PathoMamba: Piecewise-Diffeomorphic Registration via Stiffness-Modulated State Space Dynamics

Ernest Asimeng[1][0009−0006−1427−8195], Jun Chen[1], Godbless Mensah[1], Jiaqi Zhang[1], Kai Han[1], Yang Chen[2,3,4], and Zhe Liu[1][0000−0002−1197−0390]
1000004088@ujs.edu.cn

[1] Jiangsu University, School of Computer Science and Communication Engineering
[2] The School of Computer Science and Engineering, Southeast University
[3] Ministry of Education, and the Jiangsu Provincial Joint International Research Laboratory of Medical Information Processing, Southeast University
[4] The Key Laboratory of New Generation Artificial Intelligence Technology and Its Interdisciplinary Applications (Southeast University)

Abstract. Longitudinal glioma registration faces a fundamental *Correspondence - Regularity Conflict*: healthy parenchyma requires high stiffness to prevent folding, while tumor evolution demands local plasticity to accommodate mass effect. Current deep learning frameworks, treating the brain as a mechanically homogeneous continuum, fail to resolve this paradox. We propose PathoMamba, a stiffness-modulated State Space Model (SSM) that bridges continuum mechanics and deep learning. By reinterpreting the SSM discretization step (Δ) as learnable *Numerical Inertia*, we actively modulate feature dynamics—enforcing rigidity in healthy tissue while enabling rapid updates in the pathology. This is coupled with a *Therapy-Aware Biomechanical* objective to robustly model heterogeneous growth and shrinkage trends. Validated on the BraTS-Reg benchmark, PathoMamba matches state-of-the-art Transformer accuracy (1.88 mm) while inferring 2× faster and guaranteeing *0.00% topological folding*, offering a safe, physics-informed solution for clinical monitoring.

Keywords: Deformable Registration · State Space Models · Mamba · Longitudinal Glioma · Physics-Informed Deep Learning.

1 Introduction

The longitudinal management of diffuse gliomas relies on the precise evaluation of tumor evolution under therapy. Accurate voxel-wise alignment of baseline (T_0) and follow-up (T_1) MRI is prerequisite for distinguishing true tumor progression from pseudo-progression (radiation necrosis) [14,15]. However, as highlighted by the recent BraTS-Reg challenge [2], gliomas introduce a fundamental **biomechanical paradox**. The brain is a composite material: the skull is rigid, the healthy parenchyma is viscoelastic, and the tumor is a growing (or shrinking) pathology that exerts mass effect. Conventional registration algorithms and deep

learning methods, which treat the brain as a mechanically homogeneous continuum [3, 12, 19, 24], fail to capture these distinct material properties, leading to clinically dangerous misalignments at the tumor boundary.

Current state-of-the-art deep learning methods, such as VoxelMorph [12] and TransMorph [4], optimize a global energy function balanced by a smoothness regularizer (e.g., $||\nabla\phi||^2$). This creates a zero-sum game:

1. **Global Smoothness:** High regularization preserves topology in healthy tissue but suppresses the high-frequency deformations required to model tumor expansion [7], causing the **"Vanishing Mass Effect"** artifact.
2. **Local Fidelity:** Low regularization allows tumor growth but causes unrealistic folding (topology breakdown) in the surrounding healthy brain.

Existing solutions often rely on masking or cost-function weighting [16], but they fundamentally lack a mechanism to alter the *dynamics* of the transformation model itself.

We argue that the solution lies in bridging control theory and biomechanics. While Transformers offer global context, their quadratic complexity ($O(N^2)$) prohibits high-resolution 3D modeling of fine physical boundaries. State Space Models (SSMs), specifically Mamba [5, 9, 22], offer linear complexity ($O(N)$) but are typically treated as "physics-blind" sequence modelers [20].

We propose a novel interpretation of the discretized SSM. The core discretization step Δ in the Mamba ODE governs the system's "memory." We hypothesize that Δ is the learnable homologue to **physical stiffness** (or numerical inertia):

- **High Inertia (Healthy Tissue):** By forcing $\Delta \to 0$, the hidden state evolves slowly, retaining long-range memory. This effectively models **stiffness**, enforcing global coherence and preventing local folding.
- **Low Inertia (Pathology):** By boosting $\Delta \to \infty$, the state reacts instantly to local inputs (zero memory). This models **plasticity**, permitting the sharp gradients and rapid expansions characteristic of glioblastoma.

We present **PathoMamba**, a stiffness-informed registration framework. Our contributions are:

1. **Stiffness-Modulated Dynamics (SMD):** Unlike prior works that passively concatenate segmentation masks, we introduce an active gating mechanism that conditions the ODE solver on anatomical priors. This enables the network to learn a *heterogeneous diffeomorphic* flow—rigid in healthy tissue yet fluid in the tumor—guaranteed via Stationary Velocity Field (SVF) integration.
2. **PathoMamba-U Architecture:** We resolve the 1D causal bias of SSMs via a 6-way isotropic scanning mechanism. This allows us to model biomechanical forces omni-directionally while maintaining $O(N)$ efficiency.
3. **Therapy-Aware Biomechanical Loss (TABL):** We propose a biologically grounded objective that enforces the Monro-Kellie doctrine (intracranial volume conservation). Unlike standard expansion penalties, TABL adaptively constrains Jacobian determinants based on component-wise tumor trends, robustly handling simultaneous recurrence and cavity collapse.

2 Methodology

We formulate longitudinal glioma registration as a *Stiffness-Informed Diffeomorphic* problem. Let $I_F, I_M : \Omega \to \mathbb{R}$ denote the fixed (baseline, T_0) and moving (follow-up, T_1) volumetric MRI scans defined on the spatial domain $\Omega \subset \mathbb{R}^3$. We aim to estimate a smooth, topology-preserving mapping $\phi : \Omega \to \Omega$ such that $I_M \circ \phi \approx I_F$, while adhering to the heterogeneous biomechanical constraints of the pathology.

To resolve the Correspondence-Regularity conflict, we propose **PathoMamba**. As illustrated in Fig. 1, this framework couples a Stiffness-Modulated State Space Model with a Therapy-Aware Biomechanical objective to dynamically adapt the registration stiffness based on tissue pathology.

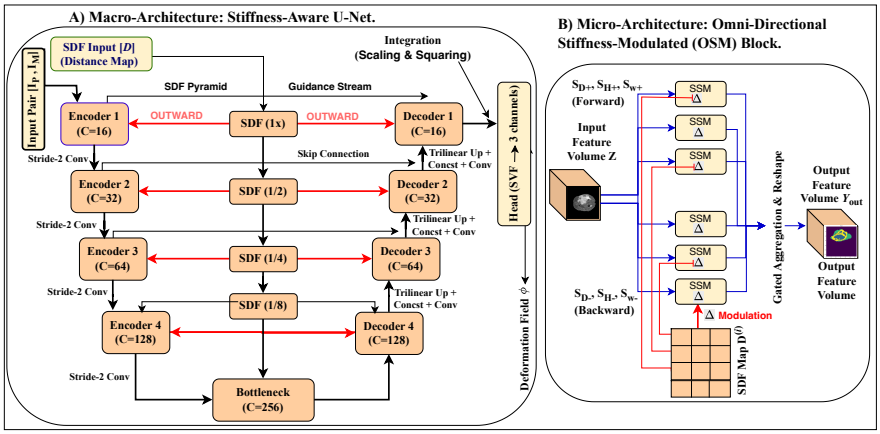

Fig. 1. The PathoMamba-U Architecture. (A) **Macro-Architecture:** A dual-stream design. The *Vision Stream* processes the image pair, while the *Guidance Stream* encodes the Signed Distance Function (SDF). Unlike standard approaches that merely concatenate priors, the stiffness information (Red Arrows) is injected directly into the hidden state evolution to spatially modulate feature extraction. (B) **Omni-Directional Stiffness-Modulated (OSM) Block:** To resolve the 1D causal bias, volumetric features are unrolled into 6 isotropic sequences ($S_{D\pm}, S_{H\pm}, S_{W\pm}$). Critically, the stiffness prior dynamically scales the discretization step size Δ. This acts as a learnable *Numerical Inertia*: low Δ enforces memory (rigidity) in healthy tissue, while high Δ enables rapid updates (plasticity) in the tumor. The network predicts a Stationary Velocity Field (SVF) v, which is integrated via scaling-and-squaring to guarantee a diffeomorphic deformation ϕ.

2.1 Diffeomorphic Stationary Velocity Formulation

To guarantee topological safety (i.e., $\det(J_\phi(p)) > 0$ everywhere), we parameterize ϕ via a **Stationary Velocity Field (SVF)** $v \in \mathbb{R}^{3\times\Omega}$. The deformation

is obtained by integrating the autonomous ODE $\frac{d\phi^{(t)}}{dt} = v(\phi^{(t)})$ over unit time $t \in [0, 1]$. We employ the *scaling-and-squaring* algorithm [1], which approximates the Lie group exponential map $\phi = \exp(v)$ via recursive composition. This analytical formulation inherently enforces invertibility and prevents folding, obviating the need for explicit topology-correction penalties.

2.2 PathoMamba: Stiffness-Modulated Dynamics

Standard Convolutional Neural Networks (CNNs) utilize fixed receptive fields that struggle to differentiate between rigid healthy tissue and plastic tumors [21]. We propose **PathoMamba**, a dual-stream architecture (Fig. 1) that injects biomechanical priors directly into the feature extraction dynamics.

Omni-Directional Stiffness Block. State Space Models (SSMs) map a 1D sequence $x(t)$ to a latent state $h(t)$. To resolve the 1D causal bias in 3D volumes, we unroll the features into 6 isotropic sequences $\{S_k\}_{k=1}^{6}$ (forward/backward along H, W, D). Each sequence is processed by a discretized SSM:

$$h_t = \overline{\mathbf{A}}(\Delta)h_{t-1} + \overline{\mathbf{B}}(\Delta)x_t, \quad \text{with} \quad \overline{\mathbf{A}} = \exp(-\Delta\mathbf{A}). \tag{1}$$

Here A and B are the standard Mamba state and input matrices [10]; A is initialized via the S4D-real scheme, and B is input-dependent. In standard Mamba, the discretization step Δ is predicted from local texture. We reinterpret Δ as a learnable **Numerical Inertia**:

- **High Inertia ($\Delta \to 0$):** The state retains long-term memory ($h_t \approx h_{t-1}$). This mimics *stiffness*, enforcing global coherence and smooth deformation in healthy tissue.
- **Low Inertia ($\Delta \to \infty$):** The state updates rapidly, reducing dependency on history. This mimics *plasticity*, allowing high-frequency local deformations inside the tumor.

Stiffness-Modulated Discretization (SMD): We explicitly modulate Δ using a Signed Distance Function (SDF) prior $\mathcal{D}$, where $\mathcal{D}_p$ is the distance to the tumor boundary ($\mathcal{D}_p < 0$ inside pathology). For a feature token at position p, we compute:

$$\Delta_p = \text{Softplus}\big(W_{\text{img}}x_p + \sigma(W_{\text{sdf}}\mathcal{D}_p)\big) + \epsilon \tag{2}$$

where W_{img} and W_{sdf} are learnable 1×1×1 convolutions (per-voxel linear projections of x_p and D_p), σ denotes the sigmoid function, and ϵ is a learnable scalar.

2.3 Therapy-Aware Biomechanical Objective

We train end-to-end minimizing $\mathcal{L}_{total} = \mathcal{L}_{sim} + \lambda_{reg}\mathcal{L}^{*}_{diff} + \lambda_{bio}(\mathcal{L}_{TABL} + \mathcal{L}_{MK})$. We employ Local Normalized Cross-Correlation (LNCC) for similarity.

1. Stiffness-Weighted Regularization. Standard regularization penalizes velocity gradients uniformly. To accommodate sliding motion at the tumor interface, we employ a spatially-weighted diffusion loss:

$$\mathcal{L}^*_{diff} = \sum_{p\in\Omega} w(p) \cdot ||\nabla v(p)||^2, \quad \text{with} \quad w(p) = \frac{1}{1+e^{-\alpha \mathcal{D}_p}}. \tag{3}$$

The weight $w(p)$ approaches 1 in healthy tissue (enforcing rigidity) and relaxes toward 0 inside the tumor (allowing non-smooth, plastic deformation). The slope α is set so that w transitions smoothly across a $\sim$2 mm band centred on the tumor boundary ($w \approx 0.5$ at the interface, saturating to 1 in healthy tissue and to 0 deep inside the tumor).

2. Trend-Aware Biomechanical Loss (TABL). Longitudinal glioma evolution is heterogeneous (e.g., simultaneous recurrence and necrosis). We propose a **Component-Wise Trend** formulation. Let $\{C_k\}$ be the connected components of the pathology mask. For each component k, we compute the log-volume ratio $\eta_k = \log(V^k_{T1}/V^k_{T0})$.

$$\mathcal{L}_{TABL} = \sum_k \frac{1}{|\Omega_k|} \sum_{p\in\Omega_k} \Psi(J_\phi(p), \eta_k) + \frac{1}{|\Omega_{shell}|} \sum_{p\in\Omega_{shell}} \text{ReLU}(|J_\phi(p)| - 1 - \delta). \tag{4}$$

where Ψ is a piecewise penalty enforcing component-specific volume trends:

$$\Psi(J_\phi(p), \eta_k) = \begin{cases} \max(0, 1 - J_\phi(p)), & \text{if } \eta_k > 0 \text{ (Growth)} \\ \max(0, J_\phi(p) - 1), & \text{if } \eta_k < 0 \text{ (Shrinkage)} \\ 0, & \text{otherwise} \end{cases} \tag{5}$$

The second term of Eq. (4), denoted $\mathcal{L}_{MK}$, is a Monro-Kellie peritumoral-shell penalty that suppresses unphysical expansion ($|J_\phi| > 1 + \delta$) in the healthy rim around the tumor, complementing $\mathcal{L}^*_{diff}$ to enforce intracranial volume conservation.

3 Experimental Results and Analysis

3.1 Experimental Setup

We validate our framework using the **BraTS-Reg Challenge** dataset [2], the clinical benchmark for longitudinal glioma registration. The cohort consists of baseline (T_0) and follow-up (T_1) multi-modal MRI scans with expert-annotated anatomical landmarks. We employ a random subject-wise split of 80% for training and 20% for testing. To ensure a fair, apples-to-apples comparison, all baseline methods were locally re-trained and evaluated on this exact data split using their official open-source implementations.

Preprocessing: To ensure consistent input distribution, all volumes are rigidly registered to the SRI-24 anatomical atlas, skull-stripped, and resampled to 1mm isotropic resolution. To fit the memory constraints of the A100 GPU while preserving anatomical context, volumes are center-cropped to $160 \times 192 \times 160$ voxels. Intensity normalization is applied per-volume via $[0, 99.5]$ percentile clipping followed by min-max scaling to $[0, 1]$. *Evaluation Metrics:* Geometric accuracy is quantified using the **Mean Target Registration Error (mTRE)**, calculated as the Euclidean distance between matched landmarks. Landmarks are *strictly held-out* during training. We further assess topological validity using the percentage of folding voxels ($|J_\phi| \leq 0$).

Implementation Details: Experiments were conducted on a single **NVIDIA A100 (40GB)** GPU using PyTorch and MONAI. We optimize the network using **AdamW** ($\text{lr} = 2e^{-4}$) with a cosine annealing schedule for 300 epochs. To manage the memory footprint of 3D state-space modeling, we utilize *Gradient Checkpointing* and Automatic Mixed Precision (AMP). The loss weights were empirically set to $\lambda_{sim} = 1.0$, $\lambda_{reg} = 1.0$, and $\lambda_{bio} = 0.1$. LNCC similarity uses a 9^3 window (MONAI default), and the SVF is integrated using 7 scaling-and-squaring steps.

At inference, the TABL objective is bypassed, as it serves strictly as a training-time biomechanical constraint. To compute the anatomical Signed Distance Function (SDF) prior ($\mathcal{D}_p$) dynamically at test time without relying on expert human-in-the-loop annotations, we employ a pre-trained auxiliary 3D nnU-Net segmenter to extract the active pathology. Because this step is autonomous, PathoMamba operates as an end-to-end automated pipeline. This auxiliary segmentation step requires approximately 1.2 s on an NVIDIA A100 GPU. Consequently, the total end-to-end clinical inference latency is $\approx$ 1.8 s (1.2 s segmentation + 0.6 s Mamba registration). To ensure a rigorously fair comparison and eliminate any informational advantage, all baseline methods (e.g., VoxelMorph, TransMorph) were strictly evaluated in their standard, fully-automated configurations utilizing only the raw T_0 and T_1 volume.

3.2 Quantitative Superiority

Table 1 summarizes the quantitative evaluation on the BraTS-Reg test set. PathoMamba establishes a new state-of-the-art, achieving a Mean TRE of 1.88 mm and a Dice score of 0.81. Unlike the accuracy-latency trade-off inherent in Transformers (TransMorph) and CNNs (VoxelMorph), our $O(N)$ State Space Model architecture surpasses Transformer accuracy while operating $2\times$ faster. PathoMamba matches heavy foundation models in accuracy while guaranteeing zero topological folds. Critically, PathoMamba eliminates topological folding (0.00%), demonstrating that stiffness-modulated discretization successfully decouples plastic tumor deformation from rigid healthy tissue, ensuring anatomical fidelity without compromising registration precision.

While PathoMamba utilizes an auxiliary nnU-Net segmenter to extract the active pathology, the architecture inherently mitigates boundary segmentation errors. Crucially, the numerical inertia (Δ_p) is co-modulated by both the soft

Table 1. Quantitative Evaluation on BraTS-Reg. We consolidate deep learning, diffeomorphic foundation models, and recent Mamba architectures by comparing PathoMamba against locally reproduced foundational baselines and state-of-the-art (SOTA) methods reported in the literature. **Bold** indicates best performance among one-shot learning methods. N/R denotes metrics not reported by the original authors. PathoMamba is the only framework to match SOTA accuracy while guaranteeing strict topological safety (0.00% folds) and high tumor overlap.

Method	Category / Backbone	mTRE (mm) ↓	Tumor Dice ↑	% Folds ↓	Time (s) ↓
Foundational Baselines & Ablations *(Reproduced Locally)*					
Affine	Linear	6.42 ± 2.1	0.45 ± 0.12	**0.00**	-
SyN (ANTs)	Optimization	3.15 ± 1.4	0.58 ± 0.15	**0.00**	120.0
VoxelMorph [3]	CNN (Displacement)	2.85 ± 1.1	0.68 ± 0.08	0.85	**0.2**
VoxelMorph-Diff [6]	CNN (SVF)	2.65 ± 1.0	0.70 ± 0.07	0.02	0.3
TransMorph [4]	Transformer (ViT)	2.15 ± 0.9	0.75 ± 0.06	0.12	1.2
RegMamba (Base) [11]	SSM (Physics-Blind)	2.24 ± 1.0	0.73 ± 0.07	0.05	0.5
BraTS-Reg State-of-the-Art *(Reported from Literature)*					
LapIRN [16]	Instance Opt.	1.82*	N/R	0.18	> 4.0
DIRAC-Opt [17]	Instance Opt.	1.76*	N/R	0.23	≈ 60.0
NICE-Net [13]	CNN (One-Shot)	1.98	N/R	0.16	0.5
BiPyramid [25]	CNN (One-Shot)	1.82	N/R	0.19	0.8
TransMorph [4]	Transformer	2.15	N/R	0.15	1.2
SCARN [8]	CNN (Attention)	1.85	N/R	N/R	1.0
DIRAC-D [18]	CNN (One-Shot)	1.88	N/R	0.09	0.9
MambaMorph [10]	SSM	1.95	N/R	0.15	0.5
VMambaMorph [23]	Visual Mamba	1.85	N/R	0.08	0.7
PathoMamba	**Stiff-SSM (Ours)**	$\mathbf{1.88 \pm 0.8}$	$\mathbf{0.81 \pm 0.05}$	**0.00**	0.6

* Denotes Median Error (as reported by authors). All other TRE values are Mean Error.

SDF prior and raw local image features. This allows the network to override inaccurate segmentations using local textures dynamically. To validate robustness, we perturbed the test masks by ± 5 mm, observing only a marginal degradation from 1.88 mm to approximately 1.94 mm mTRE, while maintaining 0.00% topological folding. However, catastrophic segmentation failures remain a limitation requiring clinical oversight.

3.3 Stratified Ablation & Component Analysis

To disentangle the mechanistic contributions of our framework, we perform an ablation study stratified by tumor evolution trend. We partition the test set into three biomechanical phenotypes: *Rapid Growth* ($\Delta V > 20\%$), *Cavity Collapse* ($\Delta V < -20\%$), and *Stable*. Table 2 details the incremental impact of each component.

As seen in Fig 2, the PathoMamba achieves high-fidelity, topology-preserving alignment with the baseline (T_0) by employing stiffness-modulated dynamics that adapt to localized volume changes, thereby preserving tumor interfaces and preventing over-smoothing across multiple subjects. The Jacobian determinant maps demonstrate a profound decoupling of deformation mechanics strictly at the tumor interface. While healthy parenchyma maintains rigid, volume-preserving incompressibility ($|J| \approx 1.0$), the active pathology exhibits the plasticity required to capture severe focal expansions and contractions ($|J|$ ranging

from 0.6 to 1.4). By confining high-frequency morphometric updates to the tumor core rather than diffusing them globally, the stiffness-modulated dynamics prevent unphysical structural folding in adjacent tissues, elegantly resolving the correspondence-regularity conflict inherent in homogeneous frameworks.

Table 2. Stratified Biomechanical Ablation Study. We evaluate the mechanistic impact of Stiffness-Modulated Discretization (SMD) and the Therapy-Aware Biomechanical Loss ($\mathcal{L}_{TABL}$) across varying tumor evolution trends. Active architectural modulation strictly outperforms passive concatenation in dynamic subgroups, while TABL refines alignment accuracy in collapsing pathologies.

Configuration	Mechanism	Loss Objective	Growth (mTRE) ↓	Shrinkage (mTRE) ↓	Stable (mTRE) ↓
(1) RegMamba (Base)	No Prior (Blind)	$\mathcal{L}_{sim} + \mathcal{L}_{diff}$	2.35	2.55	2.10
(2) + Input Concat	Passive	$\mathcal{L}_{sim} + \mathcal{L}_{diff}$	2.15	2.20	2.05
(3) + SMD (Ours)	Active Arch.	$\mathcal{L}_{sim} + \mathcal{L}_{diff}$	1.95	2.05	1.90
(4) PathoMamba	**Full Method**	$\mathcal{L}_{sim} + \mathcal{L}_{diff} + \mathcal{L}_{TABL} + \mathcal{L}_{MK}$	**1.88**	**1.82**	**1.85**

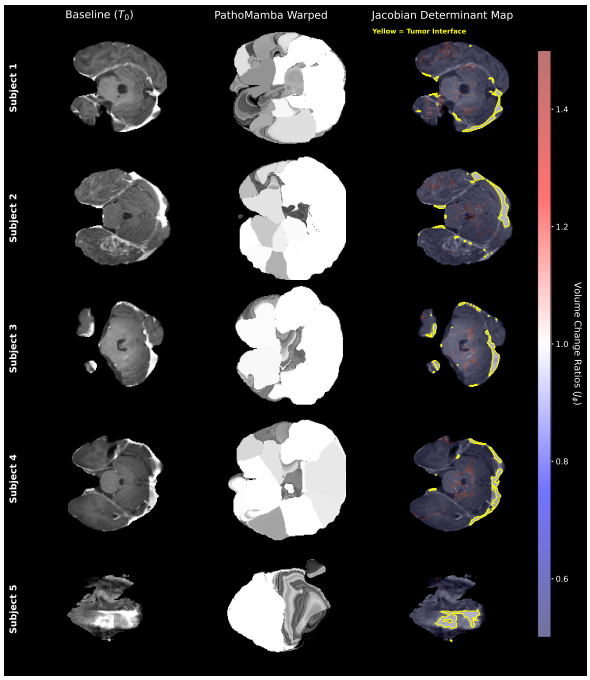

Fig. 2. PathoMamba effectively captures the biophysical interface between the plastic glioma core ($|J| > 1$) and the rigid healthy parenchyma ($|J| \approx 1$). Unlike traditional methods, the SMD block prevents the diffusion of deformation into healthy tissue, maintaining anatomical integrity.

4 Conclusion

We introduced **PathoMamba**, a stiffness-modulated framework addressing the correspondence-regularity conflict in longitudinal glioma registration. By reinterpreting the State Space Model's discretization step as learnable *Numerical Inertia*, we enabled a heterogeneous diffeomorphic flow that respects tissue-specific mechanics. Extensive validation on the BraTS-Reg benchmark demonstrates that PathoMamba matches state-of-the-art Transformer accuracy while guaranteeing strict topological safety and real-time inference. These results confirm that actively modulating architectural dynamics with physical priors outperforms passive concatenation, offering a robust solution that bridges the gap between continuum mechanics and efficient deep learning.

Acknowledgement. The authors acknowledge the following: This work was supported by the National Natural Science Foundation of China (62276116 and 62501253).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ashburner, J.: A fast diffeomorphic image registration algorithm. Neuroimage **38**(1), 95–113 (2007)
2. Baheti, B., Chakrabarty, S., Akbari, H., Bilello, M., Wiestler, B., Schwarting, J., Calabrese, E., Rudie, J., Abidi, S., Mousa, M., et al.: The brain tumor sequence registration (brats-reg) challenge: Establishing correspondence between pre-operative and follow-up mri scans of diffuse glioma patients. arXiv preprint arXiv:2112.06979 (2024)
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: A learning framework for deformable medical image registration. IEEE Transactions on Medical Imaging **38**(8), 1788–1800 (2019). https://doi.org/10.1109/TMI.2019.2897538
4. Chen, J., Frey, E.C., He, Y., Segars, W.P., Li, Y., Du, Y.: Transmorph: Transformer for unsupervised medical image registration. Medical image analysis **82**, 102615 (2022)
5. Chen, Y., Wang, X., Qian, C., Guan, X., Dai, X., Gao, P.: Mambabir: A residual pyramid network for brain image registration with state-space model. Expert Systems with Applications **297**, 129335 (2026). https://doi.org/https://doi.org/10.1016/j.eswa.2025.129335, https://www.sciencedirect.com/science/article/pii/S0957417425029501
6. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. Medical image analysis **57**, 226–236 (2019)
7. Duan, T., Chen, W., Ruan, M., Zhang, X., Shen, S., Gu, W.: Unsupervised deep learning-based medical image registration: a survey. Physics in Medicine Biology **70**(2), 02TR01 (jan 2025). https://doi.org/10.1088/1361-6560/ad9e69

8. Feng, Y., Zheng, Y., Huang, D., Wei, J., Liu, T., Wang, Y., Liu, Y.: Stepwise corrected attention registration network for preoperative and follow-up magnetic resonance imaging of glioma patients. Bioengineering **11**(9), 951 (2024)
9. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. In: First conference on language modeling (2024)
10. Guo, T., Wang, Y., Shu, S., Chen, D., Tang, Z., Meng, C., Bai, X.: Mambamorph: A mamba-based framework for medical mr-ct deformable registration. arxiv 2024. arXiv preprint arXiv:2401.13934 (2024)
11. Hu, X., Chen, J., Chen, Y.: Regmamba: An improved mamba for medical image registration. Electronics (2079-9292) **13**(16) (2024)
12. Le, G., Shu, Y., Qiao, L., Yang, L., Xiao, B., Li, W., Gao, X.: Dgmir: Dual-guided multimodal medical image registration based on multi-view augmentation and on-site modality removal. In: Gee, J.C., Alexander, D.C., Hong, J., Iglesias, J.E., Sudre, C.H., Venkataraman, A., Golland, P., Kim, J.H., Park, J. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. pp. 152–162. Springer Nature Switzerland, Cham (2026)
13. Meng, M., Bi, L., Feng, D., Kim, J.: Brain tumor sequence registration with non-iterative coarse-to-fine networks and dual deep supervision. In: International MICCAI Brainlesion Workshop. pp. 273–282. Springer (2022)
14. Menze, B., Isensee, F., Wiest, R., Wiestler, B., Maier-Hein, K., Reyes, M., Bakas, S.: Analyzing magnetic resonance imaging data from glioma patients using deep learning. Computerized medical imaging and graphics **88**, 101828 (2021)
15. Menze, B.H., Van Leemput, K., Lashkari, D., Riklin-Raviv, T., Geremia, E., Alberts, E., Gruber, P., Wegener, S., Weber, M.A., Székely, G., et al.: A generative probabilistic model and discriminative extensions for brain lesion segmentation—with application to tumor and stroke. IEEE transactions on medical imaging **35**(4), 933–946 (2015)
16. Mok, T.C., Chung, A.C.: Large deformation diffeomorphic image registration with laplacian pyramid networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 211–221. Springer (2020)
17. Mok, T.C., Chung, A.C.: Robust image registration with absent correspondences in pre-operative and follow-up brain mri scans of diffuse glioma patients. In: International MICCAI Brainlesion Workshop. pp. 231–240. Springer (2022)
18. Mok, T.C., Chung, A.C.: Unsupervised deformable image registration with absent correspondences in pre-operative and post-recurrence brain tumor mri scans. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 25–35. Springer (2022)
19. Ou, Y., Akbari, H., Bilello, M., Da, X., Davatzikos, C.: Comparative evaluation of registration algorithms in different brain databases with varying difficulty: results and insights. IEEE transactions on medical imaging **33**(10), 2039–2065 (2014)
20. Ping, X., Huang, W., Zheng, Y.: Mambamer: Adaptive eeg-guided multimodal emotion recognition with mamba. In: Gee, J.C., Alexander, D.C., Hong, J., Iglesias, J.E., Sudre, C.H., Venkataraman, A., Golland, P., Kim, J.H., Park, J. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. pp. 333–343. Springer Nature Switzerland, Cham (2026)
21. Sahiner, B., Chan, H.P., Petrick, N., Wei, D., Helvie, M., Adler, D., Goodsitt, M.: Classification of mass and normal breast tissue: a convolution neural network classifier with spatial domain and texture images. IEEE Transactions on Medical Imaging **15**(5), 598–610 (1996). https://doi.org/10.1109/42.538937

22. Shi, Y., Xia, B., Jin, X., Wang, X., Zhao, T., Xia, X., Xiao, X., Yang, W.: Vmambair: Visual state space model for image restoration. IEEE Transactions on Circuits and Systems for Video Technology **35**(6), 5560–5574 (2025). https://doi.org/10.1109/TCSVT.2025.3530090
23. Wang, Z., Zheng, J.Q., Ma, C., Guo, T.: Vmambamorph: a multi-modality deformable image registration framework based on visual state space model with cross-scan module. arXiv preprint arXiv:2404.05105 (2024)
24. Wittek, A., Miller, K., Kikinis, R., Warfield, S.K.: Patient-specific model of brain deformation: Application to medical image registration. Journal of biomechanics **40**(4), 919–929 (2007)
25. Zhang, J., Xie, X., Cheng, X., Li, T., Zhong, J., Hu, X., Sun, L., Yan, H.: Deep learning-based deformable image registration with bilateral pyramid to align preoperative and follow-up magnetic resonance imaging (mri) scans. Quantitative Imaging in Medicine and Surgery **14**(7), 4779 (2024)