

# Advancing Intra-operative Precision: Dynamic Data-Driven Non-Rigid Registration for Enhanced Brain Tumor Resection in Image-Guided Neurosurgery

Nikos Chrisochoides, Andriy Fedorov, Fotis Drakopoulos, Andriy Kot, Yixun Liu, Panos Foteinos, Angelos Angelopoulos, Olivier Clatz, Nicholas Ayache, Peter M. Black, Alex J. Golby, Ron Kikinis

**Abstract** During neurosurgery, medical images of the brain are used to locate tumors and critical structures, but brain tissue shifts make pre-operative images unreliable for accurate removal of tumors. Intra-operative imaging can track these deformations but is not a substitute for pre-operative data. To address this, we use Dynamic Data-Driven Non-Rigid Registration (NRR), a complex and time-consuming image processing operation that adjusts the pre-operative image data to account for intra-operative brain shift. Our review explores a specific NRR method for registering brain MRI during image-guided neurosurgery and examines various strategies for improving the accuracy and speed of the NRR method. We demonstrate that our implementation enables NRR results to be delivered within clinical time constraints while leveraging Distributed Computing and Machine Learning to enhance registration accuracy by identifying optimal parameters for the NRR method. Additionally, we highlight challenges associated with its use in the operating room.

**Key words:** Dynamic Data Driven Application Systems, Image-Guided Neurosurgery, Deformable Registration, Intra-operative Parametric Search, Deep Learning

---

Nikos Chrisochoides, Andriy Fedorov, Fotis Drakopoulos, Panos Foteinos, Angelos Angelopoulos, Andriy Kot, Yixun Liu  
Center for Real-Time Computing, Old Dominion University, Norfolk, VA USA e-mail: nikos@cs.odu.edu

Andriy Fedorov, Ron Kikinis, Alexandra J. Golby  
Surgical Planning Laboratory, Department of Radiology, Brigham and Women's Hospital, Boston, MA USA


Olivier Clatz, Nicholas Ayache  
INRIA Sophia Antipolis, France

Alexandra J. Golby, Peter M. Black  
Department of Neurosurgery, Brigham and Women's Hospital, Boston, MA USA

## 1 Introduction

Cancer is one of the leading causes of death in the USA and worldwide. Among the different types of cancer, brain cancer was estimated to claim over 50 thousand new victims in 2008 [2], when we first consider summarizing the preliminary results of our approach. However, brain cancer continues to be a significant health-care problem. The number of Americans living with brain tumors exceeds 700,000, surpassing the number of COVID-19 deaths in mid-summer 2021.

Neurosurgical resection is one of brain tumor patients' most common and effective treatment options. The resection must remove as much as possible of the tumor tissue while maximally preserving the vital structures of the healthy brain. Maximal tumor excision increases time to progression, and reduces symptoms and seizures. In this Chapter, we explore how the concept of Dynamic Data Driven Applications Systems (DDDAS) [88, 37], together with the advances in medical image acquisition, distributed computing, and Machine Learning can assist in enabling image guidance during neurosurgery and potentially can improve the accuracy of the procedure, allowing more complete tumor resections without additional morbidity. We focus more on the implementation aspects, while in [90], we focus more on the approach's mathematical modeling and computational aspects.



**Fig. 1** Intra-operative brain deformation. Left: pre-operative, higher quality image, showing the location of brain tumor. Right: intra-operative image showing brain shift [5].

There are many challenges in accomplishing the objectives of neurosurgery. In this Chapter, we focus on one of them: the exact locations of the brain areas responsible for critical brain function, e.g., the motor cortex, are patient-specific and cannot be identified with the naked eye. This is where medical imaging, distributed computing, and machine learning become essential.

Magnetic Resonance Imaging (MRI) is indispensable in demonstrating brain pathologies. Although not distinguishable from the naked eye, neoplastic tissues can be differentiated from brain tissue based on changes in MR signal and corresponding image intensities. MRI has also been shown to be useful in constructing functional mapping of the brain using functional MRI (fMRI) [62]. Both the structural and

functional imaging data are used for the purposes of improving the precision of the resection.

Image registration in general, is concerned with spatial alignment of corresponding features in two or more images. During image registration, a spatial transformation is applied to one image, which is called *floating*, so that it is brought into alignment with the *target*, or *reference* position of the object (brain). During rigid image registration, the floating image corresponds to the pre-operative image, which is aligned with the patient's position using translations and rotations (rigid transformations).

During the course of surgery, opening of the skull and dura causes changes in pressure inside the Intra-Cranial Cavity (ICC).

Because of this and other factors, such as drainage of cerebrospinal fluid, induced changes in brain tumors, and the effect of gravity, the brain changes its shape, introducing discrepancies in relation to the pre-operative configuration. *Non-rigid* image registration uses a spatially varying transformation to account for this deformation. In general, image registration algorithms are based on optimizing certain similarity criteria between the fixed and floating image under varying parameters of spatial transformation. The complexity of this optimization depends on the number of parameters that describe the transformation. Both rigid and non-rigid registration are open research areas in medical image processing. However, non-rigid registration is a conceptually more difficult problem, which usually requires significant computing resources and time.

Non-rigid registration recovers the deformation of the brain based on intra-operatively acquired imaging data. Advances in medical image acquisition have made it possible to acquire high-resolution images, in particular MRI during the surgery. Intra-operative MRI (iMRI) cannot substitute pre-operative images because of its limited resolution and the high processing time required to obtain functional data. However, iMRI can be used to guide the registration of the pre-operative data.

There are three main requirements to non-rigid registration (NRR) [24]. First, NRR should deliver accurate results. Second, the result should be consistently accurate, independent of the specific registered images, and not be sensitive to small variations in the parameter selection. One important requirement for Image-Guided Neurosurgery (IGNS) is that registration must be completed within the neurosurgical workflow's time constraints, typically between 4 to 5 minutes.

The prospective application of NRR is a dynamic process. iMRI is obtained periodically as requested by the surgeon. Immediately following iMRI, NRR should be used to estimate the deformation of the brain and update the pre-operative images. Usually, hospitals do not have locally available large-scale computational facilities. In this Chapter, we describe an infrastructure that enables the computation of non-rigid registration using remotely located high-performance computing resources guided by intra-operative image updates.

In the context of the application, we define the *response time* as the time between the acquisition of the intra-operative scan of the deformed tissue and the final visualization of the registered preoperative data on the console in the operating room. These steps are performed intra-operatively from the DDDAS steered by the peri-

odic acquisitions of the iMRI data. Our broad objective is to minimize the perceived (end-to-end) response time of the DDDAS component.

To our knowledge, none of the systems developed (2005 timeframe) were used prospectively during image-guided neurosurgeries. Our approach to developing such a dynamic data-driven NRR system for IGNS was to adopt an existing NRR method of established accuracy and parallelize the most time-consuming components of this method and develop an end-to-end system to facilitate image guidance during neurosurgery.

## 2 Related Work

The research in NRR for IGNS can be separated into the development of the core registration methods and the design of end-to-end systems capable of supporting NRR computation and delivering the results intra-operatively. The choice of the NRR method depends mostly on the intra-operative image modality that captures brain deformation [26]. However, the core computation components of NRR are very similar for different intraoperative imaging modalities.

Registration algorithms are based on optimizing certain similarity measures between the intensities of the reference and floating images. In non-rigid registration, the number of parameters (degrees of freedom) that are being optimized is exceedingly large compared to rigid registration. This contributes significantly to the costs of computing the similarity metric and to the evaluation of gradients required during optimization. However, optimizing the similarity measure alone can lead to unrealistic solutions since non-rigid registration is an ill-posed problem. Therefore, NRR usually includes some form of solution regularization. Biomechanical modeling of tissue deformation is one such regularization approach. Deformation of tissue is usually modeled using the Finite Element Method (FEM) [107], which requires solving a system of equations. The size of this system is proportional to the resolution of the brain biomechanical model.

Timely completion of the core NRR computations is the key component for efficient end-to-end registration systems. Several strategies have been proposed to parallelize the time-consuming steps in medical image processing. Christensen and collaborators were some of the first to discuss using parallel computing resources for solving time-consuming problems related to brain MRI processing on a massively parallel SIMD architecture [23].

Warfield et al. [102] presented some of the first results in intra-operative processing (segmentation) of iMRI. The authors demonstrate linear speedup of segmentation on a 20-processor workstation, which allows the processing of a typical dataset in about 20 seconds. The developed method was applied and evaluated prospectively during neurosurgeries and liver cryo-ablation procedures [103]. The same group later developed a high-performance method for intra-operative non-rigid registration, which uses linear biomechanical model [101] solved in parallel. Although

the authors report clinically acceptable timing results delivered by their implementation, the evaluation was restricted to off-line experimental studies.

Computation of the NRR result within the time constraints of neurosurgery is an essential requirement. To facilitate this task, support of the computation on the remote resources may be required. The community has recognized these issues, and several solutions have been proposed. Stefanescu et al. [99] describe an NRR implementation that is exposed as a web service. Ino et al. developed an end-to-end system for rigid registration computation on a remote cluster [63]. Lippman and Kruggel use a customized grid infrastructure to design an NRR system for IGNS [71].

Medical image computing is one area where machine learning (ML) is applied to solve problems [67]. Specifically, for medical image registration, many new methods currently utilize convolutional neural networks, which operate directly on the preoperative and intraoperative MRI images to produce a registered image. A notable example is MIT's VoxelMorph model [9]. VoxelMorph utilizes a CNN, which computes a registration field and warps the preoperative image with the registration field using a spatial transformation function. In VoxelMorph, ML is used directly for image registration, while in our case ML is used for optimal values of the parameters in NRR. This Chapter will look at two major applications, VoxelMorph[9] and the method proposed by J. Krebs et al. [69].

VoxelMorph [9] addresses the problem of fast deformable medical image registration with a focus on brain MRI, but it can also be used for other tissues. VoxelMorph uses a solution formulated by an unsupervised learning convolutional neural network for computing a registration field and a spatial transformation function for warping the preoperative image. The application can also use instance-specific optimization by fine-tuning the network parameters for each MR image. As noted in [9]

Application	Problem	Solution	Machine Learning Usage
VoxelMorph	Fast deformable medical image registration (focus on brain MRI, but can be used for other tissues as well)	Unsupervised learning convolutional neural network with spatial transformation function and instance-specific optimization	Computation of registration field
Deep Learning APBNRR	Deformable medical image registration with tumor resection (for brain MRI with focus on accurate geometry representation of brain in modeling deformation)	Adaptive physics-based non-rigid registration with supervised learning deep feedforward neural network and patient-specific features	Prediction of optimal input parameters for APBNRR
J. Krebs et al Method	Probabilistic deformation modeling for diffeomorphic registration (for cardiac MRI)	Unsupervised learning conditional variational autoencoder (CVAE) network with exponentiation layer	Learning of a probabilistic deformation encoding.

**Fig. 2** Compares three machine learning approaches in medical image registration for major applications: two state-of-the-art and APBNRR deep learning; the columns from left to right correspond to the application, problem, solutions, and usage.

it runs in 0.45 seconds on a top-tier GPU and 57 seconds on a CPU, with an average DICE [36] score of 75.3%. Another approach is proposed by J. Krebs et al. [69] aims to address the issue of probabilistic deformation modeling for diffeomorphic registration in cardiac MRI. Their solution employs an unsupervised learning conditional variational autoencoder (CVAE) network, which utilizes an exponentiation layer to create diffeomorphic transformations. The average execution time is 0.32 seconds on a top-tier GPU with an average DICE score of 79.9% and a mean Hausdorff distance of 7.9 mm. A comparison of the problems, the proposed solutions, and the usage of machine learning in NRR and the two applications listed above are presented in the Table of Figure 2 from [3].

### 3 Physics-Based Non-Rigid Registration

The core registration method of our dynamic infrastructure was originally developed by Clatz et al. in [24]. This Physics-Based Non-Rigid Registration (PBNRR) approach is specifically designed for registering high-resolution pre-operative data with iMRI. The NRR computation consists of preoperative and intra-operative components. Intraoperative processing starts with the acquisition of the first iMRI. However, the *time-critical* part of the intra-operative computation is initiated when a scan showing the shift of the brain is available. The basic idea of the registration method is to estimate the *sparse deformation field* that matches similar locations in the image and then uses a biomechanical model of the brain, deformation to discard unrealistic displacements and derive *dense deformation field* that defines transformation for each point in the image space.

Sparse displacement vectors are obtained at the selected points in the image, where the variability in the intensities in the surrounding region exceeds some threshold. Such *registration* points can be identified prior to the time-critical part of the computation in the floating (pre-operative) image. Once the reference (intra-operative) scan is available, the deformation vector is estimated at each of the selected points by means of block matching. Fixed-size rectangular regions (blocks) centered at the registration points are identified in the floating image. Given such a block, we next select a search region (window) in the reference image. The block's displacement that maximizes the intensity-based similarity metric between the image intensities in the block and the overlapping portion of the window corresponds to the vector value of the sparse deformation field at the registration point. The normalized cross-correlation (NCC) similarity metric is evaluated as follows:

$$NCC = \frac{\sum_{i \in B} (B_T(i) - \bar{B}_T)(B_F(i) - \bar{B}_F)}{\sqrt{(B_T(i) - \bar{B}_T)^2 (B_F - \bar{B}_F)^2}}.$$

$\bar{B}_T$  and  $\bar{B}_F$  correspond to the average intensity values within the block in the reference and floating image, respectively. We note the high computational complexity of the block-matching procedure. Considering the sizes of three-dimensional blocks

and windows are defined in pixels as  $B = \{B_x, B_y, B_z\}$  and  $W = \{W_x, W_y, W_z\}$ , the bound on the number of operations is  $O(B_x B_y B_z \times W_x W_y W_z)$  for one registration point.

Estimation of brain deformation is based on the finite element analysis (FEA) using a linear elastic model of brain deformation. The finite element mesh of the intra-cranial volume is constructed from the segmented ICC volume following the methods we evaluated in a separate studies [39] and [35]. We then iteratively seek such a position of the mesh vertices  $\mathbf{U}$  that balances the mechanical forces of the modeled tissue that resist deformation, with the external forces that correspond to the displacements  $\mathbf{D}$  estimated by block matching:

$$\mathbf{F}_i \Leftarrow \mathbf{K}\mathbf{U}_i, \mathbf{U}_{i+1} \Leftarrow [\mathbf{K} + \mathbf{H}^T \mathbf{S} \mathbf{H}]^{-1} [\mathbf{H}^T \mathbf{S} \mathbf{D} + \mathbf{F}_i].$$

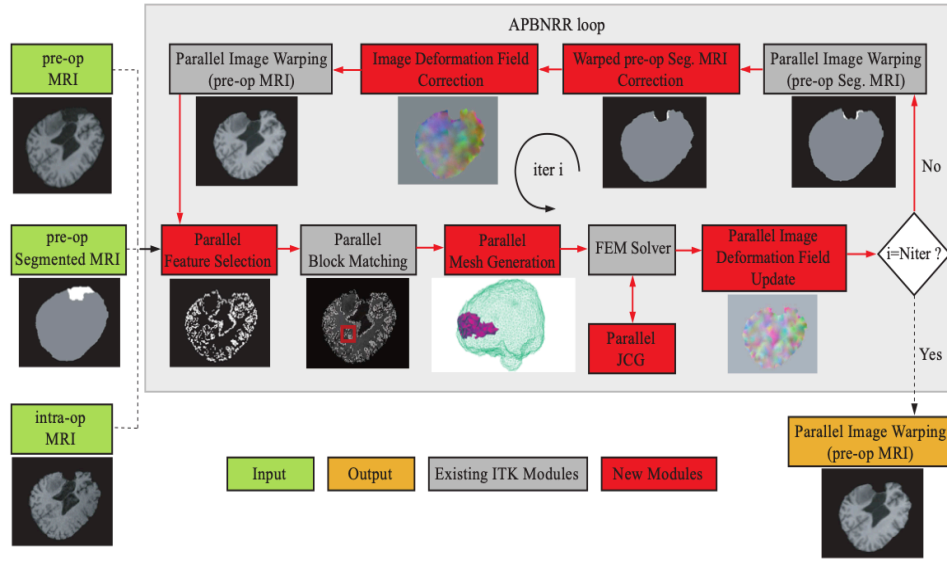
Here,  $\mathbf{K}$  is the mechanical stiffness matrix [25],  $\mathbf{H}$  is the interpolation matrix from the mesh vertices to the block matching displacements,  $\mathbf{S}$  is the matrix that captures the confidence in the block matching results.  $\mathbf{F}$  is the force that is increasing between iterations to slowly cancel the influence of the mechanical forces.

Both block matching and iterative estimation of displacements are time critical and should be performed while the surgeons are waiting. Block matching contributes most to the computation costs because of the exhaustive search for the optimum block position. Iterative estimation of mesh vertex displacements based on a biomechanical model requires a solution of a system of linear equations during each iteration. However, the size of that system is constrained by the number of mesh vertices, which cannot be arbitrarily large due to inherent properties of the NRR algorithm [39].

### 3.1 Adaptive Non-Rigid Registration (APBNRR)

However, complete resection of large brain tumors leads to large brain shifts. This creates a large cavity of elements in the tessellated brain image model which compromises the accuracy of the biomechanical model defined on pre-operative MRI. This section summarizes the extensions of the PBNRR by: (i) removing additional outliers due to tissue resection using an Adaptive Non-Rigid Registration (APBNRR) method and gradually adjusting the mesh for the FEM model to an incrementally warped segmented intraoperative MRI (iMRI), see Figure 3 and for a more detailed description see [27], [28] and [35].

APBNRR takes as input a preoperative segmented MRI (moving) and a range of twenty-seven registration and mesh generation parameters (indicated in Table of Figure 4. APBNRR augments PBNRR to accommodate soft-tissue deformation caused by tumor resection. This iterative method adaptively modifies a heterogeneous finite element model to optimize non-rigid registration in the presence of tissue resection. Using the segmented tumor and the registration error at each iteration, APBNRR gradually excludes the resection volume from the model. During each it-



**Fig. 3** The APBNRR framework. The green, red, and gray boxes represent the input, the additional parallel modules to manage real-time adaptivity (in red), and the existing ITK modules [80]. The red arrows show the execution order of the modules. Orange represents the output warped pre-operative MRI.

eration, registration is performed, the registration error is estimated, the mesh is deformed to a predicted resection volume, and the brain model (minus the predicted resection volume) is re-tessellated. Re-tessellation is required to ensure high-quality mesh elements, which is important for the convergence of the linear solver.

#### 4 High-Performance Infrastructure for Non-Rigid Registration

The baseline code used in the design of the NRR was the implementation developed and evaluated by Clatz et al. [24]. Based on the benchmarking and analysis of this implementation, we identified the following problems:

1. The execution time of the original non-rigid registration code is highly data-dependent. When computed on a high-end 4 CPU workstation, the computation time varies between 30 and 50 minutes. The scalability of the code is poor due to workload imbalances.
2. The code is designed as a single monolithic component (since it was not evaluated in the intraoperative mode), and a single failure at any point requires restarting the registration from the beginning.
3. The original code is implemented in PVM [11] which is not widely supported as compared to the use of MPI [97] for message passing.



Parameter	Description	Parameter (cont.)	Description (cont.)
<b>1. Half Block Size X</b>	Half block size (in voxels) for the X dimension.	<b>15. Mesh Method</b>	Meshing method (BCC, Delaunay, LD, Hex)
<b>2. Half Block Size Y</b>	Half block size (in voxels) for the Y dimension.	<b>16. Trade-off</b>	Trade-off between mechanical energy and matching energy in energy minimization equation. The larger the trade-off, the more weight is given to the matching energy term.
<b>3. Half Block Size Z</b>	Half block size (in voxels) for the Z dimension.	<b>17. Shape Function Type</b>	FEM shape function type (Linear, Linear with ESF, Quadratic, Quadratic with ES)
<b>4. Half Window Size X</b>	Half block matching window size (in voxels) for the X dimension.	<b>18. Linear Solver Type</b>	Linear solver type (LSQR, ITPACK, LU, BICGSTAB)
<b>5. Half Window Size Y</b>	Half block matching window size (in voxels) for the Y dimension.	<b>19. Young's Modulus (parenchyma)</b>	Young's modulus for brain parenchyma.
<b>6. Half Window Size Z</b>	Half block matching window size (in voxels) for the Z dimension.	<b>20. Poisson Ratio (parenchyma)</b>	Poisson ratio for brain parenchyma.
<b>7. Selection Fraction</b>	Percentage of selected blocks from total number of blocks.	<b>21. Young's Modulus (tumor)</b>	Young's modulus for tumor.
<b>8. Approximation Steps</b>	Number of approximation iterations.	<b>22. Poisson Ratio (tumor)</b>	Poisson ratio for tumor.
<b>9. Interpolation Steps</b>	Number of interpolation iterations.	<b>23. CBC3D Spacing</b>	Lattice spacing (in mm) for CBC3D.
<b>10. Adaptive Iterations</b>	Maximum number of adaptive iterations	<b>24. CBC3D Fidelity</b>	Mesh fidelity for CBC3D.
<b>11. Zero Blocks Fraction</b>	Minimum number of blocks with zero correspondence.	<b>25. CBC3D Iterations</b>	Smoothing iterations for CBC3D.
<b>12. Rejection Fraction</b>	Percentage of rejected outlier blocks.	<b>26. CBC3D Multiple Fidelities</b>	Multiple fidelities for CBC3D.
<b>13. Similarity Metric</b>	Similarity metric for block matching (NCC or NMI).	<b>27. Delaunay Delta</b>	Tetrahedron size for Delaunay meshing.
<b>14. Non-connectivity</b>	The pattern for feature selection.		

**Fig. 4** Shows the tunable parameters APBNRR utilizes. Parameters 1-12 are utilized in the deep learning model, while the rest are fixed. *I/O* parameters are not included in this table. CBC3D refers to the latest image-to-mesh conversion presented in [29].

Consequently, we identified the following implementation objectives in the design of the system.

**High-performance** Develop an efficient and portable software environment for parallel and distributed implementation of real-time non-rigid registration method for both small-scale parallel machines and large-scale geographically distributed Clusters of Workstations (CoWs). The implementation should be able to work on both dedicated and time-shared resources.

**Quality-of-service (QoS)** Provide functionality not only to sustain failure but also to dynamically replace/reallocate faulty resources with new ones during the real-time data acquisition and computation.

**Ease-of-use** Develop a GUI that automatically will handle exceptions (e.g., faults, resource management, and network outages), and assist in the parameter initialization.

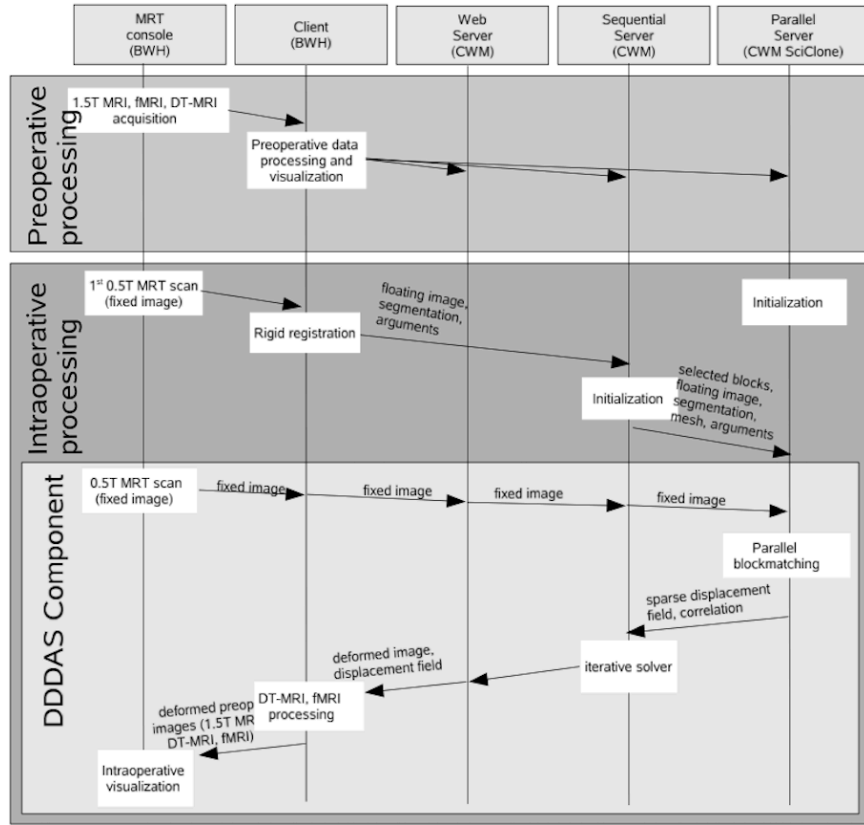
Different strategies can be explored in the high-performance implementation of the described NRR method. We first explore how this can be done using ubiquitous CoWs. During the studies of NRR at BWH, the implementation based on CoW was utilized prospectively, as mentioned in the reference [4]. We also describe our efforts to further increase the availability of the implementation by developing its components ported on Graphical Processing Units (GPUs) and studying the use of Grid resources.

We develop NRR DDDAS based on the concept of the *computational workflow*. We re-design the core NRR implementation as a coordinated set of processing components communicating by passing data. Such an approach allows to separate time-critical steps, and concentrate on the optimum parallelization strategies for each step that requires performance improvement.

#### 4.1 Cluster of Workstations

CoWs have become power plants of ubiquitous computing. Availability of such cluster at the College of William and Mary (CWM, Williamsburg, VA) motivated the development of the implementation of the CoW-based NRR DDDAS. In addition to the dedicated computing cluster, we use the shared resources of a computer lab to boost the computing power and reliability of the implemented system. The targeted users of our DDDAS are clinical researchers of Brigham and Women’s Hospital (BWH, Boston, MA). Our approach is to map the components of the workflow on the computing and communication resources of CWM and BWH and expose the DDDAS to clinical researchers by means of a web service interface. The timeline of the interaction with the complete NRR DDDAS is shown in Figure 5. The key component of this system, which requires parallelization, is block matching.

**Parallel feature selection.** PBNR and APBNRR algorithmically identify image features by analyzing voxel intensity variation across the intracranial cavity. For each feature candidate, it computes the variance within a block of size  $\mathbf{B}_s$ . It then selects  $\mathbf{F}_s$  features with the highest variance. Experimental evaluation has shown that when  $\mathbf{B}_s = 3$  or 5 and  $3\% < \mathbf{F}_s < 10\%$ , a sufficient number of image blocks ( $> 310^5$ ) can be selected. The method also uses a connectivity pattern to avoid selecting blocks that are too close to each other, thereby influencing the distribution of selected blocks in the image. Three simplex patterns are available: “vertex” (i.e., zero-order simplex implies 26 connectivity), “edge” (i.e., first-order simplex implies 18 connectivity), and “face” (i.e., second-order simplex implies 6 connectivity). The higher the order of the simplex pattern the higher the density of the selected blocks. Since the “face” pattern results in a higher density of blocks near the boundaries/interfaces of anatomical structures, features expected to be most persistent between pre-operative and intraoperative image acquisitions, it is most suitable



**Fig. 5** Timeline of the image processing steps during IGNS (the client is running at BWH, and the server is using multiple clusters at CWM, for fault-tolerance purposes).

for IGNS-based Glioma resection. The parallel implementation partitions the preoperative image into  $k$  sub-regions, where  $k$  is the number of threads. Each thread computes a variance value and an image index for each feature inside the sub-region. Then, the computed pairs are sorted in parallel based on their variance and merged into a global vector. The size of the global vector is equal to the total number of computed blocks. Finally,  $0.5 + N_{\text{Features}} \times F_s$  blocks are selected from the global vector.

**Multi-level distributed block matching** To find a match for a given block, we need the block center coordinates and the areas of the fixed and floating images bounded by the block matching window [24]. The fixed and floating images are loaded on each processor during the initialization step, as shown in Fig. 5. The total workload is maintained in a *work-pool* data structure. Each item of the work pool contains the three coordinates of the block center (the total number of blocks for a typical dataset is around 100,000), and the best match found for that block (in case the block was processed; otherwise, that field is empty).

However, because of the scarce resource availability, we have to handle computational clusters that belong to different administrative domains. We address this issue with hierarchical multi-level organization of the computation using the master-worker model. A dedicated master node is selected within each cluster. The master maintains a replica of the global work pool and is responsible for distributing the work according to the requests of the nodes within the assigned cluster and communicating the execution progress to the other master(s).

**Parallel Mesh Generation** The image segmentation generates a patient-specific finite element mesh for PBNRR and APBNRR. The tetrahedral mesh’s quality influences the solution’s numerical accuracy and the correctness of the estimated transformation. The higher the quality of the elements (i.e., the larger the minimum dihedral angle), the better the convergence of the linear solver. A parallel Delaunay meshing method is employed to tessellate the segmented brain with high-quality tetrahedral elements and to model the brain surface with geometric and topological guarantees [57]. Single-tissue (i.e., brain parenchyma) and multi-tissue (i.e., brain parenchyma and tumor) meshes are generated. Parameter  $\delta$  determines the size of the mesh, where a smaller  $\delta > 0$  generates a larger mesh.

**Multi-level Dynamic Load Balancing** The imbalance of the processing time across different nodes involved in the computation is caused by our inability or difficulty to predict the processing time required per block of data on a given architecture. The main sources of load imbalance are *platform-dependent*. These are caused by the heterogeneous nature of the PEs. More importantly, some of the resources may be time-shared by multiple users and applications, which affects the processing time unpredictably. The (weighted-) static work assignment is ineffective when some resources operate in the time-shared mode.

We have implemented a multi-level hierarchical dynamic load balancing scheme for parallel block matching. We use an initial rough estimation of the combined computational power of each cluster involved in the computation (based on CPU clock speed) for the weighted partitioning of the work pool and initial work assignment. However, this is a rough “guess” estimation, which is adjusted at runtime using a combination of master/worker and work-stealing [14, 105] methods. Each master maintains an instance of the global work pool. Initially, all these pools are identical. The portion of the work-pool assigned to a specific cluster is partitioned into meta-blocks (a sequence of blocks), which are passed to the cluster nodes using the master-worker model. As soon as all the matches for a meta-block are computed, they are communicated back to the master, and a new meta-block is requested. If the portion of the work pool assigned to a master is processed, the master continues with the “remote” portions of work (i.e., those initially assigned to other clusters). As soon as the processing of a “remote” meta-block is complete, it is communicated to all the other master nodes to prevent duplicated computation.

**Multi-Level Fault Tolerance** Our implementation is completely decoupled, which provides the first level of fault tolerance, i.e., if the failure takes place at any of the stages, we can seamlessly restart just the failed phase of the algorithm and recover the computation. The second level of fault tolerance pertains to the parallel block-matching phase. It is well-known that the vulnerability of parallel computa-

tions to hardware failures increases as we scale the size of the system. We would like to have a robust system that, in case of failure, could continue the parallel block matching without recomputing results obtained before the failure. This functionality is greatly facilitated by maintaining the previously described work-pool data structure which is managed by the master nodes.

The work-pool data structure is replicated on the separate file systems of these clusters and has a tuple for each block center. A tuple can be either empty if the corresponding block has not been processed or it contains the three components of the best match for a given block. The work pool is synchronized periodically between the two clusters, and within each cluster, it is updated by the PEs involved. As long as one of the clusters involved in the computation remains operational, we can sustain the failure of the other computational side and deliver the registration result.

**Ease-of-use** The implementation comprises the client and server components. The client is running at the hospital site and is based on a Web service, which makes it highly portable and easy to deploy. The input data and arguments are transferred to the participating sites on the server side. Currently, we have a single server responsible for this task. The computation uses the participating remote sites to provide the necessary performance and fault tolerance.

**Table 1** Response time (sec) of the intra-surgery part of the CoW-based NRR DDDAS at various stages of development.

Setup	ID						
	1	2	3	4	5	6	7
High-end workstation, using original PVM implementation	1558	1850	2090	2882	2317	2302	3130
SciClone (240 procs), no load-balancing	745	639	595	617	570	550.4	1153
SciClone (240 procs) and CS lab(29 procs), dynamic 2-level load-balancing and fault-tolerance	<b>30</b>	<b>40</b>	<b>42</b>	<b>37</b>	<b>34</b>	<b>33</b>	<b>35</b>

We applied the developed NRR DDDAS for registering seven image datasets acquired at BWH. The computations for two of these seven registration computations were accomplished during the course of surgery (at the College of William and Mary), while the rest of the computations were done retrospectively. All of the intra-operative computations utilized *SciClone* (a heterogeneous cluster of workstations located at CWM, reserved in advance for the registration computation) and the workstations of the student lab (time-shared mode). The details of the hardware configuration can be found in [19]. Data transfer between the networks of CWM and BWH (subnet of Harvard University) is facilitated by the Internet2 backbone network, with the slowest link having a bandwidth of 2.5 Gbps.

The evaluation results are summarized in Table 1. We were able to reduce the total response time to 2 minutes (4 minutes, including the time to transfer the data).

We showed that dynamic load balancing is highly effective in the time-shared environment. The modular structure of the implemented code greatly assisted in the overall usability and reliability of the code. The fault-tolerance mechanisms implemented are essential and introduce a mere 5-10% increase in the execution time.

**Table 2** Performance results for the 6 clinical cases with 1 and 12 threads. The experiments were 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

Almost 10 years later, we performed a more comprehensive experimental evaluation at Old Dominion University using 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 2 from [28] illustrates the end-to-end execution time (including I/O) and 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 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.

The Parallel APBNRR (in short PAPBNRR) refers to the additional parallelization of APBNRR modules added and presented in [28] and bring the performance of the adaptive physics-based method to nearly real-time. Nevertheless, our latest non-rigid registration technology presented in detail in [28] 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, and (ii) it exploits additional parallelism. 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.

According to our study in [28], the combination of algorithms, software, and hardware has significantly improved real-time results in 10 years since we started this project. In 2015, a single high-performance computing workstation with 6 CPUs at ODU (vs. 270 CPUs used at CWM in 2005-6) was on average 2.93 times more accurate than BSpline, 3.12 times more accurate than PBNRR, and 3.78 times more accurate than Rigid registration.

## 4.2 Grid Computing Resources

Before the widespread use of multi-core CPUs and heterogenous workstations (or nodes) the HPC community experimented with Grid Computing, which has since evolved into Cloud Computing. This idea gained popularity around 2008, the same year the first draft of this chapter was written. So during that time, the HPC community invested significant effort in developing standards and software for Grid computing, deploying production grid systems worldwide and porting applications on those systems. One such production system under continuous improvement and development is USA-based TeraGrid [89]. As of May 2007, TeraGrid was connecting 11 high-end computational sites within the USA, providing “...more than 250 TFLOPS of computing capability and more than 30 petabytes of storage” and therefore making TeraGrid ...“ the world’s largest supercomputer (at that point), most comprehensive distributed cyberinfrastructure for open scientific research” [89]. Around 2008, TeraGrid connected 11 computational centers to provide a cumulative peak performance of 1124 teraflops.

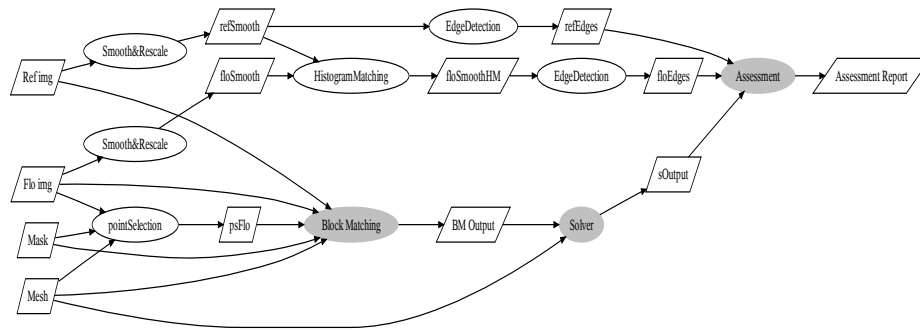
Using the Grid infrastructure for NRR DDDAS had two major advantages. First, the implementation is not restricted to run on a specific cluster resource. With the multiple computing centers participating in TeraGrid, temporary resource outages are more feasible. Second, complex image processing methods, like NRR, often require the proper setting of many parameters to achieve optimum accuracy. Identification of such parameter combinations is a non-trivial task. One approach to selecting the optimum parameter combination is to use *speculative computation* [63], when multiple instances of NRR are computed in parallel with different parameter settings. In this regard, we have developed initial accuracy assessment solutions [38] to facilitate intra-operative speculative NRR over the Grid. In our Grid NRR DDDAS, we leverage the CoW-based implementation, augment it with the automatic error estimation, and develop a framework for the speculative execution of NRR on the TeraGrid.

While TeraGrid resources could be accessed directly for individual job submission and data transfer, doing this manually on a large scale or as part of workflow execution is not practical. We adopted Swift workflow scripting and management system [106] to implement and deploy NRR workflow. Swift has been developed and evaluated to support grid implementations based on Globus Toolkit, which allows to use this system without any modifications to schedule workflows on TeraGrid. SwiftScript, the scripting language used for workflow definition, is a powerful way of abstracting interaction of the processing tasks, which allows defining composite data inputs, and dependencies between the processing tasks and provides familiar control structures like loops and conditional structures, which allow flexible control over workflow definition and execution.

Fault tolerance and dynamic load balancing are important characteristics of NRR DDDAS. Swift implements basic fault tolerance of workflow execution at the individual task level, which is critical for NRR computations. If a particular task fails to deliver the output, Swift will re-schedule its execution, possibly on a different site. The Swift infrastructure also provides task-level load-balancing. The execu-

tion traces for the same computational task are continuously collected and used to dynamically select the best-performing site when the task is scheduled again.

Swift provides the means to define and execute the workflow, which consists of individual processing tasks. Each of the processing tasks must be available as an executable at each of the sites that will be involved in the workflow computation. The details of running a specific task are provided to Swift in the so-called *translation catalog* available at the client (submission) site. The translation catalog contains the identifier of the remote site where the executable is installed and the optional information on its invocation.



**Fig. 6** NRR workflow diagram for single registration execution (shaded are the time-critical components of the workflow).

The NRR workflow diagram of a single NRR procedure, together with the accuracy assessment module, is shown in Figure 6. The block-matching task is parallelized using MPI and deployed on the TeraGrid sites for remote parallel execution. The other workflow components are executed on the local resources (single node of the CWM SciClone cluster). This NRR workflow corresponds to the base case for computation supported by the cluster-based implementation we discussed earlier, augmented with the accuracy assessment module. The accuracy assessment module of the Grid-based NRR DDDAS was developed separately [38]. This module allows to estimate the registration error automatically. The construction of the workflow for speculative execution is straightforward with the scripting capabilities of Swift. This allows us to study some parameters' impact on registration accuracy.


We considered the impact of varying the block size and outlier rejection rate on the accuracy of NRR on retrospective clinical data. Table 3 summarizes the improvement in accuracy evaluated at the expert-selected anatomical landmarks with the optimum combination of these two parameters compared to their default settings. Based on the experimental data, in most cases, good registration accuracy is achieved using the default parameters suggested by Clatz et al. [24]. However, in Case 3, the improvement in registration accuracy was significant. In both cases, however, there were landmark points where registration errors exceeded voxel di-



**Table 3** Absolute improvement in accuracy (mm) evaluated at selected landmarks using optimal values of block size and outlier rejection rate.

Case	1	2	3	4	5	6	7	8	9	10	11	12
1	0.2	0.1	0.2	0.2	0.3	0.1	0.3	0.4	0.0	0.2	0.2	0.3
2	0.2	0.0	0.0	0.0	0.2	0.3	0.3	0.2	0.5	0.3	0.1	0.2
3	0.6	1.0	0.2	2.9	0.0	0.1	0.3	-	-	-	-	-
4	0.9	0.8	0.2	0.4	0.7	0.5	0.7	0.4	0.3	0.7	0.7	-
5	0.0	0.3	0.0	0.0	0.8	0.5	0.4	0.0	0.3	0.4	-	-
6	0.2	0.1	0.2	0.1	2.0	0.1	0.0	0.1	-	-	-	-

mensions. The analyzed data also suggests that the optimum value of outlier rejection is varied in different locations of the image. For example, if we consider landmarks 5 and 12 in Case 1, the optimal combination of the studied parameters differs in each case, as shown in Figure 7.




**Fig. 7** Influence of the block size and rejection rate on landmark error: case 1, landmark 5 (left), and case 1, landmark 12 (right).

### 4.3 Graphical Processing Units

Graphics Processing Unit (GPU), an inexpensive, single-chip, massively parallel architecture, has shown higher throughput and performance per dollar orders than traditional CPUs. In addition, a GPU can be easily deployed in the Operating Room as a co-processor of the CPU without hindering routine surgical operations. Around 2008, researchers have made efforts to accelerate NRR using GPU [94, 86, 100, 70]. However, to satisfy the requirement for the accuracy and real-time in the clinic, a more advanced GPU-based NRR is imperative.

The workflow implementation of NRR DDDAS allows us to use the best parallelization strategies for individual components. The block matching component is embarrassingly parallel, which makes it highly amenable to GPU parallelization.



**Fig. 8** Parallel block matching execution time on CPU vs. GPU (Left) and the effect of GPU thread configuration on the performance of the block matching kernel code (Right).

We use CUDA programming model [91] to develop the GPU implementation of this component. CUDA organizes GPU threads in the grid, which is an array of blocks and each block is an array of threads. The kernel is the core code to be executed on each thread, which performs on different data sets using its ID in a SIMD fashion. CUDA programming model can be treated as two levels loop: block level and thread level. In the following code, the outer loop can be parallelized using GPU on the thread block level. Computation of the similarity metric for an image block (inner loop) is parallelized on GPU thread level, while the similarity metric computation is done on CPU:

- 1: **for** each image block  $fbk$  in floating image **do**
- 2:   define search window  $sw$  in fixed image
- 3:   **for** each image block  $tblk$  in  $sw$  **do**
- 4:     calculate similarity  $s$  between  $fbk$  and  $tblk$
- 5:   **end for**
- 6:   find the maximum  $s$  and corresponding displacement
- 7: **end for**

This GPU-based implementation of block matching can gain a speedup of about 10, as we show in Figure 8, compared to CPU. The speed-up is measured at seven different image block sizes. Figure 8 clearly shows that GPU running time increases linearly as we increase the block size, but CPU exhibits a super linear behavior.

Optimization of GPU codes is particularly important since there are numerous parameters of the execution environment, which can affect performance. Significant evidence exists that there can be orders of magnitude performance differences depending on the level of optimization for GPU implementations [10, 95, 96]. The search space generated by the execution configuration is so large that finding the optimal parameters by trial and error is not practical. Several studies (around 2008) tackled this problem through empirical search-based approaches [96, 84]. We utilize the method provided in [84] to optimize the GPU execution configuration for block matching and improve speedup further, as shown in Figure 8 (right). We observe a speedup of about 30 when comparing optimized and non-optimized implementations.


## 5 Deep Learning Model Architecture

Through the evolution of Grid Computing to Cloud computing and advances in AI, we have made significant progress in improving the accuracy and execution time of NRR by leveraging Machine Learning for Patient-Specific search of NRR parameters. In this section, we summarize our work appeared in [3], where we have presented a model that uses a deep feedforward neural network trained on data from past APBNRR executions. *The model takes fourteen parameters as input, including twelve APBNRR parameters with a very large range of possible values and two additional patient-specific parameters.* The output is a predicted Hausdorff distance of the registered preoperative image. The idea is to replace speculative executions over the Grid (now Cloud) we described earlier in Section 4.2. The twelve APBNRR parameters are the half block size  $\mathbf{B}_s$ , the half window size  $\mathbf{W}_s$  in the axial, coronal, and sagittal direction, the fraction  $\mathbf{F}_s$  to select the image blocks, the number of approximation (outlier rejection) steps  $\mathbf{N}_{rej}$ , the number of interpolation steps  $\mathbf{N}_{int}$ , the maximum number of adaptive iterations  $\mathbf{N}_{(iter,max)}$ , the minimum number of blocks with a zero correspondence  $\mathbf{N}_{(b0,min)}$ , and the percentage of rejected outlier blocks  $\mathbf{F}_r$ . The two patient-specific parameters are the location of the tumor in the brain (lobe-wise) and the degree of brain deformation, which can be directly inferred from the rigid registration error. These two parameters improve model performance by providing additional cues for the neural network to learn and fine-tune the model for a specific patient during an IGNS session.

The neural network was implemented using Keras [18] on a TensorFlow backend [1]. It consists of four hidden, fully connected layers, each comprised of 128 neurons. We used ReLU [7] as the activation function and stochastic gradient descent with Nesterov momentum [93] for optimization. We chose to use SGD because it has been shown to lead to better generalization compared to adaptive gradient optimizers such as Adam, due to its tendency to converge to better global minima [68]. No dropout layers or methods of preventing overfitting were used, as the training data is complex and highly varied. Furthermore, the feature selection in parallel mode is non-deterministic, meaning that the same input parameters can yield different outputs in every execution, further reducing the ability of the model to “memorize” the training data and overfit to the training set. The neural network was built in an iterative fashion, wherein at each iteration, a new patient case was added (consisting of data from about 100,000 APBNRR executions) and the network’s hyperparameters and architecture were tuned to reduce the loss marked in the previous iteration.

### 5.1 Deep Learning APBNRR

The deep learning portion of APBNRR takes place before the actual execution of APBNRR. APBNRR inputs the parameter sets predicted by the deep learning model to result in the lowest Hausdorff distances. A visualization of how the deep learning model works is shown in Figure 9. The neural network is given as input each param-



**Fig. 9** Deep learning portion of APBNRR registration.

eter set in a pool consisting of patient-specific parameter sets, which was produced by augmenting a base, general parameter set pool. We have created tools that do this automatically. The neural network iterates through each parameter set and outputs the Hausdorff distance of the registered image that APBNRR would produce if this parameter set were utilized. Out of all those predictions, the lowest ones are compiled in a file and can be used as input to APBNRR. This process takes about 15 seconds.

The deep learning model takes input parameter sets from a pool specific to each patient. This is an augmented version of a base, general parameter pool, which has been enhanced to include two patient-specific features: the tumor’s location in the brain and the size of the deformation, as can be derived from the rigid registration error. In our experiments, the patient-specific features have yielded significantly better results.

We encountered a couple of important challenges while constructing the deep learning model. The first is that APBNRR is non-deterministic due to the parallel feature selection algorithm. Given two identical parameter sets, APBNRR will yield two registered images with different Hausdorff distances. This non-determinism leads to a degree of “randomness” in results, which hinders the ability of the neural network to predict the Hausdorff distances correctly. Unfortunately, there is no practical solution to this issue other than amassing more training data to “average out” this randomness, leading us to the second challenge: collecting training data takes a long time. To gather training data, we must run APBNRR for every parameter set in an exhaustive pool for every additional patient case we want to include in our training set. Evaluating a single case exhaustively ( 1 million executions) takes about 1-2 months, depending on the severity and size of the brain tumor.

## 5.2 Final Experimental Evaluation

This section summarizes our findings in [3] and a summary in [32]. Our data set for the model consists of medical resonance images from thirteen patient cases provided by the Huashan Worldwide Medical Center and Brigham and Women’s Hospital of Harvard University and output data from over 2.6 million executions of APBNRR, collected over five months. Eleven of the thirteen cases were used for training ( 2.4 million parameter sets) and two for evaluation ( 200,000 parameter sets). APBNRR was executed with arrays of 120 parameter sets on a supercomputing cluster, utilizing 450 CPU cores and over 4 TB of RAM. The total computing time required for the executions was approximately 3600 hours.

We used the deep learning model to achieve a training root mean squared error (RMSE) of 1.41 and an evaluation RMSE of 1.21 for predicted Hausdorff distances. For further evaluation, utilizing the trained model, we again executed APBNRR for the 13 cases, using the top 120 parameter sets predicted by the model for each case to yield the lowest Hausdorff distances. The best results of those executions are displayed in Table of Figure 10, along with the results from several other registration methods as a comparison. The best results from the APBNRR parameter sweeps that were used for generating our dataset are also displayed as a reference of what could have been the lowest Hausdorff distance value.

Case	Type	Tumor Location	Rigid Registration	B-Spline Registration	PBNRR	APBNRR (Default Parameters)	APBNRR (Deep Learning)	APBNRR (Parameter Sweep)
1	PTR	Left frontal lobe	17.00	17.00	16.49	5.69	2.78	2.78
2	PTR	Left frontal lobe	10.59	5.28	10.76	2.30	2.64	1.77
3	PTR	Left parietal lobe	16.15	13.78	15.12	4.60	2.40	1.70
4	PTR (deep)	Left parietal-occipital lobes	29.93	21.34	27.76	2.83	1.84	1.46
5	CTR	Frontal-temporal lobes	25.51	25.18	22.50	4.97	2.64	2.64
6	CTR	Right frontal lobe	5.59	5.59	3.43	3.09	2.78	2.64
7	CTR	Right temporal lobe	17.90	16.94	15.56	4.11	2.94	2.00
8	CTR	Left posterior-temporal lobe	18.85	17.49	17.38	3.57	3.29	3.24
9	CTR	Left frontal lobe	17.14	7.48	15.41	4.25	2.40	1.84
10	ETR (deep)	Left frontal lobe	35.72	27.77	33.57	3.71	2.06	1.77
11	CTR	Left frontal lobe	25.72	25.72	23.90	3.42	2.64	2.30
12	PTR	Right frontal lobe	26.89	15.86	26.89	4.00	2.85	2.21
13	CTR	Left frontal lobe	19.24	14.61	19.89	2.40	2.00	1.48

**Fig. 10** Shows the thirteen patient cases that consist of our data set and the results (measured as the Hausdorff distance in mm) achieved with various methods of registration, including with APBNRR using deep learning and APBNRR using a parameter sweep. Cases with numbers 1-11 were used for training, and 12, and 13 were used for evaluation.)

On average, APBNRR with deep learning is 8.45 times better than rigid registration, 6.71 times better than B-Spline registration, and 7.9 times better than PBNRR. Overall, APBNRR with deep learning leads to superior results (for tested

cases) than any of the registration methods noted above. Choosing the correct parameters for medical image registration is a difficult task, as there are many (usually infinite) possible values and combinations of parameters that can lead to better or worse results. The deep learning portion of the APBNRR framework makes this faster and easier by greatly limiting the set of possible optimal parameters for each patient, moving APBNRR one step closer to being utilized in a real-time setting where registration accuracy and speed are critical.

## 6 Discussion

We have briefly described the use of Dynamic Data Driven Application Systems for Image-Guided Neurosurgery, enabled by the advances in medical image acquisition, parallel/distributed computing, Machine Learning, and algorithmic changes in [27, 28, 35] the original NRR method presented in [24]. This DDDAS concept, *for the first time in clinical practice [19, 4], helped us to complete and present non-rigid registration results to neuro-surgeons [4] at BWH during tumor resection procedures using image landmark tracking across the entire brain volume.* This is accomplished in three phases:

- **Phase I:** Reduced the total response time of the time-critical computation component to about 35 sec, delivering an effective speedup of nearly 100 compared to the original implementation [24]. To achieve this, we used remotely (at CWM) several CoWs with a total of 269 processors [19] (the data transfer from CWM to the operating room in BWH takes between 3 to 5 min). Our data suggest that we improved the accuracy of the NRR method by performing speculative execution [40] on the TeraGrid, which was capable of delivering about 250 TFLOPS. This computing power translates into tens of thousands of registrations (with different parameters) in almost real-time if there was proper coordination with all sides to avoid scheduling conflicts. However, this requires the availability of a network connection between the operating room and remote computing resources. Also, the imaging data must be anonymized before transfer to address confidentiality concerns.
- **Phase II:** Eliminated the use of Grid and lately Cloud computing for the speculative execution of PBNRR by using much cheaper hardware that can be located nearby the operating room. Our results in the CRTIC indicate that it is possible to complete the time-critical component of non-rigid registration within a minute—save another 3 to 5 minutes for the data transfer— using a single (or two, for fault-tolerance) high-end workstations with NVIDIA GeForce 8800 GT GPU and 2 x Intel Core2 Duo CPU 3.16GHz. We believe in using current and emerging heterogeneous hardware architectures and the coordinated use of Cloud (or TeraGrid when performing these studies). Our results show that GPU provides excellent computing capabilities without sacrificing the result's accuracy.
- **Phase III:** Improved the accuracy of the original NRR method by extending the NRR model to incorporate adaptivity to improve accuracy and replaced the

expensive speculative execution of multiple NRRs (from Phase I & II) with machine learning to improve both accuracy and performance. Our results indicate that the APBNRR with deep learning is 8.45 times better than rigid registration, 6.71 times better than B-Spline registration, and 7.9 times better than PBNRR (outcome of Phase I and II). A key element for evaluating the performance of APBNRR is determining optimal input parameters. Optimal input parameters are usually in our base parameter pool's first 800 – 39,000 parameter-sets. Determining optimal input parameters by doing a parameter sweep utilizing arrays of 120 jobs over the Cloud would take an estimated 33 minutes – 27 hours. This high variability is not acceptable in a neurosurgery session. Deep learning offers significant speedups in this area by allowing the user to limit the parameter pool to a custom value. By using deep learning, we could evaluate fewer parameter sets, potentially completing APBNRR in 5 minutes on moderate CoWs (5 to 10 HPC nodes) while also producing good results.

This Chapter presented only a few of the most important findings and results. However, several efforts at CRTIC alone address other aspects like: (1) robust and real-time image-to-mesh (I2M) conversion [44, 92, 64, 65, 42, 21, 43, 47, 48, 59, 77, 8, 15, 16, 51, 50, 53, 52, 54, 78, 17, 55, 56, 49, 57, 29, 58, 31, 46, 29, 30, 60], (2) more accurate computation of Hausdorff Distance [13, 41], (3) parallel Euclidean Distance Transform, which is critical for parallel I2M conversion [98], and many preliminary efforts that lead to the findings presented in this Chapter and we list them here for completeness without expanding on them [6, 45, 22, 20, 74, 75, 104, 76, 83, 72, 81, 79, 73, 80, 82, 28, 61, 33, 85, 66, 34, 35, 12, 32]

## 7 Conclusion and Future Work

Choosing the correct parameters for medical image registration is a difficult task, as there are many (usually infinite) possible values and combinations of parameters that can lead to better or worse results. The deep learning portion of the APBNRR framework makes this easier by greatly limiting the set of possible optimal parameters for each patient, moving APBNRR one step closer to being utilized in a real-time setting where registration accuracy and speed are critical.

Operating rooms, like Advanced Multi-modality Image Guided Operating (AMIGO) [87] suite, provide new capabilities to improve intra-operative image guidance. Advances in high-performance and distributed tools for image analysis, like the NRR DDDAS we presented in this chapter, will be essential to meet the ever-increasing computational demands of such environments. DDDAS will be critical in health care, among other areas, where this concept proved successful [88, 37].

Within AMIGO's workflow, APBNRR execution takes place after the intraoperative MR phase and is meant to assist in tumor and residual tumor assessment due to the ability of APBNRR to perform well in the context of tumor resection. The registered intraoperative image with the tumor resected allows the neurosurgeon to evaluate better how well the tumor has been resected. The time window

for APBNRR to execute as part of this workflow is a few minutes. As such, on top of accuracy, speed is also critical. Deep learning solves the problem of limiting the parameter pool that needs to be evaluated, but some issues with APBNRR need to be resolved. One of the most important is the dependence of APBNRR performance on the tumor size. The larger the tumor, the more time it takes for APBNRR to execute. Furthermore, the more adaptive iterations APBNRR goes through, the longer the execution. We have seen an execution time ranging from 5 to 30 minutes –for large and deep brain tumors, with parameter sets having more adaptive iterations and slower execution than those with lower ones but better registration accuracy on average. This suggests we still need to evaluate deep learning APBNRR on how it would perform in a real-world setting. Specifically, we will look at how APBNRR would perform in the brain tumor resection workflow in AMIGO, which represents an ideal setting for using APBNRR.

Deep learning APBNRR takes us a step closer to enabling NRR to be used in real-world scenarios. However, some issues must be resolved before that can happen. First, more training data must be collected for the deep learning model to offer more accurate predictions. Second, work needs to be done to enable the deep learning model to generate a parameter pool that is limited in size and can also be evaluated rapidly. Finally, the performance of APBNRR needs to be further improved and extensively evaluated for deep tumors which involve very large brain deformation. Finally, the overall performance of deep learning APBNRR depends on the computational resources available. In particular, the more computational resources available, the larger the parameter pool of the best parameter sets predicted by the neural network can be. As a result, there is a greater chance of achieving better accuracy and performance.

## Acknowledgments

The Richard T. Cheng Endowment supports NC. AF, YL, AK, and NC were partly supported by the NSF grants CCS-0719929, CCS-0750901, CCF-0833081, CCF-1439079, and the John Simon Guggenheim Memorial Foundation. AF was supported in part by a grant from Brain Science Foundation. AF was supported in part by a grant from Brain Science Foundation. RK, AG, and PB were supported in part by NIH grants: U41 RR019703, P01 CA67165, NIH R01NS049251, NIH 5R01EB027134, and NCIGT P41EB015898. This research was supported by an allocation through the TeraGrid Advanced Support Program. This work was performed [in part] using computational facilities at the College of William and Mary and Old Dominion University, which were enabled by grants from Sun Microsystems, the National Science Foundation, and Virginia's Commonwealth Technology Research Fund. We also thank Dr. Frederica Darema for her hard work laying the foundation, building a DDDAS community, and encouraging us to present our work in this Volume.



## References

1. Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283, 2016.
2. American Brain Tumor Association. Facts & statistics, 2008. <http://www.abta.org/siteFiles/SitePages/5E8399DBEEA8F53CBBBBF212C63AE113.pdf>, accessed 23 Dec 2008.
3. Angelos Angelopoulos and Nikos Chrisochoides. Deep learning real-time adaptive physics-based non-rigid registration for accurate geometry representation of brain in modeling deformation during glioma resection. 2019.
4. N. Archip, O. Clatz, S. Whalen, D. Kacher, A. Fedorov, A. Kot, N. Chrisochoides, F. Jolesz, A. Golby, P.M. Black, and S.K. Warfield. Non-rigid alignment of preoperative MRI, fMRI, and DT-MRI with intra-operative MRI for enhanced visualization and navigation in image-guided neurosurgery. *NeuroImage*, 35:609–624, 2007.
5. N. Archip and I.-F. Talos. Spl brain tumor resection image dataset, 2007. <http://www.spl.harvard.edu/publications/item/view/541>.
6. Neculai Archip, Andriy Fedorov, Bryn Lloyd, Nikos Chrisochoides, Alexandra Golby, Peter Black, and Simon K. Warfield. Integration of patient specific modeling and advanced image processing techniques for image-guided neurosurgery. In *SPIE Medical Imaging*, pages 422–429, 2006.
7. Raman Arora, Amitabh Basu, Poorya Mianjy, and Anirbit Mukherjee. Understanding deep neural networks with rectified linear units, 2018.
8. Michel Audette, Andrey Chernikov, and Nikos Chrisochoides. A review of mesh generation for medical simulators. 2011. Chapter in *Handbook of Real-World Applications in Modeling and Simulations*.
9. Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, Adrian V. Dalca, and John Guttag. An unsupervised learning model for deformable medical image registration. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, jun 2018.
10. M. M. Baskaran, U. Bondhugula, S. Krishnamoorthy, J. Ramanujam, A. Rountev, and P. Sadayappan. A compiler framework for optimization of affine loop nests for gpgpus. In *ICS'08: Proceedings of the 22nd Annual International Conference on Supercomputing*, page 225–234, 2008.
11. A. Belguelin, J. Dongarra, A. Geist, R. Manchek, S. Otto, and J. Walpore. Pvm: Experiences, current status, and future direction. In *Supercomputing '93 Proceedings*, pages 765–766, 1993.
12. Joi Best, Kevin Garner, Daming Feng, Fotios Drakopoulos, Yixun Liu, and Nikos Chrisochoides. Image-to-mesh conversion tool for image-driven simulations (paper). In *Virginia Space Grant Consortium*, 2019.
13. Eric Billet, Andriy Fedorov, and Nikos Chrisochoides. The use of robust local hausdorff distances in accuracy assessment for image alignment of brain mri. *Insight Journal*, 2008.
14. R. Blumofe, C. Joerg, B. Kuszmaul, C. Leiserson, K. Randall, and Y. Zhou. Cilk: An efficient multithreaded runtime system. In *Proceedings of the 5th Symposium on Principles and Practice of Parallel Programming*, pages 55–69, 1995.
15. Andrey Chernikov and Nikos Chrisochoides. Multi-tissue tetrahedral image-to-mesh conversion with guaranteed quality and fidelity. *SIAM Journal on Scientific Computing*, 33:3491–3508, 2011.
16. Andrey Chernikov and Nikos Chrisochoides. Tetrahedral image-to-mesh conversion for biomedical applications. In *ACM Conference on Bioinformatics, Computational Biology and Biomedicine*, pages 125–134, Chicago, IL, August 2011.
17. Andrey Chernikov, Panagiotis Foteinos, Yixun Liu, Michel Audette, Andinet Enquobahrie, and Nikos Chrisochoides. Tetrahedral image-to-mesh conversion approaches for surgery simulation and navigation. *Image-Based Geometric Modeling and Mesh Generation*, pages

- 69–84, 2013. Lecture Notes in Computational Vision and Biomechanics (Yongjie (Jessica) Zhang, ed.).
18. Francois Chollet et al. Keras, 2015.
  19. N. Chrisochoides, A. Fedorov, A. Kot, N. Archip, P.M. Black, O. Clatz, A. Golby, R. Kikinis, and S.K. Warfield. Toward real-time image guided neurosurgery using distributed and Grid computing. In *Proc. of IEEE/ACM SC06*, 2006.
  20. Nikos Chrisochoides. Real-time non-rigid registration for igns: Mesh generation. Invited Joint Seminar with Waterloo Institute for Health Informatics Research and The Center f, March 2007.
  21. Nikos Chrisochoides. Parallel mesh generation for medical image computing. In *SIAM Conference on Parallel Processing for Scientific Computing*. SIAM, March 2008.
  22. Nikos Chrisochoides, Andriy Fedorov, Andriy Kot, Neculai Archip, Daniel Goldberg-Zimring, Dan Kacher, Stephen Whalen, Ron Kikinis, Ferenc Jolesz, Olivier Clatz, Simon K. Warfield, Peter Black, and Alexandra Golby. Grid-enabled software environment for enhanced dynamic data-driven visualization and navigation during image-guided neurosurgery. In *International Conference on Computational Science (ICCS 2007)*, pages 980–987, 2007.
  23. G.E. Christensen, M.I. Miller, M.W. Vannier, and U. Grenander. Individualizing neuroanatomical atlases using a massively parallel computer. *Computer*, 29(1):32–38, 1996.
  24. O. Clatz, H. Delingette, I.-F. Talos, A. Golby, R. Kikinis, F. Jolesz, N. Ayache, and S.K. Warfield. Robust non-rigid registration to capture brain shift from intra-operative MRI. *IEEE Trans. Med. Imag.*, 24(11):1417–1427, 2005.
  25. H. Delingette and N. Ayache. *Soft tissue modeling for surgery simulation*, volume XII of *Handbook of Numerical Analysis: Special volume: Computational models for the human body*, pages 453–550. Elsevier, Netherlands, 2004.
  26. C. DeLorenzo. *Image-Guided Intraoperative Brain Deformation Recovery*. PhD thesis, Yale University, 2007.
  27. Fotios Drakopoulos and Nikos Chrisochoides. A parallel adaptive physics-based non-rigid registration framework for brain tumor resection. In *VMASC 2014 Capstone Conference*, April 2014. Best paper in Medical Simulations.
  28. Fotios Drakopoulos and Nikos Chrisochoides. Accurate and fast deformable medical image registration for brain tumor resection using image-guided neurosurgery. *Computer Methods in Biomechanics and Biomedical Engineering*, 2015. <http://dx.doi.org/10.1080/21681163.2015.1067869>.
  29. Fotios Drakopoulos and Nikos Chrisochoides. Tetrahedral image-to-mesh conversion for anatomical modeling and surgical simulations. In *VMASC 2015 Capstone Conference*, 2015.
  30. Fotios Drakopoulos and Nikos Chrisochoides. Lattice-based multi-tissue mesh generation for biomedical applications. Brigham and Women’s Hospital & Harvard Medical School, April 2016. SPL25.
  31. Fotios Drakopoulos, Ricardo Ortiz, Andinet Enquobahrie, and Nikos Chrisochoides. Image-to-mesh conversion for arteriovenous malformation surgical simulators. In *13th U.S. National Congress on Computational Mechanics*, July 2015. Biomechanics Modeling: Advances and Applications to Real-World Problems.
  32. Fotios Drakopoulos, Christos Tsolakis, Angelos Angelopoulos, Yixun Liu, Chengjun Yao, Kyriaki Rafailia Kavazidi, Nikolaos Foroglou, Andrey Fedorov, Sarah Frisken, Ron Kikinis, Alexandra Golby, and Nikos Chrisochoides. Adaptive physics-based non-rigid registration for immersive image-guided neuronavigation systems. *Frontiers in Digital Health*, February 2021.
  33. Fotios Drakopoulos, Michael Weissberger, Kathryn Holloway, and Nikos Chrisochoides. Biomechanical deformable registration for deep brain stimulation. In *VMASC 2016 Capstone Conference*, Suffolk, VA, April 2016.
  34. Fotios Drakopoulos, Chengjun Yao, Yixun Liu, and Nikos Chrisochoides. An evaluation of adaptive biomechanical non-rigid registration for brain glioma resection using image-guided neurosurgery. In *MICCAI Computational Biomechanics for Medicine CBM XI Workshop*, Athens, Greece, October 2016.

35. Fotios Drakopoulos, Chengjun Yao, Yixun Liu, and Nikos Chrisochoides. An evaluation of adaptive biomechanical non-rigid registration for brain glioma resection using image-guided neurosurgery. *Computational Biomechanics for Medicine*, pages 111–122, May 2017.
36. Tom Eelbode, Jeroen Bertels, Maxim Berman, Dirk Vandermeulen, Frederik Maes, Raf Bisschops, and Matthew B. Blaschko. Optimization for medical image segmentation: Theory and practice when evaluating with dice score or jaccard index. *IEEE Transactions on Medical Imaging*, 39(11):3679–3690, nov 2020.
37. Sai Ravela Alex J. Aved Erik P. Blasch, Frederica Darema. *Handbook of Dynamic Data Driven Applications Systems*. Springer Cham, 2022.
38. A. Fedorov, E. Billet, M. Prastawa, A. Radmanesh, G. Gerig, R. Kikinis, S. K. Warfield, and N. Chrisochoides. Evaluation of brain MRI alignment with the robust Hausdorff distance measures. In *Proc. of ISVC 2008*, pages 594–603, 2008.
39. A. Fedorov and N. Chrisochoides. Tetrahedral mesh generation for non-rigid registration of brain MRI: Analysis of the requirements and evaluation of solutions. In *Proc. of the 17th International Meshing Roundtable*, pages 55–72, 2008.
40. A. Fedorov, B. Clifford, S.K. Warfield, R. Kikinis, and N. Chrisochoides. Non-rigid registration for image-guided neurosurgery on the TeraGrid: A case study. Technical Report WM-CS-2009-05, Department of Computer Science, College of William and Mary, 2009.
41. Andriy Fedorov, Eric Billet, Marcel Prastawa, Alireza Radmanesh, Guido Gerig, Ron Kikinis, Simon K. Warfield, and Nikos Chrisochoides. Evaluation of brain mri alignment with the robust hausdorff distance measures. In *4th International Symposium on Visual Computing*, pages 594–603. Springer, 2008.
42. Andriy Fedorov and Nikos Chrisochoides. Adaptive mesh refinement for non-rigid registration of brain mri. In *8th World Congress on Computational Mechanics*, 2008.
43. 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 17, pages 55–72, October 2008.
44. Andriy Fedorov, Nikos Chrisochoides, Ron Kikinis, and Simon K. Warfield. Tetrahedral mesh generation for medical imaging. In *8th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2005)*, 2005.
45. Andriy Fedorov, Nikos Chrisochoides, Ron Kikinis, and Simon K. Warfield. An evaluation of three approaches to tetrahedral mesh generation for deformable registration of brain mr images. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 658–661, 2006.
46. Daming Feng, Christos Tsolakis, Andrey Chernikov, and Nikos Chrisochoides. Scalable 3d hybrid parallel delaunay image-to-mesh conversion algorithm for distributed shared memory architectures. In *24th International Meshing Roundtable*, Austin, Texas, October 2015.
47. Panagiotis Foteinos, Andrey Chernikov, and Nikos Chrisochoides. Toward real time image to mesh conversion for non rigid registration. In *IEEE International Conference on Bioinformatics & Biomedicine*, November 2009. refereed.
48. Panagiotis Foteinos, Andrey Chernikov, and Nikos Chrisochoides. Guaranteed quality tetrahedral delaunay meshing for medical images. In *7th International Symposium on Voronoi Diagrams in Science and Engineering*, pages 215–223, Laval University, Quebec City, Canada, June 2010.
49. Panagiotis Foteinos, Andrey Chernikov, and Nikos Chrisochoides. Guaranteed quality tetrahedral delaunay meshing for medical images. *Computational Geometry: Theory and Applications*, 47:539–562, 2014.
50. Panagiotis Foteinos and Nikos Chrisochoides. A dynamic parallel 3d delaunay triangulation. In *International Meshing Roundtable*, pages 9–26, October 2011.
51. Panagiotis Foteinos and Nikos Chrisochoides. High-quality multi-tissue mesh generation for finite element analysis. In *MeshMed, Workshop on Mesh Processing in Medical Image Analysis (MICCAI)*, pages 18–28, September 2011.
52. Panagiotis Foteinos and Nikos Chrisochoides. Dynamic parallel 3d delaunay triangulation. In *International Meshing Roundtable*, pages 3–20, Paris, France, 2012. Springer Berlin Heidelberg.

53. Panagiotis Foteinos and Nikos Chrisochoides. High-quality multi-tissue mesh generation for finite element analysis. *Image-Based Geometric Modeling and Mesh Generation*, pages 161–172, 2012. Editor Yongjie (Jessica) Zhang.
54. Panagiotis Foteinos and Nikos Chrisochoides. Multi-dimensional image-to-mesh conversion on massively parallel systems. In *Challenges in Massively Parallel Simulations using Unstructured Meshes, SIAM Conference on Parallel Processing for Scientific Computing*, February 2012.
55. Panagiotis Foteinos and Nikos Chrisochoides. 4d space-time delaunay meshing for medical images. In *International Meshing Roundtable*, pages 223–240. Springer International, 2013.
56. Panagiotis Foteinos and Nikos Chrisochoides. High quality real-time image-to-mesh conversion for finite element simulations. In *27th ACM International Conference on Supercomputing (ICS'13)*, pages 233–242, Eugene, Oregon, June 2013.
57. Panagiotis Foteinos and Nikos Chrisochoides. High quality real-time image-to-mesh conversion for finite element simulations. *Journal on Parallel and Distributed Computing*, 74:2123–2140, 2014.
58. Panagiotis Foteinos and Nikos Chrisochoides. 4d space-time delaunay meshing for medical images. *Engineering with Computers*, 31:499–511, July 2015. doi:10.1007/s00366-014-0380-z.
59. Panagiotis Foteinos, Yixun Liu, Andrey Chernikov, and Nikos Chrisochoides. An evaluation of tetrahedral mesh generation for non-rigid registration of brain mri. In *Computational Biomechanics for Medicine V, 13th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) Workshop*, pages 126–137, September 2010.
60. Kevin Garner, Fotios Drakopoulos, and Nikos Chrisochoides. Towards distributed 3d adaptive anisotropic image-to-mesh conversion for biomedical applications. In *NanoBioTech 2023*, July 2023.
61. Kevin Garner, Daming Feng, Fotios Drakopoulos, Yixun Liu, and Nikos Chrisochoides. Image-to-mesh conversion tool. In *24th International Meshing Roundtable*, October 2015.
62. A.J. Golby, R.A. Poldrack, J. Illes, D. Chen, J.E. Desmond, and J.D. Gabrieli. Memory lateralization in medial temporal lobe epilepsy assessed by functional MRI. *Epilepsia*, 43(8):855–863, 2002.
63. F. Ino, K. Ooyama, and K. Hagihara. A data distributed parallel algorithm for nonrigid image registration. *Parallel Computing*, 31(1):19–43, 1 2005.
64. Bhautik Joshi, Andriy Fedorov, Nikos Chrisochoides, Simon K. Warfield, and Sebastien Ourselin. Application-driven quantitative assessment of approaches to mesh generation. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1160–1163, 2007.
65. Bhautik Joshi, Andriy Fedorov, Nikos Chrisochoides, Simon K. Warfield, and Sebastien Ourselin. A quantitative assessment of approaches to mesh generation for surgical simulation. *Engineering with Computers*, 24:417–430, 2008.
66. Melina Kazakidi, Fotios Drakopoulos, Chander Sadasivan, Nikos Chrisochoides, John Eka-terinaris, and Baruch B. Lieber. Numerical simulation of cerebral aneurysm by flow diversion. In *European Congress on Computational Methods in Applied Sciences and Engineering*, June 2016. Workshop on Simulation of Cardiovascular Procedures and Devices.
67. Justin Ker, Lipo Wang, Jai Rao, and Tchoyoson Lim. Deep learning applications in medical image analysis. *IEEE Access*, 6:9375–9389, 2018.
68. Nitish Shirish Keskar and Richard Socher. Improving generalization performance by switching from adam to sgd, 2017.
69. Julian Krebs, Tommaso Mansi, Boris Mailhé, Nicholas Ayache, and Hervé Delingette. Un-supervised probabilistic deformation modeling for robust diffeomorphic registration, 2018.
70. D. Levin, D. Dey, and P. Slomka. Acceleration of 3d, nonlinear warping using standard video graphics hardware: implementation and initial validation. *Comput Med Imaging Graph*, 28(8):471–483, 2004.
71. H. Lippman and F. Kruggel. Quasi-real-time neurosurgery support by MRI processing via grid computing. *Neurosurg. Clin. N. Am.*, 16(1):65–75, 2005.

72. Yixun Liu and Nikos Chrisochoides. Using k-means clustering and mi for non-rigid registration of mri and ct. In *International Workshop on Machine Learning in Medical Imaging (MLMI), the 13th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, September 2010.
73. Yixun Liu and Nikos Chrisochoides. Heterogeneous biomechanical model on correcting brain deformation induced by tumor resection. In *Computational Biomechanics for Medicine Workshop CBM7*, 2012.
74. 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*, pages 401–404, November 2009.
75. Yixun Liu, Andriy Fedorov, Ron Kikinis, and Nikos Chrisochoides. Non-rigid registration for brain mri: Faster and cheaper. *International Journal of Functional Informatics and Personalized Medicine (IJFIPM)*, 3:48–57, 2010.
76. Yixun Liu, Andriy Fedorov, Ron Kikinis, and Nikos Chrisochoides. A novel point based non-rigid registration method and its application for brain shift. In *SPIE Medical Imaging*, 2010.
77. Yixun Liu, Panagiotis Foteinos, Andrey Chernikov, and Nikos Chrisochoides. Multi-tissue mesh generation for brain images. In *International Meshing Roundtable*, number 19, pages 367–384, October 2010.
78. Yixun Liu, Panagiotis Foteinos, Andrey Chernikov, and Nikos Chrisochoides. Mesh deformation-based multi-tissue mesh generation for brain images. *Engineering with Computers*, 8:305–318, October 2012.
79. Yixun Liu, Andriy Kot, Fotios Drakopoulos, Andriy Fedorov, Andinet Enquobahrie, Olivier Clatz, and Nikos Chrisochoides. An itk implementation of physics-based non-rigid registration method. *Insight Journal*, 2012.
80. Yixun Liu, Andriy Kot, Fotios Drakopoulos, Andriy Fedorov, Chengjun Yao, Andinet Enquobahrie, Olivier Clatz, and Nikos Chrisochoides. An itk implementation of a physics-based non-rigid registration method for brain deformation in image-guided neurosurgery. *Front. Neuroinform.*, 8, February 2014. doi: 10.3389/fninf.2014.00033.
81. Yixun Liu, Hui Xue, Christoph Guetter, Marie-Pierre Jolly, Nikos Chrisochoides, and Jens Guehring. Moving propagation of suspicious myocardial infarction from delayed enhanced cardiac imaging to cine mri using hybrid image registration. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1284–1288, March 2011.
82. Yixun Liu, Chengjun Yao, Fotios Drakopoulos, Jinsong Wu, Liangfu Zhou, and Nikos Chrisochoides. A non-rigid registration method for correcting brain deformation induced by tumor resection. *Medical Physics*, 41, August 2014. doi: 10.1118/1.4893754.
83. Yixun Liu, Chengjun Yao, Liangfu Zhou, and Nikos Chrisochoides. A point based non-rigid registration for tumor resection using imri. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1217–1220, April 2010.
84. Yixun Liu, Eddy Z. Zhang, and Xipeng Shen. A cross-input adaptive framework for gpu programs optimization. In *23rd IEEE IPDPS*, 2009. accepted.
85. Shahram Mohrehkesh, Andriy Fedorov, Arun Brahmavar Vishwanatha, Fotios Drakopoulos, Ron Kikinis, and Nikos Chrisochoides. Large scale cloud-based deformable registration for image guided therapy. In *International Workshop on Big Data Analytics for Smart and Connected Health*, Washington DC, June 2016.
86. Pinar Muyan-Ozcelik, John D. Owens, Junyi Xia, and Sanjiv S. Samant. Fast deformable registration on the gpu: A cuda implementation of demons. In *ICCSA '08: Proceedings of the 2008 International Conference on Computational Sciences and Its Applications*, pages 223–233, Washington, DC, USA, 2008. IEEE Computer Society.
87. National Center for Image Guided Therapy. Advanced Multimodality Image Guided Operating (AMIGO) Suite, 2009. <http://www.ncigt.org/pages/AMIGO>.
88. National Science Foundation. DDDAS: Dynamic Data Driven Applications Systems, 2009. [http://www.nsf.gov/cise/cns/dddas/DDDAS\\_Appendix.jsp](http://www.nsf.gov/cise/cns/dddas/DDDAS_Appendix.jsp).
89. National Science Foundation. The TeraGrid project, 2009. <http://www.teragrid.org/>.

90. Fotis Drakopoulos Yixun Liu Christos Tsolakis Emmanuel Billias Olivier Clatz Nicholas Ayache Alex Golby Peter Black Nikos Chrisochoides, Andrey Fedorov and Ron Kikinis. Real-time dynamic data driven deformable registration for image-guided neurosurgery: Computational aspects. *Frontiers in Digital Health*, 2023.
91. NVIDIA Corporation. NVIDIA CUDA programming guide, 2008. Version 2.1.
92. Maria-Cecilia Rivara, Carlo Calderon, Andriy Fedorov, and Nikos Chrisochoides. Parallel decoupled terminal-edge bisection method for 3d mesh generation. *Engineering with Computers*, 22:111–119, 2006.
93. Sebastian Ruder. An overview of gradient descent optimization algorithms, 2017.
94. Antonio Ruiz, Manuel Ujaldon, Lee Cooper, and Kun Huang. Non-rigid registration for large sets of microscopic images on graphics processors. *J Sign Process Syst*, 55(1-3):229–250, April 2008.
95. S. Ryoo, C. I. Rodrigues, S. S. Baghsorkhi, S. S. Stone, D. B. Kirk, and W. W. Hwu. Optimization principles and application performance evaluation of a multithreaded gpu using cuda. In *PPoPP'08: Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, page 73–82, 2008.
96. S. Ryoo, C. I. Rodrigues, S. S. Stone, S. S. Baghsorkhi, S. Ueng, J. A. Stratton, and W. W. Hwu. Program optimization space pruning for a multithreaded gpu. In *CGO'08: Proceedings of the Sixth Annual IEEE/ACM International Symposium on Code Generation and Optimization*, page 195–204, 2008.
97. M. Snir, S. Otto, S. Huss-Lederman, and D. Walker. *MPI The Complete Reference*. The MIT Press, 1998.
98. Robert Staubs, Andriy Fedorov, Leonidas Linardakis, Benjamin Dunton, and Nikos Chrisochoides. Parallel n-dimensional exact signed euclidean distance transform. *Insight Journal*, 2006.
99. R. Stefanescu, X. Pennec, and N. Ayache. Parallel non-rigid registration on a cluster of workstations. In *Proc. of HealthGrid'03*, 2003.
100. C. Vetter, C. Guetter, C. Xu, and R. Westermann. Non-rigid multi-modal registration on the gpu. In *Medical Imaging 2007: Image Processing*, volume 6512, page 651228, 2007.
101. S. K. Warfield, M. Ferrant, X. Gallez, A. Nabavi, and F. A. Jolesz. Real-time biomechanical simulation of volumetric brain deformation for image guided neurosurgery. In *Supercomputing '00: Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM)*, page 23, Washington, DC, USA, 2000. IEEE Computer Society.
102. S. K. Warfield, F. A. Jolesz, and R. Kikinis. Real-time image segmentation for image-guided surgery. In *Supercomputing '98: Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)*, pages 1–7, Washington, DC, USA, 1998. IEEE Computer Society.
103. S. K. Warfield, A. Nabavi, T. Butz, K. Tuncali, S. G. Silverman, P. Black, F. A. Jolesz, and R. Kikinis. Intraoperative segmentation and nonrigid registration for image guided therapy. In *MICCAI '00: Proceedings of the Third International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 176–185, London, UK, 2000. Springer-Verlag.
104. Michael Weissberger, Yixun Liu, Joseph Han, and Nikos Chrisochoides. Multi-modal non-rigid registration for image guided head and neck surgery. In *American Rhinologic Society Fall meeting*, 2010.
105. I. Wu. *Multilist Scheduling: A New Parallel Programming Model*. PhD thesis, School of Comp. Sci., Carnegie Mellon University, Pittsburg, PA 15213, July 1993.
106. Y. Zhao, M. Hategan, B. Clifford, I. Foster, G. von Laszewski, V. Nefedova, I. Raicu, T. Stef-Praun, and M. Wilde. Swift: fast, reliable, loosely coupled parallel computation. In *Proc. of 2007 IEEE Congress on Services*, pages 199–206, 2007.
107. O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu. *The Finite Element Method: Its Basis and Fundamentals*. Butterworth-Heinemann, 2005.