optima [2, 47]. However, it seems likely that gross errors can be avoided by applying continuity assumptions in the displacement estimation algorithm [6, 21, 40, 49]. In this case, the combination of NPR with small RF-windows may offer particularly high strain image quality compared to many alternative strain estimation methods, while incurring a comparatively low computational cost. This has only been demonstrated to a limited extent by the present investigation, so it remains an avenue for further investigation.

Future work will focus on applying NPR for smoothing displacement data from small RF-windows for practical purposes. In that regard, we should go beyond uniform smoothing, which fails to make the most efficient use of displacement data. However, we are interested in applying NPR within a real-time interface that displays persisted strain images during freehand two- and three-dimensional ultrasound scanning. A fully non-uniform smoothing method is unsuitable for this application, because artefacts are introduced by the process of averaging multiple images at different resolutions. A compromise is required to enable optimal use of data on a local scale, without dramatically changing resolution between different image regions. For example, variable data weights could be normalised in local neighbourhoods to ensure that resolution remains substantially constant throughout all of the strain data, while allowing NPR to preferentially follow the best displacement measurements in every locality. In any event, it would be worthwhile to investigate strategies for maximising the accuracy of strain images without effectively changing the resolution.

Appendices

A Computational issues

We want to integrate NPR within a real-time strain imaging system, so computation time is important. The time to solve Equation 3 depends on how NPR is used — *i.e.*, how often and according to what system are the data weights and the smoothing strength changed — because different solution methods apply. Uniform smoothing can be remarkably cheap, and the implementation is simple. Non-uniform smoothing, on the other hand, is potentially cheap, but that depends on choosing the right solution method with suitable optimisation.

Consider a grid of $N_x \times N_y$ displacement estimates. $\mathbf{W} + \mathbf{M}^T \mathbf{M}$ is a $N_x N_y \times N_x N_y$ symmetric positive definite matrix. Equation 3 can be solved quickly, because the matrix is extremely sparse. Traditional methods for solving Equation 3 are either direct [12], based on Cholesky factorisation, or iterative [42], based on successively modifying a trial solution, incrementally reducing and eventually minimising the cost expressed in Equation 2.

Rather than focusing on any particular solution method, we outline several techniques that are often useful. This is intended as introductory advice for anyone implementing NPR within a real-time strain imaging system. We provide an indication of computational efficiency in terms of operation counts, although this is an incomplete picture, since memory access may also be a significant overhead. A further indication is provided by example timings for a 128×115 grid smoothed on a PC with a 3 GHz processor.

A.1 Direct methods

The computation time depends on the extent to which non-zero elements in $\mathbf{W} + \mathbf{M}^T \mathbf{M}$ are clustered around the leading diagonal. Ideally, $\mathbf{W} + \mathbf{M}^T \mathbf{M}$ is a low-order “band-matrix” [12]. Direct methods offer an extremely fast route to the solution if the non-zero elements are tightly clustered. To be precise, the speed depends on the number of columns, b_i , between the left-most non-zero element in row i and the diagonal. We denote the mean and mean squared values across all rows by \bar{b} and \bar{b}^2 respectively. In these terms, matrix factorisation requires

$$\{N_x N_y (\bar{b}^2 + \bar{b}) \text{ multiplications, } N_x N_y \bar{b} \text{ divisions and } \frac{1}{2} N_x N_y (\bar{b}^2 + \bar{b}) \text{ subtractions}\}.$$

The overhead for substitution is much lower, coming out at

$$\{2N_x N_y \bar{b} \text{ multiplications, } N_x N_y \text{ divisions and } 2N_x N_y \bar{b} \text{ subtractions}\}.$$

For direct methods, the two-dimensional grid of displacement measurements must be mapped onto a one-dimensional scan index. Three examples are shown in Figure 19. The scan index combined with the grid dimensions determines \bar{b} and \bar{b}^2 . This has no influence on the final solution, but impacts significantly on the computation time. Fractal-based space-filling curves such as pseudo-Hilbert scans have intuitive appeal, so it is worth noting that their performance is poor, the overhead being substantially higher than for a raster scan. Space-filling curves mostly preserve proximity when mapping from two dimensions down to one, but this is achieved with a penalty of introducing some outliers, where neighbouring points on the original grid are very far apart in the scan index. Band-matrix manipulations are sensitive to outliers, hence the net increase in computation time.

Nevertheless, computation time can be reduced by suitable scan indexing. For a raster scan, if the smaller dimension is N_x , the overhead can be estimated from $\bar{b} \simeq 2N_x$ and $\bar{b}^2 \simeq 4N_x^2$. A diagonal scan is more efficient, especially if $N_y \simeq N_x$, in which case $\bar{b} \simeq \frac{4}{3}N_x$ and $\bar{b}^2 \simeq 2N_x^2$. The direct method with a diagonal scan applied to our example requires 2.35 s for factorisation and 39.7 ms for substitution. The delay associated with factorisation may be acceptable in the context of off-line processing, but it is unsuitable for real-time strain imaging. On the other hand, there is no need to recalculate factorisations between one frame and the next if uniform smoothing is applied, in which case the only significant overhead is the 39.7 ms for substitution. This approach might be acceptable for real-time strain imaging, depending on the available processing hardware and the required frame rate. More efficient methods are nonetheless desirable.

1	2	3	4	5
6	7	8	9	10
11	12	13	14	15
16	17	18	19	20

(a)

1	2	4	7	11
3	5	8	12	15
6	9	13	16	18
10	14	17	19	20

(b)

6	7	10	11	13
5	8	9	12	14
4	3	17	16	15
1	2	18	19	20

(c)

Figure 19: Scan indices. Three methods are demonstrated on a small grid. (a) Raster scan. (b) Diagonal scan. (c) Pseudo-Hilbert scan.

Direct methods are faster at solving one-dimensional systems, because $\bar{b} \simeq 2$ and $\bar{b}^2 \simeq 4$. A pseudo-1D method can be applied for uniform smoothing over multiple dimensions, by applying uniform smoothing over each direction independently. The overall frequency response (*cf.*, Equation 5) is

$$H(\omega_x, \omega_y) = \frac{1}{1 + \frac{r_x}{w} \omega_x^4 + \frac{r_y}{w} \omega_y^4 + \frac{r_x r_y}{w^2} \omega_x^4 \omega_y^4}. \quad (8)$$

The overhead for factorisation is

{ $6(N_x + N_y) - 20$ multiplications, $2(N_x + N_y) - 1$ divisions and $3(N_x + N_y) - 10$ subtractions}, and the overhead for substitution is

{ $4N_x N_y - 6(N_x + N_y)$ multiplications, $2N_x N_y$ divisions and $4N_x N_y - 6(N_x + N_y)$ subtractions}. For our example this translates into times of 10 μ s for factorisation and 1.1 ms for substitution.

Differences between strain images after uniform smoothing by the pseudo-1D method and the full 2D method are barely perceptible. However, the pseudo-1D method restricts our freedom to modify the roughness penalty in Equation 1, because it cannot respond to costs on derivatives of the form $\partial^{n+m}/\partial x^n \partial y^m$. Furthermore, the pseudo-1D method is inapplicable to non-uniform smoothing.

A.2 Iterative methods

Iterative methods are important for efficient NPR with non-uniform smoothing. Every frame has a unique matrix, $\mathbf{W} + \mathbf{M}^T \mathbf{M}$, which cannot be factorised quickly enough for real-time strain imaging. Instead, we can apply the point Gauss-Seidel iterative method. This is guaranteed to converge, and is faster than the related Jacobean method [42]. The cost, C_B , is minimised with respect to one displacement datum at a time, so an individual calculation at iteration k modifies displacement $\hat{\mathbf{d}}_i$ as follows:

$$\hat{\mathbf{d}}_i^{(k+1)} = \frac{w_i \hat{\mathbf{d}}_i - \sum_{j < i} S_{i,j} \hat{\mathbf{d}}_j^{(k+1)} - \sum_{m > i} S_{i,m} \hat{\mathbf{d}}_m^{(k)}}{S_{i,m} + w_i}, \quad (9)$$

where \mathbf{S} represents $\mathbf{M}^T \mathbf{M}$. This is a quick calculation, because $S_{i,j}$ is only non-zero for a small set of points around i (8 for the roughness penalty in Equation 1), and it does not matter how they are indexed. Slightly fewer calculations are needed for points at image edges, which have fewer neighbours, but to a very close approximation the per-iteration overhead is

$$\{8N_x N_y \text{ multiplications, } N_x N_y \text{ divisions and } 8N_x N_y \text{ subtractions}\},$$

which amounts in our example to 0.48 ms per iteration.

The overall computation time depends on how many iterations are required, which depends in turn on how the first iteration is initialised. Most simply, we can start with the raw displacement measurements, *i.e.*, $\hat{\mathbf{d}}_0 = \hat{\mathbf{d}}$. For a grid of 128×115 displacement data, strain images typically converge on the final strain image within 20–200 iterations for practical purposes, varying with the level of smoothing — more than X iterations will be required, if X is the length of the significant portion of NPR’s equivalent kernel. If the sampling rate is increased, the time required for convergence with a point iterative method varies with the square of the number of samples, because each iteration involves more calculations, and more iterations are required.

For some applications the point Gauss-Seidel iterative method may offer adequate speed. In other cases, some strain images may fail to converge in acceptably few iterations, but this might not matter, because pre-convergence strain images have reasonable appearance, albeit with less smoothing than required. However, for other applications convergence may be required, and this iterative method may be too slow. In many of these cases, the requirements may still be fulfilled by hybrid methods, where greater care is taken to achieve a high degree of optimisation.

A.3 Hybrid methods

Hybrid methods involve a combination of different processing steps, successively modifying a grid of data, starting with the original displacement measurements, so as to find an accurate approximation of the ideal NPR solution. Hybrid methods may be useful in many schemes with non-uniform smoothing to achieve acceptable speed. The following descriptions serve as examples of how different solution methods can be combined.

One possibility is to apply the direct method at a scale smaller than the entire grid, but larger than individual points. This is a block-level variation on the Gauss-Seidel method, in which C_B is minimised with respect to blocks of displacement data, holding the rest of the displacements constant. This can be repeated for overlapping blocks throughout the image, and iterated. The computational overhead consists of factorising sub-matrices once per frame (for which \bar{b} and \bar{b}^2 are small), and performing substitution in each block once per iteration. Fewer iterations are required compared to the point Gauss-Seidel method, so this may be a good approach for applications with heavy non-uniform smoothing.

Another approach is to apply the pseudo-1D direct method for uniform smoothing to improve the initialisation of point Gauss-Seidel iterations. If the initial grid is closer to the correct solution, fewer iterations are required. The pseudo-1D approach can also be used to apply uniform smoothing within blocks of data that are smaller than the entire grid. This could be used to approximate non-uniform smoothing, using the average precision indicator within each block. The pseudo-1D method is a quick way to apply coarse smoothing; afterwards, point iterations refine the fine-scale weighting of different displacement measurements.

B Matrices

Here we give details regarding the matrices required in the finite-difference version of the roughness penalty (*cf.*, Equation 2). Displacement data on the regression surface are expressed as a vector of length $N_x N_y$,

$$\mathbf{d}^T = \left[\mathbf{d}_{(1,1)} \quad \mathbf{d}_{(1,2)} \quad \dots \quad \mathbf{d}_{(1,N_y)} \quad \mathbf{d}_{(2,1)} \quad \dots \quad \mathbf{d}_{(N_x,N_y)} \right].$$

\mathbf{M} is a $2(N_x N_y - N_x - N_y) \times N_x N_y$ matrix defined such that the vector $\mathbf{M}\mathbf{d}$ lists every second difference throughout the grid of displacement data.

$$\mathbf{M}\mathbf{d} = \begin{bmatrix} \sqrt{r_x}(2\mathbf{d}_{(2,1)} - \mathbf{d}_{(3,1)} - \mathbf{d}_{(1,1)}) \\ \vdots \\ \sqrt{r_x}(2\mathbf{d}_{(N_x-1,N_y)} - \mathbf{d}_{(N_x,N_y)} - \mathbf{d}_{(N_x-2,N_y)}) \\ \sqrt{r_y}(2\mathbf{d}_{(1,3)} - \mathbf{d}_{(1,4)} - \mathbf{d}_{(1,2)}) \\ \vdots \\ \sqrt{r_y}(2\mathbf{d}_{(N_x,N_y-1)} - \mathbf{d}_{(N_x,N_y)} - \mathbf{d}_{(N_x,N_y-2)}) \end{bmatrix}$$

The most direct route to the optimal regression surface is by solving Equation 3, which represents a set of $N_x N_y$ simultaneous equations derived by differentiating the cost function with respect to every displacement datum. $\mathbf{M}^T \mathbf{M}$ is an important part of the simultaneous equations.

Owing to its sparsity, assuming a diagonal weighting matrix, simultaneous equations associated with coordinates (x, y) , away from the edge of the regression surface, have the form

$$\begin{aligned} \left[w_{(x,y)} + 6(r_x + r_y) \right] \hat{d}_{(x,y)} + r_x \left[\hat{d}_{(x-2,y)} - 4\hat{d}_{(x-1,y)} - 4\hat{d}_{(x+1,y)} + \hat{d}_{(x+2,y)} \right] + \\ r_y \left[\hat{d}_{(x,y-2)} - 4\hat{d}_{(x,y-1)} - 4\hat{d}_{(x,y+1)} + \hat{d}_{(x,y+2)} \right] = w_{(x,y)} \hat{d}_{(x,y)}. \end{aligned}$$

C Diagonal and lateral resolution

This appendix outlines a rationale for estimating the lateral resolution of NPR based on measurements of axial and diagonal resolution. Most of the tests in Section 4 are based on changing a single setting, r , so the axial and lateral smoothing strengths are linked by a relationship of the form, $r = \alpha_x r_x = \alpha_y r_y$, where α_x and α_y are pre-set scale-factors. As in Section 3.1, we assume that resolution in an arbitrary direction, j , is related to the frequency response,

$$H(\omega_j) = 1 / (1 + r_j \omega_j^4), \quad (10)$$

so the resolving limit is $\gamma r_j^{1/4}$, where γ is a constant that depends on the resolution criterion.

In principle, this means that the resolution in every direction can be found by measuring resolution in a single direction, and calculating multiplying factors based on the relative values of r_j . In practice, we take an experimental approach, because scale factors associated with the different sampling rates make it difficult to establish the absolute values of r_x and r_y based on the source code.

Our measurements show that the resolving limit is 1.13 times larger in the diagonal (xy) direction than in the axial direction for the settings used in Section 4.1. We can find the lateral resolution by relating the absolute smoothing strength in that direction, r_{xy} , to the absolute smoothing strength in the lateral direction, r_x . Frequency components associated with diagonally oriented inhomogeneities have lateral and axial components of $\omega_x = \omega_y = \omega_{xy}/\sqrt{2}$. Substituting this into Equation 5 gives

$$H(\omega_{xy}) = 1 / \left(1 + \frac{1}{4}(r_x + r_y)\omega_{xy}^4 \right),$$

which is similar to Equation 10, so $r_{xy} = \frac{1}{4}(r_x + r_y)$. A prediction of the ratio between lateral and axial resolution ($r_x^{1/4}/r_y^{1/4}$) in terms of the measured ratio between diagonal and axial resolution ($r_{xy}^{1/4}/r_y^{1/4}$) is found by rearrangement. On this basis, the settings used for the experiment in Section 4.1 apparently result in lateral resolution being 1.53 times coarser than axial resolution.

References

- [1] S. K. Alam, J. Ophir, and E. E. Konofagou. An adaptive strain estimator for elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 45(2):461–472, March 1998.
- [2] E. W. Barankin. Locally best unbiased estimates. *Annals of Mathematical Statistics*, 20(4):477–501, 1947.
- [3] R. G. Barr. Clinical applications of a real time elastography technique in breast imaging. In *Ultrasonic Measurement and Imaging of Tissue Elasticity*, page 112, Snowbird, Utah, October 2006.
- [4] I. Céspedes, Y. Huang, J. Ophir, and S. Spratt. Methods for estimation of subsample time delays of digitized echo signals. *Ultrasonic Imaging*, 17:142–171, 1995.

- [5] I. Céspedes and J. Ophir. Reduction of image noise in elastography. *Ultrasonic Imaging*, 15:89–102, 1993.
- [6] L. Chen, G. M. Treece, J. E. Lindop, A. H. Gee, and R. W. Prager. A quality-guided displacement tracking algorithm for ultrasonic elasticity imaging. Technical Report CUED/F-INFENG/TR 593, Cambridge University Department of Engineering, January 2008.
- [7] C. L. de Korte, E. I. Céspedes, A. F. W. van der Steen, G. Pasterkamp, and N. Bom. Intravascular ultrasound elastography: assessment and imaging of elastic properties of diseased arteries and vulnerable plaque. *European Journal of Ultrasound*, 7:219–224, 1998.
- [8] C. L. de Korte, G. Pasterkamp, A. F. W. van der Steen, H. A. Woutman, and N. Bom. Characterization of plaque components with intravascular ultrasound elastography in human femoral and coronary arteries *in vitro*. *Circulation*, 102:617–623, 2000.
- [9] J. Duchon. Splines minimizing rotation-invariant semi-norms in Sobolev spaces. In *Constructive Theory of Functions of Several Variables*, pages 85–100, 1976.
- [10] 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.
- [11] 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.
- [12] A. George and J. W. Liu. *Computer solution of large sparse positive definite systems*. Prentice-Hall, 1981.
- [13] P. J. Green and B. W. Silverman. *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*. CRC Press, 2004.
- [14] 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.
- [15] F. Kallel and J. Ophir. A least-squares strain estimator for elastography. *Ultrasonic Imaging*, 19:195–208, 1997.
- [16] K. Kaluzynski, X. Chen, S. Y. Emelianov, S. R. Skovoroda, and M. O’Donnell. Strain rate imaging using two-dimensional speckle tracking. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 48(4):1111–1123, July 2001.
- [17] G. S. Kimeldorf and G. Wahba. A correspondence between Bayesian estimation on stochastic processes and smoothing by splines. *The Annals of Mathematical Statistics*, 41(2):495–502, April 1970.
- [18] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Dynamic resolution selection in ultrasonic strain imaging. *Ultrasound in Medicine and Biology*, 2007. *In press*.
- [19] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Estimation of displacement location for enhanced strain imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 54(9):1751–1771, September 2007.
- [20] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. An intelligent interface for freehand strain imaging. Technical Report CUED/F-INFENG/TR 578, Cambridge University Department of Engineering, May 2007.
- [21] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Phase-based ultrasonic deformation estimation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2007. *In press*.

- [22] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. The general properties including accuracy and resolution of linear-filtering methods for strain estimation. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 2008. *In press*.
- [23] M. A. Lubinski, S. Y. Emelianov, and M. O'Donnell. Speckle tracking methods for ultrasonic elasticity imaging using short-time correlation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 46(1):82–96, January 1999.
- [24] J. Luo, J. Bai, P. He, and K. Ying. Axial strain calculation using a low-pass digital differentiator in ultrasound elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51(9):1119–1127, September 2004.
- [25] R. L. Maurice and M. Bertrand. Lagrangian speckle model and tissue-motion estimation — theory. *IEEE Transactions on Medical Imaging*, 18(7):593–603, July 1999.
- [26] M. O'Donnell, A. R. Skovoroda, B. M. Shapo, and S. Y. Emelianov. Internal displacement and strain imaging using ultrasonic speckle tracking. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 41:314–325, May 1994.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] G. F. Pinton, J. J. Dahl, and G. E. Trahey. Rapid tracking of small displacements with ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 53(6):1103–1117, June 2006.
- [31] 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.
- [32] C. H. Reinsch. Smoothing by spline functions. *Numerische Mathematik*, 10:177–183, 1967.
- [33] A. Savitzky and M. J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, July 1964.
- [34] B. W. Silverman. Spline smoothing: the equivalent variable kernel method. *The Annals of Statistics*, 12(3):898–916, 1984.
- [35] S. Srinivasan, F. Kallel, R. Souchon, and J. Ophir. Analysis of an adaptive strain estimation technique in elastography. *Ultrasonic Imaging*, 24:109–118, 2002.
- [36] S. Srinivasan, J. Ophir, and S. K. Alam. Elastographic imaging using staggered strain estimates. *Ultrasonic Imaging*, 24:229–245, 2002.
- [37] C. Sumi. Fine elasticity imaging utilizing the iterative RF-echo phase matching method. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 46(1):158–166, 1999.
- [38] C. Sumi, A. Suzuki, and K. Nakayama. Estimation of shear modulus distribution in soft tissue from strain distribution. *IEEE Transactions on Biomedical Engineering*, 42(2):193–202, 1995.

- [39] W. E. Svensson and D. Amiras. Ultrasound elasticity imaging. *Breast Cancer Online*, 9(6):1–7, 2006.
- [40] 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.
- [41] G. M. Treece, J. E. Lindop, A. H. Gee, and R. W. Prager. Freehand ultrasound elastography with a 3D probe. *Ultrasound in Medicine and Biology*, 2007. *In press*.
- [42] R. S. Varga. *Matrix Iterative Analysis*. Springer, 2000.
- [43] T. Varghese and J. Ophir. A theoretical framework for performance characterization of elastography: the strain filter. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(1):164–172, January 1997.
- [44] F. Viola and W. F. Walker. A comparison of the performance of time-delay estimators in medical ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 50(4):392–401, April 2003.
- [45] M. Vogt and H. Ermert. Development and evaluation of a high-frequency ultrasound-based system for *in vivo* strain imaging of the skin. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 52(3):375–385, March 2005.
- [46] G. Wahba. Improper priors, spline smoothing and the problem of guarding against model errors in regression. *Journal of the Royal Statistical Society, Series B*, 40(3):364–372, 1978.
- [47] A. J. Weiss and E. Weinstein. Fundamental limitations in passive time delay estimation — part i: narrow-band systems. *IEEE Transactions on Acoustics, Speech and Signal Processing*, ASSP-31(2):472–485, 1983.
- [48] E. Whittaker. On a new method of graduation. *Proceedings of the Edinburgh Mathematical Society*, 41:63–75, 1923.
- [49] Y. Zhu and T. J. Hall. A modified block matching method for real-time freehand strain imaging. *Ultrasonic Imaging*, 24:161–176, 2002.