

 Open access • Journal Article • DOI:10.1109/TUFFC.2008.620

Phase-based ultrasonic deformation estimation — Source link

[JE Lindop](#), [Graham M. Treece](#), [Andrew H. Gee](#), [Richard W. Prager](#)

Institutions: [University of Cambridge](#)

Published on: 15 Feb 2008 - [IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control \(IEEE\)](#)

Topics: [Deformation \(meteorology\)](#) and [Algorithm design](#)

Related papers:

- [A time-efficient and accurate strain estimation concept for ultrasonic elastography using iterative phase zero estimation](#)
- [Elastography: A Quantitative Method for Imaging the Elasticity of Biological Tissues](#)
- [Estimation of Displacement Location for Enhanced Strain Imaging](#)
- [Internal displacement and strain imaging using ultrasonic speckle tracking](#)
- [Effect of skin impedance on image quality and variability in electrical impedance tomography: a model study](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/phase-based-ultrasonic-deformation-estimation-1807ip5693>

**PHASE-BASED ULTRASONIC
DEFORMATION ESTIMATION**

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

CUED/F-INFENG/TR 555

May 2006

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

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

PHASE-BASED ULTRASONIC DEFORMATION 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

Deformation estimation is the foundation of emerging techniques that image the mechanical properties of soft tissues. We present theoretical analysis and experimental results for an investigation of phase-based ultrasonic deformation estimation. Numerous phase-based algorithm variants are tested quantitatively on simulated RF data from uniform scatterer fields, subject to a range of uniform strain deformations. Particular attention is paid to a new algorithm using Weighted Phase Separation, the performance of which is demonstrated in application to *in vivo* freehand strain imaging. The results support the theory that underlies the new algorithm, and more generally highlight the factors that should be considered in the design of high performance deformation estimators for practical applications.

1 Introduction

Ultrasonic imaging of tissue mechanical properties is a growing field in which there are many competing approaches. The majority of schemes require high accuracy estimation of the small deformations that occur between successive frames in an ultrasound scan, although a smaller set of alternatives work in conjunction with conventional Doppler motion estimation [7, 11]. Most systems employ a conventional two-dimensional ultrasound scanner, and the aim of deformation estimation is to produce an array of one- or two-dimensional displacement estimates, which may be thought of as noisy samples from the displacement field over a fine grid of locations throughout each frame. The recorded displacement estimates are sometimes displayed directly as displacement images, but it is more common to produce strain images by taking spatial derivatives of the displacement. There also exist more elaborate analyses that aim at displaying quantitative estimates of the mechanical properties of the underlying tissue. This can be tackled, for example, by solving the inverse problem for the elasticity field when deformation data has been recorded under a known static load [6]. Alternatively, elastic moduli can be estimated from a record of wavefront propagation when low frequency shear waves are transmitted through the region of interest [3].

This paper is not concerned with methods for analysing deformation data: whichever techniques become clinically important, the quality of the results from high level analysis will depend to a large degree on the accuracy of the underlying deformation estimation. Moreover, it is desirable that this estimation be performed at low computational cost. We consider the task of estimating the deformation between a pair of RF ultrasound frames acquired pre- and post-deformation. “Deformation estimation” refers to the special case of displacement estimation when displacement is a continuously varying function of location. At present, it is normal to produce each displacement estimate by positioning a window over a small section of data in the pre-deformation frame and identifying the closest matching window in the post-deformation frame. The displacement estimate is the difference between the pre- and post-deformation window positions. The task of window matching entails adjusting the post-deformation window position in order to find the optimum in a measure of signal similarity. The most commonly cited measure is the correlation coefficient, although near-identical performance has been reported from techniques employing alternative measures such as the sum of squared differences [13] and the phase of the complex cross-correlation function [4, 10]. The estimation procedure is repeated throughout a grid of locations, as above, until the displacement field has been adequately sampled.

The window matching approach to deformation estimation is sometimes problematic: pre- and post-deformation windows often match poorly, because deformation may not be negligible on the scale of the individual windows. Many techniques have been developed whereby the post-deformation signals are warped to increase the correlation between pre- and post-deformation windows. These are often termed “adaptive” strain estimators. The simplest adaptive method is to apply a uniform stretch to the post-deformation signal, aiming to reverse part of the signal transformation that has actually taken place [1]. Deformation data from adaptive strain estimators are measurably less noisy than standard displacement estimation, but the improvement is accompanied by a considerable increase in computational cost.

We have recently noted, however, that traditional window matching approaches can be enhanced. Since finite length windows are required to produce displacement estimates with low noise, the accuracy of the data can be improved by estimating the location at which the displacement estimate is valid. Thus, each deformation datum consists of an estimate of the displacement location in addition to the displacement itself. In past systems, it was implicitly assumed that the location is the window centre, which results in an “amplitude modulation” artefact with the RF signal amplitude modulating the strain image. For this reason, our location estimation technique is termed Amplitude Modulation Correction (AMC). We have demonstrated that AMC yields better performance at lower computational cost than adaptive strain estimation [9].

AMC can be implemented particularly easily in conjunction with phase-based displacement estimators. In the light of this result, the paper in hand presents further analysis of phase-based deformation estimation. Both theoretical and empirical methods are employed for the derivation and assessment of new phase-based deformation estimators. In particular, we introduce a new family of highly versatile algorithms referred to as Weighted Phase Separation (WPS¹). It is shown that the WPS framework can reproduce the performance of conventional phase-based methods, but WPS can also be adapted when different properties are required. The analysis in this report considers deformation estimation in the axial direction only, reflecting the fact that axial deformation is usually most important: the accuracy is superior because RF ultrasound signals have far lower lateral and elevational bandwidth, and in many elasticity imaging schemes the largest deformations actually occur axially. The WPS algorithm is validated by demonstration with *in vivo* freehand strain imaging. Furthermore, it has previously been shown that phase-based methods can be adapted for displacement estimation in other directions [4], so it is likely that the findings of this study can be developed in the future for multi-dimensional deformation estimation.

2 Theory

This section presents theoretical analysis of phase-based deformation estimators. It begins by reviewing Amplitude Modulation Correction (AMC): the accuracy of the deformation data is increased by estimating the location at which each displacement estimate is valid (for a full introduction see [9]). We consider the case when AMC is applied to conventional algorithms, where matching post-deformation windows are found by locating the zero crossings in the phase of the complex cross-correlation function. Later we introduce Weighted Phase Separation (WPS), which is partly motivated by the simplicity of adding AMC. In general, WPS is a framework in which the expected importance of different signal portions can be incorporated in the deformation estimator by an arbitrary selection of weightings, so as to maximise accuracy. The weighting strategy can be adjusted to reflect any theoretical or empirical experience. We present analysis to justify simple weighting strategies linked to signal amplitude and signal phase, which are employed in later experiments for an initial validation of the WPS concept.

¹WPS is the subject of UK patent application GB 0610172.9.

Throughout this paper we consider a signal model that offers a high level of generality.

$$a_1(t) = a_{r1}(t) + ja_{i1}(t) = s_1(t)e^{j\phi_{s1}(t)} \quad (1a)$$

$$= f(t)e^{j\phi(t)} + n_1(t)e^{j\phi_{n1}(t)} \quad (1b)$$

$$a_2(t+d(t)) = a_{r2}(t+d(t)) + ja_{i2}(t+d(t)) = s_2(t+d(t))e^{j\phi_{s2}(t+d(t))} \quad (1c)$$

$$= f(t)e^{j\phi(t)} + n_2(t)e^{j\phi_{n2}(t)} \quad (1d)$$

$a_1(t)$ and $a_2(t)$ are analytic representations of pre- and post-deformation RF signals with real and imaginary parts as indicated. They can alternatively be expressed in phasor notation with envelopes $s_1(t)$ and $s_2(t)$ and phases $\phi_{s1}(t)$ and $\phi_{s2}(t)$. Equations 1b and 1d encapsulate the relationship that is exploited for deformation estimation. Both signals contain a common signal, $f(t)e^{j\phi(t)}$, which undergoes an arbitrary stretch, $t \rightarrow t+d(t)$, caused by the movement of scatterers in the underlying tissue. To a varying extent the common signal is masked by noise signals $n_1(t)$ and $n_2(t)$ which encompass all signal components not pertaining to the common signal. These include electrical noise, changes to scatterer interference patterns and decorrelation due to off-axis movement.

2.1 Location estimation with conventional algorithms

For each window, AMC entails estimating the location at which the displacement estimate is valid. This is necessary unless the strain is zero, because a range of displacements are present in any finite length window. One might imagine that windows could be made very small in order to avoid this ambiguity, effectively measuring point displacements, but the size of errors in the displacement estimates is inversely related to the window length, so the overall level of estimation noise would increase.

In general, a suitable implementation of AMC for any particular algorithm is found by analysing the properties of the displacement estimator to derive the following form of approximation.

$$\hat{d}_n \simeq \frac{\sum_{t=n\Delta t}^{n\Delta t+T} W(t)d(t)}{\sum_{t=n\Delta t}^{n\Delta t+T} W(t)} \quad (2)$$

The displacement estimate for window n is \hat{d}_n ; t is axial distance measured in number of samples ($t = 0$ is the surface of the probe); Δt is the shift in starting position between consecutive analysis windows moving down the pre-deformation A-line; T is the window length; $d(t)$ is the actual displacement of tissue at pre-deformation position t ; and $W(t)$ is approximately the weighting of data at that position. The distribution of weightings over the length of the window modulates the location at which \hat{d}_n is most likely to correspond to the underlying tissue displacement, *i.e.*, amplitude modulation. Note that in general, however, the weightings may depend not only on amplitude but on all signal properties. It is often reasonable to assume that the strain, s , is uniform over the length of the window, *i.e.*, $d(t) = \alpha + st$. In this case, we define τ_n as the location at which \hat{d}_n is most likely to correspond to the actual common signal displacement, *i.e.*, $\hat{d}_n = \alpha + s\tau_n$. An estimate, $\hat{\tau}_n$, is produced by substituting these expressions into Equation 2.

$$\hat{\tau}_n = \frac{\sum_{t=n\Delta t}^{n\Delta t+T} W(t)t}{\sum_{t=n\Delta t}^{n\Delta t+T} W(t)} \quad (3)$$

The accuracy of the location estimate depends on (1) the validity of the uniform strain assumption² and (2) the accuracy of the weighted-sum approximation in Equation 2. The analysis for determining the weightings depends on the properties of the displacement estimator. Using our signal

²In some cases it is more accurate to approximate the spatial variation of $d(t)$ with a higher order polynomial. A similar procedure can be applied to estimate the polynomial coefficients from a set of weightings and displacement estimates from a group of neighbouring windows.

model from Equation 1, we will review the derivation of weightings for conventional phase-based displacement estimators [9].

At window n with trial displacement \tilde{d}_k , the cross-correlation function, $\langle a_1, a_2 \rangle$, and its phase, Φ , are as follows.

$$\langle a_1, a_2 \rangle (n\Delta t, \tilde{d}_k) = \sum_{t=n\Delta t}^{n\Delta t+T} a_1^*(t) a_2(t + \tilde{d}_k) \quad (4a)$$

$$\Phi (n\Delta t, \tilde{d}_k) = \angle \langle a_1, a_2 \rangle (n\Delta t, \tilde{d}_k) \quad (4b)$$

The displacement estimate, \hat{d}_n , is the displacement at which Φ is zero.

$$\Phi (n\Delta t, \hat{d}_n) = 0 \quad (5)$$

To simplify the following analysis we introduce $t_2(t, \tilde{d}_k)$. This expresses the pre-deformation location of the signal component with which data at t is compared under trial displacement \tilde{d}_k . The same symbol is also used in conjunction with displacement estimates, *i.e.*, $t_2(t, \hat{d}_n)$. In other words, the point in the common signal that was at location t_2 before the deformation translates to $t_2 + d(t_2)$ in the post-deformation signal, and this is the location that gets compared with t in the pre-deformation signal.

$$t_2 + d(t_2) \triangleq t + \tilde{d}_k \quad (6)$$

Consider the terms of the cross-correlation function.

$$\langle a_1, a_2 \rangle (n\Delta t, \tilde{d}_k) = \sum_{t=n\Delta t}^{n\Delta t+T} \left\{ f(t)f(t_2)e^{j(\phi(t_2)-\phi(t))} + f(t)n_2(t_2)e^{j(\phi_{n_2}(t_2)-\phi(t))} \right. \\ \left. + n_1(t)f(t_2)e^{j(\phi(t_2)-\phi_{n_1}(t))} + n_1(t)n_2(t_2)e^{j(\phi_{n_2}(t_2)-\phi_{n_1}(t))} \right\} \quad (7)$$

The terms divide into two categories. ρ_d contains terms associated with signal stretching and ρ_s contains the noise terms.

$$\langle a_1, a_2 \rangle (n\Delta t, \tilde{d}_k) = \rho_d (n\Delta t, \tilde{d}_k) + \rho_s (n\Delta t, \tilde{d}_k) \quad (8a)$$

$$\text{where } \rho_d (n\Delta t, \tilde{d}_k) = \sum_{t=n\Delta t}^{n\Delta t+T} f(t)f(t_2)e^{j(\phi(t_2)-\phi(t))} \quad (8b)$$

$$\text{and } \rho_s (n\Delta t, \tilde{d}_k) = \sum_{t=n\Delta t}^{n\Delta t+T} \left\{ f(t)n_2(t_2)e^{j(\phi_{n_2}(t_2)-\phi(t))} + n_1(t)f(t_2)e^{j(\phi(t_2)-\phi_{n_1}(t))} \right. \\ \left. + n_1(t)n_2(t_2)e^{j(\phi_{n_2}(t_2)-\phi_{n_1}(t))} \right\} \quad (8c)$$

Every term in ρ_s is a sum over the product of signals that are generally uncorrelated, so unless T is very small these tend to cancel out. Thus, unless the common signal is very weak, ρ_d is usually the major constituent of the cross-correlation function. For an insight into the mechanism of displacement estimation by cross-correlation function phase methods, we briefly consider the case when $\rho_s \ll \rho_d$, so noise terms are neglected. The phase zero condition from Equation 5 implies that the cross-correlation function has no imaginary part at the match.

$$\Im \left\{ \rho_d (n\Delta t, \hat{d}_n) \right\} = 0 \quad \Rightarrow \quad \sum_{t=n\Delta t}^{n\Delta t+T} f(t)f(t_2) \sin(\phi(t_2) - \phi(t)) = 0 \quad (9)$$

A further approximation can be made for typical window lengths and strains. The small angle approximation applies so long as $sT \ll \lambda$ (λ denotes the ultrasonic wavelength at the centre

frequency).

$$\sum_{t=n\Delta t}^{n\Delta t+T} f(t)f(t_2)(\phi(t_2) - \phi(t)) \simeq 0 \quad (10)$$

A more illuminating form is produced when we define the local mean frequency, $\bar{\omega}(t, t_2)$, and substitute from Equation 6 for $t_2 - t$ [9].

$$\bar{\omega}(t, t_2) \triangleq \frac{\phi(t_2) - \phi(t)}{t_2 - t} \quad (11a)$$

$$t_2 - t = \hat{d}_n - d(t_2) \simeq \hat{d}_n - d(t) \quad (11b)$$

This leads to an alternative expression for the relationship in Equation 10.

$$\sum_{t=n\Delta t}^{n\Delta t+T} f(t)f(t_2)\bar{\omega}(t, t_2) (\hat{d}_n - d(t)) \simeq 0 \quad (12)$$

Rearrangement yields an expression for \hat{d}_n .

$$\hat{d}_n \simeq \frac{\sum_{t=n\Delta t}^{n\Delta t+T} f(t)f(t_2)\bar{\omega}(t, t_2)d(t)}{\sum_{t=n\Delta t}^{n\Delta t+T} f(t)f(t_2)\bar{\omega}(t, t_2)} \quad (13)$$

The approximation in Equation 13 has the required form for AMC, *cf.*, Equation 2. The weightings are $W(t) = f(t)f(t_2)\bar{\omega}(t, t_2)$. These weightings can be evaluated with moderate accuracy. It is difficult to estimate $\bar{\omega}(t, t_2)$ because both pre- and post-deformation noise signals have a large impact on the recorded frequency perturbations, but a reasonable estimate is made simply by assuming that $\bar{\omega}(t, t_2)$ is equal to the centre frequency, which is constant at least on the scale of the window length. In practice, this means that $\hat{\tau}_n$ is estimated assuming $W(t) = f(t)f(t_2)$, where the signal envelope product, $s_1(t)s_2(t + \hat{d}_n)$, is taken as an estimate for $f(t)f(t_2)$.

2.2 Weighted Phase Separation

Following the analysis in Section 2.1, it seems evident that signal phase separation could be considered as a standalone displacement estimator, without recourse to the cross-correlation function. Pre- and post-deformation points must be aligned to within $\lambda/2$ to avoid phase wrapping ambiguity. When this is the case, the phase separation of the common signal is equal to the local alignment error scaled by the local frequency, as expressed in Equation 11. At alignment \tilde{d}_k , a point-wise displacement estimate $\hat{d}(t, \tilde{d}_k)$ can be evaluated by subtracting the estimated alignment error.

$$\hat{d}(t, \tilde{d}_k) = \tilde{d}_k + \frac{\hat{\phi}(t) - \hat{\phi}(t_2)}{\hat{\omega}(t, t_2)} \quad (14)$$

$\hat{\phi}(t)$ and $\hat{\phi}(t_2)$ are estimates of the common signal phase at the pre- and post-deformation points, and $\hat{\omega}(t, t_2)$ is an estimate of $\bar{\omega}(t, t_2)$. Again, in this report we will assume a constant value for $\hat{\omega}(t, t_2)$, replacing it with the nominal probe centre frequency, ω_0 . Deviations in the centre frequency introduce bias in the point-wise estimates, but this bias will be eliminated later in the analysis. As for estimating the phase separation, $\hat{\phi}(t) - \hat{\phi}(t_2)$, there may be scope for sophisticated adaptive filtering approaches for removing the noise signals, but in this investigation we simply record the phase of the overall signal.

$$\hat{\phi}(t) = \arg a_1(t) = \phi_{s1}(t) \quad (15a)$$

$$\hat{\phi}(t_2) = \arg a_2(t + \tilde{d}_k) = \phi_{s2}(t + \tilde{d}_k) \quad (15b)$$

Recall that the overall phase of an analytic signal can readily be evaluated in the range $[-\pi, +\pi]$ by taking the inverse tangent of the ratio of its real and imaginary parts.

$$\phi_s(t) = \arg a(t) = \tan^{-1} \left(\frac{a_i(t)}{a_r(t)} \right) \quad (16)$$

Assuming that the signals are aligned to within $\lambda/2$, one might envisage detecting the phase of the RF ultrasound signals and immediately applying Equation 14 to produce a point-wise displacement estimate at every sample. However, this approach would suffer from extremely high noise. Firstly, when only a single sample is used there is no chance for noise terms to cancel out, so the level of estimation noise is inevitably higher than for estimates on a coarser scale. Secondly, the alignment needs to be iteratively corrected to reduce the level of noise, since the approximation in Equation 11b is accurate only for closely aligned signals. Furthermore, the size of errors introduced by frequency perturbations is proportional to the alignment error. Intuitively it seems sensible to use each point-wise displacement estimate as a new alignment, iteratively refining the estimate. However, the alignment could actually become poorer if the point-wise estimates are noisy.

A more robust approach refines the alignment across a wider region (a window) by taking a weighted average of point-wise estimates.

$$\hat{d}(t, \tilde{d}_k) = \tilde{d}_k + \frac{\phi_{s1}(t) - \phi_{s2}(t + \tilde{d}_k)}{\omega_0} \quad (17a)$$

$$\tilde{d}_{k+1} = \frac{\sum_{t=n\Delta t}^{n\Delta t+T} W(t) \hat{d}(t, \tilde{d}_k)}{\sum_{t=n\Delta t}^{n\Delta t+T} W(t)} \quad (17b)$$

$$= \tilde{d}_k + \frac{\sum_{t=n\Delta t}^{n\Delta t+T} W(t) (\phi_{s1}(t) - \phi_{s2}(t + \tilde{d}_k))}{\omega_0 \sum_{t=n\Delta t}^{n\Delta t+T} W(t)} \quad (17c)$$

Each point-wise estimate at alignment \tilde{d}_k follows Equation 17a. The weighted sum in Equation 17c is used for iterative realignment. With each iteration the points are better aligned so the estimates are more accurate. Eventually the alignment will converge on an optimum for the window.

$$\tilde{d}_k = \hat{d}_n \quad \Rightarrow \quad \tilde{d}_{k+1} = \tilde{d}_k \quad (18)$$

In practice, iterations cease when a convergence criterion is satisfied. Thereafter it may be desirable to refine the point alignments further by proceeding with more stages of analysis using shorter windows for the weighted averaging. Alternatively, the optimal alignment, \hat{d}_n , may itself be recorded as a robust displacement estimate. This type of algorithm, called Weighted Phase Separation (WPS), is the principal subject of our investigation.

It should be noted that when convergence occurs, the error in ω_0 no longer causes bias, because the final sum of weighted phase separations is zero. Notice also that by explicitly employing weightings for displacement estimation, WPS becomes an ideal target for AMC.

The weightings, $W(t)$, can be adjusted to emphasise signal portions that are of special interest. For example, a weighting of zero is implicitly applied to data outside each window. Within each window it is simplest to use uniform weightings, $W(t) = 1$. However, if it is possible to infer the reliability of different portions of the signals, then the most reliable portions should be weighted more heavily to reduce the overall estimation error. Section 2.3 examines the factors that need to be considered in determining a weighting strategy. Section 2.3.2 also details further observations that improve the accuracy with which effective weightings are estimated when using conventional phase-based displacement estimators.

2.3 Weighting selection

Here we consider weighting selection in WPS to minimise the expected mean squared error or variance of the displacement estimates. It can be shown that the overall variance of a weighted

sum of independent estimates is minimised by choosing weightings proportional to the reciprocal of the variance for each estimate [9]. Point-wise displacement estimates in a window are not in fact independent, so ideally the level of new information provided by each point would also affect the weighting. However, we do not pursue this sophistication in the present investigation. Weightings will be chosen inversely proportional to the estimated variance.

$$W(t) = \hat{\sigma}_d^2(t, \tilde{d}_k)^{-1} \quad (19)$$

To estimate the variance we consider the four sources of error in the point-wise estimates following Equation 17a.

Tissue-signal displacement error. Displacement $d(t)$ of the common signal $f(t)e^{j\phi(t)}$ does not correspond exactly to displacement in the underlying tissue. Ultrasonic resolution is limited, so a single dominant scatterer is sometimes the primary signal source over an extended region around its actual location. The displacement in the recorded signal throughout this region is equal to the displacement of the dominant scatterer, even if the underlying tissue is subject to a high strain. Similarly, it is not possible to resolve the displacements of multiple scatterers within a single resolution cell. The effect of these phenomena in reducing the accuracy of deformation estimation certainly warrants investigation, although it will not be covered here. In effect, therefore, our analysis applies to improving the accuracy with which the *signal* displacement is estimated.

Frequency estimation error. Clearly, error in the estimation of $\bar{\omega}(t)$ introduces displacement estimation error. We will assume that the scale of these errors is fairly uniform throughout the data, however, so frequency is not considered in the present weighting strategy.

Alignment error. The location at which displacement is *actually* being estimated is t_2 . This introduces error that depends on the level of inaccuracy in the approximation $d(t_2) \simeq d(t)$ in Equation 11b. This depends in turn on the level of strain and on the accuracy of the signal alignment, thus motivating phase deweighting in Section 2.3.2.

Phase estimation error. $\phi_{s1}(t) \neq \phi(t)$ and $\phi_{s2}(t + \tilde{d}_k) \neq \phi(t_2)$ because the recorded signals are corrupted by noise. The variance of each phase estimate depends on the local ratio between common and noise signal power, demonstrated as follows. Note that this is different to the ultrasonic SNR, which only considers electrical noise.

2.3.1 Signal amplitude

We derive a simple approximation for the common signal phase estimation variance, $\sigma_\phi^2(t)$. The phase estimate follows Equation 16, for which we consider the behaviour when the common signal power is larger than the noise power. It can be shown that an uncorrelated noise signal of known power but unknown phase (with no further assumptions) when added to the common signal introduces the same variance in the real and imaginary parts of the analytic signal, although the real and imaginary errors are uncorrelated. Figure 1 shows these two signal components on an Argand diagram to illustrate the link between noise in the real and imaginary parts and noise in the phase estimate.

The noise signal contributes an error Δy to the imaginary part and Δx to the real part. The common signal power is several times greater than the noise signal power for most portions of the signal, so the phase error, $\Delta\phi$, may be estimated applying the small angle approximation.

$$\Delta\phi(t) \simeq \frac{p(t)}{f(t)} = \frac{\Delta x(t) \sin \phi(t) + \Delta y(t) \cos \phi(t)}{f(t)} \quad (20)$$




Figure 1: Phase estimation noise: at a moderate ratio of common signal power to noise signal power, the phase estimation error, $\Delta\phi$, is inversely proportional to the common signal envelope, f . Noise in the real and imaginary parts, Δx and Δy , only translates to phase estimation noise through the component perpendicular to the common signal, $p = \Delta x \sin \phi + \Delta y \cos \phi$.

We need to estimate the variance for weighting selection.

$$\sigma_{\phi}^2(t) = \text{E} \left[\Delta\phi(t)^2 \right] \quad (21a)$$

$$\simeq \text{E} \left[\frac{\Delta x(t)^2 \sin^2 \phi(t) + 2\Delta x(t)\Delta y(t) \sin \phi(t) \cos \phi(t) + \Delta y(t)^2 \cos^2 \phi(t)}{f(t)^2} \right] \quad (21b)$$

$$\simeq \frac{\text{E}[\Delta x(t)^2] \sin^2 \phi(t) + \text{E}[\Delta y(t)^2] \cos^2 \phi(t)}{\text{E} \left[f(t)^2 \right]} \quad (21c)$$

$$\simeq \frac{\sigma_n^2(t) \sin^2 \phi(t) + \sigma_n^2(t) \cos^2 \phi(t)}{\text{E} \left[f(t)^2 \right]} = \frac{\sigma_n^2(t)}{\text{E} \left[f(t)^2 \right]} \quad (21d)$$

The approximate error is taken from Equation 20, and the product of uncorrelated real and imaginary errors in Equation 21b becomes zero under the statistical expectation operator. The expected squared errors, by contrast, are equal to the noise power, which is not estimated so in Equation 21d the phase estimation variance is shown to be inversely proportional to the common signal power.

Recall from Equation 19 that the variance we require is of the point-wise displacement estimate. Inspection of Equation 17a shows that errors in the pre- and post-deformation phase estimates combine additively in the overall displacement error. Therefore, the overall variance includes the sum of both phase estimation variances, from which the reciprocal is taken in order to evaluate a weighting. Since we make no attempt to estimate the noise power it is replaced by unity in the following expressions.

$$W(t) = \left(\sigma_{\phi}^2(t) + \sigma_{\phi}^2(t_2) \right)^{-1} = \left(\frac{1}{f(t)^2} + \frac{1}{f(t_2)^2} \right)^{-1} = \frac{f(t)^2 f(t_2)^2}{f(t)^2 + f(t_2)^2} \quad (22a)$$

$$= \frac{f(t)f(t_2)}{c + c^{-1}} \quad (22b)$$

c in Equation 22b denotes the ratio of the common signal envelopes, $f(t_2)/f(t)$, which is likely to be close to unity — it is unity by definition if the alignment error is zero. Small perturbations in c are difficult to estimate, so a constant value will be assumed. Of course, the common signal envelope is not readily accessible, so for the purpose of practical weightings it is replaced with the full envelope of the recorded signal.

$$W(t, \tilde{d}_k) = s_1(t)s_2(t + \tilde{d}_k) \quad (23)$$




Figure 2: Variation of the phase contribution to weightings, $W_B(t)$, against phase separation, $\phi(t_2) - \phi(t)$, when using cross-correlation function phase.

This derivation includes several assumptions that might be avoidable, so there is scope for introducing adaptive filtering concepts to improve the estimation of the various signal components. However, it is encouraging to note that the practical weightings in Equation 23 resemble our approximation in Equation 13 for weightings in the cross-correlation function phase, since this is already known to have some practical utility.

2.3.2 Signal phase

Now we consider how phase might influence the variance on point-wise displacement estimates, and how it can be incorporated in WPS weighting strategies. The implications of signal phase variations are less obvious than signal amplitude. Therefore, we begin by considering how phase affects effective weightings with the cross-correlation function phase.

The approximation in Equation 13 indicates only amplitude and frequency contributions to the weightings. However, the small angle approximation leading to Equation 10 may not be accurate in the case of long windows. A better approximation can be made by interpreting the scaling between phase and sine value as a phase-dependent weighting. Thus with no loss of accuracy Equation 9 is rewritten in the form of Equation 24.

$$\sum_{t=n\Delta t}^{n\Delta t+T} W_A(t)W_B(t)(\phi(t_2) - \phi(t)) = 0 \quad (24a)$$

$$\text{where } W_A(t) = f(t)f(t_2) \quad (24b)$$

$$\text{and } W_B(t) = \frac{\sin(\phi(t_2) - \phi(t))}{\phi(t_2) - \phi(t)} \quad (24c)$$

W_A is the amplitude-based weighting (as before) while W_B is a new phase deweighting. Figure 2 illustrates the size of W_B for phase separations in the range $[-\pi, +\pi]$. This is of interest partly because it may yield a better implementation of AMC for cross-correlation function phase, but also because we shall consider whether a similar weighting is useful in WPS.

It is possible to implement WPS in such a way as to precisely reproduce the behaviour of cross-correlation function phase. This requires taking the weighting strategy from Equation 24 and replacing the common signal quantities with the envelope and phase of the recorded signals. However, our aim is not merely to find an alternative implementation of cross-correlation function phase, but to investigate possible improvements for more accurate estimation.

There are two arguments for deweighting large phase separations. Firstly, as the phase separation approaches $\pm\pi$ there is an increased likelihood of phase wrapping errors, where an extra quantum displacement error of $\pm\lambda$ can arise if the phase separation appears on the wrong side of the real axis due to noise. It will be appreciated that the WPS framework could be a vehicle for within-window phase unwrapping strategies, which might solve this problem, hence overcoming one of the main limitations of cross-correlation function phase. However, phase-unwrapping is not pursued in this investigation. The output of all phase arithmetic is restricted to the range $[-\pi, \pi]$ unless otherwise specified.

The second reason for deweighting large phase separations is more fundamental. We note once more that the point-wise displacement estimates become less accurate at large alignment errors. This limitation arises from the approximation in Equation 11b. Errors are more likely to be significant in the case of large phase separations indicating large alignment errors.

A rigorous probabilistic analysis of these two phenomena may be challenging, so in this investigation we restrict ourselves to heuristics as follows.

$$W_B(t) = \left| \frac{\pi - |\phi_{s2}(t + \tilde{d}_k) - \phi_{s1}(t)|}{\pi} \right|^n \quad (25)$$

We test the simple strategy of Equation 25, where n determines the severity of the deweighting. WPS variants with n in the range 0 (no phase deweighting) up to 3 (severe phase deweighting) are investigated.

3 Experimental methods

Experiments have been performed using simulated RF data frames from uniform strain fields. The strain estimation signal-to-noise ratio, SNR_e is evaluated as a measure of deformation estimation performance.

$$\text{SNR}_e = \frac{\mu_{\hat{s}}}{\sigma_{\hat{s}}} \quad (26)$$

$\mu_{\hat{s}}$ is the mean and $\sigma_{\hat{s}}$ is the standard deviation within any particular set of strain estimates. SNR_e results enable comparisons to be made of a range of deformation estimators based both on WPS and on cross-correlation function phase.

One of the effects of AMC is variation in the recorded spacing of the location estimates. The mean spacing (averaged over the image area) is therefore larger than the spacing of the windows, since a larger gap covers more of the image than does a small gap. This is not important in comparing uncorrected strain images with corrected ones where the same displacement estimates (but different location estimates) are employed, since AMC essentially both corrects the location perturbation and highlights any variability in the resolution, which goes unnoticed when AMC is not used. However, if a set of algorithms are compared, all of which are corrected with AMC, then the way that the displacement estimator changes the resolution should be taken into account to perform a fair comparison.

The spacing of windows is fixed in all the experiments, because it is known that where a strain estimate is produced by differencing a pair of displacement estimates, the simplest way to reduce strain estimation noise is by increasing the separation of the estimation locations. This reduces the extent to which any estimation noise is amplified. It is a typical example of the tradeoff between noise and resolution. The performance measure SNR_e is directly proportional to the spacing. To avoid resolution effects colouring the performance comparison of AMC-corrected deformation estimators, a preferred performance measure, SNR_e^β , is introduced for these cases.

$$\text{SNR}_e^\beta = \frac{\mu_{\hat{s}}}{\sigma_{\hat{s}}} \times \frac{\text{mean window spacing}}{\text{mean spacing of location estimates}} \quad (27)$$

3.1 Algorithms

We test two families of deformation estimators. The first is an example of cross-correlation function phase: the efficient phase zero search (EPZS) [8] has been adapted from the phase zero seeking concept of Pesavento *et al.* [12]. The second family, WPS, has been introduced in this report. Both families conform to the usual principle of placing a window over a section of pre-deformation RF ultrasound data and moving another window over the corresponding post-deformation data to produce a displacement estimate, which is assumed to apply at the centre of the window (by default) or at the location estimate (if AMC is applied). The spacing between successive windows is 2.7λ , and displacement is converted to strain by taking the difference between successive displacements

divided by the difference between their locations. The range of algorithm variants that are tested is listed in Table 1. The nature of the differences is explained as follows.

WPS and EPZS both perform short-range searches by iterative techniques, and both rely on iteration seeding so that each search begins within $\lambda/2$ of the correct alignment. Typically, each search is seeded with the displacement estimate from the preceeding window in the same A-line, but this is sometimes susceptible to error propagation. We have developed superior seeding strategies, for which the interested reader should refer to Appendix A.

EPZS is outlined in previous work [8]. In addition to the standard algorithm, we also test a variant in which the signal is selectively amplified so as to have the same envelope at all points. We call this discarded amplitude (referred to in previous work as limit log compression). Note that intermediate levels of amplitude compression could be applied to attain any level of intermediate performance between EPZS and EPZS_L. The motivation for discarded amplitude is that, evidently, this is an alternative means of mitigating amplitude modulation effects, which has different properties to AMC. On the other hand, as listed in Table 1a, discarded amplitude is still amenable to the new form of AMC. There are variants in terms of the way that AMC is applied. Tests are variously performed without AMC, with AMC (only considering the envelope, *cf.*, Section 2.1) and with the new implementation of AMC (considering both phase and envelope, *cf.*, Section 2.3.2). Note that for efficient operation the variant EPZS_A2 is implemented preferably through the WPS framework, with the correct choice of weightings so as to reproduce the correct behaviour.

The new WPS algorithms require slightly more explanation. The pre-processing is mostly unchanged from EPZS: matched FIR filters produce 5–10 MHz real and imaginary parts for the analytic signal. In the case of WPS, this is converted to arrays of phase and envelope data. Phase is detected following Equation 16, after which demodulation to the baseband is performed by subtracting $\omega_0 t$, where ω_0 is the nominal probe centre frequency. $2n\pi$ offsets are discarded, so phase values are stored in the range $[-\pi, +\pi]$.

Each iterative search by WPS begins with a rough estimate \tilde{d}_0 . Iterations similar to Equation 17c refine this until the rate of change is small (< 0.001 samples, for example). Phase is recorded at baseband, however, so the actual iteration formula in terms of baseband phase is as follows, where point-wise phase separations are expressed in the range $[\tilde{d}_k\omega_0 - \pi, \tilde{d}_k\omega_0 + \pi]$ since this leaves small risk of point-wise phase wrapping errors.

$$\tilde{d}_{k+1} = \frac{\sum_{t=n\Delta t}^{n\Delta t+T} W(t) (\phi_{s1}(t) - \phi_{s2}(t + \tilde{d}_k))}{\omega_0 \sum_{t=n\Delta t}^{n\Delta t+T} W(t)} \quad (28)$$

The WPS variants differ in terms of the weighting strategies and whether or not AMC is applied. Normally amplitude is incorporated through the envelope product, following Equation 23, although discarded amplitude variants (weightings independent of amplitude) are also tested. Phase weightings are incorporated as well, following Equation 25, where the choice of phase deweighting level, n , is indicated in Table 1b.

3.2 Simulation

Simulated RF ultrasound data was generated using Field II [5]. The simulations have 2×10^5 scatterers positioned at random according to a uniform distribution throughout a $50 \times 50 \times 6$ mm volume, with random scattering strengths distributed uniformly over the range $[0, \gamma_{\max}]$. The probe parameters model the 5-10 MHz probe of the Dynamic Imaging³ Diasus ultrasound machine, for which the point spread function has been measured experimentally — the pulse has a centre frequency of 6.0 MHz and bandwidth 2.1 MHz — and the sampling frequency is 66.7 MHz.

For each frame, 128 A-lines were simulated, spanning 40 mm in the lateral direction, recorded to a depth of 40mm. Simulations were performed at a range of compressions (0%, 0.01%, 0.1% 0.5% 1.0%, 2.0%, 4.0%) by rescaling the axial spacing of the scatterers, so that algorithms can be tested at a range of net compressions (0.01%, 0.1%, 0.5%, 1.0%, 1.5%, 2.0%, 3.0%, 4.0%). This is

³<http://www.dynamicimaging.co.uk>

(a) Amplitude	AMC		
	Off	On (envelope)	On (envelope & phase)
Normal	EPZS	EPZS_A1	EPZS_A2
Discarded	EPZS_L	-	EPZS_LA

(b) Amplitude	Phase deweighting	AMC	
		Off	On
Normal	0	WPS_0	WPS_A0
	1	-	WPS_A1
	2	-	WPS_A2
	3	-	WPS_A3
Discarded	0	WPS_L0	-
	1	-	WPS_LA1
	2	-	-
	3	-	-

Table 1: Algorithm variants in (a) EPZS and (b) WPS.




Figure 3: B-scan of simulated RF data.

important, because the performance of the deformation estimators is strain dependent. Five data sets were generated for different scatterer fields. This contributes to the reliability of the results, which record the mean across the five data sets.

The Field II output was converted to the RF ultrasound format of the Stradwin⁴ freehand 3D ultrasound system. RF samples are recorded with 16-bit signed integer precision. The signals were normalised before conversion, such that in all cases the mean power is fixed at $V_{\text{rms}} = 2^{10}$. The data were combined with additive white Gaussian noise, reducing the SNR to 20 dB. Figure 3 shows an example B-scan from the simulated data.

4 Results

Results are presented in graphical form, with plots of performance (SNR_e or SNR_e^β) against either window length or strain. Recall that the same window parameters are used across all algorithms to yield comparable data. The use of differencing for strain estimation and the absence of error removal such as median filtering mean that for basic algorithms the performance may appear poor — it should be considered that any performance could be boosted at the expense of resolution by standard filtering techniques.

Figure 4 compares the basic algorithms, EPZS and WPS_0. Performance where AMC is not applied and amplitude is not discarded is perhaps unimportant. Subsequent Figures 5–7 demonstrate the need to account for amplitude modulation. Figure 8 illustrates the benefit of phase

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

deweighting. Results for the heuristic phase deweighting strategies are shown in Figure 9, which is the first instance where SNR_e^β is the performance measure — it is found that a substantial effect of phase deweighting is to increase the spacing of the estimation locations, which makes comparisons based on SNR_e unfair.

Leading algorithms from the results so far are compared in Figure 10. The main distinction is between the performance of discarded amplitude algorithms as opposed to the retained amplitude algorithms with AMC. The leading algorithms, including the best of the discarded amplitude algorithms, are compared by means of SNR_e^β -strain characteristics in Figure 11. With short windows these result in typical “strain filters”, but a different pattern emerges when the optimal window length is chosen at each strain for each algorithm

Confusion could arise in the interpretation of some of these results, especially in the cases of long windows and high strains. SNR_e results do not indicate whether performance degradation is caused by small numbers of outliers or gradual degradation across the image. This is clarified in Figure 12 where a median filter has been applied to the strain estimates.

Finally, Figure 13 demonstrates *in vivo* strain imaging where WPS_A1 has been applied to freehand scanning using the Dynamic Imaging Diasus ultrasound machine with a 5–10 MHz probe. The raw RF ultrasound signal is sampled at 66.7 MHz by a PC running Stradwin, as described in [8]. This enables validation of the simulation work, checking that the correct behaviour is observed under real scan conditions.




Figure 4: Comparison of the basic algorithms, EPZS and WPS_0 at a range of strains and window lengths. Performance is (almost) identical at 0.1% strain, while WPS_0 performs slightly better at 1% and 4%, indicating that it is slightly less affected by amplitude modulation.




Figure 5: Comparison of the AMC variants applied to EPZS at a range of strains and window lengths. Both implementations of AMC yield far higher performance than the basic algorithm, and the new correction in EPZS_A2 (accounting for phase deweighting) outperforms the old correction at high strains with long windows.




Figure 6: EPZS_A2 is compared with the discarded amplitude variants at a range of strains and window lengths. The discarded amplitude algorithms perform less well, but the new AMC correction applied in EPZS_LA improves performance. Note that the discarded amplitude algorithms show a lesser reduction in performance when the window is longer than optimal.




Figure 7: Comparison of the WPS variants in the absence of phase weighting at a range of strains and window lengths. Again, this demonstrates the importance of handling amplitude modulation. Discarded amplitude performs less well than AMC, but it is once again less badly affected by excessive window length.




Figure 8: Comparison of the best EPZS variant with two WPS variants at a range of strains and window lengths. WPS_A0 is found to be less successful than the other algorithms because it has no phase deweighting, which is especially marked at high strains where the performance of WPS_A0 is hit far earlier by phase wrapping. WPS_A1 with moderate phase deweighting marginally outperforms EPZS_A2.




Figure 9: Comparison of WPS with different levels of phase deweighting (n) at a strain of 0.5%. (a) SNR_e. (b) Mean spacing of location estimates. (c) Performance adjusted for resolution, SNR_e^β.




Figure 10: Comparison of leading phase based algorithms at a range of strains and window lengths recording SNR_e^β. At 0.1% strain the performances of WPS_A0, WPS_A1 and EPZS_A2 are (almost) identical, with lower performance from the three amplitude compression algorithms. At 1.0% strain the best performance comes from WPS_A1 followed by EPZS_A2, although discarded amplitude algorithms WPS_LA1 and EPZS_LA perform well with long windows. Performance at 4.0% at most window lengths is best for WPS_A0, though only by a small margin. For longer windows WPS_A1 and EPZS_A2 are more robust, though still less robust than the discarded amplitude variants.

5 Discussion

5.1 Interpretation of results

In general, the results indicate that it should be possible to develop high performance strain estimation systems based on either WPS or EPZS. Following the tests on a range of variants, it is clear that subtle modifications bring marked changes in performance, both in quantitative terms (better/worse) and also in qualitative terms (better suited to particular scan conditions).

Before considering the results in detail, it is instructive to consider the meaning of the performance measures, their value for evaluating deformation estimation algorithms, and also their limitations. SNR_e is a good measure because it usually aligns closely with the level of noise that is perceived subjectively on inspecting a “uniform” image. The main limitation of SNR_e is that it does not adjust for performance changes that are introduced by resolution changes rather than changed algorithm performance. One determinant of resolution is the spacing of estimation locations. Since this is modulated by the weighting strategy, for algorithms that use AMC, more meaningful comparisons can be made by considering SNR_e^β. However, a remaining limitation is




Figure 11: Comparison SNR_e^β -strain characteristics of leading algorithms at various window lengths. The left and middle graphs show the shortest and longest window lengths at which results were recorded over all strains in the range 0.01–4%, while the right-hand graph shows the results for each algorithm where the optimal window length was employed at each strain. (Note that the discarded amplitude algorithm could not be tested at its optimum at low strains, because the range of window lengths permitted in the test software had an upper limit of 108.1λ . Interpretation is complicated since resolution clearly depends partly on window length.)




Figure 12: Effect of outliers. When phase wrapping errors begin to occur, initially there are only a few large “peak-hopping” errors. Wherever peak-hopping occurs it registers as an extremely large strain error, sufficiently large to skew the SNR_e value of the entire strain image far beyond the effect it has on subjective image quality. For example, images are shown for EPZS_A2 operating on 4% strain with various window lengths (*cf.*, Figure 5c). (a) $11.8\lambda \Rightarrow \text{SNR}_e = 10.6$, (b) $12.7\lambda \Rightarrow \text{SNR}_e = 4.7$, (c) $14.5\lambda \Rightarrow \text{SNR}_e = 3.0$. (d) 14.5λ with outliers removed by a 3.5mm lateral median filter $\Rightarrow \text{SNR}_e = 33.4$. (e) Leading algorithms at 4.0% strain are compared in conjunction with median filtering (*cf.*, Figure 10c).

the effect of window length on resolution, which is still unaccounted for. It would be best if a new performance measure, SNR_e^γ , say, could be introduced to adjust for window length as well. The relationship between window length and resolution in strain images is difficult to determine, however, and insofar as this relationship has been investigated, it seems that resolution is affected less severely by window length than by window spacing [2]. Thus, when inspecting graphs of performance against window length, it should simply be recalled that the resolution is to some extent reduced: windows of length 30λ that yield ten times higher SNR_e^β than 3λ are probably producing




Figure 13: Images from a freehand scan of human biceps *in vivo*. (a) B-scan. (b) Strain image using WPS_A1 (1D windows of length 13.5λ at an axial spacing of 2.7λ) with a 2.0 mm (*i.e.*, 15.5λ) least squares filter. Black represents zero strain and white represents 1.5% relaxation.

somewhat less than a tenfold increase in performance.

Figure 4 shows the basic algorithms where amplitude is retained without AMC. At each strain there is an optimal window length that yields the maximum SNR_e , which is longest at low strain and shortest at high strain. The height of the peak is also strain dependent, with better performance at the highest strain. This is because the deformation signal is then larger, but the small scale of the increase in the peak SNR_e also indicates an increased level of noise. The deformation signal at 4% strain is 40 times greater than at 0.1% strain, but SNR_e increases just 1.5 times. This arises partly because higher strain leads to increased signal decorrelation, and also partly because amplitude modulation (estimation location error) becomes increasingly important. WPS_0 performs slightly better than EPZS at the higher strains, because it has no phase deweighting. Phase deweighting causes exaggeration of the amplitude modulation effect.

Amplitude modulation is corrected with the algorithms in Figure 5, where AMC is applied to EPZS. AMC offers considerably better performance at all window lengths and all strains. The benefit is especially marked at low strains, where decent SNR_e is not possible with short windows though much longer windows offer substantial improvement. The results show that the new AMC implementation (EPZS_A2) is superior to the original implementation (EPZS_A1). That said, the difference only becomes apparent when the largest phase separations in a window are on the scale of 2π , *i.e.*, when the displacement across the window is on the scale of λ . In these cases, phase deweighting becomes important. Stark differences between EPZS_A1 and EPZS_A2 are therefore observed for long windows at 1.0% and 4.0% strain. On the other hand, phase deweighting is practically never an issue at low strains such as 0.1%: substantially an entire A-line needs to be covered by a single window before the displacement over the window even reaches $\lambda/4$. Arbitrarily long windows cannot be used at most strains, however, since the theory underlying both algorithms breaks down at phase separations in excess of $\pm\pi$.

Extreme phase separations eventually occur at the edges of long windows at high strains. AMC is compared with the discarded amplitude variants in Figure 6. Discarded amplitude algorithms perform less well because the preprocessing is equivalent to selectively amplifying signal portions where the ultrasonic SNR is poor. The attractive feature of discarded amplitude approaches, on the other hand, is increased robustness when excessive window lengths are employed.

Discarded amplitude still requires AMC location estimation, owing to the phase deweighting effect at the edges. It is for this reason that EPZS_LA outperforms EPZS_L. However, it is interesting to note that EPZS_L has higher SNR_e over a small region, which implies that the present implementation of AMC leaves room for future improvement. Nevertheless, with the longest windows the behaviour of EPZS_LA is eventually preferable to both EPZS_A2 and EPZS_L.

The patterns of behaviour in the WPS family with AMC and discarded amplitude are strikingly

similar to EPZS. Figure 7 shows that WPS_A0 and WPS_L0 both outperform WPS_0, as expected. WPS_A0 offers higher peak performance than WPS_L0, although it deteriorates sooner at excessive window lengths. There is a simple explanation for the relative robustness of discarded amplitude algorithms in both WPS and EPZS paradigms. The estimation location remains relatively close to the window centre when there is no amplitude, so phase wrapping at the window edges sets in later compared to the other algorithms. When amplitude is retained, the displacement estimate sometimes applies with greatest validity near the end of the window, meaning that the maximum alignment error at the far end is up to twice what it would be for an estimate at the centre. Phase wrapping occurs sooner as the greatest phase separations approach $\pm\pi$. Furthermore, even when phase wrapping is present with discarded amplitude, if the estimation location is at the centre of the window, then phase wrapping errors at either end of the window cancel out (by symmetry). By contrast, if the estimation location has shifted significantly away from the centre, then phase wrapping errors are highly asymmetric, so the onset of phase wrapping causes abrupt performance degradation.

Figure 8 compares results of the best algorithms tested so far in addition to WPS with phase deweighting, revealing that when long windows are used WPS_A0 actually performs less well than EPZS_A2. Better performance of the heuristic WPS_A1 demonstrates that phase deweighting is useful, and the phase deweighting in EPZS is clearly suboptimal. Since the phase deweighting strategy in WPS is entirely heuristic, Figure 9 compares the full set of deweighting options at 0.5% strain. At first glance, it appears that a performance improvement is produced at every step by incrementing the phase deweighting parameter n from 0 up to 3. However, Figure 9b reveals that this is caused by changes in the mean spacing of the estimation locations. Therefore, we turn in Figure 9c to the adjusted performance measure, SNR_e^β . The adjusted characteristics show a subtler impact from phase deweighting, with best performance at $n = 1$ or 2. Across a range of strains $n = 1$ most often yields the highest SNR_e^β , so this is adopted for the remaining tests.

Leading variants of WPS and EPZS are compared in Figure 10 recording SNR_e^β . There is actually little difference between WPS_A0, WPS_A1 and EPZS_A2, indeed the performances at 0.1% strain are almost identical. WPS_A0 does worse at higher strains when phase deweighting is needed. The discarded amplitude algorithms perform less well in general, although the robustness with long windows is demonstrated again. WPS_LA1 is the best discarded amplitude algorithm, marginally outperforming EPZS_LA.

It is interesting to consider the use of a fixed window length over a range of strains to compare SNR_e^β -strain characteristics of the leading algorithms. The form of the characteristics in the left and middle graphs in Figure 11 are typical of “strain filters” that are often illustrated in the literature. However, graphs at fixed window lengths neglect one of the striking features of these results as a whole: at lower strains the highest performance is achieved by using much longer windows. When optimal window lengths are employed at every strain level, the height of the peak in SNR_e^β is shown in the right hand graph. Perversely, this indicates that by a small margin WPS_A0 may be the *best* algorithm if short windows are used, or the *worst* if arbitrary window lengths can be used. With optimal window lengths, the best of the algorithms with retained amplitude is WPS_A1. The other feature is that WPS_LA1 apparently performs best for low strain estimation with optimal window lengths. A rigorous investigation of the effect of window length on resolution is required so that these interesting findings can be properly understood (and exploited, if appropriate).

A caveat is necessary to explain the sharp peaks in the graphs of SNR_e and SNR_e^β against window length for retained amplitude algorithms at 4% strain. It would be easy to assume that the peak has fundamental significance, enabling the maximum window length to be predicted by some theoretical means. In fact, the drop in performance occurs wherever large errors first arise, *i.e.*, the first instances of peak-hopping errors where windows are matched approximately one wavelength away from the actual displacement. The reality of this effect is illustrated in Figure 12b, where a single large error has dramatically reduced the value of SNR_e . However, note that this is an outlier. It is a disadvantage of SNR_e as a performance measure that outliers can have a large impact. The introduction of a single strain error in the region of 30% has a dramatic impact when the other strain estimates are all in the range 3.6%–4.4%. A large fraction of the

noise may be produced by a single estimate. Similarly, in Figure 12c there are basically few large errors, but the value of SNR_e is now very low. If a median filter is applied, these errors can be removed and the image has a high SNR_e . The results in Figure 12e further demonstrate this point by repeating the test from Figure 10c, with the addition of a lateral median filter spanning 3.5 mm / 11 A-lines. With outliers excluded, most algorithms are still performing better at 14.5λ , with the best performance from algorithms WPS_A1 and EPZS_A2.

Simulation results do not necessarily imply that the same behaviour is repeated in the context of real data. However, WPS is already being tested in conjunction with *in vivo* freehand ultrasound data, and the image quality matches expectations following the simulation results. Figure 13 shows a scan of human biceps, for example, in which the main features in the strain image correspond to different layers of muscle.

5.2 Conclusion

The tests show that both the WPS and EPZS families of deformation estimators have strong potential for supplying the first stage processing in ultrasonic elasticity imaging systems. It is relevant to point out, therefore, that these algorithms generally entail simplicity and low computational cost compared to most alternatives in the literature. Both families are suitable for real time processing on current desktop computers. This is already well known for EPZS, while WPS represents a new paradigm. The main cost in any variant of WPS is signal pre-processing, at which stage signal envelope and phase must be detected. Thereafter the iterative searches consist mostly of additions in the case of discarded amplitude, or simple additions and multiplications where the amplitude is retained. EPZS_A2 comes at a slightly higher computational cost, since the sine function is required to evaluate the weightings.

WPS_A1 combines simplicity, versatility and high performance in a framework that can readily be optimised depending on the particular properties of specific ultrasonic deformation estimation applications. There is no apparent theoretical reason why it should not also be possible to perform estimation in the lateral direction when lateral phase is detected, or indeed in three directions when 3D data is acquired. It may also be possible to adapt WPS for deformation and displacement estimation applications in other fields altogether.

Appendices

A Robust iteration seeding

A novel iteration seeding strategy was employed for all algorithms in this investigation. The rate of peak-hopping errors is close to the rate for exhaustive searches, but at a far lower computational cost. The standard approach to iteration seeding was developed by Pesavento *et al.* [12]. Each A-line is processed individually, and the first window at the top of each A-line is seeded with a trial displacement of zero. This is sensible, since tissue that remains in contact with the probe surface by definition has zero displacement in the probe reference frame. Each following window is seeded with the displacement estimate from the preceding window, since the differential displacement over the distance between windows is almost invariably well below half a wavelength. This strategy enables high accuracy tracking, provided that errors in the displacement estimates are much smaller than $\lambda/2$. However, sometimes large errors do occur owing to small regions of extreme decorrelation. If the following window is then too far from the correct alignment, it produces another error when it converges on the wrong phase zero, thus some errors propagate. Once a large error has appeared in an A-line, it is rare and purely a chance event if subsequent estimates return to the correct phase zero, otherwise the remainder of the A-line consists only of noise.

An example of this type of error propagation is shown in Figure 14b, taking a single strain image from a freehand scan of an olive/gelatin phantom (this mimics a stiff inclusion in soft tissue). The probe has rolled slightly about the elevational axis, so contact is best at the right hand edge of the image, and there is no contact on the left. One of the results is that the strain is higher on the

right. Despite the limited contact region, there is no problem with the coupling of the ultrasound beam into the phantom, since the phantom is covered in a fine layer of coupling fluid. However, the coupling fluid sometimes introduces severe decorrelation because it can flow out of plane. It is also problematic if the first windows track fluid, and a displacement discontinuity arises with the first window that tracks the solid part of the phantom. This makes alignment errors in excess of $\lambda/2$ more likely. For example, in Figure 14b large errors have occurred in three A-lines followed by error propagation.

To avoid tracking errors, we have developed new seeding strategies to support advances in phase-based deformation estimation. There are three features to the superior iteration seeding. Note that only the first of the features was required for the algorithms tested in the simulation study.

Cross-seeding. A-lines are searched in parallel, *i.e.*, a first window displacement is estimated at the top of every A-line before proceeding to the second row. The correlation coefficient (or any alternative accuracy indicator) is evaluated for each pair of matched windows. Each window in the next row working down every A-line is seeded with the displacement estimate from the nearby window in the previous row which had the highest correlation coefficient, searching laterally across l A-lines to either side. $l = 1$ means that the seed is taken from the current A-line or either immediate neighbour, which immediately eliminates almost all error propagation, *e.g.*, where it has been applied in Figure 14c. In rare cases when a large region of the image becomes erroneous, it may be useful to employ a larger value for l . This could happen if part of the image is produced from tissue with poor mechanical contact at the probe surface, or where there are large decorrelating features such as major arteries. In these cases l governs the rate of “correction propagation” (see Figure 15b). Anything up to and including every A-line in the image may be searched for cross-seeding. However, large values of l can cause errors if there are high shear strains. In practice, we find that $l = 10$ is a good choice for freehand strain imaging.

Multi-pass analysis. Owing to the limited rate of correction propagation, part of the image may remain erroneous despite cross-seeding. This can be fixed by repeating the process. Multi-pass analysis requires that the correlation coefficient (or alternative accuracy estimate) is stored alongside every displacement estimate. A second pass begins with windows at the bottom of the image, seeding each window with the best estimate from the previous row (below). This enables continued correction propagation by cross-seeding, and a displacement estimate from the previous pass is replaced if a higher correlation coefficient is recorded. Since the processing time for phase-based methods is mainly pre-processing rather than the actual search, two-pass analysis does not severely reduce the speed of the algorithm. For $l = 10$, it almost never arises that errors remain after a second pass, but the example in Figure 15 used $l = 1$. This slow rate of correction propagation left a triangle of errors in the bottom left corner after the second pass. This can be removed by a third pass (top to bottom). In general, for maximal robustness passes should continue down and up the image until none of the displacement estimates change. This is still far less computationally expensive than an exhaustive search. However, with $l = 10$ it almost never occurs that any error propagation remains after the second pass, so the best practical solution may be always to apply a two-pass strategy.

Continuity checking. Regions of data with extremely poor ultrasonic SNR or high levels of decorrelation sometimes yield a higher correlation coefficient at the wrong phase zero than near the actual displacement. This means that sometimes (albeit rarely) an analysis pass tracking one wavelength away from the actual displacement becomes interleaved with an analysis pass that produced good estimates. When this happens it causes large errors, because between-window displacements larger than λ may be recorded (*cf.*, Figure 12). To overcome this, each pass of a multi-pass analysis produces displacement estimates that either enter the display buffer (if the correlation coefficient is higher) or enter a reserve buffer (if the correlation coefficient is lower). Before the next pass, continuity checking refines the




Figure 14: Example of cross-seeding in scan of olive/gelatin phantom. (a) B-scan. The strain images employ a linear strain scale with black for 0% and white for 1% extensive strain. (b) Strain image with error propagation. (c) Strain image where cross-seeding has eliminated error propagation.




Figure 15: Example of multi-pass analysis. (a) B-scan. The strain images employ a linear strain scale with black for 0% and white for 1% compressive strain. (b) Single-pass strain image with cross-seeding (down). (c) Two-pass strain image (down, up). (d) Three-pass strain image (down, up, down).

sorting of estimates between the display and reserve buffers: sequentially, working through the display buffer in any direction, at every point an average is computed of the four adjacent displacement estimates in the display buffer (above, below, left, right), and the reserve buffer entry replaces the display buffer entry in the few cases when it is closer to the average value. A similar effect could be achieved by median filtering, but in that case resolution would be reduced.

The images in Figure 15 come from the same scan of an olive/gelatin phantom as Figure 14. In Figure 15 only the corner of the transducer casing was in contact with the phantom surface, so a large displacement discontinuity is present at every A-line (largest discontinuity on the left of the image). In this case errors have propagated in every A-line, but as soon as one of the estimates finds the the correct phase zero a higher correlation coefficient is registered, and the correction propagates into neighbouring A-lines. The low value $l = 1$ and the lateness of the correction mean that three passes are required before error propagation has been eliminated.

References

- [1] S. K. Alam and J. Ophir. Reduction of signal decorrelation from mechanical compression of tissues by temporal stretching: applications to elastography. *Ultrasound in Medicine and Biology*, 23(1):95–105, January 1997.

- [2] S. K. Alam, J. Ophir, and T. Varghese. Elastographic axial resolution criteria: an experimental study. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 47(1):304–309, January 2000.
- [3] J. Bercoff, M. Tanter, and M. Fink. Supersonic shear imaging: a new technique for soft tissue elasticity mapping. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51(4):396–409, April 2004.
- [4] X. Chen, M. J. Zohdy, S. Y. Emelianov, and M. O’Donnell. Lateral speckle tracking using synthetic lateral phase. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51(5):540–550, May 2004.
- [5] 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.
- [6] F. Kallel and J. Ophir. A least-squares strain estimator for elastography. *Ultrasonic Imaging*, 19:195–208, 1997.
- [7] R. M. Lerner, S. R. Huang, and K. J. Parker. Sonoelasticity images derived from ultrasound signals in mechanically vibrated tissues. *Ultrasound in Medicine and Biology*, 16:231–239, 1990.
- [8] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. 3D elastography using freehand ultrasound. *Ultrasound in Medicine and Biology*, 32(4):529–545, April 2006.
- [9] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Estimation of displacement location for enhanced strain imaging. Technical Report CUED/F-INFENG/TR 550, Cambridge University Department of Engineering, March 2006.
- [10] 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.
- [11] K. J. Parker, S. R. Huang, R. A. Musulin, and R. M. Lerner. Tissue response to mechanical vibrations for sonoelasticity imaging. *Ultrasound in Medicine and Biology*, 16:241–246, 1990.
- [12] 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.
- [13] 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.