

Near Real-Time Robust Nonrigid Registration of Volumetric Ultrasound Images For Neurosurgery

Hassan Rivaz, D. Louis Collins

*PERFORM Centre, Department of Electrical and Computer Engineering, Concordia University, Montreal, QC, Canada
McConnell Brain Imaging Center, Montreal Neurological Institute, McGill University, Montreal, QC, Canada*

Abstract

Ultrasound images are acquired before and after the resection of brain tumors to help the surgeon to localize the tumor and its extent and to minimize the amount of residual tumor after the resection. Since the brain undergoes large deformation between these two acquisitions, deformable image-based registration of these datasets is of substantial clinical importance. In this work, we present an algorithm for nonrigid REgistration of ultraSOUND images (RESOUND) that models the deformation with free-form cubic B-splines. We formulate a regularized cost function that uses normalized cross correlation as the similarity metric. To optimize the cost function, we calculate its analytic derivative and utilize the stochastic gradient descent technique (Klein et al., 2007) to achieve near real-time performance. We further propose a robust technique to minimize the effect of non-corresponding regions such as the resected tumor and possible hemorrhage in the post-resection image. Using manually labeled corresponding landmarks in the pre- and post-resection ultrasound volumes, we show that our registration algorithm reduces the mean target registration error from an initial value of 3.7 mm to 1.5 mm. We also compare RESOUND to the previous work of (Mercier et al., 2013) and show that it has three important advantages: (1) it is fully automatic and does not require a manual segmentation of the tumor, (2) it produces smaller registration errors, and (3) it is about 30 times faster. The clinical dataset is available online on the BITE database website.

Keywords: Ultrasound registration, Non-rigid registration, Robust estimation, Normalized cross correlation, NCC, Brain surgery, Image guided neuro-surgery, IGNS

Introduction

Not all pathological tissue is removed during surgery. In some cases, residual tumor is left behind when it involves eloquent cortex and its removal will result in a functional or cognitive deficit in the patient. In other cases, the extent of brain tumors is difficult to define during neurosurgery. As a result, different studies report that the surgeons leave residual tumor in 64% (Stummer et al., 2006) and 54% (Knauth et al., 1999) of patients. Therefore, most neurosurgical systems are based on neuro-navigation and pre-operative images. In these systems, the pre-operative images, usually magnetic resonance (MR), act as “GPS road maps”, and an optical or electromagnetic tracking device, which tracks the tools, acts as the GPS location signal. This provides the surgeon with the guidance to locate the tumor and its extents in the operating room. Unfortunately, these neuro-navigation systems do not usually provide the required accuracy for three main reasons. First the brain deforms after craniotomy and opening of the dura, and therefore, the maps are no longer accurate. This deformation is referred to brain shift in the literature (Hill et al., 1998; Roberts et al., 1998). Second, image-to-patient registration is performed by selecting homologous landmarks on the skin and in the MR image, which is subject to error. And third, the tracking devices have errors, which

further reduce the navigation accuracy. As a result, guidance with intra-operative imaging is becoming more widespread.

MR imaging has been used during or after the operation (Nabavi et al., 2001; Hartkens et al., 2003; Nimsy et al., 2004; Claus et al., 2005; Hatiboglu et al., 2009). However, intra-operative MR is expensive and is only available at a very limited number of centers. Tracked intra-operative ultrasound, on the other hand, is affordable and not cumbersome and has been successfully used in neurosurgery (Bucholz et al., 1997; Unsgaard et al., 2002b,a; Lunn et al., 2003; Keles et al., 2003; Letteboer et al., 2005; Tirakotai et al., 2006; Rygh et al., 2005). In terms of accuracy, intra-operative ultrasound has been shown to be as good as intra-operative MR (Unsgaard et al., 2005) and Gerganov et al. (2009). Our group has also focused on integrating ultrasound with neurosurgery, either using B-mode ultrasound (Arbel et al., 2004; Mercier et al., 2012, 2013; Nigris et al., 2012; Rivaz and Collins, 2012) or power Doppler (Reinertsen et al., 2007). In our institute, surgeons acquire ultrasound images before the resection, and after performing the resection, they collect a second ultrasound mainly to look for residual tumor by comparing the pre- and post-resection images. As the surgery goes on, the brain shift increases (Nabavi et al., 2001) and therefore the two sets of images will be misaligned. Figure 1 shows an example for rigid alignment of the two ultrasound images using the tracking data. In the superimposed im-

Email address: hrivaz@ece.concordia.ca (Hassan Rivaz)




Figure 1: Pre- and post-resection ultrasound images. The misalignment in the superimposed images in (c) and (d) is mainly caused by the brain shift. Structures (1) and (2) are respectively the septum and ventricles. The tumor and the resection cavity are not present in these ultrasound slices.

age, some areas have more than 7 mm of misalignment.

The registration of the two ultrasound volumes, one acquired before and another at end of the tumor resection, is challenging due to three main reasons. First, the tissue deforms during the resection, which requires the registration to be non-rigid. Second, the intensity and contrast of two ultrasound images of the same tissue target can change depending on the imaging angle, time gain control (TGC) settings, resection of some tissue in the path of the ultrasound wave, and depth of the tissue in the ultrasound image. And third, there are some regions, such as the tumor and the resection site, that do not have correspondence in the other image.

Previous work that register ultrasound images include the following: Rohling et al. (1997, 1998) and Gee et al. (2003) performed rigid registration of two volumes using cross correlation. Krucher et al. (2000) performed 3D non-rigid registration using mutual information. Poon and Rohling (2006) divided the image volumes into sub-volumes and performed rigid registration over these sub-volumes using cross correlation. Grau et al. (2007) and Rajpoot et al. (2009) used local orientation and phase differences as the similarity measure to perform rigid registration between volumetric echocardiograms. Foroughi et al. (2006) used attribute vectors (Shen and Davatzikos, 2002) to automatically select a set of leading points in the first volume. The correspondences of these points are found in the second volume and the volumes are warped accordingly. This algorithm was adapted in Leung et al. (2009) to perform non-rigid registration of 2D ultrasound images in real-time, and was also extensively validated in Khallaghi et al. (2012).

To the best of our knowledge, the recent work of Mercier et al. (2013) was the first attempt to register ultrasound images acquired before and at the end of the neurosurgery. To reduce the effect of missing data, they manually segmented the tumor and masked it out of the data. They then used the cross-correlation-based non-rigid registration algorithm of ANIMAL+INSECT (Collins et al., 1999). Manual segmentation in 3D is challenging and time consuming and can hinder widespread clinical application of such method.

In this work, we present a tool for non-rigid REgistration of ultraSOUND volumes (RESOUND) using a regularized cost function. Our cost function has an image similarity term based

on the normalized cross correlation (NCC), and a smoothness constraint. We use free-form cubic B-splines to model the deformation field. To optimize the cost function, we use the analytic derivative of NCC and exploit the computationally efficient stochastic gradient descent algorithm (Klein et al., 2007). We also perform hierarchical registration in three levels to speed the computations and prevent the algorithm from getting trapped in local minima. In the coarse levels, an approximate transformation is found, which is used as a starting point for the finer levels. Three features of RESOUND make it computationally efficient: stochastic gradient descent optimization, analytic estimation of the gradient of the cost function, and hierarchical search. As a result, our basic implementation takes 5 sec to perform non-rigid registration of volumetric data on a single core of a 3.6 GHz processor. To minimize the effect of missing correspondences, we propose a technique based on standard deviation (std) of NCC gradients. We provide results on simulation and clinical data to validate our technique. The ultrasound data along with the landmarks that we use for validation are available on the BITE database (Mercier et al., 2012).

Methods

We first set up image registration as an optimization problem and show how it can be efficiently optimized using a gradient descent method. An outline of our algorithm is shown in Figure 2. We then present a robust technique for efficiently suppressing the effect of missing correspondences.

Image Registration

Let $I_m(\mathbf{x})$ and $I_f(\mathbf{x})$: $\Omega \subset \mathbb{R}^d \rightarrow \mathbb{R}$ be respectively the moving and fixed images. Here d is the dimension of the images; it is 2 for 2D and 3 for 3D images. Registration of these two images can be formulated as (see Figure 2)

$$C = D(I_f(\mathbf{x}), I_m(\mathbf{T}_\mu(\mathbf{x}))) + \frac{\omega_R}{2} \|\nabla \mu\|^2 \quad (1)$$

$$\hat{\mu} = \arg \min_{\mu} C \quad (2)$$

where D is a dissimilarity metric, ω_R is a regularization penalty weight, ∇ is the gradient operator and \mathbf{T}_μ is the transformation

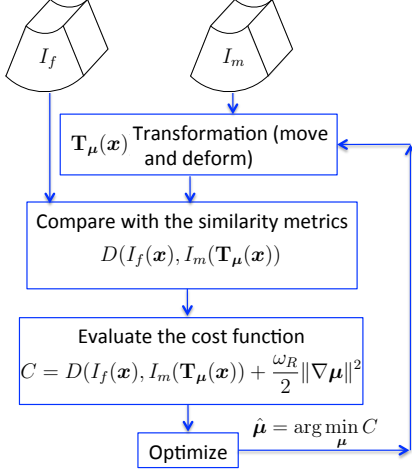


Figure 2: Iterative registration algorithm.

modeled by μ . We choose a free-form transformation parameterized by the location of cubic B-spline nodes. Therefore, μ determines the 3D displacement of all B-spline nodes. The regularization term is proportional to the norm of $\nabla\mu$, which is the difference between the translation of adjacent B-spline nodes. Therefore, this term is simply trying to produce a deformation field where neighboring B-spline nodes translate similarly. In this work, we use $-NCC^2$ as the dissimilarity metric. The negative sign is to transform the similarity metric NCC into a *dissimilarity* metric, and the power 2 is to make the derivative easier to compute. NCC works well for registering images of the same modality, i.e. ultrasound to ultrasound. Its value is 1 if the intensity relationship between the two images is linear and close to 0 if there is no intensity relationship (it can also be negative if the intensities in two images are flipped or inverted). We compute the NCC over N small cubic patches and average the results to calculate D :

$$D = -\frac{1}{N} \sum_{i=1}^N \rho_i^2, \quad (3)$$

$$\rho_i^2 = \frac{\left(\sum_{j=1}^n (f_j - \bar{f}_i)(m_j - \bar{m}_i)\right)^2}{\sum_{j=1}^n (f_j - \bar{f}_i)^2 \cdot \sum_{j=1}^n (m_j - \bar{m}_i)^2} \quad (4)$$

$$f_j = I_f(\mathbf{x}_j), \quad m_j = I_m(\mathbf{T}_\mu(\mathbf{x}_j))$$

where \bar{f} and \bar{m} are the mean values of I_f and I_m over the cubic boxes. The numerator of ρ is the covariance of I_f and I_m , and the denominator is the product of standard deviation of I_f and I_m . To further simplify the notation of Eq. 4, we set $f_j^0 = f_j - \bar{f}_i$ and $m_j^0 = m_j - \bar{m}_i$, and the column vectors $\mathbf{f} = \{f_1^0 \cdots f_n^0\}$ and $\mathbf{m} = \{m_1^0 \cdots m_n^0\}$ contain all the intensities inside the cubic patches. Using the new vector notation, Eq. 4 becomes:

$$\rho_i^2 = \frac{(\mathbf{f}^T \mathbf{m})^2}{(\mathbf{f}^T \mathbf{f})(\mathbf{m}^T \mathbf{m})} \quad (5)$$

where superscript T denotes transpose. Being compact, this equation is easier to handle for derivative computations in the next section.

Optimization

Given the current transformation parameter μ , a typical iterative optimization technique will seek an incremental $\Delta\mu$ so that $\mu + \Delta\mu$ reduces the cost function. Neglecting the regularization term of Eq. 1 to avoid clutter, the dissimilarity term

$$D(I_f(\mathbf{x}), I_m(\mathbf{T}(\mathbf{x}; \mu + \Delta\mu))) \quad (6)$$

should decrease. The gradient-based methods are faster than simplex and non-gradient based optimizations for finding $\Delta\mu$ especially if an analytic derivative can be obtained. For our D defined by Eqs. 3 and 5, the derivative is straightforward to compute as follows. We first apply the chain rule

$$\nabla_{\Delta\mu} D = \frac{\partial D}{\partial \Delta\mu} = \frac{\partial \mathbf{x}}{\partial \Delta\mu} \cdot \frac{\partial I_m}{\partial \mathbf{x}} \cdot \frac{\partial D}{\partial I_m} \quad (7)$$

where $\frac{\partial \mathbf{x}}{\partial \Delta\mu}$ is simply the transformation Jacobian, $\frac{\partial I_m}{\partial \mathbf{x}}$ is the image gradient, and $\frac{\partial D}{\partial I_m}$ is the derivative with respect to the intensity of all pixels. To compute $\frac{\partial D}{\partial I_m}$, we have to differentiate Eq. 5 with respect to the intensity of the pixels of the moving image, i.e. m . After some manipulations, we have:

$$\frac{\partial D}{\partial m} = -\frac{1}{N} \sum_{i=1}^N \frac{\partial \rho_i^2}{\partial m} \quad (8)$$

$$\frac{\partial \rho^2}{\partial m} = 2 \frac{\mathbf{f}^T \mathbf{m}}{(\mathbf{f}^T \mathbf{f})(\mathbf{m}^T \mathbf{m})} \left(\mathbf{f} - \frac{\mathbf{f}^T \mathbf{m}}{(\mathbf{m}^T \mathbf{m})} \mathbf{m} \right) \quad (9)$$

Note that here, ∂m is the *intensity* of different pixels/voxels of I_m and is independent of the transformation $T(\mathbf{x}; \mu)$.

We illustrate in Figure 3 different steps of computing $\frac{\partial D}{\partial \Delta\mu}$ using Eqs. 7 and 9. Here, I_f and I_m are images from BrainWeb (Collins et al., 1998), and the transformation $T(\mathbf{x}; \mu)$ is simply a rigid body translation in the y direction by 1 pixel. We compute NCC and its derivatives on patches of size 21×21 pixels. In the first step, we use Eq. 9 to compute the derivative of the NCC with respect to the intensity of each pixel (see part (c)). If we multiply this term by image gradients in the x and y directions, we get (d), which tells us the gradient descent direction for each pixel. Note that most pixels are moving upward, which is the ground truth transformation. Finally, multiplying the result of (d) by the transformation Jacobian $\frac{\partial \mathbf{x}}{\partial \Delta\mu}$, we get (e) which predicts the descent direction for each B-spline node. In cubic B-splines, if we denote the distance between the nodes as d and the size of the image in a direction as m , the number of nodes in that direction is $\lceil m/d \rceil + 3$. In this illustrative example, we have set the spacing between two adjacent nodes to $d = 7$ pixels. Therefore, the number of nodes for the cubic B-spline is $(21/7 + 3) \times (21/7 + 3) = 6 \times 6$, as shown in (e). This translates to $2 \times 36 = 72$ unknowns (2 because of the 2D image here) to solve for in the optimization problem of Eq. 1. Note that most nodes are moving upward in this example, which is the ground truth transformation.

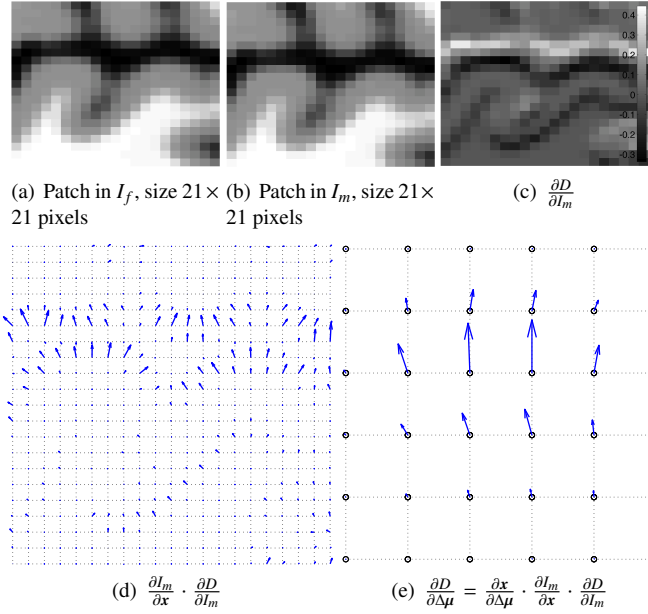


Figure 3: Computing the derivatives using the chain rule of Eq. 7. I_m is similar to I_f , except that it is moved 1 pixel down. Please see the text for details.

In our volumetric data, we compute NCC in small cubic patches, which can cover the entire image or can even have overlap. Such densely selected boxes, however, increases the computation time. Instead, we use a method similar to the stochastic gradient descent optimization method of Klein et al. (2007, 2009). In this method, the dissimilarity term is simply calculated only on a random subset of image samples \mathbf{x} in each iteration, and therefore it is computationally efficient. Since the pixel/voxel selection is random in each iteration, the result is not biased towards a subset of image samples and it converges to the same result as selecting all image samples in a fraction of the time (Klein et al., 2007, 2009). Our stochastic optimization implementation is as follows. In every iteration, we randomly select N cubic patches $i = 1 \dots N$ and compute the derivative of our cost function with respect to the B-spline nodes according to Eqs. 7 and 9. Note that these N patches are selected randomly throughout the volumes and are not necessarily centered on the B-spline nodes. Letting $\nabla_{\Delta\mu} C$ be the gradient of C in Eq. 1 with respect to $\Delta\mu$, the update equation is

$$\mu_{t+1} = \mu_t + a_t \nabla_{\Delta\mu} C. \quad (10)$$

The step size is a decaying function of the iteration number t :

$$a_t = \frac{a}{(A + t)^\tau} \quad (11)$$

with $a > 0, A \geq 0$ and $0 < \tau \leq 1$ user-defined constants (Klein et al., 2007). In the next section, we present our criteria for limiting the effect of missing correspondences.

Missing Correspondences

The sources of missing correspondences are abundant in our application. The tumor is replaced by the resection cavity, and

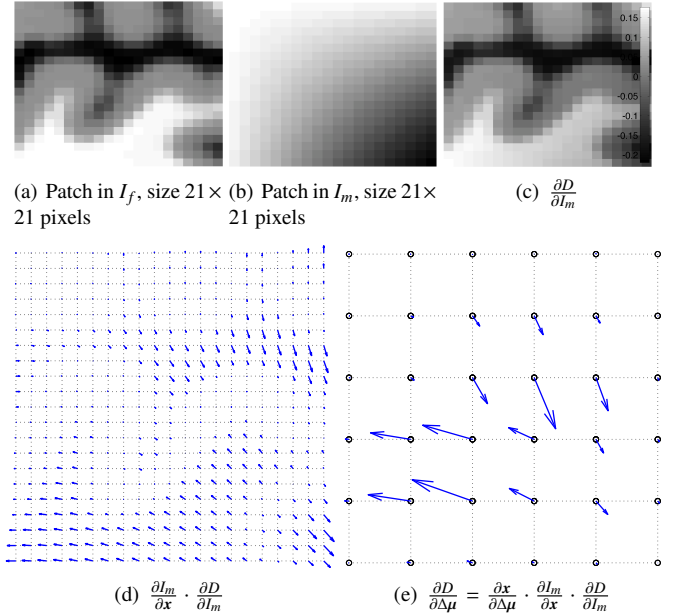


Figure 4: Computing the derivatives using the chain rule of Eq. 7 for two patches that do not correspond. Compare with Figure 3. See the text for details.

the areas around the resection are enhanced due to Surgical (Ethicon, Somerville, NJ), a hemostatic (i.e. blood-clotting) agent that is placed around the edges of the resection. Possible ruptured blood vessels are an additional source of missing correspondences. To minimize the adverse effect of missing data, we use the following approach, which is based on our chain rule estimation of the NCC derivative.

Assuming that the cubic patches are small enough, the deformation inside each patch is negligible. Therefore, the image derivatives inside the patch must have some agreement in terms of the direction of motion of the patch. Figure 3 (d) is an inlier example where the displacements for most voxels are pointing upwards, and not in random directions. Formally, let the vector $\frac{\partial D}{\partial \mathbf{x}}$ be $\frac{\partial I_m}{\partial \mathbf{x}} \cdot \frac{\partial D}{\partial I_m}$ (see also Figure 3 (d)). To keep or discard a patch, we compute the unitless metric r :

$$r = \min \left\{ \frac{\text{var}\left(\frac{\partial D}{\partial x}\right)}{\left\langle \frac{\partial D}{\partial x} \right\rangle^2}, \frac{\text{var}\left(\frac{\partial D}{\partial y}\right)}{\left\langle \frac{\partial D}{\partial y} \right\rangle^2}, \frac{\text{var}\left(\frac{\partial D}{\partial z}\right)}{\left\langle \frac{\partial D}{\partial z} \right\rangle^2} \right\} \quad (12)$$

where $\langle \cdot \rangle$ denotes mean and var is the variance. The numerator shows the variance in the orientation of the gradients of different voxels, and the denominator normalizes this value by the amplitude of the average of the gradient over the patch. A small ratio for each of the three values shows an agreement in terms of the orientation of the translation of the patch in the three $x y z$ directions. Picking the minimum value means the patch voxels agree on the translation orientation in at least one of the three directions. We illustrate in Figure 4 two patches that do not match. There is no agreement among the patch pixels in the descent direction of the translation; in part (d) the arrows for different pixels span almost 360° . We exploit this disagreement to find non-corresponding patches and discard patches that have




Figure 5: Comparison of average NCC values and the standard deviation of gradient orientations in simulated images. The edge structures in I_f and I_m is enclosed with respectively red and green squares. The edge is moved up by 2 pixels in I_m , as can be seen from the red and green squares. Most of the images, except for inside the red and green boxes, do not correspond. The value of NCC inside the box can be low, as shown in (c), and therefore a low NCC value does not necessary mean the data should be suppressed. However, the std of the gradient orientations in (d) is low inside the squares where the data can be trusted, and therefore can be used to suppress non-corresponding patches.

r values larger than a threshold T . Using one of our patient data, we found that a threshold value of $T = 0.2$ produces good results. We use this value for all of our patient data.

Experiments and Results

We first demonstrate the performance of our algorithm for finding non-corresponding patches on simulated data. We then present the results of RESOUND for registering pre-resection and post-resection ultrasound volumes.

Simulated Data

We simulate two simple 2D images I_f and I_m as shown in Figure 5. The purpose of this experiment is to only show whether RESOUND can find regions that do not correspond, and no registration is performed. Except for the parts inside the red (for I_f) and green (for I_m) rectangles, the images consist of uncorrelated strong noise. The parts inside the rectangles are identical, but are displaced in the vertical direction by 2 pixels. The goal is to consider the parts inside the rectangles as corresponding patches, and the rest of the image as non-corresponding patches. We simulate 400 different images I_f and I_m similar to the ones shown, with different instances of noise. We then compute NCC for all patches of size 5×5 pixels inside the images. Part (c) shows the average NCC values over the entire 400 pairs of images. Although the pixels inside the rectangle correspond to each other, they have a low average NCC. Part (d) shows that the r value, however, has lower values inside the rectangle and therefore can be used for finding non-corresponding patches.

Clinical Data

Data Description

The clinical data from 13 patients are acquired at the Montreal Neurological Institute, and are part of the BITE database available online (Mercier et al., 2012). The intra-operative US images are obtained using an HDI 5000 (Philips, Bothell, WA) with a P7-4 MHz phased array transducer. The pre-resection images were acquired before opening the dura, except for the

patients who had undergone a previous brain surgery and the dura was not used to close the cavity. For these patients, the images were taken directly on the cortex or the dura repair patch (Dura-Guard, Synovis, St. Paul, MN). After the resection, the surgeon acquired the post-resection ultrasound and used it to find any residual tumor. The study was approved by the Montreal Neurological Institute and Hospital Review Ethics Board, and informed consent from each participant was received.

The pixel size of 2D ultrasound images is 0.3 mm. The ultrasound probe is tracked with a Polaris tracking system (NDI, Waterloo, Canada), which provides the three locations and angles (6 DOF) of each image. Image acquisition, tracking, and synchronization of the tracking and imaging data are performed using the Intra-operative Brain Imaging System (IBIS) program, developed by Simon Drouin and Anka Kochanowska (Mercier et al., 2011). Figure 6 shows the 2D ultrasound images acquired before and after the resection. Each sweep contains between 200 and 500 2D images.

We reconstruct 3D volumetric data with an isotropic voxel size of 1 mm using a pixel based method (Solberg et al., 2007). This relatively large voxel size minimizes the effect of speckles, because, for every voxel, multiple measurements from different images are available (Rohling et al., 1997; Rivaz et al., 2011, 2007). We therefore do not perform any post-processing, such as median filtering or Gaussian blurring of the volumes. The depth, the area of the sweep and the number of 2D ultrasound images in each sweep is different. Therefore, different volumes have different sizes from 100^3 to 200^3 voxels. Figure 7 shows 2D slices from volumetric ultrasound data of four patients. The large tumor in the pre-resection and the resection cavity in the post-resection images do not correspond. In addition, the enhanced areas around the tumor are caused by Surgicel (Ethicon, Somerville, NJ), a blood-clotting agent that is placed around the edges of the resection, and do not have correspondences in the pre-resection images.

Evaluation Using mTRE Between Corresponding Landmarks

To validate the results, a neuro-radiologist has selected 10 corresponding anatomical landmarks in each of the 13 patients,




Figure 6: The ultrasound 2D images acquired before and after the resection of the tumor. (1) and (2) are respectively the longitudinal cerebral fissure and the skull. More than 200 images are acquired in each sweep; we are showing just a few here for improved visualization.

which are also available in BITE. The landmarks are anatomical structure visible in ultrasound such as sulci bifurcations, vessels, choroid plexus and septa. Figure 8 shows four such landmarks on the reconstructed 3D data for one patient. These landmarks are used to calculate mean target registration errors (mTRE) to assess the accuracy of the registration (Jannin et al., 2002). The mTRE of n corresponding marks at locations \mathbf{x} and \mathbf{x}' in the two images is calculated according to

$$\text{mTRE} = \frac{1}{n} \sum_i^n \|\mathbf{T}(\mathbf{x}_i) - \mathbf{x}'_i\| \quad (13)$$

where \mathbf{T} is the transformation (see Eq. 1) and $\|\mathbf{v}\|$ is the Euclidean length of the vector \mathbf{v} . The initial mTRE values are shown in Table 1; the mean initial mTRE over 13 patients is 3.7 mm.

To measure the accuracy of the manual landmark selection, the landmarks on the post-resection images were selected a second time. The mean distance between the two sets of landmarks in the *post-resection* images is 1.58 mm with a std of 0.60 mm. Using the two sets of landmarks in the post-resection images, two measurements for mTRE were calculated between the *pre- and post-resection* images, giving a mTRE mean absolute difference of 0.27 mm with a std of 0.25 mm. Comparing these numbers to the mean initial mTRE of 3.7 mm, the landmarks are accurate for validation purposes.

Parameter Selection

Larger values of N in Eq. 3 generally improve the results, as averaging is performed over more patches. This improvement, however, plateaus for very large values of N and also comes at the cost of increased computational cost. We found that $N = 100$ gives good results for our volumetric patient data. The size of the small cubic patches determines n ; Larger values of n increase the accuracy of NCC computation, but require larger patches which make the NCC values less local. We always set it to $n = 5^3 = 125$ pixels in our patient data.

The recommended values for optimization parameters of Eq. 11 are provided in Klein et al. (2007): A should be around 0.1 of the maximum number of iterations or less and τ should be

more than 0.6. The value of a is user-defined and is critical as it determines the step-size. If a is too small more iterations are required and it is also more likely that the optimization gets trapped in a local minima. On the other hand, the registration can diverge if a is too large. Fortunately, for large enough number of iterations the final registration result varies negligibly if a is varied by as much as 100%. a also depends on the similarity metric. We used one of our patient data and set it to values between 1 and 10^4 by multiplying it by 10 each time and evaluated the deformation at each iteration. After we found its order of magnitude, we varied it by smaller steps and finally set it to 100.

We perform the optimization hierarchically in three levels. The input volumes, which we use at the third (i.e. last) level, have a voxel size of 1 mm in all three directions. We downsample the input volumes by factor of two in all three directions twice, so that the voxel sizes in the second and first levels are 2 mm and 4 mm respectively. We set the distance between the B-spline nodes in $x y z$ directions to 32 voxels in all three levels, which translates to physical distances of 128 mm, 64 mm and 32 mm respectively in the first, second and third levels. Since the volumes have fewer voxels in the coarse levels, there are fewer B-spline nodes at these levels. The registration starts from the first level, and proceeds to the second and third levels. This multi-level optimization increases the capture range, reduces the chance of getting trapped in a local minimum.

Registration Results

We first show the displacement field, i.e. the brain shift, computed by RESOUND for patient 13 in Figure 9 (c). The magnitude of the displacements is between 0.5 mm and 2.9 mm. Note that the 2.9 mm is the maximum displacement in the slice shown; the maximum displacement in the entire volume is 4.1 mm. This value is close to the maximum initial TRE distance of 4.0 mm reported in Table 1. Also, note that the maximum deformation is located close to the cortex where the ultrasound probe is positioned. This is in accordance with the intuition that the highest displacement should happen around the cortex where the brain tissue is the least constrained.




Figure 7: Pre- and post-resection US images from four patients are shown respectively in the first and second columns. The third column shows the images superimposed. The white arrows in the first column point to the longitudinal cerebral fissure. The dashed red and solid green arrows respectively point to tumor cavity and Surgical enhancement respectively.

Figure 10 shows the substantial improvement in the alignment of the pre- and post-resection ultrasound images using RESOUND throughout the image, as identified by the arrows. The red contours are automatically calculated from the image in (a) using a Canny edge detector, and are overlaid on the ultrasound images in (b) and (c) to allow an assessment of the image alignments. Note that the initial alignment in (b) suggests a complete resection of the tumor, but (c) clearly shows (the black arrows) that residual tumor exists in the post-resection ultrasound. Therefore, deformable registration using RESOUND may result in greatly reduced residual tumor after the resection.

The final mTRE results for all patients are shown in Table 1. These results demonstrate that RESOUND produces small mTRE values even for patient 1 where the tumor size is 79.2 cm^3 , approximately equivalent to a cubic patch with 43 voxels in every dimension (each voxel is 1 mm). These areas do not correspond, and must be suppressed for accurate registration. Furthermore, RESOUND performs hierarchical multi-level registration to speed the computations, increase the convergence range and prevent the algorithm from getting trapped in local minima. Therefore, it converges even for patient 11 which has a large initial mTRE of 10.5 mm. The last three rows show a summary of results over the 13 patients:

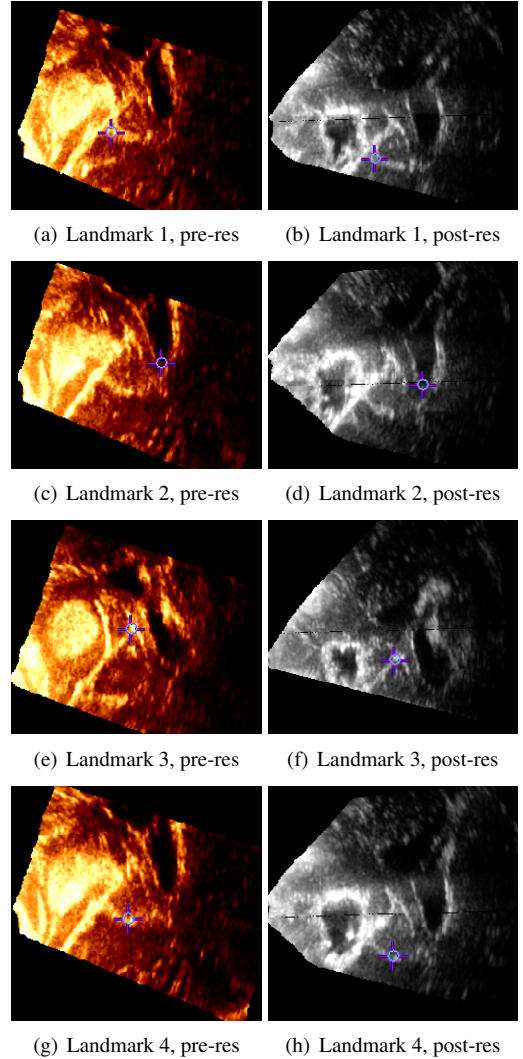


Figure 8: Four corresponding homologous landmarks shown in the pre- and post-resection ultrasound slices of patient P8. There are 10 corresponding landmarks in this patient's data.

- The average initial mTRE value over the 13 patients is 3.7 mm, which is reduced to 1.9 mm in Mercier et al. (2013) and 1.5 mm in RESOUND.
- The initial standard deviation in mTRE is 2.4 mm, which is reduced to 1.0 mm in Mercier et al. (2013) and 0.5 mm in RESOUND.
- The reduction in the mTRE values of both Mercier et al. (2013) and RESOUND are statistically significant: Mercier et al. (2013) and RESOUND give p -values of 0.002 using a paired t-test.

We also plot the distribution of maximum TRE values for all patients in Figure 11. We see that while both Mercier et al. (2013) and RESOUND reduce the maximum TRE values, this reduction is smaller for RESOUND at just 5.3 mm. The reduction of maximum TRE values from the initial values is statistically significant with both Mercier et al. (2013) and RESOUND: the

Table 1: The mTRE (mm) in initial alignment (obtained by rigidly registering US and MR using tracking information) and after registration. LGG and HGG respectively stand for low and high grade glioma. The smallest number in every row is in bold. The p -values in the last row show the statistical significance of improvement over the initial mTRE. The fifth column are the results from Mercier et al. (2013).

Patient	Tumor type	Tumor size (cm ³)	Initial	Mercier <i>et al.</i>	RESOUND
P1	LGG	79.2	2.3 (0.6-5.4)	1.3 (0.2-3.7)	1.8 (0.5-4.0)
P2	HGG	53.7	3.9 (2.8-5.1)	1.0 (0.5-1.9)	1.4 (0.7-2.3)
P3	HGG	31.6	4.6 (3.0-5.9)	3.7 (0.9-6.6)	1.4 (0.7-2.2)
P4	HGG	0.2	4.1 (2.6-5.5)	1.3 (0.5-2.8)	1.2 (0.3-2.4)
P5	HGG	32.3	2.3 (1.4-3.1)	1.1 (0.5-1.6)	1.0 (0.2-1.7)
P6	HGG	13.9	4.4 (3.0-5.4)	2.2 (1.1-4.8)	1.0 (0.4-1.7)
P7	HGG	63.1	2.7 (1.7-4.1)	1.4 (0.5-3.2)	1.7 (0.9-3.6)
P8	HGG	4.8	2.2 (1.0-4.6)	1.4 (0.3-3.9)	1.4 (0.6-3.2)
P9	HGG	10.4	3.9 (1.0-6.7)	2.9 (0.9-5.8)	1.9 (0.7-4.1)
P10	LGG	39.7	2.9 (0.8-9.0)	2.2 (0.2-9.5)	2.2 (0.6-5.3)
P11	LGG	49.1	10.5 (7.8-13.0)	3.7 (0.4-9.3)	2.5 (1.1-4.2)
P12	HGG	31.9	1.6 (1.3-2.2)	0.7 (0.4-2.0)	0.7 (0.2-1.6)
P13	LGG	37.3	2.2 (0.6-4.0)	1.4 (0.5-4.4)	1.3 (0.2-2.8)
mean	-	34.4	3.7 (2.1-5.7)	1.9 (0.5-4.6)	1.5 (0.5-3.0)
std	-	23.3	2.4 (1.9-2.8)	1.0 (0.3-2.6)	0.5 (0.3-1.1)
p -value	-	-	-	0.002	0.002




Figure 9: The brain shift in mm, along with the overlaid images before and after registration. Structures (1) and (2) in (a) are respectively the septum and ventricles. Note the misalignment in (d) pointed to by an arrow, which is corrected in (e). The results correspond to patient 13 in Table 1. The tumor and the resection cavity are not present in the shown ultrasound slices. The registration is performed in 3D.

p values are respectively 0.01 and 0.0008. The maximum TRE values of RESOUND are also smaller than that of Mercier et al. (2013) by a statistically significant amount with a p -value of 0.01.

Discussions

The importance of automatic registration of the pre- and post-resection ultrasound is twofold. First, it can help the surgeon find residual tumor tissue after the resection, and minimize the resection of the healthy and critical brain tissue. Second, it can be used to register pre-operative MR to post-resection ultrasound. Direct registration of these two images is challenging due to the large amount of the brain shift, tumor resection and missing data, and the inherent challenge in registering MRI and ultrasound. The pre-operative ultrasound can break up this hard registration problems into two simpler ones: registration of the MRI and the pre-resection ultrasound, and registration of the pre- and post-resection ultrasound. A disadvantage of such approach (compared to direct registration of the MR and post-resection US) is that composing these two simpler registrations accumulates errors. We will study the level of error accumulation in a future work.

The advantages of RESOUND over the method of Mercier et al. (2013) are threefold. First, it does not require manual segmentation of the tumor, which is a challenging and time consuming task. It rather incorporates a novel robust technique for finding missing correspondences. Second, RESOUND limits the effect of other sources of outliers such as bleeding in the tissue and shadowing in the US images. While the missing tumor correspondence is manually segmented in Mercier et al. (2013), the effect of other outliers may degrade the quality of those registration results. And third, Mercier et al. (2013) used the optimized implementation of ANIMAL+INSECT (Collins et al., 1999) which takes few minutes to non-rigidly register two volumes. Our basic implementation of RESOUND in MATLAB[®] MEX functions registers the two volumetric data in about 5 sec using a single core of a 3.6 GHz processor. This is mainly due to the multi-resolution pyramid approach and the efficient stochastic gradient descent optimization.

RESOUND can be integrated into the IBIS neuro-navigation system (Mercier et al., 2011). Our implementation is a basic implementation in MATLAB MEX, which can be improved to

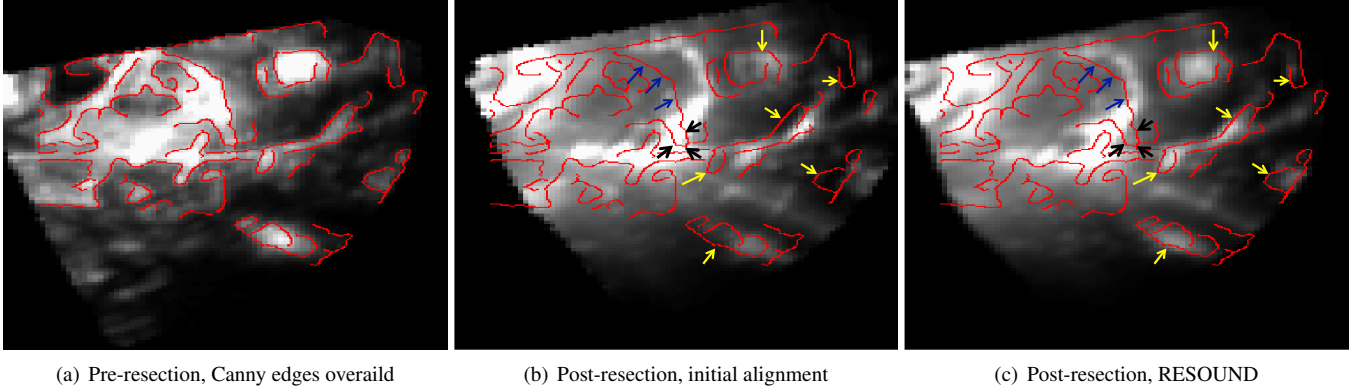


Figure 10: The poor initial alignment is improved after registration with RESOUND. The red contour is automatically calculated in (a) using a Canny edge detector and is overlaid on (b) and (c) to allow easy inspection of the alignment. The Canny edges are only used for visualization and not in image registration. The arrows point to the structures whose alignments are corrected in (c). The blue and black arrows point to the tumor boundary. The black arrows show residual tumor after resection, which is only visible *after* the registration in (c). The registration is performed in 3D.

achieve shorter run-time. Finally, the calculation of the cost function in Eq. 3 and its derivative in Eq. 7 are easily parallelizable over the $N = 100$ patches, allowing RESOUND to run in real-time.

RESOUND suppresses patches which have a large variation in the descent direction for each voxel. There is recent work that tries to match the orientation of image gradients (Karaçali, 2007; Nigris et al., 2012), i.e. aligns image edges, to register them. This line of work is fundamentally different from RESOUND outlier suppression mechanism in that it looks directly at image gradients, while RESOUND considers the gradient of the similarity metric.

In this work, we mainly focused on US registration in neurosurgery. Another interesting application for RESOUND is 2D (Rivaz et al., 2014) and 3D (Rivaz et al., 2008) elastography. In basic elastography, two images (2D or 3D) are acquired from the tissue before and after applying a deformation to the tissue. The goal is to recover the tissue deformation to get clues about tissue mechanical properties because parts of the tissue that compress less are likely tumors. Even state of the art elastography techniques can have problems recovering large tissue deformations in the presence of outliers. RESOUND can estimate an approximate deformation field, which can then be used as an initial solution for an elastography algorithm to estimate accurate deformations.

Conclusions

We presented RESOUND, a nonrigid registration algorithm that uses normalized cross correlation as the similarity metric and cubic B-splines as the deformation model. We efficiently optimized the cost function using stochastic gradient descent and multi-level hierarchical optimization. Therefore, we can registers volumetric data in 5 sec. We further proposed a robust technique to minimize the effect of missing correspondences such as the resected tumor and possible hemorrhage, and therefore does not need manual segmentation of the tumor.

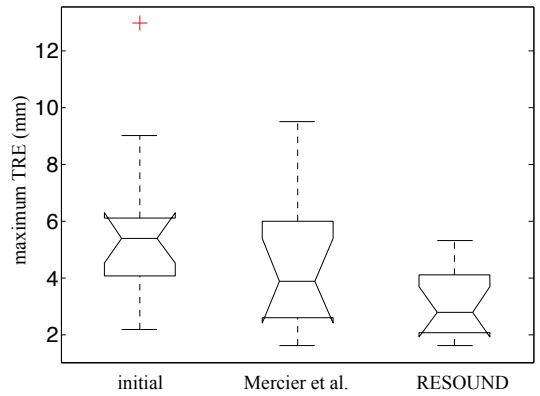


Figure 11: The boxplot of maximum TRE values from Table 1. RESOUND gives the smallest maximum TRE value of 5.3 mm.

We showed that our registration algorithm reduces the mean target registration error from an initial value of 3.7 mm to 1.5 mm. The tumor appearance in pre-resection ultrasound is more clear compared to the appearance of any residual tumor in the post-resection ultrasound. RESOUND allows registering these two volumes and therefore renders detecting tumor residuals substantially more likely, ultimately leading to a lower cancer recurrence rate.

Acknowledgements

We would like to thank the anonymous reviewers for their constructive and detailed comments. This work was financed by the Fonds Québécois de la recherche sur la nature et les technologies, the Canadian Institute of Health Research (MOP-97820), and the Natural Science and Engineering Research Council of Canada.

References

Arbel T, Morandi X, Comeau R, Collins DL. Automatic non-linear MRI-ultrasound registration for the correction of intra-operative brain deforma-

- tions. *Compt Aided Surg*, 2004;94:123–136.
- Bucholz R, Yeh D, Trobaugh J, McDurmont L. The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device. *CVRMed-MRCAS 97*, 1997:459–566.
- Claus E, Horlacher A, L H, Schwartz R, DelloIacono D, Talos F, Jolesz F, Black P. Survival rates in patients with low-grade glioma after intraoperative magnetic resonance image guidance. *Cancer*, 2005;103:1227–1233.
- Collins DL, Zijdenbos A, Baare A, Evans A. ANIMAL+INSECT: improved cortical structure segmentation. *Inform Proc Med Imag*, 1999;1613:210–223.
- Collins DL, Zijdenbos A, Kollokian V, Sled J, Kabani N, Holmes C, Evans A. Design and construction of a realistic digital brain phantom. *IEEE Trans. Medical Imag*, 1998;173:463–468.
- Foroughi P, Abolmaesumi P, Hashtrudi-Zaad K. Intra-subject elastic registration of 3D ultrasound images. *Medical Image Analysis*, 2006;10:713–725.
- Gee A, Treece G, Prager R, C C, Berman L. Rapid registration for wide field of view freehand three-dimensional ultrasound. *IEEE Trans. Med. Imag.*, 2003;22:1344–1357.
- Gerganov V, A S, Akbarian A, Stieglitz L, Samii M, Fahlbusch R. Reliability of intraoperative high-resolution 2d ultrasound as an alternative to high-field strength MR imaging for tumor resection control. *Neurosurgery*, 2009;1113:512–519.
- Grau V, Becher H, Noble JA. Registration of multiview real-time 3-d echocardiographic sequences. *Medical Imaging, IEEE Transactions on*, 2007;269:1154–1165.
- Hartkens T, Hill D, Castellano-Smith A, Hawkes D, Maurer C, Martin A HW, Liu H, Truwit C. Measurement and analysis of brain deformation during neurosurgery. *IEEE Trans. Medical Imag.*, 2003;22:82–92.
- Hatiboglu M, Weinberg J, Suki D, Rao G, Prabhu S, Shah K, Jackson E, R S. Impact of intraoperative high-field magnetic resonance imaging guidance on glioma surgery: A prospective volumetric study. *Neurosurgery*, 2009;64:1073–1081.
- Hill D, Maurer C, Maciunas R, Barwise J, Fitzpatrick M, Wang M. Measurement of intraoperative brain surface deformation under a craniotomy. *Neurosurgery*, 1998;433:514–526.
- Jannin P, Fitzpatrick JM, Hawkes DJ, Pennec X, Shahidi R, Vannier MW. Validation of medical image processing in image-guided therapy. *Neurosurgery*, 2002;212:1445–1449.
- Karaçali. Information theoretic deformable registration using local image information. *Int. J. Comp. Vis.*, 2007;723:219–237.
- Keles G, Lamborn K, Berger S. Coregistration accuracy and detection of brain shift using intraoperative sononavigation during resection of hemispheric tumors. *Neurosurgery*, 2003;53:556–562.
- Khallaghi S, Leung C, Hastrudi-Zaad K, Foroughi P, Nguan C, Abolmaesumi P. Experimental validation of an intrasubject elastic registration algorithm for dynamic-3D ultrasound images. *Medical Physics*, 2012;39:5488–5497.
- Klein S, Pluim J, Staring M, Viergever M. Adaptive stochastic gradient descent optimisation for image registration. *Int J Comput Vis*, 2009;81:227–239.
- Klein S, Staring M, Pluim J. Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines. *IEEE Trans. Image Proc*, 2007;1612:2879–2890.
- Knauth M, Wirtz CR, Tronnier VM, Aras N, Kunze S, Sartor K. Intraoperative mr imaging increases the extent of tumor resection in patients with high-grade gliomas. *American journal of neuroradiology*, 1999;209:1642–1646.
- Krucher J, Meyer C, LeCarpentier G, Fowlkes J, Carson P. 3D spatial compounding of ultrasound images using image-based nonrigid registration. *Ultrasound Med. Biol.*, 2000;22:1344–1357.
- Letteboer M, Willems P, Viergever M, Niessen W. Brain shift estimation in image-guided neurosurgery using 3-d ultrasound. *IEEE Trans Med. Imag.*, 2005;22:267–276.
- Leung C, Hashtrudi-Zaad K, Foroughi P, Abolmaesumi P. Intra-subject elastic registration of 3D ultrasound images. *Ultrasound Med. Biol.*, 2009;35:1159–1176.
- Lunn K, Paulsen K, D R, Kennedy F, Hartov A, West J. Displacement estimation with co-registered ultrasound for image guided neurosurgery: A quantitative in vivo porcine study. *IEEE Trans. Med. Imag*, 2003;22:1358–1368.
- Mercier L, Maestro R, Petrecca K, Araujo D, Haegelen C, Collins DL. Online database of clinical MR and ultrasound images of brain tumors. *Medical Physics*, 2012.
- Mercier L, Maestro R, Petrecca K, Araujo D, Haegelen C, Collins DL. Registering pre and postresection 3-dimensional ultrasound for improved visualization of residual brain tumor. *Ultrasound Med. and Biol.*, 2013;391:16–29.
- Mercier L, Maestro R, Petrecca K, Kochanowska A, Drouin S, Yan C, Janke A, Chen S, Collins DL. New prototype neuronavigation system based on preoperative imaging and intraoperative freehand ultrasound: System description and validation. *IJ CARS*, 2011;64:507–522.
- Nabavi A, Black PM, Gering DT, Westin CF, Mehta V, Pergolizzi Jr RS, Ferrant M, Warfield SK, Hata N, Schwartz RB, Wells WM, Kikinis R, Jolesz F. Serial intraoperative magnetic resonance imaging of brain shift. *Neurosurgery*, 2001;484:787–798.
- Nigris DD, Collins DL, Arbel T. Multi-modal image registration based on gradient orientations of minimal uncertainty. *IEEE Trans. Medical Imag*, 2012;30:2343–2354.
- Nimsky C, Fujita A, Ganslandt O, Von Keller B, Fahlbusch R. Volumetric assessment of glioma removal by intraoperative high-field magnetic resonance imaging. *Neurosurgery*, 2004;554:358–370.
- Poon T, Rohling R. Three-dimensional extended field-of-view ultrasound. *Ultrasound Med. and Biol.*, 2006;323:357–369.
- Rajpoot K, Grau V, Noble JA. Local-phase based 3d boundary detection using monogenic signal and its application to real-time 3-d echocardiography images. In: *Biomedical Imaging: From Nano to Macro*, 2009. ISBI'09. IEEE International Symposium on. IEEE, 2009. pp. 783–786.
- Reinertsen I, Lindseth F, Unsgaard G, Collins D. Clinical validation of vessel-based registration for correction of brain-shift. *Med. Imag. Anal.*, 2007;116:673–684.
- Rivaz H, Boctor E, Choti M, Hager G. Real-time regularized ultrasound elastography. *IEEE Trans Med Imaging*, 2011;304:928–945.
- Rivaz H, Boctor EM, Choti MA, Hager GD. Ultrasound elastography using multiple images. *Medical image analysis*, 2014;182:314–329.
- Rivaz H, Collins DL. Self-similarity weighted mutual information: A new non-rigid image registration metric. *Medical Image Computing Computer Assisted Intervention (MICCAI)*, 2012:91–98.
- Rivaz H, Fleming I, Assumpcao L, Fichtinger G, Hamper U, Choti M, Hager G, Boctor E. Ablation monitoring with elastography: 2d in-vivo and 3D ex-vivo studies. *Medical Image Computing & Computer Assisted Interventions, MICCAI*, New York, NY, 2008:458–466.
- Rivaz H, Zellars R, Hager G, Fichtinger G, Boctor E. Beam steering approach for speckle characterization and out-of-plane motion estimation in real tissue. *IEEE Int. Ultrasonics Symp.*, 2007:781–784.
- Roberts D, Hartov A, Kennedy F, Miga M, Paulsen K. Intraoperative brain shift and deformation: A quantitative analysis of cortical displacement in 28 cases. *Neurosurgery*, 1998;434:749–758.
- Rohling R, Gee A, Berman L. Three-dimensional spatial compounding of ultrasound images. *Medical Image Analysis*, 1997;13:177–93.
- Rohling R, Gee A, Berman L. Automatic registration of 3-d ultrasound images. *Ultrasound Med. and Biol.*, 1998;24:841–854.
- Rygh O, Selbekk T, Torp S, Lydersen S, Hernes T, Unsgaard G. Comparison of navigated 3D ultrasound findings with histopathology in subsequent phases of glioblastoma resection. *Acta Neurochirurgica*, 2005;15010:1033–1042.
- Shen D, Davatzikos C. HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imag.*, 2002;2111:1421–1439.
- Solberg O, Lindseth F, Torp H, Blake R, H N. Freehand 3D ultrasound reconstruction algorithms: A review. *Neurosurgery*, 2007;337:991–1009.
- Stummer W, Pichlmeier U, Meinel T, Wiestler OD, Zanella F, Reulen HJ. Fluorescence-guided surgery with 5-aminolevulinic acid for resection of malignant glioma: a randomised controlled multicentre phase iii trial. *The lancet oncology*, 2006;75:392–401.
- Tirakotai D, Miller S, Heinze L, Benes L, Bertalanffy H, Sure U. A novel platform for image-guided ultrasound. *Neurosurgery*, 2006;58:710–718.
- Unsgaard A, Gronningsaeter S, Ommedal S, Nagelhus H. Brain operations guided by real-time two-dimensional ultrasound: new possibilities as a result of improved image quality. *Neurosurgery*, 2002a;512:402–412.
- Unsgaard A, Gronningsaeter S, Ommedal S, Nagelhus H. Neuronavigation by intraoperative three-dimensional ultrasound: Initial experience during brain tumor resection. *Neurosurgery*, 2002b;504:804–812.
- Unsgaard A, Selbekk T, Gronningsaeter S, Ommedal S, Nagelhus H. Ability of navigated 3D ultrasound to delineate gliomas and metastases: Comparison of image interpretations with histopathology. *Acta Neurochirurgica*, 2005;1474:1259–1269.