A QUALITY-GUIDED DISPLACEMENT TRACKING ALGORITHM FOR ULTRASONIC ELASTICITY IMAGING

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

CUED/F-INFENG/TR 593

January 2008

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

Email: lc420/gmt11/jel35/ahg/rwp @eng.cam.ac.uk

A quality-guided displacement tracking algorithm for ultrasonic elasticity imaging

Lujie Chen, Graham M. Treece, Joel E. Lindop, Andrew H. Gee
and Richard W. Prager

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

Abstract

Displacement estimation is a key step in the evaluation of tissue elasticity by quasistatic strain imaging. An efficient displacement estimation process may incorporate a tracking strategy that initializes each point according to its neighbours' displacements and then performs a localized search. This increases the accuracy and reduces the computation expense compared with exhaustive search. However, simple tracking strategies fail when the target displacement map exhibits complex structure. For example, there may be discontinuities and regions of indeterminate displacement caused by decorrelation between the pre- and post-deformation radio frequency (RF) echo signals. This paper introduces a novel displacement tracking algorithm, with an intelligent search strategy guided by a data quality indicator. Comparisons with existing methods show that the proposed algorithm is more robust across a range of displacement distributions, RF window lengths and noise levels.

1 Introduction

After over twenty years of research and development, ultrasound elasticity imaging is gradually demonstrating its ability to characterize the mechanical properties of tissue [6]. Research activity has moved from simulation and laboratory experiments [1] to clinical trials [7]. This trend is supported by ongoing development of signal processing methods and hardware innovation.

Some sort of deformation estimation is required by all elasticity imaging methods. For example, transient deformation induced by the acoustic radiation force can be used to measure the tissue's mechanical properties [15]. Alternatively, waves of shear deformation produced by an impulsive load on either side of the tissue [3], or internally by acoustic radiation force [2], are related to tissue elasticity. Depending on the magnitude of the deformation involved, some of the above methods may benefit from the displacement tracking algorithm that we describe in this paper. However, our focus is on the quasistatic approach to elasticity imaging [16], where deformation is introduced by pressing on the probe. Two RF ultrasound frames are acquired pre- and post-deformation. Some indication of the tissue elasticity is provided by the axial strain, which is usually retrieved in two steps: *displacement estimation*, by matching pre-deformation RF data windows with post-deformation windows; and *strain estimation*, by differentiating the displacement field. The speed and accuracy of the strain imaging system depends heavily on the displacement estimation algorithm in the first step.

Among several techniques for displacement estimation [12, 24], correlation maximization was the first to be proposed and remains the most widely used [16]. The pre-deformation RF frame is divided into an axial-lateral grid of small windows, which are associated with corresponding windows in the post-deformation frame by searching for the highest correlation match. Tissue displacement is then given by the shift between the pre- and post-deformation windows. Efforts have been made to adapt the correlation algorithm to provide subsample accuracy [4]. Moreover, standard correlation [21], and its variants such as the sum of absolute differences [5, 19] and the sum of squared differences [26], have been integrated with multi-level or multi-scale schemes [5, 19, 21, 26] to improve the estimation stability. However, since these schemes rely on expensive,

brute force search to find the correlation peaks, they are less than ideal for implementation in a real-time system.

As an alternative to exhaustive search, Hall et. al. [9] proposed a tracking strategy that takes the displacement in each row of the grid as an initial guess for the next row. This assumes that the displacement is continuous in the axial direction. The benefit is a smaller search range and hence a reduction in the computation expense. However, there is an additional, implicit assumption that the displacement estimated in the previous row does not suffer from significant errors; otherwise, subsequent rows' displacements might be incorrectly calculated because the search range does not include the correct solution.

A second approach to displacement estimation is based on the phase of RF data. O'Donnell et. al. [17, 18] showed that the phase of the complex correlation between baseband signal windows is related to the relative displacement of the two windows. In particular, the correlation peak corresponds to zero phase, while a nonzero phase can be used to estimate the location of the peak without further search. Based on their work, Pesavento et. al. [20] proposed an improved method that calculates the displacement by tracking the zero phase location from point to point[1]. The efficiency of this tracking approach lies in the way the displacement map is accumulated. A point's displacement is initialized according to the value of a neighbour's previously calculated displacement. Phase zero search at the new point then refines the accuracy of this initial guess. If the guess is close to the actual value, the correlation phase can be used to estimate the zero phase location with minimal further computation.

Although they use different techniques for the point-wise displacement estimation, Hall's [9] and Pesavento's [20] approaches both rely on a tracking strategy to advance the estimation process from one point to the next. This strategy sits on top of the low-level displacement estimation process, and controls the accumulation of the displacement map for the whole frame. An intelligent tracking algorithm should be able to identify inaccurate points and avoid initializing other points with unreliable data; while a naïve algorithm may assign the initial displacement without considering the data's reliability, causing subsequent displacement estimations to fail.

Tracking strategies have not been well studied in the ultrasonic strain imaging literature. This might be because of the homogeneity of phantoms typically used to evaluate strain imaging systems. Data acquired from such phantoms is amenable to processing by simple row-to-row [20, 28], column-to-column [11] or diagonal [27] tracking strategies. However, as strain imaging emerges from the laboratory into the clinic, there is an urgent need for more sophisticated tracking algorithms. The mechanical properties of human tissue are inhomogeneous: consequently, real displacement maps often exhibit several disjoint regions, with displacement continuity only within each region. Moreover, RF echo decorrelation is commonplace when scanning *in vivo*. These factors present important challenges for displacement tracking algorithms.

Lindop et. al. [14] refer to tracking errors as *drop-outs* and proposed a two-pass tracking strategy to reduce drop-outs. Later, Lindop et. al. [13] and Treece et. al. [22, 23] described a refined drop-out correction method. In the first tracking pass (from the top of the frame to the bottom), every point's displacement is initialised from a seed point in the previous row. The algorithm examines a fixed-length lateral span in the previous row, centred on the current point, to locate the best seed point. In this context, "best" is defined in terms of the correlation between the previously matched pre- and post-deformation RF windows. The second pass of tracking (from the bottom to the top) uses a similar seeding strategy and corrects remaining errors by comparing the correlations between the two passes: better correlated matches found in the second pass overwrite poorer matches discovered in the first pass. The algorithm prevents poor displacement estimations, characterised by low correlations, from affecting subsequent estimates, and performs well in many situations. However, it has two limitations. Firstly, the search range for the seed point in the previous row is a fixed parameter that has to be set in advance. If a small distance is chosen, all points within the range might be unreliable: initialization of the current point is therefore inaccurate. If a larger distance is chosen, this increasingly violates the assumption of

[1] In the sequel, we shall refer loosely to each RF window as a *point*. Displacement is estimated at these points, then differentiated to obtain strain. A displacement tracking algorithm propagates displacement estimates from point to point.

short range (but not long range) displacement continuity. Secondly, the propagation direction of the algorithm is constrained (up-down and down-up, in this case), in much the same manner as all existing methods [11, 20, 27, 28]. If a previous row contains a large section of invalid data, the algorithm has to resort to a second pass, or even further passes, to correct these errors.

This paper introduces a novel, quality-guided displacement tracking algorithm for ultrasonic strain imaging. Its basis is a propagation scheme that is not constrained to any particular set of directions, but is instead guided by the quality of the data. The algorithm can be seeded at a single location or, for greater robustness, at multiple points. The quality measure that guides the algorithm can be any suitable metric calculated from the RF data. The general approach is inspired by similar methods used in other disciplines for phase unwrapping [8]. This paper includes a series of experiments that demonstrate the new approach's robustness in comparison with simpler alternatives. Specifically, the quality-guided algorithm is able to track through geometrically irregular and disjoint regions, and copes well with regions of poorly correlated RF data.

2 Quality-guided displacement tracking

2.1 Seed initialization

The proposed quality-guided tracking algorithm comes in single-seed and multiple-seed variants. The former demonstrates the fundamental structure of the tracking strategy, while the latter is a significant extension that adds robustness when there are disjoint regions and displacement discontinuities. For both versions, the first stage is seed initialization.

The seed initialization method used in this study is based on a simple grid test. On an RF frame, a grid of $N \times N$ equally spaced points are tested as seed candidates. At each test point, a window is defined with the point at its centre. The complex cross-correlation is used to measure the similarity between the corresponding windows in RF frames recorded before and after deformation:

$$C(d) = \frac{\sum_{t \in T} a_1(t) a_2^*(t+d)}{\sqrt{\sum_{t \in T} |a_1(t)|^2 \sum_{t \in T} |a_2(t)|^2}} \tag{1}$$

where $C(d)$ is the correlation coefficient at displacement d, a_1 and a_2 are analytic RF signals obtained before and after deformation respectively, a_2^* is the complex conjugate of a_2, and T is the window size. A brute force search that varies d within the maximum expected displacement identifies the best match displacement. The corresponding correlation coefficient and displacement are recorded for each test point. After all grid points have been processed, the one with the highest correlation is chosen as the seed for the single-seed tracking algorithm. For the multiple-seed tracking algorithm, all of the grid points are used as seeds.

It remains to discuss the choice of an appropriate value for N. For the single-seed algorithm, the larger the value of N, the greater the chance of discovering a good seed. For the multiple-seed algorithm, the larger the value of N, the greater the probability that each disjoint region benefits from at least one good seed: conversely, with small N, there is an increased likelihood that some regions might not be properly tracked. In practice, reliable seeding is readily achieved with values of N that are sufficiently small to have a negligible impact on the overall computational load, which is dominated by the subsequent displacement tracking stage.

2.2 Single-seed, quality-guided tracking

The proposed tracking strategy requires a quality indicator to determine each point's reliability before it is used to initialize a neighbouring point. A number of criteria may be used as quality indicators, including B-mode intensity, phase gradient variance and correlation coefficient. In this paper, all studies make use of the complex correlation coefficient, as given in Equation (1).

The quality map may be represented as $Q(x, y)$, where y and x are indices in the axial and lateral directions respectively. The tracking algorithm maintains a set S which contains points that have been initialized and are ready for processing. In the first step, S_1 (1 is the index of the step) starts with the one seed, $p_1 = (x_0, y_0)$. As p_1 is the only point ready for processing, it is selected and fed into any suitable displacement estimation method [12, 16, 17, 18, 20, 24] to calculate the displacement at this point. If using a correlation-based quality indicator, the calculated displacement is then used to determine $Q(x_0, y_0)$ according to Equation (1).

In the second step, the point set is updated by removing p_1 from S and adding its four neighbours, whose initial displacement estimates and qualities are assigned according to p_1's values:

$$S_2 = \{p_2, p_3, p_4, p_5\} \tag{2}$$

where, $p_2 = (x_0 + 1, y_0)$, $p_3 = (x_0 - 1, y_0)$, $p4 = (x_0, y_0 + 1)$, and $p_5 = (x_0, y_0 - 1)$. Since there are now four points ready for processing, the next point to be processed is selected according to the maximum quality criterion:

$$Pt_2 = \arg\max\left(Q(p_2), Q(p_3), Q(p_4), Q(p_5)\right) \tag{3}$$

where Pt_2 refers to the current point to be processed (2 denotes the index of the step), and the arg max() operator extracts the point with the maximum quality value from S. If there is a tie, as there will inevitably be at step 2, any one of the tied points is selected at random. Subsequently, Pt_2 is processed by a displacement estimation method, and during this procedure both its displacement and quality are refined.

Similarly, in the third step, Pt_2 is removed from the point set S. Its neighbours that have not been processed are initialized with Pt_2's displacement and quality, and added to S. A new current point is then selected, according to the maximum quality criterion. This recursive process continues and can be described as:

$$S_{k+1} = S_k + \text{Neighbour}(Pt_k) - Pt_k \tag{4}$$

$$Pt_{k+1} = \arg\max\left(Q(p) |\, p \in S_{k+1}\right) \tag{5}$$

where the Neighbour() operator extracts a point's 4-way neighbours that are not in S and have not yet been processed. For any neighbour that is already in S, a comparison is made between its quality value and that of the current point. If the current point's quality is greater, the neighbour's initial displacement estimate is replaced by the current point's displacement. This simply reflects the fact that a better displacement estimate is now available. The algorithm terminates when the point set S is empty, which indicates that all points in the RF frame have been processed.

This strategy ensures that high quality regions are processed first, while low quality regions are avoided at an early stage. The advantage is twofold. Firstly, accurate estimation in high quality regions is propagated to other points, thus facilitating subsequent processing. Secondly, initialization errors frequently encountered in low quality regions are prevented from propagating. Note, however, that the strategy does not guarantee that points are processed in order of quality. A high quality point will not be processed until the tracking path reaches it. The initial estimation of each point is obtained from one of its 4-way neighbours, with no jumps allowed, in accordance with the short range continuity assumption. Furthermore, there is no constraint on the direction in which the tracking path propagates: hence, the algorithm is able to process a region with an arbitrary geometrical shape. Depending on the type of quality indicator chosen, the quality map may be available prior to tracking (e.g. B-mode intensity) or it may have to be built dynamically during tracking (e.g. correlation coefficient). Figure 1 shows a flowchart for the single-seed, quality-guided tracking strategy with the correlation quality indicator generated dynamically.

2.3 Multiple-seed, quality-guided tracking

There are two common situations that may cause the single-seed tracking algorithm to fail, retrieving only a partially correct displacement map. Firstly, there may be two high quality regions

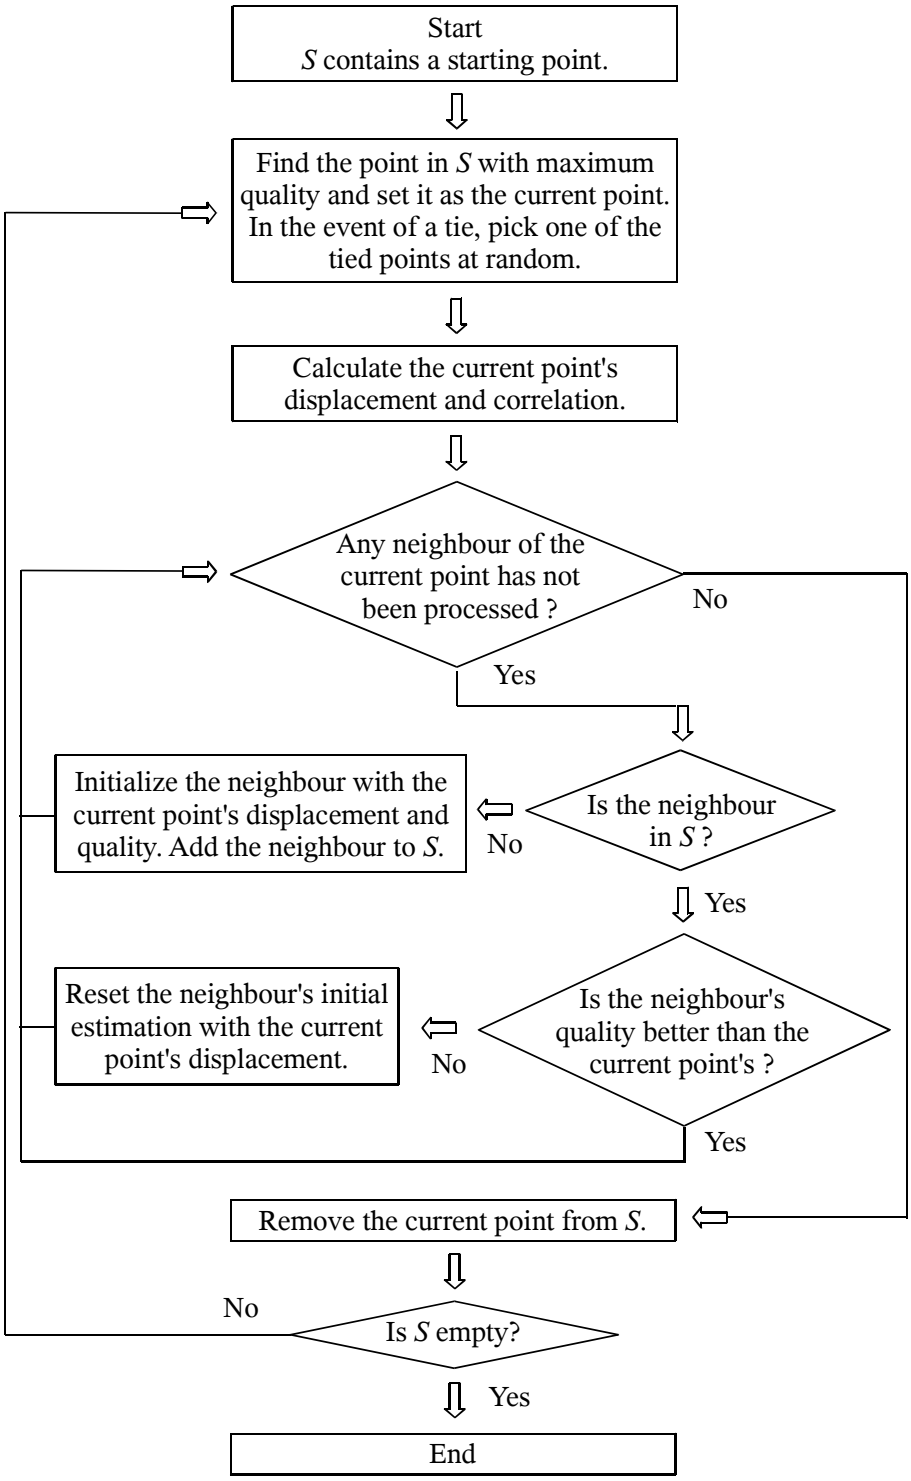

Figure 1: Flowchart of the single-seed quality-guided tracking algorithm.

separated by a region of poorly correlated RF data. Whichever region provides the initial seed, it is unlikely that the tracking algorithm will be able to maintain accurate estimation as it traverses the poor quality region. Therefore, tracking in the second high quality region will not be correctly initialized and displacement estimation will fail. Secondly, the correct displacement map may be discontinuous, as is often encountered when imaging arterial elasticity [21] or, more generally, when there are slip boundaries. When the tracking path reaches a discontinuity where there is a dramatic change in the true displacement, initialization by a 4-way neighbouring point is no longer appropriate and subsequent displacement estimation will fail.

A multiple-seed variant of the quality-guided tracking strategy is proposed to overcome these problems. As with the single-seed version, seed initialization is by the grid point test (Section 2.1), but in this case all the grid points are used as seeds and are added to the point set in the first step:

$$S_1 = \{p_1, p_2, p_3, \ldots, p_n\} \tag{6}$$

Equation (6) describes a point set S containing n seeds. The seeds are not all processed in the first step: instead, only the one with the maximum quality is selected as the current point Pt_1

$$Pt_1 = \arg\max\left(Q(p_1), Q(p_2), Q(p_3), \ldots, Q(p_n)\right) \tag{7}$$

and fed into the displacement estimation procedure. All subsequent processing — displacement estimation on the current point, initializing the current points' neighbours, updating the point set, and generating a new current point — is identical to the single-seed version (Figure 1). If at least one good seed is planted in each disjoint region, all regions should be propagated individually and successfully. This is because the quality at a region's boundary is relatively low, either because of signal decorrelation or because of discontinuity-induced tracking failure. Hence, when a seed propagates to a region boundary, it stops growing and other seeds get a chance to proceed. Eventually, regions of poorly correlated data, or points at the boundaries between discontinuous regions, are processed at the last stage of tracking. Inevitable estimation errors are confined to these small localities and not propagated elsewhere.

An issue that arises with the use of multiple seeds is the correction of tracking errors caused by bad seeds. Bad seeds arise because the standard cross-correlation metric used for seed initialization can easily pick up an incorrect best match, especially when the window is small or the signal is noisy [25]. A distinct feature of such seeds is that they are only able to propagate into a small area. The quality at the boundary drops dramatically, and subsequent points are processed by good quality estimates propagating from other seeds. The resulting displacement map, however, contains small regions of incorrect displacement data surrounding the initial poor seeds.

The bad-seed defect can be fixed by re-initializing any problematic regions using good quality points at their boundaries, and repeating the tracking procedure. Regions in need of this treatment are identified using a simple threshold on the number of points propagated from any individual seed: small regions, below the threshold, are labelled as requiring reprocessing. Then, the multiple-seed algorithm is re-run, but this time discarding the seeds that grew the small regions. During this second-pass of tracking, displacement estimation is performed only on the labelled regions. This simple method avoids considerable redundant reprocessing and obviates the need to explicitly locate the boundary points of the troublesome regions.

2.4 Implementation details

Efficient implementation of the quality-guided tracking algorithm requires each point's attributes (location, membership of S, displacement and quality) to be stored in two separate data structures. The first, a simple two-dimensional array indexed by x and y, is used to quickly identify neighbours. However, this data structure is unsuitable for one important stage of the algorithm, namely identifying the maximum quality point in the set S. Linear search of the two-dimensional array has algorithmic complexity $\mathcal{O}(n)$ at each tracking step, where n is the number of points.

Thus, the active point set S is also maintained as a one-dimensional list, sorted according to quality. The tracking algorithm then need only pick the first point from S at each step, with no

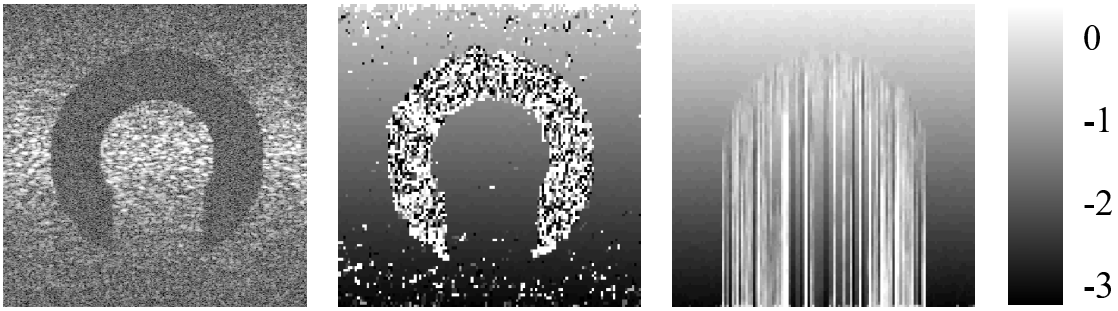

(a) (b) (c)

Figure 2: (a) B-mode image of the simulated data. The arc-shaped region contains only noise. Outside the arc-shaped region, white noise is superimposed on the RF signal, with less noise at the centre than at the axial extremities. Displacement obtained by (b) exhaustive search and (c) tracking along A-lines. The units displayed on the right are cycles of the sampled RF signal.

searching required. When the current point's neighbours are added to the list, they are inserted at the appropriate location according to their quality, thus maintaining the integrity of the sorted list. Binary search, with complexity $\mathcal{O}(\log(n))$, is used to find the insertion indices, significantly reducing the computational load. In such an implementation, the tracking overhead amounts to around 1% of the cost of the underlying displacement estimation, even with efficient phase zero search estimation techniques.

3 Results and discussion

Field II [10] simulations and an *in vivo* scan were used to evaluate the quality-guided tracking algorithm. The simulations comprised two RF frames (256 vectors, 3464 samples) with a uniform 1% strain field. The focus depth was set at 20 mm, the centre of the frame. Artificially introduced white noise corrupted certain regions of the RF data. The *in vivo* scan of a human carotid artery was recorded using a Terason[2] T3000 ultrasound system with a 6.25 MHz linear array transducer. The RF sampling frequency was 35.776 MHz, with each frame comprising 128 vectors of 2395 samples. In all the studies, we used complex cross-correlation based exhaustive search, within the maximum expected displacement range, for seed initialization; correlation coefficient as the quality indicator; and phase zero search [20], with logarithmic compressed signal amplitude, for deformation estimation.

3.1 Tracking of geometrically irregular regions

Figure 2(a) shows a B-mode image of the first simulated data set. The signal in the arc-shaped region is completely masked by white noise, providing a test of the algorithms' ability to track across decorrelated regions. Noise is also introduced outside the arc-shaped region, with a higher signal-to-noise ratio (SNR) at the centre than at the top and the bottom. Five algorithms were tested, namely exhaustive search without tracking, tracking along A-lines [20], drop-out correction [13, 22, 23], and both variants of the quality-guided tracking algorithm.

The displacement distribution obtained by exhaustive search is shown in Figure 2(b). Clearly, it is not possible to recover the displacement in the totally decorrelated arc-shaped region. Where the SNR is low (top and bottom of the frame), exhaustive search picks up a large number of false matches and produces erroneous displacement estimates. Results obtained by tracking top-to-bottom along A-lines, starting with an assumed displacement of zero at the top, are shown in Figure 2(c). When the tracking algorithm encounters the uncorrelated data, the displacement

[2] `www.terason.com`

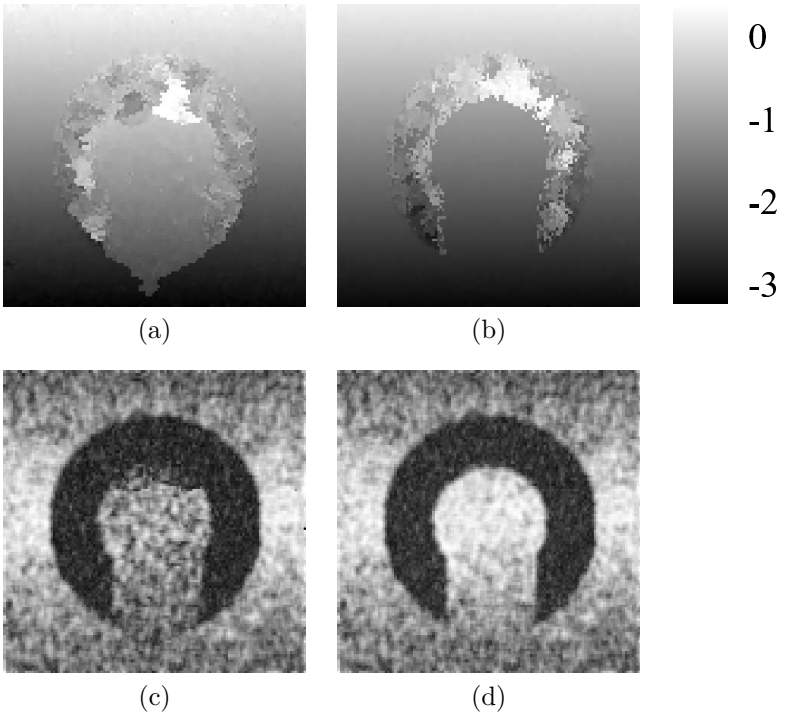

Figure 3: Same data as in Figure 2(a). Displacement obtained by (a) the first pass and (b) the second pass of the drop-out correction method. Quality map of (c) the first pass and (d) the second pass.

estimates inevitably become inaccurate. Since these incorrect estimates are used to initialise the phase zero search at subsequent points, correct displacement tracking is not re-established, even outside the arc-shaped region. The outcome is that tracking errors are propagated from the arc-shaped region downwards. However, it is interesting to note that, except for these understandable tracking errors, the simple A-line tracking strategy performs better than exhaustive search in low SNR regions at the top and bottom of the frame. This illustrates the advantage of exploiting displacement continuity inherent to all tracking approaches.

The drop-out correction method processes the RF frames in two passes. Figure 3(a) shows the results of the first pass. As with A-line tracking, the region below the arc receives spurious initial displacements and tracking fails. However, there are good displacement estimates available either side of the arc, and these gradually propagate inwards, since each point is initialised by searching a fixed-length window in the previous row, picking the displacement with the highest corresponding quality. At the end of the first pass, only the area enclosed by the arc and a tapering region below it are in need of correction. The second pass repeats the process, but this time sweeping from the bottom of the frame (where the first pass has discovered plenty of good displacement estimates to seed the second pass) to the top. If the second pass discovers matches of better quality than those of the first pass, the corresponding point's displacement and quality are updated. Results in Figure 3(b) illustrate correct displacement estimation everywhere in the frame apart from inside the arc itself. The two passes' quality maps are shown in Figure 3(c) and (d), with bright pixels indicating high correlation and dark pixels low correlation. The arc contains uncorrelated random noise, so the quality in this region is very low. The quality of the region enclosed by the arc is higher after the second pass than after the first pass, indicating that this region's displacement estimates have been corrected. The drop-out correction method illustrates the advantages of using a quality measure to guide the tracking algorithm, albeit with a highly constrained direction of propagation in this instance.

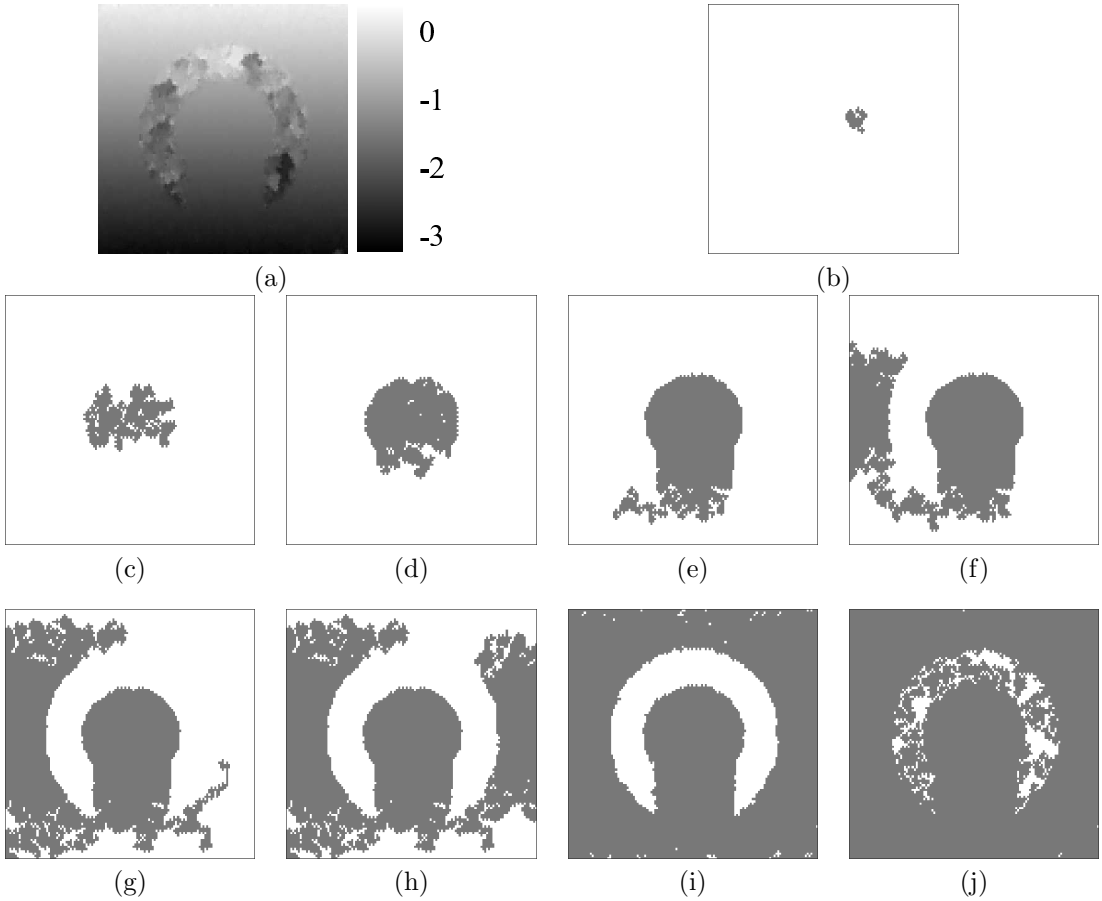

Figure 4: Same data as in Figure 2(a). (a) Displacement obtained by the single-seed quality-guided tracking algorithm. (b)–(j) show the propagation progress.

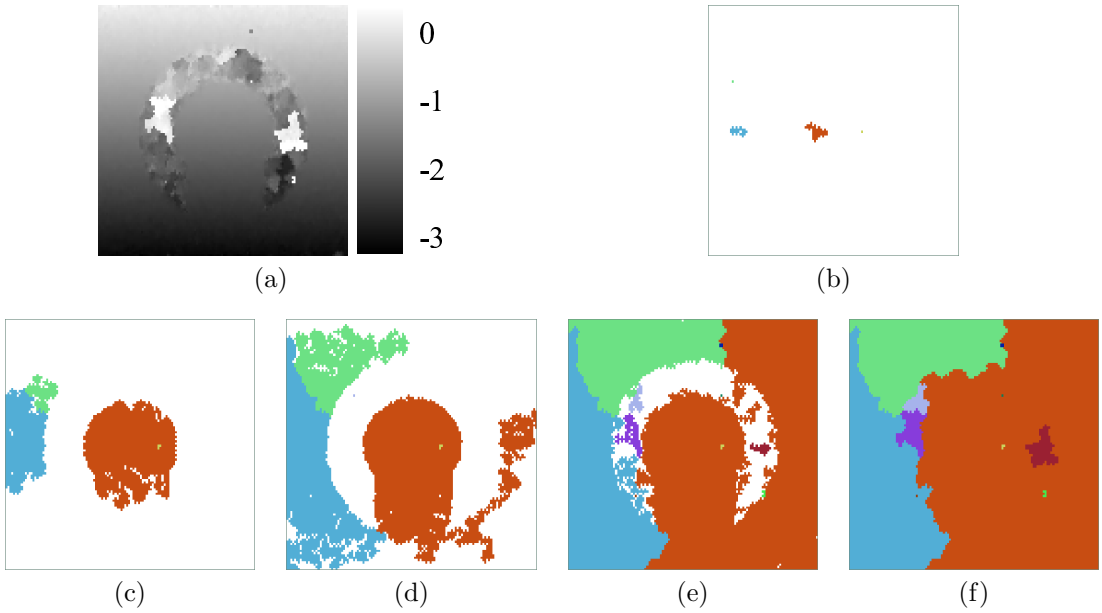

Figure 5: Same data as in Figure 2(a). (a) Displacement obtained by the first pass of the multiple-seed algorithm. (b)–(f) show the propagation progress.

The single-seed quality-guided tracking algorithm retrieves an accurate deformation distribution in one go — see Figure 4(a). Unlike all previously documented methods, the tracking path of this algorithm is dynamically generated. For this reason, it is instructive to visualize the progress of the tracking procedure, as shown in Figures 4(b)–(j). Dark pixels indicate points that have had their displacement estimated by phase zero search, while bright pixels indicate unprocessed points. At the first stage (Figure 4(b)), the seed starts from the central region enclosed by the arc. This has the best quality among all the grid points tested in the seed initialization process. In Figures 4(c)–(e), the seed propagates and fills up the region enclosed by the arc. Note in Figure 4(e) how the tracking path avoids entering the noisy arc: instead, it is guided by the quality indicator to the edge of the arc and proceeds to the well-correlated data on the left. After that, the algorithm quickly detects the strong SNR at the centre-left and processes the data with priority (Figure 4(f)). When high quality regions at the left are exhausted, the tracking path proceeds to the right (Figure 4(g)). It once again finds high quality territory and processes it with priority (Figure 4(h)). After all the high SNR regions have been processed, the tracking algorithm gradually proceeds to the lower SNR regions at the top and bottom of the frame (Figure 4(i)). Only when all the surrounding points have been processed does the tracking path penetrate the arc region itself (Figure 4(j)). Here, estimation errors are inevitable, but they are prevented from propagating into other regions of the frame. Without constraints on the propagation direction, the quality-guided tracking algorithm is capable of retrieving the displacement distribution of a geometrically irregular region in one pass, while minimizing the effects of estimation errors.

While the single-seed variant has performed flawlessly on this particular data, it is nevertheless instructive to examine the multiple-seed variant as well. In particular, we will see why this algorithm requires two passes to counter the effects of exhaustive search errors at the seed initialization stage. Figure 5(a) shows the displacement distribution obtained by the first pass. The overall distribution is correct and errors inside the arc are not propagated. However, a very small patch above the arc has not been correctly processed. To discover the root cause, one has to examine the tracking progress shown in Figures 5(b)–(f).

During seed initialization, twenty grid points were tested and a few were processed at an early stage (Figure 5(b)). Points propagated from each seed are displayed in a distinct colour. As the

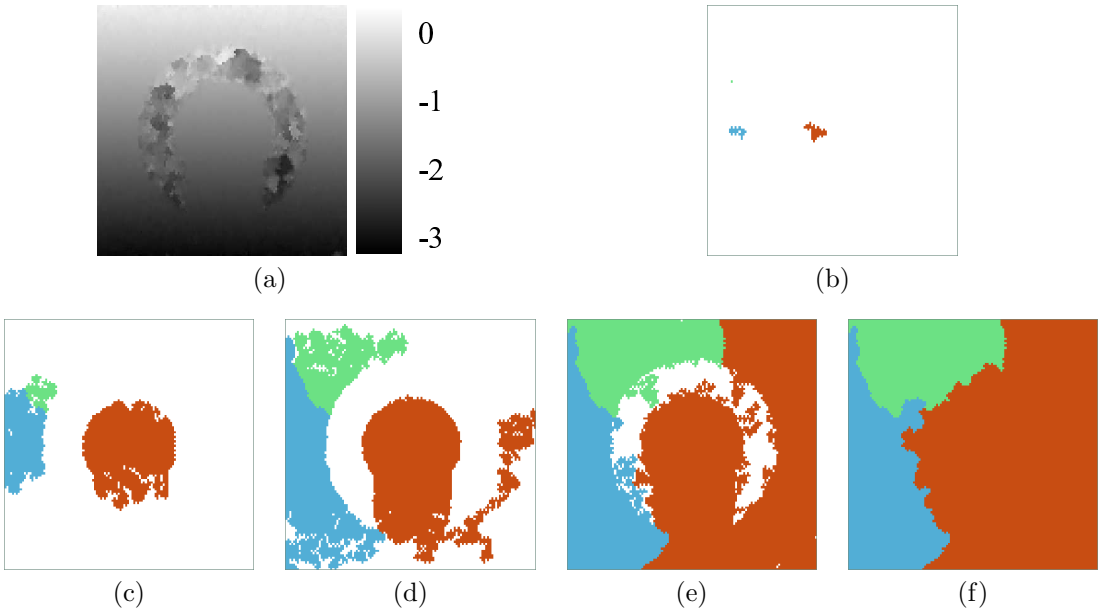

Figure 6: Same data as in Figure 2(a). (a) Displacement obtained by the second pass of the multiple-seed algorithm. (b)–(f) show the propagation progress.

algorithm proceeds (Figure 5(c)), two seeds with highly correlated data at the centre of the frame grow faster. Although the right-centre region also has good quality data, it has not been processed so far because there happens to be no good seed planted inside it. Eventually, the central tracked region is guided to the right-centre region of high quality data (Figure 5(d)). All well-correlated data is now processed and the tracking progresses to the noisy arc (Figure 5(e)). At this stage, some invalid seeds inside the arc start to grow because the maximum quality of the remaining, unprocessed points is nearly zero: hence, invalid seeds with comparable quality get selected for processing. Figure 5(f) shows the final propagation map, which is highly informative since it indicates the regions grown from each seed. For example, there are two bright patches inside the arc in Figure 5(a) that correspond to regions grown from two invalid seeds in Figure 5(f). The displacement error above the arc can be traced to a bad seed at the same position on the propagation map. The bad seed can be further traced back to a false correlation peak chosen by the exhaustive search algorithm. Since the correlation peak was nevertheless rather high, the seed had a chance to grow a little before fizzling out as the quality at its boundaries dropped.

The second pass of the multiple-seed algorithm fixes the problem with minimal extra computation. A correct displacement distribution is retrieved, as shown in Figure 6(a). The second pass repeats the displacement estimation process on points grown from seeds that form an area smaller than a threshold, in this case 5% of the whole frame. Initial displacement estimates for these reprocessed regions are obtained at their boundaries using information propagated from other seeds. For this particular data set, the 5% test eliminates all but three of the original twenty seeds. Figures 6(b)–(f) show the progress of the second pass: note how the surviving three seeds propagate regions previously processed by the discarded seeds. While Figures 6(b)–(f) hint at similar computational expense to the first pass, this is not in fact the case. The tracking is indeed repeated in its entirety, albeit from a reduced seed set, but the expensive displacement calculations are performed only in the small regions flagged for reprocessing. As explained in Section 2.4, time spent on tracking can be safely ignored in an efficient implementation. The second pass, therefore, is responsible for only a small part of the overall computational cost.

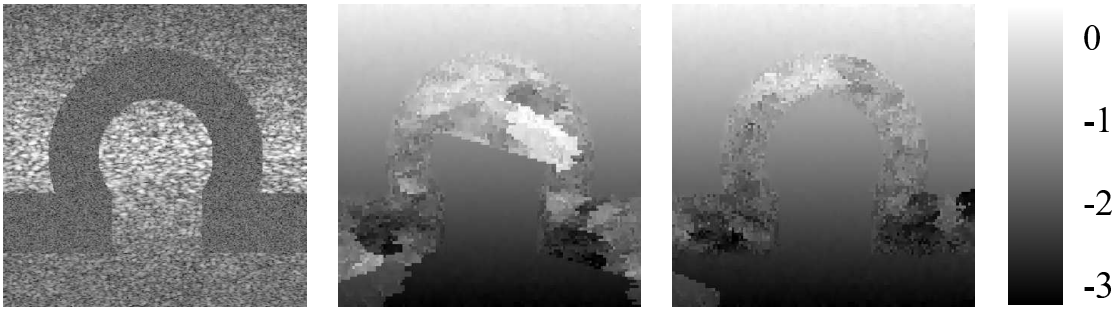

(a) (b) (c)

Figure 7: (a) B-mode image of the simulated data. The Ω-shaped region contains only noise. Outside the Ω-shaped region, white noise is superimposed on the RF signal, with less noise at the centre than at the axial extremities. Displacement obtained by (b) the first pass and (c) the second pass of the drop-out correction method.

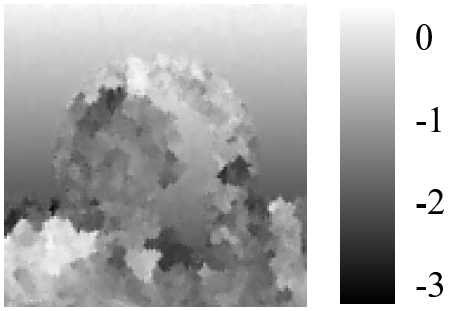

Figure 8: Same data as in Figure 7(a). Displacement obtained by the single-seed algorithm.

3.2 Tracking of disjoint regions

In the second Field II simulation (Figure 7(a)), the noisy arc was extended sideways to partition the well-correlated data into two disjoint regions. Such a situation is not uncommon when scanning *in vivo*, as we shall see in Section 3.3. To succeed, a tracking algorithm will need to recover the displacements in the two regions separately.

Results obtained by the drop-out correction method are shown in Figures 7(b)–(c). The surprising element of these results is the unexpected resumption of correct tracking in the first pass below the noisy region, caused by a lucky good match left-of-centre that subsequently propagates down and to the right. In the second pass (Figure 7(c)), bottom-to-top tracking propagates good displacement estimates to the entire region enclosed by the arc. The remaining poor displacement estimates at the bottom left corner could be corrected with a third, top-to-bottom pass.

The single-seed quality-guided tracking algorithm does not recover the correct displacement below the noisy region — see Figure 8. The seed starts propagating from above the noisy arc. When the tracking path reaches the arc, estimation errors arise and are propagated into the region below. This illustrates the danger of relying on displacement continuity over the whole frame. In such situations, a multiple-seed approach is required for successful displacement estimation.

Figure 9(a) shows the displacement distribution retrieved by the first pass of the multiple-seed tracking algorithm. Accurate displacement estimates are obtained throughout the two well-correlated regions. Figures 9(b)–(d) show how both regions benefit from at least one good seed that propagate independently from each other. When both well-correlated regions are fully processed, the algorithm penetrates the noisy region, the quality drops dramatically and spurious seeds inside the noisy region get a chance to grow, producing the bright patches in Figure 9(a).

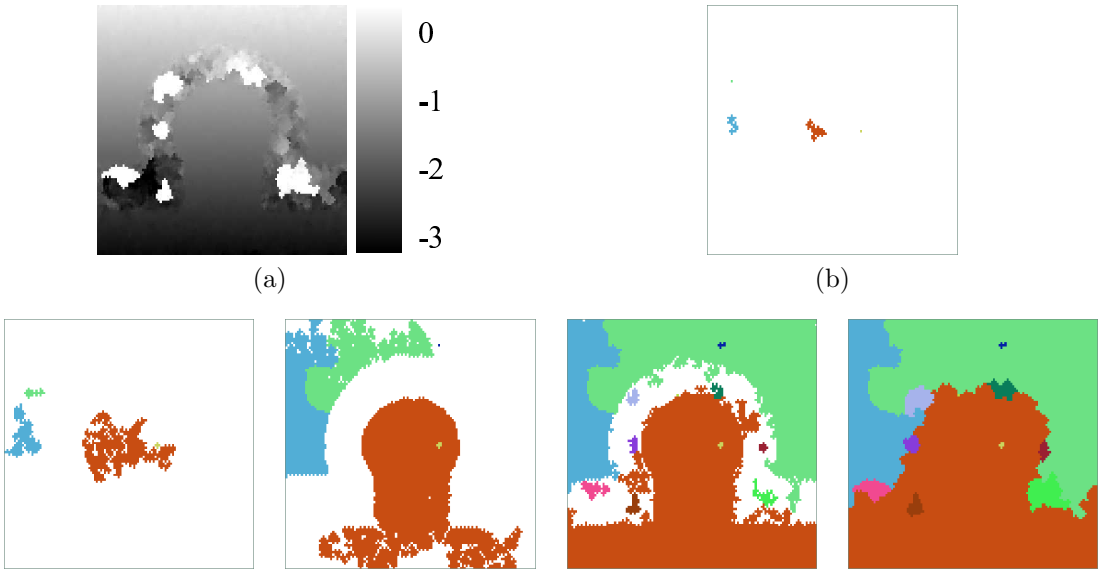

Figure 9: Same data as in Figure 7(a). (a) Displacement obtained by the first pass of the multiple-seed algorithm. (b)–(f) show the propagation progress.

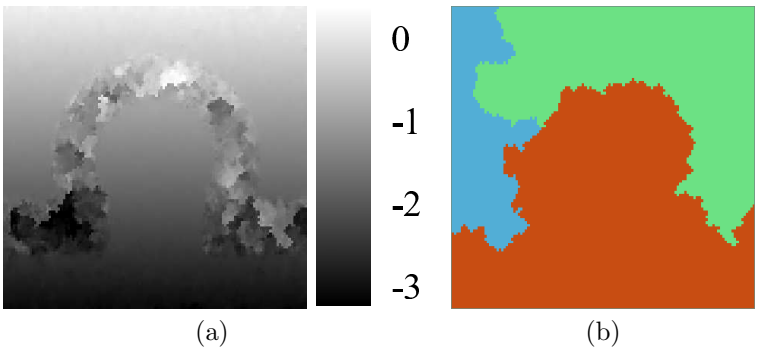

Figure 10: Same data as in Figure 7(a). (a) Displacement obtained by the second pass of the multiple-seed algorithm. (b) seed propagation map.

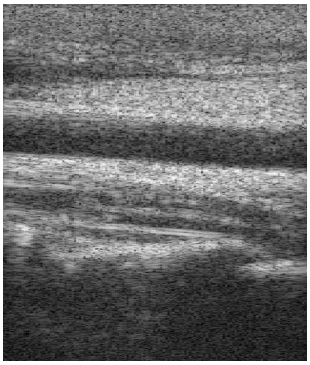

Figure 11: B-mode image of the carotid artery scan.

These white patches fall below the 5% threshold and are therefore among the regions reprocessed in the second pass of the multiple-seed algorithm (Figure 10(a)), which starts from just three seeds (Figure 10(b)). Although in this case reprocessing is not absolutely necessary, in general the second pass does help to correct discontinuous regions grown from bad seeds. Comparing the propagation maps of the two passes, it is apparent that the three seeds used in the second pass have invaded some of each other's territory, as well as that of the eliminated seeds. However, the displacement is only recalculated in those small regions flagged for reprocessing. Doing otherwise — for example, overwriting good quality displacements on a per-point basis with better quality displacements — would not only slow down the algorithm but also introduce unwanted, small discontinuities in otherwise continuously tracked regions.

3.3 *In vivo* displacement tracking

Figure 11 shows a B-mode image of a human carotid artery. The artery separates well-correlated data into two regions. Inside the artery, the scattering is weak and blood flow causes significant decorrelation between the pre- and post-deformation frames. A further challenge to displacement estimation is the pulsation of the arterial wall, which produces a discontinuous deformation distribution in its vicinity. Both these factors point to the need for a robust tracking algorithm.

Figures 12(a)–(b) shows the displacement and quality obtained by the first pass of the drop-out correction method. Displacement estimation above the artery is correct, but the algorithm loses track as it encounters the decorrelated data inside the artery. Nevertheless, a chance good match below the artery is propagated downwards left and right, with correct tracking reestablished across the full lateral range in the bottom third of the frame. The second, bottom-to-top pass, shown in Figures 12(c)–(d), corrects most of the remaining poor estimates below the artery, apart from a small region at the right of the frame, just below half way down. In this region, poor estimates from the first pass just happen to have higher quality, and are therefore not overwritten by better estimates propagating up from below. This definite tracking error should be contrasted with the spurious displacement estimates in the bottom quarter of the frame, where the weak backscatter signal is so dominated by noise that accurate displacement estimation is impossible, whatever the tracking algorithm.

Results obtained by the single-seed quality-guided algorithm are poor in this case, as shown in Figure 13. Accurate displacement estimates are obtained only in the region below the artery, which is where the seed was planted. Estimation errors inside the artery propagate upwards, with subsequent tracking failing outright.

The first pass of the multiple-seed tracking algorithm produces promising results, with plausible displacements both above and below the artery — see Figure 14(a). The early propagation map (Figure 14(b)) shows several seeds planted on either side of the artery. These propagate independently (Figure 14(c)) until all the well-correlated data is processed (Figure 14(d)). The

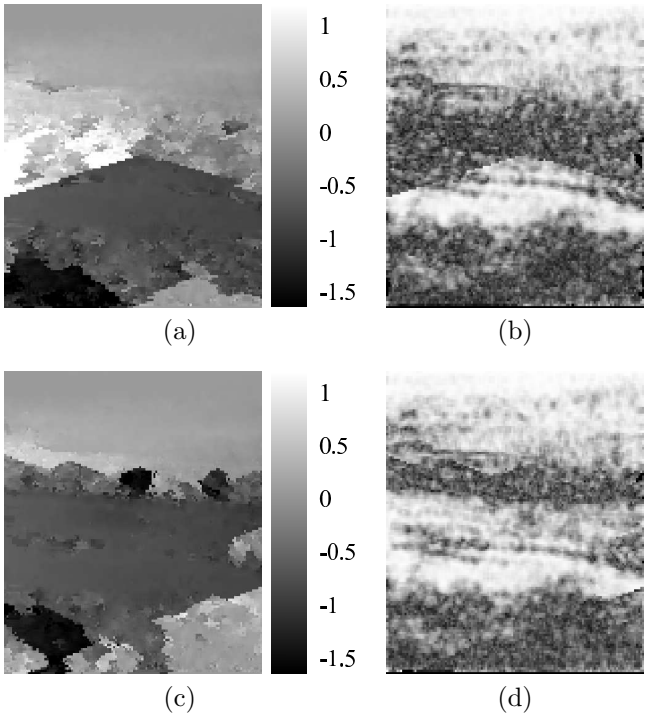

Figure 12: (a) Displacement and (b) quality map of the first pass of the drop-out correction method. (c) Displacement and (d) quality map of the second pass.

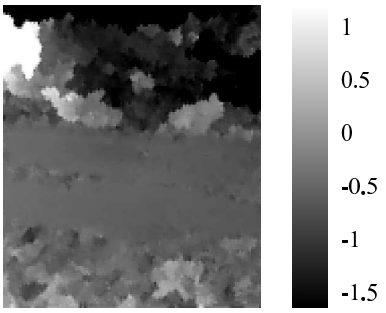

Figure 13: Same data as in Figure 11. Displacement obtained by the single-seed algorithm.

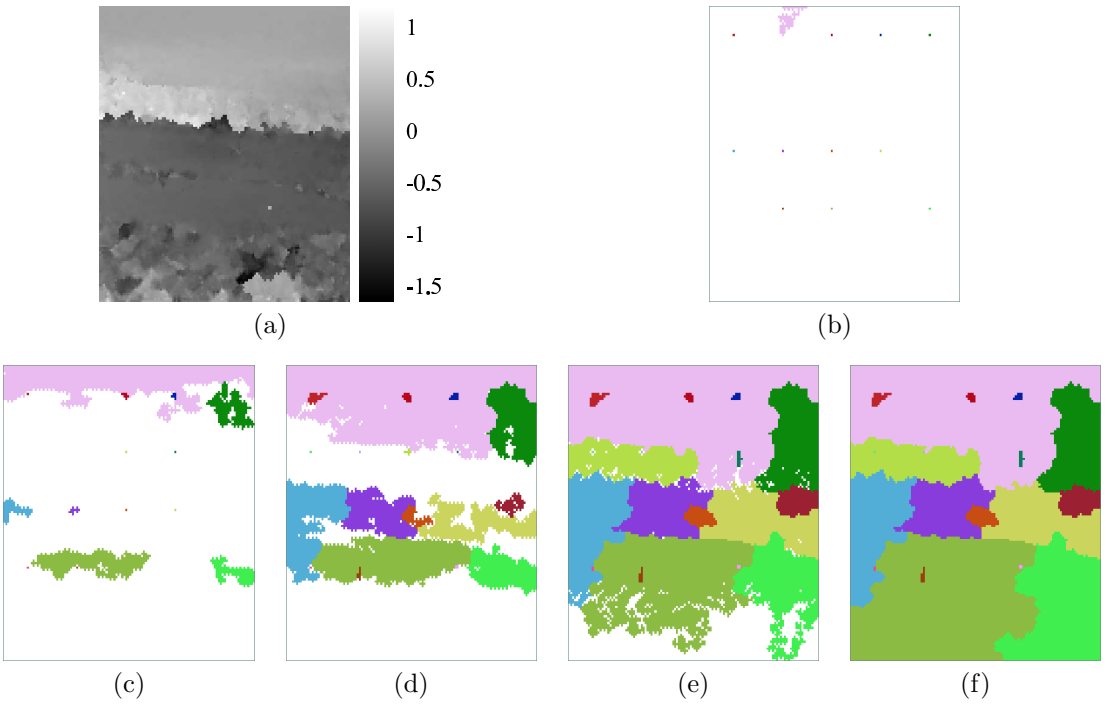

Figure 14: Same data as in Figure 11. (a) Displacement obtained by the first pass of the multiple-seed algorithm. (b)–(f) show the propagation progress.

tracking path then proceeds to the poorly correlated data inside the artery (Figure 14(e)) and finally to the even lower quality data at the bottom of the frame (Figure 14(f)). However, there is a visible error — a tiny bright dot — below the artery in Figure 14(a), just to the right of centre. The final propagation map (Figure 14(f)) shows the bad seed responsible for this error. The growth of the seed is first evident in Figure 14(d), at the stage where the tracking algorithm had exhausted most of the high quality data.

Figure 15 shows how this error is corrected, with negligible computational overhead, by the second pass. An important difference between the second pass of the multiple-seed algorithm and that of the drop-out correction method is that the former uses seed propagation area to flag possible errors, while the latter relies on a further quality comparison. As we have seen at several points in this study, and as has been noted previously elsewhere [25], high quality does not guarantee correct estimation. In the quality-guided approach, the quality is only used to indicate the sequence of processing, not as a criterion to select or remove individual displacement estimates.

3.4 The effect of SNR and window length

It remains to compare the various tracking strategies in terms of their robustness to different SNR levels and different RF window lengths. For any given amount of strain and SNR, there are lower and upper bounds on the RF window lengths that allow accurate displacement estimation [13]. Too short a window and the correct match will have much the same correlation as incorrect matches one or more RF cycles away: this is especially evident at low SNR. Conversely, too long a window and intra-window deformation will lower the correlation of the correct match, making it less distinguishable from incorrect matches: this is especially evident at high strain. We demonstrate in this section that the tracking algorithm plays its part too: more sophisticated approaches allow a wider range of window lengths for any given strain and SNR.

Two RF frames were simulated with uniform echo amplitude and a 1% strain field. Random white noise was added at six SNR levels (4 dB, 2 dB, 0 dB, −2 dB, −4 dB, and −6 dB) to evaluate

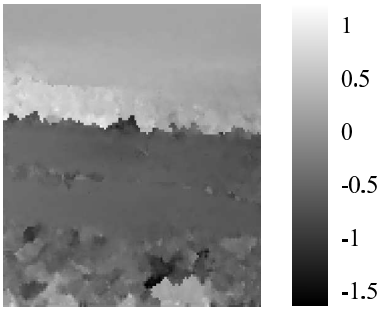

Figure 15: Same data as in Figure 11. Displacement obtained by the second pass of the multiple-seed algorithm.

the tracking capability. Unlike the previous tests, there is no completely decorrelated region and the SNR is uniform across each frame. The experiment can therefore focus on the core ability of each tracking algorithm to recover a continuous displacement field in the presence of noise, without the extra challenges of severe, local data decorrelation and displacement discontinuity.

At each SNR level, four tracking strategies were tested: tracking along A-lines, the drop-out correction method, the single-seed quality-guided algorithm and the multiple-seed quality-guided algorithm. The RF window length of the underlying displacement estimation algorithm (phase zero search) was varied between 1 and 50 cycles. We deemed each experiment a success if more than 90% of the displacement estimates were correct, where "correct" means within a quarter cycle of the known, true value. Although arbitrary, this criterion allows a meaningful comparison of the various tracking algorithms. As expected, at each noise level, there were clear upper and lower bounds on the window lengths that allowed successful recovery of the displacement field. These bounds are plotted in Figure 16. The area enclosed by the lower and upper bounds provides a measure of each algorithm's robustness to different SNRs and window lengths. We call this area the *effective tracking zone*. Although the results in Figure 16 are based on 1% strain, other strain levels produced similar patterns.

The simplest strategy, tracking along A-lines, has the smallest effective tracking zone. If an estimation error occurs, the method has no mechanism to prevent it from propagating. The single-seed algorithm performs poorly as well. At low SNR levels (below 0 dB), the single seed is often bad, even though it has the highest quality amongst all the grid points. When this happens, tracking fails immediately and never recovers. This is an obvious brittleness of the single-seed approach that we have deferred mentioning until now.

Although the single-seed algorithm is superior to the drop-out correction method in tracking a geometrically irregular region, its weakness lies in its poor tolerance of noise. The latter, on the other hand, is more robust against noise because it searches several previously processed points for the best initial guess at each new point. The fragility of the single-seed algorithm is overcome in its multiple-seed variant. If the region grown from one seed goes wrong, others can still proceed correctly. The multiple-seed algorithm has a comparable effective tracking zone to that of the drop-out correction method, with the advantage that it can process disjoint and irregularly shaped regions.

4 Conclusion

Displacement tracking is an algorithmic procedure implicit in the data processing pipelines of all ultrasonic strain imaging systems. Despite playing a large part in the efficiency and stability of the displacement estimation process, previous work on tracking has been tightly interwoven with particular displacement estimation techniques. In this paper, a generic quality-guided tracking framework has been proposed. It can be integrated with a wide variety of displacement estima-

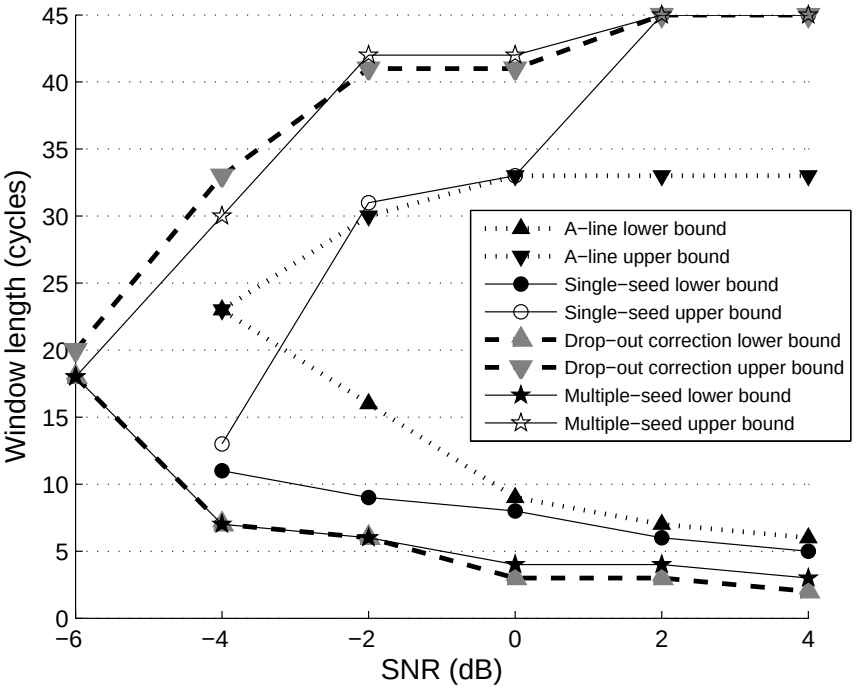

Figure 16: Effective tracking zone.

tion techniques. The single-seed algorithm demonstrates the fundamental working mechanism of the quality-guided tracking strategy. The multiple-seed variant greatly enhances the algorithm's robustness to noise and data discontinuity.

Further research on displacement tracking should explore alternative quality indicators. The measure used in this study, the complex cross-correlation coefficient, was generally effective but less so at low SNR. Alternative approaches, perhaps including a continuity element, might perform better in high noise situations. Another area for further research concerns the spatial resolution of strain imaging systems. Sophisticated tracking strategies allow the use of shorter RF windows. This should be exploited to the full, with a quantitative study of the achievable spatial resolution of the subsequent strain images. Finally, the multiple-seed version of the quality-guided tracking algorithm produces, as a by-product, an approximate segmentation of the frame into regions of continuous displacement. This segmentation could be refined and then used to improve the subsequent calculation of strain. Specifically, differentiation across displacement discontinuities could be avoided, thus eliminating some of the artefacts that commonly affect strain images.

Acknowledgements

This work was funded by Translation Award 081511/Z/06/Z from the Wellcome Trust.

References

[1] N. Belaïd, I. Cespedes, M. Thijssen, and J. Ophir. Lesion detection in simulated elastographic and echographic images: a psychophysical study. *Ultrasound in Medicine and Biology*, 20:877–891, 1994.

[2] J. Bercoff, M. Tanter, and M. Fink. Supersonic shear imaging: a new technique for soft tissue elasticity mapping. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51:396–409, 2004.

[3] S. Catheline, F. Wu, and M. Fink. A solution to diffraction biases in sonoelasticity: the acoustic impulse technique. *Journal of the Acoustical Society of America*, 105:2941–2950, 1999.

[4] I. Cespedes, Y. Huang, J. Ophir, and S. Spratt. Methods for estimation of subsample time delays of digitized echo signals. *Ultrasonic Imaging*, 17:142–171, 1995.

[5] 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.

[6] L. Gao, K. J. Parker, and S. F. Levinson. Imaging of the elastic properties of tissue — a review. *Ultrasound in Medicine and Biology*, 22:959–977, 1996.

[7] B. S. Garra, E. I. Cespedes, J. Ophir, S. R. Spratt, R. A. Zuurbier, C. M. Magnant, and M. F. Pennanen. Elastography of breast lesions: initial clinical results. *Radiology*, 202:79–86, 1997.

[8] D. C. Ghiglia and M. D. Pritt. *Two-Dimensional Phase Unwrapping*. Wiley, New York, 1998.

[9] T. J. Hall, Y. Zhu, and C. S. Spalding. In vivo real-time freehand palpation imaging. *Ultrasound in Medicine and Biology*, 29:427–435, 2003.

[10] 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.

[11] 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.

[12] S. Langeland, J. d'Hooge, H. Torp, B. Bijnens, and P. Suetens. Comparison of time-domain displacement estimators for two-dimensional rf tracking. *Ultrasound in Medicine and Biology*, 29:1177–1186, 2003.

[13] 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*, in press.

[14] J. E. Lindop, G. M. Treece, A. H. Gee, and R. W. Prager. 3D elastography using freehand ultrasound. *Ultrasound in Medicine and Biology*, 32(4):529–545, March 2006.

[15] R. K. Nightingale, M. L. Palmeri, R. W. Nightingale, and G. E. Trahey. On the feasibility of remote palpation using acoustic radiation force. *Journal of the Acoustical Society of America*, 110:625–634, 2001.

[16] J. Ophir, I. Cespedes, H. Ponnekanti, Y. Yazdi, and X. Li. Elastography: a quantitative method for imaging the elastisticity of biological tissues. *Ultrasonic Imaging*, 13:111–134, 1991.

[17] M. ODonnell, A. R. Skovoroda, and B. M. Shapo. Measurement of arterial wall motion using Fourier based speckle tracking algorithms. In *Proceedings of the IEEE Ultrasonics Symposium*, pages 1101–1104, 1991.

[18] M. ODonnell, 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.

[19] 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.

[20] 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.

[21] H. Shi and T. Varghese. Two-dimensional multi-level strain estimation for discontinuous tissue. *Physics in Medicine and Biology*, 52:389–401, 2007.

[22] 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*, in press.

[23] 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.

[24] F. Viola and W. F. Walker. A comparison of the performance of time-delay estimators in medical ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 50:392–401, 2003.

[25] 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.

[26] 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.

[27] 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.

[28] Y. Zhu and T. J. Hall. A modified block matching method for real-time freehand strain imaging. *Ultrasonic Imaging*, 24:161–176, 2002.