Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT®

Medical Image Analysis 8 (2004) 81–91

MEDICAL IMAGE ANALYSIS

www.elsevier.com/locate/media

Registration of freehand 3D ultrasound and magnetic resonance liver images

G.P. Penney [a], J.M. Blackall [a], M.S. Hamady [b], T. Sabharwal [b], A. Adam [b], D.J. Hawkes [a,*]

[a] *Division of Imaging Sciences, 5th Floor Thomas Guy House, Guy's Hospital, London SE1 9RT, UK*
[b] *Department of Interventional Radiology, 1st Floor, Lambeth Wing, St. Thomas' Hospital, Lambeth Palace Road, London SE1 7EH, UK*

Received 6 January 2003; received in revised form 10 June 2003; accepted 17 July 2003

Abstract

We present a method to register a preoperative MR volume to a sparse set of intraoperative ultrasound slices. Our aim is to allow the transfer of information from preoperative modalities to intraoperative ultrasound images to aid needle placement during thermal ablation of liver metastases. The spatial relationship between ultrasound slices is obtained by tracking the probe using a Polaris optical tracking system. Images are acquired at maximum exhalation and we assume the validity of the rigid body transformation. An initial registration is carried out by picking a single corresponding point in both modalities. Our strategy is to interpret both sets of images in an automated pre-processing step to produce evidence or probabilities of corresponding structure as a pixel or voxel map. The registration algorithm converts the intensity values of the MR and ultrasound images into vessel probability values. The registration is then carried out between the vessel probability images. Results are compared to a "bronze standard" registration which is calculated using a manual point/line picking algorithm and verified using visual inspection. Results show that our starting estimate is within a root mean square target registration error (calculated over the whole liver) of 15.4 mm to the "bronze standard" and this is improved to 3.6 mm after running the intensity-based algorithm.
© 2003 Elsevier B.V. All rights reserved.

Keywords: Ultrasound to MR registration; Sparse data; Liver metastases

1. Introduction

The registration of preoperative volumes to intraoperative modalities has been proposed to aid a number of interventional and surgical procedures. This paper investigates a method to register a preoperative MR volume to a sparse set of ultrasound images. Most previous work on MR or CT to ultrasound registration has been to aid image guided neurosurgery (Bucholz et al., 1997; Comeau et al., 2000; Jödicke et al., 1998; King et al., 2000; Roche et al., 2000). Some of these approaches make use of standard image guided surgery registration techniques, using either fiducial or anatomical landmarks, to register the preoperative and ultrasound images together. In these cases the main purpose of the ultrasound is to detect and correct for tissue deformations (Bucholz et al., 1997; Comeau et al., 2000; King et al., 2000). Other approaches directly register the image volumes together, either using feature- (Jödicke et al., 1998) or voxel-based (Roche et al., 2001) approaches. The registration of MR or CT to ultrasound has also been carried out on forearm (Porter et al., 2001), carotid artery (Slomka et al., 2001) and liver images (Aylward et al., 2002; Porter et al., 2001; Voirin et al., 2002). The approaches range from feature-based matching (Porter et al., 2001; Voirin et al., 2002) to hybrid (Aylward et al., 2002) and voxel intensity-based (Slomka et al., 2001) methods. Most of these algorithms base their registrations on vascular structures (Aylward et al., 2002; Porter et al., 2001; Slomka et al., 2001) while Voirin et al. (2002) use liver surface features. These approaches have all used a densely sampled 3D ultrasound volume. This may not be practical in some clinical

* Corresponding author.
E-mail address: david.hawkes@kcl.ac.uk (D.J. Hawkes).

1361-8415/$ - see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2003.07.003

situations as it is difficult to acquire densely sampled 3D volumes with tracked 2D probes when respiratory gating is applied. Two-dimensional phased-array systems are able to acquire dense 3D volumes very rapidly, however, they currently produce lower quality images than conventional 2D probes (Fenster et al., 2001) and may not have the field of view to acquire good coverage of large organs in a single acquisition.

This paper describes a system to register MR images to a sparse set of tracked ultrasound slices; only 10 slices were used. Our motivation is to aid image guidance during percutaneous radiofrequency ablation of liver metastases. Registration of MR or CT to abdominal US images has a number of additional difficulties compared to using intra-cranial US images. More deformation, due to patient breathing and probe pressure, is typically observed in abdominal ultrasound. Also, there are usually more artifacts present in abdominal US images and regions where no US signal is received as they lie behind highly reflective acoustic interfaces. However, larger errors can be tolerated for the treatment of liver metastases than for most neurological procedures. We estimate that, for our system to be clinically useful, registration errors should be smaller than 5 mm. Larger errors, around 10 mm, could be acceptable for the treatment of large lesions (>5 cm in diameter).

Our method is a voxel-based approach. Initially the MR image and the ultrasound slices are converted into images of vessel probability, where the image intensity represents the probability that a particular voxel or pixel contains a vascular structure. These probability images are then registered together using normalised cross-correlation as a similarity measure. Registrations are carried out on images acquired from five volunteers. Results are compared to a "bronze standard" (Jannin et al., 2002) registration which was calculated using a feature-based method that required an expert to manually identify a large number of corresponding features in both modalities.

2. Theory

Multi-modal image registration techniques have been shown to be effective for matching a number of different modality pairs. One area of multi-modal image registration which, up until recently, has seen a relatively little attention is the registration of ultrasound images to different modalities (Maintz and Viergever, 1998). This is probably due to the large differences between ultrasound images and those acquired using other modalities. Standard B-mode ultrasound images are two-dimensional, have a lower signal-to-noise ratio and more image artifacts than most other modalities, also the image is formed as a combination of speckle from homogeneous regions and reflections from boundaries.

Image registration can be split up into a framework of four main stages, namely establishing a feature space, choosing a similarity measure, defining a search or parameter space and designing a search strategy (Brown, 1992). In most intensity-based registration algorithms, the feature space is comprised of the raw intensity values from the two images. A suitable similarity measure is then used to compare these intensity values and the transformation between the two images is altered in order to optimise the value of the similarity measure. Measures based on the joint entropy histogram have been shown to be very effective in registering MR, CT and PET images (Maes et al., 1997; Studholme et al., 1996, 1997; Viola and Wells, 1995). More recently other measures such as local correlation (Weese et al., 1999) and correlation ratio (Roche et al., 1998) have been shown to perform well on these modalities. In all these cases very little image processing occurs in the formation of the feature space and so it is the role of the similarity measure to process these raw intensities in such a way so as to indicate a registration position. Typically, as the differences between the images increase the more difficult a task is asked of the similarity measure.

In order to register images which differ greatly (especially in image artifacts and null signal behind strong acoustic interfaces), we propose a strategy by which the formation of the feature space is carried out using algorithms to enhance corresponding features between the images and suppress image features that do not correspond. The algorithms produce images where the intensity values represent the probability of a particular feature being present: for the images used in this paper the features were hepatic vessels. These images are then registered together using a simple similarity measure, normalised cross-correlation (*Pearson's r*) (Press et al., 1992). In this way we are pre-processing the raw image intensity values to gather additional knowledge of the spatial distributions of intensities and particular attributes of the structures we are using to guide registration. Previous work has relied on the similarity measures being sufficiently robust to operate on raw (or in some cases gradient) image intensities.

3. Method

An overview of the registration system is given in Fig. 1. We wish to obtain a correspondence between pixel positions in the ultrasound images x_{US} and voxel positions within the MR volume x_{MR}. To achieve this, it is necessary to calculate the transformation $\mathbf{T}$ such that $x_{MR} = \mathbf{T}x_{US}$. This transformation is computed from three separate transformations, $\mathbf{T}_{cal}$, $\mathbf{T}_{track}$ and $\mathbf{T}_{reg}$ as shown in Fig. 1. The probe calibration ($\mathbf{T}_{cal}$), which transforms positions from the US image to positions relative to the infrared-emitting diodes (IREDs) attached

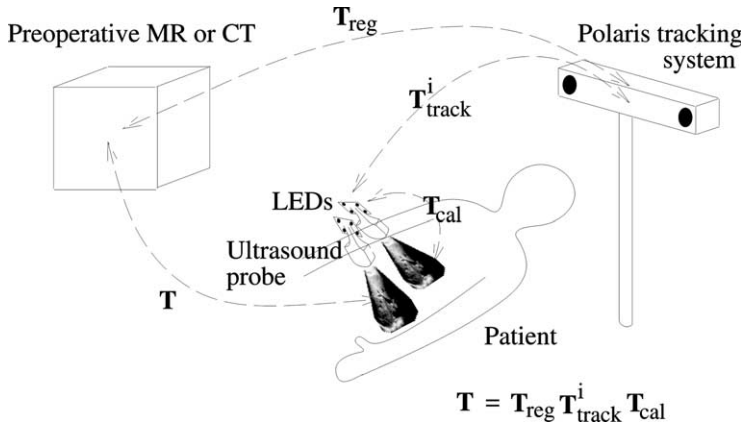

Fig. 1. Overview of system and transformations.

to the probe was calculated using a cross-wire phantom (Prager et al., 1998), this calibration technique should be accurate to 1–1.7 mm (Prager et al., 1998; Blackall et al., 2000). The tracking transformation, $\mathbf{T}_{track}$, transforms positions relative to the IREDs to positions relative to the cameras on the Polaris tracking system, this should produce a very accurate transformation as individual IREDs can be localised with a fiducial localisation error (Fitzpatrick et al., 1998) of 0.2–0.4 mm (Hajnal et al., 2001). This paper describes how we calculate $\mathbf{T}_{reg}$, the matrix which transforms positions relative to the tracking system to positions in the MR volume.

3.1. Accounting for breathing motion

It is important that all the ultrasound images are acquired at the same stage of the breathing cycle. If this is not the case then inconsistencies will arise between the position of features in different slices which may introduce significant errors into the final registration. To eliminate this potential problem, we use only US images which were acquired at maximum exhalation (the exhale position was chosen, instead of inhale, as it is more reproducible (Balter et al., 1998)).

We use the following method to acquire images at maximum exhalation. First, the US probe is placed on the inferior end of the volunteer's sternum, so that the ultrasound slice plane is approximately parallel to an axial section through the body and an US image is acquired. This image will be referred to as the "orientation initialization" (OI) image. The positional information (derived from the Polaris tracking device) associated with the OI image is used to calculate a line which goes through the OI image centre and is perpendicular to the OI image plane. This line is assumed to define an approximate superior–inferior (SI) axis of the patient. In order to acquire a single US slice at maximum exhalation the probe is positioned on the volunteer and 40 slices are acquired at a rate of approximately 10 slices/s. The tracking information for each slice is used to calculate the minimum distance between the US transducer face and the SI axis of the patient. The US image for which this distance is the smallest should correspond to the maximum exhale position, see Fig. 2. It is not appropriate to use a breath hold technique to acquire images at the same stage of the breathing cycle as patients are sedated during liver metastases procedures and cannot control their breathing.

3.2. Formation of the feature space

The US and CT images are converted from intensity images $I(x)$ into probabilistic images, $P(x)$ (where x represents the image position) in the following way. One or more features $\mathbf{F}(x)$ are calculated at each image position. Features are chosen which should allow an accurate classification between hepatic vessels and liver parenchyma. These features are transformed into vessel

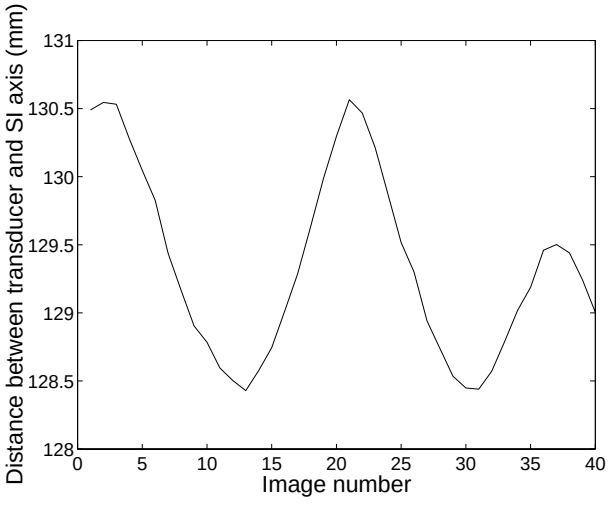

Fig. 2. Graph showing how the distance between the US transducer face and the superior–inferior axis of the patient alters over a series of 40 ultrasound images acquired at a rate of approximately 10 frames/s. For this set of images, image number 13 was calculated to be (and subsequently saved as) the image acquired at maximum exhalation.

probability images (not binary segmented images) using a probability density function p which, given a set of features $\mathbf{F}(x)$, returns an estimate of the probability that position x contains a vessel. Our aim is not to produce a perfect segmentation of the vessels, but to produce vessel probability maps of sufficient quality to allow accurate and robust registrations. This is summarised as

$$P(x) = p(\mathbf{F}(x)). \quad (1)$$

The function p was calculated using a set of training data which was kept completely separate from the data sets used for the registration experiments. For the MR images only a single feature, image intensity $I(x)$ was used. For the US images, two features were used, image intensity $I(x)$ and a feature which represents dip intensity $I_{\text{dip}}(x)$. These features, and the calculation of p, are discussed in more detail in the following sections.

3.2.1. Converting the MR image

An MR image $I_{\text{MR}}(x)$ of the liver was used as a set of training data. The liver was manually segmented from this image (and from the MR images used in the subsequent registration experiments). There were two main reasons for using segmented data. First, it greatly simplified the process of converting MR intensities to vessel probabilities as the segmented MR images contained only two tissue types, hepatic vessels and liver parenchyma, which can be separated reasonably well using an intensity threshold. Second, we would like the registration to be based solely on hepatic vessels as these are expected to move in an "approximate" rigid body relationship with the liver parenchyma, whereas a number of other vessels in the abdomen (e.g. the aorta) do not. By only matching on the hepatic vessels our hypothesis is that we should reduce the effect of errors induced due to liver movement between MR and US image acquisition. This segmentation does not need to be exact and can be thought of as defining a region of interest within the MR volume which contains voxels useful to our registration process and omits voxels which may induce errors due to non-rigid motion.

A further manual segmentation was then carried out on the set of training data to create an image which contained only hepatic vessels. The MR probability density function p_{MR} (which is essentially a look-up table converting intensities I_{MR} into probabilities p_{MR}) was calculated from these images as

$$p_{\text{MR}}(I) = \frac{\text{Number of vessel voxels of intensity } I}{\text{Number of liver voxels of intensity } I}. \quad (2)$$

3.2.2. Converting the ultrasound images

The first step for the ultrasound images is an artifact removal stage. Our method uses simple knowledge of the ultrasound image formation process – that most strong reflections and artifacts cause a loss of signal along the direction of the beam. We have calibrated our ultrasound images so that we know the beam directions. This is achieved by picking three points along the transducer face in one ultrasound image as shown in Fig. 3. Once these values are known, the US image (and all subsequent US images acquired on the same US machine with the same depth setting) can be described by a set of lines (beams) going through $\oplus$ at angles between $\pm\theta$.

The US images are Gaussian blurred (SD of four pixels) to reduce the effect of noise. Then, for each beam direction, we begin at the bottom of the image and move back towards the transducer face. The image is labelled as artifact until a threshold T_{art} value is reached. Approximately 2 cm of the US image directly adjacent to the transducer face is also labelled as artifact, this is to reduce the effect of deformation caused by probe pressure. Image positions labelled as artifact are not used in any subsequent image processing. In this and subsequent sections I_{US} will refer to US images which have been put through this simple artifact removal process.

After artifact removal, two features from the US images are used to produce the probability images. The first is image intensity. The second is the size of intensity dips along the beam directions. Intensity dips were chosen as a feature because, from visual inspection of the US images, they were seen to be a characteristic property of hepatic vessels. A dip image $I_{\text{dip}}(x)$ is generated from each US image using the following method. The algorithm searches for vessels between a minimum $v_{\min}$ and maximum $v_{\max}$ diameter. For a given vessel width v the algorithm scans along the direction of the beam and calculates the mean image intensity $\bar{c}$ around position x' (in Fig. 4) up to a distance of $\pm v/2$. The mean intensities on either side of the vessel along the direction of the beam ($\bar{a}$ and $\bar{b}$) are also calculated, see Fig. 4.

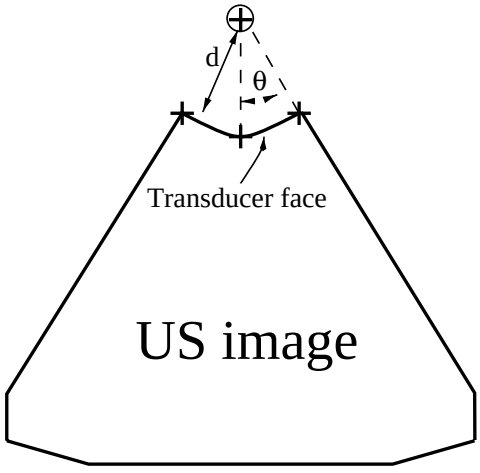

Fig. 3. Calculation of beam directions. Three points are manually picked on the transducer face +s. These points are used to calculate the centre of the arc transcribed by the transducer face $\oplus$, the radius of this arc d and a maximum angle θ.

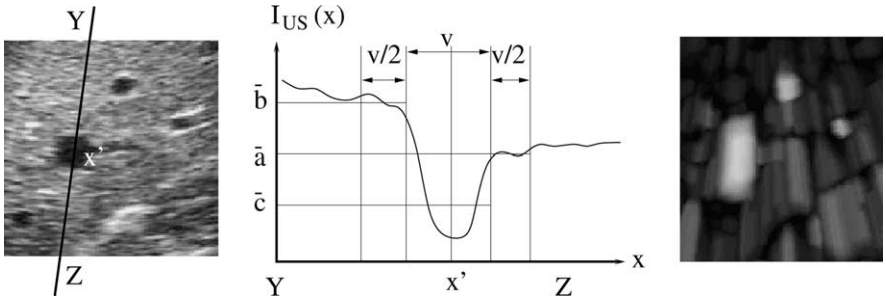

Fig. 4. Formation of the dip image. Left is ultrasound image after artifacts have been removed $I_{US}(x)$. Centre shows intensity profile along line YZ and the regions where the central $\bar{c}$, before $\bar{b}$ and after $\bar{a}$ mean intensity values are calculated. Right shows resultant dip image $I_{dip}(x)$.

Initially all the image values in $I_{dip}(x)$ are set equal to zero. Then values of $I_{dip}(x)$ along the ray YZ up to a distance $\pm v/2$ from x' (i.e. at positions $x'' = \{x' - v/2, \ldots, x' + v/2\}$) are set equal to $d = \min\{\bar{a} - \bar{c}, \bar{b} - \bar{c}\}$ if $\bar{b} > \bar{c} < \bar{a}$ and d is larger than the current value of I_{dip}. Also, at each position, x', the vessel size is varied between v_{min} and v_{max}.

This process is summarised as the following steps.

1. Set all image values in $I_{dip}(x)$ equal to zero.
2. For each position along the beam x', using vessel diameter v, calculate $\bar{a}$, $\bar{b}$ and $\bar{c}$.
3. Calculate the dip size $d = \min\{\bar{a} - \bar{c}, \bar{b} - \bar{c}\}$.
4. For each position $x'' = \{x' - v/2, \ldots, x' + v/2\}$ if $\bar{b} > \bar{c} < \bar{a}$ and $d > I_{dip}(x'')$ set $I_{dip}(x'') = d$.
5. Repeat steps 2–4 for all points along the beam and for all v between v_{min} and v_{max}.
6. Repeat steps 2–5 for each beam.

A set of 40 US images of the liver was used as training data. These images were acquired from both patients and volunteers. There was no overlap between the images used as training data and the images used in subsequent experiments. The images were manually segmented to give a set of images which only contained the hepatic vessels and the US probability density function p_{US} (which is essentially a 2D look-up table which converts I and I_{dip} into p_{US}) was calculated as shown in Eq. (3).

Fig. 5 shows a sample US image at a number of stages from the original image to a probability image.

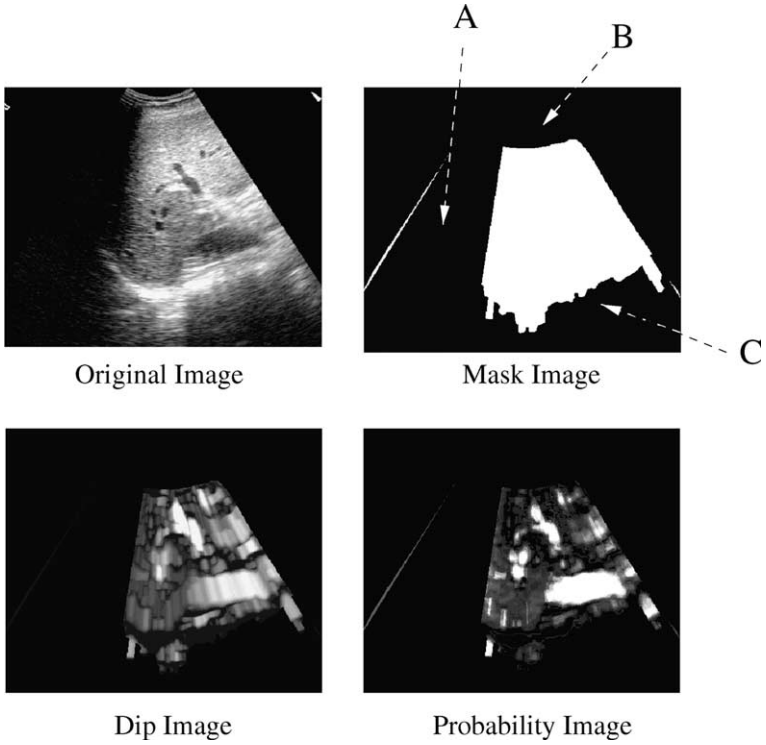

Fig. 5. Four images showing the process from the original US to the probability image. The mask image shows regions masked out due to a signal loss caused by poor probe contact (A), the removal of the area close to transducer face to reduce the effect of deformation caused by the US probe (B) and shadowing (C). Some non-vessel pixels can be seen to have been allocated a high value in the probability image. However, we would like to reiterate that our aim is not to produce a perfect segmentation, but to produce probability images of sufficient quality to allow accurate and robust registrations.

$$p_{\mathrm{US}}(I, I_{\mathrm{dip}}) = \frac{\text{Number of vessel pixels with intensity } I \text{ and dip intensity } I_{\mathrm{dip}}}{\text{Number of US pixels with intensity } I \text{ and dip intensity } I_{\mathrm{dip}}}. \quad (3)$$

3.3. The registration algorithm

Initially the algorithm Gaussian blurs (SD 2 mm) the probability MR image P_{MR} and the probability US slices P_{US}.

The similarity measure for the algorithm is calculated as follows. The current estimate of the rigid body parameters is used to reslice the P_{MR} image in the plane of each US slice. The pixel values in these reformatted slices and in the P_{US} images are then compared using the normalised cross-correlation similarity measure.

We use the following method to optimise the similarity measure. Each of the rigid body parameters is altered in turn and the value of the similarity measure calculated. The algorithm then moves the current estimate of the solution in the direction which produced the greatest improvement in the similarity measure. This optimisation scheme can be greatly affected by the choice of initial stepsize. An excessively large initial stepsize can result in an algorithm stepping outside its capture range. If the initial stepsize is too small then the optimisation may become trapped in a local optimum (Studholme et al., 1996). For all the experiments carried out for this paper, an initial stepsize of 2 mm or 2° was used, which was reduced by a factor of 2 when movements in all of the rigid body degrees of freedom produced a decrease in the value of the similarity measure. The minimum stepsize used was 0.0625 mm or 2°. This search strategy is a zeroth-order estimation scheme. A good review of alternative optimisation methods can be found in (Press et al., 1992).

3.4. Experiments

3.4.1. Data acquisition

Sets of between 15 and 28 ultrasound slices were collected by an interventionist from 5 volunteers, using a Siemens SONOLINE Versa Pro ultrasound machine and a 3.5-MHz probe. Images were captured using a video frame grabber (Snapper card, DataCell Ltd.). Ten images were manually selected from each set of images. The selection criterion was to choose images which contained a large number of clearly visible vessels and that the set of images included a wide range of views, enabling good coverage of the internal structure of the liver. These images were filtered to remove artifacts ($T_{\mathrm{art}} = 40$) and then converted into vessel probability images (max and min vessel sizes were $v_{\max} = 16$ mm, $v_{\min} = 2$ mm). The positions of the 10 slices from volunteer C can be seen in Fig. 6.

MR images were also acquired from the same volunteers on a Philips GYROSCAN ACS-2. The volunteers were asked to hold their breath at maximum exhale while the images were acquired. A typical image size was $256 \times 256 \times 28$ with voxel dimensions of $1.328 \times 1.328 \times 10$ mm. The liver was segmented from each of these volumes. A manual segmentation process was used, though it should be possible to use a semi (Schenk et al., 2000) or fully (Soler et al., 2000) automatic process to carry out this segmentation. Voxels which were not in the liver were labelled as missing data and took no further part in subsequent image processing. These images were then converted into probability images.

3.4.2. Calculation of a "bronze standard" transformation

All the registrations have been compared to a "bronze standard" transformation. This has been calculated using a modified version (Penney et al., 2001) of the iterative closest point (ICP) algorithm (Besl and McKay, 1992).

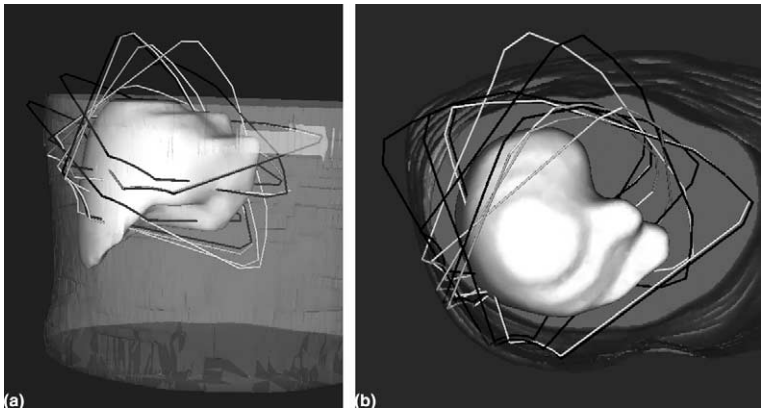

Fig. 6. Anterior–posterior (left) and cranial–caudal (right) views showing the relative positions of the 10 US images from volunteer C as black and white outlines (2 colours have been used for ease of visualisation). Also shown are renderings of the liver and skin surface from MR. The US outlines and MR renderings have been aligned using the "bronze standard" registration.

The features on which the registrations were based were the surface of the inferior vena cava (IVC) and the midlines of all the other visible hepatic vessels segmented from MR and a set of points manually picked in the US images. The graphical user interface (GUI) shown in Fig. 7 was used to allow the operator to pick points in each US image and draw lines to represent the midline of the hepatic vessels in the MR volume. The surface of the IVC was manually segmented using ANALYZE (Mayo Foundation) and loaded into the GUI. The interface allowed the user to run the modified ICP registration algorithm after picking a few points to allow an initial registration, the MR volume could then be resliced to show an approximately corresponding slice. This feature was found to help feature identification in the two images, allowing points and lines to be marked on peripheral vessels. Successive registrations could also be carried out to further improve feature identification. To aid point correspondence, lines and points were labelled as either IVC, portal vein or hepatic vein when they were picked.

All available ultrasound images were used to calculate the "bronze standard" and the aim was to use as many image features as possible. The point and line picking procedure typically took about 2 h to complete and so is impractical to use on a routine clinical basis. The final residual errors from the registrations were between 2.7 and 3.5 mm.

3.4.3. Accuracy experiments

A starting position was found using the following method. The OI image was used to find a starting estimate for the rotational parameters, i.e. the OI image was assumed to lie parallel to an axial slice in the MR volume. The translational parameters were found by manually identifying a single point in the MR and in one US image. This point was chosen to be the midline of the IVC where the hepatic veins join. Optimisation then used this position as a starting estimate.

The final rigid body parameters were compared to the "bronze standard" registration in the following way. Each voxel within the liver was transformed using both the "bronze standard" registration and the final registration parameters and then the root mean square (RMS) distance between these positions was calculated to give an RMS target registration error (TRE) (Fitzpatrick et al., 1998).

For comparison purposes, registrations were also carried out between the MR greyscale images (with non-liver voxels removed) and the greyscale US images (after the artifact removal stage) using a normalised mutual information (NMI) similarity measure (Studholme et al., 1999).

3.4.4. Precision and robustness experiments

To assess registration precision and robustness, the "bronze standard" registration was perturbed by adding random Gaussian noise to each of the rigid body parameters (SD 5 mm or 5°) to produce 50 starting positions.

Failed registrations were defined using the following method. The 50 registrations to each dataset were used to produce an average transformation (Moakher, 2002). The RMS distance (calculated over all the liver voxels) between the average and each individual transformation was calculated and a histogram plotted, as shown in Fig. 8. In three data sets there were no outliers. For the other two data sets, a threshold was chosen from visual inspection of the histograms and registrations which

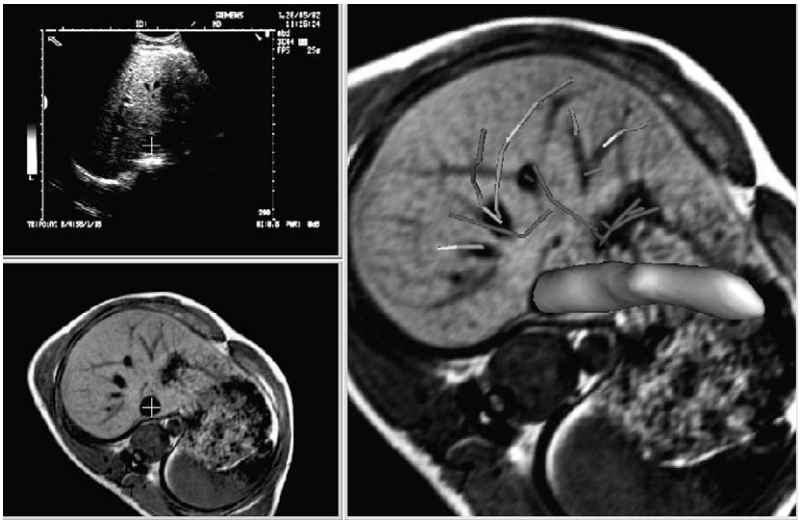

Fig. 7. Graphical user interface used to pick points and lines to calculate "bronze standard" transformation. Top left shows an ultrasound slice on which vessel points are picked, bottom left shows the current best estimate for the corresponding slice through the MR volume. Right shows the 3D image features which have already been picked and the current best estimate for the MR slice. 3D line features can be manually picked in this window.

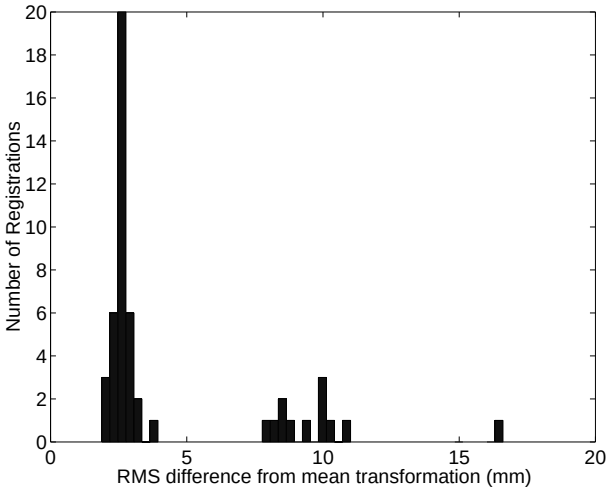

Fig. 8. Histogram plot showing the distribution of the RMS distance between a mean transformation and each of the 50 individual registration positions for volunteer E. A threshold of 4 mm was chosen, above which registrations were classified as a failure.

were further away from the mean transformation than this threshold value were classified as failures.

Two error measures were calculated from the successful results. The first was an RMS distance to the "bronze standard" registration calculated as explained above: this gives an RMS TRE value. The second was an RMS distance to an average transformation: this gives information on the precision or SD of the results.

4. Results

Registration accuracy results are given in Table 1. They show that on average the starting position, found using the OI image and a single point, was within 15.4 mm of the "bronze standard" registration and this improved to 3.6 mm after running the registration algorithm. The algorithm did not converge towards the "bronze standard" when the NMI similarity measure and greyscale images were used.

Fig. 9 shows three US slices and the corresponding slices from the MR volume after it has been resliced and reformatted using the final registration matrix. Corresponding features can clearly be seen to be well aligned in the respective images indicating an accurate registration. A number of image artifacts – in particular void signal due to poor contact between part of the transducer face and the volunteer – can be seen in the images. A linked cursor (white cross) shows corresponding positions in each image pair.

Table 1
Registration accuracy

Volunteer	Initial RMS TRE	Final RMS TRE	RMS TRE result using NMI
A	22.5 mm	4.3 mm	24.5 mm
B	17.7 mm	2.9 mm	23.2 mm
C	16.2 mm	2.9 mm	22.8 mm
D	9.8 mm	5.5 mm	39.6 mm
E	10.7 mm	2.3 mm	14.1 mm
RMS value	15.4 mm	3.6 mm	24.8 mm

Table 2
Precision and robustness experiments

Volunteer	Initial RMS TRE	Final RMS TRE	SD	Failure rate
A	11.3 mm	3.6 mm	0.5 mm	0%
B	11.3 mm	3.1 mm	0.8 mm	4%
C	11.1 mm	2.8 mm	0.9 mm	0%
D	11.6 mm	5.6 mm	0.6 mm	0%
E	10.7 mm	2.9 mm	1.1 mm	24%
RMS value	11.2 mm	3.6 mm	0.8 mm	5.6% (mean)

Table 2 shows the precision and robustness results. The results in Table 2 show the algorithm to be precise (maximum SD 1.1 mm) and robust for 4 out of the 5 data sets (maximum failure rate of 4%) for registrations which start within $\pm 5^\circ$ or 5 mm of the "bronze standard". The failure rate for data set E was much larger (24%), this was due to the presence of a local minimum which captured all except two of the failed registrations.

5. Discussion

We have presented an algorithm which can register a set of tracked 2D ultrasound slices to an MR volume. Our method differs from previous voxel-based registration methods as only a sparse set of US slices is used. The greyscale intensity images are transformed into probability images where the intensity values represent the probability of a voxel or pixel containing a vessel. We have shown that a method based on NMI and using the original intensity values, which has been highly successful in other registration applications, was not able to register the images used in this paper. We developed a strategy in which images are pre-processed to produce maps of probabilities of corresponding features; in our case vascular structures. Registration is then undertaken using these image maps.

The probability density functions used in this paper were calculated using a set of training data. One drawback to this method is that subsequent images must be similar to the set of training data. This would require a standard protocol to be used for MR. Our method should also be applicable for use with contrast enhanced CT images. In this case, as CT image values represent linear attenuation coefficients, there should be good

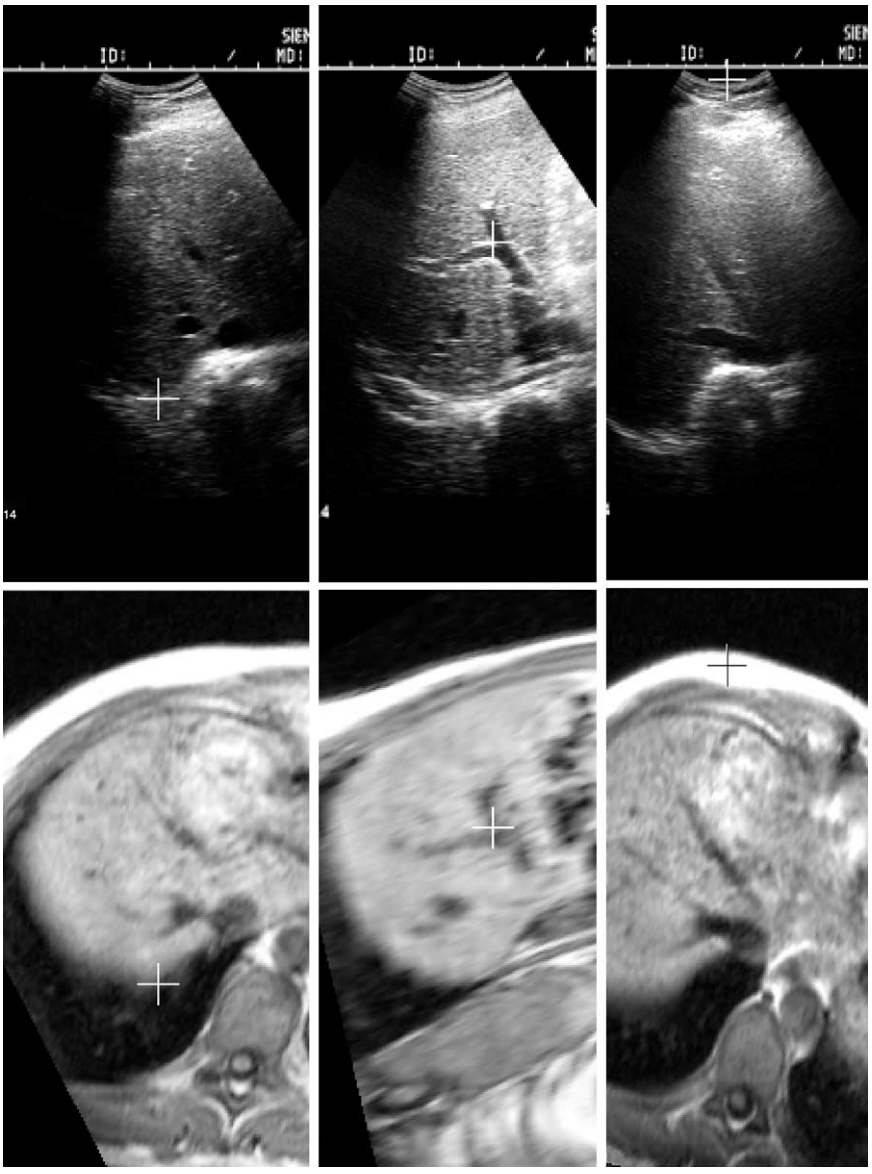

Fig. 9. Corresponding US and MR slices from volunteer C.

consistency between calculated probability density functions and subsequent sets of data. The US image acquisition can also be affected by a number of operator defined variables, these should be kept as similar as possible to the settings used when the training data were acquired. One solution to this problem would be to collect the RF data from the ultrasound machine. These data have not yet been converted into an image and so are unaffected by a number of operator variables. Our method to convert US images to probability images has been designed to work along the line of the US beam and so it is directly applicable for use with the RF data.

Our probability density functions are formed from relatively simple features, image intensity and (for US) intensity of a dip image. These features have proved to be sufficient to allow accurate and robust registrations on the images used in this paper. However, in the future, other image processing steps may be found which can extract image features that are more able to clearly delineate the differences between hepatic vessels and liver parenchyma. Such features could either be used in conjunction with our current feature sets, or could be used as a replacement. The introduction of such features should improve the quality of the probability images which in turn should improve the accuracy, robustness and capture range of the algorithm.

Our registration method assumes that the liver moves as a rigid body, whereas the liver is known to deform, both due to pressure from the US probe and the breathing motion of the patient. Deformation from the US probe is not expected to extend a long distance into

the patient (Voirin et al., 2002) and so we account for this by labelling the first few centimeters of the US image as missing data. Deformation from patient breathing has been kept to a minimum by gating both the US and MR acquisition at the same position in the breathing cycle – maximum exhale. Although the liver moves by a number of cm due to respiration, the non-rigid component has been measured to be much smaller than this (6 mm) (Rohlfing et al., 2001). Methods have also been proposed to account for this deformation using a patient specific statistical shape model (King et al., 2001).

6. Conclusions

We have presented a method to register a sparse set of 10 tracked US images to an MR volume of a liver. Our method is based on a strategy in which we pre-process images to provide a probability map of corresponding features. This pre-processing step uses information gathered from a set of manually segmented training data. Our method has been compared to a "bronze standard" registration calculated by manually picking points in both modalities. Results show that our method is accurate to within an RMS error of between 2.3 and 5.5 mm (with respect to a "bronze standard" registration) which is accurate enough for most liver procedures.

We are now integrating this method with a model of liver motion and deformation over the breathing cycle so as to track the shape and position of the liver during free breathing. The results of this work will form the basis of an image guidance system that will use the intraoperative ultrasound augmented with planning information from preoperative contrast enhanced CT (or MR) to guide needle placement in the ablation of liver metastases.

Acknowledgements

We thank the EPSRC (Grant No. GR/M53752) for funding this project and our volunteers for their patience.

References

Aylward, S.R., Jomier, J., Guyon, J.P., Weeks, S., 2002. Intraoperative 3D ultrasound augmentation. In: Proceedings of the IEEE International Symposium on Biomedical Imaging, pp. 421–424.

Balter, J.M., Lam, K.L., McGinn, C.J., Lawrence, T.S., Ten Haken, R.K., 1998. Improvement of CT-based treatment-planning models of abdominal targets using static exhale imaging. Int. J. Radiat. Oncol. Biol. Phys. 41 (4), 939–943.

Besl, P.J., McKay, N.D., 1992. A method for registration of 3-D shapes. IEEE Trans. Pattern Anal. Machine Intell. 14 (2), 239–256.

Blackall, J.M., Rueckert, D., Maurer Jr., C.R., Penney, G.P., Hill, D.L.G., Hawkes, D.J., 2000. An image registration approach to automated calibration for freehand 3D ultrasound. In: Medical Imaging Computing and Computer-Assisted Intervention – MICCAI '00, pp. 462–471.

Brown, L.G., 1992. A survey of image registration techniques. ACM Comput. Surv. 24, 325–375.

Bucholz, R.D., Yeh, D.D., Trobaugh, J., McDurmont, L.L., Sturm, C.D., Baumann, C., Henderson, J.M., Levy, A., Kessman, P., 1997. The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device. In: Proceedings of CVRMed-MRCAS, pp. 459–466.

Comeau, R.M., Sadikot, A.F., Fenster, A., Peters, T.M., 2000. Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. Med. Phys. 27 (4), 787–800.

Fenster, A., Downey, D.B., Cardinal, H.N., 2001. Three-dimensional ultrasound imaging. Phys. Med. Biol. 46 (5), R67–R99.

Fitzpatrick, J.M., West, J.B., Maurer Jr., C.R., 1998. Predicting error in rigid-body point-based registration. IEEE Trans. Med. Imag. 17 (5), 694–702.

Hajnal, J.V., Hill, D.L.G., Hawkes, D.J. (Eds.), 2001. Medical Image Registration. CRC Press, Boca Raton, Chapter 12, pp. 253–278.

Jannin, P., Fitzpatrick, J.M., Hawkes, D.J., Pennec, X., Shahidi, R., Vannier, M.W., 2002. Validation of medical image processing in image-guided therapy. IEEE Trans. Med. Imag. 21 (12), 1445–1449.

Jödicke, A., Deinsberger, W., Erbe, H., Kriete, A., Böker, D.K., 1998. Intraoperative three dimensional ultrasonography: an approach to register brain shift using multidimensional image processing. Minim. Invas. Neurosurg. 41, 13–19.

King, A.P., Blackall, J.M., Penney, G.P., Edwards, P.J., Hill, D.L.G., Hawkes, D.J., 2000. Bayesian estimation of intra-operative deformation for image-guided surgery using 3-D ultrasound. In: Medical Imaging Computing and Computer-Assisted Intervention – MICCAI '00, pp. 588–597.

King, A.P., Blackall, J.M., Penney, G.P., Hawkes, D.J., 2001. Tracking liver motion using 3-D ultrasound and a surface based statistical shape model. In: Proceedings of Mathematical Methods in Biomedical Image Analysis, pp. 145–152.

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Seutens, P., 1997. Multimodality image registration by maximization of mutual, information. IEEE Trans. Med. Imag. 16 (2), 187–198.

Maintz, J.B.A., Viergever, M.A., 1998. A survey of medical image registration. Med. Image Anal. 2 (1), 1–36.

Moakher, M., 2002. Means and averaging in the group of rotations. SIAM J. Matrix Anal. Appl. 24 (1), 1–16.

Penney, G.P., Edwards, P.J., King, A.P., Blackall, J.M., Batchelor, P.G., Hawkes, D.J., 2001. A stochastic iterative closest point algorithm (stochastICP). In: Medical Imaging Computing and Computer-Assisted Intervention – MICCAI '01, pp. 762–769.

Porter, B.C., Rubens, D.J., Strang, J.G., Smith, J., Totterman, S., Parker, K.J., 2001. Three-dimensional registration and fusion of ultrasound and MRI using major vessels as fiducial markers. IEEE Trans. Med. Imag. 20 (4), 354–359.

Prager, R.W., Rohling, R.N., Gee, A.H., Berman, L., 1998. Rapid calibration for 3-D freehand ultrasound. Ultrasound Med. Biol. 24 (6), 855–869.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1992. Numerical Recipes in C, second ed. Cambridge University Press, Cambridge.

Roche, A., Malandain, G., Pennec, X., Ayache, N., 1998. The correlation ratio as a new similarity measure for multimodal image registration. In: Medical Imaging Computing and Computer-Assisted Intervention – MICCAI '98, pp. 1115–1124.

Roche, A., Pennec, X., Rudolph, M., Auer, D.P., Malandain, G., Ourselin, S., Auer, L.M., Ayache, N., 2000. Generalized correlation ration for rigid registration of 3D ultrasound with MR images. In: Medical Imaging Computing and Computer-Assisted Intervention – MICCAI '00, pp. 567–577.

Roche, A., Pennec, X., Malandain, G., Ayache, N., 2001. Rigid registration of 3-D ultrasound with MR images: a new approach combining intensity and gradient information. IEEE Trans. Med. Imag. 20 (10), 1038–1049.

Rohlfing, T., Maurer Jr., C.R., O'Dell Walter, G., Zhong, J., 2001. Modeling liver motion and deformation during the respiratory cycle using intensity-based free-form registration of gated MR images. In: Proceedings of SPIE Medical Imaging 2001: Image Processing, pp. 337–348.

Schenk, A., Prause, G., Peitgen, H.-O., 2000. Efficient semiautomatic segmentation of 3D objects in medical images. In: Medical Imaging Computing and Computer-Assisted Intervention – MICCAI '00, pp. 186–195.

Slomka, P.J., Mandel, J., Downey, D., Fenster, A., 2001. Evaluation of voxel-based registration of 3-D power doppler ultrasound and 3-D magnetic resonance angiographic images of carotid arteries. Ultrasound Med. Biol. 27 (7), 945–955.

Soler, L., Delingette, H., Malandain, G., Montagnat, J., Ayache, N., Clément, J.-M., Koehl, C., Dourthe, O., Mutter, D., Marescaux, J., 2000. A fully automatic anatomical, pathological and functional segmentation from CT-scans for hepatic surgery. In: Medical Imaging 2000: Image Processing, pp. 246–255.

Studholme, C., Hill, D.L.G., Hawkes, D.J., 1996. Automated 3D registration of MR and CT images of the head. Med. Image Anal. 1 (2), 163–175.

Studholme, C., Hill, D.L.G., Hawkes, D.J., 1997. Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. Med. Phys. 24 (1), 25–35.

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

Viola, P., Wells, W.M., 1995. Alignment by maximization of mutual information. In: Proceedings of the 5th International Conference on Computer Vision (ICCV'95). Boston, USA, pp. 16–23.

Voirin, D., Payan, Y., Amavizca, M., Létoublon, C., Troccaz, J., 2002. Computer-aided hepatic tumour ablation: requirements and preliminary results. C.R. Biol. 325, 309–319.

Weese, J., Röche, P., Netsch, T., Blaffert, T., Quist, M., 1999. Gray-value based registration of CT and MR images by maximization of local correlation. In: Medical Imaging Computing and Computer-Assisted Intervention – MICCAI '99, pp. 656–663.