Title Page

Title
Non-Rigid Registration for Image Guided Surgery

Character 47

Authors
Michael Weissberger[1], Yixun Liu[1], Joseph K Han MD[2], Nikos Chrisochoides PhD[1]

[1] *The College of William and Mary*,
[2] *Eastern Virginia Medical School*

Corresponding Author
Joseph K Han, MD
Eastern Virginia Medical School
Department of Otolaryngology & Head and Neck Surgery
600 Gresham Drive, Suite 1100
Norfolk, Virginia 23507
Phone (757)388-6229
Fax (757)388-6241
Email: HanJK@evms.edu

Financial Support or Funding
None

Key words:
CT images, MRI images, intra-operative, registration, image guidance, non rigid registration, CT – MRI fusion

Abstract

Objective:

Endoscopic skull base or cranial procedures can cause brain shift. Currently useful intra-operative soft tissue information is obtained via intraoperative MRI (iMRI). Intraoperative CT (iCT) has been introduced but soft tissue information is limited. Therefore our objective is to determine if non-rigid registration (NRR) between iCT images and pMRI images is possible for intraoperative soft tissue movement.

Methods:

Patients with skull base pathology with compression of soft tissue were recruited. pCT and pMRI images were obtained. Segmentation of the radiographic images was performed to create a tetrahedral volume mesh. Anatomic features were identified for "block matching". iCT images were acquired and rigid registration (RR) for iCT to pCT and pMRI were performed. Block matching between the iCT and the pCT was executed solving for the dense deformation field. NRR was achieved with the pMRI according to the dense deformation field. Matching and registration error were calculated and compared between RR and NRR.

Results:

Six patients were identified for the study. Segmentation, mesh generation, point selection, block matching, and dense deformation field estimation was successfully performed. The NRR accuracy improved on an average of 4.14 times compared to RR. The registration accuracy for acquired deformation improved by 19% for NRR over that of RR.

Conclusion:

This is the first study evaluating the NRR of iCT with pMRI. NRR is possible between the pCT and iMRI. Successful application of this technique would allow iCT to become a pseudo-iMRI at a fraction of the cost, however clinical application has not been determined.

Word 250

1. Introduction

Pre operative soft tissue, such as the brain, deformations can occur during a procedure and result in misalignment between preoperatively acquired datasets and the intra-operative acquired images. Intra-operative acquisition of fMRI (functional) and Diffusion Tensor Imaging (DTI) with intra-operative MRI is impractical today due to long image acquisition and processing times. Existing commercial systems use rigid registration algorithms to map pre-operatively acquired data like CT or MRI images into the navigational system. Recently pre-operative images have also been merged with intra-operative CT (iCT). However, the rigid algorithms between preoperative MRI and iCT do not provide sufficient accuracy when the brain deformation is greater than 10-15% of the adult brain size and therefore can not account for changes in brain shift[1]. Therefore, the study of real-time non-rigid registration algorithms for preserving the accuracy of preoperative data is an important problem.

Non-rigid registration (NRR) techniques are capable of utilizing transformations that model local deformations [2-4]. The numeric approximation of these transformations (i.e., determining their coefficients) is a computationally intensive task that has so far prohibited their intra-operative application. This technologic advancement will allow for precise fusion of a preoperative MRI (pMRI) with iCT images and would essentially allow an intraoperative CT scanner to become a pseudo intraoperative MRI (iMRI) machine at a fraction of the cost of a true iMRI. In other words: a less expensive version of an iMRI machine would be developed. This would then allow a technology closely resembling iMRI to be more readily available to many more patients, physicians, and hospitals owing to the decreased cost and lack of inconvenience of a magnetic device with ferrous instruments. Accurate non-rigid registration has been shown to increase the margin of tumor-tissue removal, which is correlated with improved patient outcomes [16,17]

In this paper we address an important obstacle that can help us to bridge the gap between the research prototypes and methods we are developing and their broad clinical evaluation. We present

1. Results from a limited clinical study using the NRR technology to fuse iCT and pMRI including both real and synthetically generated intra-operative data

2. Synthetic benchmarks which we developed to fine tune our algorithm and evaluate its accuracy automatically
3. Open problems and directions for future work.

2. Materials and Method

2.1 Overview

The process of non-rigid registration involves multiple steps. Several of these steps require significant manual effort and significant computational resources. Both of these requirements previously made the complete application of non-rigid registration impossible during surgery. By leveraging supercomputing facilities, and alternatively the commercial GPU, we overcome these logistical problems to accomplish the goal of real-time NRR [3]. We divide our process into two overall phases (Figure 1). In the pre-operative phase, we perform:

1. Physical acquisition of the pre-operative images.
2. Segmentation of the pre-operative MR to obtain a binary mask of the intra-cranial cavity (ICC).
3. Creation of a patient-specific model consisting of a tetrahedral volume mesh of this segmented image.
4. Salient feature detection, where we select optimal candidates for "block matching"
5. Distribution of these images to the various remote (computational cluster) or local (GPU) computational sites for use in the intra-operative phase.

The intra-operative stage has time-critical and non-time-critical stages. These include:

1. Acquisition of the intra-operative image
2. Rigid registration of iCT to the pre-operative CT.
3. Sparse deformation field detection (Block Matching) between the iCT and the pre-operative CT
4. Solving for the dense deformation field which describes the non-rigid movement of the tissue
5. Deformation of the pre-operative images (CT, DWI, MRI, etc.) according to the dense deformation field.

Figure 1

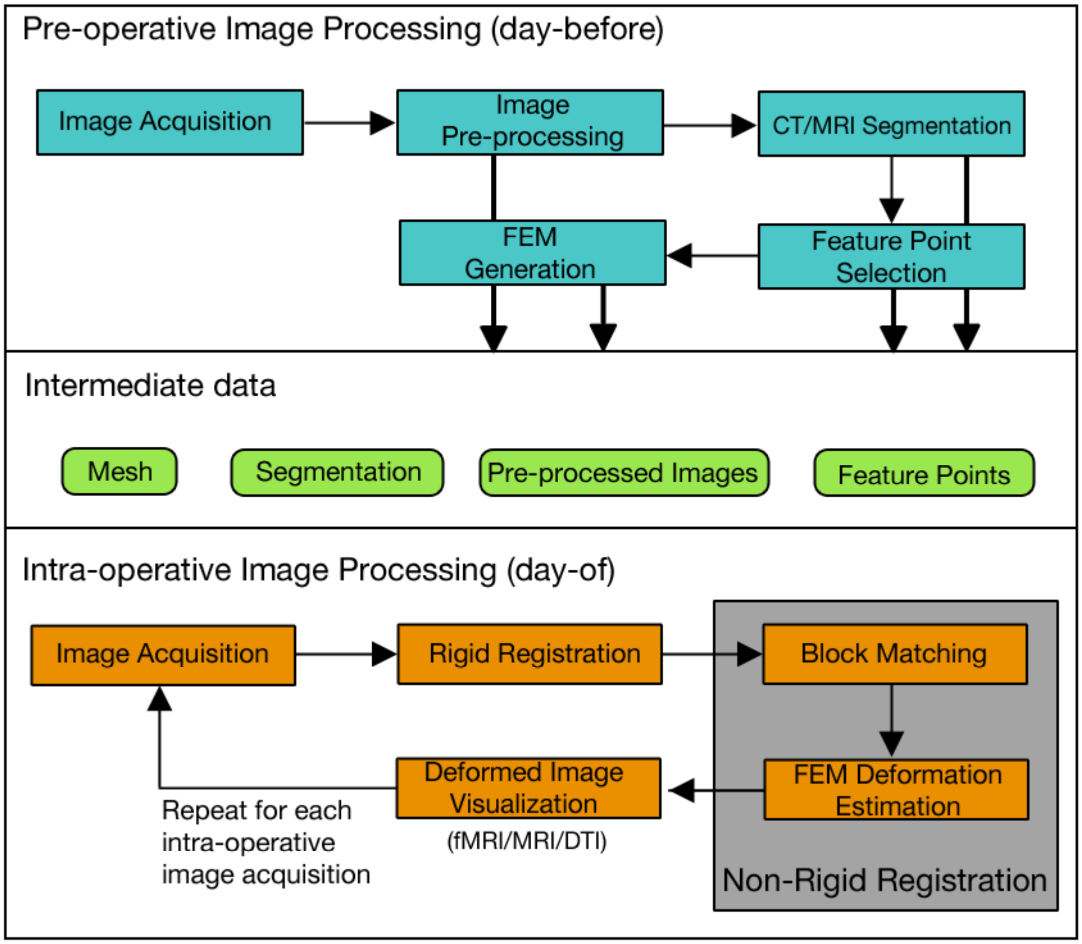

2.2 Pre-operative Phase

2.2.1 Image Acquisition

Physical acquisition of preoperative axial CT images are collected from a GE Lightspeed CT scanner (Pewaukee, WI) in DICOM format with the thickness of the axial slices being at the most 0.8 mm with no gantry tilt.

2.2.2 Segmentation

We separate the image into two regions in order to isolate the ICC, a process that is required to restrict the elastic motion of the brain model to ensure it does not deform outside of the cranial cavity. Allowing such deformations would be physically unrealistic but otherwise possible without the use of the patient-specific model in the mechanical energy term of the NRR formulation [4]. Segmenting the intra-cranial cavity is a semi-automatic process. We employed a combination of automatic operators, like region growing and level-set filters, with slice-by-slice manual segmentation to correct any erroneously included regions. We used Insight's Slicer3D software, and ITK's SNAP, parts of the Segmentation and Registration Toolkit [5], as well as in one case a novel model-based segmentation method [6].

2.2.3 Generation of the Patient-specific model

Our NRR uses a patient-specific model to tailor the computation to an individual patient's anatomy. This model uses a tetrahedral mesh (figure 4) generated from the segmentation produced in the previous step. Not all meshes are equal, and this mesh has several properties that make it well suited to our problem [7].

First, the mesh must provide a reasonably close approximation of the surface of the meshed object. A mesh which conforms perfectly to the image boundary would provide great fidelity, but at the cost of an excessive level of detail and an overwhelming number of mesh elements. We strive for an equidistribution of registration points. Next, the gradation of the mesh is an important property that may reduce interpolation error by having elements of smaller size. However, it is important to be able to adaptively refine the mesh in selected regions of interest to control computation time. Next, the shape of mesh elements is important, as small angles may increase the condition number of the stiffness matrix. Mesh generation time is important, as short times allow real-time mesh refinement. Extensive previous work was done to refine our meshing procedure to produce a mesh best-suited to our application of Image Guided Neurosurgery [8]. We use the red-green mesh algorithm discussed in [9].

2.2.4 Salient Feature Point Detection

The goal of the registration method is to recover the movement of the brain anatomy between the acquisition of the pre-op image and the intra-op image. The heart of the algorithm is a window-bounded block-matching algorithm that measures the sparse initial deformation field. We perform block-matching between the pre-operative (or floating) image (MRI or CT) and the intra-operative (or fixed) image (CT). Before we do this, we must target block matching to areas of the image that have a reasonable degree of structural information. We do this because block matching between structured sub-image regions will be less error-prone than block matching between regions with relatively less structural information. We call this stage point selection. We begin by placing a point (hereafter called a registration point) at the centers of all voxels inside the region of interest (ROI) of the floating image. This ROI is defined as the intersection of the floating image with the tetrahedral mesh. The registration points define the centers of sub-regions called blocks. The variance of the voxel intensities of each block is calculated, and a fraction of the blocks with the lowest variance is rejected. The remaining blocks are considered to be better candidates for block-matching with the intra-operative fixed image, which occurs in the next step.

2.3 Intra-operative Phase

The intra-operative phase begins with acquisition of the intra-operative CT images from a CereTom by NeuroLogica (Danvers, MA). The DICOM images are 1 mm axial images.

2.3.1 Rigid registration

The first computational step in the intra-operative phase is the rigid registration of the pre-operative data to the just-arrived intra-operative image. This step is required to remove the global rigid difference between the images that is due to different patient position and acquisition times. The time between the first pre-resection image and the first intra- resection image is generally long enough that the computational requirements of rigid registration is easily satisfied. We use a rigid-registration by maximization of normalized mutual information [10,11], implemented in the 3D slicer software package. This step generally takes less than five minutes.

2.3.2 Block matching

In the block-matching phase we begin to estimate the difference between the two images. The inputs to block matching are the pre-operative and intra-operative images to be registered, and the set of registration points selected during the pre-operative point selection phase. In order to find the deformation at the registration points we choose to perform a bounded exhaustive search by maximizing a measure of similarity between small cubic image regions (or blocks). The block selection procedure is repeated for each registration point in the pre-operative image. The resulting sparse deformation field is used as a rough estimation of the displacement between the two images.

The optimal similarity measure is a function of several factors. In previous studies [NI08] we used the Normalized Cross-Correlation. NCC is appropriate for mono-modal image registration, or image registration between images of the same modality. NCC is suitable for iCT-to-iCT registration since there is a linear or affine relationship between the two images' joint probability density functions. We can expect the intensity of any voxel in one image to be an affine function of the value of the corresponding voxel in the other image. This assumption does not necessarily hold for images from different modalities. The intensity of a voxel in an MRI is not an affine function of the corresponding voxel in a CT image. In fact, the intensity relationship between these two images isn't even functional - it is statistical [12]. A more complicated relationship often requires a more sophisticated similarity measure, such as the Correlation Ratio or one of many measures derived from information theory, like normalized mutual information (NMI). The most appropriate measure, however, is dictated by many inputs: the modality of the images, and the specifics of the images such as image subject, detector position, and lighting conditions [1]. We have compared the performance of several different similarity measures in order to choose an optimal measure for our purposes of image-guided head and neck surgery. To do this, we developed a set of synthetic benchmarks that are used to select the optimal similarity measure for two CT or MRI images.

2.3.3 Solving for the dense deformation field

The sparse deformation field measured with block matching suffers from two problems. First, it is sparse and discontinuous (defined only at the registration points). We wish to have a deformation which is defined everywhere, and which is smooth. Secondly, it is noisy because the block-matching algorithm is imperfect. The block matching measure may "choose" the wrong putative block, simply because it happens to be numerically "closer" in terms of the pre-defined similarity metric to the target block. This is especially true in areas with homogenous tissue intensity. The registration points with such erroneous block matchings are called “outliers”. We previously considered both an approximation and interpolation method. This competition is formulated as energy minimization problem:

$$W = (HU - D)^T S(HU - D) + U^T KU \text{ (1)}$$

where K is the mesh stiffness matrix, H is the linear interpolation matrix between the displacements recovered by block matching and those at the mesh vertices, S is the block matching stiffness matrix (matches with higher confidence are assigned higher weights), D is the vector for the block displacements, and U is the unknown displacement vector at the mesh vertices. The stiffness matrix K reflects an estimation of the physical properties of the brain tissue. At each iteration the following system of equations is solved:

$$F_0 = 0, F_{i-1} = KU_{i-1}, \left[K + H^T SH\right] U_i = H^T SD + F_{i-1} \text{ (2)}$$

where H^TSD is essentially the weighted displacement at mesh vertex, estimated from the displacements at the surrounding registration points. The system of linear equations is solved using biconjugate gradient stabilized method (implemented within the Gmm++ library [ ] with the diagonal preconditioner). The matrix H^TSH is block matrix and has a non-zero entry for each mesh vertex and edge, the cells of which contain registration points. A vertex or edge cell is defined as the set of tetrahedra incident on the mesh vertex or edge. The distribution of the registration points in the cells of the mesh vertices/edges affects the relative magnitude of coefficients of this matrix and in turn the condition number of the linear system.

The tetrahedral finite element mesh produced by the image-to-mesh conversion has a dual role in the above formulation. First, it is used in the mechanical energy (U^TKU) of the system to model deformation of the brain as a physical body. This is used to discover and discard the outlier registration points. At the end of each iteration a certain fraction of outlier block matches are discarded. The discarded block-matchings are those which were found to have low numerical confidence during block-matching, or which disagree with its neighbors. Second, the mesh is also used to regularize (or smooth) the displacements estimated from the minimization of the matching energy ($(HU- D)^TS(HU- D)$) from block matching. Once the solution converges, the algorithm is completed, and we have calculated a dense volumetric deformation that conforms to the measured block-matchings. For a more precise treatment of this algorithm, see [4,13].

2.4 Computational Framework

In order for these results to be useful during surgery they must be available in real time. In practical terms, this means a matter of minutes. There are several of constraints during surgery, and pausing the procedure for more than a few minutes is unacceptable in terms of risk to the patient. This NRR method takes a matter of hours to complete on a single computational node. The two different phases, pre-operative and intra-operative, have different time constraints. The pre-operative stage (includes image preparation, segmentation, mesh generation, and feature block/point selection) is not subject to constraints, and is successfully completed on a single hot processing node. The intra-operative phase consists of rigid registration, block matching and solving steps. This phase is subject to real time constraints. Block matching is the slowest of the steps in the intra-operative phase. Our parallel version of block-matching, which optionally employs real-time load balancing, allows execution of the block-matching stage in a matter of a few minutes on co-operative workstations with 50 processors. We have evaluated the use of a cluster supercomputer, grid computing resources (the TeraGrid), and a heterogenous GPU architecture to solve overcome these computational requirements. Our experience has shown us that scheduling problems make large distributed systems like the TeraGrid inappropriate for real-time use (Technical Report). While the cost of maintaining a computational cluster is prohibitive, cluster computing is a suitable environment for this work. We also investigated using a heterogeneous architecture making use of the GPU to perform block matching. The advantage of employing the GPU is that it moves the supercomputer next to the OR, and significantly reduces the cost of this NRR solution. In [13, 14] we present a detailed treatment of the computational framework we employed in our evaluation, described next.

2.5 Construction of a Synthetic Intra-operative Image

We use software described in [9] to generate the synthetic deformation fields. The software follows a method described by Rogelj et al in [12]. We create isotropic lattice (grid), and assigns a deformation vector at each node of the lattice. This deformation vector is drawn from a Gaussian distribution with parameters mean μ = 0 and variance σ [5-25], where higher σ results in larger deformation. These deformation vectors at the nodes of the lattice form a sparse and discontinuous deformation field. In order to obtain a more realistic continuous field, and to generate a dense deformation field, we then regularize these deformation vectors using thin-plate splines across the entire image volume. This dense deformation field can then be applied to the pre-op CT, resulting in a synthetically deformed CT which may play the role of the intra-operative CT during image registration. This synthetic intra-operative CT may then be used to evaluate the performance of the registration algorithm, compared with this known (synthetic) deformation field. An example deformation field is shown in figure 2.

Figure 2: An example deformation field volume. The deformation field shown is a 3D deformation volume, on the left, with magnitude of deformation shown as color. A 2D slice is extracted in the center. The field vectors for the same deformation field are shown on the right. All deformations are in units of voxels.

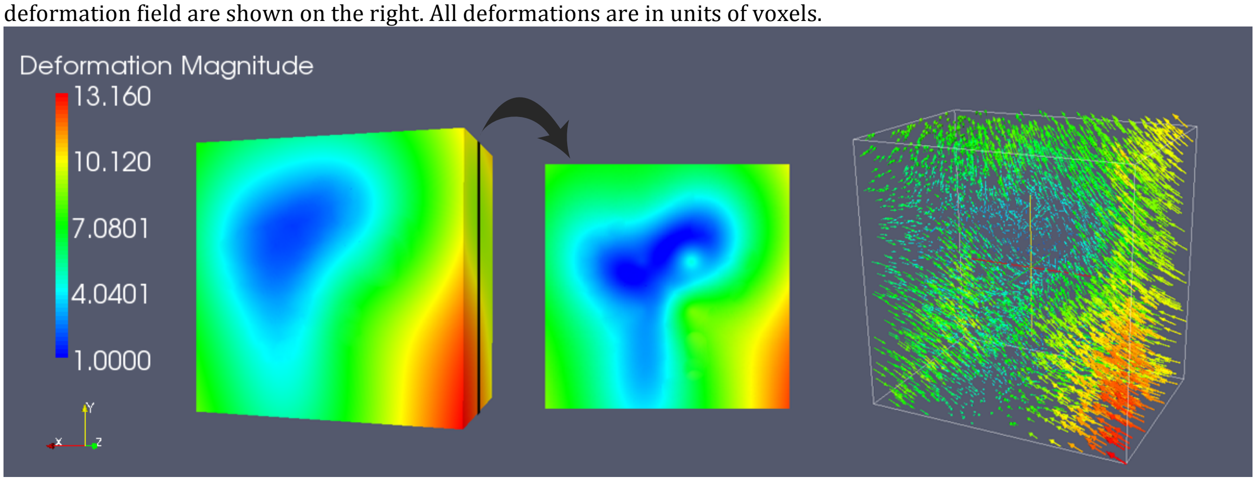

Results

We include two studies in this work. The first uses pre-operative data exclusively. We use these pre-operative CT images to answer a single question: *does CT have sufficient definition in the soft tissue regions to be a good candidate for intra-operative use as a fixed image during non-rigid registration*? To answer this question, we require a source of ground truth with which we will assess the performance of the registration algorithm. We use the previously described method to generate a synthetic deformation that serves as our ground truth. We then deform each pre-operative CT according to the synthetic deformation, and treat this new deformed CT as the intra-operative image. We perform segmentation, mesh generation, point selection, block matching and dense deformation field estimation as previously described. We perform the final two steps three times, once for each similarity metric: NMI, PMI and CC. We also include here results from using a new heterogeneous architecture for block matching that makes use of the GPU. The timings produced while using the GPU are sufficiently short to satisfy these strict requirements, however only single-precision arithmetic is available currently in commercial GPUs. This drawback will soon be rectified.

We evaluate the quality of registration in several ways. First, we calculate the matching error of block matching by subtracting the recovered displacement from the true displacement for each measured displacement vector, located at the registration points. This measure is reported in millimeters. Next, we calculate the registration error by subtracting the final dense deformation calculated at the end of each registration with the true deformation, for each mesh vertex of the finite element mesh on which the dense deformation is defined (figure 4). This measure is also reported in millimeters. In addition, the symmetric Hausdorff distance (SHD) is a measure of the distance between two point-sets as previously described [15]. We calculated the ratio of the SHD of the images rigidly registered to the SHD of the same images non-rigidly registered. Finally, mutual information may be used to assess the quality of a registration, as it measures how much information is contained in the relationship between the two intensity distributions. Here we use the normalized mutual information variant.

The second study performed uses actual intra-operative data. These images are captured intra-operatively with the previously described intra-operative CT scanner. We present three clinical cases that correspond to pre-operative synthetic cases 4-6. The registration method is the same as for the synthetic case, however the intra-operative data is not synthetically deformed in any way. In this study we are attempting to recover the real deformation caused by brain shift and mass effect during surgery. This presents a challenge for accuracy assessment and validation. In the synthetic case we have a readily available source of ground truth: the synthetic deformation. In these clinical cases there is no ground truth. For each case, we selected five points in the pre-operative image, and found the same five homologous points in the intra-operative images. We then treat this deformation as the true deformation and compare it to the deformation recovered at the same points with this non-rigid registration algorithm. The selection of these specific points represent a range of anatomical landmarks, some near the resection margin, and some farther from the resection margin in deep brain tissues. Because there is some subjectivity in the selection of these landmarks, and because the task is a difficult one with these images, we used two independent experts, a neurosurgeon and a radiologist, to independently select landmarks.

Synthetic Study

For each case in table 1 we report:

(i) The maximum and average of the block matching error distribution

(ii) The maximum and average of the overall registration error distribution

(iii) The ratio of the Hausdorff distance between the two images when non-rigidly registered to that of the images rigidly registered

(iv) The ratio of the normalized mutual information between the two images when non- rigidly registered to that of the images rigidly registered

We show improvement over rigid registration alone for all patient cases, and significant improvement for cases 1, 2, 5 and 6. We ran registrations using GPU block matching for synthetic cases 1 and 3. Case 1 GPU results are illustrated in figure 3. These results are similar in quality to the results obtained with the traditional co-operative workstation environments. This is significant as it implies that the (far cheaper) GPU is a viable alternative to the co-operative environment.

Clinical Study

Synthetic data in the form of generated deformations that produce synthetic intra-operative images is an important and standard way of assessing the performance of a registration method. Synthetic cases, however,

Figure 3 - Case 3 with GPU block-matching: From left to right: Pre-operative, intra-operative, deformed pre-op, pre-op overlaid on intra-op, deformed pre-op overlaid on intra-op

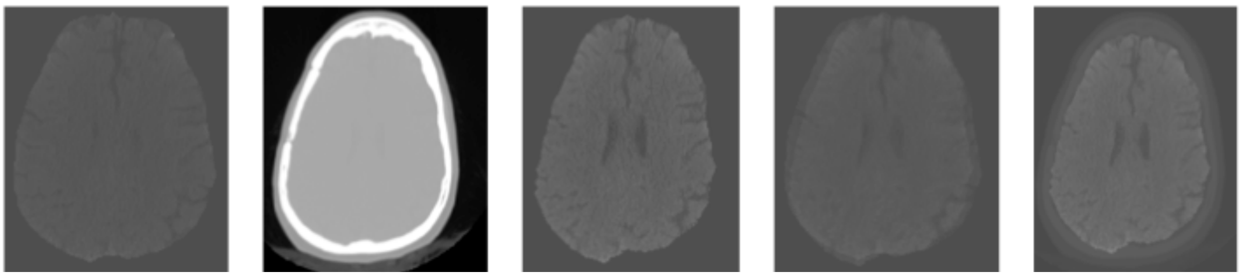

are not sufficient simulators of clinical experience. Real tissue has varying elastic properties (and thus different reactions to the same applied force) that are not effectively simulated. Real intra-operative brain shift is very different than the best synthetically derived deformations. The problem of producing good ground-truth from real intra-operative images, however, is a serious impediment to the evaluation and validation of registration techniques such as ours. This section details the results of registration of three clinical cases using real intra-operative CT, and in one case post-operative CT.

Our pre-operative CT are generally detailed enough and have sufficient spatial resolution to identify – with sub-millimeter accuracy – a number of anatomical landmarks. The intra-operative images we use, however, present a real challenge. Owing to various causes images acquired during surgery suffer from several undesirable phenomena. The intra-operative images we received had a significantly lower signal-to-noise ratio than the pre-operative images of the corresponding patient. A second source of difficulty, which exacerbates the first, is the presence of movement artifacts due to the (unconscious) patient's breathing motion in the device. This creates motion artifact that manifest as streaks or concentric circular shapes. These two phenomena combined make the reliable identification of homologous anatomical landmarks between

Figure 4: Registration error for case 4.

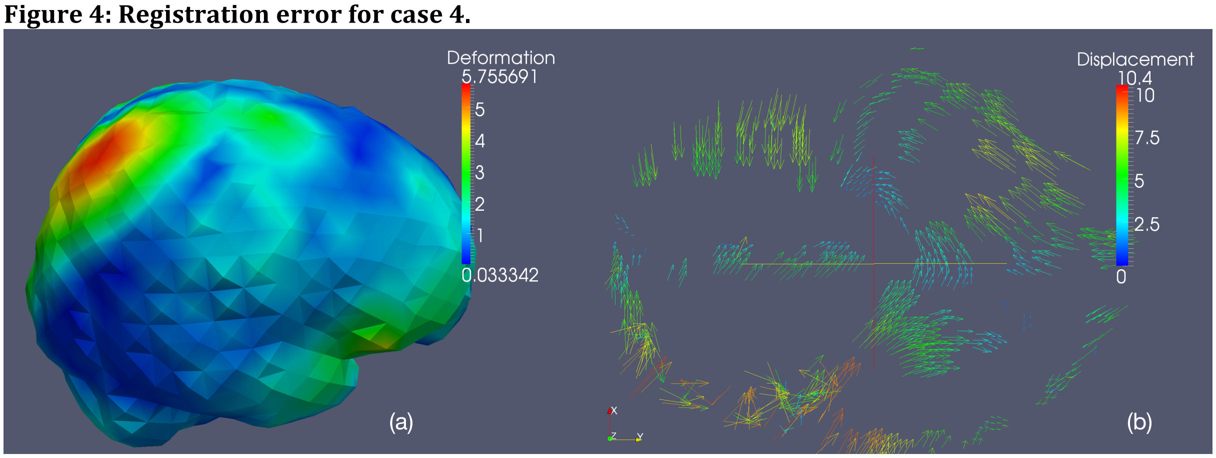

intra-operative and pre-operative images difficult and error-prone, and sometimes impossible. We captured five anatomical landmarks on each set of images. The result of this registration method as compared with their anatomical landmarks is listed in table 2.

Discussion

In this study we have performed several practical feasibility studies of a FEM-based non-rigid registration algorithm using intra-operative CT. This study includes six sets of pre-operative CT and MRI, of which half are supplemented by intra-operative or post-operative CT follow-ups. We used these data to rigorously evaluate the non-rigid registration method. To this end, we developed two mutual-information based metrics to supplement the existing similarity metric, and created a synthetic benchmark to evaluate the performance of these metrics applied to all six cases.

The synthetic benchmarks show that the normalized cross-correlation (NCC) similarity metric based block matching and registration algorithm out-performed both Normalized Mutual Information (NMI)[10,11] and Pointwise Mutual Information (PMI) [12,18]. The NCC registration algorithm achieved an average accuracy improvement of 4.14 times, compared to 1.84 times improvement for NMI, and a 3.48 times average improvement for PMI. This result is not surprising. The mono-modal nature of the CT-to-CT registration allows us to be confident in the assumption of an affine relationship between the intensity distributions of the pre- and intra-operative images. The NMI, while more flexible, is more sensitive to the small block size we use. Our results show that the PMI metric is nearly as accurate as NCC, and clearly does not suffer from the shortcomings of the NMI metric. We suggest that the PMI be evaluated further as a multi-modal block matching metric in this framework.

One of the questions we attempt to answer in this study is whether CT has enough soft tissue definition to successfully complete our relatively sparse displacement-recovering block-matching approach. We believe that we have shown that, given a sufficiently high-resolution pre-operative CT, we can recover a deformation to within 1.39 mm of the actual synthetic value on average, which is much better than rigid registration alone (4.35 mm average error). This study suggests strongly that, given an intra-operative image of similar articulation, we may expect similar accuracy. We conclude that this synthetic study implies that block-matching between a pre-operative and intra-operative CT is possible using either the Normalized Correlation Coefficient or the Pointwise Mutual Information.

We next explored the three cases that present clinically acquired intra-operative CT. The three cases involved correspond to cases 4-6 of the synthetic study. That is, the pre-operative CT for cases 4-6 in the synthetic study are the same images as the pre-operative CT for cases 1-3 in the clinical study. Cases 1 and 2 involve cancerous neoplasms of the frontal and ethmoidal sinus cavities, while case 3 involves a defect of the skull base resulting in a nasoethmoidal encephalocele. In cases 1 and 2 we attempt to register the pre-operative CT with a CT acquired intra-operatively. In case 3 the fixed image is a post-operative CT acquired on a similar scanner as the pre-operative CT. There was no CT contrast present in the images, with the exception of the pre-operative case 3 image. We first note that the pre-operative images in all cases have superior fidelity than the intra- (and post-) operative images. This manifests itself in several ways: the intra-operative images of cases 1 and 2 have a relatively low signal-to-noise ratio and are visually very noisy and grainy (the intra-operative image acquisition delivers a lower dose, which explains this phenomenon). The post-operative scan available for case 3 has a visually similar noise profile to the pre-operative scans, but has a lower soft-tissue articulation. In addition, the post-operative scan has lower resolution in the Z-direction. These are all undesirable properties in an image registration problem. If the noise in an image overpowers the features of the image, any similarity metric we use will perform poorly. The registration is also sensitive to the resolution of the input images, as it will be unable to detect block matchings that reside in-between the available slices. In practice we increase the resolution of the registered image by interpolating at points between the available data. This interpolation will obviously lead to inaccuracies as the interpolation nodes (available image slices) have larger slice spacing. The fidelity of a medical image such as those we deal with

Figure 4

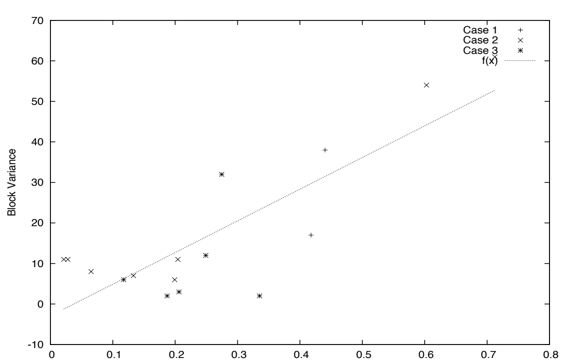

Figure 5

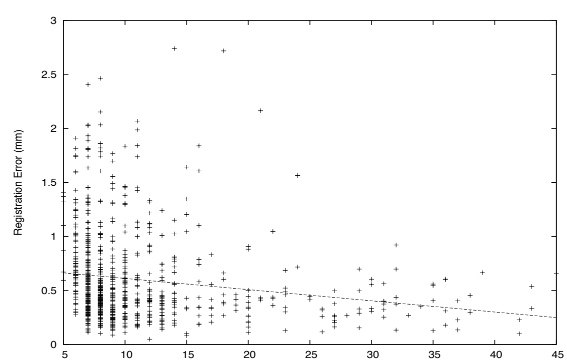

depends on many scan parameters, and setting those parameters is a problem of mutual accommodation of sometimes competing goals: there are many outcomes to optimize.

The landmark acquisition for the three clinical cases was difficult. The quality of the intra-operative scans made finding reliable landmark pairings a hard task. Some of the landmarks captured have a relatively high confidence: the ventricles are an easy feature to discern, even in a very bad image, and some of the landmarks (particularly near the tumor of interest) lie in relatively ct-homogenous tissue. That said, the difficulty of landmark matching in CT is an elocution of the motivation for our study. Nevertheless, expert-acquired landmarks are the best ground truth available to us.

Our results for the clinical study show that more work remains to be done to use this registration method with intra-operative CT. The improvement in landmark error (measured at each individual pair of landmarks) was between 3.5 times better and 4.18 times worse than rigid registration, with an average improvement across all landmarks of only 1.19 times, or 19%. That said there is a correlation between the block variance around an individual landmark and the registration accuracy at that landmark. The higher the block variance is at a landmark, the higher the relative improvement is in registration accuracy. This relationship is described in figure 4.

Figure 6

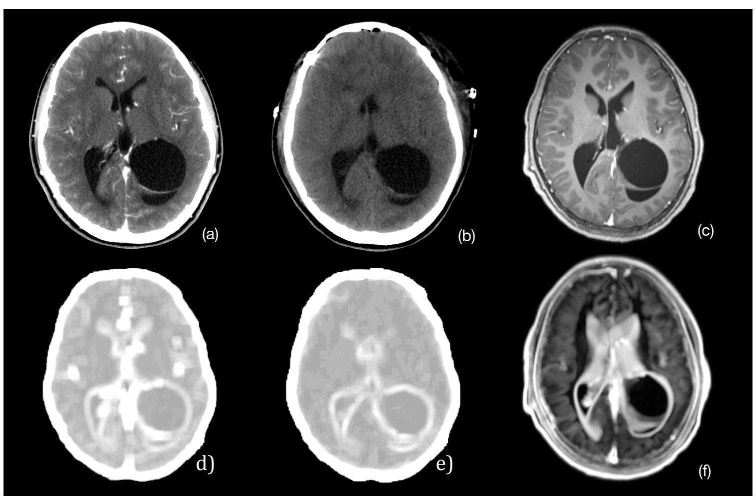

We ran a similar study of a synthetic deformation, which show the same trend in figure 5: increased block variance yields lower registration error. There are two relevant sources of the variance in our images. First is the variance of a block due to the presence of an anatomical or functional feature in the image. We call this source the feature variance. The second source of variance in a block is noise, either random quantum mottle or artifact noise, as previously described. We call this noise variance. Refer to FIGURE 6 for the pre-op CT, pre-op MRI, post-op CT, and block variance images for each of these images from case 6. These block variance images are produced by calculating the variance of an isotropic block of 3-dimensional centered at each voxel coordinate. We see that in the pre-operative CT (a, with contrast), the block variance (d) is largely coincident with the edge features in the image. The ventricles, large cyst, sulci, and midline are all displayed in the block variance image. In addition, regions that are largely structurally void in the CT are largely homogenous in the block variance image. This is a characteristic conducive to high registration accuracy. On the other hand, in the intra-operative image (b), the intra-operative block variance image (e) shows less feature variance and more noise variance than the pre-operative images. In other words, the variance pattern largely does not coincide with structural edges or features in the image, but rather the dominant source of variance is noise. This strongly suggests a hypothesis for the relatively poor performance of the clinical study: the intra-operative images used have poor articulation of soft tissues, and low signal-to-noise ratio, which causes block matching to generate a large number of erroneous matches, or outliers. Contrast this with the pre-operative MRI (c) and its block variance image (f), which shows a much larger range of feature variance, and less homogenous and noisy variance regions. The same is illustrated in FIGURE 7 where we see the distribution of block variance is much broader in MRI (red) than CT (green). We draw the conclusion that the major limiting factor in the improvement of this non-rigid registration method is the acquisition of superior quality intra-operative CT scans. The dual of this problem is the development of noise-robust similarity metrics, and similarity metrics that perform well in the absence of highly articulated soft tissue. Alternatively, it would be worth investigating to extend Andronache et. al.'s method to detect structureless patches to choose registration points, rather than simply choosing the highest variance blocks. Andronache uses Moran's spatial autocorrelation coefficient to do exactly this task [cite]. Alternatively, it may be beneficial to use edge detection images to weight possible feature points by their inclusion or proximity to a detected edge.

Figure 7

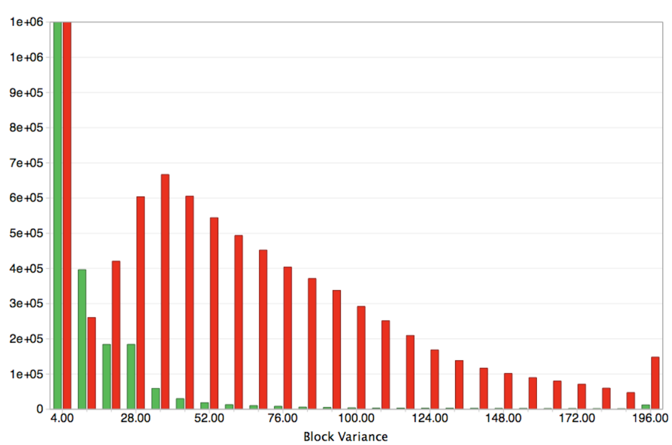

Conclusion

The studies performed in this paper are a good first step toward the advancement of our state-of-the-art non-rigid registration algorithm from mono-modal MRI-to-MRI to multi-modal CT-to-MRI. The objective of achieving superior performance with intra-operative CT is important for reducing the cost and increasing the accessibility of real-time non-rigid registration for image guided neurosurgery. We believe the studies included here, particularly the synthetic studies, show the feasibility of this goal.

Table 5.2

Case	Metric	Matching error		Registration error		SHD	NMI Ratio	Avg
		max	avg	max	avg			Improvement
1	CC	**17.92**	**0.92**	**6.34**	**1.02**	1.49	8.84	4.27
	NMI	18.97	5.05	10.16	2.96	1.91	1.81	1.47
	PMI	**12.85**	**0.66**	**6.38**	**1.16**	1.51	8.60	3.76
2	CC	**20.10**	**1.68**	**6.94**	**1.56**	1.66	9.33	3.86
	NMI	20.64	6.59	12.98	3.71	1.93	1.56	1.62
	PMI	**18.60**	**2.66**	**12.71**	**1.65**	1.58	7.56	3.65
3	CC	**15.03**	**1.12**	**8.49**	**2.62**	1.94	3.16	1.47
	NMI	15.78	5.28	9.59	3.04	2.02	1.23	1.26
	PMI	**16.91**	**1.87**	**8.79**	**2.77**	1.94	2.84	1.39
4	CC	**11.75**	**0.60**	**4.80**	**1.36**		1.60	1.80
	NMI	15.03	4.03	9.77	2.06		1.25	1.19
	PMI	**11.70**	**0.37**	**5.04**	**1.52**		1.55	1.61
5	CC	**11.70**	**0.52**	**2.74**	**0.58**	2.50	3.80	8.58
	NMI	16.03	4.29	5.68	1.13	2.59	2.50	3.54
	PMI	**10.05**	**0.40**	**2.83**	**0.67**	2.49	3.65	7.15
6	CC	**22.23**	**2.73**	**5.76**	**1.23**	1.62	1.72	4.86
	NMI	22.69	6.26	15.91	4.34	1.98	1.44	2.01
	PMI	**21.98**	**2.59**	**7.96**	**2.89**	1.61	1.71	3.32

Table 5.3: A summary of the results for the clinical study including the magnitude of expert-derived deformation, the rigid registration landmark error, and the non-rigid registration landmark error.

Case	Point	True Deformation (mm)	RR Error	NRR Error (mm)		
				CC	NMI	PMI
4	1	(-0.43,-0.87, 1.25)	1.59	0.89	1.81	0.71
	2	(-0.43,6.54, -2.50)	7.02	7.39	7.71	9.59
	3	(-0.87, 0.00,-0.50)	0.87	3.64	4.72	2.56
	4	(-3.49, -3.06, -3.75)	5.96	3.47	2.45	5.99
	5	(-1.75,-0.45, 0.00)	1.80	2.10	0.79	2.40
	6	(-6.55, -0.87, 0.00)	6.60	6.81	5.41	7.26
	7	(0.43, 0.87, 0.00)	0.98	1.11	2.48	2.56
	8	(0.44, 1.30, 0.00)	1.38	3.55	6.06	3.49
	9	(-0.44, 1.31, 1.25)	1.86	4.33	6.66	3.67
5	1	(-1.31, 0.87, -1.00)	1.86	1.82	1.60	2.95
	3	(-0.44,1.74, 3.00)	3.50	1.39	4.57	1.73
	4	(-0.44,0.43, 3.00)	3.06	0.88	2.53	1.39
	5	(-3.49, 0.87,1.00)	3.73	4.46	4.20	5.25
	6	(-0.87, 4.37, 2.00)	4.88	4.23	4.83	4.39
	7	(2.18, 4.36, 2.00)	5.27	4.22	5.12	3.76
	8	(0.44, 3.49, 0.00)	3.52	3.29	4.56	3.21
	9	(0.00,4.80, 0.00)	4.80	4.98	6.28	4.62
	10	(-0.87, 3.06, 0.00)	3.18	2.53	2.99	2.72
	11	(0.44, 5.23, 1.00)	5.35	5.20	6.09	5.12
6	1	(-3.93 2.18,-5.00)	6.72	5.92	4.58	6.04
	2	(1.74,3.05, 2.50)	4.32	2.66	6.26	4.36
	3	(2.18,-0.44, 2.50)	3.35	2.51	8.04	5.74
	4	(1.75,0.00, 2.50)	3.05	2.19	4.91	3.29
	5	(-1.75, 5.67, 5.00)	7.76	6.74	10.69	7.88
	6	(1.75, 4.36, 5.00)	6.86	4.60	7.17	4.79
	7	(0.00,-1.31, -2.50)	2.82	3.44	2.43	3.38
	8	(1.74, 3.49, 0.00)	3.90	2.98	5.92	2.94
	9	(5.24, 2.62, 0.00)	5.85	6.25	6.80	7.61

References

[1] Lisa Gottesfeld Brown. A survey of image registration techniques. ACM Comput. Surv., 24(4):325–376, 1992.

[2] C. Nimsky, O. Ganslandt, S. Cerny, P. Hastreiter, G. Greiner, and R. Fahlbusch. Quantification of, visualization of, and compensation for brain shift using intraoperative magnetic resonance imaging. Neurosurgery, 47(5):1070– 9; discussion 1079–80, Nov 2000.

[3] Neculai Archip, Olivier Clatz, Stephen Whalen, Dan Kacher, Andriy Fedorov, Andriy Kot, Nikos Chrisochoides, Ferenc Jolesz, Alexandra Golby, Peter M. Black, and Simon K. Warfield. Non-rigid alignment of pre- operative mri, fmri, and dt-mri with intra-operative mri for enhanced visualization and navigation in image-guided neurosurgery. NeuroImage, 35(2):609 – 624, 2007.

[4] O. Clatz, H. Delingette, I.F. Talos, A.J. Golby, R. Kikinis, F.A. Jolesz, N. Ayache, and S.K. Warfield. Robust non-rigid registration to capture brain shift from intra-operative MRI. In IEEE Transactions on Medical Imaging, volume 24, pages 1417–1427, 2005.

[5] L. Ibanez, W. Schroeder, L. Ng, and J. Cates. The ITK Software Guide. Kitware, Inc. ISBN 1-930934-15-7, http://www.itk.org/ItkSoftwareGuide.pdf, second edition, 2005.

[6] Yixun Liu, Andriy Fedorov, Ron Kikinis, and Nikos Chrisochoides. A data parallel non-rigid medical image registration method on cooperative architecture. Parallel Computing, preprint.

[7] Andriy Fedorov and Nikos Chrisochoides. Tetrahedral mesh generation for non-rigid registration of brain mri: Analysis of the requirements and evaluation of solutions. In International Meshing Roundtable, number 17th. Springer Verlag, October 2008.

[8] Andriy Fedorov, Nikos Chrisochoides, Ron Kikinis and Simon K. Warfield. Tetrahedral mesh generation for medical imaging. 8th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2005), 2005.

[9] Andriy Y. Fedorov. Enabling Technology for Non-Rigid Image Registration during Image-Guided Neurosurgery. PhD thesis, The College of William & Mary, 2009.

[10] T. Gaens, F. Maes, D. Vandermeulen, and P. Suetens. Non-rigid multi-modal image registration using mutual information. In MICCAI '98: Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Computer Science. Jan 1998.

[11] W. Wells, P. Viola, H. Atsumi, and S. Nakajima. Multi-modal volume registration by maximization of mutual information. Medical Image Analysis, Jan 1996.

[12] P. Rogelj and S. Kovacic. Point similarity measure based on mutual information. Lecture Notes in Computer Science. Jan 2003.

[13] Nikos Chrisochoides, Andriy Fedorov, Andriy Kot, Neculai Archip, Peter Black, Olivier Clatz, Alexandra Golby, Ron Kikinis, and Si- mon K. Warfield. Toward real-time, image guided neurosurgery using distributed and grid computing. In IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (Supercomputing), 2006.

[14] Yixun Liu, Andriy Fedorov, Ron Kikinis, and Nikos Chrisochoides. Real-time non-rigid registration of medical images on a cooperative parallel architecture. In IEEE International Conference on Bioinformatics Biomedicine, November 2009.

[15] Eric Billet, Andriy Fedorov, and Nikos Chrisochoides. The use of robust local Hausdorff distances in accuracy assessment for image alignment of brain mri. The Insight Journal, January 2008.

[16] Marvin P. Fried, Sanjay R. Parikh, and Babak Sadoughi. Image-guidance for endoscopic sinus surgery. The Laryngoscope, 118(7):1287–1292, 2008.

[17] Stephen J. Hentschel and Raymond Sawaya. Optimizing outcomes with maximal surgical resection of malignant gliomas. Cancer Control, 10(2), 2003.

[18] D. B. Russakoff, C. Tomasi, T. Rohlfing, and C. R. Maurer Jr. Image similarity using mutual information of regions. In Computer Vision - ECCV 2004, volume 3023 of Lecture Notes in Computer Science, pages 596–607. 2004.