

# Fast rigid registration of pre-operative magnetic resonance images to intra-operative ultrasound for neurosurgery based on high confidence gradient orientations

Dante De Nigris · D. Louis Collins · Tal Arbel

Received: 8 November 2012 / Accepted: 1 March 2013 / Published online: 21 March 2013  
© CARS 2013

## Abstract

**Purpose** We present a novel approach for the registration of pre-operative magnetic resonance images to intra-operative ultrasound images for the context of image-guided neurosurgery.

**Method** Our technique relies on the maximization of gradient orientation alignment in a reduced set of high confidence locations of interest and allows for fast, accurate, and robust registration. Performance is compared with multiple state-of-the-art techniques including conventional intensity-based multi-modal registration strategies, as well as other context-specific approaches. All methods were evaluated on fourteen clinical neurosurgical cases with brain tumors, including low-grade and high-grade gliomas, from the publicly available MNI BITE dataset. Registration accuracy of each method is evaluated as the mean distance between homologous landmarks identified by two or three experts. We provide an analysis of the landmarks used and expose some of the limitations in validation brought forward by expert disagreement and uncertainty in identifying corresponding points.

**Results** The proposed approach yields a mean error of 2.57 mm across all cases (the smallest among all evaluated techniques). Additionally, it is the only evaluated technique

that resolves all cases with a mean distance of less than 1 mm larger than the theoretical minimal mean distance when using a rigid transformation.

**Conclusion** Finally, our proposed method provides reduced processing times with an average registration time of 0.76 s in a GPU-based implementation, thereby facilitating its integration into the operating room.

**Keywords** Neurosurgery · Multi-modal registration · GPU · Ultrasound

## Introduction

A common setup for image-guided neurosurgery (IGNS) is one that allows for the visualization of a tracked pointer onto a pre-operative 3D magnetic resonance (MR) image in order to expose internal regions of interest (e.g., tumors) and use the pre-operative image for guidance during the procedure. However, brain movements during open-skull operations as well as inaccurate image-to-patient registrations are known to reduce the utility of using the images for guidance. Intra-operative imaging modalities have thus been proposed for improved guidance. Nonetheless, some modalities, such as intra-operative MR, involve a prohibitive cost, as well as requiring major modifications to the operating room and surgical procedure. For such reasons, intra-operative ultrasound (iUS) continues to be a highly appealing option given its ease of use and relatively inexpensive costs. The iUS is typically tracked and superimposed on the pre-operative MR with the use of a navigation system, allowing the surgeon to assess a misalignment or deformation with relation the MR. Unfortunately, the comparatively poor image quality of iUS allows for the exposure of only a limited (yet important) set of coarse anatomical structures such as the brain tumor,

---

D. De Nigris (✉) · T. Arbel  
Centre for Intelligent Machines, McGill University,  
Montreal, QC H3A 0E9, Canada  
e-mail: dante@cim.mcgill.ca

T. Arbel  
e-mail: arbel@cim.mcgill.ca

D. L. Collins  
Montreal Neurological Institute, McConnell Brain Imaging Centre,  
McGill University, Room WB-314, 3801 University Street,  
Montreal, QC H3A 2B4, Canada  
e-mail: louis.collins@mcgill.ca

the lateral ventricle boundaries, and the falx. For example, the accuracy of iUS for identifying brain tumor boundaries such as gliomas and metastases has been well documented in previous work [23]. However, there are many other important anatomical references that are left undepicted, thus limiting the direct use of iUS for guidance purposes.

The approach we follow in this work consists of matching a tracked iUS to a pre-operative MR (pMR) image, thereby permitting the updating of the pMR during surgery. The clinician can therefore benefit from the soft-tissue imaging detail found in the MR volume, as well as any other preoperative images, while taking advantage of the practical ease of US acquisition and the physical reality represented by the iUS. The critical challenge thus lies in quickly and accurately registering the MR image to an iUS image, thereby minimizing the time required to obtain an updated MR image and allowing for seamless integration in the operating room.


Registration of MR to iUS is a problem faced with multiple challenges brought forward by the widely different image formation models of each modality. Magnetic Resonance Imaging can be largely characterized as a tissue-type based modality, where the image intensity of a given voxel is mainly a function of the tissue found within the voxel's volume. On the other hand, US images illustrate different acoustic impedance transitions encountered by the ultrasonic wave. Figure 1a–d shows four different neurosurgical cases and illustrates the differences between the two modalities. Notice that the MR image allows for an accurate identification of multiple soft-tissue types such as gray matter, white matter, and bone and hence also permits the localization of many anatomical structures on the brain (e.g., lateral ventricles, falx, sulci, gray and white matter surface, etc.). Alternatively, the corresponding US image exposes the tumor tissue and its boundary (with some degree of uncertainty) and also depicts part of some key structures like the falx and the lateral ventricles.

In the particular context of brain tumors, we are also affected by the fact that pathologies from different cases can have quite unique image features in each modality. Figure 1a–c contrasts three cases with significantly different depictions in each modality. In particular, we can observe that the inner tumor tissue and the tumor boundaries in the iUS are exposed with quite distinctive image characteristics in each case. For example, Fig. 1c depicts tumor tissue with very high US intensity values, but does not allow for an accurate identification of its boundary. Alternatively, in Fig. 1b, the tumor tissue is depicted with a low US intensity and the boundary can be identified with increased certainty. Furthermore, Fig. 1b, c provides a prominent depiction of the lateral ventricles in the iUS, a highly informative anatomical structure for identifying a match across modalities. On the other hand, Fig. 1a provides a very weak depiction of the boundaries of the lateral ventricles. This

kind of variability in exposing anatomical structures in iUS is a clear challenge for conventional multi-modal registration approaches that assume a global hard mapping between image features of one modality to image features of a second modality.

The registration of MR to US images has been previously addressed by various groups [1, 10, 12, 16, 25]. Some approaches [3, 20] rely on gradient magnitude as an image feature of interest in conjunction with a conventional multi-modal similarity metric such as mutual information (MI) or normalized cross-correlation (NCC). Other proposed techniques rely on local-phase [14, 25] as a feature in conjunction with MI. A major challenge encountered by such approaches is that the image intensity response found in US is significantly non-homogeneous. Consider the coronal view found in Fig. 1d where we can clearly observe how the US intensity decays with relation to the distance from the probe. In particular, notice that the US pixel intensities corresponding to white matter are far from consistent and will likely result in a degradation of registration performance when using a similarity metric that involves the full image domain. There have also been approaches [10, 12] that propose a preprocessing stage in which imaging artifacts (e.g., speckle, noise) are reduced and consequently register the preprocessed images with a multi-modal metric like normalized mutual information (NMI) evaluated over a sampling mask that typically covers the tumor volume and part of its surrounding region. Previously, our group has proposed [1, 16] the generation of a pseudo-US from segmented structures in the MR and then registering the pseudo-US with the acquired iUS. Similarly, in [24], the registration of US to CT is addressed through a simulated US obtained from the CT. This technique makes use of a hard mapping of the CT intensity to the tissue's echogeneity, which relies on a functional relationship between those two values. Unfortunately, the strong presence of speckle and the wide variability across IGNS cases do not allow for a straightforward implementation of such strategy in this highly variable context.

In this paper, we propose a new MR to US registration framework, with the goal of providing substantially improved robustness and computational performance in the context of image-guided neurosurgery (IGNS). Registration is based on gradient orientation alignment, motivated by the fact that gradient orientations are considered to characterize the underlying anatomical boundaries found in the scene and are more robust to the effect of non-homogeneous intensity response found in US. However, orientation estimates can be noisy in this context and using the full set of orientations would lead to undue computational cost and increase the risk of errors. An important distinction with other approaches based on gradient orientations [2, 7, 9, 11, 14, 18, 19] is that we restrict the evaluation of the similarity metric to a small set of locations of interest with low uncertainty gradient orientations,



**Fig. 1** Pre-operative MR volume and the initial unregistered iUS volume in a brain tumor resection surgery, corresponding to cases 4, 6, 8, and 9. The first column shows the MR image in grayscale with the corresponding iUS overlapped and heat map colored. The second column shows the MR image. The third column shows the iUS. The *top*

*row* shows a coronal view, the *middle row* shows a sagittal view, and the *bottom row* shows a transverse view. Anatomical structures found in the iUS are identified by a *green arrow* and label, while structures found in the MR are identified by a *white arrow* and label. **a** Case 4, **b** Case 6, **c** Case 8, **d** Case 9

which allows for improved robustness and computational performance. In order to address such limitations, our technique [5,6] first selects locations whose gradient orientations have high certainty (i.e., minimal noise) and likely correspond to structures of interest. Once such locations have been

identified, we maximize their alignment with corresponding orientations. Our approach is thus asymmetric in the sense that the fixed image (US) plays a different role than the moving image (MR), allowing the registration task to focus on the alignment of boundaries that appear in the US and likely

have a counterpart in the MR image, and thereby improving registration robustness against the variability in iUS found throughout the cases evaluated.

Experimental findings show that our approach brings forward gains in computational performance and registration accuracy, as evaluated over fourteen clinical cases obtained from the publicly available MNI BITE dataset. In particular, we can achieve a robust performance, in the sense that all fourteen cases used for validation have a resulting mean distance between corresponding points that is larger than the smallest possible mean distance (under a rigid transformation) by no more than 1 mm. Furthermore, such performance is achieved with a highly reduced subset of voxels (e.g., 2% of the image) and a GPU-based implementation, which leads to an average processing time of 0.76 s. This achievement should permit the strategy to be easily embedded in the clinical IGNS system, minimizing the delay suffered every time an updated MR is required.

Validation of registration performance is measured with a significant number of clinical cases, which we argue provides a much more informative indication of real registration accuracy and robustness in comparison with controlled brain phantom-based setups. However, the lack of a gold standard does involve some limitations in the validation strategy. In particular, given the limited number of homologous landmarks, the uncertainty in identifying the precise location of such points, and the potential disagreement between experts, there is some ambiguity as to how accurate the performance metric is. In this work, we provide an analysis of the landmarks identified by each expert in order to demonstrate the degree of agreement between experts, the potential need for a non-rigid registration and to illustrate the challenges of an accurate registration validation based on real clinical cases. Given this ambiguity in measuring registration accuracy, we choose to evaluate the performance of a rigid registration only, and we demonstrate that even though a non-rigid registration may be required, a rigid registration can characterize most of the deformation encountered.

### Clinical data

We make use of fourteen clinical neurosurgical cases obtained from the Montreal Neurological Institute's Brain Images of Tumors for Evaluation (MNI BITE) [15], an open access online<sup>1</sup> dataset of clinical MR and US images of brain tumors. In particular, we evaluate the registration of pre-operative MR images to intra-operative US images obtained prior to tumor resection, identified as Group 2 of the MNI BITE dataset. The cases involve low- and high-grade gliomas (LGG and HGG, respectively), at different depths and locations in the brain and with tumor volumes ranging between

0.2 and 79.2 cm<sup>3</sup>. The initial location of each case corresponds to a preliminary registration involving the manual identification of corresponding points on the skin and the MR image, as is common in standard clinical procedures.

All pre-operative images used consist of T1-weighted gadolinium-enhanced MR images. All cases, except case 8, were acquired with an axial SPGR sequence. The pre-operative MR image of case 8 was acquired with a 3D MPRAGE acquisition. Two-dimensional US images were acquired on the dura with an HDI 5000 (ATL/Phillips, Bothell, WA, USA) machine using a P7-4 phased array transducer at a depth of 6.5 and 8 cm. The US probe was tracked with reflective spheres and the 2D images were captured using a Pinnacle PCTV frame-grabbing card. Each acquisition includes between 200 and 600 frames. A 3D US volume of voxel spacing of 1.0 × 1.0 × 1.0 mm is then reconstructed with a distance-weighted pixel-based method [22]. Note that the reconstructed volume used is different than the one found in the MNI BITE dataset. In particular, we adopt a coarser voxel spacing which reduces the number of holes found in the volume and tends to decrease the presence of noise. The processing time of reconstructing a 3D US volume currently represents a significant delay with an average processing time of 90 s, largely due to lack of software optimization. We note that an implementation of 3D US reconstruction performed in real-time was recently demonstrated in [4].

### Gradient orientation-based registration

In this section, we describe the proposed registration approach [6,5], characterized by three major components: (1) a local similarity metric based on gradient orientation alignment, (2) a multi-scale selection strategy that identifies locations of interest with gradient orientations of low uncertainty, and (3) a computationally efficient technique for evaluating gradient orientations of the transformed moving image. The registration pipeline consists of two stages. First, a pre-processing stage involving the computation of image derivatives of both volumes and the identification of locations with low uncertainty gradient orientations. The second stage consists of an optimization strategy which maximizes the average value of the local similarity metric evaluated on the locations of interest. Optimization is performed with a covariance matrix adaptation evolution strategy (CAE) [8], a non-gradient-based optimization strategy.

#### Local similarity measure

The similarity metric employed evaluates the alignment of corresponding gradient orientations as:

$$S(\nabla I_F, \nabla I_M) = \cos(\Delta\theta)^2 \quad (1)$$

<sup>1</sup> <http://www.bic.mni.mcgill.ca/Services/ServicesBITE>.

where  $\Delta\theta$  is the inner angle between corresponding image gradients,  $\nabla I_F$  and  $\nabla I_M$ . Note that the fixed image,  $I_F$ , is set to be the reconstructed 3D US volume, while the moving image,  $I_M$ , is set to be the pre-operative MR image.

In order to improve registration robustness and computational performance, we restrict the evaluation of the similarity metric to a small set of locations with low uncertainty gradient orientations. In the following section, we describe the principles behind the selection strategy employed.

### Gradient orientations from a noisy image

The main motivation behind the use of gradient orientations as an image feature of interest lies in the fact that they are directly related to the direction of anatomical boundaries. Thus, it is of critical importance to give greater importance to locations whose orientations correspond to tissue transitions and disregard all locations whose gradient orientations are brought forward by image artifacts or noise.

In this work, we adopt a selection strategy that identifies locations of interest whose corresponding gradient orientations are considered reliable or of high certainty, and we then restrict the evaluation of the local similarity metric to such locations. It is important to note that such locations are identified on the fixed image and there is no selection made on the moving image. This asymmetrical nature of this approach is very well-suited for this multi-modal context. In particular, we note that for the cases found in the MNI BITE dataset, most of the strong tissue-type transitions exposed by the fixed image (US) can also be found in the moving image (MR), while the converse is typically not true. For such purposes, we demonstrate an approach for obtaining a reliable indicator of the certainty of a given gradient orientation and consequently select a reduced set of locations with high certainty gradient orientations.

In this work, we consider the main source of uncertainty in gradient orientations to be additive Gaussian noise on voxel intensities from the reconstructed US volume. It is important to state that commonly used US noise models are typically more complex. However, they are generally defined with relation to the raw “envelope” (i.e., before log-compression) US intensity or on the log-compressed US intensity (i.e., US intensity displayed on screen). For example, a proposed noise model for log-compressed US intensity, commonly known as the Loupas model [13], characterizes the noise as additive Gaussian with a location-dependent variance proportional to the original (i.e., undegraded) intensity value. We highlight the fact that intensities found in the reconstructed US volume are obtained by a linear combination (i.e., weighted mean) of a large number of pixel intensities from log-compressed US slices and thus involve a different noise model from the one in log-compressed US.

Our simplifying assumption is to characterize such noise as an additive Gaussian model,

$$I[i] = F[i] + \varepsilon[i] \tag{2}$$

where  $F$  is the *undegraded* (i.e., noiseless) image,  $i$  is a voxel index, and  $\varepsilon[i]$  is an i.i.d. Gaussian random variable with variance  $\sigma^2$ .

The corresponding probability density of a voxel’s intensity is expressed as,

$$p(I[i] | F[i]) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(I[i] - F[i])^2}{2\sigma^2}\right) \tag{3}$$

If we also consider that the image gradient at a particular location is obtained by convolution with linear operators that act solely on the axis (dimension) of interest, we can derive the posterior probability of,  $\phi$ , the gradient orientation of the undegraded image,  $F$ , given  $m$  and  $\theta$  the observed gradient magnitude and orientation of  $I$ . The resulting expression is

$$p(\phi | m, \theta) = \frac{e^{-\frac{m^2 \sin^2 \Delta}{2\|\mathbf{K}\|^2 \sigma^2}} \cdot \Phi\left(\frac{m \cos \Delta}{\|\mathbf{K}\| \sigma}\right)}{\pi \sqrt{\sigma \|\mathbf{K}\|} \cdot I_0\left(\frac{m^2}{4\|\mathbf{K}\|^2 \sigma^2}\right) \cdot e^{-\frac{m^2}{\|\mathbf{K}\|^2 \sigma^2}}} \tag{4}$$

where  $\Delta = \phi - \theta$  and  $\Phi()$  is the cumulative distribution function of a Gaussian random variable and  $\mathbf{K}$  is the discrete kernel characterizing the derivative linear operator.

The posterior density mainly expresses a *unimodal directional density* whose variance is a monotonically decreasing function of  $\frac{m}{\|\mathbf{K}\|}$ . Thus, we obtain that  $\frac{m}{\|\mathbf{K}\|}$  represents a common indicator of the precision (i.e., inverse of variance) of a given gradient orientation and which we can use as a cue for identifying locations of interest.

In this work, we automatically identify a gradient orientation selection threshold which corresponds to the 80th percentile of gradient magnitude (i.e., the threshold above which the top 20% locations with highest gradient magnitude are found) for each particular case. To allow for reduced processing times, we also consider a random selection of a reduced number (e.g., 8,000) of voxels found within such mask.

### Evaluation of transformed moving image gradients

Instead of adopting a straightforward approach for evaluating the gradient of the transformed moving image, in which pixel intensities are evaluated first by interpolation and the image derivatives are consequently computed by convolution, we adopt a more computationally efficient approach [6] which simply maps the gradients computed (once) by convolution on the original moving image.

The linear mapping of gradients can be easily found by expanding the expression for the derivative of a transformed moving image. Consider a D-dimensional moving image,

$I_M$ , whose coordinate space is *inversely* mapped by a continuous transformation function,  $\mathbf{T}$ , to the fixed image coordinate space. In other words, a location,  $\mathbf{x} = (x_1, \dots, x_D)$ , in the fixed image coordinate space corresponds to location  $\mathbf{T}(\mathbf{x}) = (T_1, \dots, T_D)$  in the coordinate space of the original moving image. The derivative of the moving image with respect to a particular dimension,  $x_j$ , of the fixed image coordinate space is expressed as,

$$\left. \frac{\partial I_M}{\partial x_j} \right|_{\mathbf{T}(\mathbf{x})} = \sum_i^D \left. \frac{\partial I_M}{\partial T_i} \right|_{\mathbf{T}(\mathbf{x})} \cdot \left. \frac{\partial T_i}{\partial x_j} \right|_{\mathbf{x}} \quad (5)$$

where the term  $\left. \frac{\partial T_i}{\partial x_j} \right|_{\mathbf{x}}$  corresponds to the  $(i, j)$ -th component of the spatial Jacobian matrix of the transformation function.

$$J_{\mathbf{T}} = \begin{bmatrix} \frac{\partial T_1}{\partial x_1} & \dots & \frac{\partial T_1}{\partial x_D} \\ \vdots & \ddots & \vdots \\ \frac{\partial T_D}{\partial x_1} & \dots & \frac{\partial T_D}{\partial x_D} \end{bmatrix} \quad (6)$$

Re-arranging terms we obtain the expression for the transformed image gradient,

$$\nabla_{\mathbf{x}} I_M(\mathbf{T}(\mathbf{x})) = J_{\mathbf{T}}^T(\mathbf{x}) \cdot \nabla_{\mathbf{T}} I_M(\mathbf{T}(\mathbf{x})) \quad (7)$$

where  $\nabla_{\mathbf{T}} I_M(\mathbf{x}) = \left( \frac{\partial I_M(\mathbf{T}(\mathbf{x}))}{\partial T_1}, \dots, \frac{\partial I_M(\mathbf{T}(\mathbf{x}))}{\partial T_D} \right)$  is the gradient of the original moving image. The gradient of the transformed moving image is therefore evaluated as the product of the transpose of the spatial Jacobian matrix and the gradient of the original (undeformed) moving image.

In other words, the gradient of the transformed moving image can be evaluated as the product of the transpose of the spatial Jacobian matrix and the gradient of original (undeformed) image. Hence, the derivatives of both images are computed only once and the computational complexity of the method is significantly reduced.

It is of interest to note that in the case of a *rigid* transformation, the spatial Jacobian matrix is the same for all locations in the domain and simply reflects the rotational component of the transformation.

#### Alternative techniques for comparison

We evaluated the registration performance of a variety of intensity-based registration techniques with and without relevant pre-processing stages. The first set of experiments involve conventional multi-modal intensity-based approaches such as the maximization of normalized cross-correlation (NCC), the maximization of mutual information (MI), and the maximization of normalized mutual information (NMI). Since the optimization strategy can potentially play a critical role in the registration performance, we evaluated the registration with the same non-gradient-based optimization strategy employed in our technique, as well as a gradient descent (GD) strategy with an adaptive gain.

The second set of experiments involve the use of gradient magnitude images as input images instead of the original volumes, and employing the same intensity-based similarity metrics from the first set of experiments. Finally, we also evaluate the performance of an approach largely based on the registration pipeline proposed in [10]. This technique relies on a filtering pre-processing stage which aims to reduce the presence of speckle, noise, and other image artifacts in both modalities and thereby improve registration robustness. In particular, the US volume is first Gaussian blurred to reduce the effect of speckle. Then, high-intensity regions in the US are automatically identified as those whose intensity values lie above the Otsu threshold [17]. The mask is then dilated by a few voxels so as to allow for an increased registration aperture range. The MR image is similarly processed with a median filter, which tends to increase the intensity homogeneity in regions with a common tissue type. Finally, the processed images are registered by maximization of NMI.

#### Registration validation

Given the lack of a gold standard, we measure the registration accuracy of each method as the mean distance between homologous landmarks independently identified by two or three experts,<sup>2</sup> commonly referred as the mean target registration error (mTRE). Each case has between 19 and 40 landmarks in total. It is of critical importance to note that the minimal mTRE under a rigid transformation has a unique non-zero value for each case, both for each expert's landmarks as well as for the combined set of all experts' landmarks. There are two main reasons behind this phenomenon. The first is the inherent uncertainty from the experts in accurately identifying anatomical locations in both modalities (particularly in the US volume). Thus, large errors in the identification of landmarks reduce the accuracy of the performance metric (i.e., mTRE) and also result in a potentially false large value for the minimal mTRE under a rigid transformation. The second is the potential presence of non-rigid deformations. In particular, in the case of perfectly accurate landmarks, a non-zero minimal mTRE effectively quantifies the "residual" part of the deformation that is not fully explained by a rigid transformation. Hence, a large minimal mTRE can reflect the presence of errors in the identified landmarks and/or significant non-rigid components in the true deformation. For the purposes of illustrating the variability of landmark identification between experts and the potential need for a non-rigid registration, we provide a quantitative analysis of the chosen points.

<sup>2</sup> Note that the third expert in the MNI BITE dataset identified homologous landmarks for the first six cases only.

## Analysis of expert landmarks

For each expert's landmarks, we evaluate the initial mTRE of each case. We also report the mTRE evaluated with the optimal rigid transformation. The optimal rigid transformation given a set of landmarks is obtained analytically by solving the corresponding Orthogonal Procrustes problem [21]. Note that the mTRE evaluated with the optimal rigid transformation corresponds to the minimal mTRE that can be obtained with a rigid transformation given that set of landmark points. In order to analyze the agreement between different experts, we also evaluate the mTRE of a given expert's landmarks with relation to the optimal rigid transformation obtained with *another expert's* landmarks. Such measure is particularly informative in cases that do not seem to require a non-rigid transformation.

Table 1 lists the results from such analysis. For example, Table 1a evaluates the mTRE based on the landmarks identified by Expert 1. The column labeled *Initial* lists the mTRE values computed at the initial (unregistered) location of each case. The column labeled *Expert 1 Solution* lists the mTRE values computed with the optimal rigid transformation (analytically) obtained from Expert 1's landmarks. Hence, the column labeled *Expert 1 Solution* in Table 1a also lists the minimal possible mTRE. We identify all values of the minimal mTRE larger than 2 mm (highlighted in bold) as cases with potential need for a non-rigid model. Notice that between Expert 1 and 3, only Expert 1 yields one case with a minimal mTRE larger than 2 mm. In direct contrast, Expert 2 yields six cases with a minimal mTRE larger than 2 mm. We therefore conclude that for most of the evaluated cases and their corresponding landmarks, a rigid transformation can characterize most of the deformation encountered.

In Table 1, all mTRE values larger than 3 mm and evaluated with relation to a different expert's solution are identified as cases with significant disagreement between experts (highlighted in italic). The most striking case of disagreement is found in Case 1, where each expert has a corresponding minimal mTRE of 1.23, 1.23, and 1.35 mm. Thus, there seems to be no strong need for a non-rigid transformation. However, the rigid solution obtained with the landmarks from Expert 2 results in a mTRE value of 11.55 mm with Expert 1's landmarks. In contrast, the rigid solution obtained with the landmarks from Expert 3 results in a mTRE value of 2.11 mm with Expert 1's landmarks. Additionally, the rigid transformation from Expert 2 results in a mTRE of 10.69 mm when evaluated with the landmarks chosen by Expert 3. In contrast, the rigid transformation from Expert 1 results in a mTRE of 2.66 mm when evaluated with Expert 3's landmarks. Hence, we can state that Expert 1 and 3 seem to somewhat agree on the deformation encountered in Case 1, while having a strong disagreement with Expert 2.

The landmarks identified by Expert 1 and 2 for Case 1 are illustrated in Fig. 2 for further analysis. Notice the significant difference between the spatial distribution of each set and the difference between the apparent transformation for each set. In particular, the points identified by Expert 1 (shown in blue and white) are already quite close in distance, while the ones identified by Expert 2 (shown in green and yellow) are significantly farther apart. For reference purposes, in Fig. 3, we illustrate the landmarks identified by Expert 1 and 2 for Case 13, which exposes a less prominent disagreement between experts. In such case, the distributions of the two landmark sets are relatively similar when compared to the ones found in Case 1.

The relevance of the analysis on the experts' landmarks lies in highlighting the challenges involved in the validation of a registration method in real clinical cases with no gold standard. In particular, we feel it is important to underline that though the use of manually identified points allows for a quantitative evaluation of performance, there is still a significant degree of subjectivity behind such validation strategy and the numerical results should not be accepted blindly. In particular, based on the exposed variability between experts and corresponding landmarks, we state that the validation of a non-rigid registration based on this particular dataset is rather compromised. Hence, we choose to evaluate registration performance restricted to a rigid transformation.

## Results


Our proposed approach was evaluated with three configurations. The first configuration, referred to as *GOA Full Mask*, involves the maximization of gradient orientation alignment of the top 20% locations with the highest gradient magnitudes in the reconstructed 3D US volume. The second configuration, referred to as *GOA Subset*, provides reduced processing times and involves the maximization of gradient orientation alignment of 8,000 locations randomly selected from the previously defined top 20% mask. It is important to note that 8,000 locations correspond to approximately 2% of the voxel locations found in the US volume (the exact ratio varies from case to case). The first two configurations were implemented in C++ and run on a computer with an Intel Core 2 Quad Q6700 CPU. The third configuration, *GOA Subset on GPU*, is implemented to run on a NVIDIA GTX 670 video card and was developed to provide highly reduced processing times. It involves the maximization of gradient orientation alignment of 16,000 locations randomly selected from the top 20% mask. Figure 4 illustrates the images from Case 3 before and after registration with *GOA Subset on GPU*. Notice how key structures like the falx and the lateral ventricles are closely aligned after registration.

**Table 1** Analysis of homologous landmarks identified by each expert


mTRE based on Expert 1 landmarks		Expert 1 solution	Expert 2 solution	Expert 3 solution
Case	Initial	(Minimal)		
(a)				
1	3.62	1.23	<i>11.55</i>	2.11
2	6.29	1.38	1.86	1.47
3	8.79	1.63	1.98	2.25
4	4.07	1.10	1.58	1.48
5	3.20	<b>2.37</b>	2.51	2.47
6	2.89	0.96	2.11	1.93
7	3.70	1.71	2.52	–
8	4.70	0.89	2.86	–
9	4.82	1.89	2.71	–
10	3.83	1.26	2.15	–
11	1.90	1.39	1.72	–
12	4.89	1.67	<i>3.84</i>	–
13	7.30	1.45	<i>4.79</i>	–
14	4.23	1.09	2.58	–
Average	4.59	1.43	3.20	1.95
Median	4.15	1.38	2.52	2.02
mTRE based on Expert 2 landmarks		Expert 1 solution	Expert 2 solution	Expert 3 solution
Case	Initial		(Minimal)	
(b)				
1	7.20	<i>10.28</i>	1.23	<i>8.11</i>
2	5.93	2.29	1.91	2.27
3	9.02	<i>3.10</i>	<b>2.86</b>	3.92
4	3.68	2.23	1.61	1.88
5	1.96	2.12	1.67	2.21
6	1.86	2.45	1.04	2.16
7	2.31	<i>3.10</i>	1.74	–
8	3.02	<i>3.83</i>	<b>2.49</b>	–
9	5.33	<i>3.36</i>	<b>2.39</b>	–
10	2.43	2.48	1.71	–
11	1.26	1.53	1.22	–
12	2.61	<i>3.66</i>	<b>2.04</b>	–
13	3.49	<i>5.40</i>	<b>2.56</b>	–
14	3.45	<i>3.62</i>	<b>2.57</b>	–
Average	3.82	3.53	1.93	3.42
Median	3.23	3.10	1.83	2.24
mTRE based on Expert 3 landmarks		Expert 1 solution	Expert 2 solution	Expert 3 solution
Case	Initial			(Minimal)
(c)				
1	2.39	2.66	<i>10.69</i>	1.35
2	6.87	1.10	1.47	0.88
3	10.69	2.49	<i>3.06</i>	1.50
4	4.10	2.11	1.87	1.70
5	2.76	1.79	1.86	1.30
6	2.45	2.16	1.95	1.10
Average	4.88	2.05	3.48	1.31
Median	3.43	2.13	1.91	1.33

Reported numbers include the initial mTRE (Initial) and the minimal mTRE found with the analytical rigid transformation obtained from each expert's landmarks. Minimal mTRE values larger than 2 mm are highlighted in bold and expose a potential need for a non-rigid registration. Values of mTRE larger than 3 mm are shown in italic and indicate a notable difference between experts






**Fig. 2** Homologous landmarks identified by Expert 1 and 2 for Case 1. Points identified by Expert 1 are colored in *blue* (MR) and *white* (US). Points identified by Expert 2 are colored in *green* (MR) and *yellow* (US). A coronal and transverse slice of the MR are also shown for reference, as well as a translucent rendering of the skin's surface



**Fig. 3** Homologous landmarks identified by Expert 1 and 2 for Case 13. Points identified by Expert 1 are colored in *blue* (MR) and *white* (US). Points identified by Expert 2 are colored in *green* (MR) and *yellow* (US). A coronal and transverse slice of the MR are also shown for reference, as well as a translucent rendering of the skin's surface

The registration results for the three configurations and each of the fourteen cases are shown in Figs. 5 and 6. In Fig. 5, we illustrate the performance of each configuration with rela-




**Fig. 4** Overlapped pMR and iUS slices from Case 3 before and after registration. The first row shows the slices at their initial location (coronal, sagittal, and transverse). The second row shows the slices after registration with the proposed approach

tion to a superset of landmarks that includes all experts' landmarks. It is important to note that all configurations consistently improve the registration accuracy with relation to their corresponding initial location. Additionally, we can also observe that the resulting mTRE is very close to the minimal mean distance under a rigid transformation (depicted as a red dashed line). The same set of results are enlisted in Table 2, where we also report the processing times for each configuration. The first configuration, (GOA Full Mask), involves a processing time that ranges from 36 to 76 s, which is comparable to the processing time of conventional intensity-based methods. Alternatively, the second configuration, (GOA Subset), obtains a similar registration accuracy with significantly reduced processing times that range from 7 to 14 s. Finally, the third configuration, (GOA Subset on GPU), also obtains a similar registration accuracy but with highly reduced processing times that range from 0.61 to 0.93 s.

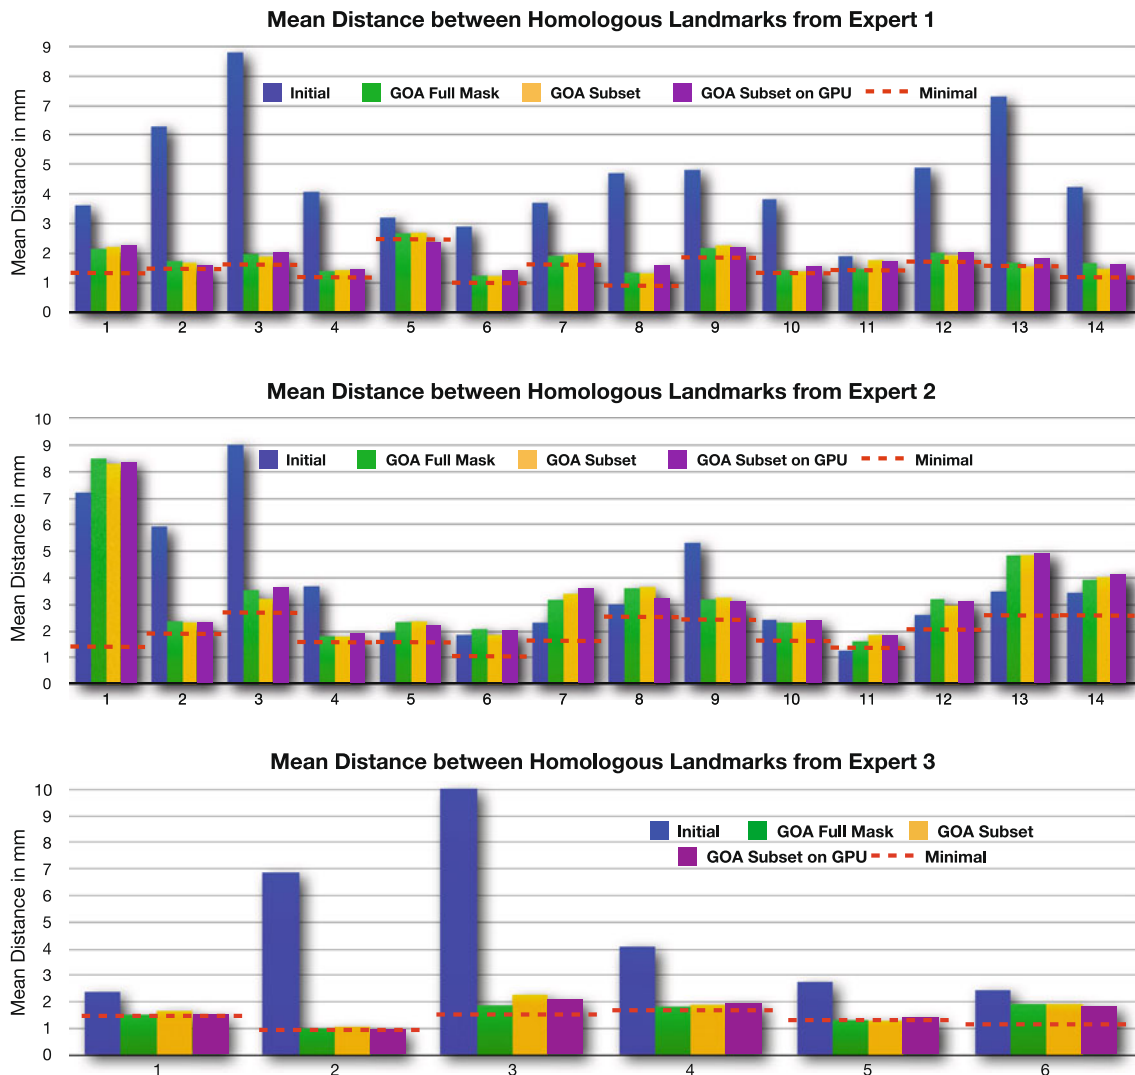
In Fig. 6, we illustrate the performance of all configurations with relation to the set of landmarks identified by each expert. It is immediately clear that our proposed method is in very close agreement with the landmarks selected by Expert 1 and Expert 3, since the resulting mTRE is quite close to their minimal mTRE. On the other hand, when evaluating with relation to Expert 2 landmarks, we tend to encounter a slightly larger value than the minimal mTRE. However, this comes as no surprise, since in the previous section we already demonstrated that Expert 2 had significant disagreements with Expert 1 and 3. Thus, any method that tends to align more with landmarks from Expert 1 and 3 will inevitably show a degradation with relation to the landmarks from Expert 2.

We also present a statistical summary of the registration performance of our proposed approaches as well as of competing techniques in Table 3. The evaluated methods are characterized by a choice of similarity metric: mutual information (MI), normalized cross-correlation (NCC), normalized



**Fig. 5** Rigid registration results with proposed method evaluated with the set of *all* landmarks independently identified by three different experts. The *x* axis corresponds to the clinical case, while the *y* axis corresponds to the mTRE between manually identified corresponding anatomical points. Also shown are the initial mTRE (Initial) in *blue*, as

well as the minimal mTRE (Minimal) possible under a rigid transformation, shown as a *red dashed line*. Notice that the three configurations yield a mTRE just slightly larger than the minimal mTRE. Table 2 also lists the resulting mTRE values for each configuration as well as the corresponding processing times



**Fig. 6** Rigid registration results with proposed method evaluated with each of three sets of landmarks identified by a particular expert. The *x* axis corresponds to the clinical case, while the *y* axis corresponds to the mTRE between manually identified corresponding anatomical points. Also shown are the initial mTRE, (Initial), in *blue*, as well as the minimal mTRE (Minimal) possible under a rigid transformation, shown as a

*red dashed line*. Notice that for Expert 1 and 3, both configurations yield a mTRE just slightly larger than the minimal mTRE and significantly decreased with relation to the initial mTRE. However, evaluation with relation to Expert 2 yields various cases with poor performance (e.g., Case 1, 7, 8, 12, 13 and 14)

**Table 2** Rigid registration results with proposed method evaluated as the mean distance (i.e., mTRE in mm) and range (i.e., minimal and maximal value) between all the landmarks independently identified by all experts combined

GOA Case	Full mask			Subset			Subset on GPU		
	mTRE in mm	(min, max)	Time in s	mTRE in mm	(min, max)	Time in s	mTRE in mm	(min, max)	Time in s
1	4.89	(0.72, 11.0)	53	4.86	(1.06, 10.51)	12	4.86	(0.81, 10.55)	0.71
2	1.79	(0.16, 5.08)	50	1.78	(0.32, 5.13)	9	1.71	(0.25, 5.11)	0.71
3	2.73	(0.37, 6.70)	76	2.65	(0.60, 6.89)	16	2.83	(0.67, 6.91)	0.93
4	1.68	(0.41, 3.91)	37	1.72	(0.46, 3.97)	7	1.76	(0.51, 3.95)	0.61
5	2.12	(0.33, 6.80)	36	2.13	(0.21, 6.40)	9	2.00	(0.22, 6.15)	0.88
6	1.81	(0.36, 3.22)	48	1.71	(0.45, 3.22)	9	1.81	(0.41, 3.21)	0.68
7	2.51	(0.76, 5.38)	48	2.64	(0.87, 5.67)	14	2.75	(0.89, 5.71)	0.79
8	2.63	(0.64, 5.22)	56	2.65	(0.55, 5.22)	10	2.51	(0.51, 5.11)	0.82
9	2.70	(0.88, 6.83)	40	2.79	(0.73, 6.98)	7	2.67	(0.78, 6.74)	0.86
10	1.95	(0.25, 4.62)	73	1.94	(0.36, 4.94)	9	2.05	(0.37, 4.97)	0.67
11	1.56	(0.21, 3.45)	52	1.82	(0.34, 3.89)	9	1.78	(0.31, 3.79)	0.73
12	2.64	(1.01, 5.35)	41	2.47	(0.96, 5.52)	8	2.58	(1.00, 5.41)	0.72
13	3.47	(0.89, 7.30)	43	3.42	(0.47, 7.05)	12	3.57	(0.63, 7.15)	0.89
14	2.94	(0.84, 5.31)	62	2.92	(0.71, 5.74)	9	3.04	(0.86, 5.42)	0.69
Average	2.53	(0.56, 5.73)	51	2.54	(0.58, 5.80)	10	<b>2.57</b>	(0.59, 5.73)	<b>0.76</b>
Median	2.57	(0.53, 5.33)	49	2.55	(0.51, 5.60)	9	<b>2.54</b>	(0.57, 5.42)	<b>0.73</b>

The results are also illustrated in Fig. 5. Also shown is the processing time for each case. Notice that the GPU implementation of our proposed approach (whose average and median performance values are highlighted in bold) yields an average processing time of 0.76 s

**Table 3** Statistical summary of rigid registration results with all evaluated techniques

Method	mTRE in mm		mTRE minus minimal mTRE in mm		Number of successes	
	Average	Median	Average	Median	<1mm	<2mm
MI+ORI+CAE	22.62	9.20	20.42	7.51	3	3
NMI+ORI+CAE	21.58	8.63	19.39	6.89	3	4
NCC+ORI+CAE	70.58	76.30	68.38	74.11	0	0
MI+ORI+GD	11.83	6.84	9.64	4.44	3	5
NMI+ORI+GD	11.68	6.88	9.64	4.44	3	5
NCC+ORI+GD	37.83	29.73	35.63	28.21	0	0
MI+GM+CAE	3.00	2.90	0.81	0.64	9	12
NMI+GM+CAE	3.01	2.78	0.82	0.59	10	12
NCC+GM+CAE	6.59	3.09	4.39	0.79	8	11
MI+GM+GD	2.90	2.67	0.71	0.61	11	14
NMI+GM+GD	2.87	2.66	0.67	0.60	12	13
NCC+GM+GD	3.77	2.91	1.57	0.69	9	12
NMI+PRE+CAE	13.47	9.22	11.28	6.83	0	1
NMI+PRE+GD	10.78	8.67	8.59	5.83	0	0
GOA full mask	2.53	2.57	0.33	0.27	14	14
GOA Subset	2.54	2.55	0.34	0.22	14	14
GOA Subset on GPU	2.57	2.54	0.37	0.33	14	14

The first column identifies the method used which are characterized by a similarity metric (e.g., mutual information (MI), normalized cross-correlation (NCC), normalized mutual information (NMI), and gradient orientation alignment (GOA)), input images (e.g., original images (ORI), gradient magnitude images (GM), and median filtered MR in conjunction with a Gaussian blurred US (PRE)), and optimization strategy (e.g., covariance matrix adaptation evolution strategy (CAE) and a gradient descent (GD) optimizer with adaptive gain). The first two columns of registration results show the average and median value of the mTRE between manually identified corresponding points. Also shown are the average and median value of the mean distance minus the minimal mTRE. The last columns indicate the number of cases in which the resulting mTRE is larger than the minimal mTRE by no more than 1 or 2 mm, respectively

mutual information (NMI), and gradient orientation alignment (GOA); input images: original images (ORI), gradient magnitude images (GM), and median filtered MR in conjunction with a Gaussian blurred US (PRE); and optimization strategy: covariance matrix adaptation evolution strategy (CAE) and a gradient descent (GD) optimizer with adaptive gain. We report the mean and median value of two registration accuracy measures: the mTRE of the homologous landmarks from all experts combined, and the difference between such mTRE and the minimal mTRE. Additionally, we demonstrate the number of cases that had a successful registration, where we define success as an instance when the mTRE is larger than the minimal mTRE by no more than 1 or 2 mm.

The results obtained clearly indicate that conventional multi-modal intensity-based metrics, like MI, NCC, and NMI, generally show very poor performance in this particular context. Nonetheless, if we make use of gradient magnitudes images and employ the same multi-modal similarity metrics, we obtain slightly improved results. A particularly good-performing configuration involves the maximization of NMI between gradient magnitude images with a gradient descent strategy, which successfully registers 12 of the 14 cases with an mTRE less than 1 mm larger than the minimal mTRE.

In comparison with other competing techniques evaluated, our proposed approach is the only one where all cases have a mTRE less than 1 mm larger than the minimal mTRE. In particular, the median value of the difference between the mTRE and the minimal mTRE is of 0.27 mm for the first configuration, (GOA Full Mask); 0.22 mm for the second configuration, (GOA Subset); and 0.33 mm for the third configuration, (GOA Subset on GPU).

## Discussion

We have presented a rigid registration method for MR to iUS and evaluated it with fourteen clinical cases obtained from the MNI BITE dataset. Registration accuracy was evaluated with the use of homologous landmarks identified by two or three experts, which provide a reliable indication of the positive performance of the algorithm with relation to competing methods. It is important to note that, while our validation with real clinical cases is much more preferable than a synthetic setup which relies on brain phantoms, there are still some serious limitations with such a validation strategy. In particular, we have pointed out that there is inherently an uncertainty in the points identified by the experts. We have exposed some of the limits of such validation strategy by reporting the minimal mTRE obtainable with a rigid transformation (with *all* available landmarks or with the landmarks associated with a particular expert), as well as the mTRE with relation to another expert's solution. Such analysis shows that there is

a significant variability in the performance metric dependent on the particular expert used for landmark identification.

While for some cases a rigid body registration might be sufficient for a successful alignment of all regions of interest, there are also cases that might require correction of non-rigid deformations. It is important to note that the evaluation of a non-rigid registration approach magnifies the challenge of adopting a proper validation scheme. Future work in this direction should not only propose a suitable non-rigid registration scheme, but also present a principled approach for evaluating its accuracy.

It is important to note that we evaluated our proposed approach on real clinical cases *prior* to resection. Future work will directly address the task of registration given the presence of resection cavities and evaluate its performance. Note that such context involves the additional challenge of uncorresponded tissue across modalities.

We have also specified that our current setup suffers a significant processing delay related to the reconstruction of a 3D US volume with the set of acquired 2D US slices. It is likely that a reduced processing time can be obtained by software optimization or by adopting a GPU-based solution [4]. Alternatively, we can adopt a slice-to-volume registration strategy which effectively bypasses the need of a reconstruction algorithm and provides almost immediate visual feedback. Future work will evaluate the performance of such techniques.

## Conclusion

We have presented a new and robust approach for the rigid registration of pMR to iUS which provides fast and robust performance evaluated over fourteen clinical cases. Our proposed approach registered all cases with a median difference between the mTRE and the minimal mTRE (i.e., lower bound) of 0.22, 0.27, or 0.33 mm (depending on the configuration). Furthermore, we have shown the computational efficiency of our technique which allows for registration times as low as 0.61 s by using a reduced set of locations and adopting a GPU-based implementation. Finally, in order to expose some of the limitations of our validation strategy, we have also reported an analysis on the manually identified landmarks used for measuring registration accuracy, where we have shown that there is a significant variability between experts.

**Acknowledgments** We would like to thank Dr. Laurence Mercier for her work putting together the MNI BITE dataset, and Simon Drouin and Anka Kochanowska for their technical support for rendering Figs. 2 and 3. We acknowledge funding from the Natural Sciences and Engineering Research Council of Canada (NSERC), and the Canadian Institutes of Health Research (CIHR MOP 97820, MOP-74725 & CHRP).

**Conflict of interest** None.

## References

- Automatic non-linear mri-ultrasound registration for the correction of intra-operative brain deformations. In: Niessen W, Viergever M (eds) *Medical image computing and computer-assisted intervention- MICCAI 2001, Lecture Notes in Computer Science*, vol. 2208, pp 913–922. Springer, Berlin (2001)
- Beyond mutual information: a simple and robust alternative. In: Meinzer HP, Handels H, Horsch A, Tolxdorff T (eds) *Bildverarbeitung für die Medizin 2005, Informatik aktuell*, pp 350–354. Springer, Berlin (2005)
- Brooks R, Collins D, Morandi X, Arbel T (2008) Deformable ultrasound registration without reconstruction. In: *Medical image computing and computer-assisted intervention-MICCAI 2008, Lecture Notes in Computer Science*, vol. 5242, pp 1023–1031. Springer, Berlin
- Dai Y, Tian J, Dong D, Yan G, Zheng H (2010) Real-time visualized freehand 3d ultrasound reconstruction based on gpu. *IEEE Trans Inf Technol Biomed* 14(6):1338–1345
- De Nigris D, Collins D, Arbel T (2012) Fast and robust registration based on gradient orientations: case study matching intra-operative ultrasound to pre-operative mri in neurosurgery. Springer, Berlin, pp 125–134
- De Nigris D, Collins DL, Arbel T (2012) Multi-modal image registration based on gradient orientations of minimal uncertainty. *IEEE Trans Med Imaging* PP(99):1
- De Nigris D, Mercier L, Del Maestro R, Collins DL, Arbel T (2010) Hierarchical multimodal image registration based on adaptive local mutual information. In: *Medical image computing and computer-assisted intervention-MICCAI 2010, Lecture Notes in Computer Science*, vol. 6362, pp 643–651. Springer, Berlin
- Hansen N, Ostermeier A (2001) Completely derandomized self-adaptation in evolution strategies. *Evolut Comput* 9:159–195
- Heinrich M, Schnabel J, Gleeson F, Brady F, Jenkinson M (2010) Non-rigid multimodal medical image registration using optical flow and gradient orientation. In: *Proceedings of medical image understanding and analysis*
- Ji S, Wu Z, Hartov A, Roberts DW, Paulsen KD (2008) Mutual-information-based image to patient re-registration using intra-operative ultrasound in image-guided neurosurgery. *Med Phys* 35(10):4612–4624
- Karaçalı B (2007) Information theoretic deformable registration using local image information. *Int J Comput Vis* 72(3):219–237
- Letteboer M, Willems P, Viergever M, Niessen W (2005) Brain shift estimation in image-guided neurosurgery using 3-d ultrasound. *IEEE Trans Biomed Eng* 52(2):268–276
- Loupas T, McDicken W, Allan P (1989) An adaptive weighted median filter for speckle suppression in medical ultrasonic images. *IEEE Trans Circuits Syst* 36(1):129–135
- Mellor M, Brady M (2004) Non-rigid multimodal image registration using local phase. In: *Medical image computing and computer-assisted intervention—MICCAI 2004*, vol. 3216, pp 789–796
- Mercier L, Del Maestro R, Petrecca K, Araujo D, Haegelen C, Collins DL (2012) Online database of clinical mr and ultrasound images of brain tumors. *Med Phys* 39(6):3253–3261
- Mercier L, Fonov V, Haegelen C, Del Maestro R, Petrecca K, Collins D (2012) Comparing two approaches to rigid registration of three-dimensional ultrasound and magnetic resonance images for neurosurgery. *Int J Comput Assist Radiol Surg* 7(1):125–136
- Otsu N (1979) A threshold selection method from gray-level histograms. *IEEE Trans Syst Man Cybern* 9(1):62–66
- Pluim J, Maintz J, Viergever M (2000) Image registration by maximization of combined mutual information and gradient information. *IEEE Trans Med Imaging* 19(8):809–814
- Pluim J, Maintz J, Viergever M (2000) Image registration by maximization of combined mutual information and gradient information. In: *Medical image computing and computer-assisted intervention- MICCAI 2000, Lecture Notes in Computer Science*, vol. 1935, pp 103–129. Springer, Berlin
- Roche A, Pennec X, Malandain G, Ayache N (2001) Rigid registration of 3D ultrasound with MR images: a new approach combining intensity and gradient information. *IEEE Trans Med Imaging* 20:1038–1049
- Sibson R (1978) Studies in the robustness of multidimensional scaling: procrustes statistics. *J R Stat Soc Ser B (Methodological)* 40(2):234–238
- Trobaugh JW, Trobaugh DJ, Richard WD (1994) Three-dimensional imaging with stereotactic ultrasonography. *Comput Med Imaging Graph* 18(5):315–323
- Unsgaard G, Selbekk T, Brostrup Mller T, Ommedal S, Torp SH, Myhr G, Bang J, Nagelhus Hernes TA (2005) Ability of navigated 3d ultrasound to delineate gliomas and metastases comparison of image interpretations with histopathology. *Acta Neurochir* 147:1259–1269
- Wein W, Brunke S, Khamene A, Callstrom MR, Navab N (2008) Automatic ct-ultrasound registration for diagnostic imaging and image-guided intervention. *Med Image Anal* 12(5):577–585
- Zhang W, Noble J, Brady J (2007) Adaptive non-rigid registration of real time 3d ultrasound to cardiovascular mr images. In: Karssemeijer N, Lelieveldt B (eds) *Information processing in medical imaging, Lecture Notes in Computer Science*, vol. 4584, pp 50–61. Springer, Berlin