
**OPTIMISATION STRATEGIES
FOR ULTRASOUND
VOLUME REGISTRATION**

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

CUED/F-INFENG/TR 635

August 2009

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

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

Optimisation strategies for 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 paper considers registration of 3D ultrasound volumes. One way to acquire 3D data is to use a mechanically-swept 3D probe. However, the usefulness of these probes is restricted by their limited field of view. While this problem can be overcome by attaching a position sensor to the probe, an external position sensor can be an inconvenience in a clinical setting and does not align the volumes correctly when there is tissue displacement or deformation. The objective of this paper is to replace the 6 degree of freedom (DOF) sensor with a combination of 3 DOF image registration and an integrated inertial sensor for measuring orientation. We examine a range of optimisation algorithms and similarity measures for registration and compare them in *in vitro* and *in vivo* experiments. We register based on multiple reslice images rather than a whole voxel array. In this paper, we use a large number of reslices for improved reliability at the expense of computational speed. We have found that the Levenberg-Marquardt method is very fast but is not guaranteed to give the correct solution all the time. We conclude that normalised mutual information used in the Nelder-Mead simplex algorithm is potentially suitable for the registration task with an average execution time of around five minutes, in the majority of cases, with two restarts in a C++ implementation on a 3.0 GHz Intel Core 2 Duo CPU machine.

1 Introduction

In recent years, it has become increasingly common to produce intrinsically 3D ultrasound (US) probes. Volumetric imaging has now matured because of the possibility of capacitive micro-machined US transducers and 2D phased arrays. It is anticipated that some clinical applications offered by 3D US will replace X-ray computed tomography (CT) with greater safety and will replace magnetic resonance imaging (MRI) at lower cost [1]. By acquiring 3D volumes from different viewpoints and subsequently aligning them to create mosaics, volumes can be created with an extended field-of-view, as well as potentially better data quality, both of which can give more useful information to clinicians. Essentially, there are two ways to acquire 3D US data. The first option is to use a 3D probe that when placed on the skin scans a small fixed volume underneath. The alternative is to use a freehand system [2, 3] which sweeps a conventional 2D probe over the region-of-interest resulting in a 3D volume. The probe has an additional position sensor (Fig. 1) to record the position and orientation of each image frame. Even though the position sensor, if calibrated well, is able to locate the B-scans within a sweep, there are additional problems associated with determining the relative positions of B-scans in an overlapping sweep. The main reason is the displacement and deformation of the underlying anatomy when scanned from different directions in multiple views. Furthermore, the position sensor is mounted on the probe, which makes it somewhat inconvenient to use. For example, line-of-sight (Fig. 1) may need to be maintained between the sensor and a reference point. It is also difficult to use a position sensor in certain 3D US screening methods, for example, with endoscopic, intravascular, trans-rectal, and trans-vaginal 3D US probes. A possible solution is to consider an image-based automatic registration technique that uses a matching algorithm to register overlapping data. Such a technique could also take into account the presence of speckle; directionality in the tissue backscatter properties; and variation in signal strength at different depths in the B-scans.

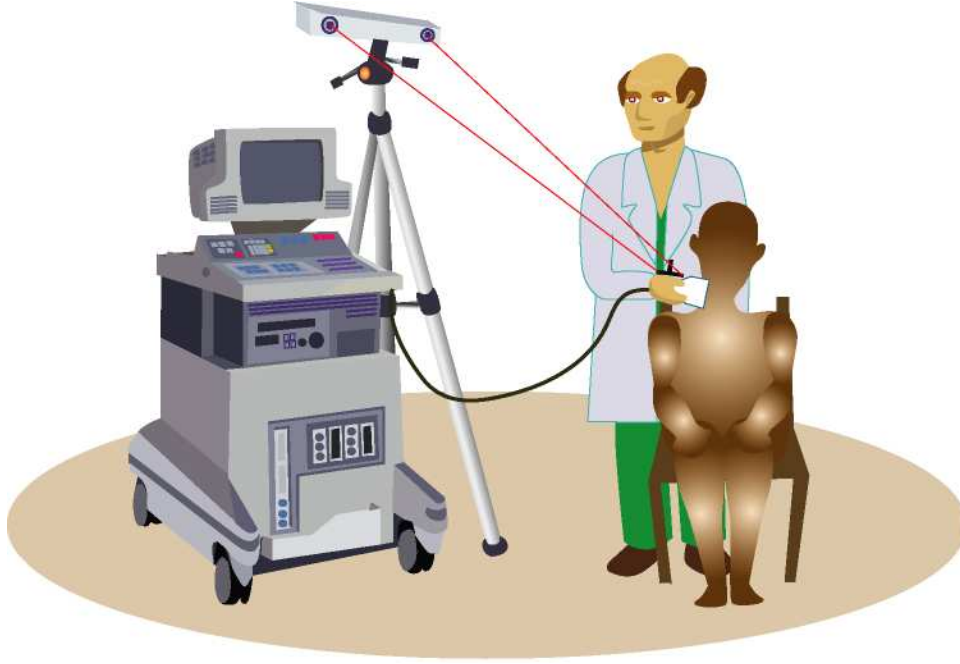


Figure 1: **A typical US system with an optical position sensor mounted on the US probe.** In this configuration, a line-of-sight has to be maintained between the sensor and the reference point. Other types of sensors for 3D US [4, 5, 6, 7, 8] may be magnetic, mechanical or acoustic.

There is a lot of literature on medical image registration, but very little of it concerns 3D US volumes although there are examples of US registered to other modalities. The main difficulties in US registration are that there is an irregular sampling of the acquisition space and the anatomy's appearance varies with the direction of insonification. Furthermore, there is a low signal to noise ratio. The main culprit is the coherent wave interference known as speckle which limits low-resolution image contrast and may even obscure true structures in high-contrast regions. While it is difficult to formulate a coherent classification of existing registration techniques, they can nevertheless be described in terms of common themes. For example, most of the registration algorithms deal with matching techniques that employ some kind of similarity measure. The similarity measures are used to calculate the six rigid-body transformation parameters (three for rotation and three for translation) and optionally include scale factors (including coarse-to-fine scale [10]) and deformation maps (including thin-plate splines [11] and B-splines [12]). There are many similarity measures available including mutual information [3, 13, 14, 15, 16], cross-correlation [17, 18], Bhattacharyya's coefficient [19], mean square of intensities [20], correlation ratio [21, 22, 23], and sum of squared differences [18]. There also exists some studies that consider the above mentioned similarity measures for multimodality fusion of US images with MRI [4, 17, 21] and CT [22, 24, 25]. Other noteworthy approaches include similarity energy for free-form deformation [26], variational approaches [27], segmentation-based approaches [28], probe pressure correction [29], and joint volume reconstruction with image alignment [30].

Some of the US-specific similarity measures are based on a maximum likelihood framework [21, 31, 32, 33, 34] and model ultrasonic speckle as multiplicative Rayleigh distributed noise (in the case of envelope-detected US B-mode data), or as signal-dependent Gaussian noise (in the case of displayed US image data that has undergone logarithmic compression, low and high-pass filtering, postprocessing, etc.). A recent study by Wachinger *et. al.* (2008)[18] considered these US-specific similarity measures for US image sequence registration. By weighting the similarity

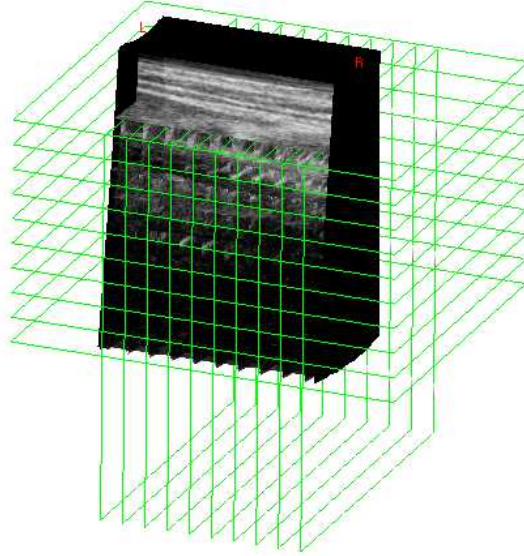


Figure 2: **Horizontal and vertical image boundaries for corresponding image pairs in two overlapping volumes obtained using a mechanically-swept 3D US probe.** The probe is positioned at the top and the direction of insonification is downwards. The image pairs are generated uniformly over the source and target volumes and located with reference to the first B-scan in the target volume. The vertical image pairs are generated over the whole volume up to a 10 pixel margin at either side and a 70 pixel margin at the top. The margins for horizontal image pairs are 70 pixels at the top and 10 pixels at the bottom. The dimensions of the image slices are 470×470 pixels, whereas the dimensions of the B-scans are 266×352 pixels. The scale is 0.01 cm/pixel. The image slices are calculated directly from the B-scans using an algorithm described in [9]. There is no intermediate voxel representation.

measures by a factor dependent on the overlap between volumes, a maximum likelihood framework based on the logarithm of Rayleigh noise was formulated and was shown to perform better than the others considered. However, these similarity measures were found to have problems with partial overlap, clearly favouring a total overlap of the volumes.

Since the images being registered often have outlier image samples due to the presence of unexpected problems such as noise, deformations and variation in the appearance of the anatomy, it is necessary to find a measure that is robust and provides statistical efficiency. After analysing several similarity measures, we have restricted this study to probabilistic similarity measures [35] which treat the intensity pairs taken from corresponding spatial locations in two images as independent and identically distributed samples of two random variables. Statistical concepts such as correlation, joint entropy and mutual information are used as similarity measures by estimating the statistical properties of the samples. However, these similarity measures come with a computational burden. In comparison to automatic-landmark-based image registration techniques, where the data must be interpolated only once to generate the landmarks and only their spatial locations are matched for corresponding volumes, the similarity-measure-based registration techniques require interpolation at each iteration of the search algorithm. Despite this advantage, the automatic-landmark-based techniques, which have been quite successful in the computer vision community, are not well suited to the registration of 3D US volumes [19]. This is due to the presence of speckle which generates many false-positives in the automatic corner detection.

In this paper, we make the following assumptions. First, we assume that any distortion of the tissue is small and therefore that only a rigid body transformation is required for successful registration. Second, there is reasonable overlap between the volumes. Third, the probe is held

stationary while recording each volume so there is no need for intra-volume registration. Finally, we assume that the relative orientation of the volumes is already known. Even though it is possible to estimate six rigid-body transformation parameters (i.e., three for translation and three for rotation) based on our recent studies [35, 36], we reached the conclusion that the similarity measures are quite sensitive to changes in rotation parameters. The convergence of the search algorithm depends highly on the initial orientation information. This must be within 1-2 degrees of the correct solution. We are aware that such a restriction may not be an acceptable limit in clinical practise and to meet more stringent requirements for physiological experimentation, we propose a hybrid approach.

2 Proposed Approach

Inertial sensors (accelerometers and rate sensors) that do not need line-of-sight to operate correctly, can be integrated relatively easily with the ultrasound probe. Thus computing 3 DOF orientation is conceptually straightforward i.e., angular velocities measured by rate sensors are single-integrated to provide orientation information. However, these inertial sensor measurements are often not perfect; gyros and accelerometers used in the inertial sensor often have measurement bias. This bias shows itself as angular drift which increases linearly over time. The accelerometer signals are even worse. The signals with their drifts are transformed and so the drifts are present in the reference co-ordinate system and hence the errors in orientation and position increase proportionally with elapsed time. In order to compensate for the sensor drift, a process called augmentation is used in inertial sensors, whereby gyros and acceleration errors are compensated by utilising other measures. For system internal augmentation, cheap inclinometers are often used, and for external augmentation, odometers, speedometer, or a magnetic compass can be used to improve the system's performance.

By using the sensor to measure orientations, we are left with a simpler problem with a lower-dimensional search space, which improves both the speed and the robustness of the registration algorithm. The complete solution is determined from both the inertial sensor and the 3 DOF registration strategies discussed in this paper. There are several commercially available inertial sensors which have been used for a wide variety of military and commercial applications [37]. They typically have a documented accuracy of less than 0.05 degrees.

In order to make the algorithm acceptably fast, we consider only corresponding image slices (Fig. 2) for registration as opposed to complete voxel arrays. We use vertical and horizontal slices through the region of interest based on nearest-neighbour interpolation of the B-scans for which we are using a very efficient reslicing algorithm [9]. Using this approach, two sets of images are obtained for two sweeps across the region-of-interest. On these, we apply some of the most popular similarity measures. It should be noted that our approach is similar to the one considered in [3], where the registration is done based on a single reslice plane. In [3], the overlap was typically very small, with little room for more than one plane on which to compare the volumes. A fundamental limitation of registration in such circumstances is that a good match on the single reslice plane does not guarantee good global alignment. In this work, we consider more extensive overlaps into which we can fit more slices. This should increase robustness and eliminate the need for any manual intervention.

The probabilistic similarity measures are particularly useful for multi-modality image registration since they do not assume any functional relationship between the two image values and have an inherent degree of robustness. These measures may not be statistically efficient, i.e. the registration variability due to noise can be larger than the sample correlation coefficient, but since US images have low signal-to-noise ratio, they are well suited to the registration task. The two images are thus considered as random variables, taking greyscale values between 1 and 256. Probabilities are then denoted with: $p_i = P_f(f(T(\mathbf{x})) = i)$, $p_j = P_F(F(\mathbf{x}) = j)$, and $p_{ij} = P_{f,F}(f(T(\mathbf{x})) = i, F(\mathbf{x}) = j)$. Here T represents a transformation over spatial coordinates \mathbf{x} .

The most widely known measure of the diversity of a distribution X is the Shannon entropy

defined by

$$H(X) = - \sum_i p_i \log p_i \quad (1)$$

Shannon [38] defines the mutual information of two random variables by the reduction in diversity of the first variable brought by the knowledge of the second variable:

$$I(X) = H(X) + H(Y) - H(X, Y) \quad (2)$$

where

$$H(X, Y) = - \sum_i \sum_j p_{ij} \log p_{ij} \quad (3)$$

The mutual information of two images expresses how much the uncertainty on one of the images decreases when the other one is known. It is assumed to be maximum when the images are registered. In this study, we have used the normalised mutual information (NMI) defined as

$$\text{NMI} = \frac{H(X) + H(Y)}{H(X, Y)} \quad (4)$$

NMI is particularly helpful in those cases where there is a small overlap between the images. As the overlap decreases, the marginal entropies $H(X)$ and $H(Y)$ may not necessarily decrease relative to the joint entropy $H(X, Y)$. Therefore, we divide the marginal entropies by the joint entropy and any change in the marginal entropies will be counteracted by a change in the joint entropy.

We also considered Kolmogorov's distance as an alternative similarity measure. It belongs to a family of f-divergences which can be used as a measure between an observed joint distribution (p_{ij}) and a computed joint distribution ($p_i p_j$) in case of independence. The Kolmogorov's distance (Basseville [39], Sarrut and Miguet [40]) is defined as

$$K = \frac{1}{2} \sum_i \sum_j |p_{ij} - p_i p_j| \quad (5)$$

In order to find the probabilities, a joint histogram of intensities is calculated. Each entry in the histogram denotes the number of times intensity i in one image coincides with intensity j in another image. Dividing each entry by the total number of entries gives the joint probability distribution p_{ij} . The probability distributions (p_i and p_j) for each image are found by summing over the rows and columns respectively of the joint histogram.

The 3 DOF position information can then be estimated by using an appropriate search algorithm. One can choose between methods that need only evaluations of the function to be minimised and methods that require evaluations of the derivative of the function. We have restricted our study to the following search algorithms: Nelder-Mead simplex algorithm [41], Powell's direction-set method [42], particle swarm optimiser [43], and Levenberg-Marquardt method [44]. Details of each are given in the appendix. We believe that the comparison of the above mentioned algorithms will offer new perspectives for a real-time 3D ultrasound volume registration. From these four algorithms, we want to select the algorithm which is both reliable and cost-effective in clinical settings.

3 Experiments and Results

US data was recorded with 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 (<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/>).

3.1 *In Vitro* Assessment

The search algorithms were tested on a series of *in vitro* scans of a speckle phantom with several 5mm spherical inclusions. The ultrasound probe was mounted on a motorised Zaber’s linear slide TLSR-300B gantry (<http://www.zaber.com>) to obtain precise translation offsets (as shown in Table 1) between two recorded volumes with known relative orientation. Each acquired volume consisted of 50 frames of data swept over 10 degrees. Since the probe was aligned to the gantry by eye, it was necessary to determine the transformation between gantry and probe coordinates, in order to correctly compare the 3 DOF offset of the registered volumes to the gantry translations. This transformation was determined by a calibration procedure similar to the single point calibration method described in [45]. Here, we are treating the centre of sphere as a point target. Since all the volumes were recorded with the probe at the same orientation, the translational components of the calibration are not required. To determine the relative orientation of the two coordinate systems, we solved the 6 DOF calibration with the point target at an arbitrary location in the gantry coordinate system and discarded the resulting translation parameters of the calibration. The calibration parameters were calculated using the same six-volume gantry data set. The point target was the centre of the uppermost bright sphere visible in the six recorded volumes. In each volume, the sphere centre was found by taking many reslices through the sphere in three orthogonal directions and automatically segmenting the circular section of the sphere visible in each image. The least-squares intersection point of lines passing through the centre of each segmented circle in directions perpendicular to the plane of the reslice provides an estimate of the sphere centre location in probe coordinates. This enabled us to establish a ground truth with which to compare the performance of the search algorithms.

Dataset	x(cm)	y(cm)	z(cm)
P1	0.2	0	0.2
P2	0.3	0	0
P3	0	0	0.3
P4	0.4	0	0.4
P5	0.5	0	0.5

Table 1: Offsets between the two volumes for *in vitro* datasets measured by the motorised Zaber’s linear slide. x is the lateral direction, y is axial and z is elevational. The y-offset is zero in each case so that there is no probe pressure distortion.

We have considered 10 horizontal and 10 vertical image pairs and 32 histogram bins in the similarity measures. These values were determined from *in vivo* data in a preliminary study [35] in order to achieve the necessary robustness in the algorithms. Even though the similarity measure works well with at least 5 horizontal and 5 vertical image pairs, we use 10 of each in these experiments for additional robustness. Similarly, based on [35], we choose 32 histogram bins because the peaks are more prominent at the ground truth. In order to reduce the effect of speckle on image registration performance, we use a 5×5 Gaussian smoothing filter with a standard deviation of 1 pixel. However, using this also increases the computational burden.

In all the search algorithms, the centres of the volumes were first aligned and then the positions of the source volume were initialised to random values in the range ± 0.5 cm from the centre in all coordinates. This accounts for more than 50% overlap of the two volumes and in practical scenarios gives enough features to register properly. In the Nelder-Mead simplex algorithm, we used the following parameter settings: $\alpha = 1$, $\beta = 0.5$, $\gamma = 2$, and $\delta = 0.5$ (refer to the appendix). The total number of iterations was set to 70. The algorithm can also terminate when the decrease in the similarity measure as a fraction of the similarity value is less than the tolerance of 10^{-5} . Since the Nelder-Mead simplex algorithm can get stuck in local minima, it is often a good idea to restart the procedure at a point where it claims to have found a minimum. In the current implementation of the Nelder-Mead simplex algorithm, we have allowed two restarts by considering the local minimum as one vertex of the simplex and reinitialising the other N vertices within ± 0.25 cm of

Datasets	Normalised mutual information				Kolmogorov's distance			
	Nelder-Mead simplex method	Levenberg Marquardt method	Particle swarm optimiser	Powell's direction-set algorithm	Nelder-Mead simplex method	Levenberg Marquardt method	Particle swarm optimiser	Powell's direction-set algorithm
P1	0.281(0)	0.220(2)	0.283	0.280(3)	0.313(0)	0.354(1)	0.315	0.315(2)
	0.281(0)	0.360(2)	0.281	0.283(1)	0.315(0)	0.362(2)	0.310	0.315(2)
	0.280(1)	0.228(3)	0.300	0.280(1)	0.308(0)	0.330(1)	0.314	0.315(1)
	0.281(0)	NA	0.277	0.281(1)	0.315(0)	NA	0.308	0.311(1)
	0.281(0)	0.326(2)	0.281	0.283(2)	0.315(0)	0.411(2)	0.320	0.315(3)
P2	0.385(0)	0.412(3)	0.384	0.385(1)	0.385(0)	0.370(2)	0.387	NA
	0.386(0)	0.401(2)	0.387	0.386(1)	0.385(0)	0.522(1)	0.387	0.385(2)
	0.385(0)	0.439(3)	0.390	0.386(3)	0.385(0)	NA	0.385	0.385(1)
	0.385(1)	0.415(1)	0.388	0.386(1)	0.386(0)	0.366(2)	0.386	0.385(1)
	0.386(0)	0.414(2)	0.387	NA	0.385(0)	NA	0.388	0.385(1)
P3	0.123(1)	0.153(2)	0.124	0.104(1)	0.130(0)	0.164(2)	0.130	0.132(1)
	0.114(0)	0.201(2)	0.116	0.124(2)	0.129(0)	0.199(3)	0.128	0.131(1)
	0.108(0)	0.153(2)	0.117	0.123(2)	0.129(1)	0.157(1)	0.130	0.133(2)
	0.114(0)	0.134(2)	0.114	0.123(1)	0.131(2)	NA	0.129	0.128(1)
	0.114(0)	NA	0.122	0.123(3)	0.129(0)	0.171(3)	0.126	0.129(1)
P4	0.327(0)	0.383(3)	0.346	0.348(1)	0.318(1)	0.402(2)	0.320	NA
	0.336(0)	NA	0.340	0.349(2)	0.318(0)	0.395(2)	0.317	NA
	0.329(0)	NA	0.345	0.349(1)	0.333(2)	NA	0.321	NA
	0.342(0)	NA	0.336	0.345(1)	0.341(0)	NA	0.322	NA
	0.327(0)	0.363(3)	0.345	0.348(3)	0.318(2)	NA	0.317	0.318(3)
P5	0.387(0)	NA	0.378	NA	NA	NA	NA	NA
	0.384(1)	0.454(1)	0.356	0.360(1)	NA	NA	NA	NA
	0.387(1)	0.421(1)	0.357	0.364(3)	NA	NA	0.392	NA
	0.365(0)	0.353(1)	0.361	0.364(1)	NA	NA	0.401	NA
	0.375(0)	NA	0.367	0.365(1)	0.391(1)	0.436(3)	NA	NA

Table 2: **RMSE(mm) for the *in vitro* experiments.** A total of 200 trials were performed on 5 datasets. NA represent the cases when the algorithms didn't converge at all (i.e. RMSE > 0.6 mm). In the case of the Nelder-Mead simplex algorithm, the terms in parentheses represent the number of restarts required to achieve the convergence. In the case of the Powell's direction-set method and Levenberg-Marquardt method, the term in the parenthesis is the most successful of the three attempts. On average, it can be seen that the misalignment was less than 0.4 mm. It can also be seen here that when there is less overlap, i.e., P4, and P5, normalised mutual information performs much better than Kolmogorov's distance.

the claimed minimum. For a 3 DOF estimation problem, this simplex is a tetrahedron and requires 4 initial position estimates. In Powell’s direction-set method, the bracket for the minimum for both Brent’s method and the gold section search were set to $[-0.5 \text{ cm}, 0.5 \text{ cm}]$. The direction matrix \mathbf{U} was taken to be an identity matrix, \mathbf{I}_3 , and the tolerance was set to 10^{-5} . In the Levenberg-Marquardt method, we consider $\beta = 2$ for normalised mutual information, and $\beta = 0.50$ for Kolmogorov’s distance, respectively. In order to calculate the Jacobian \mathbf{J} , we use the finite difference approximation by perturbing the current estimates by 1 pixel. In both Powell’s direction-set method, and the Levenberg-Marquardt method, we have considered two restarts (as considered in the Nelder-Mead simplex algorithm), however, we choose the best run among the three. For the PSO algorithm, the total number of iterations was set to 50 and we used 30 particles. The PSO algorithm was initialised with random offsets of the source volume within the vicinity of the target volume.

As described in Section 2, we have considered two similarity measures: normalised mutual information and Kolmogorov’s distance. For each similarity measure, 5 simulations were performed per dataset (see Table 2) giving a total of 200 simulations over four optimisation algorithms and two similarity measures. The accuracy and precision of the alignment algorithms were assessed using the root mean square error (RMSE) between the ground truth and positions obtained from the search algorithm. RMSE is defined as

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum d_i^2} \quad (6)$$

where d_i is the Euclidean distance between the actual position and estimated position of the B-scan corner points as shown in Fig. 3. We designate the volumes to be correctly registered if the RMSE is less than 1 mm. In addition to using RMSE, we considered a visual inspection of the registered data as a more qualitative indication of the algorithm’s accuracy.

The results of this sets of experiments are given in Table 2 and Fig. 4. Using normalised mutual information, the Nelder-Mead simplex algorithm converged in all cases out of which 80% required no restarts and 20% required one restart. Similarly PSO converged in all cases. However, the success rate of Powell’s direction-set algorithm, and the Levenberg-Marquardt method were only 92% and 72%, respectively. For the case of Kolmogorov’s distance, the success rate of Nelder-Mead’s simplex algorithm was only 84% out of which 72% required no restarts, 14% required one restart, and 14% required two restarts. PSO was convergent in 88% of cases and Powell’s direction set algorithm was convergent in 60% of the cases. In terms of reliability, the worst was the Levenberg-Marquardt method, which had a success rate of only 56%.

It can be seen from Table 2 that the RMSE is less than 1 mm. There are several possible reasons for it being greater than zero. First, we have used nearest neighbour interpolation to generate the candidate reslices. This introduces small errors into the apparent position of features. Second, the *in vitro* data has only a few features which may be partly lost in the sparse reslicing. Finally, we have applied a smoothing filter which blurs the features to some extent. Despite these issues, on visual inspection of the data (e.g. Fig. 5), it can be seen that the results are quite encouraging with an obviously good alignment between the volumes.

3.2 *In Vivo* Examples

While *in vitro* experiments give a precise assessment of registration accuracy, such experiments are of limited use considering that B-scans of phantoms are very different to those of the human body. Therefore, we performed some *in vivo* experiments using the same settings as before. For these experiments, each B-scan’s position and orientation was recorded using a Northern Digital Polaris optical tracking system (<http://www.ndigital.com>), with spatial and temporal calibration performed according to the techniques in Treece *et al.* [29]. A total of 8 *in vivo* datasets were acquired consisting of scans of the neck and calf muscles. We retained the orientation information from the position sensor and aligned the centres of the volumes before application of the search algorithms.

We used normalised mutual information for the *in vivo* datasets and allowed one run of each algorithm per dataset as it is obvious from the RMSE results given in Table 2 that algorithms are almost always convergent in the first run. The results are shown in Figs. 6-10 and Table 3. One would not expect the *in vivo* datasets to be registered quite as well as the *in vitro* datasets. This is because of probe pressure and mispositioning of features due to variations in the speed of sound in the tissue. Nevertheless, we found the RMSE to be less than 1 mm in those cases when we were quite careful during data acquisition i.e., no sudden movements, pressure or respiration artifacts (Class R in Table 3). Three of the datasets were acquired in a less careful manner and allowed some movement of the scanning subject (Class M in Table 3). Even though the RMSE results show that the optimal solution is a long way from the actual solution, one can see that the registration worked well on close visual inspection of the results (Figs. 8-10).

The overall registration time for the algorithms depends on several factors including the number of iterations, the number of image pairs, and whether we apply the smoothing filter. As a result, it is not possible to specify an exact time for each algorithm to run. For the configuration chosen in this paper, we are reporting an average execution time of approximately 5 mins for the Nelder-Mead simplex algorithm with two restarts, 2 mins for the Levenberg-Marquardt method with two restarts, 15 mins for Powell’s direction-set method with two restarts, and 17 mins for the PSO algorithm (as shown in Fig. 4). These execution times are for the algorithms implemented in Stradwin, running single-threaded on a 3.0GHz Intel Core 2 Duo CPU without using any advanced code optimisation technique.

In summary we have found that the Nelder-Mead simplex algorithm with normalised mutual information is the most appropriate choice having a high success rate with restarts. The PSO algorithm is also very successful but takes longer to run. The Levenberg-Marquardt method is the fastest of the considered algorithms, however, it is not reliable. One has to use more than two restarts in the algorithms to make it effective. With Powell’s direction-set method and the Levenberg-Marquardt method, it is not guaranteed that a restart will produce a better solution. This is because where the Nelder-Mead keeps track of the previous best solution as one of its multiple solutions, the Powell’s direction-set method, and Levenberg-Marquardt method only use the best solution out of three. The time for the Nelder-Mead algorithm could be further reduced with fewer restarts since the majority of datasets are successfully registered in one attempt. In this case, the average time is approximately 1.5 mins. It could be made even faster by using fewer iterations and a smaller search range. In the absence of probe pressure, we can expect the y-direction to be close to zero-offset.

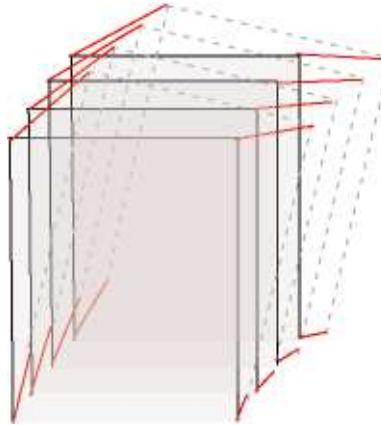


Figure 3: **Difference between true and estimated volume position.** The probe is applied at the top and the direction of insonification is downwards. The lines connecting the corner points represent the Euclidean distance between the actual positions obtained using the position sensor (solid boundaries of B-scan) and those obtained using similarity measures (dotted boundaries of B-scan) and are used in the calculation of root mean square error.

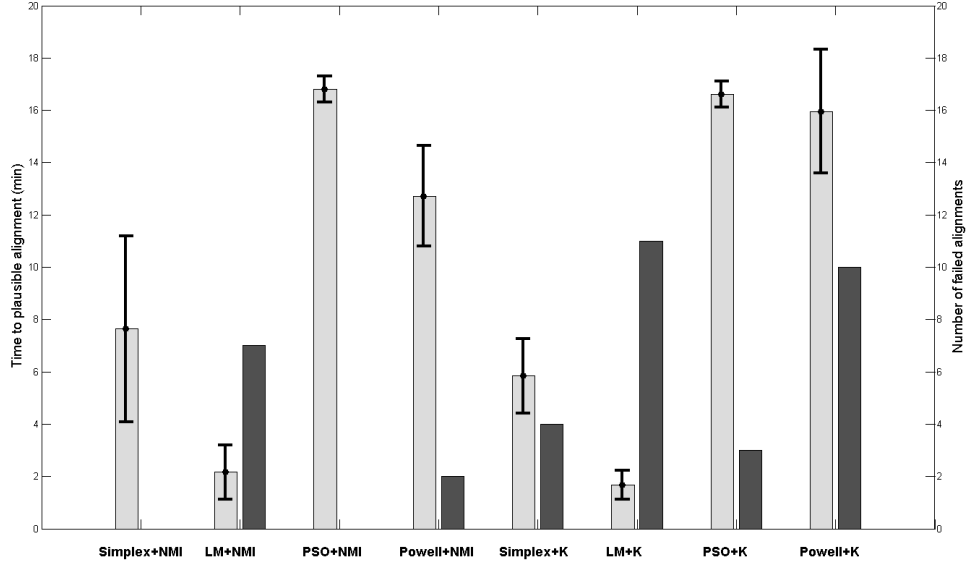


Figure 4: **Comparison of performance of the different algorithms for the *in vitro* datasets.** For each algorithm, the left bar shows the execution time and the right bar shows the number of failed alignments (refer to Table 2). For the Nelder-Mead simplex algorithm, the Levenberg-Marquardt method, and Powell’s direction-set method, the execution time is shown including the two restarts. The error bars show \pm one standard deviation of the execution time. In the majority of cases, the Nelder-Mead simplex algorithm took around 5 minutes. However, the execution time for the P2 dataset was 13 minutes on average which explains why there is a larger error bar for the Nelder-Mead simplex algorithm with normalised mutual information. The Levenberg-Marquardt method is the fastest among the considered algorithms. However, it is not very reliable. Normalised mutual information is better suited for image registration as there are fewer misregistrations. The majority of the failed alignments are on those datasets that have a small amount of overlap between volumes (P4 and P5). Since Kolmogorov’s similarity measure is less robust to small overlaps, there are a larger number of misalignments with this measure.

4 Conclusions

The goal of this paper was to present and assess several search procedures for rigid registration of three-dimensional ultrasound volumes when the orientation information is known. By considering reslice image pairs instead of a regular voxel grid, a fast computational procedure can be obtained. Using this technique, we have presented several *in vitro* and *in vivo* experimental results to show that the Nelder-Mead simplex algorithm is potentially suitable for the registration task. In general, we found that normalised mutual information is better suited for image registration in three-dimensional ultrasound as it has the ability to register datasets with partial overlaps. We have also provided a comparative analysis with Powell’s direction set algorithm, Levenberg-Marquardt method and the particle swarm optimisation algorithm. We were able to register all the *in vivo* datasets even when they had very little overlap (say 60%). Also, we are reporting a minimum registration time of less than 2 mins. It may also be possible to reduce the execution time further by considering a GPU implementation of the registration algorithm, either to interpolate the slices in parallel or to let each vertex of the simplex work on a separate parallel core. In the future, several criteria should be studied to compare precisely the behaviour of each step of the

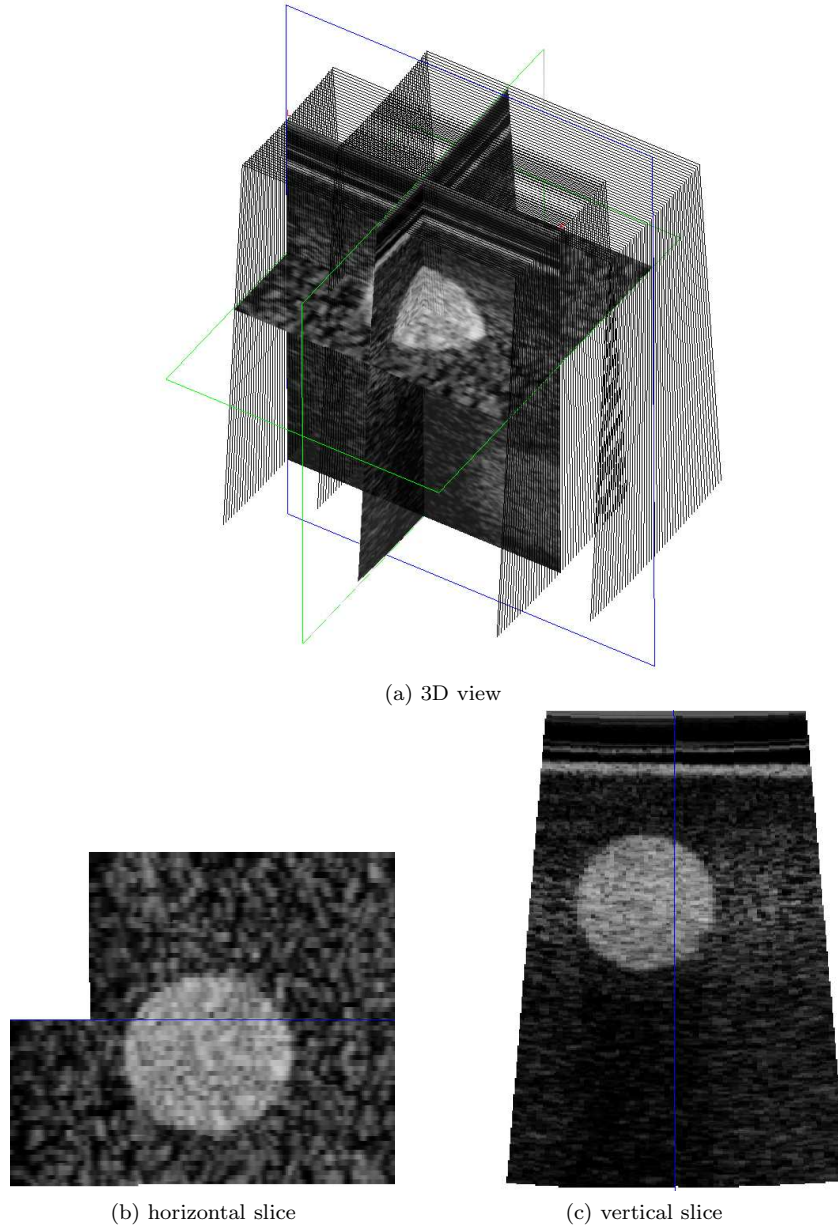


Figure 5: **Alignment results for P5 with Nelder-Mead's simplex algorithm.** (a) A 3D view showing the location of each image and a vertical and a horizontal slice through the registered volumes. The third outline (blue) roughly parallel to the individual images indicates the location of the *dividing plane* which approximately bisects the overlap region. (b) and (c) are the reslice images. The line that passes through the inclusion is the intersection with the dividing plane. One side of the line shows data from one volume and the other side shows the other volume.

registration procedure. For instance, the influence of the interpolation procedure, the histogram binning algorithms (automatic bin size estimation [46, 47, 48] for probabilistic similarity measures), probe pressure correction, and multi-scale approach [3] may all lead to further improvements in reliability as well as speed.

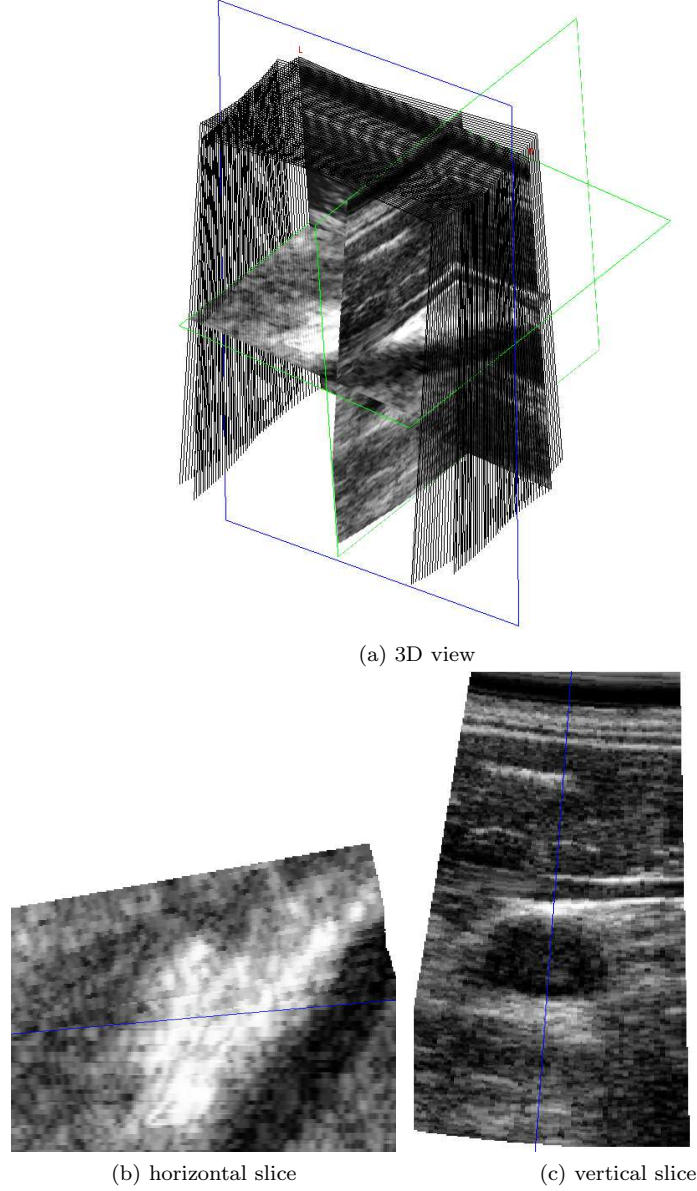


Figure 6: **Alignment results for R1 with Nelder-Mead's simplex algorithm.** The detailed description of the figure is given in the caption to Fig. 5.

Acknowledgements

The authors wish to thank Ms. Norma Umer for drawing the diagrams for Figs. 1 and 3, and Dr. James Housden for helping in the *in vivo* experiments. Umer Zeeshan Ijaz is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/F016476/1.

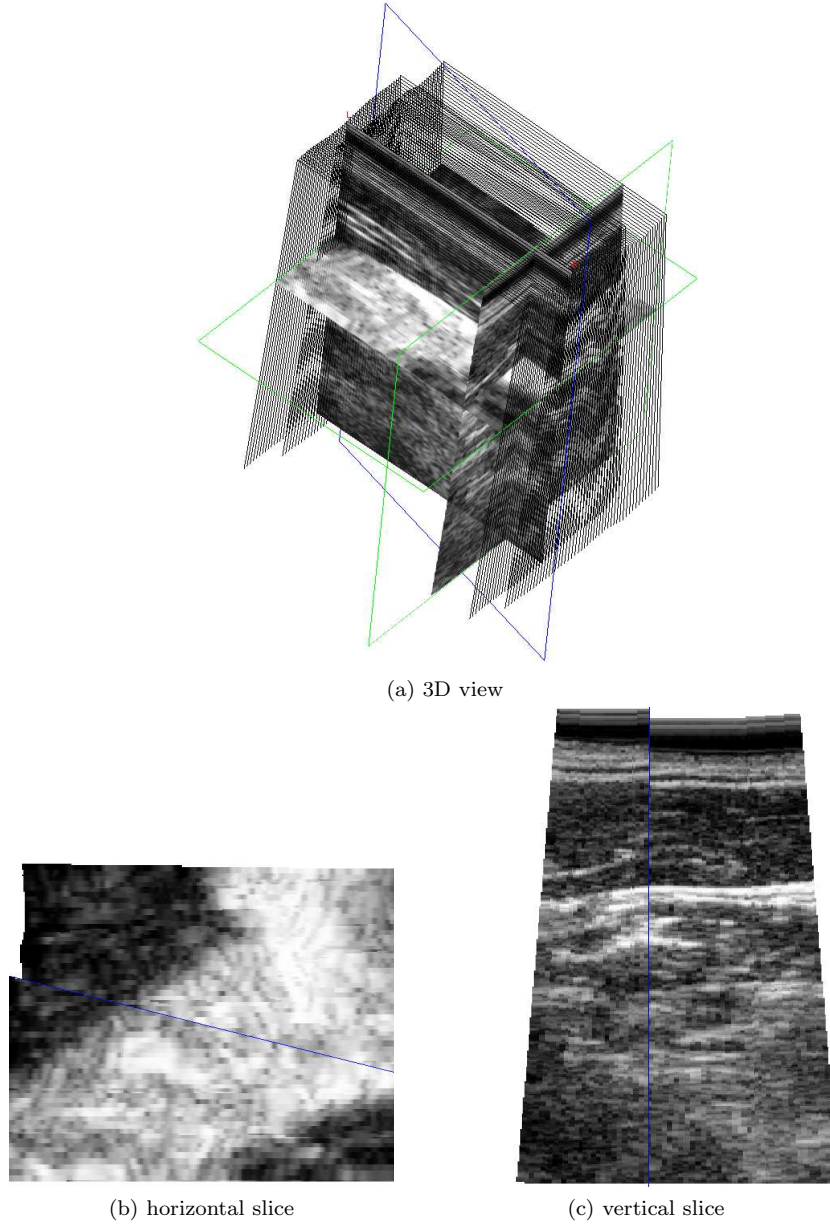


Figure 7: **Alignment results for R2 with Nelder-Mead's simplex algorithm.** The detailed description of the figure is given in the caption to Fig. 5. Here, minor misregistration is evident at the top due to probe pressure artifacts. It should be noted that we do not consider the top region of the image in the reslices as these are partly imaging the probe face rather than the tissue.

APPENDIX: Search Algorithms

A Nelder-Mead Simplex Algorithm

The downhill simplex method is due to Nelder and Mead [41, 49]. The method requires only function evaluations, not derivatives, although it is not very efficient in terms of the number of function evaluations that it requires. The procedure is as follows:

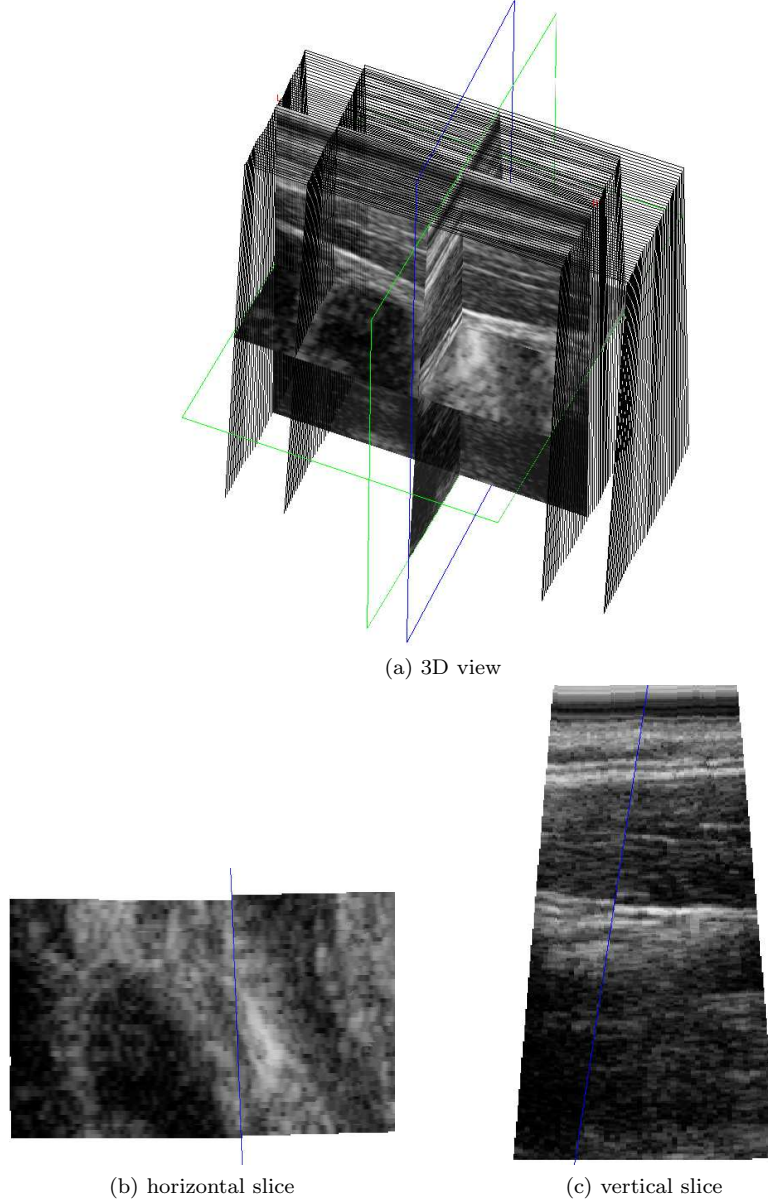


Figure 8: **Alignment results for M1 with particle swarm optimiser.** The detailed description of the figure is given in the caption to Fig. 5.

- Step 1 (Order): For N variables, order $N + 1$ vertices as $f(\mathbf{x}_1) \leq f(\mathbf{x}_2) \dots \leq f(\mathbf{x}_{N+1})$ by evaluating the similarity measure at each vertex point of the simplex. Let \mathbf{x}_0 be the centre of gravity of all points except \mathbf{x}_{N+1} . Here, we are considering the minimisation case by pre-multiplying the similarity measure by -1.
- Step 2 (Reflection): Generate a new vertex \mathbf{x}_r by reflecting the worst point according to

$$\mathbf{x}_r = (1 + \alpha)\mathbf{x}_0 - \alpha\mathbf{x}_{N+1} \quad (7)$$

where α is the reflection coefficient ($\alpha > 0$). If the reflected point is better than the second worst, but not better than the best, i.e. $f(\mathbf{x}_1) \leq f(\mathbf{x}_r) < f(\mathbf{x}_N)$, accept the reflection by replacing \mathbf{x}_{N+1} with \mathbf{x}_r and go to step 1.

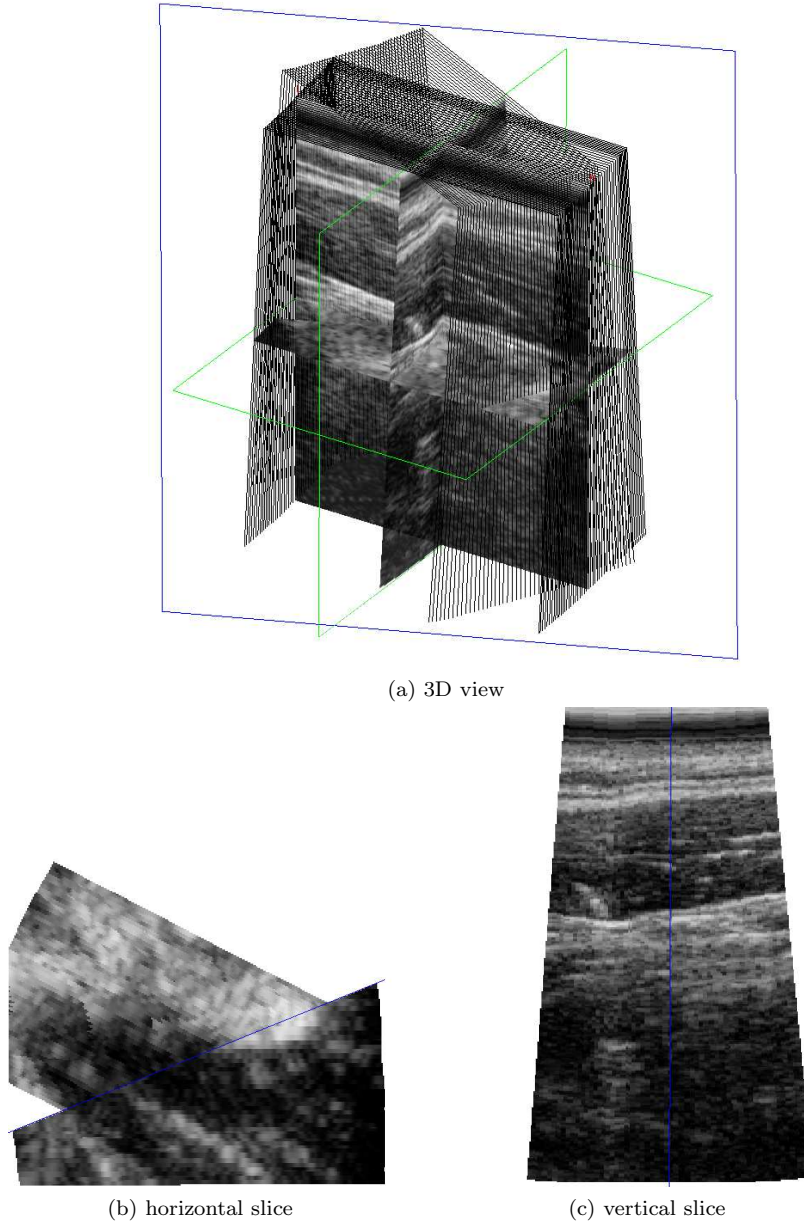


Figure 9: **Alignment results for M2 with Powell's direction-set algorithm.** The detailed description of the figure is given in the caption to Fig. 5.

- Step 3 (Expansion): If the reflection produces a function value smaller than $f(\mathbf{x}_1)$ (i.e., $f(\mathbf{x}_r) < f(\mathbf{x}_1)$), the reflection is expanded in order to extend the search space in the same direction and the expansion point is then calculated by

$$\mathbf{x}_e = \mathbf{x}_0 + \gamma(\mathbf{x}_0 - \mathbf{x}_{N+1}) \quad (8)$$

where γ is the expansion coefficient ($\gamma > 1$). If the expanded point is better than the reflected point, $f(\mathbf{x}_e) < f(\mathbf{x}_r)$, then obtain a new simplex by replacing the worst point \mathbf{x}_{N+1} with the expanded point \mathbf{x}_e and go to step 1. Otherwise, obtain a new simplex by replacing the worst point \mathbf{x}_{N+1} with the reflected point \mathbf{x}_r , and go to step 1. If the reflected point is worse than the second worst, continue to step 4.

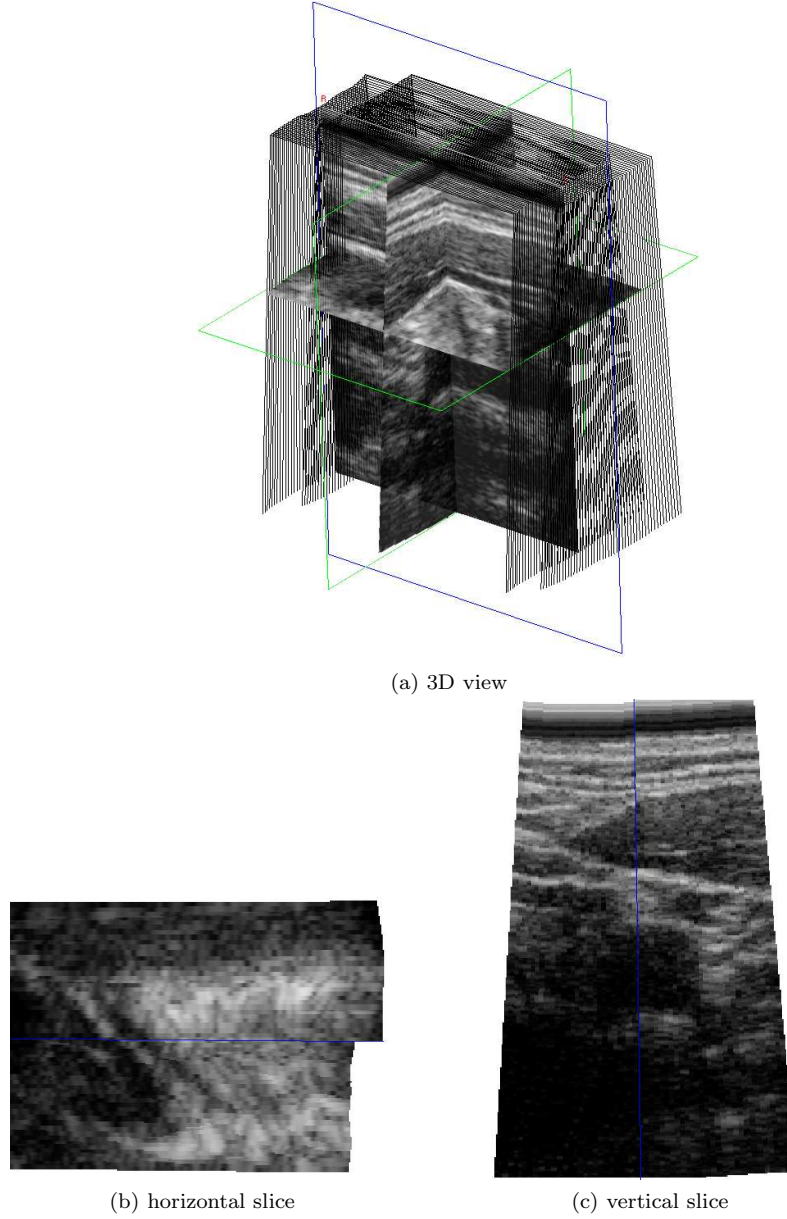


Figure 10: **Alignment results for M3 with Levenberg-Marquardt method.** The detailed description of the figure is given in the caption to Fig. 5. This is a particularly difficult dataset to register because the anatomy is such that the features tend not to vary in the lateral direction. As a result it is difficult to find a peak in the similarity measure in this direction.

- Step 4 (Contraction): When $f(\mathbf{x}_r) \geq f(\mathbf{x}_n)$, compute the contraction point

$$\mathbf{x}_c = \mathbf{x}_{N+1} + \beta(\mathbf{x}_0 - \mathbf{x}_{N+1}) \quad (9)$$

where β is the contraction coefficient ($0 < \beta < 1$). If the contracted point is better than the worst point, i.e. $f(\mathbf{x}_c) \leq f(\mathbf{x}_{N+1})$, then obtain a new simplex by replacing the worst point \mathbf{x}_{N+1} with the contracted point \mathbf{x}_c , and go to step 1. Otherwise, proceed to step 5.

- Step 5 (Reduction): For all but the best point ($i \in \{2, \dots, N+1\}$), replace the point with

$$\mathbf{x}_i = \mathbf{x}_1 + \delta(\mathbf{x}_i - \mathbf{x}_1) \quad (10)$$

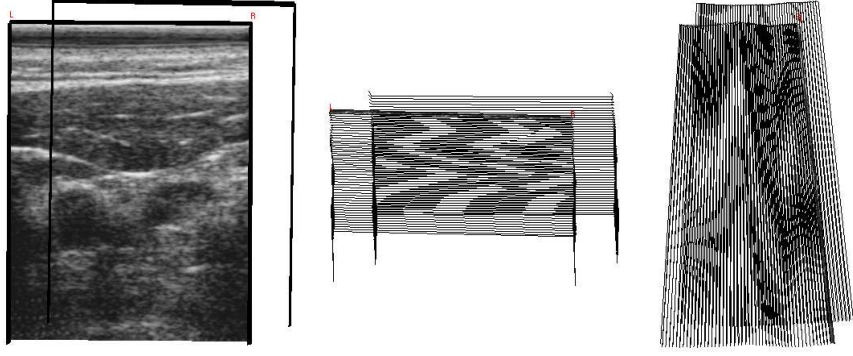


Figure 11: **Position artifacts.** The front-view, top-view, and side-view of the unregistered M1 dataset is shown here. During the acquisition, there was a rigid movement of the scanning subject between the two volumes. As a result the skin surface according to the position sensor appears to have moved downwards and therefore it is not correctly registered using the position sensor.

Datasets	Nelder-Mead simplex method	Particle swarm optimiser	Powell's direction-set algorithm	Levenberg Marquardt method
R1	0.726	0.684	21.3	0.731
R2	0.630	0.634	15.438	8.215
R3	0.994	0.984	21.757	1.037
R4	0.788	0.789	8.445	0.724
R5	0.821	0.795	0.809	0.729
M1	2.771	2.784	2.767	2.790
M2	4.931	4.950	4.919	4.941
M3	1.541	1.537	2.261	1.246

Table 3: **RMSE(mm) for *in vivo* experiments using normalised mutual information as the similarity measure.** The datasets consist of scans of the neck and calf muscles. We have used two restarts in the Nelder-Mead simplex algorithm, the Levenberg-Marquardt method, and Powell's direction-set method. Since it is important in clinical practise that the algorithms are successful in only one attempt, we have shown the result of only one trial per algorithm. Except the PSO, all algorithms were allowed two restarts per trial. Two classes of *in vivo* datasets are shown: those starting with R represent the case where every effort was made to ensure that the scanning subject was stationary (no breathing or movement) at the time of acquisition so that the position sensor readings are close to the correct alignment; those starting with M represent the cases when small movement (see Fig. 11) or probe pressure were allowed. Apart from the highlighted results (which didn't converge based on visual inspection), the R datasets have a much lower RMSE than the M datasets.

and go to step 1. Here δ is the shrinkage coefficient ($0 < \delta < 1$). This step is necessary when the contraction has failed and it shrinks the entire simplex.

B Powell's Direction-set Method

The multidimensional method consists of sequences of one-dimensional line minimisation in an appropriate direction in the search space [42, 49]. It is a direction-set method that consist of prescriptions for updating the set of directions as the method proceeds. It determines a set of directions that take us far along narrow valleys in the cost function and also along the non-interfering

directions in which the minimisation along one is not spoiled by subsequent minimisation along another. The procedure is as follows:

Let \mathbf{x}_0 be an initial guess at the location of the minimum of the function $f(x_1, x_2, \dots, x_N)$, $\mathbf{E}_k = [0, 0, \dots, 0, 1_k, 0, \dots, 0] : k = 1, 2, \dots, N$ be the set of standard basis vectors $\mathbf{U} = [\mathbf{E}_1^T, \mathbf{E}_2^T, \dots, \mathbf{E}_N^T]$ and $i = 0$, then

- Step 1: Set $\mathbf{P}_0 = \mathbf{x}_i$.
- Step 2: For $k = 1, 2, \dots, N$ find the value of γ_k that minimises $f(\mathbf{P}_{k-1} + \gamma_k \mathbf{U}_k)$ and set $\mathbf{P}_k = \mathbf{P}_{k-1} + \gamma_k \mathbf{U}_k$.
- Step 3: Set r and \mathbf{U}_r equal to the maximum decrease in f and the direction of the maximum decrease, respectively, over all the direction vectors in step 2. It should be noted that like in the previous algorithm, we are considering the minimisation case so we pre-multiply the similarity measure by -1.
- Step 4: Set $i = i + 1$.
- Step 5: If $f(2\mathbf{P}_N - \mathbf{P}_0) \geq f(\mathbf{P}_0)$ or $2(f(\mathbf{P}_0) - 2f(\mathbf{P}_N) + f(2\mathbf{P}_N - \mathbf{P}_0))(f(\mathbf{P}_0) - f(\mathbf{P}_N) - r)^2 \geq r(f(\mathbf{P}_0) - f(2\mathbf{P}_N - \mathbf{P}_0))^2$, then set $\mathbf{x}_i = \mathbf{P}_N$ and return to step 1. Otherwise, go to step 6.
- Step 6: Set $\mathbf{U}_r = \mathbf{P}_N - \mathbf{P}_0$.
- Step 7: Find the value of γ that minimises $f(\mathbf{P}_0 + \gamma \mathbf{U}_r)$. Set $\mathbf{x}_i = \mathbf{P}_0 + \gamma \mathbf{U}_r$.
- Step 8: Repeat steps 1 through 7.

Here the first inequality in step 5 indicates that there is no further decrease in the value of f in the average direction $\mathbf{P}_N - \mathbf{P}_0$. The second inequality indicates that the decrease in the function f in the direction of greatest decrease \mathbf{U}_r was not a major part of the total decrease in f in step 2. If the conditions in step 5 are not satisfied, then the direction of greatest decrease \mathbf{U}_r is replaced with the average direction from step 2, i.e., $\mathbf{P}_N - \mathbf{P}_0$. In step 7, the function is minimised in this direction. In step 2 and step 7, for one-dimensional minimisation without calculation of derivatives, we use *Brent's method*. However, this method is not quite as effective as it should be [49] and when the function has a discontinuous second (or lower) derivative, the *golden section search* is used as an alternative. In the current implementation, we have used the two minimisation algorithms together by initialising Brent's method with the results obtained from the golden section search for every one-dimensional line minimisation.

C Particle Swarm Optimiser

The particle swarm optimisation (PSO) algorithm is an evolutionary computation technique proposed by Kennedy and Everhart in 1995 [43] and is based on the social behaviour of a swarm. It has been widely used in a variety of optimisation problems [50, 51]. The PSO algorithm is described as:

$$\mathbf{v}_{k+1} = a\mathbf{v}_k + b_1\mathbf{r}_1 \otimes (\mathbf{p}_1 - \mathbf{x}_k) + b_2\mathbf{r}_2 \otimes (\mathbf{p}_2 - \mathbf{x}_k) \quad (11)$$

$$\mathbf{x}_{k+1} = c\mathbf{x}_k + d\mathbf{v}_{k+1} \quad (12)$$

The symbol \otimes denotes element-by-element vector multiplication. Equation (11) is used to calculate the velocity \mathbf{v}_{k+1} of the particle \mathbf{x}_k according to its current position from its own best experience \mathbf{p}_1 and the group's best experience \mathbf{p}_2 . b_1 and b_2 represent the strength of attraction towards its own best experience and the group's best experience, respectively. The momentum factor a is used to control the influence of the previous history of velocities on the current velocity. The particle moves toward a new position according to Eq. (12) affected by the coefficients c and

d. Here, \mathbf{x}_k represents the three rigid-body translation parameters for the source volume. The random numbers \mathbf{r}_1 and \mathbf{r}_2 are introduced for good state space exploration. They are selected in the range $[0,1]$. At each iteration, the similarity measure is used to update the best position \mathbf{p}_1 of each particle if it scores higher than the similarity measure for the previous best position. Similarly, at every iteration, the globally best position in the whole swarm \mathbf{p}_2 is also saved.

Trelea [52] performed the theoretical analysis of PSO, by considering the deterministic version in one-dimension. It was obtained by setting the random numbers to their expected values, i.e., $r_1 = r_2 = \frac{1}{2}$. Also, Eq. (11) was simplified by considering $b = \frac{b_1 + b_2}{2}$, and $p = \frac{b_1}{b_1 + b_2}p_1 + \frac{b_2}{b_1 + b_2}p_2$. The newly introduced coefficient b represents the average of the own and social attraction coefficients b_1 and b_2 . The attraction point p is the weighted average of p_1 and p_2 . It was found that the PSO is convergent if the following sets of conditions are satisfied:

$$a < 1, b > 0, 2a - b + 2 > 0 \quad (13)$$

These tuning parameters can greatly influence the performance of the PSO algorithm, often stated as the exploration-exploitation tradeoff. Exploration is the ability to test various regions in the problem space to reach a global optimum and exploitation is the ability to search around the candidate solution. We have previously tested different parameter settings for the PSO algorithm and found Trelea parameter set 2 [36] to be appropriate in the current context. These parameters originated from a study by Clerc [53] and favour exploration of the search space over exploitation. Details are given in Trelea [52] where various parameter sets were tested on different cost functionals, and it was shown that the PSO algorithm has a much higher success rate with Trelea parameter set 2. Therefore, in the current implementation, we have only considered this parameter set, i.e., $a = 0.729$, $b_1 = b_2 = 1.494$, and $c = d = 1$.

D Levenberg-Marquardt Method

The levenberg-Marquardt method [44] is primarily used as a non-linear least squares minimisation technique and reduces the sum of square of residuals.

$$g(\mathbf{x}) = \frac{1}{2} \sum_{j=1}^m r_j^2(\mathbf{x}) \quad (14)$$

Here, $\mathbf{x} = (x, y, z)$, and r_j is the residual function defined from \Re^n to \Re as

$$r(\mathbf{x}) = \beta - f(\mathbf{x}) \quad (15)$$

In the current implementation $f(\mathbf{x})$ is the similarity measure score for a single reslice pair from the corresponding volumes and β is the maximum possible value of similarity measure. Assembling the residual vector for all image pairs, we have

$$\mathbf{r}(\mathbf{x}) = (r_1(\mathbf{x}), r_2(\mathbf{x}), \dots, r_m(\mathbf{x})) \quad (16)$$

Now g can be rewritten as $g(\mathbf{x}) = \frac{1}{2} \|\mathbf{r}(\mathbf{x})\|^2$. The derivative of g can be written using the Jacobian matrix \mathbf{J} of \mathbf{r} w.r.t. \mathbf{x} defined as $\mathbf{J}(\mathbf{x}) = \frac{\partial r_j}{\partial x_i}$, $1 \leq j \leq m$, $1 \leq i \leq n$. By defining $\nabla g(\mathbf{x}) = \mathbf{J}(\mathbf{x})^T \mathbf{r}(\mathbf{x})$, and the Hessian \mathbf{H} as $\mathbf{J}(\mathbf{x})^T \mathbf{J}(\mathbf{x})$, the Levenberg-Marquardt update is given as

$$\mathbf{x}_{j+1} = \mathbf{x}_j - (\mathbf{H} + \lambda \text{diag}[\mathbf{H}])^{-1} \nabla g(\mathbf{x}_j) \quad (17)$$

The Levenberg-Marquardt method is thus

- Do an update as directed by the rule above.

- Evaluate the error at the new parameter vector.
- If the error has increased as a result of the update, then reset the increments to their previous value and increase λ by 10 or some other significant factor. Then go to the first step again.
- If the error has decreased as a result of the update, then accept the step by keeping the weights at their new values and decrease λ by 10 or so.

References

- [1] A. Fenster and D. B. Downey. 3-D ultrasound imaging-a review. *IEEE Eng. Med. Biol. Mag.*, 15.
- [2] A. H. Gee, R. Prager, G. M. Treece, and L. Berman. Engineering a freehand 3D ultrasound system. *Pattern Recognition Letters*, 24(4-5):757–777, 2003.
- [3] 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.
- [4] N. Pagoulatos, W. S. Edwards, D. R. Haynor, and Y. Kim. Interactive 3D registration of ultrasound and magnetic resonance images based on a magnetic position sensor. *IEEE Transactions on Information Technology in Biomedicine*, 3(4):278–288, 1999.
- [5] M. Nakamoto, Y. Sato, M. Miyomato, Y. Nakamjima, K. Konishi, M. Shimada, M. Hashizume, and S. Tamura. 3D ultrasound system using a magneto-optic hybrid tracker for augmented reality visualization in laproscopic liver surgery. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 2002*, volume 2489/-1/2002, pages 148–155, 2002.
- [6] P. R. Detmer, G. Bashein, T. Hodges, K. W. Beach, E. P. Filer, D. H. Burns, and D. E. Strandness Jr. 3D ultrasonic image feature localization based on magnetic scanhead tracking: In vitro calibration and validation. *Ultrasound in Medicine and Biology*, 20(4):923–936, 1994.
- [7] D. F. Leotta, P. R. Detmer, and R. W. Martin. Performance of miniature magnetic position sensor for three-dimensional ultrasound imaging. *Ultrasound in Medicine and Biology*, 23(4):597–609, 1997.
- [8] C. D. Barry, C. P. Allot, N. W. John, P. M. Mellor, P. A. Arundel, D. S. Thomson, and J. C. Waterton. Three-dimensional freehand ultrasound: image reconstruction and volume analysis. *Ultrasound in Medicine and Biology*, 23(8):1209–1224, 1997.
- [9] 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.
- [10] J. Pratikakis, C. Barillat, and P. Hellia. Robust multi-scale non-rigid registration of 3D ultrasound images. In *Scale-space and Morphology in Computer Vision*, volume 2106, pages 389–397, 2001.
- [11] J. F. Krücker, C. R. Meyer, G. L. LeCarpontier, J. B. Fowlkes, and P. L. Carson. 3D spatial compounding of ultrasound images using image-based nonrigid registration. *Ultrasound in Medicine and Biology*, 26(9):1475–1488, 2000.
- [12] T. Lange, S. Eulenstein, M. Hünerbein, H. Lamecker, and P. M. Schlag. Augmenting intraoperative 3D ultrasound with preoperative models for navigation in liver surgery. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 2004*, volume LNCS 3217/2004, pages 534–541, 2004.

- [13] J. M. Blackall, D. Rueckert, C. R. Maurer Jr., G. P. Penney, D. L. G. Hill, and D. J. Hawkes. An image registration approach to automated calibration for freehand 3D ultrasound. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 2000*, volume 1935/2000, pages 465–471, Pittsburgh, PA, USA, 2000.
- [14] M. M. J. Letteboer, M. A. Viergever, and W. J. Nessen. Rapid registration of 3D ultrasound data of brain tumor. In *Proceedings of the 17th International Congress and Exhibition on Computer Assisted Radiology and Surgery-CARS 2003, International Congress Series*, volume 1256, pages 433–439, 2003.
- [15] M. M. J. Letteboer, P. W. A. Willems, M. A. Viergever, and W. J. Niessen. Non-rigid registration of 3D ultrasound images of brain tumors acquired during neurosurgery. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 2003*, volume LNCS 2879/2003, pages 408–415, 2003.
- [16] R. Shekar and V. Zagrodsky. Mutual information based rigid and nonrigid registration of ultrasound volumes. *IEEE Transactions on Medical Imaging*, 21(1):9–22, 2002.
- [17] T. Abel, X. Morandi, R. M. Comeau, and D. L. Collins. Automatic non-linear MRI-ultrasound registration for the correction of intra-operative brain deformation. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 2001*, volume LNCS 2208, pages 913–922, 2001.
- [18] C. Wachinger, W. Wein, and N. Navab. Registration strategies and similarity measures for three-dimensional ultrasound mosaicing. *Academic Radiology*, 15:1404–1415, 2008.
- [19] R. Francois, R. Fablet, and C. Barillot. Robust statistical registration of 3D ultrasound images using texture information. In *Proceedings of the 2003 International Conference on Image Processing*, volume 1, pages 581–584, Barcelona, Spain, 2003.
- [20] S. Gao, Y. Xiao, and S. Hu. A comparison of two similarity measures in intensity-based ultrasound image registration. In *Proceedings of the 2004 International Symposium on Circuits and Systems*, volume 4, pages 61–64, Vancouver, British Columbia, May 2004.
- [21] A. Roche, X. Pennec, M. Rudolph, D. P. Auer, G. Malandain, S. Ourselin, L. M. Aueor, and N. Ayache. Generalized correlation ratio for rigid registration of 3D ultrasound with MR images. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 2000*, volume LNCS 1935, pages 203–220, 2000.
- [22] A. Leroy, P. Mozer, Y. Payan, and J. Troccaz. Rigid registration of free hand 3D ultrasound and CT scan kidney images. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 2004*, volume LNCS 3216, pages 837–844, 2004.
- [23] G. Xia, M. Brady, J. A. Noble, M. Burcher, and R. English. Nonrigid registration of 3-D free-hand ultrasound images of the breasts. *IEEE Transactions on Medical Imaging*, 21(4):405–412, 2002.
- [24] B. Brendel, S. Winter, A. Rick, M. Stockheim, and H. Ermeret. Registration of 3D CT and ultrasound datasets of the spine using bone structures. *Computer Aided Surgery*, 7(3):146–155, 2002.
- [25] S. Winter, B. Brendel, and C. Igel. Registration of bone structure in 3D ultrasound volume and CT data: Comparison of different optimization strategies. In *Proceedings of Computer Assisted Radiology and Surgery-CARS 2005, International Congress Series*, volume 1281, pages 242–247, May 2005.
- [26] X. Pennec, P. Cachier, and N. Ayache. Tracking brain deformations in time sequences of 3D US images. *Pattern Recognition Letters*, 24(4-5):801–813, 2003.

- [27] D. Zikic, W. Wein, A. Khamene, D. A. Clevert, and N. Navab. Fast deformable registration of 3D-ultrasound data using a variational approach. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 2006*, volume 4190, pages 915–923, 2006.
- [28] G. Ionescu, S. Lavalée, and J. Demongeot. Automatic registration of ultrasound with CT images: Application to computer assisted prostate radio therapy and orthopedics. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention-MICCAI 99*, volume LNCS 1679, pages 768–779, 1999.
- [29] G. M. Treece, R. W. Prager, A. H. Gee, and L. Berman. Correction of probe pressure artifacts in freehand 3D ultrasound. *Medical Image Analysis*, 6(3):199–214, 2002.
- [30] J. M. Sanches and J. S. Marques. Joint image registration and volume reconstruction for 3D ultrasound data. *Pattern Recognition Letters*, 24(4-5):791–800, 2003.
- [31] D. Boukerroui, J. Noble, and M. Brady. Velocity estimation in ultrasound images: A block matching approach. In *IPMI 2003*, volume LNCS 2732, pages 586–598, 2003.
- [32] C. Kotropoulos, X. Magnisalis, I. Pitas, and M. G. Strintzis. Nonlinear ultrasonic image processing based on signal-adaptive filters and self-organizing neural networks. *IEEE Transactions on Image Processing*, 3(1):65–77, 1994.
- [33] M. G. Strintiz and I. Kokkinidis. Motion estimation in ultrasound image sequences. *IEEE Signal Processing Letters*, 4(6):156–157, 1997.
- [34] B. Cohen and I. Dinstein. New maximum likelihood motion estimation schemes for noisy ultrasound images. *Pattern Recognition*, 35:455–463, 2002.
- [35] U.Z. Ijaz, R.W. Prager, A.H. Gee, and G.M. Treece. A study of similarity measures for in vivo 3D ultrasound volume registration. In *30th International Acoustical Imaging Symposium*, Monterey, CA, USA, March 2009.
- [36] U.Z. Ijaz, R.W. Prager, A.H. Gee, and G.M. Treece. Particle swarm optimization for in vivo 3D ultrasound volume registration. In *30th International Acoustical Imaging Symposium*, Monterey, CA, USA, March 2009.
- [37] N. Barbour and G. Schmidt. Inertial sensor technology trends. *IEEE Sensors Journal*, 1(4):332–339, 2001.
- [38] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(379-423):623–656, 1948.
- [39] M. Basseville. Information: Entropies, divergence et moyennes. Technical Report 1020, IRISA-35042, Rennes Cedex France, May 1996.
- [40] D. Sarrut and S. Miguet. Similarity measures for image registration. In *First European Workshop on Content-Based Multimedia Indexing*, pages 263–270, Toulouse, France, October 1999.
- [41] J. A. Nelder and R. Mead. A simplex method for function minimization. *Computer journal*, 7:308–313, 1965.
- [42] M. J. D. Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Computer journal*, 7:155–162, 1964.
- [43] J. Kennedy and R. C. Eberhart. Particle swarm optimization. In *Proceedings of IEEE International Conference on Neural Networks*, pages 1942–1948, 1995.
- [44] J. J. Moré. The Levenberg-Marquardt algorithm: Implementation and theory. *Lecture Notes in Mathematics*, 630.

- [45] R. W. Prager, R. N. Rohling, A. H. Gee, and L. Berman. Rapid calibration for 3-D freehand ultrasound. *Ultrasound in Medicine and Biology*, 24(6):855–869, July 1998.
- [46] Z. Zhang. Estimating mutual information via Kolmogorov distance. *IEEE Transactions on Information Theory*, 53(9):3280–3282, 2007.
- [47] P. A. Legg, P. L. Rosin, D. Marshall, and J. E. Morgan. Improving accuracy and efficiency of registration by mutual information using Sturges’ histogram rule. July 2007.
- [48] M. P. Wand. Data-based choice of histogram bin width. *The American Statistician*, 51(1):59–64, 1997.
- [49] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C++: The Art of Scientific Computing*. Cambridge University Press, February 2002.
- [50] A. Banks, J. Vincent, and C. Anyakoha. A review of particle swarm optimization. Part I: background and development. *Natural Computing*, 6:467–484, 2007.
- [51] A. Banks, J. Vincent, and C. Anyakoha. A review of particle swarm optimization. Part II: hybridisation, combinatorial, multicriteria and constrained optimization, and indicative applications. *Natural Computing*, 7:109–124, 2008.
- [52] I. C. Trelea. The particle swarm optimization algorithm: convergence analysis and parameter selection. *Information Processing Letters*, 85:317–325, 2003.
- [53] M. Clerc. The swarm and the queen: towards a deterministic and adaptive particle swarm optimization. In *ICEC*, pages 1951–1957, Washington DC, USA, 1999.