RAPID HYBRID ULTRASOUND VOLUME REGISTRATION

U. Z. Ijaz, R. W. Prager,
A. H. Gee and G. M. Treece

CUED/F-INFENG/TR 644

March 2010

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

Email: uzi20/rwp/ahg/gmt11 @eng.cam.ac.uk

Rapid hybrid ultrasound volume registration

Umer Zeeshan Ijaz, Richard W. Prager, Andrew H. Gee and Graham M. Treece

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

Abstract

This work is concerned with registration of ultrasound volumes acquired using a mechanically-swept 3D probe to produce extended-field-of-view images. While the registration can be achieved by attaching a position sensor to the probe, this can be an inconvenience in a clinical setting and aligns the volumes incorrectly if there is tissue motion due to subject movement or respiratory motion. The objective of this work is to replace the 6 degree of freedom (DOF) sensor with a combination of 3 DOF image registration and an unobtrusively integrated inertial sensor for measuring orientation. We aim to produce a highly reliable system that is able to register a pair of volumes very quickly, making it suitable for clinical use. As such, this is an extension of our previous work [1] and provides further results for a fast implementation of the Nelder-Mead simplex algorithm using reslices. Additionally, we provide some results for a voxel-array based multiresolution search algorithm. The multiresolution approach has a more complete coverage of the search space at the expense of speed and is more exhaustive in comparison to the Nelder-Mead Simplex algorithm. With similar reliability statistics for the two algorithms, both are possible candidates for a clinical system.

1 Introduction

Currently, 3D ultrasound offers superior and efficient volumetric imaging at a lower cost than other imaging modalities. However, with the limited field of view of the acquired volumes, there is a great deal of interest in creation of 3D mosaics by stitching together consecutively recorded volumes. This mosaicing will provide the sonographer with a compounded volume of higher quality and the ability to visualise anatomical structures from a variety of different angles. For this purpose, we propose a hybrid intensity-based image registration strategy. In multi-volume registration, we consider the notion of a similarity measure between two volumes, where the target volume is kept stationary while the position of the source volume is changed iteratively until an extreme value of the similarity measure designates a registration point. In this regard, we have made the following assumptions:

- The probe is held stationary while recording each volume so there is no need for intra-volume registration.
- There is a reasonable overlap between the volumes.
- The relative orientation of the volumes is already known through the use of an inertial sensor.
- Any distortion of the tissue is small and therefore only a rigid body transformation is required for successful registration.
- Although it is reasonable to assume that the subsequent volumes will be in roughly the same direction of movement as the first pair of volumes, we have considered an undirected approach.

Even though we have whole volumes of data available, we have chosen to use only a subset of this data in order to speed up the computation in the search algorithm. As a result, we consider a slice-based strategy in which we evaluate the similarity measure on a subset of images across the region-of-interest. This reduces the 3D data to a set of 2D images and makes the algorithms acceptably fast. In order to provide invariance to overlap statistics, normalised mutual information is used as a similarity measure [2].

Furthermore, we have considered two approaches to tackle image based volume registration. In the first approach, we have used a very efficient reslicing algorithm [3] to generate vertical and horizontal reslices orthogonal to the central frame of the source volume (see Fig. 9). Thus we obtain two sets of images from two sweeps across the region-of-interest. Henceforth, they are refered to as *reslice pairs*. Using this approach, we compared the performance of several iterative algorithms on *in vitro* datasets. Fig. 1 reproduces some of the results from [1] to provide the context for this development. Based on this study, we found the particle swarm optimiser to be the most reliable algorithm. However, it is inherently slow as it maintains a higher number of multiple solutions. Therefore, we rejected it along with Powell's direction-set algorithm mainly on account of computational time. The Levenberg-Marquardt method is the fastest of the algorithms considered, but it is not very reliable. It works on the principle of reducing the sum of squares of residuals. In order to use this algorithm, we must calculate the similarity measure separately for individual slice pairs which makes it less robust. Thus we are left with the Nelder-Mead simplex algorithm for further consideration which is not only reliable but also fast. The Nelder-Mead simplex algorithm for a 3 DOF search requires only four initial position estimates (which can be visualised as a tetrahedron). It requires only function evaluations, not derivatives and is faster than other multiple solution algorithms. However, it can get stuck in a local minimum. We have therefore used a heuristic approach based on multiple restarts, considering the local minimum as one vertex of the simplex and reinitialising the other vertices within a certain range of the claimed minimum. In comparison to [1], we have considered an efficient implementation of the Nelder-Mead simplex algorithm in this paper. We consider a larger number of datasets and analyse the reliability of the Nelder-Mead simplex algorithm against the amount of data used to calculate the similarity measure and the initialisation range in the axial direction. This is done with the intention of speeding up the algorithm.

Although, the Nelder-Mead simplex algorithm satisfies most of our requirements, we also consider a voxel-array based approach similar to the one considered in [4] as an alternative approach. One of the disadvantages of the algorithms considered in [1] was that they required the reslices to be calculated at each new position. Since the orientation information of the volumes is known *a priori*, a one-time generation of voxel-arrays and a multiresolution approach can avoid the need for recalculating reslices and may have a speed advantage. The details are given in the next section.

2 Voxel-array based multiresolution search algorithm

This algorithm was originally used in [4] for rigid correction of misaligned pairs of sweeps initially aligned using a spatial locator. It made use of a single reslice through the overlap region which required some manual intervention to position correctly. However, in our implementation, we want to avoid any manual intervention and have no initial information on the alignment. Therefore, we must use multiple slices distributed over the volume to be sure that a few slices fall within the overlap region. The original algorithm made use of 2D smoothing of the reslice image. We require 3D smoothing of the voxel-arrays because slices are taken in all three directions.

The working of the multiresolution search algorithm is given in Figs. 2 and 3. The algorithm works by first aligning the centres of the source and target volumes, and then by fitting a minimum-voxels bounding reslice. The size and orientation of this reslice is chosen automatically such that it encompasses all of the data in the minimum number of voxels[1]. The voxel array is obtained

---

[1] An alternative approach would have been to create voxel-arrays within the extent of one of the volumes. We have chosen not to do this because larger voxel-arrays give more data samples to calculate the similarity measure. However, larger voxel-arrays will also introduce null-data from outside the volumes (see Fig. 4 for further clarification). Furthermore, useful data is found mostly in the centre of the voxel-arrays particularly when the volumes have different orientations (an extreme case could be volumes that are 90 degrees apart), therefore, the

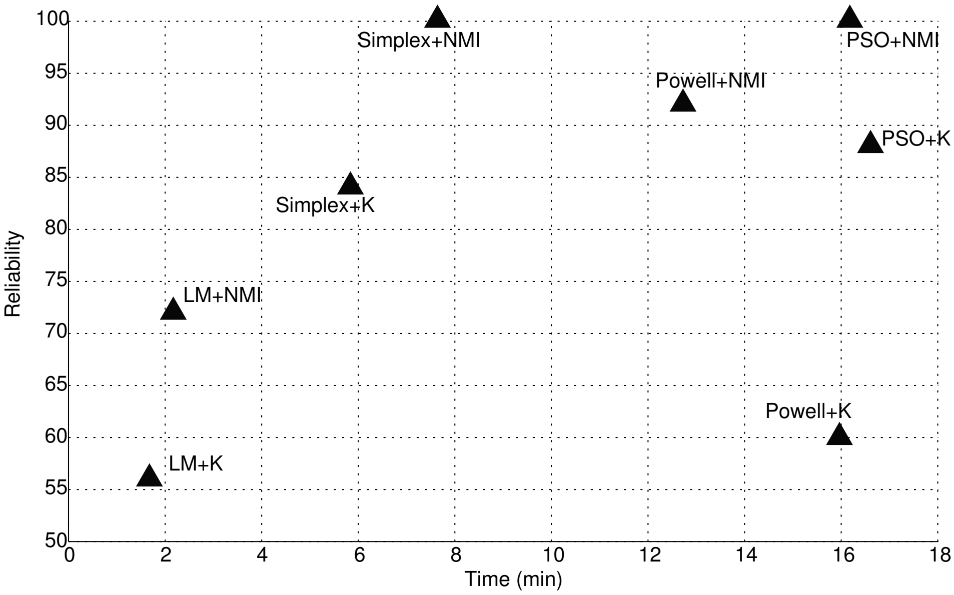

Figure 1: Comparison of performance of different algorithms for 5 *in vitro* datasets with 5 trials per dataset for each algorithm. For the Nelder-Mead simplex algorithm (Simplex), the Levenberg-Marquardt method (LM), and Powell's direction-set method (Powell), the execution time is shown including the two restarts. Here, the restarts are necessary to avoid local optimum and to increase reliability. In case of the particle swarm optimiser (PSO), we have not allowed any restarts. The comparison is provided on two similarity measures, namely, normalised mutual information (NMI), and Kolmogorov's distance (K).

by taking reslices parallel to the minimum-voxels bounding reslice between the minimum (*min_d*) and maximum (*max_d*) extents of the data. In the reslice plane, the resolution is specified by considering every *resxy* pixel(s). In the third direction, the resolution is specified by *resz* slice spacing in cm. After obtaining the two voxel-arrays, a volume pyramid is then constructed by smoothing and downsampling the original voxel-arrays. The search works in 4 levels with data being coarse at L4 and fine at L1. An exhaustive search is done at L4 (considering a minimum of 25% overlap in each dimension) by calculating the similarity measure on the 2D image slices from corresponding volumes specified in all three directions. It should be emphasised here that the locations of the slices are fixed with respect to the source volume. Also, as opposed to the reslices used in the Nelder-Mead simplex algorithm, the width and height of the slices are bounded by the dimensions of the voxel-arrays. After obtaining the optimal solution in L4, the subsequent searches are limited to a narrower window (*step_x*, *step_y*, and *step_z* around the optimum results obtained from the previous level). Finally, the relative difference in the positions of the two voxel arrays is used to update the position of the frames of the source volume.

3 Comparative analysis of the two approaches

For evaluation of the two approaches, ultrasound data sets were recorded using a GE RSP6-12 mechanically-swept 3D probe interfaced to a Dynamic Imaging Diasus US machine. The depth setting was 3 cm with a single focus at 1.5 cm. The B-scan resolution was 0.01 cm/pixel. Analogue RF echo signals were digitised after receive focusing and time-gain compensation, but before log-compression and envelope detection, using a Gage Compuscope CS14200 14-bit digitiser

---

data around the central region should be given more importance.

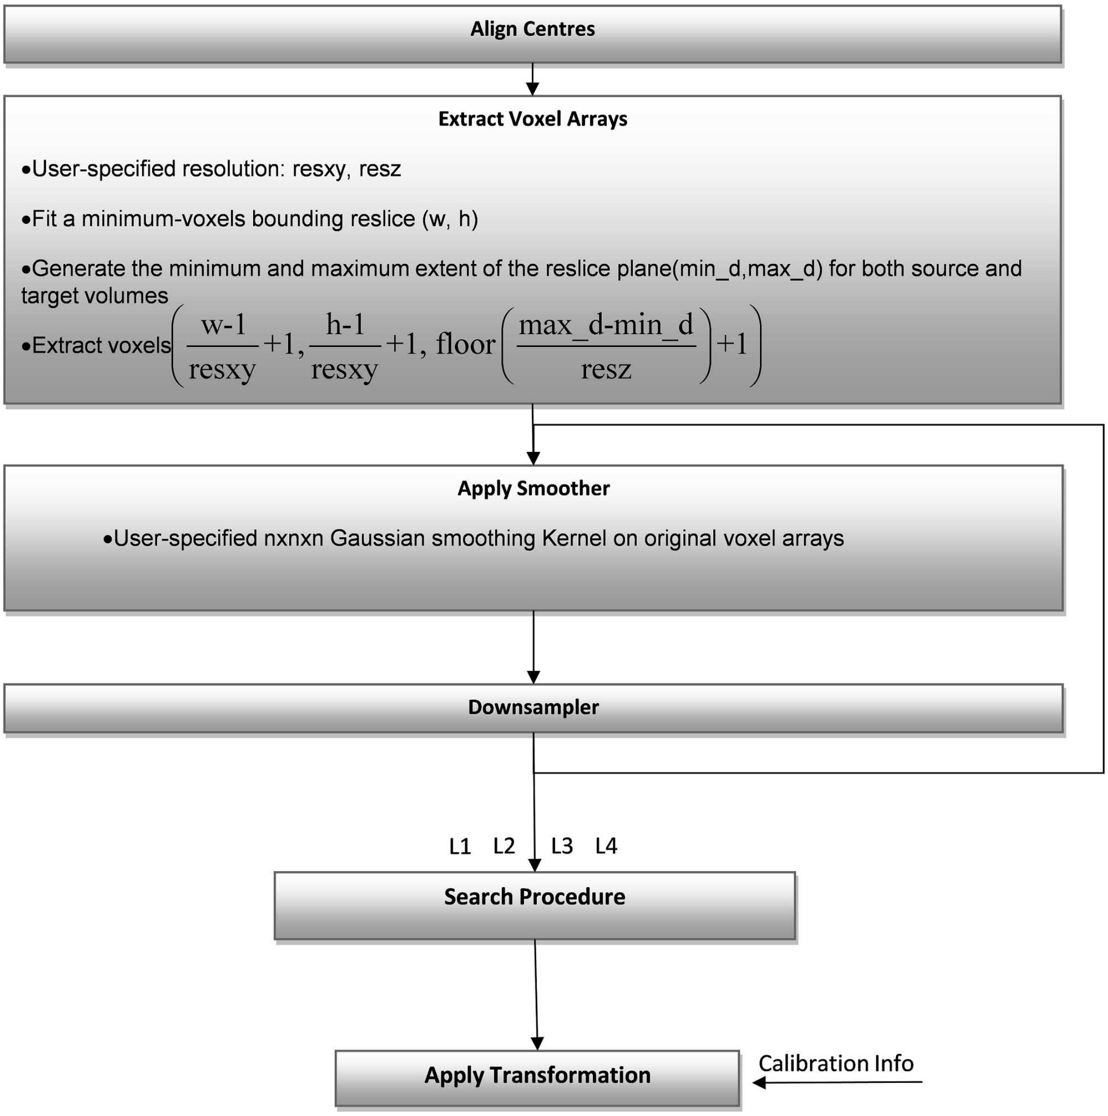

Figure 2: Flow diagram of the voxel-array based multiresolution search algorithm. Here, w, h, and d represent the width, height, and depth of the voxel-array.

(`http://www.gage-applied.com`). The RF data was then converted to B-scan images using Stradwin software (`http://mi.eng.cam.ac.uk/~rwp/stradwin/`). Each data set comprised two volumes, each of 50 frames swept over 10 degrees. We recorded data under three sets of conditions. First, we acquired five *in vitro* scans of a speckle phantom containing several 5mm spherical inclusions (P Class). Here, the ultrasound probe was mounted on a motorised Zaber's linear slide TLSR-300B gantry (`http://www.zaber.com`) to obtain precise translation offsets. We then acquired two classes of *in vivo* data sets, five in which every effort was made to ensure the scanning subject was stationary (R class in which breathing or movement was minimised at the time of acquisition) and six belonging to M class in which small movement and probe pressure variations were allowed. For these experiments, each B-scan's position and orientation was recorded using a Northern Digital Polaris optical tracking system (`http://www.ndigital.com`).

Comparing to the known offsets in the *in vitro* data, we measured the accuracy of the registration to be within 0.4 mm root mean square error (RMSE) when using 10 horizontal and 10

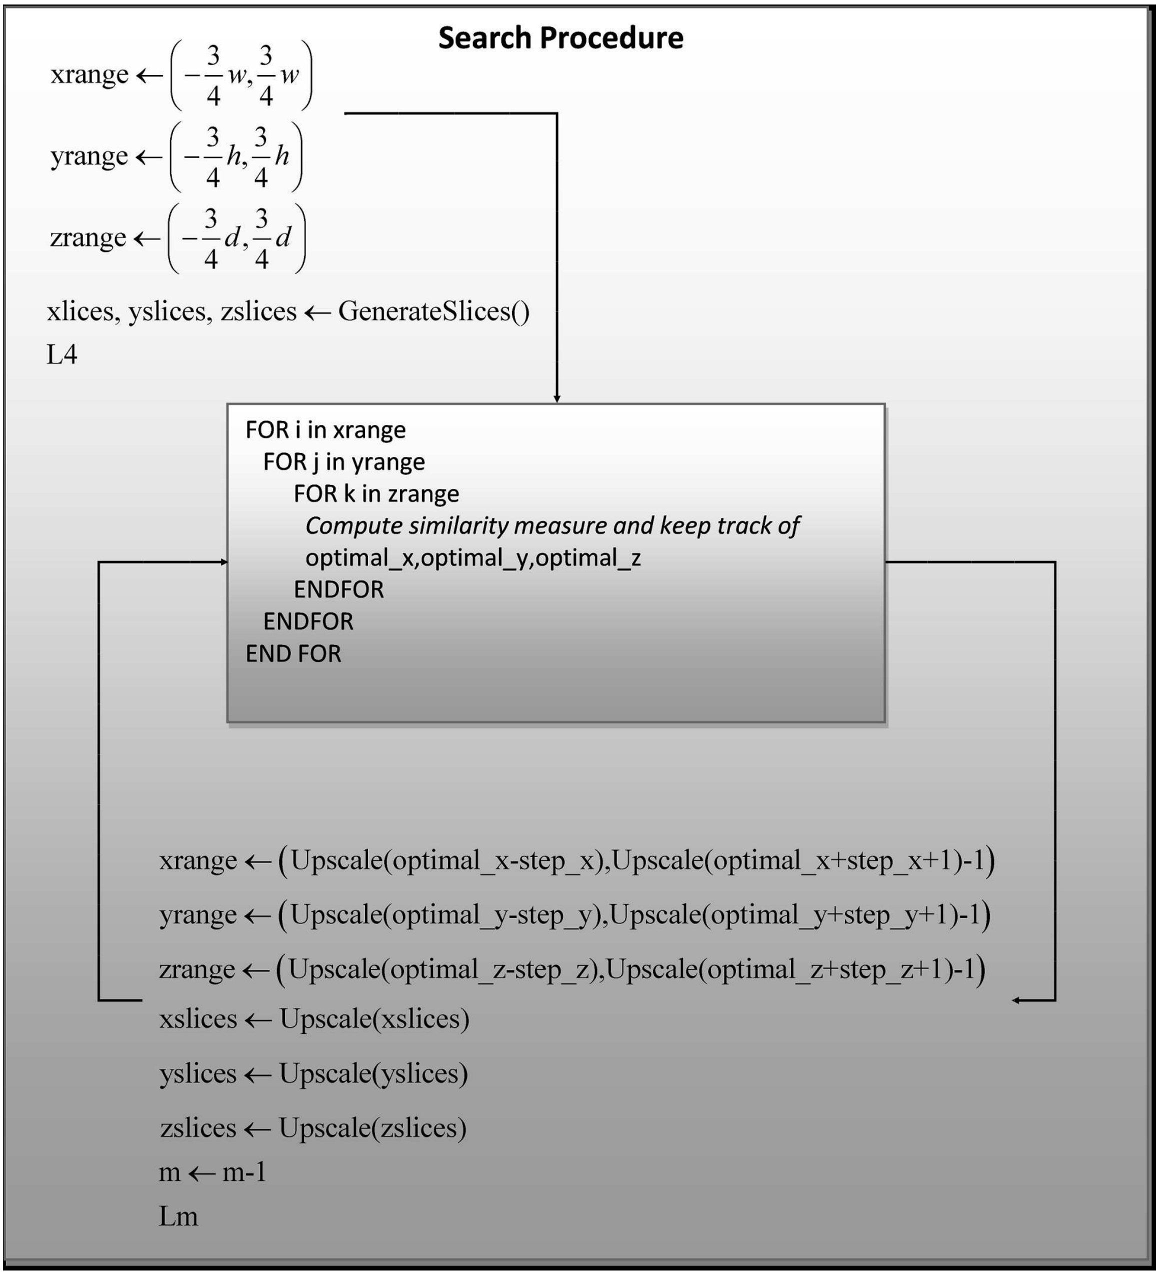

```
xrange ← $\left(-\frac{3}{4}w, \frac{3}{4}w\right)$
yrange ← $\left(-\frac{3}{4}h, \frac{3}{4}h\right)$
zrange ← $\left(-\frac{3}{4}d, \frac{3}{4}d\right)$
xlices, yslices, zslices ← GenerateSlices()
L4

FOR i in xrange
  FOR j in yrange
     FOR k in zrange
       Compute similarity measure and keep track of
       optimal_x,optimal_y,optimal_z
     ENDFOR
  ENDFOR
END FOR

xrange ← (Upscale(optimal_x-step_x),Upscale(optimal_x+step_x+1)-1)
yrange ← (Upscale(optimal_y-step_y),Upscale(optimal_y+step_y+1)-1)
zrange ← (Upscale(optimal_z-step_z),Upscale(optimal_z+step_z+1)-1)
xslices ← Upscale(xslices)
yslices ← Upscale(yslices)
zslices ← Upscale(zslices)
m ← m-1
Lm
```

Figure 3: Flow diagram of search procedure in the voxel-array based multiresolution search algorithm. Here, the left pointing arrow is the assignment operator and the *Upscale* function is used to map the indices of the voxel arrays from current level to next higher level in the volume pyramid.

vertical reslices and two restarts of the Nelder-Mead simplex algorithm [1]. There are two ways that we can increase the reliability: we can increase the number of reslices (see Fig. 6), and we can decrease the initialisation range in the axial direction (see Fig. 7). Increasing the number of reslices increases the chance that the reslices will contain significant features of the data making it easier to find the correct alignment. We observed that in typical datasets, the features were more defined axially giving a sharper peak in the similarity metric in this direction (see Fig. 8). In order to align the data, at least one of the vertices of the Nelder-Mead simplex algorithm must be

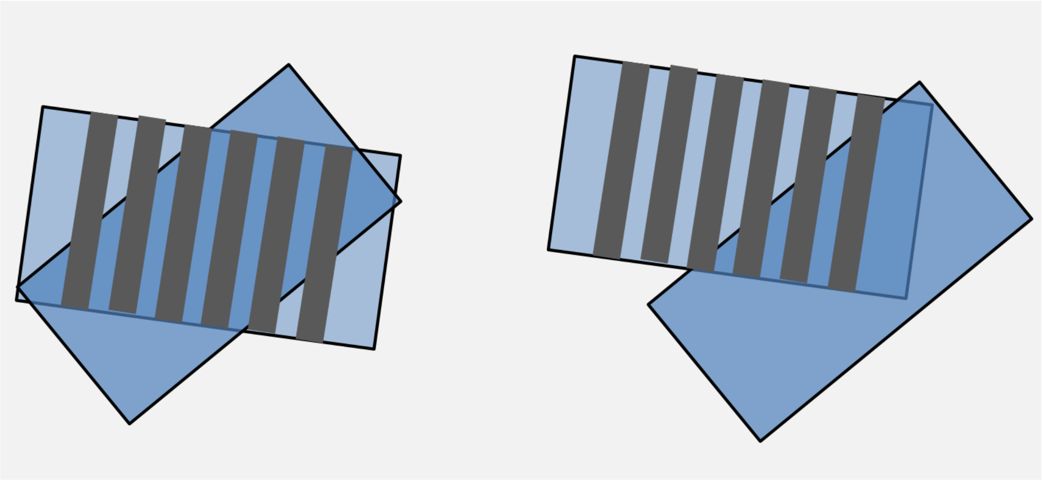

(a) Nelder-Mead simplex algorithm using reslices

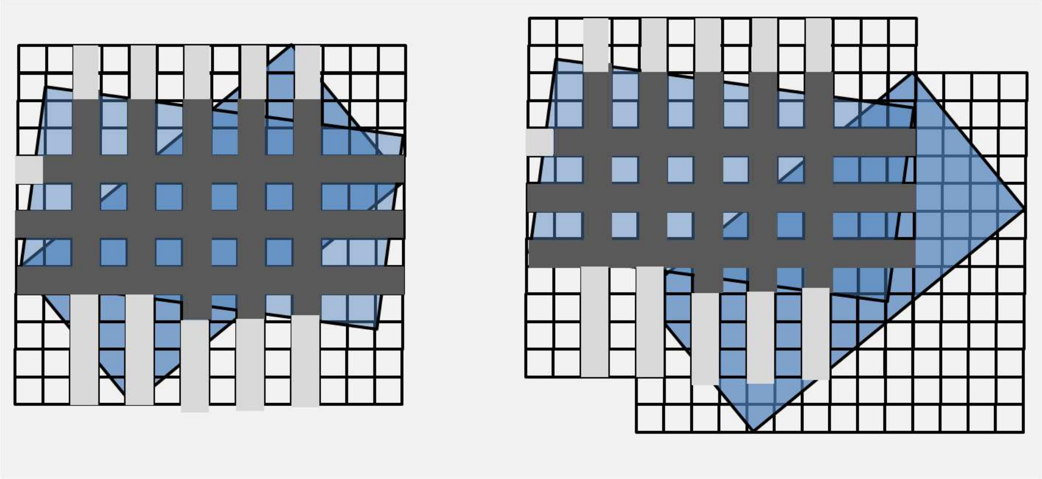

(b) Voxel-array based multiresolution search algorithm

Figure 4: **Explanation of generation of slices in both approaches**. Assuming that the direction of insonification is towards the page and we are viewing the two volumes from the top, the slices are represented by darker strips on the source volume. In both approaches, the left figure shows how the data will look at the start when the two volumes are centrally aligned, and the right figure shows an arbitrary position of the source volume in the vicinity of the target volume. Note that in the current implementation, the similarity measure is calculated on the portions of the reslice that coincides with the source volume. It may also be possible to include the lighter region of reslices in the calculation of the similarity measure. In such a case, the normalised mutual information remains unaffected when the field-of-view is varied i.e. when the lighter region ([2] discusses this invariance in detail) dominates the darker region for the case when the two volumes are further apart.

initialised close to the peak. Reducing the initialisation range in the axial direction increases the chance of starting near the peak and hence improves reliability. Even with only a few reslices (4 horizontal and vertical reslice pairs) and just one restart, we were able to register all the datasets in 8–15 seconds using a C++ implementation of our approach in Stradwin (see Fig. 5) on a 3.0 GHz Intel Core 2 Duo CPU machine.

While it may be a good idea to compare the voxel-array based multiresolution search algorithm with the Nelder-Mead simplex algorithm with slices in the same positions at the highest resolution, it is not possible in the current implementation. The position and orientation of the slices depend

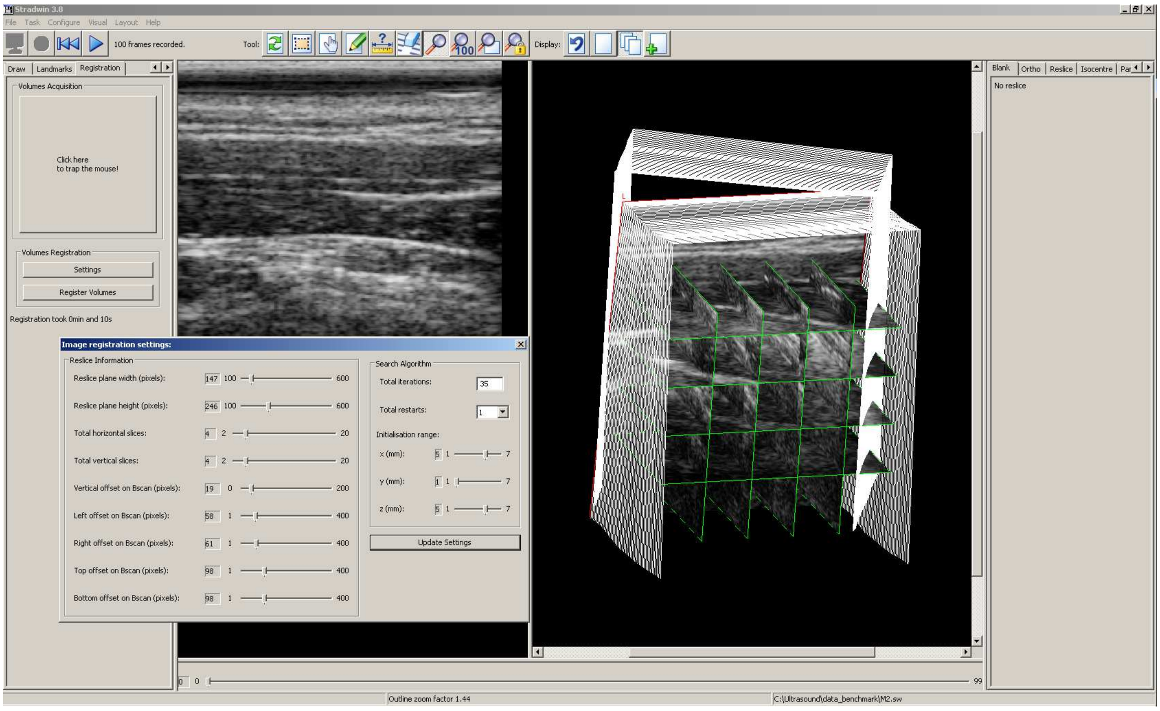

Registration took 0min and 10s

Image registration settings:

Reslice Information

Reslice plane width (pixels): 147 (100–600)
Reslice plane height (pixels): 246 (100–600)
Total horizontal slices: 4 (2–20)
Total vertical slices: 4 (2–20)
Vertical offset on Bscan (pixels): 19 (0–200)
Left offset on Bscan (pixels): 58 (1–400)
Right offset on Bscan (pixels): 61 (1–400)
Top offset on Bscan (pixels): 98 (1–400)
Bottom offset on Bscan (pixels): 98 (1–400)

Search Algorithm

Total iterations: 35
Total restarts: 1
Initialisation range:
x (mm): 5 (1–7)
y (mm): 1 (1–7)
z (mm): 5 (1–7)

(a) Nelder-Mead simplex algorithm using reslices

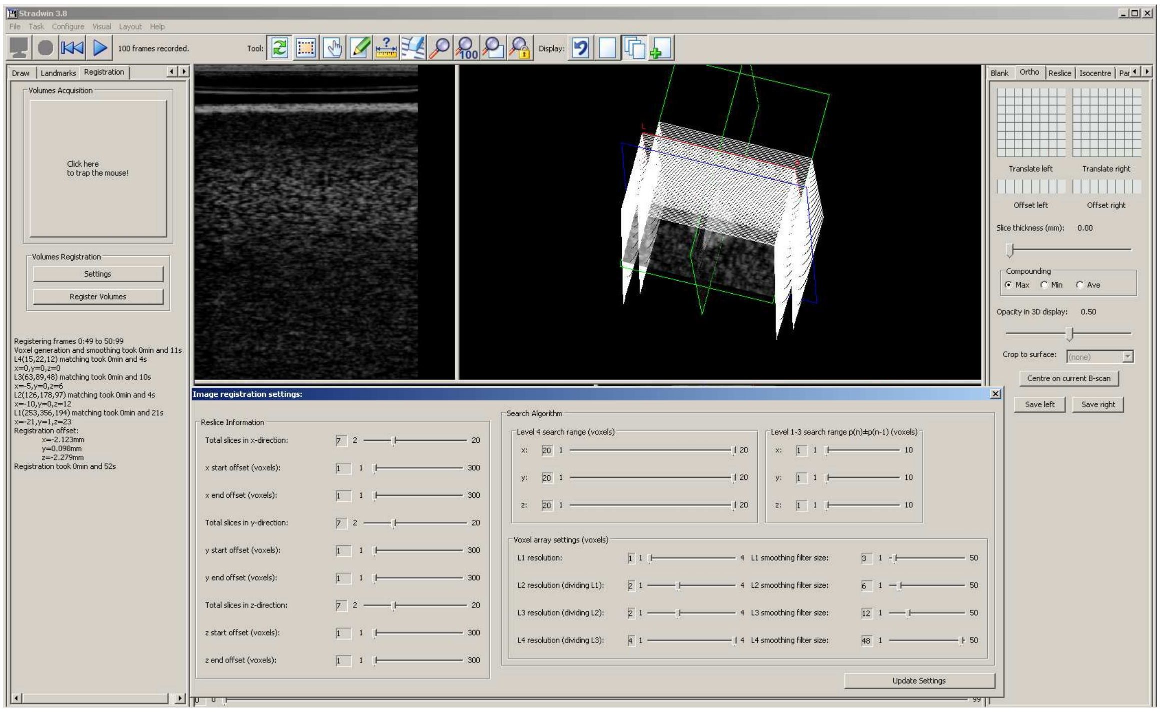

```
Registering frames 0:49 to 50:99
Voxel generation and smoothing took 0min and 11s
L4(15,22,12) matching took 0min and 4s
x=0,y=0,z=0
L3(63,89,48) matching took 0min and 10s
x=-5,y=0,z=6
L2(126,178,97) matching took 0min and 4s
x=-10,y=0,z=12
L1(253,356,194) matching took 0min and 21s
x=-21,y=1,z=23
Registration offset:
          x=-2.123mm
          y=0.098mm
          z=-2.279mm
Registration took 0min and 52s
```

Slice thickness (mm): 0.00
Compounding: ● Max ○ Min ○ Ave
Opacity in 3D display: 0.50
Crop to surface: (none)

Image registration settings:

Reslice Information

Total slices in x-direction: 7 (2–20)
x start offset (voxels): 1 (1–300)
x end offset (voxels): 1 (1–300)
Total slices in y-direction: 7 (2–20)
y start offset (voxels): 1 (1–300)
y end offset (voxels): 1 (1–300)
Total slices in z-direction: 7 (2–20)
z start offset (voxels): 1 (1–300)
z end offset (voxels): 1 (1–300)

Search Algorithm

Level 4 search range (voxels)
x: 20 (1–20)
y: 20 (1–20)
z: 20 (1–20)

Level 1-3 search range p(n)±p(n-1) (voxels)
x: 1 (1–10)
y: 1 (1–10)
z: 1 (1–10)

Voxel array settings (voxels)
L1 resolution: 1 (1–4)
L2 resolution (dividing L1): 2 (1–4)
L3 resolution (dividing L2): 2 (1–4)
L4 resolution (dividing L3): 4 (1–4)
L1 smoothing filter size: 3 (1–50)
L2 smoothing filter size: 6 (1–50)
L3 smoothing filter size: 12 (1–50)
L4 smoothing filter size: 48 (1–50)

(b) Voxel-array based multiresolution search algorithm

Figure 5: **GUI developed for testing the two approaches**. Note that both versions have the functionality for 3 DOF calibration of the orientation sensor.

on the position and orientation of the minimum-voxels bounding reslice around the two volumes. This can change for each dataset depending on the relative orientation of the two volumes. On the other hand, in the Nelder-Mead simplex algorithm, the slices are always generated with respect to the source volume and their sizes can be altered (see the position of reslices in Fig. 4). Nonetheless, we can still compare the algorithms in terms of reliability. The preliminary results for the voxel-array based multiresolution search for two different configurations are shown in Fig. 10. It can be seen here that the reliability values are quite high when the slices are concentrated in the central region (Configuration A). With configuration B, some of the *in vitro* datasets failed because the slices didn't pass through the main features located in the centre. With limited features available in the coarse level, having more slices in the central region of the voxel-arrays also ensures that fewer slices go out of scope (i.e. slices that overlap with null-data of the target volume) when there are large differences in the position between the volumes. With configuration B, one of the failures (M3 dataset) was successfully registered later by reducing the smoothing, which suggests that smoothing may be obscuring some of the features in the data and would be better at a lower setting. Also, one dataset (M4) has two possible positions that give a convincing alignment by visual inspection (see Fig. 11). The Nelder-Mead simplex algorithm found each of these solutions with equal probability, whereas the voxel-array based multiresolution search algorithm found one solution or the other depending on the chosen configuration of the slices.

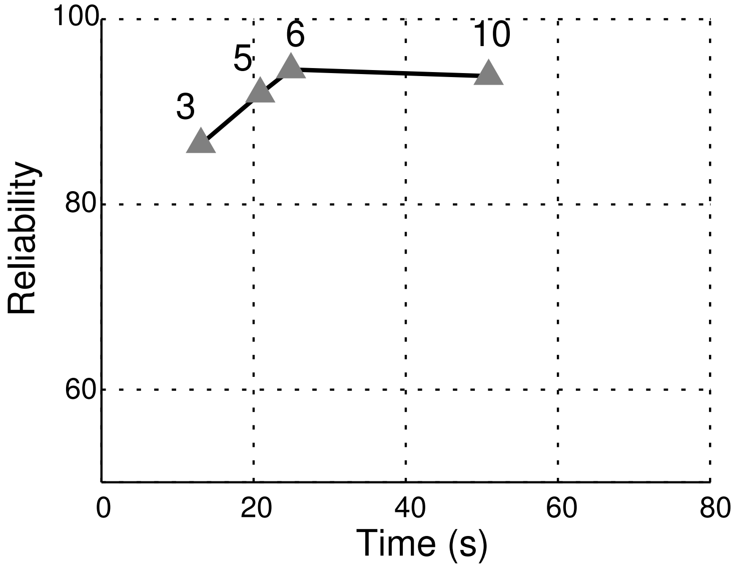

Figure 6: **The effect of the number of vertical and horizontal image reslice pairs in the Nelder-Mead Simplex algorithm**. The vertical image pairs were generated over the whole volume up to a 10 pixel margin at either side and a 16 pixel margin at the top. The margins for horizontal image pairs were 70 pixels at the top and 10 pixels at the bottom. The dimensions of the image slices were therefore 183 × 250 pixels (B-scan dimensions are 266 × 352 pixels). The image scale is 0.01 cm/pixel. The source volume location was initialised randomly in the range ±5mm from the centre in all three coordinates. The Nelder-Mead Simplex algorithm was used, with the total number of iterations set to 35. The algorithm was allowed 2 restarts by using the current optimum as one vertex of the simplex and reinitialising the remaining vertices within ±5mm from the centre in all coordinates. For the generation of a single data point, 160 trials were performed on 16 datasets (thus totaling 640 trials to generate 4 data points). The labels on the data points show the number of horizontal and vertical image pairs in each case (both numbers being equal).

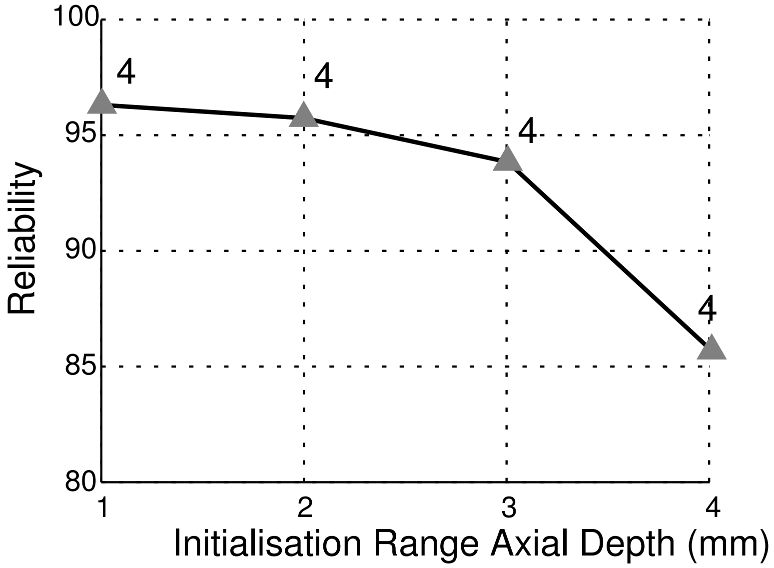

Figure 7: **The effect of axial initialisation range in the Nelder-Mead Simplex algorithm**. The vertical image pairs were generated over the whole volume up to a 10 pixel margin at either side and a 14 pixel margin at the top. The margins for horizontal image pairs were 70 pixels at the top and 70 pixels at the bottom. The dimensions of the image slices were 130×250 pixels. For the generation of a single data point, 160 trials were performed on 16 datasets. The label on the data points represent the number of horizontal and vertical image pairs. The source volume location was initialised randomly in the fixed range ±5mm from the centre in the x and z coordinates. The initialisation range was varied in the y (axial) coordinate. The B-scan dimensions, the image scale, and total iterations are exactly the same as mentioned in the caption to Fig. 6. However, the algorithm was allowed only 1 restart with the same initialisation ranges. On average, the algorithm took 8–9 seconds to converge.

4 Open issues and further work

This section discusses the current problems with the registration strategies and what needs to be done to achieve the goal of a clinical system. A significant bottleneck in the voxel-array approach is the 3D smoothing[1]. The 3D voxel-arrays are significantly larger than the 2D reslices required in the original multiresolution algorithm [4] and therefore smoothing takes a large portion of the total time (with total time of the algorithm somewhere between 30 seconds and 1.5 minutes). One possible solution is to use a different, less-intensive smoothing filter such as a mean filter that is based on the concept of an accumulation buffer, although the smoothing time would still be significant.

Another problem with the voxel-array based approach is that as the relative orientation of the two volumes changes, the size of the voxel-array increases because it has to cover both volumes. Currently, we are using a sweep angle of 10 degrees and if this increases, the voxel-arrays will be

---

[1] For example, with a faster implementation of a Gaussian blur, convolving a 31 point kernel with a 256x256x256 voxel-array takes 4.3 seconds, a 31 point kernel with a 512x256x256 voxel-array takes 8 seconds, and a 5 point kernel with a 512x256x256 voxel-array takes 2 seconds. As an alternative, we have also used a *Fast Gaussian Blur* `http://www.geometrictools.com/Documentation/FastGaussianBlur.pdf` which is based on solving an initial value problem for a particular partial differential equation (heat equation) and is based on a finite difference scheme. For a larger value of standard deviation σ, the finite difference scheme has to be iterated. This implementation of the Gaussian filter might be fast for applications that require a higher level of smoothing (higher σ-scales) but is quite slow compared to the conventional Gaussian filter with smaller σ. After using this smoothing algorithm in the multiresolution search algorithm and testing it for multiple sigma values, we rejected it as it is still slower than the Gaussian filter currently in use.

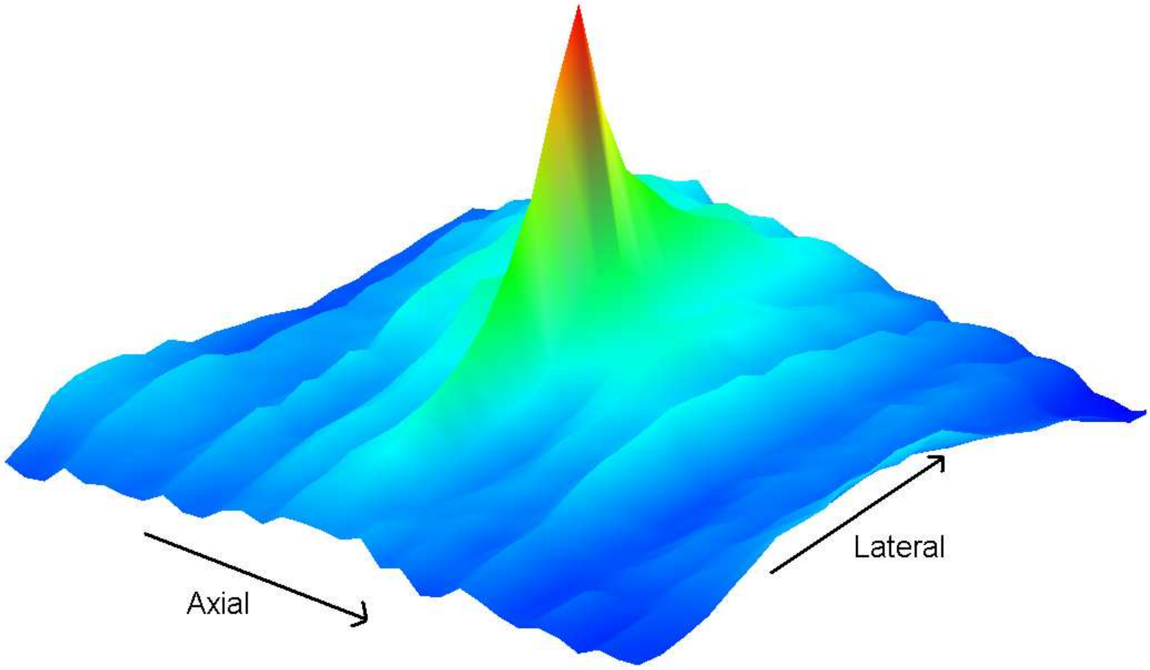

Figure 8: **Behaviour of normalised mutual information in R2 dataset under rigid translation of the volumes**. Our experience is that most *in vivo* data contains mainly horizontal features from layers of tissue. The similarity measure therefore has a sharper peak in the axial direction compared with the lateral and elevation directions. Also, the similarity measure has smaller peaks on either side of the correct peak in the axial direction. This causes the failures of the Nelder-Mead simplex algorithm which can get trapped in suboptimal local peaks.

even larger and therefore registration will be slower. Increase in size won't affect the Nelder-Mead simplex algorithm as the size of the reslices are independent of the underlying data and can always be fixed. Currently, the multiresolution algorithm works on 4 levels. There is a speed disadvantage between the 1st and 2nd level where the downsampling changes by 4, giving an increased search range at level 2. Inserting another level will require the generation of another pair of smoothed and downsampled voxel-arrays and will therefore not necessarily be faster. This is something that needs to be tried once the smoothing is optimised.

A possibility for making the system easier to use is to allow some limited manual intervention at appropriate points in the process. For example, when registering a sequence of volumes, one could allow the clinician to specify which pairs have failed. In the current system (see Fig. 12), when there are more than two volumes, volumes are compared in subsequent pairs, and when one pair fails, the whole dataset needs to be reregistered. Instead, we could ask the clinician to do a quick manual registration of failed volumes. This may be more convenient for the clinician than repeatedly telling the algorithm to try again using the automatic methods. In the clinical system, we do not necessarily have to choose one of the algorithms to use. Instead, we can use the faster algorithm most of the time but have the other available as a backup if the first fails. Alternatively, the backup could be to run the same algorithm again but with different settings. If we follow the manual intervention approach, we will need to be very careful to make the interface as simple as possible, and we will need to think further about how to do this. Feedback from the clinician would be very helpful in this. We can develop a clinical system that at this point requires a lot of human interaction (for instance, choosing the algorithm, reslice planes, similarity measures, etc.) and as we get experience of using it in clinical settings, we can finalise some of the settings. There will be valuable information in the clinical feedback which will ensure that we develop a system

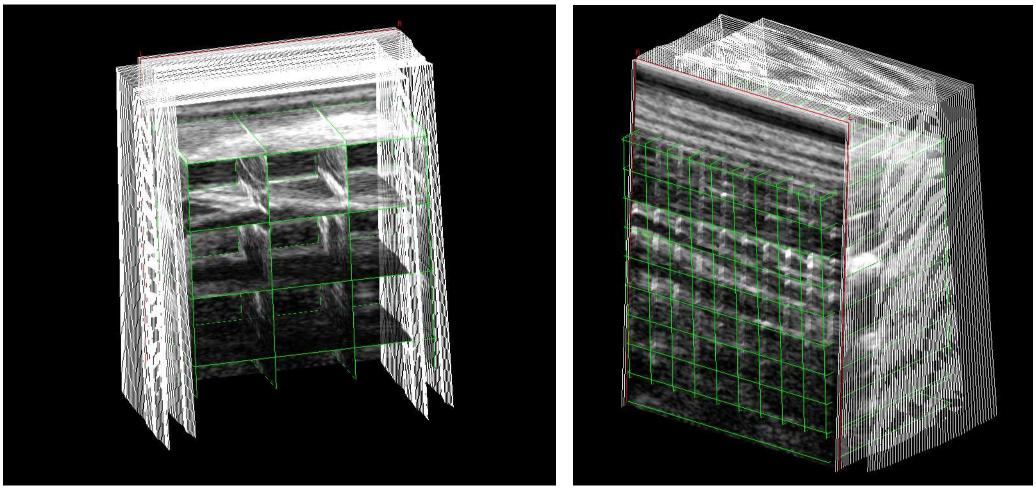

(a) 4 horizontal and vertical reslice pairs (b) 10 horizontal and vertical reslicepairs

Figure 9: Examples of reslice plane configurations used for experimental analysis in Figs. 6 and 7.

which is clinically useful.

References

[1] U. Z. Ijaz, R. W. Prager, A. H. Gee, and G. M. Treece. Optimisation strategies for ultrasound volume registration. Technical Report CUED/F-INFENG/TR 635, Cambridge University Department of Engineering, August 2009.

[2] C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32:71–86, 1999.

[3] R. W. Prager, A. H. Gee, and L. Berman. Stradx: real-time acquisition and vizualization of freehand three-dimensional ultrasound. *Medical Image Analysis*, 3(2):129–140, 1999.

[4] A. H. Gee, G. M. Treece, R. W. Prager, C. J. C. Cash, and L. Berman. Rapid registration for wide field-of-view freehand three-dimensional ultrasound. *IEEE Transactions on Medical Imaging*, 22(11):1344–1357, 2003.

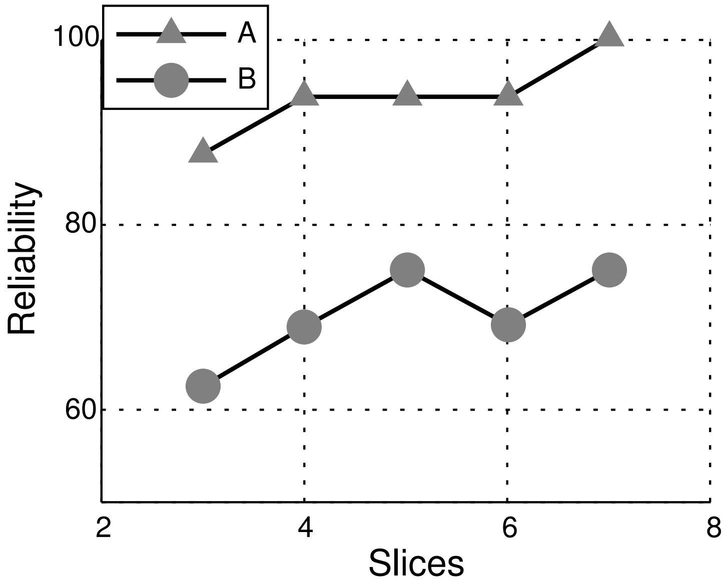

Figure 10: **The effect of the number of image reslice pairs in the voxel-array based multiresolution search algorithm**. Here, the same number of slices are considered in all three directions. At L4, L3, and L2, the voxel arrays are smoothed by convolution with 3D Gaussian kernels of 48, 12, and 6 points and downsampled by a factor of 16, 4, and 2, respectively. L1 employs no downsampling (i.e. *resxy* is 1 pixel, and *resz* is 0.01 cm) but is smoothed with a 3D Gaussian kernel of 3 points. L4 performs an exhaustive search in the range $(-(3/4)\text{n}, (3/4)\text{n})$(where n is the dimension of the voxel array in each coordinate) and subsequent levels of the algorithm are limited to a narrower range (±1 step in all dimensions centred on the optimum location obtained in the previous level). There are two configurations shown for the generation of reslices. In configuration A, the reslices are created within voxel margins of 4 voxels at either side concentrating in the central region, and in configuration B, voxel margins at either side of 1 voxel. Configuration A improves the reliability as the majority of the *in vitro* datasets have features that are concentrated in the centre. For the generation of a single data point, 16 trials were performed on 16 datasets.

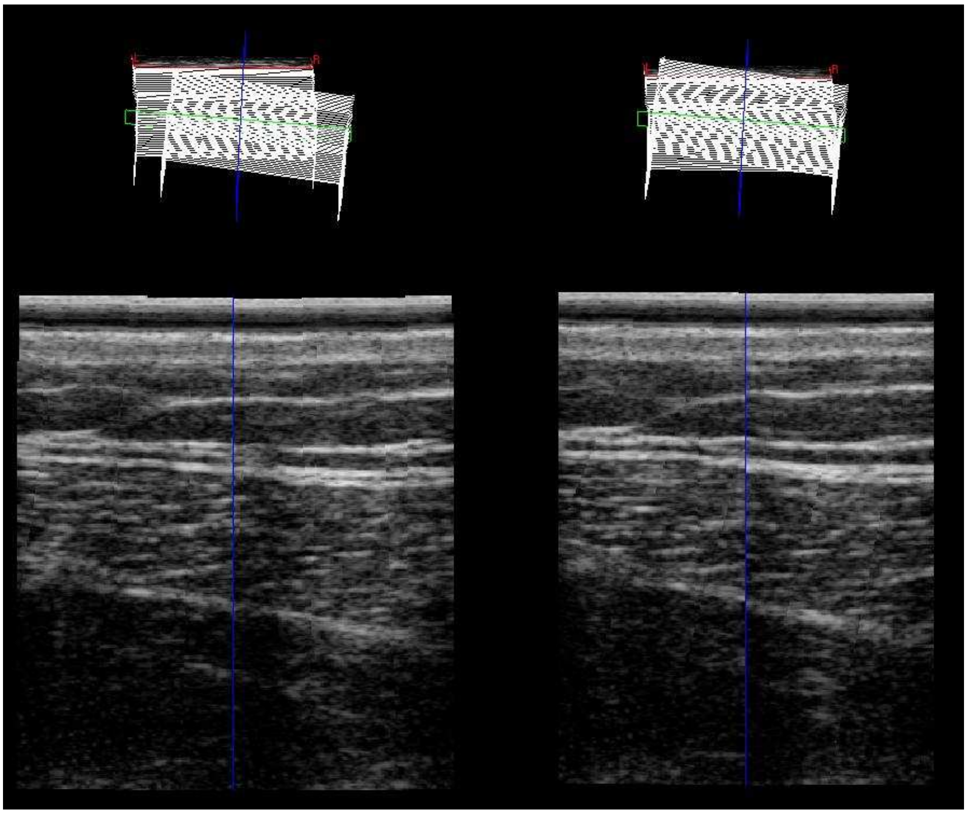

Figure 11: **Two frequent alignments for M4 dataset**. The left column shows the alignment at the global optimum whereas the right column shows the alignment at a local minimum. The Nelder-Mead simplex algorithm converges to both configurations whereas the voxel-array based multiresolution search algorithm converges to the local minimum only. On visual inspection both alignments seem to be correct.

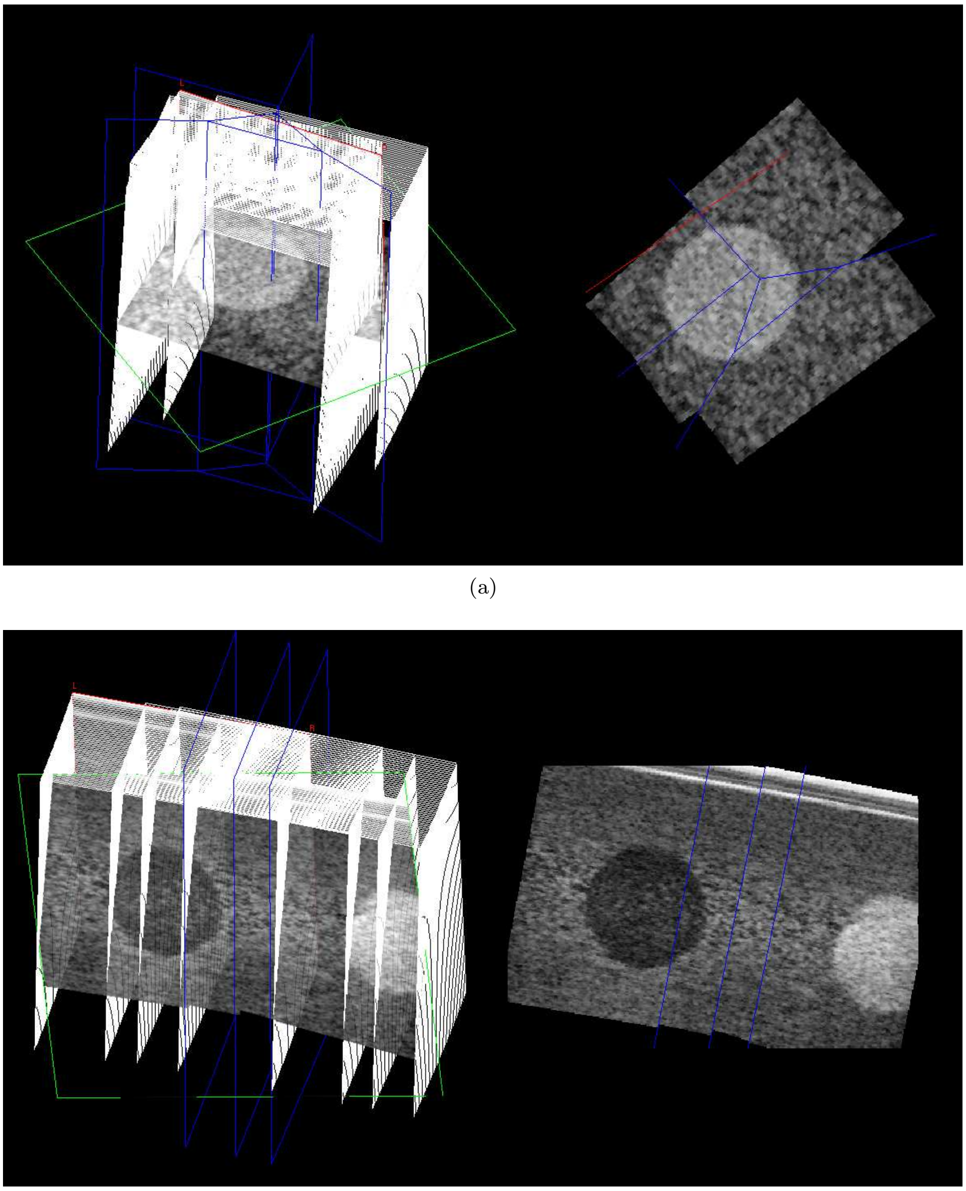

(a)

(b)

Figure 12: **Registration results using the Nelder-Mead simplex algorithm on datasets containing more than two volumes**. The slice across the region is represented by a green outline and is shown on the right side with blue lines representing the dividing planes. The data on either side of each dividing plane comes from different volumes.