Published in final edited form as:

IEEE Trans Ultrason Ferroelectr Freq Control. 2010 April ; 57(4): 866–882. doi:10.1109/TUFFC.2010.1491.

A hybrid displacement estimation method for ultrasonic elasticity imaging

Lujie Chen, R. James Housden, Graham M. Treece, Andrew H. Gee, and Richard W. Prager
University of Cambridge Department of Engineering Trumpington Street Cambridge CB2 1PZ

Abstract

Axial displacement estimation is fundamental to many freehand, quasistatic, ultrasonic strain imaging systems. In this paper, we present a novel estimation method that combines the strengths of quality-guided tracking, multi-level correlation and phase-zero search to achieve high levels of accuracy and robustness. The paper includes a full description of the hybrid method, *in vivo* examples to illustrate the method's clinical relevance, and finite element simulations to assess its accuracy. Quantitative and qualitative comparisons are made with leading single- and multi-level alternatives. In the *in vivo* examples, the hybrid method produces fewer obvious peak-hopping errors, while in simulation, the hybrid method is found to reduce displacement estimation errors by 5–50%. With typical clinical data, the hybrid method can generate more than 25 strain images per second on commodity hardware: this is comparable with the alternative approaches considered in this paper.

I. INTRODUCTION

The mechanical properties of tissue have long been recognised as a useful indicator of disease. In what is a natural development of manual palpation, elastography involves the imaging of tissue deformation induced by some sort of applied mechanical stress. The measured deformation, taken alongside any knowledge of the stress, allows estimation of the tissue's mechanical properties. This general paradigm embraces a wide range of different techniques, the principal distinguishing factors being how the stress is applied and how the deformation is measured. The former ranges from external palpation [1] to internal acoustic radiation force [2], while the latter might involve imaging modalities as diverse as MRI [3] and ultrasound [4].

Within this broader context, the subject of this paper is freehand quasistatic axial strain imaging [1], which has shown great promise in clinical trials [5] and has recently been commercialized¹. The clinician holds an ultrasound transducer over the area of interest while gently varying the contact pressure, thus inducing mechanical stress in the targeted tissue, predominantly in the axial direction. Consecutive frames in the RF data sequence are compared to estimate the resulting axial displacement which, in turn, is differentiated to arrive at the axial strain field. This strain field provides useful information about the tissue's mechanical properties.

© 2010 IEEE.

Lujie Chen (lc420@eng.cam.ac.uk), R. James Housden (rjh80@eng.cam.ac.uk), Graham M. Treece (gmt11@eng.cam.ac.uk), Andrew H. Gee (ahg@eng.cam.ac.uk) and Richard W. Prager (rwp@eng.cam.ac.uk).

¹e.g. ACUSON Antares, Siemens Medical Systems, Issaquah, WA, USA; Sonix, Ultrasonix Medical Corporation, Richmond, BC, Canada; HI VISION 900, Hitachi Medical Systems, Japan.

At the heart of most quasistatic strain imaging systems is an algorithm for estimating tissue deformation from one RF frame to the next. These are special cases of more general speckle tracking algorithms, with a particular emphasis on axial displacement estimation. Although axial compression generally results in tissue deformation in all three dimensions, it is only the axial displacement that is required for axial strain imaging. Lateral and (in the case of three-dimensional imaging) elevational motion are also of interest, but only inasmuch as they can be used to improve the accuracy of the axial displacement field. Lateral strain imaging *per se*, for instance to estimate the Poisson's ratio, is a less common endeavour that is not the subject of this paper.

While it is difficult to formulate a coherent classification of existing displacement estimation algorithms, they can nevertheless be described in terms of certain common themes. For example, all displacement estimation algorithms deal at some stage with the matching of blocks, or windows, of data between the pre- and post-deformation frames. Similarity metrics used for this purpose include cross-correlation [1], [4], [6], sum of absolute differences [7], [8], sum of squared differences [9], correlation phase [10], [11] and phase separation [12]. Windows are typically aligned to subsample precision [9], [11]–[21], perhaps by way of continuous, cubic spline signal representations [22]–[24]. Matching might be carried out by independent exhaustive searches at each window [4], or some inter-window continuity could be imposed, either by minimizing a global cost function that penalises discontinuous displacements [9], [19] (perhaps within the context of an optical flow model [25]), or by tracking the displacements from one window to the next [12], [16], [20], [21], [26]–[28]. The search at each window might be fully two-dimensional [15], or there may be separate, one-dimensional axial and lateral searches, perhaps with some iteration between the two [17], [28]. There could be a single, high resolution matching process [11], [12], [14], [18], [28], or a multi-level approach with a coarse matching stage initializing subsequent stages at progressively higher resolutions [8], [9], [19], [29]–[31].

A well-designed displacement estimation algorithm aims to strike an appropriate balance between accuracy, robustness and speed. Single-level exhaustive search is neither fast nor robust: a large search range is required at each window, and false-positive matches are common when the window size is small (as required for fine scale displacement estimation) and the noise level significant [32]. Robustness is generally improved by imposing some sort of continuity on the displacement estimates. If this is achieved through tracking, there is also a benefit in terms of speed, since the search can be confined to a small range around the displacement estimate propagated from the previous window.

A particularly efficient way to track axial displacements is to use some sort of phase-based similarity metric between the RF signals in the pre- and post-deformation windows. Optimal alignment is indicated by zero correlation phase [10], or zero phase separation [12], but a non-zero result, coupled with knowledge of the transducer's centre frequency, allows rapid iteration towards the correct shift provided the initial alignment is within half a wavelength of the zero-phase shift [11]. Phase-zero search is typically employed within a tracking framework, with the displacement estimated at each window providing the initial alignment at its immediate neighbours. Such a scheme is effective, since the true displacement will not normally change by more than half a wavelength from one window to the next. An inaccurate estimate of the centre frequency does not affect the accuracy of the displacement estimate, just the speed of convergence [11].

However, phase-zero tracking is no panacea. It is essentially a one-dimensional, axial process, so lateral motion must be compensated using some other technique [18], [28]. In common with all approaches that use previous estimates to constrain search windows, it is vulnerable to catastrophic failure through error propagation. Errors might be present from

the outset if the initial seeds are incorrect, or they might be introduced mid-course through a noise-induced bad match or through a discontinuity in the underlying deformation field, for instance at a slip plane. Error propagation can be arrested if the tracking path is allowed to evolve dynamically according to some quality measure: so windows with high quality matches, as judged by, for instance, the post-alignment correlation, are preferred to less promising windows when it comes to initializing the searches at their neighbours [12], [26]–[28]. Alternatively, *post hoc* continuity checks orthogonal to the tracking direction can be successful in suppressing error propagation [16], [21]. Displacement can even be tracked in disjoint regions separated by slip planes, provided multiple seeds are employed, at least one seed is planted in each region, and the tracking wavefronts evolve in a quality-guided manner [26]. But everything depends on the correctness of the initial seeds, which is difficult to guarantee when there is decorrelation caused by elevational displacement or noise.

Compared with single-level tracking, multi-level strategies are intrinsically more robust to decorrelation and error propagation. They start by estimating coarse-scale motion on a sparse grid of large windows, typically using the correlation of the RF envelope as the similarity metric. This process is relatively noise tolerant, since the similarity metric depends on the data's macro-structure, which does not deteriorate so rapidly in the presence of noise and fine-scale RF signal decorrelation [31]. Moreover, axial and lateral displacement can be assessed by the same metric. The resulting coarse motion estimates are interpolated onto a finer grid of windows at the next level, where they seed higher resolution searches, and so on through several levels, until the desired resolution is reached. Should an error occur at any stage, it affects subsequent stages only within its local region, with no risk of propagation throughout the entire frame. The principal shortcoming of multi-level approaches is the processing overhead of the higher level searches. Fast, phase-based techniques are not well-suited to large windows, since they are relatively intolerant of intra-window strain [12].

In this paper, we present a hybrid displacement estimation method that combines elements of multi-level search and single-level tracking. At the finest level, where computational cost is paramount, we use quality-guided, phase-zero tracking to estimate the axial displacements. This process is preceded by a two-dimensional, correlation based, multi-level search that provides the initial axial and lateral offsets at a subset of the fine level windows. The expense of lateral search is avoided at the fine level: instead, the initial estimates are simply interpolated to provide lateral motion compensation at all the fine level windows.

Details of the hybrid method are presented in Section II. In Section III, we describe *in vivo* and simulation experiments designed to compare the hybrid method with leading multi-level and single-level alternatives. The results are discussed in Section IV, with attention to speed, accuracy, robustness to signal decorrelation and resilience in the presence of slip planes. Finally, in Section V we present our conclusions and some suggestions for further work.

II. HYBRID DISPLACEMENT ESTIMATION

Figure 1 provides an overview of the hybrid displacement estimation method. There are three levels of processing that we shall refer to as L1, L2 and L3. In common with other multi-level approaches, successive levels are seeded by their immediate predecessors and recover an increasingly fine-scale grid of displacement estimates by matching windows between pre- and post-deformation frames. There are only nine L1 windows, whose centres coincide with a subset of the L2 windows. Likewise, the L2 window centres coincide with a subset of the L3 windows. In this way, displacement estimates can be passed from one level to the next, with no need for interpolation.

L1 and L2 operate on RF envelope data and utilize coarse-to-fine block matching with a standard cross-correlation metric. Each L1 window is searched independently within predetermined axial and lateral bounds, while L2 employs quality-guided axial-lateral displacement tracking. L3 operates on baseband analytic data and performs fine scale axial displacement estimation using the weighted phase separation (WPS) algorithm [12], again with quality-guided tracking. Lateral displacement estimates are not refined at L3. In common with other methods of this class, the hybrid method is best suited to linear array transducers, with A-lines running parallel to the direction of applied stress.

A. Preprocessing

Pre- and post-deformation baseband analytic signals are computed by filtering the raw RF data with a Hilbert filter (antisymmetric coefficients) and a symmetric filter with an appropriate passband, then multiplying by $\exp(-j\omega_0 t)$, where ω_0 is the transducer's centre frequency. The signals are then downsampled by a factor of S_y . For optimal processing speed, S_y should be set to as high a value as possible without violating the Nyquist sampling condition for the baseband signal. The downsampled, baseband analytic signals are used in all subsequent processing. The amplitude a provides the envelope data for L1 and L2, while the phase ϕ is required by the WPS algorithm at L3.

B. Level 1 search

At level 1, nine sparsely distributed windows in the pre-deformation frame are matched with corresponding windows in the post-deformation frame by searching for the peak of the correlation coefficient

$$C(d_x, d_y) = \frac{\sum_{(x,y) \in T} [a_1(x, y) - \bar{a}_1] [a_2(x+d_x, y+d_y) - \bar{a}_2]}{\sqrt{\sum_{(x,y) \in T} [a_1(x, y) - \bar{a}_1]^2 \sum_{(x,y) \in T} [a_2(x+d_x, y+d_y) - \bar{a}_2]^2}} \quad (1)$$

where x and y denote position in the lateral and axial directions respectively, d_x is the lateral displacement, d_y is the axial displacement, a_1 and a_2 are the envelope data before and after deformation respectively, $\bar{a}_1$ and $\bar{a}_2$ are the intra-window averages of a_1 and a_2 , and T is the window size. The search at each window is independent of its neighbours (ie. no tracking) and is constrained within predetermined axial and lateral bounds. The search is not exhaustive, but instead adopts the faster multi-resolution strategy illustrated in Figure 2. The parameters of this search are the window dimensions (3×12 in Figure 2), the search region dimensions (6×24 in Figure 2) and the initial skip factor, which must be a power of two (4 in Figure 2). Note that some sampling asymmetry must be tolerated at the lower resolutions, as is evident in Figure 2(a).

It is only at the lowest resolution that the number of candidate match locations scales with the size of the search area. Since computation at this level is relatively cheap, the algorithm can search a large region with an acceptable overhead: we provide typical timings in Section IV. At finer resolutions, where the correlation coefficients require progressively more computation, the number of candidate locations will not exceed three. The output of the L1 search is nine independent axial-lateral displacement estimates, one for each of the L1 search windows.

C. Level 2 search

The L2 search uses windows and search ranges that are smaller than those at L1. The number of L2 windows is a further independent parameter of the algorithm. This number is

constrained by the need for a regular distribution of L2 windows with respect to the L1 windows. We suggest one or two L2 windows in the axial and lateral gaps between each L1 window, as in Figure 1. In our experience, performance is not particularly sensitive to whether there are one or two windows in each gap.

As shown in Figure 1, nine of the L2 windows are concentric with the L1 windows and can therefore inherit initial axial and lateral displacement estimates from the output of the L1 search. These and the remaining L2 windows are matched using the same multi-resolution technique as in Figure 2, though with a smaller initial skip factor. However, this time the searches are not independent. Instead, a tracking strategy is used to propagate displacement estimates from one window to the next, with the search at each window confined to a narrow range centred around the estimate propagated from its neighbour (or from the parent L1 window in the case of the initial nine L2 windows). The direction of propagation is not fixed in advance (eg. up-down or left-right), but is generated dynamically according to a quality metric. Specifically, after the initial nine L2 windows have been processed by multi-resolution search, their match qualities, as given by the post-alignment correlation coefficient in Equation (1), are put into descending rank order. The displacements from the highest quality match are propagated to that window's immediate neighbours, which are then processed, generating a further set of displacement and quality scores that are inserted into the ranked list. The next highest ranking match is then propagated to its neighbours, and so on, until all L2 windows have been processed.

Full details of this tracking strategy can be found in [26]. In common with all tracking approaches, the quality-guided method is fast, with implicit enforcement of displacement continuity in both axial and lateral directions. The advantage of the quality-guided approach is that any tracking failures soon lead to matches with low correlation coefficients that end up at the foot of the ranked list: these misaligned windows do not, therefore, have a chance to propagate their poor displacement estimates to other windows. Furthermore, since the L2 search is seeded at nine evenly distributed locations, the strategy is able to track across displacement discontinuities: it is only necessary that there should be at least one L1 window in each discontinuous region of the underlying displacement field [26]. Highly fragmented fields may require more than nine L1 windows, at the expense of slower computation.

At the end of the L2 stage, the refined axial displacement estimates are used to seed a subset of the L3 windows for the final, high resolution search. However, the lateral estimates are treated differently. Our goal is not to recover lateral displacement to a high degree of precision. Indeed, this may not even be possible with current mid-range scanning technology, since the lateral resolution, limited by the transducer element spacing, is typically worse than the axial resolution². For axial strain imaging, the role of the lateral displacement estimates is merely to compensate for the tissue's lateral motion and hence obtain better axial displacement estimates [8], [18], [19], [31]. We suggest that the L2 lateral displacement estimates are sufficient, and perhaps even optimal, for this purpose. Any attempt to further refine the lateral estimates at L3 may be dangerous, since the correlation of small L3 windows is highly susceptible to noise, especially in the presence of laterally orientated specular reflection. This hypothesis will be tested empirically in Section IV.

At this point, the lateral displacement estimates are optionally smoothed by minimizing the cost function

²For higher precision lateral strain imaging, novel beamforming techniques can potentially increase the lateral resolution to levels comparable with the axial resolution [33]. Furthermore, compared with the hardware available for this study, some top-of-the-range systems employ transducers with significantly higher element densities. We consider such a transducer in Section IV-E.

$$E = \sum_{(x,y) \in U} Q(x,y) [f(x,y) - g(x,y)]^2 + k \sum_{(x,y) \in U} \left[\left(\frac{\partial g}{\partial x} \right)^2 + \left(\frac{\partial g}{\partial y} \right)^2 \right] \quad (2)$$

where f is the lateral displacement field obtained at L2, g is the smoothed lateral displacement field, U is the set of L2 windows and k is a smoothness weight. Q is the displacement estimation precision given by $C(1-C)$ [34], where C is the post-alignment correlation coefficient. The Q term ensures that g approximates f closely in the vicinity of high quality matches, with more smoothing allowed around low quality matches. Numerous techniques, such as Gauss-Seidel iteration [35], the iterative conditional modes algorithm [19], the conjugate gradient algorithm [36], and the multigrid method [37], exist to solve this standard regression problem. We use the multigrid method [37]: see [38] for details of our implementation.

D. Level 3 search

The L3 search operates on the downsampled, baseband analytic signals. The axial displacements are refined to subsample precision by matching the small L3 windows, with weighted phase separation (WPS) [12] as the similarity metric. Phase-based techniques are especially suitable for low-level processing since they can exploit knowledge of the transducer's centre frequency to rapidly converge on the optimal alignment [11]. The lateral displacement of each L3 window is obtained by bilinear interpolation³ of the L2 windows' lateral displacements and subsequently fixed: there is no lateral displacement refinement at L3.

WPS-based matching is applied first to the subset of L3 windows that are concentric with the L2 windows (see Figure 1), since it is here that axial displacement estimates are available from L2. The weighted phase separation of pre- and post-deformation L3 windows is given by

$$\Delta\phi = \frac{\sum_{(x,y) \in T} [a_1(x,y) + a_2(x+d_x, y+d_y)] [\phi_1(x,y) - \phi_2(x+d_x, y+d_y)]}{\sum_{(x,y) \in T} [a_1(x,y) + a_2(x+d_x, y+d_y)]} \quad (3)$$

where x and y denote position in the lateral and axial directions respectively, a_1 and a_2 are the pre- and post-deformation signal amplitudes, ϕ_1 and ϕ_2 are the pre- and post-deformation signal phases, d_y and d_x are subsample axial and lateral displacements, and T is the window size. Here we weight each sample's phase separation by the sum of the amplitudes a_1 and a_2 : alternative weighting schemes are discussed in [12].

In Equation (3), a_2 and ϕ_2 are expressed in terms of subsample axial and lateral displacements. They are obtained by bilinear interpolation of the four neighbouring integer alignments. This is an important and necessary detail, since the lateral displacement interpolated from L2 will not typically be an integer number of A-lines, and we wish to estimate the axial displacement to subsample precision.

With knowledge of the transducer's centre frequency, $\Delta\phi$ can be used to infer a refined axial displacement where the weighted phase separation is closer to the desired value of zero [11], [12]. Next, a new value of $\Delta\phi$ is calculated at this new axial displacement, and the process is repeated. This iterative procedure is a key attraction of phase-based alignment techniques: it

³Cubic interpolation was also investigated, but it comes at greater cost and did not improve the results in Section IV.

converges rapidly to the zero-phase displacement [11], [12]. Finally, the axial displacement estimates are propagated to the remaining L3 windows using the same quality-guided tracking approach as in L2 [26], with WPS-based axial refinement at each window.

III. EXPERIMENTAL METHODS

Several displacement estimation methods are compared *in vivo* and using finite element simulations. The methods under consideration are:

hybrid

the technique described in Section II, using the parameters in Table I. Since L1 and L2 windows are used for coarse, cross-correlation based comparisons, their sizes are fixed with respect to the RF frame size. The L3 windows pertain to the fine level zero-phase search, and therefore have their sizes fixed with respect to the RF wavelength. The parameters are thus set according to the formulae in Table I, with no further tuning to optimize performance on individual data sets.

no lateral offset

as above, but without offsetting the L3 windows by the lateral offsets found at L2. This method provides a useful baseline, illustrating the consequences of not compensating for lateral motion.

L3 lateral tracking

a modification of the hybrid method with fine scale lateral displacement estimation at L3. Identical to the hybrid method at L1 and L2, but at L3, after each WPS-based axial displacement estimate, the correlation between the pre- and post-deformation windows is computed within a search range of ± 2 A-lines, with quadratic regression to locate the peak to sub-A-line precision. The new axial and lateral estimates are propagated to neighbouring windows within the quality-guided tracking framework. This method is included in our study to reveal any artefacts caused by fine scale lateral search.

single-level tracking

the single-level tracking strategy implemented in the Stradwin system [28]. Error propagation is limited by quality-guided tracking, though the tracking direction is constrained to be predominantly axial. A second tracking pass is used to correct remaining estimation errors. This method is included in the comparison since it is free-to-download and a good example of a high performance, single-level tracking strategy.

multi-level

a recently published multi-level approach [31]. The displacement estimates at each level are interpolated onto the finer grid of windows at the next level. This contrasts with the hybrid method, where displacement estimates are passed directly to a subset of the lower level windows, then propagated to the remainder by quality-guided tracking. Full details of our implementation can be found in Appendix A. This is the most recent example of multi-level displacement estimation techniques [8], [9], [19], [29]–[31]. It is included in our study since it provides a natural contrast to single-level tracking and can cope with discontinuous displacement distributions.

Two *in vivo* scans were performed using a T3000 (Terason Ultrasound, Massachusetts, USA) ultrasound system with a 7.75 MHz linear array transducer (128 elements at a pitch of 0.3 mm). The RF sampling frequency was 40 MHz and each frame comprised 128 A-lines.

There was a single transmit focus at 15 mm and dynamic receive focusing. Tissue deformation was induced by manually pressing on the probe. These scans serve to illustrate pertinent features of the various displacement tracking methods and their ability to perform *in vivo*.

Each of the three finite element simulations in Figure 3 was designed to mimic a mode of tissue deformation commonly encountered in clinical practice. In the first two simulations, there is a stiff inclusion palpated with normal axial probe displacement (first simulation, minimal lateral displacement) and tilted axial probe displacement (second simulation, more significant lateral displacement). In the third simulation, the central slab of material is sandwiched between slip planes and will therefore displace laterally under pressure: the resulting discontinuities in the displacement field extend right across the RF frame and pose a real challenge to displacement tracking algorithms. Finite element analysis (Abaqus 6.7, Simulia, Rhode Island, USA) was used to calculate the three three-dimensional displacement fields.

For each simulation, two frames of simulated RF echo data were obtained, pre- and post-deformation, using Field II [39]. The simulation parameters were not untypical of real-world scanning scenarios: 128 elements/A-lines at a pitch of 0.3 mm, element width 0.275 mm, 40 mm axial scanning depth, 6.5 MHz centre frequency and an RF sampling rate of 66.7 MHz. The aperture was 32 elements wide with triangular apodization, there was a single transmit focus at 20 mm, and dynamic receive focusing. The elevational focus radius was 23 mm for both transmit and receive. The scatterer density of 13.8 mm^{-3} was identical for all materials, rendering the inclusion and slip planes invisible in the B-mode images, as shown in Figure 4(a). Examples of axial strain images, obtained using the hybrid method, can be found in Figures 4(b)–(d).

While the simulations allowed for slip planes, there are other factors that are liable to induce gross displacement estimation errors. These include signal drop-out and other phenomena, such as reverberation, that are less easy to model directly. As a proxy for these effects, the experimental protocol involved corrupting the simulated RF echo data with additive noise at a fixed signal-to-noise ratio (SNR) of 30 dB. Out-of-plane tissue displacement, whether caused by the applied mechanical stress or physiological motion, will cause further signal decorrelation. We therefore repeated the Field II simulations with inter-frame elevational transducer offsets of 0, 0.1, 0.2, 0.3 and 0.4 mm.

IV. RESULTS AND DISCUSSION

A. In vivo experiments

The first scan, shown in Figure 5(a), captures a human testis. The axial strain image obtained by the hybrid method in Figure 5(m) suggests that the elliptical testis region contains relatively uniform stiffness tissue. Comparing the lateral displacement images in Figures 5(b) and (c), we see the expected slowly varying distribution with the hybrid method (c), interpolated from just 77 L2 windows, and a much higher frequency distribution with L3 lateral tracking (b). It is difficult to argue that the high frequency components in Figure 5(b) reflect true tissue motion: more plausible is that the correlation peak is rather shallow in the lateral direction (since the lateral focusing is poor), and the process of finding its peak within a ± 2 A-line search range is highly sensitive to noise in the small L3 windows. The noisy lateral displacement estimates result in incorrectly offset windows that corrupt the L3 axial estimates, leading to prominent peak-hopping errors in Figure 5(g) that, when differentiated, produce the dark-to-light artefacts in Figure 5(l). Similar errors are apparent with single-level tracking, Figure 5(n), and multi-level estimation, Figure 5(o), since they both attempt to estimate fine-scale lateral displacement. The “no lateral offset” assumption

is evidently incorrect too, since the axial displacement image (f) also exhibits a number of prominent error patches, resulting in characteristic dark-to-light artefacts in Figure 5(k).

The B-mode image in Figure 6(a) shows a breast fibroadenoma, which all the strain images in Figures 6(k)–(o) suggest is stiff. The RF signal at the bottom left of the B-mode image is dominated by noise, so the displacement and strain images in this region are unreliable. The axial displacement fields in Figures 6(f)–(j) exhibit discontinuities immediately above and below the fibroadenoma, which show up as bright bands in the strain images. These phenomena are most probably a result of slip between the adjacent tissue layers. Comparing the performances of the various displacement estimation methods, we come to the same conclusions as with the testis example. L3 lateral tracking, single-level tracking and the multi-level method all produce implausibly high frequency lateral displacement fields. Subjective assessment of the background noise level and the apparent lesion contrast suggests that the hybrid strain image (m) is the best of the five.

B. Finite element simulations

Each displacement estimation method was evaluated 100 times at each elevational offset in the range 0 mm to 0.4 mm. For each individual trial, there was a different RF signal realization (i.e. different Field II scatterer locations) and different random noise at a fixed SNR of 30 dB. After each test, the estimated axial displacement field was compared with the known ground truth and the average absolute point-wise difference d was recorded. The results in Figures 7(a,c,e) are presented in terms of the mean and standard deviation of d across the 100 trials. Figures 7(b,d,f) show the proportion of peak-hopping errors: these are of interest since they tend to produce visually discernible artefacts in strain images. An axial displacement estimate is counted as a peak-hopping error if it differs from the ground truth by half a wavelength or more. The discussion that follows concentrates on the main features of the results. Secondary effects, such as the role of downsampling and the relative merits of correlation and phase-based similarity metrics, are addressed in Appendix B.

Figures 7(a) and (b) confirm that, in the presence of signal decorrelation and noise, performance can suffer with increased amounts of lateral search. The “no lateral offset” method is the most accurate. This is not surprising, given that the ground truth lateral displacement field reaches a maximum of just 0.1 mm, which is approximately one third of the A-line separation. The performance of the hybrid method is comparable at elevational offsets below 0.4 mm, though at 0.4 mm it is evident that decorrelation-induced, incorrect lateral displacement estimates are affecting the axial tracking. With this transducer (not untypical of current mid-range hardware), L3 lateral tracking degrades performance significantly.

Figures 7(c) and (d) demonstrate the benefits of lateral motion compensation in all but the most contrived cases. The ground truth lateral displacement field reaches a maximum of 0.7 mm, which is 2.3 times the A-line separation. Without offsetting the L3 windows laterally, axial displacement estimation is inevitably compromised: the “no lateral offset” method produces inaccurate displacements even with highly correlated RF data (zero elevational offset). With this non-uniform compression, the single level tracking method becomes particularly brittle at 0.4 mm elevational offset: the large error bars are indicative of occasional catastrophic tracking failures.

The third simulation verifies the hybrid method’s resilience when the displacement field is discontinuous. Provided that at least one of the L1 windows falls within each discontinuous region, the quality-guided tracking algorithm is capable of recovering the displacement field within each region separately, without vulnerability to tracking failures across discontinuities [26]. Figure 8 shows examples of axial (a–c) and lateral (d–f) displacement

fields for this simulation. The ground truth in Figure 8(d), obtained from the finite element analysis, reveals lateral displacement of between one and two A-lines in the centre layer, and less than half an A-line in the top and bottom layers. This explains the “no lateral offset” method’s ability to recover the axial displacement in the top and bottom layers but not in the centre layer: see Figures 8(a) and (b). In contrast, the hybrid method performs well in all three layers (Figure 8(c)). The lateral displacement fields in Figures 8(e)–(f) reveal the expected noisy result with L3 lateral tracking compared with a low resolution approximation with the hybrid method. The single-level tracking and the multi-level method also produce noisy lateral displacement fields, like the one in Figure 8(e).

Figures 8(e) and (f) illustrate the trade-off between lateral displacement resolution and precision. While lateral displacement fields are normally slowly varying, favouring the low resolution, high precision approach taken by the hybrid method, the balance of the trade-off is less obvious when the true lateral displacement field has sharp discontinuities, as in Figure 8(d). At low resolution (Figure 8(f)), there will inevitably be errors in the vicinity of the slip planes. However, at high resolution (Figure 8(e)), the precision is compromised across the entire image. The quantitative results in Figure 7(e) suggest that the trade-off favours the hybrid method, with performance significantly better than L3 lateral tracking at elevational offsets of 0.3 and 0.4 mm, and comparable at 0 to 0.2 mm.

The performance of the hybrid method can be attributed to three key factors. First, compared with the single-level tracking strategy, the L1 and L2 searches of the hybrid method provide more accurate initial seeds to prime the fine-scale displacement tracking. Second, compared with the multi-level approach, the hybrid method incorporates a quality-guided tracking strategy to propagate estimates from one window to the next, which avoids inappropriate interpolation across discontinuities. Third, lateral displacement estimation is carefully balanced in the hybrid method. The lateral estimates provided by the L1 and L2 searches are relatively insensitive to noise and decorrelation by virtue of their low resolution, but are nevertheless sufficiently accurate to compensate for lateral motion. There is no lateral search at the finest resolution: this results in a faster and (for this transducer) more robust axial displacement estimation procedure. A state-of-the-art transducer, with superior lateral sampling density, is considered in Section IV-E.

Table II shows execution times for the displacement estimation methods. It should be noted that the preprocessing overhead is unlikely to be necessary in any commercial implementation, since the required analytic signals would most probably be available *a priori*.

C. Regression parameter for lateral estimates

Figure 9 explains our choice of 0.1 for the hybrid method’s regression parameter k , where k is defined in Equation (2). We re-ran the hybrid method on one of the simulated data sets, at several values of k . As seen from the results in Figure 9, smoothing lateral estimates does not have much impact on accuracy in general. Heavy smoothing can be detrimental when the lateral displacement field contains discontinuous regions caused by slip planes ($k = 0.5$), since heavily smoothed estimates are inaccurate around the region boundaries. Light smoothing ($k = 0.1$) reduces estimation errors compared with the no smoothing case ($k = 0$). Although the difference is small, the regression step has negligible computational cost and is therefore recommended as a worthwhile feature of the hybrid algorithm.

D. Search region and window size effects

The default parameters of the hybrid method, given in Table I, are designed to track a maximum strain of 5%. This limit depends predominantly on the size of the L1 search

region ($\lambda/10$ allows for up to 5% axial deformation in either direction). The strain limit can be modified by changing the size of the L1 search region and making a proportionate adjustment to the size of the L2 search region. Figure 10(a) shows the effects of such adjustments on the performance of the hybrid method. It is not surprising that reducing the search distance improves the algorithm's robustness (as reflected in the error bars) but limits the maximum strain that can be tracked. The 5% default setting was not chosen for optimal robustness, but rather to give a reasonable strain limit for typical clinical practice.

The window size parameters also affect the performance of the hybrid method. Figure 10(b) shows the effects of varying the L1 window size in the range 15% to 25% of the RF frame size, and making a proportionate adjustment to the L2 window size. Increasing the window size improves the algorithm's robustness at the expense of slower computation. Once again, the default 20% setting does not optimize robustness, but provides a reasonable compromise between robustness and speed.

Note that the L3 search range is implicitly limited to half the RF wavelength [11]. The effects of L3 window size were explored in [12]. Our default setting of eight RF cycles is rather small and therefore tolerant of relatively high strains, where larger windows would suffer significant intra-window strain and thus decorrelation.

E. Transducer element density

It is reasonable to ask whether low-level lateral tracking might be less error-prone with state-of-the-art transducers that boast superior element densities. Figure 11 shows results of the slip plane simulations with a 256-element transducer (0.15 mm element spacing). As expected, all techniques that attempt low-level lateral tracking show improved performance. However, the hybrid algorithm is still the marginal winner. There will, of course, become a point at which low-level lateral tracking becomes worthwhile. What this paper has shown, however, is that displacement estimation algorithms must take full account of the system's lateral resolution and limit their lateral searches accordingly. Attempting to track at too fine a resolution is not only wasteful of computation but also detrimental to performance.

V. CONCLUSIONS AND FURTHER WORK

The hybrid displacement estimation method combines multi-level and tracking strategies to achieve an appealing balance between accuracy, robustness, and speed. In the course of recovering the axial displacement field, lateral motion is also estimated but not to the same resolution. Instead, a sparse grid of lateral displacements is interpolated to provide subsample offsets for the fine grid of windows used at the highest resolution of axial search. This strategy has been shown to be relatively robust to signal decorrelation and fast enough for real-time implementation on commodity hardware.

The multiple high-level seeds and the quality-guided tracking framework enable the method to recover the sort of discontinuous displacement fields that are not uncommon in clinical practice. Future work will extend the technique to three-dimensional quasistatic strain imaging applications. Elevational focusing in commercial three-dimensional probes is generally no better than lateral focusing, suggesting that the strategy presented here for lateral motion compensation should be appropriate in the elevational direction as well.

Acknowledgments

This work was funded by Translation Award 081511/Z/06/Z from the Wellcome Trust. The authors would like to thank the anonymous referees, whose constructive comments helped to improve the original version of the manuscript.

APPENDIX A: IMPLEMENTATION OF THE MULTI-LEVEL ALGORITHM

The multi-level method in [31] features four levels that we shall refer to as L'3, L'2, L'1 and L'0. The first three levels, L'3, L'2, and L'1, operate on downsampled envelope data. The last level, L'0, considers the raw RF data, with no downsampling. The standard cross-correlation metric is employed at all levels. The experiments in Section IV used the parameters in Table III, which were selected to optimize the algorithm's performance.

At each level, poor displacement estimates are rejected based on a threshold of the cross-correlation metric, so that they do not affect the next level: instead, better estimates are interpolated from nearby areas of the image. However, the fixed threshold in [31] is unsuitable for our studies, since we have varying degrees of decorrelation and therefore require a variable threshold. Instead, we rejected the poorest 25% of the displacement estimates. The 25% level was chosen to optimize performance.

[31] describes a combination of cubic spline interpolation and least-squares line fitting, to smooth the displacement estimates after each level of processing. There is a parameter p , where $p = 0$ for natural cubic spline interpolation and $p = 1$ for least-squares straight line fitting. No particular value of p is recommended in [31]. We experimented with various values of p and found that $p = 0$ produced the best results for the experiments in Section IV.

TABLE III

PARAMETERS USED IN THE MULTI-LEVEL DISPLACEMENT ESTIMATION METHOD.

parameter (lateral $\times$ axial)	in vivo	simulations
L'3 image size (pixels)	128 $\times$ 161	128 $\times$ 429
number of L'3 windows	3 $\times$ 3	3 $\times$ 3
L'3 window size (mm ²)	7.8 $\times$ 7.8	7.8 $\times$ 7.8
L'3 search region (top, mm ²)	9.0 $\times$ 9.0	9.0 $\times$ 9.0
L'3 search region (middle, mm ²)	10.2 $\times$ 10.2	10.2 $\times$ 10.2
L'3 search region (bottom, mm ²)	11.4 $\times$ 11.4	11.4 $\times$ 11.4
L'2 image size (pixels)	128 $\times$ 322	128 $\times$ 858
number of L'2 windows	7 $\times$ 9	7 $\times$ 9
L'2 window size (mm ²)	3.9 $\times$ 3.9	3.9 $\times$ 3.9
L'2 search region (mm ²)	5.1 $\times$ 5.1	5.1 $\times$ 5.1
L'1 image size (pixels)	128 $\times$ 645	128 $\times$ 1717
number of L'1 windows	21 $\times$ 27	21 $\times$ 27
L'1 window size (mm ²)	1.8 $\times$ 1.8	1.8 $\times$ 1.8
L'1 search region (mm ²)	2.4 $\times$ 2.4	2.4 $\times$ 2.4
L'0 image size (pixels)	128 $\times$ 1319	128 $\times$ 3464
number of L'0 windows	64 $\times$ 35	64 $\times$ 47
L'0 window size (mm ²)	0.9 $\times$ 0.8	0.9 $\times$ 1.0
L'0 search region (mm ²)	1.5 $\times$ 0.9	1.5 $\times$ 1.1

APPENDIX B: SIMILARITY METRICS AND DOWNSAMPLING

The multi-level method was originally presented using correlation at all levels [31], and we have duly performed a faithful comparison with this method as published. Nevertheless, it is

reasonable to ask how much of any performance disparity can be attributed to different similarity metrics at the lowest level. We therefore implemented a version of the multi-level method that uses WPS [12] at the finest level, allowing a more direct comparison with the hybrid approach. To further level the playing field, we also evaluated the hybrid method with no downsampling, as is the case for the multi-level method [31].

The results in Figure 12 reveal little difference between the two variants of the multi-level method. This suggests that the hybrid method owes its superior performance to high level strategy and not low level nuance. The hybrid method without downsampling performs at least as well as all the other methods in all tests, while remaining competitive in terms of speed (Table IV). Nevertheless, we would suggest that the significant speed advantages of downsampling (Table II) more than compensate for any small performance deficit. Phase-based methods are particularly robust to downsampling since they operate on baseband analytic signals. In contrast, direct downsampling of raw RF data may violate the Nyquist sampling criterion [31].

TABLE IV

EXECUTION TIMES OF THE DISPLACEMENT ESTIMATION METHODS WITHOUT DOWNSAMPLING THE ANALYTIC RF DATA, RUNNING SINGLE-THREADED ON A 2.4 GHZ INTEL CORE 6600 PROCESSOR. THE POSTPROCESSING STAGE, COMMON TO ALL METHODS, INVOLVES THE COMPUTATION OF THE AXIAL STRAIN IMAGE FROM THE DISPLACEMENT FIELD.

method	procedure	time (ms)	
		in vivo	simulations
hybrid	preproc.	33	90
	L1 and L2	24	85
	regression	< 0.1	< 0.1
	L3	24	60
	postproc.	10	16
	total	91	251
multi-level	preproc.	13	34
	L'3– L'0	65	178
	postproc.	10	16
	total	88	228
multi-level WPS	preproc.	37	98
	L'3– L'0	66	135
	postproc.	10	16
	total	113	249

REFERENCES

- [1]. Ophir J, Céspedes I, Ponnekanti H, Yazdi Y, Li X. Elastography: a quantitative method for imaging the elasticity of biological tissues. *Ultrasonic Imaging*. 1991; vol. 13:111–134. [PubMed: 1858217]
- [2]. Sugimoto T, Ueha S, Itoh K. Tissue hardness measurement using the radiation force of focused ultrasound. *Proceedings of IEEE Ultrasonics Symposium 1990*. Dec.1990 vol. 3:1377–1380.

- [3]. Muthupillai R, Lomas DJ, Rossman P, Greenleaf JF, Manduca A, Ehman RL. Magnetic resonance elastography by direct visualization of propagating acoustic strain waves. *Science*. 1995; vol. 269(no. 5232):1854–1857. [PubMed: 7569924]
- [4]. Dickinson RJ, Hill CR. Measurement of soft tissue motion using correlation between A-scans. *Ultrasound in Medicine and Biology*. 1982; vol. 8(no. 3):263–271. [PubMed: 7101574]
- [5]. Burnside ES, Hall TJ, Sommer AM, Hesley GK, Sisney GA, Svensson WE, Fine JP, Jiang J, Hangiandreou NJ. Differentiating benign from malignant solid breast masses with US strain imaging. *Radiology*. Nov; 2007 vol. 245(no. 2):401–410. [PubMed: 17940302]
- [6]. Bilgen M, Insana MF. Deformation models and correlation analysis in elastography. *The Journal of the Acoustical Society of America*. 1996; vol. 99(no. 5):3212–3224. [PubMed: 8642127]
- [7]. Bohs L, Trahey G. A novel method for angle independent ultrasonic imaging of blood flow and tissue motion. *IEEE Transactions on Biomedical Engineering*. 1991; vol. 38(no. 3):280–286. [PubMed: 2066142]
- [8]. Chaturvedi P, Insana MF, Hall TJ. 2-D companding for noise reduction in strain imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 1998; vol. 45:179–191.
- [9]. Yeung F, Levinson SF, Parker KJ. Multilevel and motion model-based ultrasonic speckle tracking algorithms. *Ultrasound in Medicine and Biology*. 1998; vol. 24:427–441. [PubMed: 9587997]
- [10]. O'Donnell M, Skovoroda AR, Shapo BM, Emelianov SY. Internal displacement and strain imaging using ultrasonic speckle tracking. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 1994; vol. 41:314–325.
- [11]. Pesavento A, Perrey C, Krueger M, Ermert H. A time efficient and accurate strain estimation concept for ultrasonic elastography using iterative phase zero estimation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. Sep; 1999 vol. 46(no. 5):1057–1067.
- [12]. Lindop JE, Treece GM, Gee AH, Prager RW. Phase-based ultrasonic deformation estimation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. Jan; 2008 vol. 55(no. 1):94–111.
- [13]. Alam SK, Ophir J, Konofagou EE. An adaptive strain estimator for elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 1998; vol. 45(no. 2):461–472.
- [14]. Céspedes I, Huang Y, Ophir J, Spratt S. Methods for estimation of subsample time delays of digitized echo signals. *Ultrasonic Imaging*. 1995; vol. 17:142–171.
- [15]. Ebbini ES. Phase-coupled two-dimensional speckle tracking algorithm. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 2006; vol. 53(no. 5):972–990.
- [16]. Jiang J, Hall TJ. A parallelizable real-time motion tracking algorithm with applications to ultrasonic strain imaging. *Physics in Medicine and Biology*. 2007; vol. 52:3773–3790. [PubMed: 17664576]
- [17]. Konofagou E, Ophir J. A new elastographic method for estimation and imaging of lateral displacements, lateral strains, corrected axial strains and Poisson's ratios in tissues. *Ultrasound in Medicine and Biology*. 1998; vol. 24(no. 8):1183–1199. [PubMed: 9833588]
- [18]. Lubinski AM, Emelianov SY, O'Donnell M. Speckle tracking methods for ultrasonic elasticity imaging using short time correlation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 1999; vol. 46(no. 1):82–96.
- [19]. Pellot-Barakat C, Frouin F, Insana MF, Herment A. Ultrasound elastography based on multiscale estimations of regularized displacement fields. *IEEE Transactions on Medical Imaging*. 2004; vol. 23:153–163. [PubMed: 14964561]
- [20]. Zahiri-Azar R, Salcudean SE. Motion estimation in ultrasound images using time domain cross correlation with prior estimates. *IEEE Transactions on Biomedical Engineering*. 2006; vol. 53:1990–2000. [PubMed: 17019863]
- [21]. Zhu Y, Hall TJ. A modified block matching method for real-time freehand strain imaging. *Ultrasonic Imaging*. 2002; vol. 24:161–176. [PubMed: 12503771]
- [22]. Viola F, Walker WF. A spline-based algorithm for continuous time-delay estimation using sampled data. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 2005; vol. 52(no. 1):80–93.

- [23]. —. A comparison between spline-based and phase-domain time-delay estimators. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 2006; vol. 53(no. 3):515–517.
- [24]. Pinton GF, Trahey GE. Continuous delay estimation with polynomial splines. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 2006; vol. 53(no. 11):2026–2035.
- [25]. Maurice RL, Bertrand M. Lagrangian speckle model and tissue-motion estimation–theory. *IEEE Transactions on Medical Imaging*. 1999; vol. 18(no. 7):593–603. [PubMed: 10504093]
- [26]. Chen L, Treece GM, Lindop JE, Gee AH, Prager RW. A quality-guided displacement tracking algorithm for ultrasonic elasticity imaging. *Medical Image Analysis*. Apr; 2009 vol. 13(no. 2): 286–296. [PubMed: 19081285]
- [27]. Treece, GM.; Lindop, JE.; Gee, AH.; Prager, RW. Efficient elimination of dropouts in displacement tracking. *Proceedings of the Fifth International Conference on the Ultrasonic Measurement and Imaging of Tissue Elasticity*; Snowbird, Utah. October 2006; p. 68
- [28]. —. Freehand ultrasound elastography with a 3-D probe. *Ultrasound in Medicine and Biology*. 2008; vol. 34(no. 3):463–474.
- [29]. Bai J, Ding C, Fan Y. A multi-scale algorithm for ultrasonic strain reconstruction under moderate compression. *Ultrasonics*. 1999; vol. 37:511–519.
- [30]. Chen H, Shi H, Varghese T. Improvement of elastographic displacement estimation using a two-step cross-correlation method. *Ultrasound in Medicine and Biology*. 2007; vol. 33(no. 1):48–56. [PubMed: 17189046]
- [31]. Shi H, Varghese T. Two-dimensional multi-level strain estimation for discontinuous tissue. *Physics in Medicine and Biology*. 2007; vol. 52:389–401. [PubMed: 17202622]
- [32]. Walker F, Trahey E. A fundamental limit on delay estimation using partially correlated speckle signals. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 1995; vol. 42:301–308.
- [33]. Liebgott H, Fromageau J, Wilhelm JE, Vray D, Delachartre P. Beamforming scheme for 2D displacement estimation in ultrasound imaging. *EURASIP Journal on Applied Signal Processing*. 2005; vol. 8:1212–1220.
- [34]. Céspedes I, Ophir J, Alam SK. The combined effect of signal decorrelation and random noise on the variance of time delay estimation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*. 1997; vol. 44(no. 1):220–225.
- [35]. Varga, RS. *Matrix Iterative Analysis*. Springer; 2000.
- [36]. Press, WH.; Teukolsky, SA.; Vetterling, WT.; Flannery, BP. *Numerical recipes in C: The art of scientific computing*. 2nd ed. Cambridge University Press; 2002.
- [37]. Briggs, WL.; Henson, VE.; McCormick, SF. *A multigrid tutorial*. Society for Industrial and Applied Mathematics; 2000.
- [38]. Treece, GM.; Lindop, JE.; Gee, AH.; Prager, RW. Uniform precision ultrasound strain imaging. Vol. 624. University of Cambridge; 2009. Tech. Rep. CUED/FINFENG/TR
- [39]. Jensen, JA. Field: A program for simulating ultrasound systems. *Proceedings of the 10th Nordic-Baltic Conference on Biomedical Imaging*; Tampere; 1996. p. 351–353.
- [40]. Lindop JE, Treece GM, Gee AH, Prager RW. An intelligent interface for freehand strain imaging. *Ultrasound in Medicine and Biology*. Jul; 2008 vol. 34(no. 7):1117–1128. [PubMed: 18440122]

Fig. 1.
Overview of the hybrid displacement estimation method.

Fig. 2.

The multi-resolution search used at L1 and L2. (a) At the lowest resolution, only every fourth sample is considered in the pre- and post-deformation frames, as indicated by the red and blue squares. The centre of the pre-deformation window (red square, top) is aligned with sixteen putative locations in the post-deformation search region (red squares, bottom), and the correlation coefficient is calculated at each alignment using just the red and blue samples. The location of the best match (gold square, bottom) is recorded. (b) At the next resolution, every other sample contributes to the processing. Three candidate locations centred around the gold square are considered, and the location of the best match (green square, bottom) is recorded. (c) At the finest resolution, three candidate locations centred around the green square are considered, and the correlation coefficient is now calculated without skipping any samples. The location of the best match (magenta square, bottom) completes the search.

Fig. 3.

Finite element simulations. The domain is a 140 mm diameter by 60 mm high cylinder filled with material of Young's modulus 10 kPa. All simulated materials (background and inclusion) have a Poisson's ratio of 0.495. The dimensions of the probe face are 40 mm (lateral) by 10 mm (elevation). The boundary conditions are no slip between the probe face and the top surface of the cylinder, frictionless slip at the bottom surface, with all other surfaces unconstrained. (a) A 15 mm diameter spherical inclusion with Young's modulus 40 kPa is embedded in the material. In the first simulation, the material is compressed by the ultrasound probe translating 0.6 mm in the y direction. In the second simulation, there is an additional 1° rotation around the z -axis passing through the centre of the probe face. (b) In the third simulation, there are three layers of material, each of Young's modulus 10 kPa, separated by frictionless slip planes. A 5 mm cylindrical ring with Young's modulus 0.3 kPa surrounds the three slabs, to ensure that they do not slide off each other completely. A 0.6 mm displacement along the y -axis is applied to the probe.

Fig. 4.

Finite element/Field II simulations. (a) B-mode image. Axial strain images obtained using the hybrid method for (b) the first simulation with normal probe compression, (c) the second simulation with tilted probe compression and (d) the third simulation with slip planes: dark is hard and bright is soft. Note that the axial strain is normalized by a spatially varying estimate of the stress [40] and then displayed on a nonlinear greyscale. For this reason, strain images in this paper are not accompanied by the usual intensity-to-strain colour bar.

Fig. 5.

(a) B-mode image of the testis. The first column (b–e), the second column (f–j) and the third column (k–o) show respectively the lateral displacement field, the axial displacement field and the axial strain image recovered by the various methods. Results for each method are aligned in rows, from the no lateral offset method (top) to the multi-level method (bottom). In the axial strain images (k–o), dark is hard and bright is soft. Arrows in the axial displacement fields and the strain images indicate plausible peak-hopping artefacts.

Fig. 6.

(a) B-mode image of the breast with stiff fibroadenoma (indicated). The first column (b–e), the second column (f–j) and the third column (k–o) show respectively the lateral displacement field, the axial displacement field and the axial strain image recovered by the various methods. Results for each method are aligned in rows, from the no lateral offset method (top) to the multi-level method (bottom). In the axial strain images (k–o), dark is hard and bright is soft. Arrows in the axial displacement fields and the strain images indicate plausible peak-hopping artefacts.

Fig. 7.

Results for the finite element simulations. (a,b) Normal compression with stiff inclusion. (c,d) Tilted compression with stiff inclusion. (e,f) Normal compression with slip planes. (a,c,e) show errors calculated from the average absolute point-wise difference between the estimated axial displacement field and the ground truth from the finite element analysis. (b,d,f) show the percentage of peak-hopping errors. A displacement estimate is counted as peak-hopping if it differs from the ground truth by half a wavelength or more. The ± 1 standard deviation error bars are based on 100 repetitions.

Fig. 8.

Displacement fields for the slip-plane simulation at zero elevational offset. (a) Ground truth axial displacement field. The axial displacement field obtained by (b) the “no lateral offset” method and (c) the hybrid method, with no additive noise. (d) Ground truth lateral displacement field. The lateral displacement field obtained by (e) L3 lateral tracking and (f) the hybrid method.

Fig. 9.

The role of the regression parameter k in the hybrid method. Simulation results for the slip planes at several values of k . $k=0$ represents no smoothing, $k=0.1$ light smoothing and $k=0.5$ heavy smoothing. The error is obtained in the same way as in Figure 7(e).

Fig. 10.

The effect of the search region (a) and the window size (b) parameters of the hybrid method. In (a), the percentage indicates the maximum strain that can be tracked: this depends directly on the L1 search region. In (b), the percentage indicates the ratio of the L1 window size to the RF frame size. The results are based on the tilted compression simulation: results for the other two simulations show very similar trends. The error is obtained in the same way as in Figure 7(c).

Fig. 11. Simulation results for the slip planes and a 256-element transducer (i.e. double the element density of the transducers considered elsewhere in this paper). Compare with Figure 7 (e) and (f).

Fig. 12.

Results for the finite element simulations obtained by the multi-level correlation [31], the multi-level WPS and the hybrid methods. Two versions of the hybrid method are included, one with downsampled RF data as in Figure 7, the other without downsampling. The two multi-level methods do not use downsampling. (a,b) Normal compression with stiff inclusion. (c,d) Tilted compression with stiff inclusion. (e,f) Normal compression with slip planes. The errors are obtained in the same way as in Figure 7.

TABLE I

PARAMETERS USED IN THE HYBRID DISPLACEMENT ESTIMATION METHOD. THE L1 AND L2 WINDOWS AND SEARCH REGIONS ARE SQUARE. THEIR SIZES DEPEND ON l , THE LATERAL WIDTH OF THE IMAGE. THE DOWNSAMPLING RATE S_y MUST SATISFY THE NYQUIST SAMPLING LIMIT OF THE BASEBAND ANALYTIC RF SIGNAL: INCREASING S_y MAKES THE ALGORITHM RUN FASTER. THE LATERAL DIMENSIONS OF L3 WINDOWS ARE CHOSEN TO MATCH, AS CLOSELY AS POSSIBLE, THE AXIAL DIMENSION OF EIGHT RF CYCLES.

parameter	formula	in vivo	simulations
downsampling rate S_y	–	3	6
image size (lat. $\times$ ax. mm ²)	–	38.4×25.0	38.4×40.0
number of L1 windows (lat. $\times$ ax.)	–	3×3	3×3
L1 window size (mm ²)	$l/5$	7.8×7.8	7.8×7.8
L1 search region (top, mm ²)	$l/5 + l/30$	9.0×9.0	9.0×9.0
L1 search region (middle, mm ²)	$l/5 + l/15$	10.2×10.2	10.2×10.2
L1 search region (bottom, mm ²)	$l/5 + l/10$	11.4×11.4	11.4×11.4
L1 search skip factor (samples)	–	4	4
number of L2 windows (lat. $\times$ ax.)	–	11×7	11×11
L2 window size (mm ²)	$l/15$	2.7×2.7	2.7×2.7
L2 search region (mm ²)	$l/15 + l/50$	3.3×3.3	3.3×3.3
L2 search skip factor (samples)	–	2	2
smoothing weight k in Equation (2)	–	0:1	0:1
number of L3 windows (lat. $\times$ ax.)	–	64×35	64×47
L3 window lateral size (mm)	–	0.9	0.9
L3 window axial size (mm)	8 RF cycles	0.8	1.0

TABLE II

EXECUTION TIMES OF THE DISPLACEMENT ESTIMATION METHODS, RUNNING SINGLE-THREADED ON A 2.4 GHZ INTEL CORE 6600 PROCESSOR. THE POSTPROCESSING STAGE, COMMON TO ALL METHODS, INVOLVES THE COMPUTATION OF THE AXIAL STRAIN IMAGE FROM THE DISPLACEMENT FIELD.

method	procedure	time (ms)	
		in vivo	simulations
hybrid	preproc.	11	14
	L1 and L2	4	4
	regression	< 0.1	< 0.1
	L3	10	14
	postproc.	10	16
	total	35	48
single-level tracking	total	37	44
multi-level	total	88	228