

To appear in *Computer Methods in Biomechanics and Biomedical Engineering*
Vol. 00, No. 00, Month 20XX, 1–20

Accurate and Fast Deformable Medical Image Registration for Brain Tumor Resection using Image-Guided Neurosurgery

Fotis Drakopoulos and Nikos P.Chrisochoides*

Department of Computer Science, Old Dominion University, Norfolk, VA, USA

(Received 00 Month 20XX; accepted 00 Month 20XX)

We present a Parallel Adaptive Physics-Based Non-Rigid Registration framework for aligning pre-operative to intra-operative brain Magnetic Resonance Images (MRI) of patients who have undergone a tumor resection. This framework extends our earlier work on the Physics-Based methods by using an adaptive, multi-material, parallel Finite Element (FE) biomechanical model to describe the physical deformations of the brain. Our registration technology incorporates fast image-to-mesh convertors for remeshing the brain model in real-time eliminating the poor quality elements; various linear solvers to accurately estimate the volumetric deformations; efficient block-matching techniques to compensate for the missing/unrealistic matches induced by the tumor resection.

Our evaluation is based on six clinical volume MRI data sets including: (i) isotropic and anisotropic image spacings, and (ii) partial and complete tumor resections. We compare our framework with four methods: a Rigid and BSpline deformable registration implemented on 3D Slicer v4.4.0, a Physics-Based non-rigid registration available on ITK v4.7.0, and an Adaptive Physics-Based non-rigid registration. We show that the proposed technology provides the finest MRI alignments among all the methods. The Hausdorff Distance is on average up to 3.78 and 3.12 times more accurate compared to the rigid and the other non-rigid registration methods, respectively. Additionally, it brings the end-to-end execution within the real-time constraints imposed by the neurosurgical procedure. In a Linux Dell workstation with 12 Intel Xeon 3.47 GHz CPU cores and 96GB of RAM, it registers the anisotropic volume data in less than 93 seconds and the isotropic data in less than 21 seconds.

Keywords: non-rigid registration; ITK; FEM; biomechanical model; adaptivity;

1. Introduction

Image registration is generally concerned with spatial alignment of corresponding features between two or more images. During registration, a spatial transformation is applied to one image, which is called moving, such that it is brought into alignment with the other image, which is called fixed. In the context of Image-Guided Neurosurgery (IGNS), the moving image corresponds to a brain pre-operative MRI, which is aligned with the position of the patient before the surgery using a rigid transformation, and the fixed image corresponds to an intra-operative MRI (iMRI) acquired during the surgery. After the opening of the skull and dura, the brain changes its shape because of the cerebrospinal fluid leakage, the effect of gravity, and the administration of osmotic diuretics, introducing discrepancies in relation to the pre-operative configuration. The Non-Rigid Registration (NRR) can usually compensate for the brain shifts caused by all these factors.

However, in the case of a tumor resection, the NRR becomes an ill-posed problem because: (i) not all features in the moving image have a correspondence in the fixed image; (ii) the discrepancies after a rigid alignment are large because of the tissue removal; (iii) the NRR requires significant computing resources and time. In this paper, we focus on all these aspects and propose a framework

*Corresponding author. Email: nikos@cs.odu.edu

to address one of the most challenging problems of NRR for IGNS: brain deformation due to tumor resection. The other challenge is the tissue retraction (Miga, Roberts, et al. (2001); Platenik, Miga, et al. (2002)), which is outside of the scope of this paper.

Modeling the behavior of the brain remains a key issue in providing navigation for IGNS. In our previous work Liu, Kot, et al. (2014), we developed ITK¹ filters for physics-based non-rigid registration (PBNRR), for brain shift only, which take into account the tissue properties, improve the accuracy compared to the rigid registration, and reduce the execution time using GPU and multi-core accelerators. The purpose of this method is to use a known displacement vector associated with sparse feature points inside the cranial cavity, to estimate the entire brain deformation with a brain biomechanical model (Clatz, Delingette, et al. (2005)). In this paper we augment the PBNRR method by employing an adaptive heterogeneous biomechanical model for tissue/tumor removal depicted in the iMRI and exploiting additional parallelism to satisfy the time constraints of the neurosurgical procedure.

Warfield, Ferrant, et al. (2000) presented a two step method for the biomechanical simulation of volumetric brain deformations, achieving reasonably accurate image alignments within the IGNS time constraints. In the first step, an active surface algorithm iteratively deforms the surface of the first brain volume to match that of the second volume (Ferrant, Cuisenaire and Macq (1999)). In the second step, the volumetric brain deformation implied by the surface changes is computed in parallel via a biomechanical model. The method was evaluated on two neurosurgery cases. In our work, we match image blocks located in the entire volume of the cranial cavity instead of only the brain surface. The selection and the matching of the blocks, as well as the volumetric brain deformations, are all computed in parallel.

A prospective study on NRR of pre-operative data (T1,fMRI, DTI) with intra-operative data (T1) was reported in Archip, Clatz, et al. (2007). The clinical cases obtained from eleven patients enrolled over twelve months. The computationally intensive components of this method were parallelized in Chrisochoides, Fedorov, et al. (2006) with a cluster of 300 workstations, and the image alignments were calculated in less than seven minutes. In this paper, our method is implemented for shared memory multiprocessor architectures, and the image alignments are computed between 21 and 93 seconds in a Dell workstation with twelve cores and 96GB of RAM.

In summary, these methods and others (Mostayed, Garlapati, et al. (2013); Farnia, Ahmadian, et al. (2015)) compensate only for small brain deformations (shifts). The complex neurosurgical procedure of the tumor resection, which invalidates the biomechanical model defined on the pre-operative MRI and compromises the fidelity of the IGNS, is not addressed.

Risholm, Melvr, et al. (2009) presented an adaptive FE multi-level grid registration scheme that accommodates a superficial tumor resection. The method uses a planar FEM grid and a Delaunay mesh split strategy. It was evaluated on 2D medical and synthetic images. In our work, rather than splitting the elements, we generate in parallel a new global mesh from a warped segmented image. Our parallel refinement scheme eliminates the poor quality elements with very small ($< 5^\circ$) and very large ($> 170^\circ$) dihedral angles occurring after each model deformation and accurately captures the complex morphological changes imposed by the tumor resection.

Gu and Qin (2009) proposed a global-to-local NRR scheme to compensate for the brain shift and unmatched outliers caused by the tumor resection. Initially the method aligns the images globally using a mutual information metric enhanced by a joint saliency map (JSM). Then it uses paired pools of scale-invariant keypoints in combination with the JSM to guide a local control-point correspondence detection and an outlier keypoint rejection. The algorithm was evaluated on five 2D brain tumor resection MRI data sets. Given the mono-modal (MR-T1 weighted) nature of the registration problem in this paper, it is reasonable to assume that there is an affine relationship between the two image intensity distributions. However, for completeness, in addition to a Normalized Cross Correlation (NCC) metric, which is suitable for matching images from the same modality, we use a Normalized Mutual Information (NMI) metric, which is traditionally used in

¹<http://www.itk.org/>

the case of a multi-modal registration. Besides, we use a two-step block-matching technique to eliminate the blocks with a missing target correspondence, and a parallel outlier rejection module that removes the blocks with the larger error between the matching displacements and the computed mesh deformations. The percentage of the selected and the rejected blocks can be adjusted by the user. In addition, we evaluate our framework on clinical 3D MRI data.

Periaswamy and Farid (2006) presented a robust Expectation-Maximization (EM) framework for the simultaneous segmentation and registration of volume MRIs with partial or missing data. A Matlab implementation of this method required about 30 minutes to register a pair of $64 \times 64 \times 64$ volumes on a 2.8GHz Linux machine. In this paper, we evaluate the performance of our method on larger adult brain MRI with anisotropic or isotropic spacings (Table 3). We show that on a Linux Dell workstation with twelve Intel Xeon 3.47 GHz CPU cores, our method accurately aligns a pair of anisotropic volume MRIs in less than 93 seconds.

Our previous work (APBNRR) Drakopoulos, Liu, et al. (2014) extended the PBNRR Liu, Kot, et al. (2014) to compensate for the brain deformations induced by a tumor resection using a parallel adaptive heterogeneous biomechanical FE model for tissue/tumor removal depicted in the iMRI. The APBNRR was evaluated on twelve patients and reduced the alignment error on average 2.57 times compared to the PBNRR; however, the evaluation was conducted only on isotropic volume MRI data and the parallelization fraction of the program was about 70%. The proposed method is evaluated on both isotropic and anisotropic clinical MRI data and achieves a higher parallelization fraction ($\approx 98\%$).

In summary, this paper presents a fast and accurate Parallel Adaptive Physics-Based Non-Rigid Registration framework to compensate for the brain deformations induced by a tumor resection. Our registration scheme computes a sparse displacement field associated with image blocks inside the cranial cavity, and relies on an adaptive, multi-material, parallel FE brain biomechanical model to estimate a dense deformation field for image warping. The three major problems of the NRR of brain tumor resection MRI data are addressed as follows:

- (1) The problem of missing/unrealistic matches between blocks in the moving and the target MRI is addressed with a two-step block-matching scheme. The first step searches for a position in the target image that maximizes a similarity metric for a block in the moving image. The second step matches the blocks without a target correspondence to their closest point on the surface of the target brain object. The final matching displacements are scaled and the outlier blocks are rejected.
- (2) The problem of large image discrepancies is addressed with an adaptive FE biomechanical model. The model is remeshed from a warped segmented image to capture not only the brain deformations but also the complex geometric changes nearby the tumor margins, while maintaining throughout the process meshes with good quality elements –critical for the accuracy and convergence rate of the solver.
- (3) The problem of real-time constraints imposed by IGNS is addressed with the parallelization of all the computational-intensive components of the framework.

We evaluate our method qualitatively and quantitatively on six clinical volume isotropic and anisotropic MRI data sets, including partial and complete tumor resections. We compare this framework with four other methods: (i) a Rigid registration, (ii) a BSpline deformable registration implemented on 3D Slicer v4.4.0 (Johnson, Harris and Williams (2007)), (iii) a Physics-Based non-rigid registration available on ITK v4.7.0 (Liu, Kot, et al. (2014)), and (iv) an Adaptive Physics-Based non-rigid registration (Drakopoulos, Liu, et al. (2014)).

The evaluation results indicate that the proposed scheme is the most accurate among all the methods and satisfies the hard time constraints of the neurosurgical procedure. In the rest of the paper we will describe the framework and outline all its modules in more detail.

2. Method

Figure 1 illustrates the proposed framework. The method iteratively estimates a dense deformation field by incrementally and accurately incorporating small changes in the geometry of the domain (brain depicted in the iMRI) resulted by a partial, complete or extensive tumor resection during image-guided neurosurgery. The computation of the dense deformation field is facilitated by a sparse displacement vector associated with highly discriminant blocks inside the cranial cavity, as well as a heterogeneous FE biomechanical model (Figure 2) that describes the physical deformation of the brain. During the deformable (or non-rigid) registration, the model is deformed to reflect the changes in the brain morphology caused by the tumor resection, gravity, and other factors. After each deformation, the quality of the elements deteriorates, thus the model is globally remeshed from a warped segmented image to avoid the heavily distorted or flipped elements. PAPBNRR uses an additive deformation field for image warping. As a result, the input images are interpolated only once, independently of the number of iterations. The model deformation and the image warping stop when the number of image blocks without a correspondence in the intra-operative image is less than a threshold, or when the maximum number of iterations has been reached (Figure 1).

In the remainder of this section, we describe the segmentation of the pre-operative MR image, and the modules in Figure 1, which altogether improve on average by 30% the accuracy and by 66% the performance of the method we initially presented in Drakopoulos, Liu, et al. (2014).

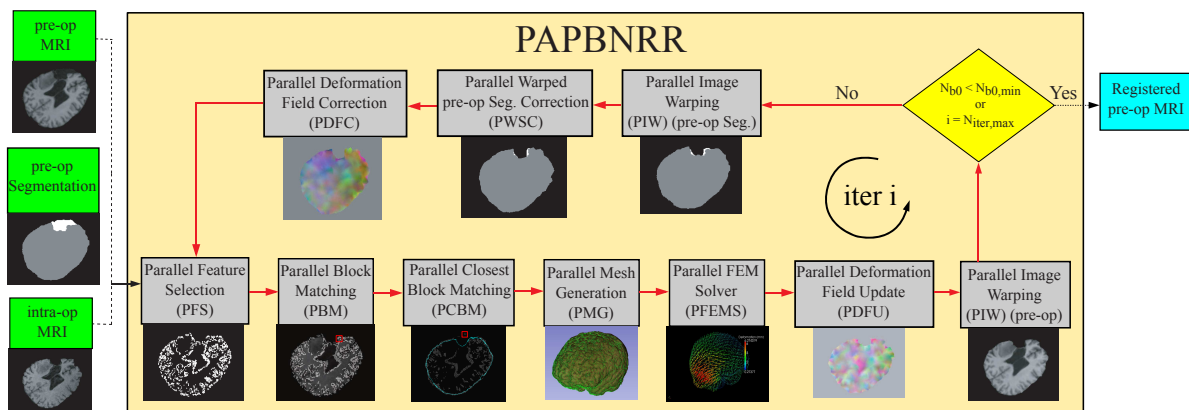


Figure 1. The PAPBNRR framework. Green, cyan, and gray represent the input, the output, and the modules of the framework, respectively. The modules are built on the ITK open-source system and are parallelized with the POSIX thread library for shared memory multiprocessor architectures. The red arrows show the execution order of the different modules in the loop. The loop breaks when the number of blocks with zero displacements (N_{b_0}) is less than a minimum ($N_{b_0,min}$), or when the maximum number of iterations ($N_{iter,max}$) has been reached.

2.1 MRI Segmentation

The pre-operative image is usually acquired a week before the first intra-operative image, thus, any computational requirements of the pre-operative segmentation are easily satisfied. Before segmenting the brain parenchyma and tumor tissues depicted in the pre-operative image, we apply a skull stripping algorithm that removes the skull from the image and isolates the intra-cranial cavity. We employ the Brain Extraction Tool (BET) for this purpose (Smith (2002)). Segmenting the intra-cranial cavity is a semi-automatic process. Automatic segmentation is typically performed with thresholding, region growing and connectivity operators, which select regions based on similar voxel intensities may be in continuous regions. In this study, we employed a combination of automatic operators implemented in 3D Slicer, like region growing and level-set filters (Antiga, Piccinelli, et al. (2008)), with slice-by-slice manual segmentation to correct any erroneously included regions.




Figure 2. A heterogeneous brain FE biomechanical model visualized with 3D Slicer. Green and red represent the surface of the brain parenchyma and tumor sub-mesh, respectively. Top row: the generated mesh superimposed on the (cropped) volume MRI. Bottom row: axial, sagittal and coronal slice intersections between the MRI and the mesh surfaces.

Our registration framework relies on a warped pre-operative segmentation to: (i) restrict the selection of image blocks with high variance inside the intra-cranial cavity (subsection 2.2), and (ii) generate a multi-tissue mesh (subsection 2.4) which is the basis of a patient-specific, heterogeneous biomechanical model, that describes the physical deformation of the brain (subsection 2.5). Therefore, the accuracy of the pre-operative segmentation is crucial for the correctness of the produced transformation, which is used to update the pre-operative image, and deliver it to the surgeon during the procedure. However, an evaluation of the registration based on the segmentation accuracy is beyond the scope of this paper, and it will be considered in our future work.

2.2 Parallel Feature Selection (PFS)

The input of the module is an image and a mask: as an input image we use the warped pre-op MRI of iteration i ; as an input mask we use the corrected warped segmentation of iteration i (see PWSC). At the beginning of iteration $i = 1$, the input image is equal to the pre-op MRI, and the input mask is set to be the pre-op segmentation. The output of the module is a set of features in the input image with the highest intensity variance in the surrounding region. Each feature is a block of voxels restricted inside the mask. Parameter F_s (Table 5) determines the percentage of the selected blocks from the total number of image blocks. Parameter B_s (Table 5) determines the size of the block in the axial, coronal, and sagittal image directions.

Three non-connectivity patterns are used (Liu, Kot, et al. (2014)), to control the distribution of the selected blocks in the image. The “vertex” (26 non-connectivity), “edge” (18 non-connectivity), and “face” (6 non-connectivity) patterns, prevent the selection of blocks which are connected with neighboring blocks via a vertex, an edge, and a face, respectively. Figure 3 illustrates how the non-connectivity patterns influence the selection of blocks in a brain MRI. During a brain tumor resection procedure, we anticipate large volumetric deformations nearby the resection margins, thus, in our experimental evaluation we use a “face” non-connectivity pattern (Table 5). This pattern results in a higher block density nearby the tissue boundaries/interfaces and a lower block density in the rest of the brain regions (Figure 3). The parallelization of the module has been described in Drakopoulos, Liu, et al. (2014).




Figure 3. The distribution of the selected blocks in a brain MRI, using different non-connectivity patterns. The results are depicted on six consecutive MRI slices. From top to bottom row: axial pre-op MRI slice (left) and volumetric rendering (right), selected blocks with “vertex” non-connectivity, selected blocks with “edge” non-connectivity, selected blocks with “face” non-connectivity. The “vertex” pattern results in a more uniform distribution, while the “face” pattern results in a higher block density nearby the tissue boundaries/interfaces (i.e., intracranial brain, ventricles and tumor surfaces).

2.3 Parallel Block Matching (PBM) - Parallel Closest Block Matching (PCBM)

Considering a block in the pre-operative image and a window W_s (Table 5), the block matching algorithm searches for a position inside the window in the intra-operative image that maximizes a similarity measure. By assembling the individual block displacement vectors computed from this step, one can create a sparse displacement field, which is used by the FEM Solver (see PFEMS) to estimate the unknown mesh displacements U associated with the mesh vertices.

Similarity measures in registration include Mean Square Difference of intensity (MSD), Mutual Information (MI), and Cross Correlation (CC) (Brown (1992)). In this paper, we enhance the PBM module described in Liu, Kot, et al. (2014), introducing a NMI metric (Studholme, Hill and Hawkes (1999)). The value of the metric is normalized between zero and one. A value close to one indicates that two blocks are very similar while a value close to zero indicates that the blocks are very different. After partitioning the selected blocks to a specifying number of hardware threads, we maximize the following equation (Studholme, Hill and Hawkes (1999)):

$$NMI(B_{pre}, B_{intra}) = 1 - \frac{H(B_{pre}, B_{intra})}{H(B_{pre}) + H(B_{intra})} \quad (1)$$

where $H(B_{pre})$ and $H(B_{intra})$ are the marginal entropies of a pre-operative and an intra-operative image block, respectively, and $H(B_{pre}, B_{intra})$ is the joint entropy of the two blocks. $0 \leq NMI(B_{pre}, B_{intra}) \leq 1$ because $H(B_{pre}) + H(B_{intra}) \geq H(B_{pre}, B_{intra})$. The maximization of

(1) seeks a transformation where joint entropy is minimised with respect to the marginal entropies.

The PBM approach works well only in the case of a brain shift. In the case of a tumor resection, the blocks might not have a target correspondence inside the specified window. Besides, a very large window is opposed to the assumption that a complex non-rigid transformation can be approximated by point-wise translations of small image regions, often leading to: a) unrealistic matches that deteriorate the quality of the produced transformation and b) performance degradation due to an exhaustive search on a large number of image blocks.

The PCBM module uses a Euclidean distance point metric to match the blocks with a missing correspondence to their closest point on the surface \mathcal{S} of the target brain object. Figure 4 illustrates such an example. Initially \mathcal{S} is extracted from the intra-operative image. Then, for each block i




Figure 4. The PCBM module on a pair of brain MRIs. Left: pre-operative MRI. The white points on the top right of the left figure represent the blocks without a correspondence. Point A represents one of those blocks. Right: intra-operative MRI. \mathcal{S} : the points on the target brain surface; $d_{i,\mathcal{S}}$: Euclidian distance vector from the center of the block to its closest point on \mathcal{S} ; d_i : Euclidian distance vector bounded by the block matching window W .

with displacement $\|d_i\| = 0$, its minimum Euclidean distance to \mathcal{S} is computed; that is, $\|d_{i,\mathcal{S}}\|$ where $\|\cdot\|$ is the Euclidian norm. The maximum displacement due to the block matching window is:

$$\|d_{max}\| = \sqrt{\left(s_x \left\lfloor \frac{W_{s,x}}{2} \right\rfloor\right)^2 + \left(s_y \left\lfloor \frac{W_{s,y}}{2} \right\rfloor\right)^2 + \left(s_z \left\lfloor \frac{W_{s,z}}{2} \right\rfloor\right)^2} \quad (2)$$

where s is the image spacing (Table 3) and W_s the window size (Table 5). The block displacement d_i is restricted inside the window with the equation:

$$d_i \leq \frac{\|d_{max}\|}{\|d_{i,\mathcal{S}}\|} \cdot d_{i,\mathcal{S}} \quad (3)$$

where $\|d_{max}\|/\|d_{i,\mathcal{S}}\| < 1$ and $\|d_{i,\mathcal{S}}\| \neq 0$. In practice, we always choose the upper bound from (3). Figure 4 depicts a block displacement before (blue) and after (green) it is bounded by the block-matching window.

Parameter $N_{b_0,min}$ (Figure 1, Table 5) controls the minimum allowed number of blocks with a missing target correspondence. Our experimental evaluation showed that $N_{b_0,min} = 1\% \cdot N_b$ results to a high registration accuracy within an acceptable number of adaptive iterations. N_b is the number of selected blocks from the PFS module. The PCBM module computes the block displacements in

parallel by assigning $\lfloor N_{b_0}/k \rfloor$ blocks to each one of the $k - 1$ threads and $N_{b_0} - (k - 1) \cdot \lfloor N_{b_0}/k \rfloor$ blocks to the k th thread.

2.4 Parallel Mesh Generation (PMG)

Our framework incorporates an ITK structured hexahedral mesh generator and different multi-material unstructured tetrahedral image-to-mesh convertors developed by our group (Foteinos and Chrisochoides (2013); Chernikov and Chrisochoides (2011); Liu, Foteinos, et al. (2012)). An evaluation of the registration based on different meshing techniques is beyond the scope of this paper and was addressed in our previous work Foteinos, Liu, et al. (2011). In this paper, we employ a parallel Delaunay method (Foteinos and Chrisochoides (2014)) to generate a tetrahedral mesh from a multi-labeled segmented image in real-time, and accurately capture the morphological changes of the brain during the resection. This meshing approach recovers the tissue boundaries defined in the segmented image with geometric and topological guarantees, and meshes the underlying volume with high quality tetrahedra. Each tetrahedron is associated with a single label which is assigned to its centroid using the segmented image (see subsection 2.1).

Due to changes in the geometry of the resected tissue in and around the tumor area, we incrementally approximate the changes in the resected tumor using multiple mesh generation phases. Let us denote the mesh in the beginning of iteration i (Figure 1) with M_i . Initially, a multi-material mesh is generated from the input segmentation. After the completion of iteration i , the displacements are computed on the vertices of M_i , which yields the deformed mesh M'_i . M'_i is not used for the next loop because it consists of distorted elements of poor quality in terms of dihedral angles and aspect ratio. Instead, a corrected warped image segmentation is used in the next iteration as the input of the mesher, which produces in parallel a new global mesh M_{i+1} . Parameter δ (Table 5) determines the size of the mesh ($\delta > 0$). The smaller the δ , the larger the mesh.

Table 1 illustrates how the remeshing eliminates the elements with small and large dihedral angles for an example of 4 adaptive iterations. “Before” corresponds to the deformed mesh at $i - 1$, and the “After” to the new generated mesh at i . The #Tets refers to the number of tetrahedral elements in the new generated mesh at i . Generally, the higher the quality of the elements, the better the conditioning of the stiffness matrices and thus the convergence of the linear solver.

Table 1. The minimum and maximum dihedral angles α , before and after the Parallel Mesh Generation ($\delta = 5$).

Iteration i	#Tets	α_{min}		α_{max}	
		Before	After	Before	After
1	14278	-	5.00°	-	169.68°
2	13482	0.34°	5.24°	179.28°	169.92°
3	13497	0.14°	4.91°	179.79°	169.71°
4	12957	0.10°	5.01°	179.80°	171.32°

2.5 Parallel FEM Solver (PFEMS)

Clatz, Delingette, et al. (2005) originally presented the FEM Solver (FEMS). It estimates the deformations U on the mesh vertices using a hybrid approach that starts with an approximation and produces an interpolation-based formulation. In Liu, Fedorov, et al. (2009), we parallelized this method by employing ParMETIS² to implement a balanced parallel partitioning of the tetrahedral mesh among the processing cores and PETSc (Balay, Abhyankar, et al. (2015)) to assemble and solve the linear system of equations. In Liu, Kot, et al. (2014) we integrated FEMS in ITK using the ITPACK matrix structures (Kincaid, Respass, et al. (1982)). In the parallel implementation of FEMS (PFEMS) (Drakopoulos and Chrisochoides (2014)), we avoid using the ParMetis or PETSc

²<http://glaros.dtc.umn.edu/gkhome/metis/parmetis/overview>

libraries because their integration into the ITK is very cumbersome, and they present maintenance and portability issues.

Figure 5(a) depicts PFEMS. Green represents the input (i.e. mesh, N_b selected blocks, block matching displacements, parameters $E_b, E_t, \nu_b, \nu_t, \lambda, F_r, N_{appr}, N_{int}$ listed in Table 5); gray represents the main steps (i.e. Initializations, Assembly, Outlier Rejection, Interpolation); cyan represents the components of each step; yellow represents the output (i.e. mesh deformations U). K_b, K_m are the $N \times N$ stiffness matrices of the blocks and the mesh, respectively. $K_g = K_b + K_m$ is the $N \times N$ stiffness matrix of the biomechanical model. F is the block displacement vector of size N . $N = 3 \cdot N_n$ where N_n is the number of the mesh vertices.

The parallel outlier rejection loop depicted in Figure 5(a), for each iteration j , removes $(N_b \cdot F_r)/N_{appr}$ blocks with the largest error between the computed mesh deformations and the block matching displacements. N_b is the number of selected image blocks from the PFS module, N_{appr} is the number of outlier rejection steps, and F_r controls the fraction of the rejected blocks (Table 5). Based on our experimental evaluation on clinical adult brain MRIs, we suggest that, in order to provide accurate and fast alignments, one should use the following thresholds for parameters N_b, F_r, N_{appr} :

- $2 \cdot 10^5 \leq N_b \leq 4 \cdot 10^5$, which implies approximately $3\% \leq F_s \leq 10\%$. A fraction $F_s = 5\%$ usually selects a sufficient number of image blocks ($\approx 3 \cdot 10^5$).
- $10\% \leq F_r \leq 40\%$. A fraction $F_r = 25\%$ is sufficient to reject every significant outlier. Values of $F_r > 40\%$ reject some relevant matches, and when $F_r < 10\%$ a considerable amount of outliers might survive.
- $5 \leq N_{appr} \leq 20$, for gradually smoothing the estimated displacement field under the time constraints of the neurosurgical procedure.

In the future we will investigate how other types of image modalities (e.g. Computed Tomography, Ultrasound) influence the values of the above thresholds. Figure 5(b) illustrates the selected blocks before and after the parallel outlier rejection loop in an example with $N_b = 318026$, $N_{appr} = 10$ and $F_r = 10\%, 25\%, 40\%$. The details of the parallel FEM Solver implementation can be found in Drakopoulos and Chrisochoides (2014). In this paper we integrate into PFEMS the:

- (1) VNL storage matrix structures.
- (2) LSQR Paige and Saunders (1982) and BICGSTAB Barrett, Berry, et al. (1994) iterative linear solvers.
- (3) LU decomposition Kenneth (1986) direct linear solver.
- (4) Elements with quadratic interpolation functions and/or Extra Shape Functions (ESF) Bathe (1996).

The ITPACK and VNL storage structures are described in Kincaid, Respass, et al. (1982) and the VNL library³, respectively. The isoparametric linear/quadratic, tetrahedron/hexahedron elements, as well as the hexahedron ESF element, are described in commercial FEM software packages like Ansys⁴ and Abaqus⁵.

Table 2 compares the accuracy and the speed of the available linear solvers for a mildly and a highly non-symmetric sparse linear system of equations $F = K_g \cdot U$ of size 8391×8391 . In the case of a mildly non-symmetric system, the iterative solvers JCG, LSQR, BICGSTAB converge to the same solution $\|U\|$ with the direct solver LU; however, in the case of a highly non-symmetric system, the JCG doesn't provide an accurate solution because its residual norm r is significantly larger than the other norms (Table 2). Generally, the JCG method is not guaranteed to converge in the case of general non-symmetric systems (Kincaid, Respass, et al. (1982)).

Figure 6 illustrates the matrices K_m, K_b, K_g for the two types of systems described in Table 2. K_m

³<https://public.kitware.com/vxl/doc/release/core/vnl/html/>

⁴<http://www.ansys.com/>

⁵<http://www.3ds.com/products-services/simulia/products/abaqus/>

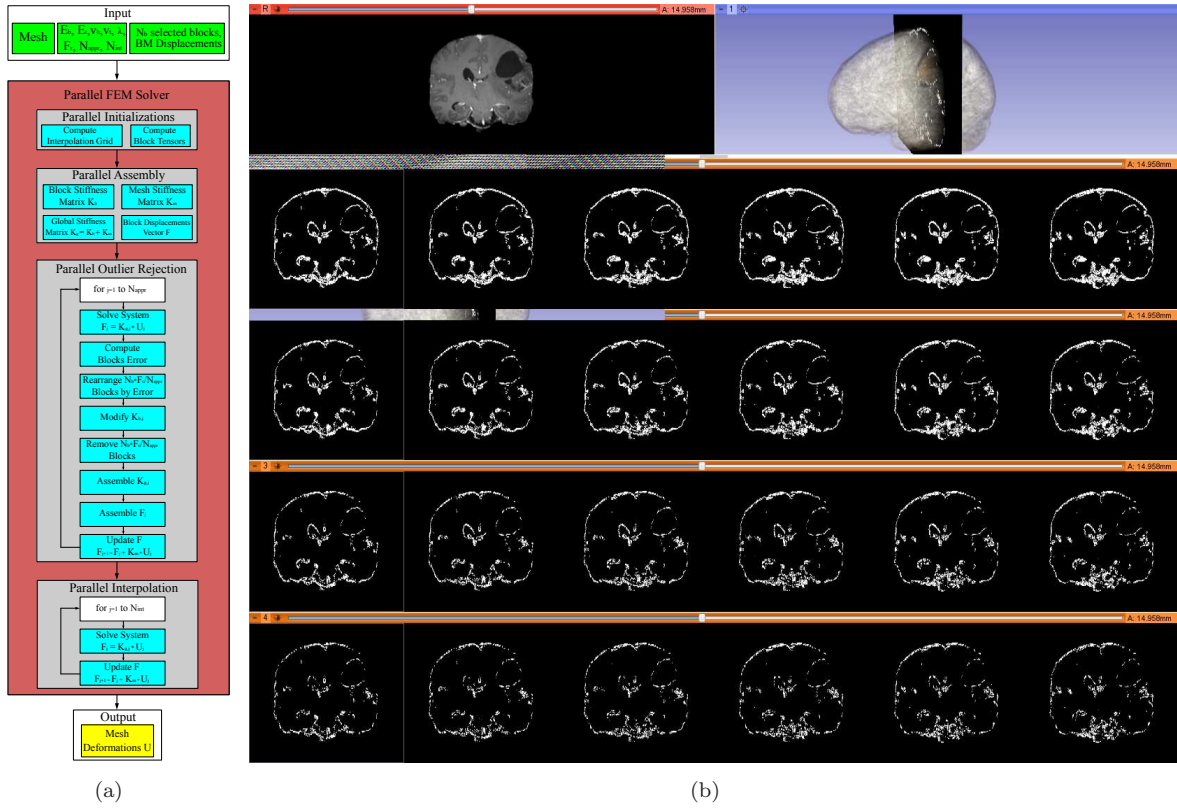


Figure 5. (a): The Parallel FEM Solver (PFEMS). Green represents the input, gray the PFEMS steps, cyan the components of each step, and yellow the output. (b): The selected blocks before and after the Parallel Outlier Rejection step. The image blocks are depicted on six consecutive coronal MRI slices. From top to bottom row: coronal pre-op MRI slice (left) and its position in the volume MRI (right); blocks before rejection ($N_b = 318026$); blocks after rejection with $F_r = 10\%$ (286226); blocks after rejection with $F_r = 25\%$ (238521); blocks after rejection with $F_r = 40\%$ (190816).

is symmetric and independent of the registration blocks. It depends on the geometry-quality of the mesh, the type of the finite elements, and the mechanical properties of the brain materials. K_b, K_g are non-symmetric matrices. Also, they depend on the number and the positions of the selected registration blocks and on the similarity of a block with its corresponding target block. Generally, different values of F_s (Table 2) result in different entries at matrices K_b and K_g (Figures 6(b)-6(c) and 6(d)-6(e), respectively).

In practice, when an adequate number of image blocks is selected ($F_s > 1\%$), the matrix K_g is most likely mildly non-symmetric. The LSQR, BICGSTAB, LU solvers are designed for the general case of non-symmetric linear systems. For the purposes of our registration scheme, the BICGSTAB method seems to be the most suitable choice because: (i) it provides a similar solution with a direct LU solver, and (ii) it converges rapidly (Table 2).

2.6 Parallel Deformation Field Update (PDFU)

The PDFU module takes as an input the mesh deformations U from PFEMS and produces an additive image deformation field. The additive field holds the sum of the previous image fields (iterations $1, 2, \dots, i-1$) and the current image field (iteration i), according to the equation:

$$DF'_i = DF'_{i-1} + DF_i \quad (4)$$

where $i \geq 1$ (Figure 1) and DF'_i, DF'_{i-1} are the additive image deformation fields at iterations i and $i-1$, respectively. DF_i is the produced image deformation field at iteration i . The DF_i computation and the addition at (4) are implemented in parallel in two steps. For the first step,

Table 2. Comparison of the available linear solvers in PFEMS for a sparse $N \times N$ linear system of equations $F = K_g \cdot U$ with N : 8391; K_g : biomechanical stiffness matrix; r : residual norm; $\|U\|$: ℓ^2 norm of the solution vector; F_s : % selected blocks from the total number of blocks; JCG: Jacobi Conjugate Gradient; LSQR: Least Squares; BICGSTAB: Biconjugate Gradient Stabilized; LU: Lower-Upper decomposition with forward elimination and back substitution.

K_g	F_s	Linear solver	r	$\ U\ $	Time (sec)
Mildly non-symm.	5% ^a	JCG	$4.90 \cdot 10^{-6}$	45.530	0.12
		LSQR	$1.02 \cdot 10^{-12}$	45.530	24.03
		BICGSTAB	$8.49 \cdot 10^{-6}$	45.530	0.12
		LU	n/a	45.530	213.50
Highly non-symm.	0.1% ^b	JCG	$6.95 \cdot 10^{-2}$	101.871	0.18
		LSQR	$4.86 \cdot 10^{-5}$	43.134	46.12
		BICGSTAB	$9.89 \cdot 10^{-6}$	43.108	0.12
		LU	n/a	43.136	207.52

^a282123 image blocks

^b5222 image blocks




Figure 6. The stiffness matrix values of the system described in Table 2. K_m : mesh stiffness matrix; K_b : blocks stiffness matrix; $K_g = K_m + K_b$: biomechanical stiffness matrix; F_s : % selected blocks from the total number of blocks.

the deformed mesh is partitioned into k sub-meshes, where k is the number of threads. Each thread then transforms the sub-mesh deformations into image deformations. For the second step, the image fields DF'_{i-1} and DF_i are partitioned into k subfields, and the addition at (4) is performed in parallel.

2.7 Parallel Image Warping (PIW) - Parallel Warped pre-op Segmentation Correction (PWSC) - Parallel Deformation Field Correction (PDFC)

We employ ITK's multi-threaded filter to warp the input image using an additive deformation field and an interpolator to estimate the image values at the non-integer pixel positions. For warping a pre-operative MR image, we use a linear interpolator; for warping a pre-operative segmentation, we use a nearest neighbor interpolator.

The PWSC module takes into account the resected tumor by overlapping two images: a warped segmentation of iteration i (Figure 1) and a mask of the intra-operative image. Figure 7 illustrates

an example of this process where the same representative slice is shown in each subfigure. After overlapping the two images (Figure 7(c)), the portion of the tumor that falls into the mask background is removed (Figure 7(d)). For the parallelization of the module, the segmented image is partitioned into k pieces, and the overlapping/relabeling is performed concurrently by k threads.




Figure 7. The correction of a warped segmentation at iteration i . White and gray represents the tumor and the rest of the brain, respectively. Blue represents the mask. (a): warped segmentation, (b): intra-op mask, (c): (a) and (b) overlapped, (d): corrected warped segmentation.

After PWSC, the additive deformation field at (4) is not consistent with the corrected warped segmentation. For this reason, the values of DF_{add_i} that correspond to the corrected region depicted in Figure 7(c), are set equal to zero. The deformation field correction is performed in parallel after the partition of the field into k pieces, where k is the number of threads.

3. Results

We evaluate the proposed framework qualitatively and quantitatively on six patients (two males, four females). All patients underwent surgery at the Department of Neurosurgery at Shanghai Huashan Hospital Chen, Ying, et al. (2008) for parietal, frontal, and temporal brain tumors. Two cases are Partial Tumor Resections (PTR), and four are Complete Tumor Resections (CTR). Cases 1-3 are anisotropic, while cases 4-6 are isotropic (after resampling). Table 3 lists the MRI data. The original image sizes and spacings (mm) before the resampling are: $448 \times 512 \times 176$ and $0.488 \times 0.488 \times 1.00$ (case4); $512 \times 448 \times 176$ and $0.488 \times 0.488 \times 1.00$ (case 5); $448 \times 448 \times 58$ and $0.496 \times 0.496 \times 1.622$ (case 6). The study was carried out with Institutional Review Board approval.

Table 3. The clinical MRI data of this study. PTR: Partial Tumor Resection; CTR: Complete Tumor Resection; x: axial; y: coronal; z: sagittal.

Spacing Type	Case	Genre	Type	Tumor Location	Image Size (voxels)	Image Spacing $s_x \times s_y \times s_z$ (mm)
Anisotropic	1	M	PTR	parietal	$512 \times 448 \times 176$	$0.488 \times 0.488 \times 1.000$
	2	F	PTR	frontal	$384 \times 512 \times 176$	$0.488 \times 0.488 \times 1.000$
	3	F	CTR	frontal	$448 \times 512 \times 176$	$0.488 \times 0.488 \times 1.000$
Isotropic	4	F	CTR	parietal	$219 \times 250 \times 176$	$1.000 \times 1.000 \times 1.000$
	5	M	CTR	temporal	$250 \times 219 \times 176$	$1.000 \times 1.000 \times 1.000$
	6	F	CTR	temporal	$240 \times 240 \times 134$	$1.000 \times 1.000 \times 1.000$

In this paper we compare the proposed framework (PAPBNRR) with the following methods:

- (1) A Rigid registration implemented on 3D Slicer v4.4.0 BRAINSFit module (Johnson, Harris and Williams (2007)).
- (2) A BSpline deformable registration implemented on 3D Slicer v4.4.0 BRAINSFit module (Johnson, Harris and Williams (2007)).
- (3) A Physics-Based Non-Rigid Registration (PBNRR) implemented on ITK v4.7.0 (Liu, Kot, et al. (2014)).
- (4) An Adaptive Physics-Based Non-Rigid Registration (APBNRR) (Drakopoulos, Liu, et al. (2014)).

On the one hand, the Rigid and BSpline methods use histogram bins and spatial samples in estimating a MMI (Mattes Mutual Information) cost metric for the registration. The larger the number of sampled voxels, the slower and more precise the fit. The default number of histogram levels is 50. For a fast registration, the default number of samples is approximately 0.2% of the total number of samples. For a higher accuracy, we set 100 histogram levels and sampling percentages 0.5% and 1% (Table 4). On the other hand, the physics-based methods (PBNRR, APBNRR, PAPBNRR) rely on the selection of image blocks with the highest intensity variance. A prohibited connectivity pattern is used to avoid the selection of blocks that are too close to each other (Figure 3, Table 5). The larger the number of the selected blocks, the slower and more precise the fit. The PBNRR, APBNRR use a NCC metric to measure the similarity of a pair of blocks in the pre-operative and the intra-operative image. The PAPBNRR employs a NCC or a NMI metric. For the evaluation results in subsections 3.1-3.3 we rely on the NCC for all the physics-based methods (Table 5). In subsection 3.4 we compare the results of our framework using different metrics (NCC, NMI).

The BSpline method requires the tessellation of the image with a 3D lattice, and uses piecewise cubic polynomials to optimize the registration. The physics-based methods require the tessellation of the image with an unstructured mesh, which uses either hexahedral or tetrahedral elements. In this paper, we employ meshes consisting of tetrahedral elements because they can adapt to highly curved, complex anatomical structures (e.g. tumor) with less distortion on their shape, compared to the hexahedral elements. Generally, the less distorted the elements, the better the conditioning of the mesh stiffness matrix. Additionally, our framework incorporates finite elements with either linear or quadratic interpolation basis-functions. In this study, we employ a linear (4-node) tetrahedron because it is less computationally intensive and thus, more suitable for time critical applications. PBNRR (Liu, Kot, et al. (2014)) uses a static, homogeneous biomechanical model to estimate the volumetric brain deformations. However, the adaptive physics-based methods (APBNRR, PAPBNRR), incorporate a heterogeneous (brain parenchyma, tumor) model, which is remeshed to eliminate the distorted or poor quality elements after each deformation. In all physics-based methods, we simulate the tissues as linear isotropic materials with mechanical properties (i.e., Young’s modulus, Poisson ratio) obtained from Miga, Roberts, et al. (2001). Table 5 lists the material properties of the tissues.

In the conducted experiments, the rigid registration was initiated using a manual linear transformation. The non-rigid registration was initiated using the produced rigid transformation. Tables 4, 5 list the input parameters for the BRAINSFit module and the physics-based methods, respectively. The 3D Slicer documentation⁶ and our previous work (Liu, Kot, et al. (2014); Drakopoulos, Liu, et al. (2014)) provide additional details for the parameters.

Table 4. The BRAINSFit input parameters for the Rigid and BSpline deformable registration methods. MMI: Mattes Mutual Information, VR3DT: Versor Rigid 3D Transform, LBFGSB: Limited memory Broyden Fletcher Goldfarb Shannon minimization with simple bounds.

Parameter	Value		Description
	Rigid	BSpline	
Initialization transform	Manual	Rigid	-
Sampling Percentage	0.5%	1%	% of voxels sampled for MMI
Histogram bins	100	100	Number of bins for MMI
Grid size	-	$15 \times 15 \times 15$	Number of grid points
Optimizer type	VR3DT	LBFGSB	-
Max iterations	1500	1500	Max iterations for optimizer
Min step length	10^{-3}	10^{-3}	Threshold for optimizer

⁶<http://www.slicer.org/slicerWiki/index.php/Documentation/4.3/Modules/BRAINSFit>

Table 5. The input parameters for the non-adaptive (PBNRR) and the adaptive (APBNRR, PAPBNRR) physics-based non-rigid registration methods. x: axial; y: coronal; z: sagittal; N_b : total number of selected blocks.

Parameter	Units	Value	Description	Physics-Based method
Initialization transform	-	Rigid	-	All
Linear solver	-	JCG	-	PBNRR-APBNRR
		BICGSTAB	-	PAPBNRR
Matrix scheme	-	ITPACK	-	PBNRR-APBNRR
		VNL	-	PAPBNRR
Element type	-	Linear tetrahedron	-	All
Similarity metric	-	NCC	-	All
F_s	-	5%	% selected blocks (results to N_b blocks)	All
Non-connectivity	-	“face”	Block non-connect. pattern	All
$B_{s,x} \times B_{s,y} \times B_{s,z}$	voxels	$3 \times 3 \times 3$	Block size	All
$W_{s,x} \times W_{s,y} \times W_{s,z}$	voxels	$13 \times 13 \times 5$ (Case 1-3)	Window size	All
		$7 \times 7 \times 5$ (Case 4-6)		
δ	-	5	Mesh size	All
E_b	Pa	$2.1 \cdot 10^3$	Brain Young’s modulus	All
E_t	Pa	$2.1 \cdot 10^4$	Tumor Young’s modulus	APBNRR-PAPBNRR
ν_b	-	0.45	Brain Poisson’s ratio	All
ν_t	-	0.45	Tumor Poisson’s ratio	APBNRR-PAPBNRR
λ	-	1	Trade off parameter	All
F_r	-	25%	% of rejected outlier blocks	All
N_{appr}	-	10	Num of outlier rejection steps	All
N_{int}	-	5	Num of interpolation steps	All
N_{iter}	-	4	Num of iterations	APBNRR
$N_{iter,max}$	-	4	Max num of iterations	PAPBNRR
$N_{b0,min}$	-	$1\%N_b$	Min num. of blocks with zero displacement	PAPBNRR

3.1 Qualitative Results

Figure 8 presents the qualitative results of this study. For comparative purposes, we depict the same representative MRI slice for all the images belonging to the same patient. A qualitative evaluation of the five registration methods can be performed with a visual comparison/inspection of: (i) the registered pre-operative MRI (columns 8(b)-8(f)) and the corresponding intra-operative MRI (column 8(a)), (ii) the subtracted MRI (8(g)-8(j)). The more similar the registered pre-operative and the intra-operative MRI, the more precise the alignment. The black and white regions in the subtracted images indicate larger discrepancies, while the gray regions indicate smaller discrepancies.

On the one hand, columns 8(e)-8(f) and 8(i)-8(j) illustrate the higher accuracy of the adaptive physics-based methods. The APBNRR and PAPBNRR preserve the morphology of the brain after the neurosurgical resection, especially near the tumor margins. On the other hand, the BSpline (columns 8(c), 8(g)) and the PBNRR (columns 8(d), 8(h)) fail to capture the geometry of the tumor cavities, but they exhibit higher accuracy on brain regions with small deformation (shift).

3.2 Quantitative Results

We employ a publicly available implementation of the Hausdorff Distance (HD) metric (Com-mandeur, Velut and Acosta (2011)) to compute the alignment error after a rigid and a non-rigid registration. The smaller the HD value, the better the alignment. The HD ratio demonstrates how many times more accurate (ratio > 1) the PAPBNRR alignment compared to the other methods. For each case study, the HD value is computed between two point sets: A' and B . Initially, A is extracted from the rigid registered pre-operative image and B from the intra-operative image. Then A' is computed by warping A according to the produced deformation field of each method. For the point extraction, we employ ITK’s Canny edge detection method (Canny (1986)).

Table 6 depicts the quantitative results. The proposed framework provides the finest alignments among all the methods; is on average 1.51, 2.93, 3.12, and 3.78 times more accurate than APBNRR, BSpline, PBNRR, and Rigid registration, respectively. Compared to our previous adaptive scheme (APBNRR), the new method reduces the alignment error on average by $(7.39 - 5.14)/7.39 \approx 30\%$. Figure 9 illustrates the HD alignment error for all the experiments. Clearly, the accuracy of the physics-based non-rigid registration methods has been continuously improving since the release of our first framework (PBNRR). Furthermore, our latest scheme requires fewer adaptive iterations

to register the MRI volumes compared to the APBNRR (Table 6). This is an important fact considering the hard time constraints imposed by the IGNS.




Figure 8. Qualitative results for the Rigid, BSpline, PBNRR, APBNRR, and PAPBNRR registration methods. Each row represents a single case. From top to bottom row: case 1-6. From left to right column: (a) intra-op MRI; (b) rigid pre-op MRI; (c) non-rigid BSpline pre-op MRI; (d) non-rigid PBNRR pre-op MRI; (e): non-rigid APBNRR pre-op MRI; (f): non-rigid PAPBNRR pre-op MRI; (g): (c) subtracted from (a); (h): (d) subtracted from (a); (i): (e) subtracted from (a); (j): (f) subtracted from (a).

Table 6. Quantitative evaluation results for the 6 clinical cases. HD_{RIGID} , $HD_{BSPLINE}$, HD_{PBNRR} , HD_{APBNRR} , and $HD_{PAPBNRR}$ is the alignment error after a Rigid, BSpline, PBNRR, APBNRR, and PAPBNRR registration, respectively. All HD are in mm. The number in the parenthesis denotes the number of adaptive iterations for the adaptive physics-based methods. PTR: Partial Tumor Resection, CTR: Complete Tumor Resection.

Case	Type	HD_{RIGID}	$HD_{BSPLINE}$	HD_{PBNRR}	HD_{APBNRR}	$HD_{PAPBNRR}$	$\frac{HD_{RIGID}}{HD_{PAPBNRR}}$	$\frac{HD_{BSPLINE}}{HD_{PAPBNRR}}$	$\frac{HD_{PBNRR}}{HD_{PAPBNRR}}$	$\frac{HD_{APBNRR}}{HD_{PAPBNRR}}$
1	PTR	16.90	14.97	12.99	9.91 (4)	5.73 (3)	2.95	2.61	2.27	1.73
2	PTR	15.81	15.10	13.71	8.96 (4)	7.89 (2)	2.00	1.91	1.74	1.14
3	CTR	18.69	11.56	15.90	3.57 (4)	3.56 (2)	5.25	3.25	4.47	1.00
4	CTR	14.45	11.57	12.20	5.83 (4)	3.74 (2)	3.86	3.09	3.26	1.56
5	CTR	17.72	17.72	13.49	6.00 (4)	6.16 (2)	2.88	2.88	2.19	0.97
6	CTR	21.42	14.45	17.91	10.04 (4)	3.74 (3)	5.73	3.86	4.79	2.68
Average		17.50	14.23	14.37	7.39	5.14	3.78	2.93	3.12	1.51

3.3 Performance Results

We run all the experiments in a Dell Linux workstation with two sockets of six Intel Xeon X5690@3.47 GHz CPU cores each, totaling twelve cores and 96GB of RAM. For the Rigid and BSpline, we run the BRAINSFit module from the terminal otherwise Slicer's GUI degrades the performance significantly. Table 7 illustrates the end-to-end execution time (including I/O) and




Figure 9. The Hausdorff Distance (HD) alignment error for the 6 clinical cases.

the speed-up. Slicer’s BRAINSFit module exhibits a real-time performance with twelve threads. The BSpline is the fastest among all the methods requiring on the average 12.25 seconds to complete. Also, it is highly parallelizable with an average speed-up of 6.54. The corresponding values for the Rigid registration are 27.96 seconds and a speed-up of 5.32. Regarding the performance

Table 7. Performance results for the 6 clinical cases with 1 and 12 threads. The experiments conducted in a workstation with 2 sockets of 6 Intel Xeon X5690@3.47 GHz CPU cores each, totaling 12 cores and 96GB of RAM. The I/O time is included.

Case	Time (sec)										Speed-up				
	RIGID		BSPLINE		PBNRR		APBNRR		PAPBNRR		RIGID	BSPLINE	PBNRR	APBNRR	PAPBNRR
	1T	12T	1T	12T	1T	12T	1T	12T	1T	12T					
1	60.18	16.20	156.28	20.08	138.71	81.18	579.62	213.97	483.63	93.14	3.71	7.78	1.71	2.71	5.19
2	280.55	45.31	128.68	17.80	123.76	73.99	509.89	192.70	275.26	54.12	6.19	7.23	1.67	2.65	5.09
3	555.65	77.33	135.91	18.71	112.79	68.72	486.95	185.86	265.61	52.46	7.19	7.26	1.64	2.62	5.06
4	45.51	9.76	33.75	6.31	23.80	20.90	107.43	63.17	53.97	18.51	4.66	5.35	1.14	1.70	2.92
5	52.25	10.76	31.26	5.97	23.91	20.62	109.17	64.08	53.94	18.85	4.86	5.24	1.16	1.70	2.86
6	44.53	8.40	29.36	4.61	19.46	16.80	81.90	48.85	61.32	21.10	5.30	6.37	1.16	1.68	2.91
Average	173.11	27.96	85.87	12.25	73.74	47.04	312.49	128.11	198.96	43.03	5.32	6.54	1.41	2.18	4.00




Figure 10. The end-to-end execution time for the 6 clinical cases using 1 thread (left) and 12 threads (right).

of the adaptive physics-based methods, we characterize them as nearly real-time. Nevertheless, our latest non-rigid registration technology provides the image alignments extremely fast (between 18.51s and 93.14s) for two reasons: (i) it requires fewer adaptive iterations to register the MRI volumes compared to the APBNRR method (Table 6), and (ii) it exploits additional parallelism (Figure 11). Indeed, the PAPBNRR completes the registration on average $47.04/43.03 \approx 1.1$ and $128.11/42.60 \approx 3$ times faster than the PBNRR and APBNRR, respectively.

Figures 10 and 11 illustrate the end-to-end execution time (one and twelve threads) and the end-to-end speed-up, respectively. The APBNRR is the most computationally intensive method. Its performance is limited when twelve cores are used because only 70% of the program was parallelized (Drakopoulos, Liu, et al. (2014)). On the contrary, in this paper, we further reduce the overheads due to adaptivity mainly on account of a new parallel FEM Solver (Figure 5(a)) and a parallel Closest Block Matching technique (Figure 4). Currently the only sequential component in the PAPBNRR is the BICGSTAB solver, which is part of the PFEMS module (see component *Solve System* in Figure 5(a)).

Table 8 illustrates the performance of all the PAPBNRR modules for the clinical case 1. The time that corresponds to the sequential portion of the program is 9.04 seconds. The time that corresponds to the parallel portion is $483.63 - 9.04 = 474.59$ seconds. As a result, the sequential and the parallel fraction of the PAPBNRR is $9.04/483.63 \approx 1.86\%$ and $474.59/483.63 \approx 98.13\%$, respectively; much higher than APBNRR’s parallelization fraction (70%). Besides, our multi-threaded implementation significantly reduces the execution time of the computationally intensive modules PFS, PBM, PCBM, and PFEMS. Specifically, the two block-matching schemes (PBM, PCBM) exhibit nearly linear speed-up (Table 8). In the future, we will parallelize the BICGSTAB solver to further reduce the overheads due to adaptivity. Generally, the proposed method is on average $(128.11 - 43.03)/128.11 \approx 66\%$ faster compared to the APBNRR (Table 7).

Finally, we should point out that the anisotropic data are at least 2-3 times more computationally intensive than the isotropic data because of their higher resolution.




Figure 11. The end-to-end speedup for the 6 clinical cases using 12 threads.

Table 8. Performance of the PAPBNRR modules with 1 and 12 threads (clinical case 1).

Module	Time (sec)		Speed-up
	1T	12T	
PFS	50.54	9.97	5.06
PBM	198.49	17.10	11.60
PCBM	57.08	5.50	10.37
PMG	18.12	12.43	1.45
PFEMS	133.09*	38.93*	3.41
PDFU	3.31	0.81	4.08
PIW	10.84	3.03	3.57
PDFC+PWSC	12.16	5.37	2.26
Total	483.63	93.14	5.19

* includes 9.04 sec for the sequential BICGSTAB.

3.4 Quantitative PAPBNRR results using different similarity metrics

Additionally, we evaluate our registration framework using a NMI metric, and we compare the alignment accuracy and the block matching execution time of the NCC and NMI metrics. Table 9 illustrates the results. The average alignment error is 13.6% smaller, when the NCC is used instead of NMI. In the future, we will explore the influence of a different block size and number of histogram levels on the NMI metric, which in turn, will address the accuracy of the registration. We also observe significant overheads in the parallel block matching performance when the NMI is employed (Table 9). Most likely, the computation of the marginal and joint entropies from (1) is more expensive, compared to the covariance and the standard deviation needed for the NCC. Nevertheless, in a multi-modal registration, where the affine hypothesis between the intensity distributions of the two images is hardly valid, the NMI metric is more suitable.

Table 9. Comparison of the HD alignment error and the block matching execution time for the NCC-NMI metrics, in PAPBNRR. PBM: Parallel Block Matching; NCC: Normalized Cross Correlation; NMI: Normalized Mutual Information; Number of histogram bins for NMI: 64; 12T: 12 Threads.

Case	HD _{PAPBNRR} (mm)		PBM Time (12T) (sec)	
	NCC	NMI	NCC	NMI
1	5.73	7.52	17.10	960.45
2	7.89	8.35	15.73	886.44
3	3.56	3.59	14.51	823.04
4	3.74	3.31	1.24	68.90
5	6.16	9.38	1.22	71.45
6	3.74	3.60	1.40	73.68
Average	5.14	5.95	10.14	480.66

4. Summary and Conclusion

We presented a fast and accurate biomechanical non-rigid registration method to compensate for the deformation induced by a brain tumor resection. The presented framework builds upon the ITK open-source system. It incorporates a variety of mesh generators and linear solvers, and relies on a parallel, adaptive, multi-material FE biomechanical model to accurately capture not only the small volumetric deformations, but also the complex geometric changes associated with the neurosurgical tumor resection.

Our method addresses the problem of missing/unrealistic matches occurring during the resection, with a two step block-matching scheme and an outlier rejection procedure, all implemented in parallel. The first block-matching step maximizes a similarity metric (i.e. NCC, NMI) over a specified window in the target image. The second step handles the blocks that do not have a correspondence in the target image. It matches them to their closest point on the surface of the target brain object. The final computed block displacements are scaled and an adjustable fraction of blocks with the largest error between the matching displacements and the estimated mesh deformations is rejected.

We compared the proposed scheme with various methods available on the visualization and image analysis packages 3D Slicer v4.4.0 and ITK v4.7.0, including: a Rigid registration, a BSpline deformable registration, a Physics-Based non-rigid registration, and an Adaptive Physics-Based non-rigid registration. The qualitative and quantitative results obtained with the six patients indicate that our registration technology can be used to overcome the intra-operative acquisition problem and enhance the surgeon’s view during a neurosurgical procedure for several reasons. First, it achieves the most precise alignments so far, especially nearby the tumor resection margins. The quality of the alignments is not affected by the portion (partial or complete) of the resected tumor depicted in the iMRI or the type of the image data (isotropic or anisotropic). Second, it reduces the overheads due to the finite element adaptivity, satisfying the hard time constraints imposed by IGNS. As a result, it can register a pair of isotropic or anisotropic adult brain MRI in less than 21 or 93 seconds, respectively.

In our future work, we will investigate how the segmentation accuracy and the use of additional tissues (e.g. falx cerebri, ventricles) in the heterogeneous model, affect the accuracy and the performance of the deformable registration. Also, we are planning to integrate a hybrid meshing algorithm, for generating a mixed heterogeneous mesh model, consisting of tetrahedral, pentahedral, and hexahedral elements. In that way, we will improve the mesh gradation and focus our computational resources in regions of interest. Our preliminary results indicate that we could reduce the vertices (i.e. degrees of freedom) by 20%, and thus, the computational cost of the Parallel FEM solver.

Finally, we will evaluate our framework on other applications (i.e. Deep Brain Stimulation (DBS)). Modern DBS surgery makes use of stereotactic systems and image guidance from different modalities (e.g. MRI, CT) to facilitate the accurate placement of electrode leads into precise locations in

the deep brain structures, and guide the surgery.

4.1 Acknowledgements

This work is supported in part by NSF grants: CCF-1139864 and CCF-1439079 and by CRCF grant No. MF15-002-LS, John Simon Guggenheim Foundation, ODU's Modeling and Simulation fellowship, and the Richard T.Cheng Endowment. Special thanks for helping with edits and proofreading to Alexis Brueggeman.

References

- Antiga L, Piccinelli M, Botti L, Ene-Iordache B, Remuzzi A, Steinman D. 2008. An image-based modeling framework for patient-specific computational hemodynamics. *Medical & Biological Engineering & Computing*. 46(11):1097–1112.
- Archip N, Clatz O, Whalen S, Kacher D, Fedorov A, Kot A, Chrisochoides N, Jolesz F, Golby A, Black PM, Warfield SK. 2007. 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.
- Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, Dalcin L, Eijkhout V, Gropp WD, Kaushik D, et al. 2015. PETSc users manual. Argonne National Laboratory. Report No.: ANL-95/11 - Revision 3.6. Available from: <http://www.mcs.anl.gov/petsc>.
- Barrett R, Berry M, Chan TF, Demmel J, Donato J, Dongarra J, Eijkhout V, Pozo R, Romine C, der Vorst HV. 1994. Templates for the solution of linear systems: Building blocks for iterative methods, 2nd edition. SIAM.
- Bathe KJ. 1996. Finite element procedures. New Jersey: Prentice-Hall.
- Brown LG. 1992. A survey of image registration techniques. *ACM Comput Surv*. 24(4):325–376.
- Canny J. 1986. A computational approach to edge detection. *IEEE Trans Pattern Anal Mach Intell*. 8(6):679–698.
- Chen W, Ying M, Jian-Hong Z, Liang-Fu Z. 2008. The department of neurosurgery at shanghai huashan hospital. *Neurosurgery*. 62(4):947–953.
- Chernikov AN, Chrisochoides NP. 2011. Multitissue tetrahedral image-to-mesh conversion with guaranteed quality and fidelity. *SIAM J Sci Comput*. 33(6):3491–3508.
- Chrisochoides N, Fedorov A, Kot A, Archip N, Black P, Clatz O, Golby A, Kikinis R, Warfield S. 2006. Toward real-time image guided neurosurgery using distributed and grid computing. In: SC 2006 Conference, Proceedings of the ACM/IEEE; Nov. p. 37–37.
- Clatz O, Delingette H, Talos IF, Golby A, Kikinis R, Jolesz F, Ayache N, Warfield S. 2005. Robust nonrigid registration to capture brain shift from intraoperative mri. *Medical Imaging, IEEE Transactions on*. 24(11):1417–1427.
- Commandeur F, Velut J, Acosta O. 2011. A vtk algorithm for the computation of the hausdorff distance. *VTk Journal*.
- Drakopoulos F, Chrisochoides N. 2014. A parallel adaptive physics-based non-rigid registration framework for brain tumor resection. In: Zhang Y, Tavares J, editors. Computational modeling of objects presented in images. fundamentals, methods, and applications. Lecture notes in computer science; vol. 8641. Springer International Publishing; p. 57–68.
- Drakopoulos F, Liu Y, Foteinos P, Chrisochoides NP. 2014. Towards a real time multi-tissue adaptive physics based non-rigid registration framework for brain tumor resection. *Frontiers in Neuroinformatics*. 8(11).
- Farnia P, Ahmadian A, Shabanian T, Serej N, Alirezaie J. 2015. Brain-shift compensation by non-rigid registration of intra-operative ultrasound images with preoperative mr images based on residual complexity. *International Journal of Computer Assisted Radiology and Surgery*. 10(5):555–562.
- Ferrant M, Cuisenaire O, Macq BMM. 1999. Multiobject segmentation of brain structures in 3d mri using a computerized atlas. vol. 3661. p. 986–995.
- Foteinos P, Chrisochoides N. 2013. High quality real-time image-to-mesh conversion for finite element simulations. In: Proceedings of the 27th International ACM Conference on International Conference on Supercomputing; Eugene, Oregon, USA. New York, NY, USA: ACM; p. 233–242. ICS '13.
- Foteinos P, Liu Y, Chernikov A, Chrisochoides N. 2011. An evaluation of tetrahedral mesh generation for

- nonrigid registration of brain mri. In: Wittek A, Nielsen PM, Miller K, editors. Computational biomechanics for medicine. Springer New York; p. 131–142.
- Foteinos PA, Chrisochoides NP. 2014. High quality real-time image-to-mesh conversion for finite element simulations. *Journal of Parallel and Distributed Computing*. 74(2):2123–2140.
- Gu Z, Qin B. 2009. Nonrigid registration of brain tumor resection mr images based on joint saliency map and keypoint clustering. *Sensors*. 9(12):10270–10290.
- Johnson H, Harris G, Williams K. 2007. Brainsfit: Mutual information registrations of whole-brain 3d images, using the insight toolkit.
- Kenneth KS. 1986. Sparse matrix techniques. *Circuit Analysis, Simulation and Design*. 3(1).
- Kincaid DR, Respass JR, Young DM, Grimes RR. 1982. Algorithm 586: Itpack 2c: A fortran package for solving large sparse linear systems by adaptive accelerated iterative methods. *ACM Trans Math Softw*. 8(3):302–322.
- Liu Y, Fedorov A, Kikinis R, Chrisochoides N. 2009. Real-time non-rigid registration of medical images on a cooperative parallel architecture. In: *Bioinformatics and Biomedicine, 2009. BIBM '09. IEEE International Conference on*; Nov. p. 401–404.
- Liu Y, Foteinos P, Chernikov A, Chrisochoides N. 2012. Mesh deformation-based multi-tissue mesh generation for brain images. *Engineering with Computers*. 28(4):305–318.
- Liu Y, Kot A, Drakopoulos F, Yao C, Fedorov A, Enquobahrie A, Clatz O, Chrisochoides NP. 2014. An itk implementation of a physics-based non-rigid registration method for brain deformation in image-guided neurosurgery. *Frontiers in Neuroinformatics*. 8(33).
- Miga M, Roberts D, Kennedy F, Platenik L, Hartov A, Lunn K, Paulsen K. 2001. Modeling of retraction and resection for intraoperative updating of images. *Neurosurgery*. 49(1):75–84; discussion 84–5.
- Mostayed A, Garlapati R, Joldes G, Wittek A, Roy A, Kikinis R, Warfield S, Miller K. 2013. Biomechanical model as a registration tool for image-guided neurosurgery: Evaluation against bspline registration. *Annals of Biomedical Engineering*. 41(11):2409–2425.
- Paige CC, Saunders MA. 1982. Lsqr: An algorithm for sparse linear equations and sparse least squares. *ACM Trans Math Softw*. 8(1):43–71.
- Periaswamy S, Farid H. 2006. Medical image registration with partial data. *Medical Image Analysis*. 10(3):452 – 464. Special Issue on The Second International Workshop on Biomedical Image Registration (WBIR03).
- Platenik L, Miga M, Roberts D, Lunn K, Kennedy F, Hartov A, Paulsen K. 2002. In vivo quantification of retraction deformation modeling for updated image-guidance during neurosurgery. *IEEE Trans Biomed Eng*. 49(8):823–35.
- Risholm P, Melvr EL, Mrken K, Samset E. 2009. Intra-operative adaptive fem-based registration accommodating tissue resection. vol. 7259. p. 72592Y–72592Y–11.
- Smith SM. 2002. Fast robust automated brain extraction. *Human Brain Mapping*. 17(3):143–155.
- Studholme C, Hill D, Hawkes D. 1999. An overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognition*. 32(1):71–86.
- Warfield SK, Ferrant M, Gallez X, Nabavi A, Jolesz FA. 2000. Real-time biomechanical simulation of volumetric brain deformation for image guided neurosurgery. In: *Proceedings of the 2000 ACM/IEEE conference on Supercomputing*; Dallas, Texas, USA. Washington, DC, USA: IEEE Computer Society. Supercomputing '00.