
**A HYBRID DISPLACEMENT
ESTIMATION METHOD FOR
ULTRASONIC ELASTICITY IMAGING**

L. Chen, R. J. Housden, G. M. Treece,
A. H. Gee and R. W. Prager

CUED/F-INFENG/TR 615

November 2008

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

Email: lc420/rjh80/gmt11/ahg/rwp @eng.cam.ac.uk

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 all freehand, quasistatic, ultrasonic strain imaging systems. In this paper, we present a novel estimation method that combines elements of multi-level correlation and phase-zero search to achieve the noise tolerance of the former and the speed of the latter. Given typical clinical B-scans, the hybrid method can generate more than 25 strain images per second on commodity hardware. 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 demonstrate its superior accuracy compared with previously published alternatives.

1 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 [22] to internal ultrasonic radiation force [26], while the latter might involve imaging modalities as diverse as MRI [20] and ultrasound [12].

Within this broader context, the subject of this paper is freehand quasistatic axial strain imaging [22], which has shown great promise in clinical trials [6] and has recently been commercialized. 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 B-scan 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.

At the heart of every quasistatic strain imaging system is an algorithm for estimating tissue deformation from one B-scan 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 [3, 12, 22], sum of absolute differences [4, 9], sum of squared differences [30], correlation phase [21, 24] and phase separation [18]. Windows are typically aligned to subsample precision [1, 7, 13, 15, 16, 18, 19, 23, 24, 30, 31, 32]. Matching might be carried out by independent

exhaustive searches at each window [12], or some inter-window continuity could be imposed, either by minimizing a global cost function that penalises discontinuous displacements [23, 30], or by tracking the displacements from one window to the next [11, 15, 18, 27, 28, 31, 32]. The search at each window might be fully two-dimensional [13], or there may be separate, one-dimensional axial and lateral searches, perhaps with some iteration between the two [16, 28]. There could be a single, high resolution matching process [7, 18, 19, 24, 28], or a multi-level approach with a coarse matching stage initializing subsequent stages at progressively higher resolutions [2, 9, 10, 23, 25, 30].

A well-designed displacement estimation algorithm aims to strike an appropriate balance between accuracy, robustness to noise, 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 [29]. Noise tolerance 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 [21], or zero phase separation [18], 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 [24]. The displacement will not normally change by more than half a wavelength from one window to the next, so such a tracking scheme is feasible provided it is carefully seeded [11, 28].

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 [19, 28]. In common with all tracking approaches, it is vulnerable to catastrophic failure through error propagation. Errors might be present from the outset if the 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 [11, 18, 27, 28]. Alternatively, *post hoc* continuity checks orthogonal to the tracking direction can be successful in suppressing error propagation [15, 32]. 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 [11]. But everything depends on the correctness of the initial seeds, which is difficult to guarantee in high-noise scenarios.

Compared with single-level tracking, multi-level strategies are intrinsically more robust to noise 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. 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.

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, 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 2. Section 3 contains qualitative *in vivo* and quantitative simulation comparisons of the hybrid technique with leading multi-level and single-level tracking alternatives, with attention to speed, accuracy, robustness to noise and resilience in the presence of slip planes. Finally, in Section 4 we present our conclusions and some suggestions for further work.

2 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 magnitude 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 [18], again with quality-guided tracking. Lateral displacement estimates are not refined at L3.

2.1 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 pass-band, 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 magnitude data for L1 and L2, while the phase ϕ is required by the WPS algorithm at L3.

2.2 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 d_x is the lateral displacement, d_y is the axial displacement, a_1 and a_2 are the magnitude images 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).

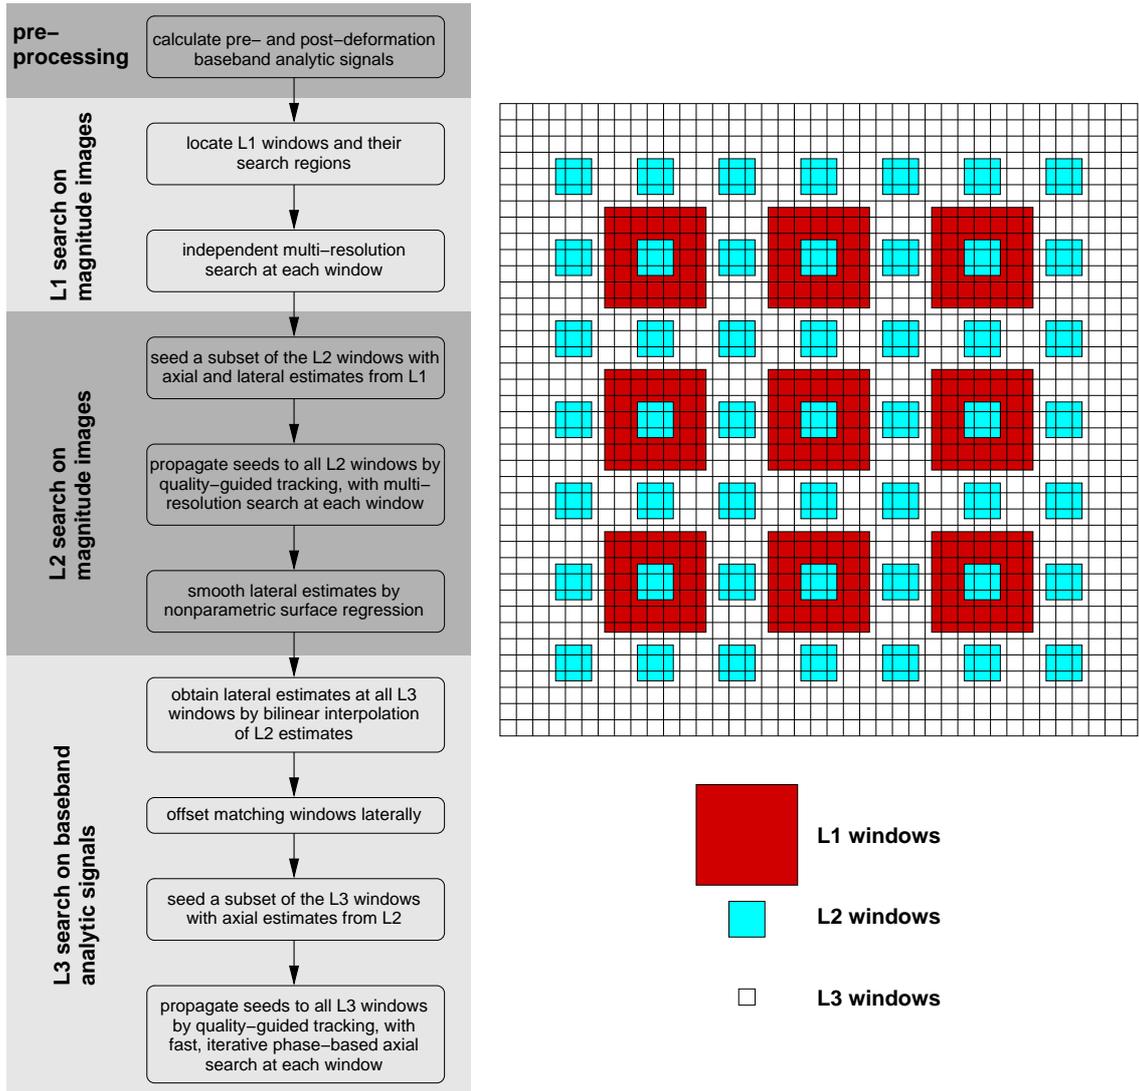


Figure 1: Overview of the hybrid displacement estimation method.

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 3. 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.

2.3 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 is the last stage of lateral displacement estimation, and our experience is that there should be at least five L2 windows in either direction to give a sufficiently dense sampling of the lateral displacement field.

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.

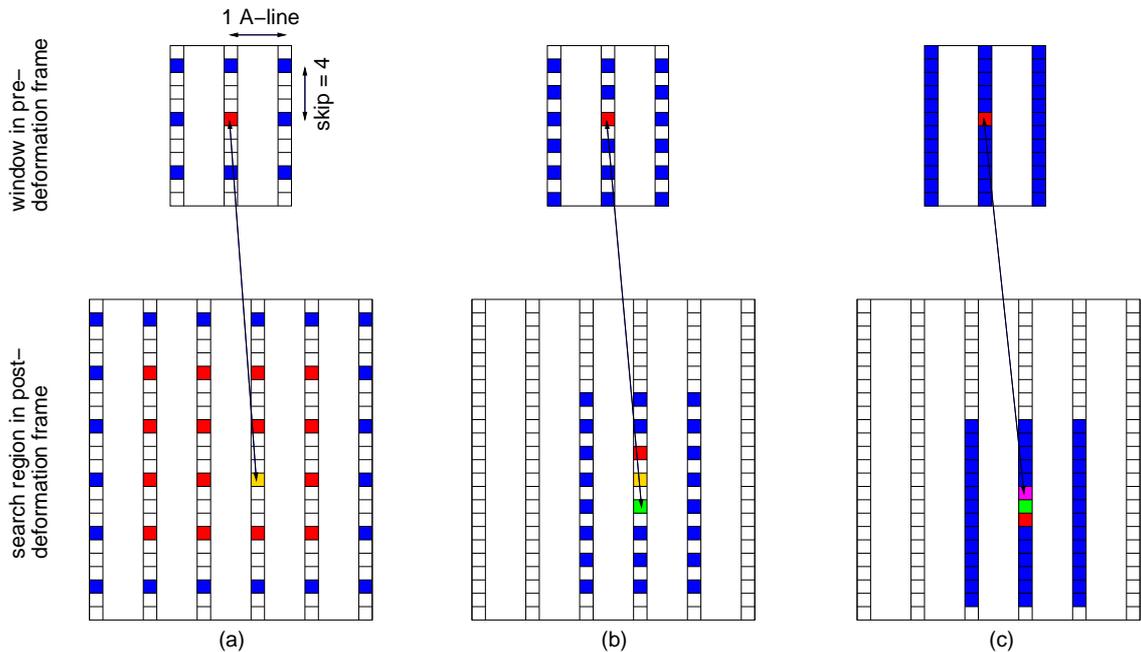


Figure 2: The multi-resolution search used at L1 and L2. (a) At the lowest resolution, only every fourth sample is considered in the pre- and post-deformation frames, as indicated by the red and blue squares. The centre of the pre-deformation window (red square, top) is aligned with sixteen putative 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.

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 [11]. 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 [11].

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 is not even possible with conventional beamforming methods, since the lateral resolution, dictated by the transducer element spacing, is poor¹. 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. 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 is 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 3.

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

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

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 $C/(1 - C)$ [8], 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 exist to solve this standard regression problem [5, 23, 30]: we use the multigrid technique [5].

2.4 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) [18] 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 [24]. The subsample lateral displacement of each L3 window is obtained by bilinear interpolation² 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 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 integer 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 [18].

Although Equation (3) is expressed in terms of integer axial and lateral displacements, $\Delta\phi$ can be estimated at subsample shifts by bilinear interpolation of the four neighbouring integer alignments. This is an important and necessary detail, since the lateral displacement interpolated

¹For higher precision lateral strain imaging, novel beamforming techniques can potentially increase the lateral resolution to levels comparable with the axial resolution [17].

²Cubic interpolation was also investigated, but it comes at greater cost and did not improve the results in Section 3.

parameter	formula	in vivo	simulations
downsampling rate S_y	—	3	6
downsampled image size (A-lines \times samples)	—	128×510	128×573
number of L1 windows (lateral \times axial)	—	3×3	3×3
L1 window size (A-lines \times samples)	$l/5$	26×102	26×115
L1 search region (top, A-lines \times samples)	$l/5 + l/30$	30×120	30×135
L1 search region (middle, A-lines \times samples)	$l/5 + l/15$	34×136	34×153
L1 search region (bottom, A-lines \times samples)	$l/5 + l/10$	38×154	38×173
L1 search skip factor (samples)	—	4	4
number of L2 windows (lateral \times axial)	—	7×11	7×11
L2 window size (A-lines \times samples)	$l/15$	9×34	9×38
L2 search region (A-lines \times samples)	$l/15 + l/50$	11×44	11×50
L2 search skip factor (samples)	—	2	2
smoothing weight k in Equation (2)	—	0.1	0.1
number of L3 windows (lateral \times axial)	—	64×42	64×47
L3 window axial size (samples)	8 RF cycles	14	14
L3 window lateral size (A-lines)	—	3	3

Table 1: Parameters used in the hybrid displacement estimation method. The L1 and L2 window sizes and search ranges depend on l , which is the number of A-lines for lateral dimensions, or the number of (downsampled) samples for axial dimensions. 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 L3 windows were approximately square, with lateral dimensions chosen to match, as close as possible, the axial dimension of eight RF cycles.

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 [18, 24]. Next, a new value of $\Delta\phi$ is calculated at this new axial displacement, and the process is repeated. This iterative procedure is fundamental to all phase-based alignment techniques: it converges rapidly to the zero-phase displacement [18, 24]. Finally, the axial displacement estimates are propagated to the remaining L3 windows using the same quality-guided tracking approach as in L2 [11], with WPS-based axial refinement at each window.

3 Results and discussion

In this section, several displacement estimation methods are compared qualitatively *in vivo* and quantitatively using finite element simulations. The methods under consideration are:

hybrid: the technique described in Section 2, using the parameters in Table 1

no lateral offset: as above, but without offsetting the L3 windows by the lateral offsets found at L2

L3 lateral tracking: the hybrid method with fine scale lateral displacement estimation 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 axial displacement is then re-estimated at this new lateral offset. The new axial and lateral estimates are propagated to neighbouring windows within the quality-guided tracking framework.

Stradwin: the single-level tracking strategy implemented in the free-to-download Stradwin system [28]

multi-level: a recently published multi-level approach [25] with no tracking. Full details of our implementation can be found in Appendix A.

3.1 Qualitative *in vivo* experiments

Two *in vivo* scans were performed using a T3000 (Terason Ultrasound, Massachusetts, USA) ultrasound system with a 7.75 MHz linear array transducer. The RF sampling frequency was 40 MHz and each frame comprised 128 A-lines. These scans serve to illustrate pertinent features of the various displacement tracking methods and their ability to perform *in vivo*. A more rigorous, quantitative comparison of the methods follows in Section 3.2.

The first scan of a human thyroid captures a benign thyroid nodule, indicated by the arrow in Figure 3(a). The axial strain images in Figures 3(k)–(o) suggest that the nodule is of similar stiffness to the surrounding tissue. Comparing the lateral displacement images in Figures 3(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 3(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 error patches in Figure 3(g) that, when differentiated, produce the dark-to-light artefacts in Figure 3(l). Similar errors are apparent in the Stradwin and multi-level results, 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 3(k).

The B-scan in Figure 4(a) shows a breast fibroadenoma, which all the strain images in Figures 4(k)–(o) suggest is stiff. The RF signal at the bottom left of the B-scan is dominated by noise, so the displacement and strain images in this region are unreliable. The axial displacement fields in Figures 4(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 thyroid example. L3 lateral tracking, Stradwin and the multi-level method all produce implausibly high frequency lateral displacement fields. Subjective assessment of the background noise level and apparent lesion contrast suggests that the hybrid strain image (m) is the best of the five.

3.2 Quantitative finite element simulations

Each of the three finite element simulations in Figure 5 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 B-scan 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 [14]. The simulation parameters were not untypical of real-world scanning scenarios: 128 A-lines spanning 40 mm laterally, 40 mm axial scanning depth, 6.5 MHz centre frequency and an RF sampling rate of 66.7 MHz. The scatterer density was identical for all materials, rendering the inclusion and slip planes invisible in the B-mode images, as shown in Figure 6(a). Examples of axial strain images, obtained using the hybrid method, can be found in Figures 6(b)–(d).

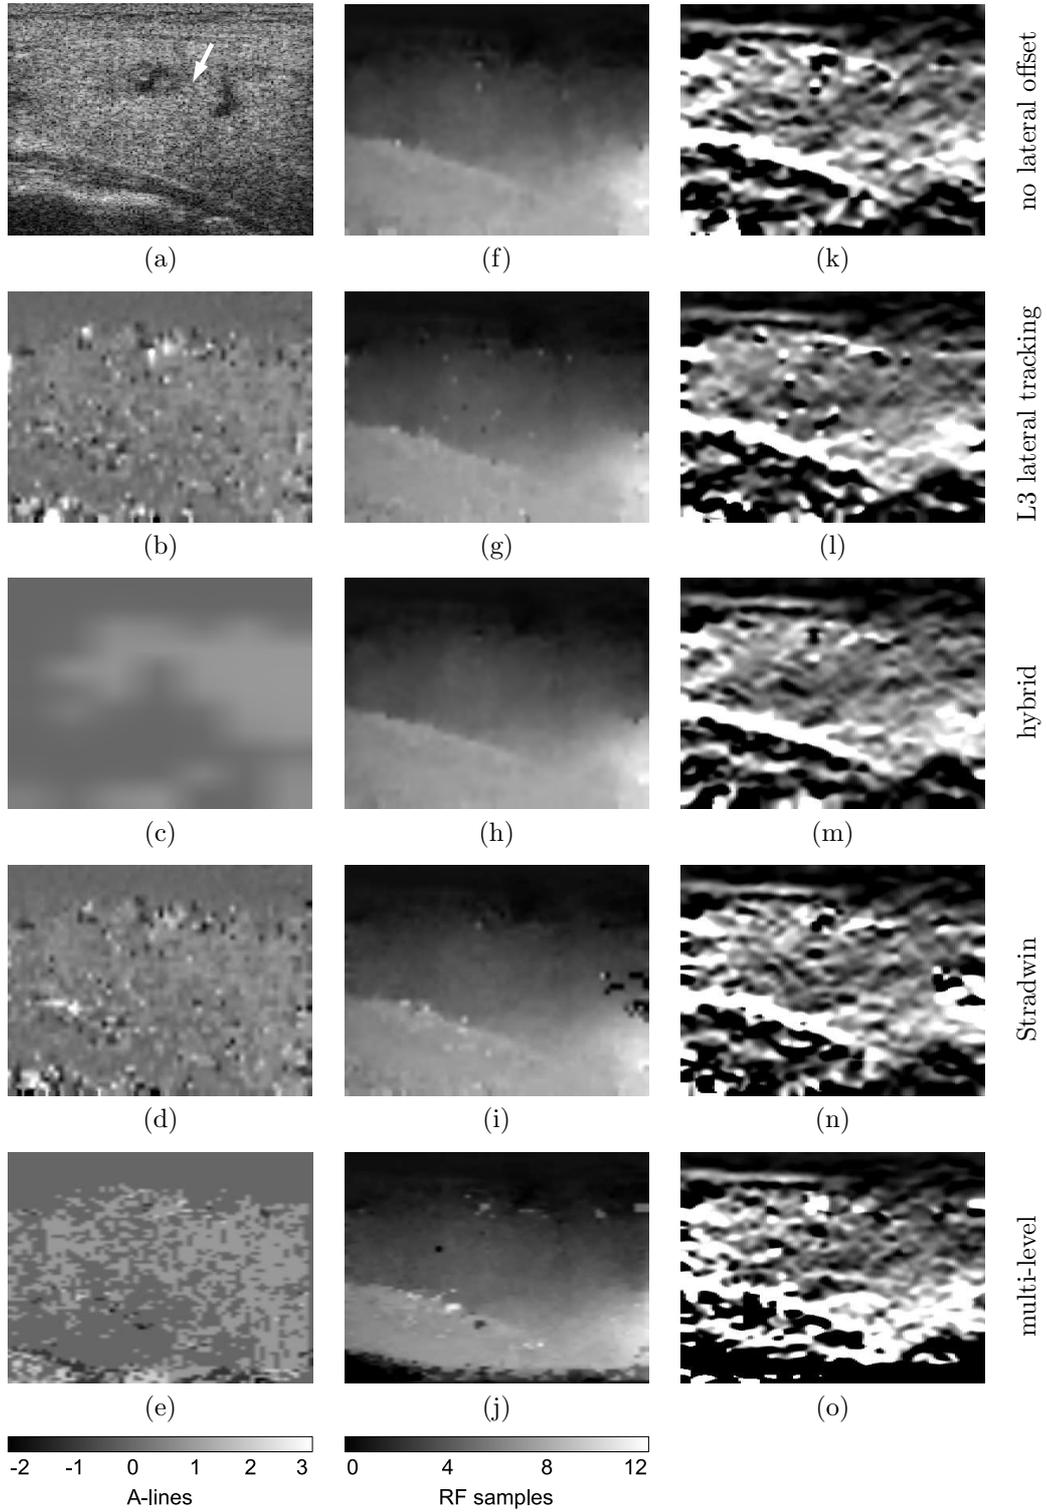


Figure 3: (a) B-mode image of the thyroid, with a benign nodule indicated by the arrow. Lateral (b–e) and axial (f–j) displacement fields recovered by the various methods. In the axial strain images (k–o), dark is hard and bright is soft.

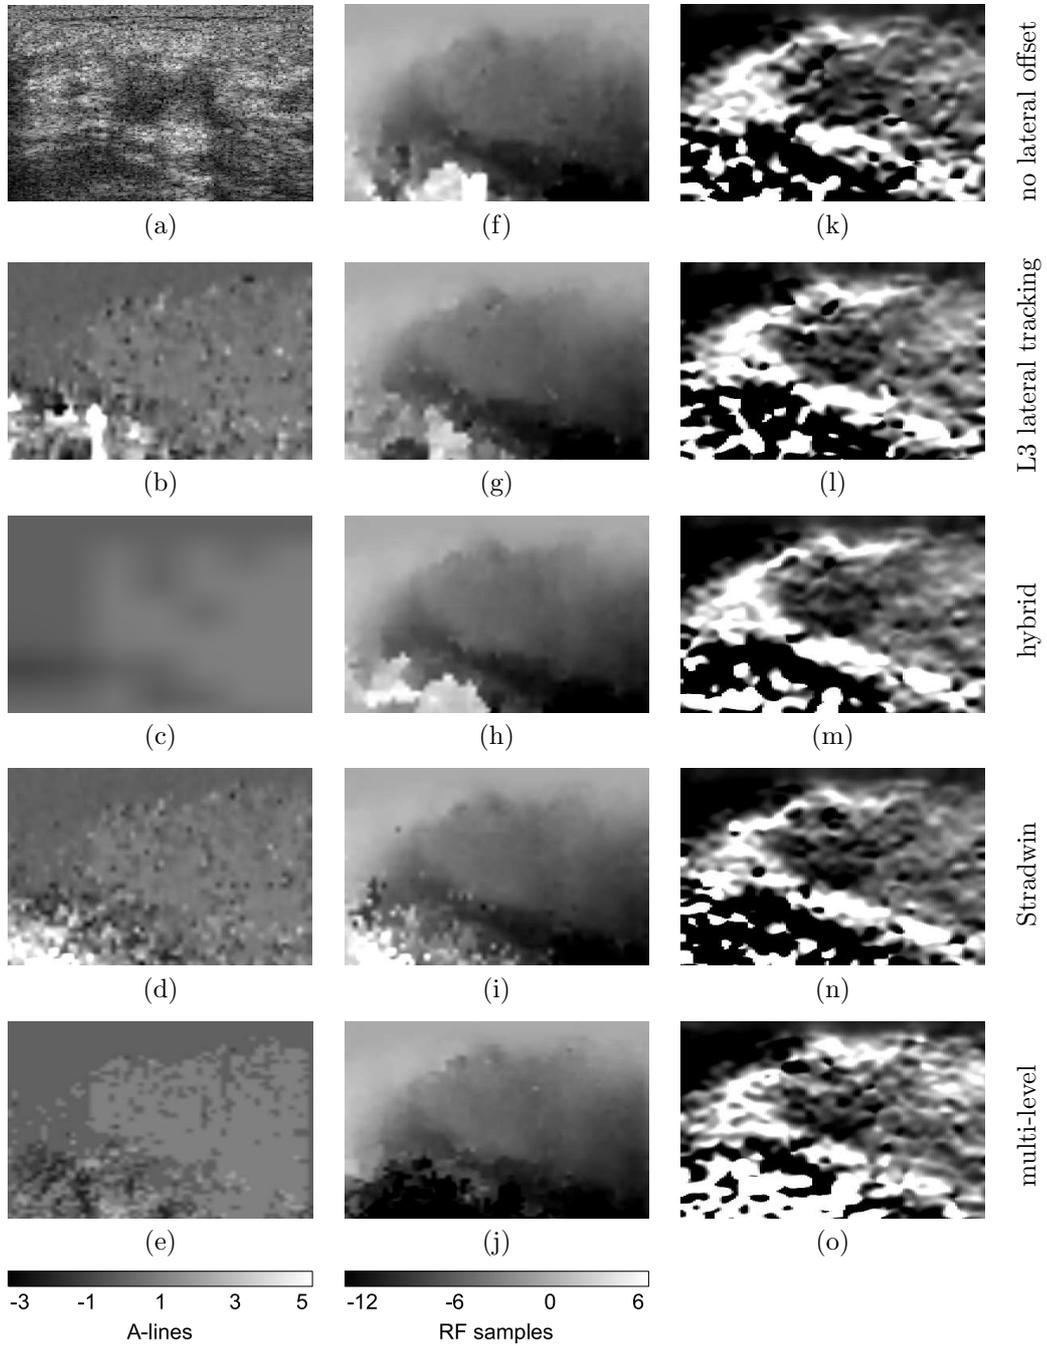


Figure 4: (a) B-mode image of the breast with stiff fibroadenoma. Lateral (b–e) and axial (f–j) displacement fields recovered by the various methods. In the axial strain images (k–o), dark is hard and bright is soft.

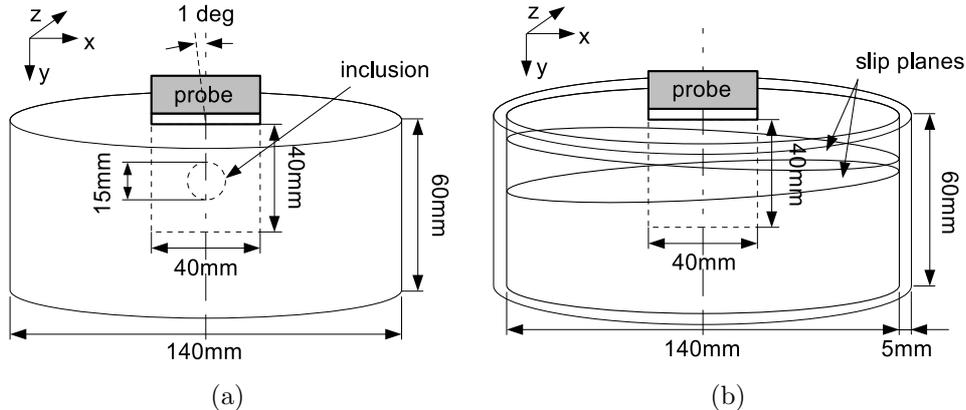


Figure 5: Finite element simulations. The domain is a 140 mm diameter by 60 mm high cylinder filled with material of Young’s modulus 10 kPa. All simulated materials (background and inclusion) have a Poisson’s ratio of 0.495. The dimensions of the probe face are 40 mm (lateral) by 10 mm (elevational). 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 0.5 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.

The experimental protocol involved corrupting the simulated RF echo data with additive noise at six signal-to-noise ratios (SNR) in the range 4 dB to -1 dB. Each displacement tracking method was evaluated 200 times (with different random noise) at each SNR level. 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 are presented in terms of the mean and standard deviation of d across the 200 trials.

Figure 7(a) confirms that, in the presence of noise, performance suffers 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 SNR levels above 0 dB, though at -1 dB it is clear that incorrect, noise-affected lateral displacement estimates are affecting the L3 axial tracking. The worst performer is the hybrid method with L3 lateral tracking, further evidence of the futility of fine resolution lateral search in the presence of poor lateral focusing and noise.

Figure 7(b) demonstrates the benefits of lateral motion compensation in all but the most contrived cases. The ground truth lateral displacement field reaches a maximum of 0.81 mm, which is 2.6 times the A-line separation. Without offsetting the L3 windows laterally, axial displacement estimation is severely compromised, as evidenced by the poor performance of the “no lateral offset” method. The other four methods show much the same relative performance as in Figure 7(a).

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 [11]. The ground truth in Figure 8(d) reveals lateral displacement of between 1 and 2 A-lines in the centre layer, and less than 0.5 A-lines in the top and bottom layers. This explains the “no lateral offset”

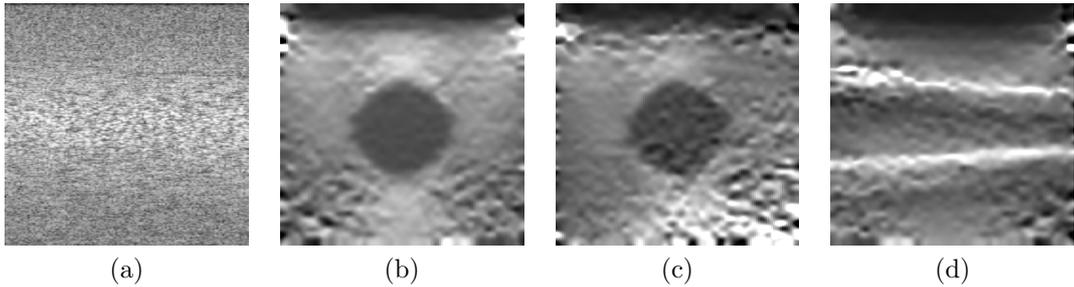


Figure 6: Finite element/Field II simulations. (a) B-mode image. Axial strain images obtained using the hybrid method for (b) the first simulation with normal probe compression, (c) the second simulation with tilted probe compression and (d) the third simulation with slip planes: dark is hard and bright is soft.

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 good, albeit low resolution, approximation with the hybrid method. Stradwin 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(c) suggest that the trade-off favours the hybrid method in all but the lowest noise scenarios. In fact, the apparent advantage of the multi-level method at high SNR is largely attributable to other factors, which we will discuss in Section 3.3. In other respects, the results in Figure 7(c) are similar to those for the tilted compression in Figure 7(b), except that the “no lateral offset” method does not suffer to quite the same extent, since only the centre layer is problematic.

Figure 9 explains our choice of the hybrid method’s regression parameter k , where k is defined in Equation (2). Setting $k = 0.1$ improves the results compared with the no regression 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.

3.3 Computational complexity

Table 2 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*.

The hybrid method and Stradwin’s tracking algorithm owe much of their speed to their use of downsampled data: the same amount of downsampling was used for each. In contrast, we implemented the multi-level technique according to the description in [25], with no downsampling of the RF data. This explains the multi-level method’s relatively slow performance.

An alternative comparison can be found in Figure 10(a), which shows a repeat of the third finite element simulation with slip planes, but this time with no downsampling in the Stradwin and hybrid methods ($S_y = 1$). We chose this simulation since Figure 7(c) demonstrates the multi-level method outperforming the hybrid method at high SNR. With no downsampling (Figure 10), the hybrid and multi-level methods are virtually indistinguishable at high SNR, and the hybrid

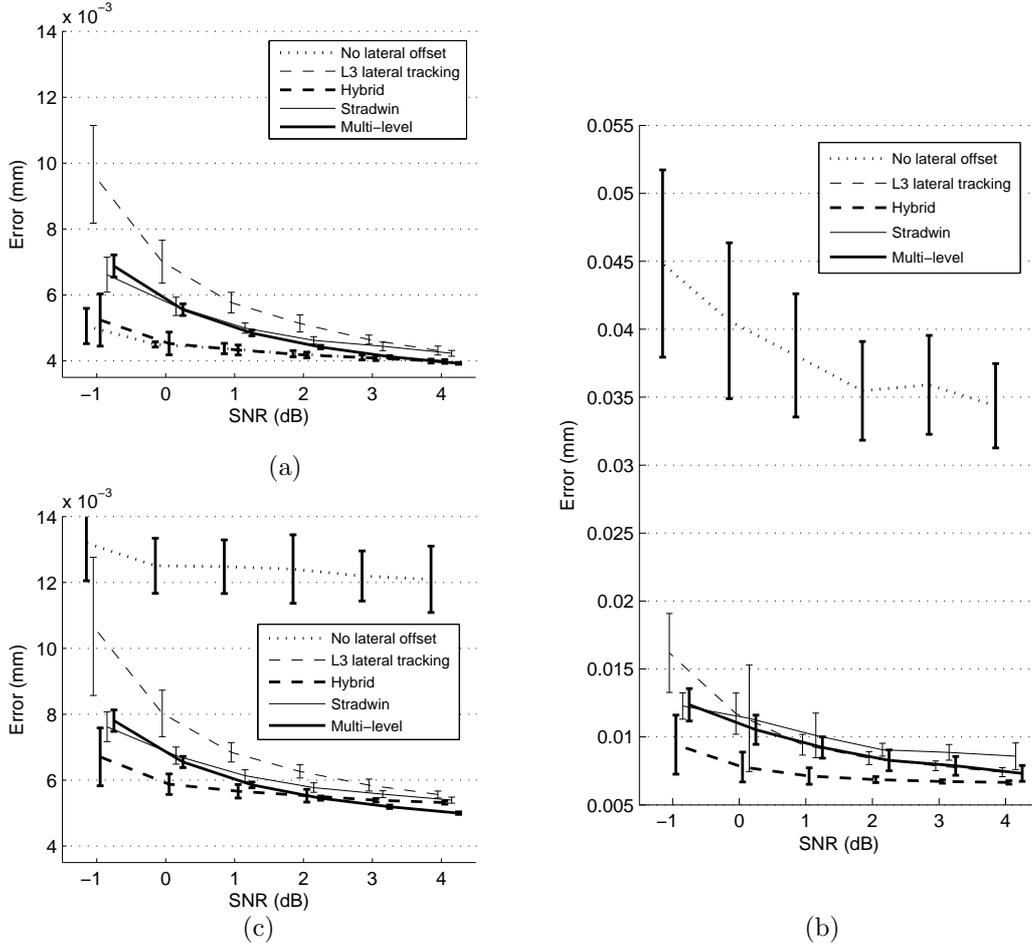


Figure 7: Results for the finite element simulations. (a) Normal compression with stiff inclusion. (b) Tilted compression with stiff inclusion. (c) Normal compression with slip planes. The ± 1 standard deviation error bars are based on 200 repetitions.

method	procedure	time (ms)	
		in vivo	simulations
hybrid	preprocessing	12	14
	L1 and L2 search	3	3
	regression	< 0.1	< 0.1
	L3 search	13	16
	postprocessing	10	13
	total	38	46
Stradwin	total	41	44
multi-level	total	109	307

Table 2: 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.

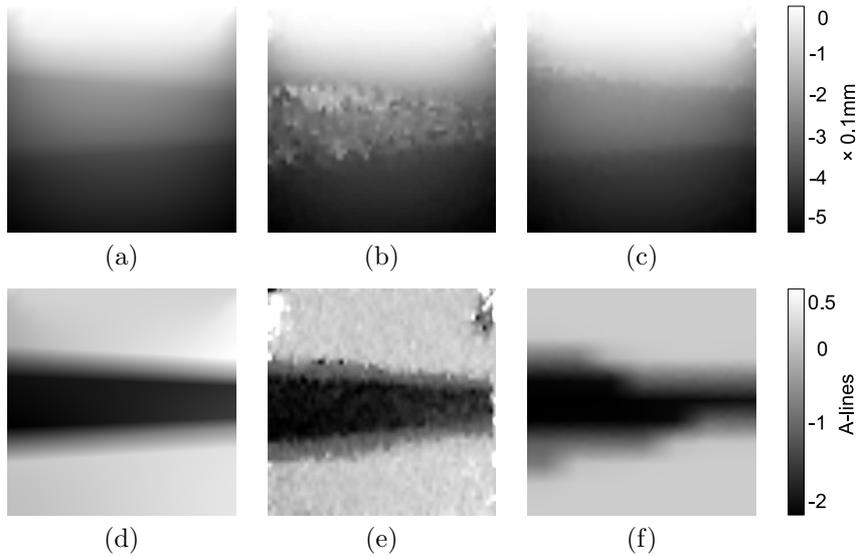


Figure 8: Displacement fields for the third simulation. (a) Ground truth axial displacement field. The axial displacement field obtained by (b) the “no lateral offset” method and (c) the hybrid method, with no additive noise. (d) Ground truth lateral displacement field. The lateral displacement field obtained by (e) L3 lateral tracking and (f) the hybrid method.

method’s advantage is further consolidated at low SNR. In particular, the hybrid method’s stability, as indicated by the error bars, appears to depend strongly on the amount of downsampling.

We have presented most of the results with downsampling, so as to demonstrate the hybrid method’s superior accuracy at real-time, clinically practical execution rates. Nevertheless, Figure 10 shows that its performance can be improved even further, should this be required, while still exceeding 10 frames per second with typical *in vivo* data.

4 Conclusions and further work

The hybrid displacement estimation method combines multi-level and tracking strategies to achieve an appealing balance between speed and robustness. 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 noise 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 sector probes is of comparable quality to lateral focusing, suggesting that the strategy presented here for lateral motion compensation should be appropriate in the elevational direction as well.

References

- [1] S. K. Alam, J. Ophir, and E. E. Konofagou. An adaptive strain estimator for elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 45(2):461–472, 1998.

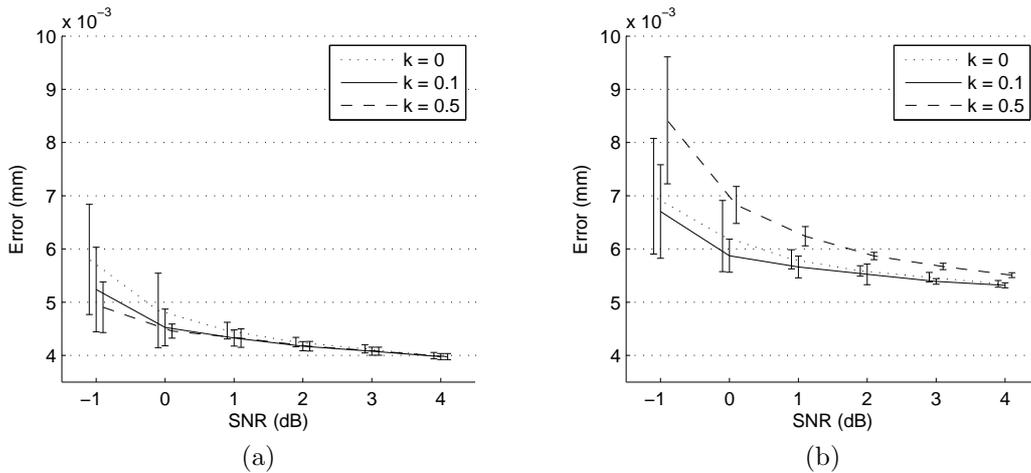
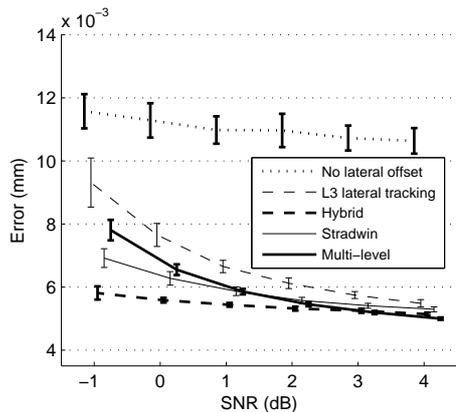


Figure 9: The role of the regression parameter k . Simulation results for (a) the stiff inclusion with normal compression and (b) the slip planes with normal compression, at several values of k .

- [2] J. Bai, C. Ding, and Y. Fan. A multi-scale algorithm for ultrasonic strain reconstruction under moderate compression. *Ultrasonics*, 37:511–519, 1999.
- [3] M. Bilgen and M. F. Insana. Deformation models and correlation analysis in elastography. *The Journal of the Acoustical Society of America*, 99(5):3212–3224, 1996.
- [4] L.N. Bohs and G. Trahey. A novel method for angle independent ultrasonic imaging of blood flow and tissue motion. *IEEE Transactions on Biomedical Engineering*, 38(3):280–286, 1991.
- [5] W. L. Briggs, V. E. Henson, and S. F. McCormick. *A multigrid tutorial*. Society for Industrial and Applied Mathematics, 2000.
- [6] E. S. Burnside, T. J. Hall, A. M. Sommer, G. K. Hesley, G. A. Sisney, W. E. Svensson, J. P. Fine, J. Jiang, and N. J. Hangiandreou. Differentiating benign from malignant solid breast masses with US strain imaging. *Radiology*, 245(2):401–410, November 2007.
- [7] I. Céspedes, Y. Huang, J. Ophir, and S. Spratt. Methods for estimation of subsample time delays of digitized echo signals. *Ultrasonic Imaging*, 17:142–171, 1995.
- [8] I. Céspedes, J. Ophir, and S. K. Alam. The combined effect of signal decorrelation and random noise on the variance of time delay estimation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(1):220–225, 1997.
- [9] P. Chaturvedi, M. F. Insana, and T. J. Hall. 2-D companding for noise reduction in strain imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 45:179–191, 1998.
- [10] H. Chen, H. Shi, and T. Varghese. Improvement of elastographic displacement estimation using a two-step cross-correlation method. *Ultrasound in Medicine and Biology*, 33(1):48–56, 2007.
- [11] L. Chen, G. M. Treece, J. E. Lindop, A. H. Gee, and R. W. Prager. A quality-guided displacement tracking algorithm for ultrasonic elasticity imaging. *Medical Image Analysis*, in press.
- [12] R. J. Dickinson and C. R. Hill. Measurement of soft tissue motion using correlation between A-scans. *Ultrasound in Medicine and Biology*, 8(3):263–271, 1982.



(a)

method	procedure	time (ms)	
		in vivo	simul.
hybrid	preproc.	40	90
	L1 and L2	18	75
	regression	< 0.1	< 0.1
	L3 search	28	68
	postproc.	10	13
	total	96	246
Stradwin	total	99	182
multi-level	total	109	307

(b)

Figure 10: The effects of downsampling. (a) A repeat of the third simulation with slip planes, but without downsampling the analytic RF signal in the hybrid and Stradwin methods. These results should be compared with the downsampled equivalents in Figure 7(c). (b) Execution times of the displacement estimation methods without downsampling; compare with Table 2.

- [13] E. S. Ebbini. Phase-coupled two-dimensional speckle tracking algorithm. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 53(5):972–990, 2006.
- [14] J. A. Jensen. Field: A program for simulating ultrasound systems. In *Proceedings of the 10th Nordic-Baltic Conference on Biomedical Imaging*, pages 351–353, Tampere, 1996.
- [15] J. Jiang and T. J. Hall. A parallelizable real-time motion tracking algorithm with applications to ultrasonic strain imaging. *Physics in Medicine and Biology*, 52:3773–3790, 2007.
- [16] E. Konofagou and J. Ophir. 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*, 24(8):1183–1199, 1998.
- [17] H. Liebgott, J. Fromageau, J. E. Wilhelm, D. Vray, and P. Delachartre. Beamforming scheme for 2D displacement estimation in ultrasound imaging. *EURASIP Journal on Applied Signal Processing*, 8:1212–1220, 2005.
- [18] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. Phase-based ultrasonic deformation estimation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 55(1):94–111, January 2008.
- [19] A. M. Lubinski, S. Y. Emelianov, and M. O’Donnell. Speckle tracking methods for ultrasonic elasticity imaging using short time correlation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 46(1):82–96, 1999.
- [20] R. Muthupillai, D. J. Lomas, P.J. Rossman, J. F. Greenleaf, A. Manduca, and R. L. Ehman. Magnetic resonance elastography by direct visualization of propagating acoustic strain waves. *Science*, 269(5232):1854–1857, 1995.
- [21] 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, 1994.
- [22] J. Ophir, I. Céspedes, H. Ponnekanti, Y. Yazdi, and X. Li. Elastography: a quantitative method for imaging the elasticity of biological tissues. *Ultrasonic Imaging*, 13:111–134, 1991.

- [23] C. Pellot-Barakat, F. Frouin, M. F. Insana, and A. Herment. Ultrasound elastography based on multiscale estimations of regularized displacement fields. *IEEE Transactions on Medical Imaging*, 23:153–163, 2004.
- [24] A. Pesavento, C. Perrey, M. Krueger, and 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*, 46(5):1057–1067, September 1999.
- [25] H. Shi and T. Varghese. Two-dimensional multi-level strain estimation for discontinuous tissue. *Physics in Medicine and Biology*, 52:389–401, 2007.
- [26] T. Sugimoto, S. Ueha, and K. Itoh. Tissue hardness measurement using the radiation force of focused ultrasound. In *Proceedings of IEEE Ultrasonics Symposium 1990*, volume 3, pages 1377–1380, December 1990.
- [27] G. M. Treece, J. E. Lindop, A. H. Gee, and R. W. Prager. Efficient elimination of dropouts in displacement tracking. In *Proceedings of the Fifth International Conference on the Ultrasonic Measurement and Imaging of Tissue Elasticity*, page 68, Snowbird, Utah, October 2006.
- [28] G. M. Treece, J. E. Lindop, A. H. Gee, and R. W. Prager. Freehand ultrasound elastography with a 3-D probe. *Ultrasound in Medicine and Biology*, 34(3):463–474, 2008.
- [29] F. Walker and E. Trahey. A fundamental limit on delay estimation using partially correlated speckle signals. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42:301–308, 1995.
- [30] F. Yeung, S. F. Levinson, and K. J. Parker. Multilevel and motion model-based ultrasonic speckle tracking algorithms. *Ultrasound in Medicine and Biology*, 24:427–441, 1998.
- [31] R. Zahiri-Azar and S. E. Salcudean. Motion estimation in ultrasound images using time domain cross correlation with prior estimates. *IEEE Transactions on Biomedical Engineering*, 53:1990–2000, 2006.
- [32] Y. Zhu and T. J. Hall. A modified block matching method for real-time freehand strain imaging. *Ultrasonic Imaging*, 24:161–176, 2002.

A Implementation of the multi-level algorithm

The multi-level method in [25] 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 magnitude images. 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 3 used the parameters in Table 3, 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 [25] is unsuitable for our studies, since we have varying noise levels and therefore require a variable threshold. Instead, we rejected the poorest 25% of the displacement estimates. The 25% level was chosen to optimize performance.

[25] 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 [25]. We experimented with various values of p and found that $p = 0$ produced the best results for the experiments in Section 3.

parameter	in vivo	simulations
L'3 image size (pixels)	128 × 191	128 × 429
number of L'3 windows (lateral × axial)	3 × 3	3 × 3
L'3 window size (pixels)	26 × 38	26 × 86
L'3 search region (top, pixels)	30 × 44	30 × 100
L'3 search region (middle, pixels)	34 × 52	34 × 114
L'3 search region (bottom, pixels)	38 × 58	38 × 128
L'2 image size (pixels)	128 × 382	128 × 859
number of L'2 windows (lateral × axial)	7 × 9	7 × 9
L'2 window size (pixels)	13 × 38	13 × 86
L'2 search region (pixels)	17 × 48	17 × 106
L'1 image size (pixels)	128 × 765	128 × 1718
number of L'1 windows (lateral × axial)	21 × 27	21 × 27
L'1 window size (pixels)	6 × 38	6 × 86
L'1 search region (pixels)	8 × 42	8 × 96
L'0 image size (A-lines × samples)	128 × 1559	128 × 3465
number of L'0 windows (lateral × axial)	63 × 81	63 × 81
L'0 window size (A-lines × samples)	3 × 38	3 × 86
L'0 search region (A-lines × samples)	5 × 42	5 × 96

Table 3: Parameters used in the multi-level displacement estimation method.