Joint Decomposition and Learned Deformation for Part-to-Whole MRI to Histology Registration

Chinmayee Athalye, Winifred Trotman, Lisa M. Levorse, Eric Teunissen-Bermeo, Noah Capp, Sheina Emrani, Ranjit Ittyerah, Karthik Prabhakaran, John L. Robinson, Edward B. Lee, Daniel T. Ohm, M. Dylan Tisdall, Sandhitsu R. Das, David A. Wolk, David J. Irwin, and Paul A. Yushkevich

University of Pennsylvania, Philadelphia, PA, USA
cathalye@engineering.upenn.edu

Abstract. Linking postmortem histopathology with MRI is essential for validating neuroimaging biomarkers at the cellular level. However, registering single histology sections to 3D MRI volumes presents compounded challenges: localizing the correct slice plane within the volume, solving a part-to-whole problem caused by tissue fragmentation, and bridging fundamental differences in resolution and contrast between modalities. We present an end-to-end framework using implicit neural representations (INRs) that jointly addresses these challenges through three coordinated SIREN networks: a support network that extracts shared anatomical structure, a residual network that captures stain-specific appearance, and a deformation network that models tissue displacements. A two-phase optimization first learns global affine alignment and image decomposition, then estimates local, non-rigid, and 3D deformations with frozen affine parameters. We validate this pipeline on 18 postmortem MRI-histology pairs with 7T T2 MRI and LFB-CV stained histology slides. Our method achieves the highest robustness across initialization conditions (94–100% success rate vs. 67–100% for baselines) while maintaining comparable accuracy to established optimization-based methods on successful cases. This joint optimization avoids error accumulation from sequential pipelines and enables accurate, high-throughput registration suitable for large-scale clinico-pathological studies. We open source our code here.

Keywords: Implicit neural representations · Histology-MRI registration · Part-to-whole alignment.

1 Introduction

Linking postmortem histopathology with MRI provides a powerful tool for biological research and clinical validation. It is essential for validation of neuroimaging biomarkers against gold-standard histopathology and understanding radiological findings at the cellular level. However, registering histology sections

to MRI volumes is a challenging problem that needs to account for (1) *dimensionality and resolution gap* while mapping sub-micron resolution 2D histology sections into millimeter resolution 3D MRI volumes; (2) *large differences in signal contrasts* of the same tissue; (3) *modeling non-linear tissue deformations* like non-uniform shrinking, tearing, stretching, and geometric inconsistencies from curved slicing planes that all require modeling out-of-plane deformations; (4) *handling part-to-whole registration* problem caused by larger tissue pieces being cut into smaller fragments.

Prior work addresses these challenges through different strategies, each tackling a subset of the problem. Serial histology reconstructs adjacent sections into a 3D volume for volumetric registration [26,15], while blockface photography [24,23] and rescanning of cut tissue slabs [14] provide intermediate images that simplify slice localization. However, these approaches are resource-intensive and limit throughput. Without serial sectioning, the problem becomes severely underconstrained, as all challenges must be solved from a single, dimensionally-mismatched image pair [16,17,22]. Style transfer or generative models can bridge the appearance gap by converting histology to "MRI-like" images [26,19], but since conversion and registration are not jointly optimized, errors propagate downstream [5]. Tissue deformations are typically handled by classical non-rigid registration as a separate stage [12], compounding errors across the pipeline. No existing approach addresses all four challenges within a single jointly optimized framework.

Traditional optimization-based methods rely on intensity-based objectives like mutual information [2,18], which perform poorly for deformable registration, while keypoint-based methods struggle with tissue artifacts [10]. Both require strong initialization that is difficult to obtain for MRI-to-histology pairs [8]. CNN-based approaches offer faster inference but need large annotated datasets [7], are susceptible to domain shift [6], and operate on fixed voxel grids [25] that discard the resolution advantage of histology. Implicit neural representations (INRs) offer a compelling alternative: they parameterize images and deformation fields as continuous functions [21], naturally bridging the resolution gap while retaining per-pair optimization flexibility with differentiable regularization [3,9]. Recent work has shown INRs can jointly decompose and register 2D gene expression images [4]. We extend this to the substantially harder problem of registering a 3D MRI volume to a single histology slice in a cross-modal, part-to-whole setting.

Our contributions are threefold: (1) an end-to-end INR framework that jointly performs image decomposition and 3D-to-2D deformable registration, eliminating error propagation from sequential pipelines; (2) a support–residual decomposition that disentangles shared anatomical structure from stain-specific appearance, enabling cross-modal alignment without style transfer or intermediate modalities; and (3) a two-phase optimization with learnable affine initialization that handles slice localization in 3D MRI volume, part-to-whole alignment, and non-rigid deformation in a unified formulation. We validate on postmortem brain tissue with 7T whole-hemisphere MRI and LFB-CV staining.

2 Methods

Given a 3D postmortem MRI volume $M : \mathbb{R}^3 \to \mathbb{R}$ (moving image) and a 2D histology section $F : \mathbb{R}^2 \to \mathbb{R}^3$ (fixed image, RGB), we want to find a spatial transformation that localizes the correct oblique slice plane within M, locates the subregion of the histology slide in the coronal slice, and models regional, out-of-plane tissue deformations introduced during histological processing.

We employ three SIREN-based [20] implicit neural networks: support $\mathcal{S}$, residual $\mathcal{R}$, and deformation $\mathcal{D}$. Support $\mathcal{S}$ and residual $\mathcal{R}$ network share an identical architecture that maps normalized 2D spatial coordinates $\bar{\mathbf{x}} \in [-1, 1]^2$ to 3-channel RGB pixel values. The histology image is decomposed additively as $F(\bar{\mathbf{x}}) \approx \mathcal{S}(\bar{\mathbf{x}}) + \mathcal{R}(\bar{\mathbf{x}})$, where $\mathcal{S}$ captures anatomical structure shared with MRI and $\mathcal{R}$ captures stain-specific appearance. The deformation network $\mathcal{D}$ maps 2D normalized coordinates to a 3D displacement bounded as $\mathbf{u}(\bar{\mathbf{x}}) = \tanh(\mathcal{D}(\bar{\mathbf{x}})) \cdot s$, where s controls the maximum displacement magnitude.

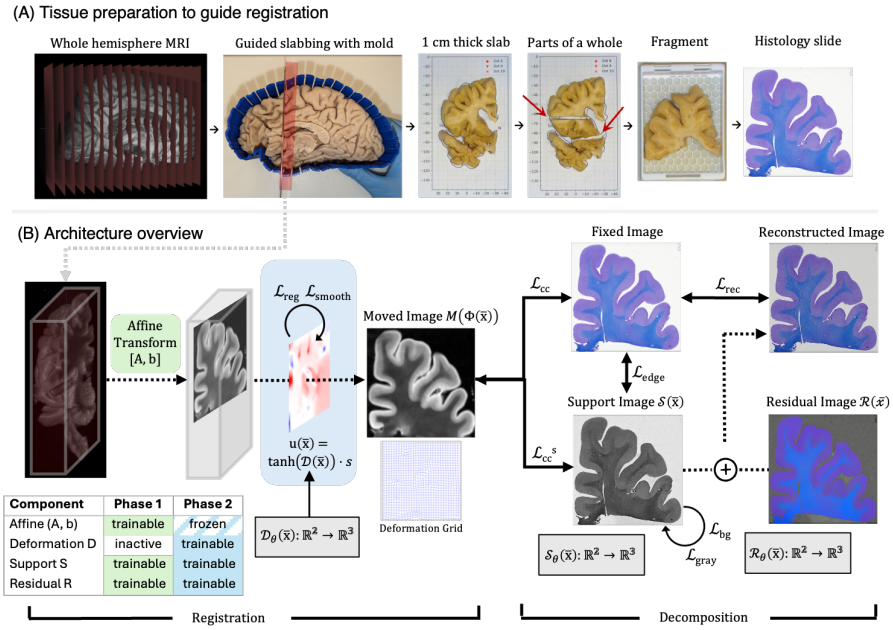

Fig. 1. Overview of the joint decomposition and registration framework. (A) Tissue preparation pipeline from whole-hemisphere 7T MRI through mold-guided slabbing to LFB-CV stained histology fragments. (B) Three SIREN networks jointly decompose the histology into shared anatomy ($\mathcal{S}$) and stain-specific residual ($\mathcal{R}$) while estimating a 3D deformation field ($\mathcal{D}$) to sample the MRI volume. The MRI is registered to both the fixed and support images and the $\mathcal{L}_{cc}$ drives the registration. Phase 1 (green) learns affine alignment and decomposition; Phase 2 (blue) adds non-rigid deformation with frozen affine parameters and continues improving the decomposition.

2.1 Data

Postmortem human brain hemispheres were fixed for 60 days and scanned at ultra-high resolution 7T MRI producing scans with 0.28 mm isotropic resolution. A patient-specific 3D printed mold, generated from the MRI segmentation, guided uniform coronal slabbing at 1 cm intervals and established a spatial frame of reference between the MRI volume and the cut tissue. The bigger slabs were cut into further 2 or 3 small pieces to fit the histology cassette (Figure 1). Slabs were paraffin-embedded, sectioned at 20 μm thickness, and stained with Luxol Fast Blue-Cresyl Violet (LFB-CV), which highlights myelin in white matter and cell bodies in gray matter, providing contrast analogous to MRI tissue boundaries. The digitized histology slides had a resolution of 0.4×0.4 μm^2. Refer to [1] for full details of the imaging and histology tissue preparation protocol.

2.2 Initialization

Using the mold-based frame of reference, we extracted a 1 cm MRI slab centered on the approximate coronal plane of each histology slide (dotted gray arrows in Figure 1). We computed fine-grained binary tissue masks for the histology slide and the MRI [13], then performed coarse alignment via moments of inertia on the masks followed by rigid registration with brute-force search over rotation angle and flip, using the open-source library Greedy. For tissue that was cut into smaller pieces, we partitioned the MRI segmentation using METIS graph partition [11] and selected the appropriate part. The resulting composite rigid transform served as initialization for the INR-based optimization.

2.3 Two-Phase Training

Training proceeds in two phases to separate global alignment from local deformation estimation.

Phase 1 (Affine + Decomposition): The affine parameters $(\mathbf{A}, \mathbf{b})$ are initialized by aligning the binary masks of the histology and MRI images and set as learnable parameters with an elevated learning rate ($10\times$ base). Networks $\mathcal{S}$ and $\mathcal{R}$ are trained simultaneously. The deformation network $\mathcal{D}$ is excluded from optimization. This phase jointly refines the slice plane in the MRI volume while learning the support–residual decomposition, i.e., style transfer.

Phase 2 (Deformation + Decomposition): After phase 1, affine parameters $(\mathbf{A}, \mathbf{b})$ are frozen and $\mathcal{D}$ is introduced with a fresh optimizer. The deformation network produces a bounded 3D displacement $\Phi(\bar{\mathbf{x}}) = \mathbf{A}\bar{\mathbf{x}} + \mathbf{b} + \mathbf{u}(\bar{\mathbf{x}})$. We add the displacement after the affine transform rather than to the input coordinates: displacement is a 3D field, and the input coordinates are 2D. It models local, non-rigid tissue deformations like warped slicing while preserving the global alignment from Phase 1.

We treat the histology as fixed so the deformation is parameterized over its 2D domain; registering in the opposite direction would define the warp over the full 3D MRI domain while only a 2D submanifold contributes to the loss, which is computationally wasteful.

2.4 Loss Functions

In total, we calculate eight different loss functions, each weighted by α_i with weights linearly interpolated between start and end values across training, and with phase-aware scheduling. We used the same hyperparameters across all cases without per-pair tuning.

Cross-correlation similarity: Following [4], we combine global normalized cross-correlation (NCC) with local NCC (LNCC, window 32×32):

$$\mathcal{L}_{\text{cc}} = -\frac{1}{2}\left[\text{NCC}(F, M(\Phi(\bar{\mathbf{x}}))) + \text{LNCC}_{\Omega}(F, M(\Phi(\bar{\mathbf{x}})))\right] \tag{1}$$

where $M(\Phi(\bar{\mathbf{x}}))$ is the moved MRI slice. This loss is computed both between the moved MRI and the fixed histology image ($\mathcal{L}_{\text{cc}}$, $\alpha_{\text{cc}} = 1$) and between the moved MRI and the support image ($\mathcal{L}^S_{\text{cc}}$, $\alpha^S_{\text{cc}} = 1$ ramped to 2). Early in Phase 1, support network $\mathcal{S}$ is uninformative and cannot drive registration alone, so $\mathcal{L}_{\text{cc}}$ against the fixed histology image anchors the alignment. Ramping α^S_{cc} lets $\mathcal{L}^S_{\text{cc}}$ take over once $\mathcal{S}$ matures. Ω denotes the binary histology tissue mask, restricting the LNCC computation to the tissue region. This masking is critical for part-to-whole registration where the histology tissue covers only a portion of the MRI slice.

Reconstruction loss: The additive decomposition is enforced via computing the mean squared error (MSE) between the fixed histology and the reconstructed image:

$$\mathcal{L}_{\text{rec}} = \frac{1}{N}\sum_{\bar{\mathbf{x}}} \|F(\bar{\mathbf{x}}) - [\mathcal{S}(\bar{\mathbf{x}}) + \mathcal{R}(\bar{\mathbf{x}})]\|^2 \tag{2}$$

and is weighted heavily ($\alpha_{\text{rec}} = 100$) to anchor the decomposition.

Grayscale loss: Since the MRI is intrinsically grayscale, we encourage the support image to be achromatic by penalizing inter-channel differences:

$$\mathcal{L}_{\text{gray}} = \frac{1}{|\Omega|}\sum_{\bar{\mathbf{x}}\in\Omega}\left[(S_R - S_G)^2 + (S_G - S_B)^2 + (S_R - S_B)^2\right] \tag{3}$$

where S_R, S_G, S_B are the RGB channels of the support image. This loss is ramped from $\alpha_{\text{gray}} = 1$ to 40, progressively enforcing grayscale structure as the decomposition matures. Having support $\mathcal{S}$ generate a 3-channel image with grayscale penalty lets $\mathcal{S}$ and residual $\mathcal{R}$ share identical architectures.

Edge matching loss: To ensure that the main anatomical GM–WM boundaries are captured by the support rather than the residual image, we match Sobel gradients of the support to those of the fixed histology:

$$\mathcal{L}_{\text{edge}} = \frac{1}{|\Omega|}\sum_{\bar{\mathbf{x}}\in\Omega}\left[|\nabla_x \tilde{S} - \nabla_x \tilde{F}| + |\nabla_y \tilde{S} - \nabla_y \tilde{F}|\right] \tag{4}$$

where $\tilde{S}$ and $\tilde{F}$ denote the luminance (grayscale) of the support and fixed images, and ∇_x, ∇_y are Sobel operators. This loss is ramped from $\alpha_{\text{edge}} = 1$ to 20.

Background variance loss: The support image should produce a constant value for the background, outside the tissue mask Ω, to avoid encoding background staining artifacts:

$$\mathcal{L}_{\text{bg}} = \text{Var}\left[\mathcal{S}(\bar{\mathbf{x}}) \mid \bar{\mathbf{x}} \notin \Omega\right] \tag{5}$$

ramped from $\alpha_{\text{bg}} = 10$ to 20 throughout training. Without this loss, the support network $\mathcal{S}$ models background artifacts stochastically whereas we want the artifacts confined to the residual image.

Together, the edge matching and the background variance losses prevent the support network from encoding extraneous tissue visible in the full MRI coronal slice but absent from the histology section. Without these masked losses, the NCC loss between the moved MRI and support image would drive the support to reproduce MRI content outside the tissue mask.

Jacobian regularization: Active only in Phase 2, this penalizes deviations of the Jacobian determinant from unity to prevent topological folding of the transformation:

$$\mathcal{L}_{\text{reg}} = \frac{1}{|\Omega|} \sum_{\bar{\mathbf{x}} \in \Omega} |1 - \det(J_{\Phi(\bar{\mathbf{x}})})| \tag{6}$$

where J is computed on the composed affine + deformation coordinates via automatic differentiation through the SIREN and weighted by a constant $\alpha_{\text{reg}} = 1$.

Displacement gradient penalty: Active only in Phase 2, this encourages spatial smoothness of the 3D deformation field by penalizing the L_2 norm of finite differences between neighboring displacement vectors:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{N} \sum_{\bar{\mathbf{x}}} \|\partial_x \text{u}(\bar{\mathbf{x}})\|^2 + \|\partial_y \text{u}(\bar{\mathbf{x}})\|^2 \tag{7}$$

where $\text{u}(\bar{\mathbf{x}}) = \tanh(\mathcal{D}(\bar{\mathbf{x}})) \cdot s$ is the scaled displacement from the deformation network. Displacement scale s is set to 0.1 to bound the maximum deformation to 10% of the image extent. The weight is ramped from $\alpha_{\text{smooth}} = 1$ to 50 during Phase 2 training.

3 Implementation Details

Both the histology image and the coronal plane of the MRI were resized to 256×256. This resizing was only for training efficiency. As the INRs use normalized spatial coordinates, the images can be sampled at original resolution during inference. All networks use sinusoidal activation layers $\sin(\omega_0 \mathbf{W}\mathbf{x}+\mathbf{b})$ with a mid-network skip connection that concatenates the input to the hidden representation at the middle layer. Support network $\mathcal{S}$ and residual network $\mathcal{R}$ share identical architectures with 5 hidden layers of width 256, frequency parameter $\omega_0 = 30$, and Fourier positional encoding with 6 frequency bands. Deformation network $\mathcal{D}$ also has 5 hidden layers of width 256 but a lower frequency $\omega_0 = 5$ and no

positional encoding, which encourages a smoother displacement field. Its final layer weights are initialized to near zero to avoid large, unstable deformations at the start of training.

We used an Adam optimizer with a base learning rate of 10^{-4} for support and residual networks and 10^{-3} for affine parameters for 200 epochs in phase 1. The deformation network was trained for 200 epochs in phase 2 with displacement scale $s = 0.1$ that bounds the maximum deformation to 10% of the image extent. The training was carried out on an NVIDIA A6000 GPU and took less than a minute each for phases 1 and 2.

4 Results

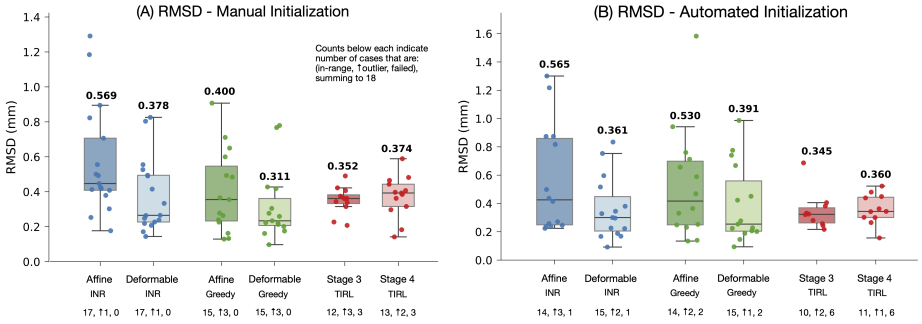

Fig. 2. Registration accuracy (S-RMSD, mm) across three methods with (A) fine-grained manual and (B) approximate automated initialization. Lower is better. Boxplots show the interquartile range. Counts below each method indicate number of cases with (in-range S-RMSD, ↑ outlier S-RMSD, failed registration) out of 18 total cases. Numbers above boxes indicate mean S-RMSD excluding outliers and failed cases. TIRL stage 3 is in-plane deformations and stage 4 adds out-of-plane deformations.

Validation scheme: We manually traced two anatomical curves (GM-WM boundary) on histology slides and corresponding MRI slabs to evaluate registration accuracy. On the MRI side, curves were traced on the best-aligned coronal slice plus four adjacent slices in each direction to account for out-of-plane deformations from cutting and histological processing. After registration, MRI contours were resampled into histology space and the symmetric nearest-neighbor root mean squared distance (S-RMSD) was calculated between the manually traced histology contour and the resampled MRI contour. We compared our method with greedy-based deformable registration [1] and Tensor Image Registration Library (TIRL) [10,24], all given the same initialization. We also compared results using fine-grained manual initialization and automated initialization based on matching moments of inertia of binary masks. We evaluated the methods on 18 MRI-histology slab pairs from 16 subjects.

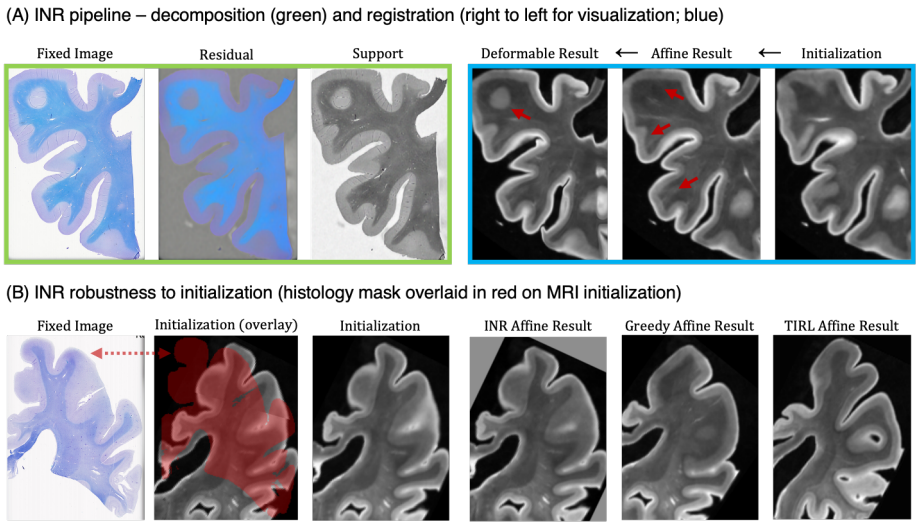

Fig. 3. Qualitative Results: (A) INR pipeline for a representative case. The decomposition shown in green. Registration progresses from initialization through affine to deformable alignment (blue); red arrows highlight regions improved by each step of the registration. (B) Robustness to initialization quality. Given a skewed automated initialization (histology mask overlaid in red to visualize the degree of skew), greedy and TIRL produce catastrophically failed registrations. INR localizes the correct coronal plane after affine registration but doesn't fully recover the rotation.

Robustness: Both INR and greedy methods achieved 100% success rate with manual initialization while TIRL completed registration for only 15/18 subjects (83%). INR was the most robust method to automated initialization completing registration on 17/18 cases (94%) compared to greedy's 16/18 (89%) and TIRL's 12/18 (67%). Notably, TIRL separates its deformable step into in-plane (stage 3) and out-of-plane (stage 4) deformations. Its performance decreases from stage 3 to stage 4, i.e., it fails to meaningfully model out-of-plane deformations which are a common artifact of histological tissue preparation. TIRL is also very sensitive to initialization as noted by its high failure rate with approximate automated initialization. The INR framework is better equipped to recover from imprecise initializations, a critical property for high-throughput settings where manual correction is infeasible. An example of this robustness is shown in Figure 3(B).

Accuracy: Among successful cases, pairwise Wilcoxon signed-rank tests revealed no significant differences between INR deformable and either TIRL Stage 4 ($p = 0.39$) or Greedy deformable ($p = 0.89$) with manual initialization. However, INR was the only method for which the deformable stage yielded a statistically significant improvement over its affine stage ($p = 0.012$, manual; $p = 0.017$, automated), indicating that the learned deformation refines alignment beyond the affine fit. The learned displacement fields are nearly diffeomorphic: across the 18 pairs the mean Jacobian determinant is 1.002 ± 0.21, only 0.28%

of points have a negative Jacobian determinant, and 5 of 18 pairs are fully diffeomorphic (no negative Jacobians). The improvements in registration at each stage of the INR pipeline can be clearly seen in Figure 3(A).

5 Conclusion

In this work, we present an end-to-end postmortem 7T T2 3D hemisphere MRI to 2D single section LFB-CV histology registration pipeline based on joint decomposition and registration using INRs. We show that our method is robust to imprecise initializations while achieving comparable accuracy to traditional optimization-based methods, making it more suitable for high-throughput registration tasks. The joint formulation allows decomposition and registration to mutually inform each other which also makes the framework naturally extensible to new stains without redesigning the pipeline. Future work includes improving automated initialization using fine-grained anatomical masks instead of binary mask alignment, testing decomposition and registration performance on other stains such as hematoxylin and eosin (H&E), and exploring loss functions and regularization to drive more precise deformable registration outside the histology tissue area. We open source our code here.

Acknowledgments: This work was supported by NIH grants R01 AG056014, P30 AG072979, and R01 AG069474.

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

References

1. Athalye, C., Bahena, A., Khandelwal, P., et al.: Operationalizing postmortem pathology-MRI association studies in Alzheimer's disease and related disorders with MRI-guided histology sampling. Acta Neuropathologica Communications (1) (May 2025). https://doi.org/10.1186/s40478-025-02030-y
2. Borovec, J., Kybic, J., Arganda-Carreras, I., et al.: ANHIR: Automatic Non-Rigid Histological Image Registration Challenge. IEEE Transactions on Medical Imaging (10) (Oct 2020). https://doi.org/10.1109/TMI.2020.2986331
3. Byra, M., Poon, C., Rachmadi, M.F., et al.: Exploring the performance of implicit neural representations for brain image registration. Scientific Reports (1) (Oct 2023). https://doi.org/10.1038/s41598-023-44517-5
4. Byra, M., Poon, C., Shimogori, T., et al.: Implicit Neural Representations for Joint Decomposition and Registration of Gene Expression Images in the Marmoset Brain. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-43999-5_61
5. Casamitjana, A., Mancini, M., Iglesias, J.E.: Synth-by-Reg (SbR): Contrastive Learning for Synthesis-Based Registration of Paired Images. In: Simulation and Synthesis in Medical Imaging. Springer International Publishing, Cham (2021). https://doi.org/10.1007/978-3-030-87592-3_5

6. Chen, J., Liu, Y., Wei, S., et al.: A survey on deep learning in medical image registration: New technologies, uncertainty, evaluation metrics, and beyond. Medical Image Analysis (Feb 2025). https://doi.org/10.1016/j.media.2024.103385
7. Eppenhof, K.A.J., Lafarge, M.W., Moeskops, P., et al.: Deformable image registration using convolutional neural networks. In: Medical Imaging 2018: Image Processing. SPIE (Mar 2018). https://doi.org/10.1117/12.2292443
8. Ferrante, E., Paragios, N.: Slice-to-volume medical image registration: A survey. Medical Image Analysis **39**, 101–123 (Jul 2017). https://doi.org/10.1016/j.media.2017.04.010
9. van Harten, L., Herten, R.L.M.V., Stoker, J., et al.: Deformable Image Registration with Geometry-informed Implicit Neural Representations. In: Medical Imaging with Deep Learning. PMLR (Jan 2024)
10. Huszar, I.N., Pallebage-Gamarallage, M., Bangerter-Christensen, S., et al.: Tensor image registration library: Deformable registration of stand-alone histology images to whole-brain post-mortem MRI data. NeuroImage (Jan 2023). https://doi.org/10.1016/j.neuroimage.2022.119792
11. Karypis, G., Kumar, V.: METIS: A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices (1997)
12. Khandelwal, P., Duong, M.T., Levorse, L.M., et al.: VIOLET: Volumetric Image registration via Optimization and Learning for Efficient image Translation. In: Simulation and Synthesis in Medical Imaging. Springer Nature Switzerland, Cham (2026). https://doi.org/10.1007/978-3-032-05573-6_9
13. Khandelwal, P., Duong, M.T., Sadaghiani, S., et al.: Automated deep learning segmentation of high-resolution 7 Tesla postmortem MRI for quantitative analysis of structure-pathology correlations in neurodegenerative diseases. Imaging Neuroscience (May 2024). https://doi.org/10.1162/imag_a_00171
14. Lahna, D., Roese, N., Woltjer, R., et al.: Postmortem 7T MRI for guided histopathology and evaluation of cerebrovascular disease. Journal of Neuropathology & Experimental Neurology (1) (Jan 2023). https://doi.org/10.1093/jnen/nlac103
15. Mancini, M., Casamitjana, A., Peter, L., et al.: A multimodal computational pipeline for 3D histology of the human brain. Scientific Reports (1) (Aug 2020). https://doi.org/10.1038/s41598-020-69163-z
16. Pichat, J., Iglesias, J.E., Nousias, S., et al.: Part-to-Whole Registration of Histology and MRI Using Shape Elements. In: 2017 IEEE International Conference on Computer Vision Workshops (ICCVW). IEEE, Venice, Italy (Oct 2017). https://doi.org/10.1109/ICCVW.2017.21
17. Pichat, J., Iglesias, J.E., Yousry, T., et al.: A Survey of Methods for 3D Histology Reconstruction. Medical Image Analysis (May 2018). https://doi.org/10.1016/j.media.2018.02.004
18. Pielawski, N., Wetzer, E., Öfverstedt, J., et al.: CoMIR: Contrastive Multimodal Image Representation for Registration. In: Advances in Neural Information Processing Systems. Curran Associates, Inc. (2020)
19. Qin, C., Shi, B., Liao, R., et al.: Unsupervised Deformable Registration for Multi-modal Images via Disentangled Representations. In: Information Processing in Medical Imaging. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-20351-1_19
20. Sitzmann, V., Martel, J., Bergman, A., et al.: Implicit Neural Representations with Periodic Activation Functions. In: Advances in Neural Information Processing Systems. Curran Associates, Inc. (2020)

21. Sun, S., Han, K., You, C., et al.: Medical image registration via neural fields. Medical Image Analysis (Oct 2024). https://doi.org/10.1016/j.media.2024.103249
22. Unberath, M., Gao, C., Hu, Y., et al.: The Impact of Machine Learning on 2D/3D Registration for Image-Guided Interventions: A Systematic Review and Perspective. Frontiers in Robotics and AI (Aug 2021). https://doi.org/10.3389/frobt.2021.716007
23. Ushizima, D., Chen, Y., Alegro, M., et al.: Deep learning for Alzheimer's disease: Mapping large-scale histological tau protein for neuroimaging biomarker validation. NeuroImage (Mar 2022). https://doi.org/10.1016/j.neuroimage.2021.118790
24. Wang, Y., Ho, W., Huszar, I.N., et al.: Precise MRI-histology coregistration of paraffin-embedded tissue with blockface imaging. Imaging Neuroscience (Aug 2025). https://doi.org/10.1162/IMAG.a.106
25. Wolterink, J.M., Zwienenberg, J.C., Brune, C.: Implicit Neural Representations for Deformable Image Registration. In: Proceedings of The 5th International Conference on Medical Imaging with Deep Learning. PMLR (Dec 2022)
26. Yushkevich, P.A., Muñoz López, M., Iñiguez de Onzoño Martin, M.M., et al.: Three-dimensional mapping of neurofibrillary tangle burden in the human medial temporal lobe. Brain (9) (Sep 2021). https://doi.org/10.1093/brain/awab262