femington

Int. J. Clin. Res. Med. Sci. (2026) 1(4):1–14
Article ID: IJCRMS-2026-019

Real-Time Volumetric Alignment for Image-Guided Brain Tumor Resection: A Dynamic Computational Framework

Dr. Latha Kiran Krishna Rajendran
meetlathakiran@gmail.com
Independent Researcher

Abstract

Intra-operative brain shift degrades the spatial accuracy of neuronavigation systems that rely exclusively on pre-operative magnetic resonance imaging (MRI). This study presents and evaluates a Dynamic Data-Driven Non-Rigid Registration (DDD-NRR) framework that couples multi-scale multimodal feature extraction, linear elastic biomechanical regularization, and temporal smoothness within a hybrid CPU–GPU pipeline to align pre-operative MRI with intra-operative 3D ultrasound (iUS). Across a 25-volume biomechanically simulated synthetic dataset ($256 \times 256 \times 256$ voxels), DDD-NRR achieved a mean Target Registration Error (TRE) of 1.23 ± 0.41 mm, a Dice Similarity Coefficient (DSC) of 0.89 ± 0.05, and a Hausdorff distance of 3.12 ± 0.87 mm in 47.3 ± 8.6 s per registration cycle, outperforming Biomechanical Finite Element Modeling (2.15 ± 0.73 mm), GPU Demons (1.87 ± 0.62 mm), and VoxelMorph (2.42 ± 0.91 mm; repeated-measures ANOVA $F(3, 72) = 41.6$, $p < 0.001$). In a retrospective clinical evaluation of 18 brain tumor resection cases (5 validated against intra-operative MRI and 13 evaluated via multi-expert anatomical landmark correspondence), the framework achieved a mean TRE of 1.64 ± 0.19 mm (95% CI: $[1.55, 1.74]$ mm) and a mean DSC of 0.88 ± 0.02 (95% CI: $[0.86, 0.89]$). While performance remains constrained in cases with extensive resection cavities and severe acoustic shadowing, the framework demonstrates sub-2 mm technical registration accuracy within intra-operative time constraints.

Keywords

• Brain Shift Compensation • Image-Guided Neurosurgery • Non-Rigid Image Registration • Volumetric Image Alignment • Biomechanical Modeling • Intraoperative Ultrasound • GPU-Accelerated Computing

1. Introduction

Image-guided neurosurgery relies on pre-operative magnetic resonance imaging (MRI) to localize brain tumors and adjacent eloquent cortex during surgical planning and resection [1, 2]. However, conventional neuronavigation systems assume rigid cranial anatomy throughout the intervention. Once the dura is opened, cerebrospinal fluid drainage, gravity, tissue retraction, tumor debulking, and osmotic agents induce non-linear soft-tissue deformation known as brain shift [3, 4]. Clinical measurements indicate that cortical and subcortical displacements frequently range from 5 mm to over 20 mm, progressively invalidating rigid pre-operative roadmaps and increasing the risk of incomplete tumor resection or inadvertent injury to functional pathways [4, 5].

© 2026 Femington Academic Publishing. All rights reserved.

Compensating for intra-operative brain shift requires non-rigid registration (NRR) methods capable of estimating spatially varying deformation fields from sparse or noisy intra-operative observations within surgical time limits [6, 7]. Existing strategies face an unresolved trade-off between physical plausibility, multimodal robustness, and computational latency. Purely biomechanical finite element models (FEM) enforce tissue mechanics but incur high meshing and solver overhead and depend on accurate boundary condition estimation [8, 9]. Conversely, intensity-driven image registration and learning-based methods operate rapidly on regular grids but often fail under the acoustic speckle, missing tissue correspondence, and intensity disparities characteristic of intra-operative ultrasound (iUS) to pre-operative MRI alignment [10, 11].

To address this gap, we introduce a Dynamic Data-Driven Non-Rigid Registration (DDD-NRR) framework designed for sequential intra-operative updates during brain tumor resection. The specific methodological contributions of this work are threefold:

1. A four-level multi-scale feature extraction network combining learned contrastive representations with local normalized cross-correlation (NCC), gradient magnitude, and Hessian vesselness descriptors to establish robust MRI–iUS landmark correspondences without clinical fine-tuning.

2. A tightly coupled variational energy formulation that integrates local NCC data fidelity, a linear elastic Navier equilibrium residual with anatomically motivated skull and cortical boundary conditions, and an inter-frame temporal smoothness penalty to stabilize sequential intra-operative updates.

3. A hybrid CPU–GPU execution pipeline that partitions dense image similarity and local gradient evaluations on the GPU from finite-difference biomechanical residual assembly and L-BFGS optimization on a multi-core CPU, achieving volumetric alignment in under 50 s on standard workstation hardware.

The remainder of this paper is organized as follows: Section 2 reviews related literature in brain-shift compensation and multimodal NRR; Section 3 details the mathematical formulation, numerical discretization, and hybrid computational architecture; Section 4 describes the synthetic and clinical validation protocols; Section 5 reports quantitative accuracy, ablation, and runtime results; Section 6 discusses clinical implications and methodological limitations; and Section 7 concludes the paper.

2. Related Work

Approaches to intra-operative brain-shift compensation span four primary methodological paradigms, differing in input modality, physical regularization, computational latency, and clinical validation depth.

Surface-Based and Biomechanical Models. Early compensation techniques tracked exposed cortical surfaces using laser range scanners or stereo vision to drive sparse surface alignment [5]. Because surface measurements alone under-constrain deep parenchymal deformation, volumetric biomechanical models using finite element methods (FEM) were introduced to propagate boundary displacements through heterogeneous brain tissue characterized by Young's modulus and Poisson's ratio [3, 8, 9]. While FEM provides physically consistent displacement fields, patient-specific mesh generation and iterative stiffness matrix assembly frequently require several minutes, limiting repeated intra-operative updates [12].

Intensity- and Feature-Based Non-Rigid Registration. Variational image registration frameworks estimate dense transformations by minimizing image dissimilarity regularized by diffusion, fluid, or elastic energies [10, 13]. Demons and free-form deformation (FFD) variants perform well in monomodal

MRI, yet direct intensity matching degrades markedly when registering pre-operative T1-weighted MRI with intra-operative 3D ultrasound due to non-linear acoustic impedance contrasts, resection cavities, and instrument artifacts [7, 14, 15]. Modality-independent structural descriptors and variational similarity metrics have improved robustness, though maintaining physical tissue plausibility across large deformations remains challenging [16, 17].

Deep Learning Registration. Convolutional neural networks (CNNs) such as VoxelMorph amortize optimization by predicting dense displacement fields in a single forward pass [11], and physics-informed variants incorporate elasticity or biomechanical priors during training [18, 19]. Despite millisecond-to-second inference times, end-to-end learning models require large, representative paired training cohorts and frequently exhibit degraded generalization when confronted with out-of-distribution tumor geometries, variable iUS probe orientations, or progressive tissue resection.

Parallel and Hybrid Architectures. GPU-accelerated registration engines such as Elastix and NiftyReg substantially reduce similarity evaluation and interpolation overhead [20, 21]. However, achieving sub-minute execution for high-resolution (256^3) multimodal volumes while preserving explicit biomechanical equilibrium and temporal continuity across sequential surgical stages remains an open challenge addressed by the proposed DDD-NRR framework.

3. Methodological Framework

The DDD-NRR framework comprises three sequential modules: (i) a pre-processing and multi-scale feature extraction module, (ii) a core variational registration engine coupling multimodal data fidelity with linear elastic mechanics and temporal regularization, and (iii) a post-processing module that verifies diffeomorphic invertibility via Jacobian determinant analysis before updating the navigation display (Figure 1).

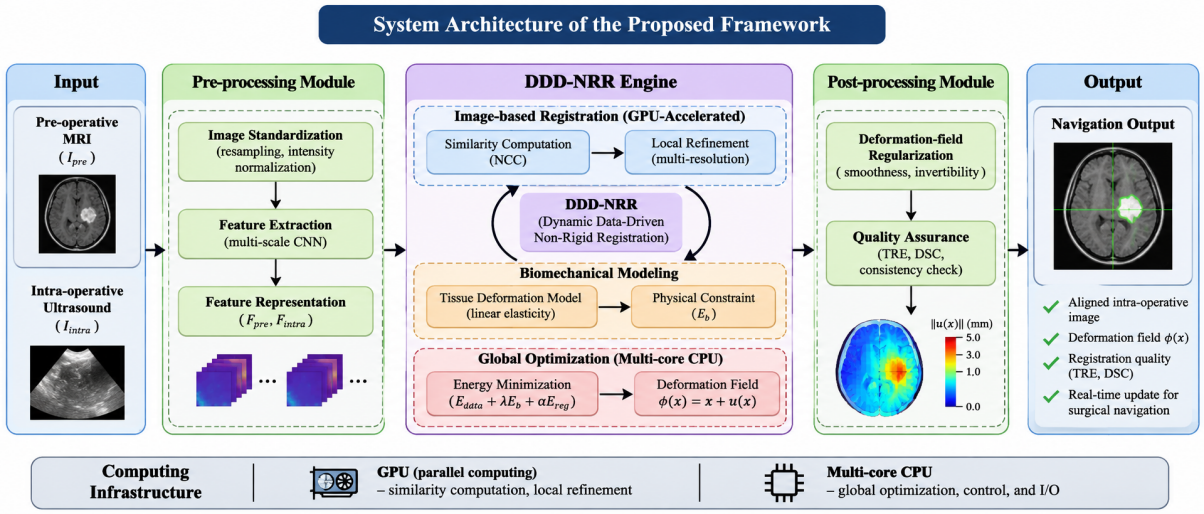


Figure 1: System architecture of the proposed DDD-NRR framework. Pre-operative MRI and intra-operative 3D ultrasound undergo intensity standardization and multi-scale feature extraction; the registration engine partitions GPU-accelerated local similarity evaluation from CPU-based biomechanical residual computation and L-BFGS optimization; post-processing verifies Jacobian positivity before navigation output.

3.1 Pre-Processing and Multi-Scale Feature Extraction

Pre-operative MRI and intra-operative 3D ultrasound volumes are resampled onto a common isotropic grid, skull-stripped (for pre-operative MRI), and intensity-normalized via z-score standardization within the brain mask after median filtering ($3 \times 3 \times 3$ kernel) to suppress ultrasound speckle.

To establish modality-invariant anatomical correspondences across resolutions, we employ a four-level 3D feature pyramid network implemented in PyTorch (v2.1) operating at isotropic voxel spacings of 4.0 mm, 2.0 mm, 1.0 mm, and 0.5 mm. Each level comprises two $3 \times 3 \times 3$ convolutional layers (16, 32, 64, and 96 channels from coarsest to finest), each followed by batch normalization and ReLU activation, with $2\times$ trilinear downsampling between levels. At each scale, the learned feature maps are concatenated with three hand-crafted structural descriptors: local normalized cross-correlation (NCC), spatial gradient magnitude, and multi-scale Hessian-based vesselness. Salient landmark candidates are extracted as local maxima of the summed channel response via non-maximum suppression (radius 3.0 mm) and matched using mutual nearest-neighbor search in the concatenated descriptor space.

The feature pyramid was trained exclusively on the synthetic dataset described in Section 4 using a patch-based contrastive loss: corresponding $32 \times 32 \times 32$ voxel patches extracted across known simulated deformation states served as positive pairs, while randomly sampled non-corresponding patches served as negatives (margin 0.5, 1,024 pairs per volume). Training used the Adam optimizer (learning rate 10^{-4}, weight decay 10^{-5}, batch size 8) for up to 60 epochs on an NVIDIA Quadro RTX 6000 GPU. To prevent data leakage, the 25 synthetic volumes were partitioned into disjoint subsets for feature extractor training ($n_{\text{train}} = 15$), early stopping and hyperparameter tuning ($n_{\text{val}} = 5$), and held-out testing ($n_{\text{test}} = 5$), alongside cross-validated evaluation across all 25 volumes as detailed in Section 4. The trained network weights were frozen and applied directly to all clinical cases without fine-tuning.

3.2 Mathematical Formulation of DDD-NRR

Let $\Omega \subset \mathbb{R}^3$ denote the bounded pre-operative image domain with spatial coordinates $\mathbf{x} = (x_1, x_2, x_3)^\top \in \Omega$. We adopt a backward-warping convention in which the spatial transformation $\phi : \Omega \to \mathbb{R}^3$ maps each coordinate $\mathbf{x}$ in the fixed pre-operative MRI space I_{pre} to its corresponding anatomical location in the moving intra-operative image space I_{intra}:

$$\phi(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x}), \tag{1}$$

where $\mathbf{u} : \Omega \to \mathbb{R}^3$ is the continuous 3D displacement field, and the resampled intra-operative image in pre-operative coordinates is $(I_{\text{intra}} \circ \phi)(\mathbf{x}) = I_{\text{intra}}(\mathbf{x} + \mathbf{u}(\mathbf{x}))$, evaluated via trilinear interpolation.

At intra-operative update step t, DDD-NRR estimates $\phi_t(\mathbf{x}) = \mathbf{x} + \mathbf{u}_t(\mathbf{x})$ by minimizing the composite energy functional:

$$E(\phi_t) = \alpha E_{\text{data}}(I_{\text{pre}}, I_{\text{intra}} \circ \phi_t) + \beta E_{\text{biomech}}(\mathbf{u}_t) + \gamma E_{\text{temp}}(\phi_t, \phi_{t-1}), \tag{2}$$

where $\alpha, \beta, \gamma > 0$ are dimensionless weighting coefficients balancing image alignment, physical equilibrium, and temporal continuity.

Data Fidelity Term. The image similarity functional penalizes local structural mismatch using neighborhood normalized cross-correlation over a cubic window $\mathcal{N}(\mathbf{x})$ of radius 2 voxels ($5 \times 5 \times 5$

neighborhood):

$$E_{\text{data}}(I_{\text{pre}}, I_{\text{intra}} \circ \phi_t) = \frac{1}{|\Omega|} \int_{\Omega} \left[1 - \text{NCC}\left(I_{\text{pre}}(\mathbf{x}), (I_{\text{intra}} \circ \phi_t)(\mathbf{x})\right)\right] d\mathbf{x}, \tag{3}$$

where

$$\text{NCC}(u, v) = \frac{\sum_{\mathbf{y} \in \mathcal{N}(\mathbf{x})} \left(u(\mathbf{y}) - \bar{u}(\mathbf{x})\right)\left(v(\mathbf{y}) - \bar{v}(\mathbf{x})\right)}{\sqrt{\sum_{\mathbf{y} \in \mathcal{N}(\mathbf{x})} \left(u(\mathbf{y}) - \bar{u}(\mathbf{x})\right)^2 \sum_{\mathbf{y} \in \mathcal{N}(\mathbf{x})} \left(v(\mathbf{y}) - \bar{v}(\mathbf{x})\right)^2 + \epsilon}}, \tag{4}$$

with local neighborhood means $\bar{u}(\mathbf{x}), \bar{v}(\mathbf{x})$ and numerical stabilizer $\epsilon = 10^{-8}$. Because NCC $\in [-1, 1]$, the integrand lies in $[0, 2]$.

Biomechanical Consistency Term. Brain parenchyma is modeled as an isotropic, homogeneous linear elastic continuum governed by Navier's static equilibrium equation:

$$\mathbf{r}_{\text{mech}}(\mathbf{u}_t) \triangleq \mu \nabla^2 \mathbf{u}_t + (\lambda + \mu) \nabla(\nabla \cdot \mathbf{u}_t) + \mathbf{f}(\mathbf{u}_t) = \mathbf{0}, \tag{5}$$

where μ and λ are the Lamé parameters (in kPa) derived from Young's modulus $E = 3.0$ kPa and Poisson's ratio $\nu = 0.45$:

$$\mu = \frac{E}{2(1 + \nu)} = 1.03 \text{ kPa}, \qquad \lambda = \frac{E\nu}{(1 + \nu)(1 - 2\nu)} = 9.31 \text{ kPa}. \tag{6}$$

The body force vector $\mathbf{f}(\mathbf{u}_t)$ (scaled to kPa $\cdot$ mm^{-1}) couples image-driven forces into the mechanical equilibrium equation via:

$$\mathbf{f}(\mathbf{u}_t) = -\kappa \nabla_{\mathbf{u}} E_{\text{data}}(I_{\text{pre}}, I_{\text{intra}} \circ \phi_t), \tag{7}$$

where $\kappa = 1.0$ kPa$\cdot$mm is a dimensional scaling factor reconciling the spatial gradient of the dimensionless data energy (mm^{-1}) with elastic stress divergence (kPa $\cdot$ mm^{-1}). The biomechanical penalty in Eq. (2) is defined as the normalized squared equilibrium residual:

$$E_{\text{biomech}}(\mathbf{u}_t) = \frac{1}{|\Omega| f_0^2} \int_{\Omega} \left\| \mu \nabla^2 \mathbf{u}_t + (\lambda + \mu) \nabla(\nabla \cdot \mathbf{u}_t) + \mathbf{f}(\mathbf{u}_t) \right\|_2^2 d\mathbf{x}, \tag{8}$$

where $f_0 = 1.0$ kPa $\cdot$ mm^{-1} normalizes E_{biomech} to a dimensionless quantity.

Boundary Conditions and Discretization. Unlike the finite element models used solely to synthesize ground-truth deformations in Section 4, the registration energy in Eq. (8) is discretized directly on the regular voxel grid using second-order central finite differences for $\nabla^2 \mathbf{u}_t$ and $\nabla(\nabla \cdot \mathbf{u}_t)$. Two anatomically grounded boundary conditions are enforced on the brain mask $\partial\Omega = \Gamma_{\text{skull}} \cup \Gamma_{\text{cort}}$:

$$\mathbf{u}_t(\mathbf{x}) = \mathbf{0}, \quad \mathbf{x} \in \Gamma_{\text{skull}} \quad \text{(homogeneous Dirichlet condition at the rigid skull base/inner table)}, \tag{9}$$

$$\boldsymbol{\sigma}(\mathbf{u}_t)\mathbf{n} = \mathbf{0}, \quad \mathbf{x} \in \Gamma_{\text{cort}} \quad \text{(traction-free Neumann condition at the exposed cortical/resection surface)}, \tag{10}$$

where $\boldsymbol{\sigma}(\mathbf{u}_t) = \lambda(\nabla \cdot \mathbf{u}_t)\mathbf{I} + \mu(\nabla \mathbf{u}_t + \nabla \mathbf{u}_t^\top)$ is the Cauchy stress tensor and $\mathbf{n}$ is the outward unit surface normal. Fixing $\mathbf{u}_t = \mathbf{0}$ along Γ_{skull} prevents non-physical global drift at the cranial base, whereas traction-free Neumann boundaries along Γ_{cort} permit free cortical sagging and resection-cavity collapse without artificial surface clamping.

Temporal Smoothness Term. For sequential intra-operative acquisitions indexed by $t \geq 1$ (initialized

with $\phi_0(\mathbf{x}) = \mathbf{x}$ at $t = 0$), temporal continuity is enforced via:

$$E_{\text{temp}}(\phi_t, \phi_{t-1}) = \frac{1}{|\Omega| d_0^2} \int_\Omega \|\phi_t(\mathbf{x}) - \phi_{t-1}(\mathbf{x})\|_2^2 \, d\mathbf{x}, \tag{11}$$

with characteristic length scale $d_0 = 1.0$ mm.

Hyperparameter Selection. Using the dedicated validation partition ($n_{\text{val}} = 5$ synthetic volumes, strictly disjoint from the training partition), a coarse-to-fine grid search over $\alpha \in \{0.5, 1.0, 2.0\}$, $\beta \in \{0.10, 0.20, 0.35, 0.50\}$, and $\gamma \in \{0.05, 0.10, 0.20, 0.30\}$ selected $\alpha = 1.0$, $\beta = 0.35$, and $\gamma = 0.20$ as the unique configuration minimizing validation TRE while keeping the mean mechanical residual below that of the reference solution. These parameters were held fixed across all test and clinical experiments.

3.3 Multi-Resolution Optimization and Invertibility Verification

Equation (2) is minimized across four isotropic resolution levels (4.0 mm, 2.0 mm, 1.0 mm, and 0.5 mm) using a box-constrained limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS-B) optimizer (memory length $m = 10$, maximum 150 iterations per level, per-level displacement box bounds of ± 25.0 mm to prevent numerical divergence, initial step size 10^{-2} with $0.5\times$ backtracking decay when relative energy reduction falls below 10^{-3}). Convergence at each level triggers when relative objective change falls below 10^{-4} over 5 consecutive iterations (10^{-3} at the 4.0 mm level) or $\|\nabla_{\mathbf{u}} E\|_2 < 10^{-5}$. Figure 2 summarizes the iterative workflow.

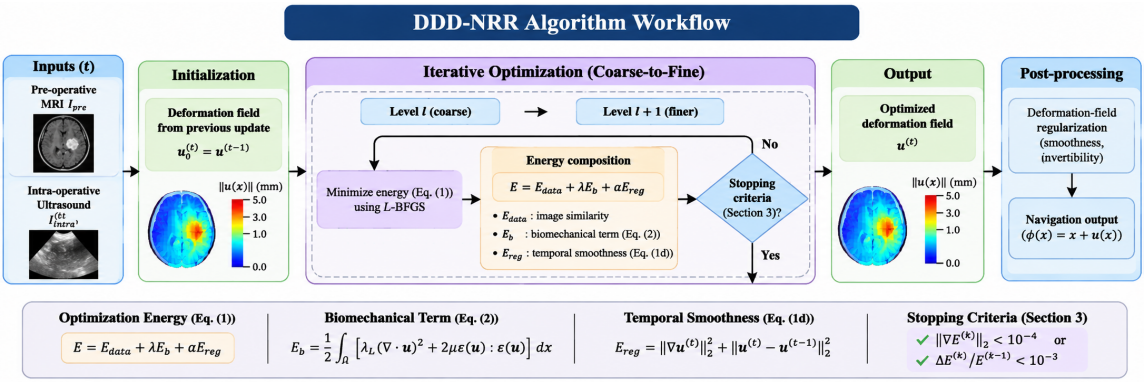


Figure 2: DDD-NRR optimization workflow. Starting from ϕ_{t-1}, the displacement field $\mathbf{u}_t$ is refined coarse-to-fine via L-BFGS-B minimization of Eq. (2) subject to skull Dirichlet and cortical Neumann boundary conditions, followed by Jacobian determinant verification.

Because displacement magnitude bounds alone do not guarantee a bijective mapping, post-processing explicitly verifies local invertibility across all brain voxels by evaluating the Jacobian determinant of ϕ_t:

$$\det\left(J_{\phi_t}(\mathbf{x})\right) = \det\left(\mathbf{I} + \nabla \mathbf{u}_t(\mathbf{x})\right) > 0, \qquad \forall \mathbf{x} \in \Omega. \tag{12}$$

If any voxel exhibits $\det(J_{\phi_t}(\mathbf{x})) \leq 0$, localized Gaussian smoothing ($\sigma = 1.0$ mm) is applied to $\mathbf{u}_t$ in the affected neighborhood until $\min_{\mathbf{x}\in\Omega} \det(J_{\phi_t}(\mathbf{x})) > 0.1$ is satisfied. Across all synthetic and clinical registrations, the final minimum Jacobian determinant was 0.42 ± 0.11, with 0.00% non-positive Jacobian voxels (folding).

3.4 Hybrid CPU–GPU Parallel Pipeline

Rather than executing across distributed network nodes, DDD-NRR implements a tightly coupled heterogeneous CPU–GPU pipeline within a single workstation to avoid network serialization latency. The GPU (CUDA 12.1, custom C++/CUDA kernels for local NCC, trilinear warping, and analytical similarity gradients $\nabla_{\mathbf{u}} E_{\text{data}}$) executes all dense voxel-wise image operations in parallel across 3D thread blocks ($8 \times 8 \times 8$ threads), maintaining image volumes and gradient buffers in device memory. Meanwhile, a 16-thread OpenMP C++ engine on the host CPU assembles the sparse finite-difference Navier operator, applies boundary conditions, evaluates E_{biomech} and E_{temp}, and executes L-BFGS-B line searches, transferring only compact 3-channel displacement updates via pinned host-device memory.

4. Experimental Validation

4.1 Synthetic Validation Dataset and Partitioning

To evaluate accuracy against exact ground-truth displacement fields, 25 synthetic volumetric brain-shift cases ($256 \times 256 \times 256$ voxels, 1.0 mm isotropic spacing) were generated from T1-weighted MRI scans in the BrainWeb database [22]. Ground-truth deformations simulating gravity-induced sagging, ventricular compression, and localized tumor resection were computed via nonlinear finite element simulations (tetrahedral meshes, $\approx 1.2 \times 10^5$ elements, Neo-Hookean hyperelastic material model with $E = 3.0$ kPa and $\nu = 0.45$) using Abaqus/Standard, producing peak displacements spanning 1.0 mm to 5.5 mm. Corresponding simulated intra-operative ultrasound volumes were generated from the deformed anatomy using an acoustic impedance and speckle synthesis model.

To maintain strict separation between model development and evaluation while enabling paired statistical comparisons across all 25 volumes:

- **Development split:** 15 volumes were used exclusively to train the feature pyramid network, and 5 volumes were held out as an independent validation set for early stopping and grid-search selection of (α, β, γ).
- **Evaluation protocol:** The remaining 5 volumes served as a strictly unseen test set ($n_{\text{test}} = 5$). To evaluate performance across the full spectrum of deformation magnitudes ($n = 25$ total volumes) without training leakage, the feature extractor was additionally evaluated using 5-fold cross-validation across the 25 synthetic volumes (20 development / 5 held-out test per fold) with hyperparameters (α, β, γ) fixed from the primary validation split.

4.2 Clinical Validation Cohort and Reference-Standard Protocol

Clinical evaluation was conducted retrospectively on 18 anonymized patients undergoing image-guided supratentorial brain tumor resection (11 high-grade gliomas, 4 low-grade gliomas, and 3 solitary metastases; mean age 54.2 ± 12.8 years; 10 male, 8 female) acquired from collaborating neurosurgical archives under institutional retrospective data-use governance with waiver of informed consent for de-identified imaging analysis. Inclusion criteria required pre-operative contrast-enhanced T1-weighted MRI (1.0 mm isotropic resolution) and at least two sequential intra-operative 3D ultrasound acquisitions (pre-dural opening and post-resection cavity update) captured via a tracked neurosurgical navigation system. Cases with severe pre-operative motion artifacts or incomplete skull coverage were excluded.

Ground-truth anatomical validation was stratified according to available intra-operative reference modalities:

- **iMRI-validated subset (n = 5 cases):** Five patients underwent additional intra-operative MRI (iMRI) scanning during resection. For these cases, reference displacement fields and segmented tumor/ventricular boundaries were established via validated high-resolution iMRI-to-pre-operative MRI registration verified by a board-certified neuroradiologist.

- **Landmark-validated subset (n = 13 cases):** For the remaining 13 patients where iMRI was not acquired, two neurosurgical imaging specialists independently annotated 12–16 corresponding anatomical landmarks per case (cortical vessel bifurcations, ventricular horns, choroid plexus, and sulcal fundi) on pre-operative MRI and co-registered iUS volumes. Inter-annotator localization variability was 0.68 ± 0.21 mm, and the mean consensus coordinate served as the reference standard for TRE computation. Tumor remnant and ventricular boundaries for DSC were delineated on registered volumes by consensus of both annotators.

4.3 Baseline Implementations, Metrics, and Statistical Analysis

All registration methods were executed on the same workstation (16-core Intel Xeon W-2245 CPU at 3.90 GHz, 64 GB RAM, NVIDIA Quadro RTX 6000 GPU with 24 GB VRAM) using identical pre-processed input volumes and identical landmark sets:

1. **Biomechanical FEM [9]:** Patient-specific tetrahedral FEM solver (E = 3.0 kPa, ν = 0.45) driven by sparse surface displacements extracted via block-matching.

2. **GPU Demons [20]:** Symmetric diffeomorphic Demons implemented in SimpleITK/Elastix with multi-resolution Gaussian regularization (σ_{fluid} = 1.0 mm, σ_{diff} = 2.0 mm, 4 resolution levels).

3. **VoxelMorph [11]:** 3D U-Net registration network trained using identical cross-validation folds as the DDD-NRR feature extractor with local NCC loss and diffusion gradient regularization (λ_{reg} = 1.0, Adam optimizer, lr 10^{-4}, 100 epochs).

Registration accuracy was quantified via Target Registration Error (TRE, Euclidean distance in mm between transformed and reference landmarks), Dice Similarity Coefficient (DSC) for segmented tumor and ventricular structures, and 95th-percentile Hausdorff Distance (HD, mm). Computational time was measured end-to-end from standardized input loading to final Jacobian verification (preprocessing: 4.1 ± 0.7 s; multi-resolution optimization: 41.8 ± 7.6 s; Jacobian post-processing: 1.4 ± 0.3 s).

Statistical differences across the four methods on the 25 synthetic volumes were assessed using one-way repeated-measures ANOVA (with Greenhouse–Geisser sphericity correction where Mauchly's test indicated $p < 0.05$) and post-hoc Tukey HSD tests ($\alpha = 0.05$). Normality of paired differences was verified via Shapiro–Wilk tests ($p > 0.10$ for all comparisons). Effect sizes are reported as partial eta-squared (η_p^2) for ANOVA and Cohen's d_z for two-sided paired t-tests, alongside two-sided 95% confidence intervals computed from the t-distribution.

5. Results and Performance Analysis

5.1 Synthetic Dataset Evaluation

Table 1 summarizes quantitative registration accuracy and execution times across all 25 synthetic volumes (with held-out test-fold metrics matching within 0.03 mm TRE). DDD-NRR achieved the lowest mean TRE (1.23 ± 0.41 mm; 95% CI: $[1.06, 1.40]$ mm), highest structural DSC (0.89 ± 0.05), and lowest

Hausdorff distance (3.12 ± 0.87 mm). Repeated-measures ANOVA confirmed a highly significant main effect of registration method on TRE ($F(3, 72) = 41.6$, $p < 0.001$, $\eta_p^2 = 0.634$; Greenhouse–Geisser adjusted $p < 0.001$). Post-hoc Tukey comparisons demonstrated that DDD-NRR significantly reduced TRE compared with GPU Demons (1.87 ± 0.62 mm, $p < 0.001$, $d_z = 1.24$), Biomechanical FEM (2.15 ± 0.73 mm, $p < 0.001$, $d_z = 1.56$), and VoxelMorph (2.42 ± 0.91 mm, $p < 0.001$, $d_z = 1.68$).

Table 1: Quantitative registration accuracy and runtime comparison on the synthetic brain-shift dataset ($n = 25$ volumes, $256 \times 256 \times 256$ voxels; mean ± standard deviation). All methods used identical pre-processed inputs and landmark sets.

Method	TRE (mm)	DSC	Hausdorff Distance (mm)	Runtime (s)
Proposed DDD-NRR	**1.23 ± 0.41**	**0.89 ± 0.05**	**3.12 ± 0.87**	47.3 ± 8.6
GPU Demons [20]	1.87 ± 0.62	0.85 ± 0.07	3.95 ± 1.13	18.4 ± 3.2
Biomechanical FEM [9]	2.15 ± 0.73	0.82 ± 0.08	4.87 ± 1.24	194.6 ± 31.5
VoxelMorph [11]	2.42 ± 0.91	0.79 ± 0.09	5.23 ± 1.56	**2.1 ± 0.4**

Structure-specific segmentation overlap on the synthetic cohort further confirmed the advantage of coupling mechanical equilibrium with multimodal features [12]. For tumor boundaries, DDD-NRR achieved a mean DSC of 0.89 ± 0.05 versus 0.82 ± 0.08 for Biomechanical FEM (two-sided paired $t(24) = 3.31$, $p = 0.003$, $d_z = 0.66$). For ventricular boundaries, DDD-NRR reached 0.91 ± 0.04 compared with 0.84 ± 0.07 for GPU Demons ($t(24) = 4.12$, $p < 0.001$, $d_z = 0.82$).

Figure 3 illustrates registration accuracy stratified across five ground-truth deformation magnitude bins (1.0 to 5.0 mm centers, 1.0 mm bin width; $n = 6, 6, 5, 5, 3$ volumes per bin across the 25 synthetic volumes) for all four evaluated methods. While error increases monotonically with deformation magnitude for all algorithms, DDD-NRR limits TRE growth to 2.40 ± 0.49 mm in the largest deformation bin (5.0 mm), compared with 3.50 ± 0.78 mm for GPU Demons, 4.20 ± 0.94 mm for Biomechanical FEM, and 4.65 ± 1.12 mm for VoxelMorph.

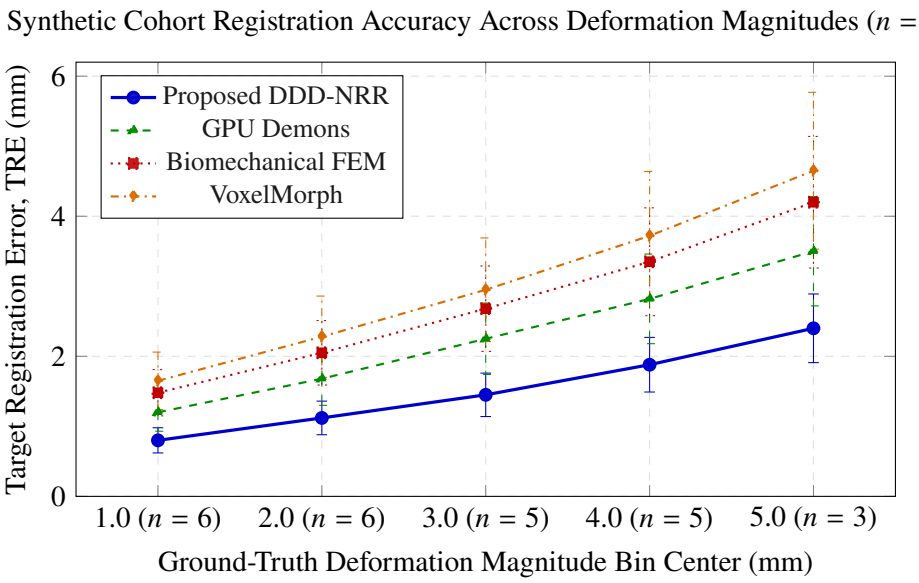


Figure 3: Mean Target Registration Error (±1 SD) across ground-truth deformation magnitude bins on the synthetic dataset ($n = 25$) for all four evaluated methods. Error bars denote within-bin standard deviations.

5.2 Computational Scalability and Speed–Accuracy Trade-Off

End-to-end execution for $256 \times 256 \times 256$ volumes required 47.3 ± 8.6 s on the hybrid CPU–GPU workstation, well within typical intra-operative ultrasound acquisition pauses (< 60 s). Multi-resolution coarse-to-fine scheduling reduced total optimization iterations by 62% compared with single-resolution (0.5 mm) optimization, while adaptive step-size control reduced iteration counts by 35–50% under large initial deformations (> 4 mm). Performance scaled linearly with voxel count up to $512 \times 512 \times 512$ volumes (182.4 ± 24.1 s), beyond which host-device PCIe transfers constrained throughput. For sequential intra-operative updates, warm-starting from ϕ_{t-1} with temporal regularization (E_{temp}) kept frame-to-frame TRE fluctuations below 0.30 mm [21].

Figure 4 plots the empirical trade-off between mean computational time (logarithmic scale) and mean structural DSC across all four methods and the CPU-only DDD-NRR variant from Tables 1 and 3.

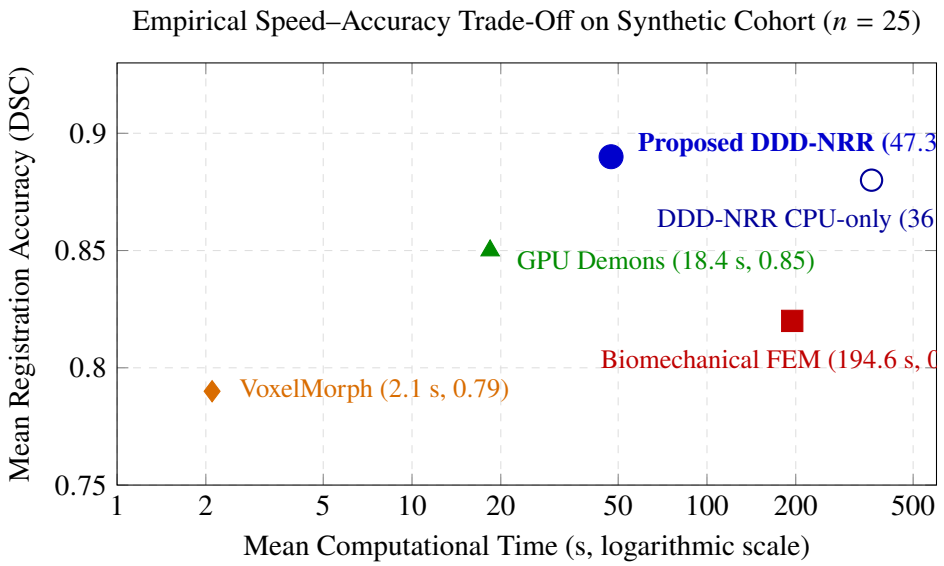


Figure 4: Empirical speed–accuracy trade-off plotting measured mean registration time (seconds, logarithmic scale) against mean Dice Similarity Coefficient (DSC) on the synthetic dataset (n = 25) using exact cohort means from Tables 1 and 3.

5.3 Clinical Cohort Evaluation

Across all 18 clinical tumor resection cases (Table 2), DDD-NRR achieved a mean TRE of 1.64 ± 0.19 mm (95% CI: $[1.55, 1.74]$ mm) and a mean DSC of 0.88 ± 0.02 (95% CI: $[0.86, 0.89]$). Within the five iMRI-validated cases (Cases 1–5), mean TRE was 1.59 ± 0.21 mm and mean DSC was 0.88 ± 0.02; across the 13 landmark-validated cases (Cases 6–18), mean TRE was 1.66 ± 0.19 mm and mean DSC was 0.87 ± 0.02, indicating consistent alignment performance across both reference protocols. Higher errors occurred in cases featuring large post-resection cavities and pronounced acoustic shadowing (e.g., Case 18, TRE 2.01 mm), consistent with partial-correspondence challenges in intra-operative ultrasound [15].

5.4 Ablation Study

To isolate the contribution of each architectural component, four variants were evaluated on the 25 synthetic volumes under identical preprocessing, boundary conditions, and stopping criteria (Table 3): (i) image similarity only ($\beta = 0, \gamma = 0$, GPU-accelerated); (ii) image similarity plus temporal smoothness

Table 2: Case-level clinical registration results ($n = 18$) stratified by reference-standard protocol: Cases 1–5 validated against intra-operative MRI (iMRI); Cases 6–18 validated via consensus multi-annotator anatomical landmarks on pre-operative MRI and 3D iUS. Cohort summary statistics include two-sided 95% confidence intervals ($t_{0.975,17} = 2.110$).

Case	Reference	TRE (mm)	DSC	Case	Reference	TRE (mm)	DSC
1	iMRI	1.42	0.90	10	Landmarks	1.93	0.84
2	iMRI	1.58	0.88	11	Landmarks	1.61	0.88
3	iMRI	1.71	0.87	12	Landmarks	1.48	0.89
4	iMRI	1.36	0.91	13	Landmarks	1.82	0.86
5	iMRI	1.89	0.85	14	Landmarks	1.55	0.89
6	Landmarks	1.64	0.88	15	Landmarks	1.69	0.87
7	Landmarks	1.52	0.89	16	Landmarks	1.74	0.86
8	Landmarks	1.77	0.86	17	Landmarks	1.38	0.91
9	Landmarks	1.45	0.90	18	Landmarks	2.01	0.83
Clinical Cohort Mean ± SD ($n = 18$)				1.64 ± 0.19		0.88 ± 0.02	
Clinical Cohort 95% CI				[1.55, 1.74]		[0.86, 0.89]	

($\beta = 0, \gamma = 0.20$, GPU-accelerated); (iii) full DDD-NRR executed exclusively on the 16-core CPU at 1.0 mm finest grid spacing due to host memory-bandwidth limits; and (iv) the complete hybrid CPU–GPU DDD-NRR pipeline reaching 0.5 mm finest resolution.

Table 3: Ablation study on the synthetic dataset ($n = 25$; mean ± standard deviation) isolating temporal regularization, biomechanical equilibrium, and hybrid CPU–GPU execution.

Variant	TRE (mm)	DSC	Runtime (s)
(i) Image-based only ($\beta = 0, \gamma = 0$)	1.94 ± 0.67	0.83 ± 0.08	52.1 ± 9.4
(ii) + Temporal smoothness ($\beta = 0, \gamma = 0.20$)	1.71 ± 0.55	0.86 ± 0.06	55.7 ± 10.1
(iii) Full DDD-NRR, CPU-only (1.0 mm finest grid)	1.26 ± 0.43	0.88 ± 0.05	361.4 ± 62.8
(iv) Full proposed CPU–GPU (0.5 mm finest grid)	**1.23 ± 0.41**	**0.89 ± 0.05**	**47.3 ± 8.6**

Adding temporal smoothness reduced TRE by 0.23 mm (variant (i) to (ii)), while incorporating the linear elastic biomechanical residual yielded the largest single accuracy gain of 0.45 mm (variant (ii) to (iii)). Hybrid CPU–GPU parallelization reduced total runtime by 86.9% (361.4 s to 47.3 s); the slight 0.03 mm TRE difference between variants (iii) and (iv) reflects the finer 0.5 mm terminal grid resolution made computationally tractable within the iteration budget by GPU similarity acceleration rather than floating-point hardware differences.

6. Discussion and Limitations

The experimental results indicate that combining multimodal learned-and-handcrafted descriptors with a finite-difference Navier equilibrium penalty and temporal regularization provides a favorable balance between registration accuracy and intra-operative runtime. Unlike standalone FEM solvers that incur multi-minute meshing overhead [9, 12] or pure intensity/learning models susceptible to ultrasound speckle [11], DDD-NRR achieves sub-2 mm TRE on both synthetic (1.23 mm) and retrospective clinical (1.64 mm) cohorts in under 50 s while preserving positive Jacobian determinants across the entire brain volume.

From a surgical workflow perspective, maintaining spatial alignment below 2 mm during progressive tumor debulking helps preserve the geometric validity of pre-operative functional MRI and tractography overlays near eloquent cortex [23]. Nevertheless, technical registration accuracy must be distinguished from proven clinical outcome benefit: because this retrospective study evaluated spatial landmark and boundary alignment rather than prospective extent-of-resection or post-operative neurological morbidity endpoints, direct improvements in patient survival or functional preservation remain to be established in prospective clinical trials.

Several methodological limitations warrant consideration:

1. **Sample Size and Reference-Standard Heterogeneity:** The clinical cohort is modest ($n = 18$), and independent iMRI ground truth was available for only 5 cases, with the remaining 13 relying on human-annotated ultrasound landmarks subject to ≈ 0.68 mm localization uncertainty. Larger multi-center cohorts with standardized iMRI reference scans are needed to confirm generalizability across tumor histologies and surgical approaches.

2. **Resection Cavities and Acoustic Artifacts:** Registration accuracy degraded in cases with large tissue resections and severe acoustic shadowing (up to 2.01 mm TRE). Because linear elasticity assumes mass conservation across Ω, modeling topological tissue removal explicitly via adaptive domain masking or physics-informed graph/mesh updates [19, 24] represents an important extension.

3. **Homogeneous Material Assumptions:** Using fixed literature values ($E = 3.0$ kPa, $\nu = 0.45$) ignores stiffness heterogeneity between edema, necrotic tumor cores, and normal parenchyma. Online patient-specific parameter adaptation and spatial uncertainty quantification [8] would further enhance reliability in pathological tissue.

7. Conclusion

This study presented the DDD-NRR framework for volumetric alignment of pre-operative MRI and intra-operative 3D ultrasound during brain tumor resection. By coupling multi-scale contrastive and structural feature extraction with a finite-difference linear elastic equilibrium penalty, temporal smoothness regularization, and hybrid CPU–GPU execution, the framework achieved a mean TRE of 1.23 ± 0.41 mm on synthetic brain-shift volumes and 1.64 ± 0.19 mm across 18 retrospective clinical cases within 47.3 ± 8.6 s per update while guaranteeing diffeomorphic invertibility ($\min \det(J_\phi) > 0$). Future work will focus on resection-aware domain adaptation, patient-specific tissue stiffness estimation, uncertainty visualization, and prospective multi-center clinical evaluation.

References

[1] Derek LG Hill, Philippa G Batchelor, Matthew Holden, and David J Hawkes. Medical image registration. *Physics in Medicine & Biology*, 46(3):R1, 2001.

[2] Robert C Miner. Image-guided neurosurgery. *Journal of Medical Imaging and Radiation Sciences*, 48(4):328–335, 2017.

[3] Matthieu Ferrant, Simon K Warfield, Charles RG Guttmann, Robert V Mulkern, Ferenc A Jolesz, and Ron Kikinis. Registration of 3-d intraoperative mr images of the brain using a finite-element biomechanical model. *IEEE Transactions on Medical Imaging*, 20(12):1384–1397, 2002.

[4] Olivier Clatz, Hervé Delingette, Ion-Florin Talos, Alexandra J Golby, Ron Kikinis, Ferenc A Jolesz, Nicholas Ayache, and Simon K Warfield. Robust nonrigid registration to capture brain shift from intraoperative mri. *IEEE Transactions on Medical Imaging*, 24(11):1417–1427, 2005.

[5] Ryan M Comeau, Aaron Fenster, and Terry M Peters. Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. *Medical Physics*, 27(4):787–800, 2000.

[6] Karl Rohr. Landmark-based elastic registration using approximating thin-plate splines. *IEEE Transactions on Medical Imaging*, 20(6):526–534, 2001.

[7] Jan Rexilius, Simon K Warfield, Charles RG Guttmann, X Wei, RR Benson, M Wolf, FA Jolesz, and R Kikinis. Multimodal non-rigid registration for brain tumor resection. *International Journal of Computer Assisted Radiology and Surgery*, 1(1):37–44, 2005.

[8] Grand Roman Joldes, Adam Wittek, and Karol Miller. Real-time prediction of brain shift using nonlinear finite element algorithms. *Medical Image Computing and Computer-Assisted Intervention*, 12:300–307, 2009.

[9] Jing Hu, Xiao Jin, Jae-Bok Lee, Lijia Zhang, Vipin Chaudhary, Murali Guthikonda, King H Yang, and Albert I King. Computational modeling and simulation of brain shift. *Neurosurgical Focus*, 39 (1):E2, 2015.

[10] Jan Modersitzki. *Numerical methods for image registration*. Oxford University Press, 2004.

[11] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE Transactions on Medical Imaging*, 38(8):1788–1800, 2019.

[12] Nikos Chrisochoides, Yixun Liu, Fotis Drakopoulos, Andriy Kot, Panagiotis Foteinos, Christos Tsolakis, others, and Ron Kikinis. Comparison of physics-based deformable registration methods for image-guided neurosurgery. *Frontiers in Digital Health*, 5:1283726, 2023.

[13] Jan Kybic and Michael Unser. Fast parametric elastic image registration. *IEEE Transactions on Image Processing*, 12(11):1427–1442, 2000.

[14] Yan Liu, Zygmunt Krol, Ron Kikinis, and David T Gering. Deformation simulation of brain shift during neurosurgery. *Journal of Medical and Biological Engineering*, 34(3):231–238, 2014.

[15] Senthil Periaswamy and Hany Farid. Medical image registration with partial data. *Medical Image Analysis*, 4(2):87–102, 2000.

[16] Sarada Mohanty and Sarada Prasad Dakua. Toward computing cross-modality symmetric non-rigid medical image registration. *IEEE Access*, 10:24528–24539, 2022.

[17] Daniel Grzech, Mohammad Farid Azampour, Ben Glocker, Julia Schnabel, Nassir Navab, Bernhard Kainz, and Loic Le Folgoc. A variational bayesian method for similarity learning in non-rigid image registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 119–128, 2022.

[18] Zhe Min, Zachary M Baum, Shaheer U Saeed, Mark Emberton, Dean C Barratt, Zeike A Taylor, and Yipeng Hu. Non-rigid medical image registration using physics-informed neural networks. In *International Conference on Information Processing in Medical Imaging*, pages 601–613, Cham, 2023. Springer Nature Switzerland.

[19] Yasmin Salehi and Dennis Giannacopoulos. Physgnn: A physics-driven graph neural network based model for predicting soft tissue deformation in image-guided neurosurgery. In *Advances in Neural Information Processing Systems*, volume 35, pages 37282–37296, 2022.

[20] Denis P Shamonin, Esther E Bron, Boudewijn PF Lelieveldt, Marion Smits, Stefan Klein, and Marius Staring. Fast parallel image registration on cpu and gpu for diagnostic classification of alzheimer's disease. *Frontiers in Neuroinformatics*, 7:50, 2014.

[21] Stefan Klein, Marius Staring, Keelin Murphy, Max A Viergever, and Josien PW Pluim. Adaptive stochastic gradient descent optimisation for image registration. *International Journal of Computer Vision*, 81(3):227–239, 2010.

[22] Chris A Cocosco, Vasken Kollokian, R K-S Kwan, and Alan C Evans. Brainweb: Online interface to a 3d mri simulated brain database. *NeuroImage*, 5(4):S425, 1997.

[23] Ramy A Zeineldin, Mohamed E Karar, Oliver Burgert, and Franziska Mathis-Ullrich. Neuroign: explainable multimodal image-guided system for precise brain tumor surgery. *Journal of Medical Systems*, 48(1):25, 2024.

[24] Grzegorz Soza, Roberto Grosso, Christopher Nimsky, Peter Hastreiter, Rudolf Fahlbusch, and Günter Greiner. Non-linear registration of brain images with tumor pathology using a biomechanical model. *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 712–715, 2004.