



Published in final edited form as:

Neuroimage. 2025 January ; 305: 120981. doi:10.1016/j.neuroimage.2024.120981.

RegBoost: Enhancing mouse brain image registration using geometric priors and Laplacian interpolation

Atchuth Naveen Chilaparasetti^a, Andy Thai^{a,*}, Pan Gao^b, Xiangmin Xu^{a,b}, M. Gopi^{a,*}

^aDepartment of Computer Science, University of California, Irvine, Irvine, CA 92617, USA

^bDepartment of Anatomy and Neurobiology, School of Medicine, University of California, Irvine, Irvine, CA 92617, USA

Abstract

We show in this work that incorporating geometric features and geometry processing algorithms for mouse brain image registration broadens the applicability of registration algorithms and improves the registration accuracy of existing methods. We introduce the preprocessing and postprocessing steps in our proposed framework as RegBoost. We develop a method to align the axis of 3D image stacks by detecting the central planes that pass symmetrically through the image volumes. We then find geometric contours by defining external and internal structures to facilitate image correspondences. We establish Dirichlet boundary conditions at these correspondences and find the displacement map throughout the volume using Laplacian interpolation. We discuss the challenges in our standalone framework and demonstrate how our new approaches can improve the results of existing image registration methods. We expect our new approach and algorithms will have critical applications in brain mapping projects.

Keywords

Registration; Geometry processing; Laplacian; Mouse brain; Neuroimaging

1. Introduction

Image registration methods are fundamental techniques in biomedical imaging and brain research. These tasks involve estimating the deformation field between a pair of images that will transform one image to align with another while maintaining accuracy and regularity.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

*Corresponding authors. andy.thai@uci.edu (A. Thai), gaop2@uci.edu (P. Gao).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Atchuth Naveen Chilaparasetti: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis. **Andy Thai:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis. **Pan Gao:** Writing – review & editing, Writing – original draft, Resources, Data curation. **Xiangmin Xu:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Conceptualization. **M. Gopi:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Conceptualization.

These alignments, in turn, enable researchers to superimpose and integrate data obtained from different samples or diverse modalities into a shared reference frame. This facilitates easier comparisons between brain datasets, opening avenues for comprehensive analysis and interpretation of diverse datasets.

Neuroscience research has evolved from the study of individual brain slices to the adoption of whole-brain analyses. This broader approach includes comprehensive methods such as cell type profiling (Kim et al., 2017), neural circuit mapping (Oh et al., 2014), and disease pathology characterization (Liebmann et al., 2016) performed at the whole-brain level. This transition to holistic brain studies reveals the inadequacy of 2D registration in capturing the brain's complex three-dimensional architecture and function. Consequently, the use of 3D registration techniques has become essential for accurately aligning and integrating diverse datasets, ensuring more accurate representations and understanding of brain structure and function.

However, achieving meaningful and accurate cross-modal comparisons or data integration from different samples requires aligning all data onto a common coordinate atlas. In light of this imperative, our proposed image registration approach seeks to align brain section image data derived from serial two-photon tomography (STPT) (Ragan et al., 2012) with the Allen Mouse Brain Common Coordinate Framework (CCF) (Wang et al., 2020).

Our study focuses on image data obtained from STPT techniques, which combine high-speed multi-foci two-photon imaging with state-of-the-art tissue sectioning. This combination allows for rapid ex vivo imaging of organs and tissues, including whole-brain imaging at submicron resolutions. These advanced imaging platforms inherently provide 2D images that can be stacked to construct a coherent 3D volume. This eliminates challenging steps in aligning stacks of 2D slices that would be necessary with conventional section-based histology (Agarwal et al., 2017).

The Allen Brain Institute has taken on a global initiative to provide a common coordinate reference system known as the Allen Mouse Brain Common Coordinate Framework (CCF) (Wang et al., 2020). This atlas includes an average intensity profile template derived from 1,675 animal subjects imaged using STPT and expert-annotated markings. We spatially align anatomical structures in brain data from image slices to the Allen CCF, enabling precise inter-subject comparisons and analysis.

One of the motivations behind the need for image registration lies within the task of labeling and segmentation. Manually annotating individual datasets is cumbersome, and automatic segmentation methods may struggle to account for the large number of annotations required to label an imaged dataset fully to the detail of a template-annotated brain image. Registration allows users to easily apply labels defined in an annotated template to an input dataset by transforming one dataset to align and match with another. Our work is motivated by the need to annotate an input dataset by transforming it to match an expertly annotated template from the Allen CCF.

This paper introduces a brain image registration framework integrating geometric and intensity-based approaches to align whole-brain mouse image data spatially. Building on

the work of Agarwal et al. (2016), which focused on registering 2D image slices, we extend this methodology to align 3D image volumes. We present a geometry-based method that can register images using curves and surfaces that define anatomical regions of the brain. We describe techniques for extracting these geometric structures from the average intensity template, demonstrate the method's performance, and compare them with existing methodologies. Furthermore, we demonstrate how our methods can be applied to enhance current registration methods by applying our proposed techniques as part of a two-step process.

1.1. Main contributions

Our main contributions can be summarized into the following:

Automated image registration framework: We introduce our method, RegBoost, a series of preprocessing and postprocessing operations as part of an image registration framework that enhances the performance of existing methods by incorporating geometric features.

Rotation correction algorithm: We propose a novel rotation correction algorithm that estimates the rotation parameter necessary to align brain volumes with the coordinate axes. This step is fundamental for a more precise image registration.

Correspondence finding methods: We present diverse methods for establishing spatial correspondences between image stacks representing brain volumes, including mapping between point clouds, curve matching, and section matching. These methods enhance the accuracy of registration.

Non-linear registration with 3D Laplacian interpolation: Building upon the work of Agarwal et al. (2016) in 2D registration, we extend the concept to 3D registration. We demonstrate that piece-wise 2D registration can result in non-smooth deformations throughout the volume, leading to the development of a 3D Laplacian interpolation method to correct this issue.

In summary, our research offers a comprehensive and automated image registration framework that combines geometric features to improve the accuracy of existing brain image registration methods. We also explore potential directions for further research in this field.

2. Related work

The biomedical image registration problem requires deforming one input image, defined as the moving image, into the spatial features of another input image, defined as the fixed image. The work in this field has developed much over the years and has been extensively covered in various reviews (Maintz and Viergever, 1998; Zitová and Flusser, 2003; Klein et al., 2009).

Optimization methods:

Earlier methods approach registration as an optimization problem with energy functions for each image pair to model deformation fields; a review of the registration algorithms that take this approach is presented in Song et al. (2017). While medical imaging technology originated in the 1970s, it was not until the 1990s that improvements in imaging technology enabled image-based processing algorithms to offer substantial analysis. Early applications in medical image data, specifically nonrigid registration with contrast-enhanced MRI data, can be seen in Rueckert et al. (1999), where the authors use affine transformations and cubic B-splines to model breast deformation. This classic approach is still used and iterated upon in modern practices. Tustison et al. (2009) uses intensity-based registration and randomly selected points within moving and fixed images to interpolate using B-splines. Qiu et al. (2009), Sun et al. (2011) approach registration with CT imaging data by applying thin plate spline (TPS) methods. Various toolboxes such as Klein et al. (2010) and ANTs (Avants et al., 2008) also use these optimization methods to achieve registration.

Diffeomorphic methods:

Shen and Davatzikos (2002) introduces geometric moment invariants to establish anatomical correspondences, different from other deformation methods that rely on optimizing image similarity metrics. Voxels are assigned an embedding of features, and each vector is used to compute a similarity score based on displacements. Later advancements developed diffeomorphic mapping systems for registration, with large deformation diffeomorphic metric mapping (LDDMM) (Beg et al., 2005) being the first to be applied to volumetric data. Demons-based approaches, such as diffeomorphic demons (Vercauteren et al., 2009) and spherical demons (Yeo et al., 2009), approach registration by identifying key voxel boundaries as “demons” and utilizing diffusing models to displace them. Patch-based registration methods (Dalca et al., 2016) map voxels to new possible locations and introduce Markov random fields to obtain the optimal mapping between voxels. Tward et al. (2020) investigates these applications where intensity profiles differ in cross-modal MRI brain data. Chandrashekhar et al. (2021) builds off of an LDDMM mapping algorithm to present a cloud-based cross-modal pipeline to perform intensity correction and registration in data with intensity inhomogeneity.

Learning-based methods:

With the rise of convolutional networks (Krizhevsky et al., 2012) in the past decade, learning-based approaches for registration problems have increased significantly. Initially used for image recognition and object segmentation, applications have since spread to medical image-based analysis. More modern frameworks leverage learning-based approaches such as VoxelMorph (Balakrishnan et al., 2018) and SynthMorph (Hoffmann et al., 2022) with significant improvements in speed and run time but come with their downsides; such models often have difficulty correlating long-range spatial relationships between moving and fixed images and may also generalize poorly outside of their initial scope of training data. Additionally, medical image datasets tend to be small and focused on specific pathologies, which may cause difficulty in learning the position-coding of images, alongside the high computational complexity of neural networks and long training times.

The rise of transformer networks in natural language processing (Vaswani et al., 2017), then in image-based tasks, have gained popularity with applications in image registration (Trappolini et al., 2021). However, these models tend to require large amounts of GPU memory. The most recent developments, such as TransMorph (Chen et al., 2022), use hybrid models to mix convolutional neural networks and vision transformers-based architecture. Manifold learning methods also can be used to approach registration from the scope of dimensionality reduction; Bashiri et al. (2019) expands upon the use of the Laplacian eigenmaps, introduced in Wachinger and Navab (2010) for learning structural representation, to apply to partially overlapped multi-modal images.

Point cloud-based methods:

Point cloud-based methods (Besl and McKay, 1992), Zhang (1994) introduce iterative closest point (ICP) registration to account for affine transformations between point clouds, along with several variants (Fitzgibbon, 2003), Rusinkiewicz and Levoy (2001) that have followed. Myronenko and Song (2010) further iterates on these methods by utilizing probabilistic methods with Gaussian mixture models to register point clouds. Their method introduces the coherent point drift (CPD) method, which provides registration robust to noise, outliers, and missing points but can be computationally expensive and slow for larger inputs. These point cloud-based methods typically downsample the volume and only deal with surface registration.

Invertibility:

The invertibility of transformations within image deformation is also a critical topic for registration practices. Specifically, invertible transformations establish that points mapped between moving and fixed images have a one-to-one correspondence. This one-to-one correspondence ensures that information is preserved and not lost when transforming image data through these deformation mappings. Fulfilling this property also results in outputs where an inverse deformation field can be applied to retrieve the original moving image. Additionally, these one-to-one correspondences can be pivotal in avoiding artifacts and irregular distortions that may be present in noninvertible transformations. These topics of invertibility and smooth transformations have been studied as a variational formulation of an image-matching problem in Dupuis and Grenander (1998). Reuter et al. (2010) also discusses the necessity of robust inverse registration to avoid bias and artifacts in computational neuroimaging and the drawbacks of non-invertible mappings in alignment with longitudinal MRI data.

Our method:

Compared to these previous algorithms, we build on previous work in optimization-based methods and differentiate our method from these in that we register the point clouds represented by these surfaces and then get a smooth deformation field over the entire volume using Laplacian interpolation. We base our method on the previous work of Agarwal et al. (2016), where they have proposed using Laplacian interpolation and point correspondences to get a smooth deformation between 2D images. We improve upon this method by extending Laplacian interpolation to a 3D space of 3D correspondences of both internal and boundary feature points to achieve improved and more coherent registration

and alignment across the image axes. There exist similar methods that match based on geometric features, such as block matching (Ourselin et al., 2000), which uses image patches and cross-correlation/mutual information to find close patches, and mBrainAligner (Qu et al., 2022), which detects high curvature corner points and matches them together. Our method is similar to mBrainAligner, but our method contrasts by attempting to get the best possible close alignment before correspondence finding, then searching in neighborhoods to find a local match to each point. Although some matches can be inaccurate, our iterative ICP operation will eventually converge to a good match; our nonlinear registration ignores some of the effects of incorrect matching as we filter outliers and minimize large sets of linear equations. Lastly, our method can complement other methods to improve existing registration results.

3. Background

The input to our algorithm consists of two image stacks — averaged template images from the Allen CCF 2017 (Allen Institute for Brain Science, 2017; Lein et al., 2007) and brain section images acquired through advanced microscopy. Although our algorithm can perform bidirectional registration, in this context, we will designate the microscopic images as the moving image I_m and the template images as the fixed image I_f .

We first define the image registration task by denoting the moving and fixed images as a function from coordinate space to intensity values, such that the moving image is represented as $I_m: \mathbb{R}^3 \rightarrow \mathbb{R}$ and the fixed image is represented as $I_f: \mathbb{R}^3 \rightarrow \mathbb{R}$. Medical and neuroimaging data can be made up of various formats, ranging from 4D multichannel volumes to singular 2D grayscale images. For our study, we assume that I_m and I_f are standard 3D single-channel volumes. Though we use single-channel volumes, the transformations done on this channel can be applied to other channels. Our implementation uses the red channel as it contains the most prominent signal information for our STPT data.

We define the transformation $T: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ as a function that maps every voxel within I_m to a voxel in I_f and computes a displacement for each voxel. We define the function $d: \mathbb{I}_m \times \mathbb{I}_f \rightarrow \mathbb{R}$ to estimate the similarity between I_m and I_f . Given I_m and I_f , we want to calculate the transformation T to maximize the similarity between the transformed I_m and I_f . This problem statement is represented as the function r in Eq. (1).

$$r(I_m, I_f) = \max_{T(I_m)} d(T(I_m), I_f) \quad (1)$$

The output of registration is the transformed image $T(I_m)$, which is acquired after I_m is displaced and deformed to align with I_f spatially. This transformed image can easily label and segment the moving image using the annotated volume from the Allen Common Coordinate Framework (CCF), as depicted in Fig. 15.

3.1. Data acquisition

3.1.1. Ethics statement: All experiments were conducted according to the National Institutes of Health guidelines for animal care and use and were approved by the University of California, Irvine Institutional Animal Care and Use Committee (IACUC, protocol #: AUP-22–163) and the Institutional Biosafety Committee (IBC).

3.1.2. Acquisition protocols: The data samples I_m from our study consist of brains from mice between 2 to 18 months of age. Mice were anesthetized with a continuous flow of 1%–2% isoflurane. Then, using a stereotaxic injection machine, the viral tracer was injected into the mouse brain. Specifically, the brains of our mice were injected with a combination of Adeno-associated virus (AAV) helpers (pENN.AAV.CamKII 0.4.Cre.SV40, Addgene viral prep #105558-AAV1, 5.3×10^{13} GC/ml + AAV8-DIO-TC66T-2A-eGFP-2A-oG, Salk Institute, CA, USA, 2.36×10^{13} GC/ml) and engineered pseudo-typed rabies virus (EnvA-RV-SAD G-DsRed, CNCM, 2.1×10^9 IU/ml), which selectively label specific cell types within the injection brain region and monosynaptic afferent cells for the whole brain.

After an incubation period of around one month (three weeks of AAV helper and nine days of engineered pseudo-typed rabies virus), the mice underwent transcardial perfusion using 1x PBS first and 4% PFA later. Following dissection, the mouse brains were post-fixed in 4% PFA overnight. Subsequently, the brains were transferred to a 1x PB solution with 0.01% sodium azide.

Next, the mouse brains were embedded in a 4% oxidized agarose solution using an appropriate mold. After the agarose solidified, the brain, still encased in agarose, was immersed in acrylamide overnight. The agarose-embedded brain was baked at 40 degrees Celsius for two hours on the subsequent day. After baking, the brains were carefully removed from the oven and stored in a container filled with 1x PB solution, preparing them for imaging on a TissueCyte 1600FC platform.

We secured the agarose-embedded brain on a glass slide using adhesive to facilitate the imaging process. Magnets on the opposite side were adhered to a metal plate at the base of the TissueCyte sample container filled with 1x PB solution. TissueCyte employs serial two-photon tomography (STPT), an automated block-face imaging technique. It captures repetitive images of the brain's surface through multiphoton illumination and physically sections the imaged area with the assistance of an integrated vibratome. This process continues until the entire brain is covered. Subsequently, the collected individual tile images, saved in the TIFF image format, were assembled in two dimensions to construct a mosaic, resulting in a comprehensive cross-sectional image. These mosaic images were then further stacked in three dimensions and converted to the NIfTI file format, culminating in a volumetric representation of the entire brain.

For our moving images, our image stack I_m consists of around 280 stitched brain sections of size 8716×11236 pixels with a pixel-to-pixel spacing of $1.25 \mu\text{m}$ and a section-to-section spacing of $50 \mu\text{m}$. Unlike (Agarwal et al., 2016), the TissueCyte platform allows scanned volumes to achieve good vertical alignment between sections for all the sections stacked over each other. These sections do not require further processing to ensure further alignment.

We downsample each section so that stacking all of those sections results in a moving image volume of pixel dimensions $280 \times 432 \times 568$ with a voxel spacing of $50 \mu\text{m} \times 25 \mu\text{m} \times 25 \mu\text{m}$. Because our data is anisotropic, we upsample the z -axis by a factor of two to maintain isotropy at 25μ along each axis. Each image contains three color channels. We perform registration across the red channel associated with the pseudo-typed rabies virus. For this study, we used a 2-month-old male mouse specimen (C57BL/6J) injected with a modified rabies virus in the dorsal hippocampus CA1 (dCA1), targeting excitatory cells, as our moving data.

We utilized the Allen Mouse Brain Common Coordinate Framework as our fixed image. The Allen CCF 2017 (Allen Institute for Brain Science, 2017; Lein et al., 2007) contains an average template containing the computed average intensities and an annotation template where the entire volume is segmented and labeled. For all of our experiments, we used the average template at a 25-micron resolution with pixel dimensions $528 \times 320 \times 456$ as our I_f , with a voxel spacing representing a spatial volume of $25 \mu\text{m} \times 25 \mu\text{m} \times 25 \mu\text{m}$.

3.2. Overview

The primary objective of our study is to establish a robust correspondence between the imaged volumes I_m and I_f to enhance the registration outcomes progressively. Our registration algorithm has the following steps, which can be used independently or combined with an existing method for better registration.

Our RegBoost framework starts with a series of preprocessing steps, defined in Section 4.1 and depicted in Fig. 1, to perform an affine alignment for I_m to the space of I_f . The initial problem of ensuring the brain sample orientation perfectly parallels the microscope imaging sensor is nontrivial. As a result, sectional images may not exhibit perfect symmetry with respect to the placement of the axes in the image domain. To compensate for any angular deviations, a post-imaging axis alignment is computed for the rotation transformation in Section 4.1.1. Then, we estimate the scaling parameters needed for initial bounding box alignment in Section 4.1.2.

Next, we define the features of each image by extracting the contours that define the prominent structures of each brain section. These serve as the basis for establishing correspondences between the imaging volumes. Such contour extraction methods are explained in Section 4.1.3. We identify the point clouds from these contours that define the boundary and internal surfaces in Section 4.1.4. Using these point clouds, we estimate the translation parameters in Section 4.1.5. We then further estimate the best possible affine registration by finding the correspondences in Section 4.1.6. Using these same correspondences, these point clouds are refined through an iterative closest point operation in Section 4.1.7.

We define the Laplacian functions used to compute deformation fields across the volume in Section 4.2. Using these 3D Laplacian operators, we present nonlinear registration in Section 4.3 as part of the RegBoost postprocessing step. Between the preprocessing and postprocessing steps of RegBoost, any section matching algorithm can be used. We explain our method, LaplacianReg, along with other methods for section matching in Section 4.4.

Finally, we present the results of RegBoost with different section matching methods in Section 5. We discuss our generalized registration pipeline and summarize how incorporating our method into existing registration methods can improve registration outcomes on our data in Section 6.

3.3. Data/code availability

Data used in this registration study will be publicly available on Mendeley data. The Allen CCF annotations and averaged templates can be downloaded directly from the Allen Brain Map Community website using the AllenSDK at <https://community.brain-map.org/t/allen-mouse-ccf-accessing-and-using-related-data-and-tools/359>

All code in this study is publicly available and can be reproduced from the GitHub repository at <https://github.com/naiveen/UCI-ALLEN-BrainRepositoryCode>.

4. Methods and materials

We define our framework, RegBoost, as a set of methods that can be applied to enhance registration results. We categorize these methods into preprocessing and postprocessing steps. The preprocessing steps aim to compute a good affine registration, which is then passed to a first-pass nonlinear registration where any reasonable registration method can be used. The result from this nonlinear registration is then used to reslice and map sections between the moving and fixed images. Using these matched sections, a second pass nonlinear registration in the postprocessing phase using Laplacian interpolation. These preprocessing and postprocessing steps combined constitute RegBoost.

4.1. RegBoost preprocessing: affine registration

4.1.1. Rotation correction: When working with algorithms that optimize using gradient descent, providing input with good initial starting points for optimization is good practice to avoid local minima. Particularly with registration, we want to provide a moving data image aligned to the image coordinate axes as closely as possible. However, it is unlikely that a given set of brain tissue slices will be perfectly aligned and symmetrical to the image axes during the imaging process; some imperfections may occur during data acquisition, and the retrieved imaging data may be tilted by a few degrees from the ideal axis. Optimizing for this step can improve registration outcomes, especially when aligning to perfectly symmetrical template volumes.

We first compute a rotation matrix to help correct the alignment of our moving image so it is much closer to the template data. Fig. 2 shows an example of this misalignment and alignment correction. It is important to note that, unlike (Agarwal et al., 2016), our images do not exhibit any significant damage and can be stacked along the coronal axis to create a 3D volume. However, the plane of symmetry of the template and the dataset may have minor misalignments that may cause suboptimal starting positions for gradient descent optimization operations, creating the need to correct these deviations by finding an appropriate rotation matrix. This issue does not arise within the data used in Agarwal et al. (2016) because the stacking algorithm forces the plane of symmetry to be perpendicular to the horizontal plane, whereas in our case, a pre-stacked set of images is used as the input.

There is no guarantee that the brain was baked into the medium in an exact perpendicular orientation during the data acquisition process.

We have observed that many coronal mouse brain images contain a fissure that neatly divides the slice into two perfectly symmetrical halves. Given that we anticipate only slight rotations, we can reasonably assume this dividing line is vertical or near-vertical.

To compute and approximate a plane that fits the brain fissure, we process each image slice and find the overall bounding box of the brain. We then utilize edge detectors to retrieve features from the center-most regions of the brain. We detect vertical or near-vertical lines near the bounding box center using these edge features. These lines are represented as line segments made up of two points. Since multiple lines may be detected throughout a single section, we take the two points of each detected line segment in the section and average the coordinates of the top-most and bottom-most points to create a new set of points. This new set of points makes up a line segment for each section in the volume. We can approximate the sagittal plane by aggregating all the points that make up the averaged line segments across the volume sections and applying singular value decomposition to these points. This resultant plane effectively bisects the brain volume into two symmetrical halves.

Without loss of generality, assume that the sagittal plane of the template volume lies within the XY-axis and its normal vector aligns with the z -axis. By aligning the normal of the sagittal plane of the moving image volumes with the z -axis, we can estimate the rotation matrix necessary to align the coordinate axes of I_m to I_f . Applying this rotation matrix to the moving image dataset I_m gives us the now-aligned images, as seen in Fig. 2. Table 1 also shows that applying rotation correction improves scores across all methods. We note that our rotation correction algorithm only intends to correct rotations between images with slight deviations and expects the fixed and moving images to be of the same orientation.

4.1.2. Scaling estimation: After the initial alignment concerning rotation, the next step in the affine registration is to get the best alignment in terms of scaling and translation. This is done to achieve a good initial alignment, decreasing the search space size needed to find correspondences later. We already account for the rotation component in the preceding section with our rotation-correction operation.

The next step involves estimating the scaling and translation parameters for the coarse alignment of the moving image I_m with the fixed image I_f . To estimate the scaling parameters in all three dimensions, we determine a bounding box that encapsulates the brain volume by examining 2D slices across each dimension of the 3D volume. As the tissue and background are easily distinguishable, a simple thresholding process can be performed, transforming it into a binary image. Regions with high-intensity values corresponding to brain tissue are set to 1, while darker-intensity regions like the background are set to 0. The active pixels in each section are then quantified as the ratio between the number of nonzero pixels and the total number of pixels in the section. These quantified active pixel numbers for each section along each axis are plotted against their respective image slice numbers, as demonstrated in Figs. 3 and 4. The plots show the span along each axis where brain tissue is present. It can be observed that there is a sudden increase and sudden decrease where brain

tissue is present. The active pixel count is close to zero for regions with no tissue. With this observation, we used a threshold of 0.05 to calculate the range of sections that comprise the brain. Once the spatial extents for I_m and I_f are established across each axis, the scaling parameter for each axis is calculated by dividing the spatial extent of I_m by that of I_f along the corresponding axis.

This proposed algorithm provides an accurate, noise-resistant method to infer the spatial extent of the brain volume within the imaging data, aiding in accurately determining the scaling parameters. Traditional methods involve computing minimum and maximum positions of non-zero pixel values and fitting bounding rectangles to each section. However, this method is susceptible to noise. Our approach is more robust to noisy pixels as we use non-zero pixel counts averaged across many sections. Additionally, we can compute tighter, more accurate bounding boxes over simpler, traditional methods.

4.1.3. Contour extraction: Following rotation correction, we extract relevant contours from the dataset. These contours represent important features throughout the brain and the boundary and inner surfaces, as depicted in Fig. 5. Calculating contours that accurately represent the internal regions is a challenging task due to the wide range of intensities present in the image, but it is necessary to retrieve and establish the correspondences needed for accurate registration. Attempting to compute contours using standard edge detectors often results in the detection of numerous spurious edges. We have observed that the robustness and reliability of the computation of the outer boundary contours are much higher than that of the computed internal contours. Thus, we separately compute these two types of contours as follows.

Our approach to eliminating spurious edges involves retrieving and segmenting specific structures of interest from the rest of the image. These segmented regions can be represented as binary images and, when smoothed, yield clean contours that accurately represent structures of interest. Our initial step in obtaining the outer surface contours segments the brain region from the surrounding noise. We achieve this segmentation by applying a suitable intensity threshold. The intensity values within our STPT data are measured in arbitrary units (a.u.), where values are dependent on factors like the excitation source, detector sensitivity, and sample preparation. Given that our imaging settings consistently produce brain regions with intensities greater than 200, thresholding is straightforward, as values above this intensity threshold will correspond with brain tissue. In contrast, regions below this threshold will typically be the background. However, it is important to note that extensive research is dedicated to accurately segmenting brain regions from surrounding noise in more complex scenarios (Fawzi et al., 2021).

To compute the inner contours, we observe that each slice contains noticeable dark regions, primarily consisting of air and bright tissue regions. To binarize the image, we apply Otsu's thresholding method and then employ a Canny edge detector on the binary image. Additionally, we filter out smaller edges, retaining only those edges with a length of at least n times the image height. We chose $n = \frac{1}{4}$ for this value. This filtering process yields clean edges defining the boundary and inner structures.

4.1.4. Point cloud representation: After extracting the relevant contours from the brain sections, these contours will be organized to represent the spatial information inherent in 3D volumes. We choose to represent our collected contours in the form of point clouds, where these curves and surfaces can be effectively represented as points so that each point can be characterized by its coordinate values and normal vectors. When assembled, they form a point cloud that represents the brain volume. These point clouds serve as the basis for establishing correspondences between I_m and I_f . Examples of these point clouds are visualized in Fig. 7.

In addition to their Cartesian coordinates, each point within our dataset is characterized by a normal vector of unit magnitude. To calculate the normal vector for each point, we utilize a k-d tree; this tree structure facilitates efficient querying of neighboring points for any given point within the point cloud. We consider a search radius of 3 pixels for each point and query all points falling within this radius from the target point. These neighboring points constitute the set used to calculate the normal vector for the focal point. Then, we calculate the normalized vectors extending from the current point to each of its neighbors within the radius. Next, we perform singular value decomposition (SVD) on these vectors (Gopi et al., 2000). This method can be generalized for 2D and 3D spaces to calculate the normal vectors.

For each point, we accumulate the points in the direction of the normal vector in both directions and consistently orient the normal vector to point toward the lower-intensity regions. This strategic orientation allows for proper calibration of normal vectors in scenarios where contours are near each other, such as within narrow tunnels. By incorporating these normal vectors and their orientations alongside their Cartesian coordinates, we enhance the information and precision available for each point within our dataset, facilitating better correspondences in our registration framework.

Point clouds can be used to represent surface boundaries. However, due to the noisy nature of points from the inner contours, combining the inner and outer contours into a single point cloud is impractical. This approach introduces excessive noise, particularly when computing normals. To circumvent this issue, we create two separate point clouds: one comprised of features from the outer contours and the other from the inner contours as seen in Fig. 14. Making use of the fact that we compute inner and outer point clouds independently, we use these point clouds separately while computing the normal vectors. We use only the outer point cloud for affine registration and use outer and inner contours together for nonlinear registration. This iterative approach makes the registration process more accurate.

4.1.5. Translation parameters: We estimate the translation components for our affine registration matrix once we extract the point clouds in the preceding section. To calculate this, we align the centroids of the point clouds representing I_m and I_f .

4.1.6. Correspondences: The next step in our registration framework is establishing correspondences. However, matching every voxel in our moving image I_m with a corresponding voxel in the fixed image I_f is challenging and computationally impractical. We avoid doing this for every voxel and simplify the correspondence problem by finding

correspondences to only the point clouds computed earlier, which typically represent prominent structures in I_m and I_f . In 2D, this equates to curve matching, while in a 3D context, this process equates to surface matching. We begin by working with the outer point cloud, then progressively moving towards matching the point clouds representing the inner surfaces. Correspondences obtained at each step are accumulated and used for nonlinear interpolation of the deformation field.

First, to calculate correspondences, we take advantage of the fact that both point clouds are coarsely aligned after an initial affine registration. Our approach involves constructing a k-d tree for the fixed image I_f 's point cloud to enable efficient querying. For every point in the moving image I_m , we query the 100 nearest neighbors in the I_f point cloud. Among these 100 neighbors, we assess the angle between their normal vectors. We exclude any outliers with angles of more than a 5-degree difference. We select the nearest point from the remaining candidates as the correspondence for the I_m point. We refine the correspondences further by considering the distance between them. We filter the outliers by removing the points corresponding to the top 1% largest distances. This method can be generalized for correspondences in 2D and 3D space. For our affine registration, we calculate the point correspondences in 3D. Fig. 6 shows example correspondences in a 2D slice.

4.1.7. ICP registration: In Sections 4.1.1, 4.1.2 and 4.1.5, we find parameters to aid in creating a coarse alignment between I_m and I_f . Once the coarse alignment of both fixed and moving images is achieved, a second affine transformation operation is needed to create a final fine alignment from I_m to I_f . To accomplish this task, we utilize iterative closest point (ICP) registration to compute this matrix.

We define a set of points from I_m as (x_m, y_m, z_m) and their correspondence points from I_f as (x_f, y_f, z_f) . From these points, we want the best affine matrix P that can transform I_m to the space of I_f . The resultant matrix P accounts for additional scaling, translation, and rotation components and a small shear component. We aim to decompose the matrix and make the transformation properly rigid. A slight shear component is acceptable because of the nonlinear registration that follows this step.

4.2. Laplacian function-based deformation field

We use the Laplacian to represent the deformation field. Here, we describe the formulation of computation of the Laplacian, which is eventually used in the postprocessing of RegBoost, as well as for preliminary non-linear registration, which we call LaplacianReg. This algorithm computes a deformation field that transforms the space of I_m to that of I_f using a set of correspondences. This vector field consists of displacement vectors for each voxel in I_m , which, when applied, transforms the voxel to its corresponding voxel in I_f .

Let the deformation field be represented by ϕ such that for a voxel with coordinates x in I_f , its corresponding voxel will be $x + \phi(x)$. We ensure that this displacement field representing the deformation is smooth by constraining it using Laplace's equation in Eq. (2):

$$\nabla^2 \phi = 0. \quad (2)$$

We consider the displacements in each axis independent of the other and derive the following definitions in Eq. (3).

$$\nabla^2 \phi_x = 0, \quad \nabla^2 \phi_y = 0, \quad \nabla^2 \phi_z = 0. \quad (3)$$

When using Laplacian interpolation, we approximate the discrete Laplacian operator ∇^2 using the finite difference method. This gives Eq. (4), Eq. (5), and Eq. (6) for 1D, 2D, and 3D spaces, respectively.

$$\nabla^2 \phi = \frac{\phi(x+1) + \phi(x-1) - 2 \cdot \phi(x)}{2} = 0 \quad (4)$$

$$\nabla^2 \phi = \frac{\phi(x+1, y) + \phi(x-1, y)}{2} + \frac{\phi(x, y+1) + \phi(x, y-1) - 4 \cdot \phi(x, y)}{2} = 0 \quad (5)$$

$$\begin{aligned} \nabla^2 \phi = & \frac{\phi(x+1, y, z) + \phi(x-1, y, z)}{2} + \frac{\phi(x, y+1, z) + \phi(x, y-1, z)}{2} \\ & + \frac{\phi(x, y, z+1) + \phi(x, y, z-1) - 6 \cdot \phi(x, y, z)}{2} = 0. \end{aligned} \quad (6)$$

The above equations hold through for every point in the respective 1D, 2D, and 3D spaces. We further describe the construction of a linear system of equations of the form $A\phi = B$, of which, upon solving, we get the displacement value at each point in the volume space. Though we describe the process for 3D space, it can be easily generalized for 1D and 2D spaces. In addition, we have additional constraints from the correspondences. Let $C_m = (u, v, w)$ represent the set of points in I_m that have correspondences $C_f = (u', v', z')$ in I_f . This gives the Dirichlet boundary conditions, as depicted in Eq. (7).

$$\begin{aligned} \phi C_f = (d_x, d_y, d_z) &= C_m - C_f \\ \phi_x C_f &= u - u', \\ \phi_y C_f &= v - v' \\ \phi_z C_f &= w - w'. \end{aligned} \quad (7)$$

Combining the above equations, we can form a linear system of equations in Eq. (8).

$$\begin{aligned} A\phi_x - B_x &= 0, \\ A\phi_y - B_y &= 0, \\ A\phi_z - B_z &= 0. \end{aligned} \tag{8}$$

The row vectors of A take the coefficients of terms in Eq. (4), Eq. (5), and Eq. (6), except for the rows corresponding to I_m points, in which case, they represent the Dirichlet boundary condition. The vectors B_x , B_y , and B_z are zero everywhere except for the rows corresponding to C_m , in which case it is d_x , d_y , and d_z , respectively. Note that A is a sparse matrix with dimensions easily in the range of millions. Even with its highly sparse nature, it is challenging to solve.

We have observed acceptable execution times of around 250 s for computing displacements across one axis using the linear solver restarted GMRES (LGMRES) algorithm (Baker et al., 2005). When considering computational speed and accuracy, we aim to allow for an error of 1 voxel unit for our correspondences. Experimenting with different values yielded a relative tolerance of 0.01 for optimizing the Laplace equations. With approximately 30,000 correspondences, we have a final MSE of about 300. We computed the absolute error between correspondences before and after registration and observed errors of 1–5 pixels, with values in the latter range being rare.

4.3. RegBoost postprocessing: nonlinear registration

After all of the affine registration steps in preprocessing, we aim to get matched sections by doing a first-pass nonlinear registration before applying our postprocessing step, depicted in Fig. 8. By performing this first-pass nonlinear registration, we can match the sections in the moving image to those in the fixed image as an identity function. 2D correspondences are then calculated between these matched sections. These 2D correspondences can be incorporated as 3D points by incorporating the section number as the third coordinate. We find correspondences between matched sections by aligning segmented contours within I_m to their corresponding contours in I_f . For this, we query the 50 nearest points and exclude neighbors whose normal vectors deviate beyond a 5-degree threshold, selecting the closest point as a correspondence, similar to in Section 4.1.6. This method is similar to the curve-matching method introduced in Agarwal et al. (2016). However, when characterizing the normal vectors, we also use their orientation as described in Section 4.1.4. This helps us get better matches when curves are close to each other, such as when delineating between blood vessels.

The 2D nonlinear registration, as introduced in Agarwal et al. (2017), operates by performing piecewise interpolation between each pair of matched sections. This process sequentially stacks the deformation fields to construct the final 3D deformation map. However, a critical limitation of this method is its inability to ensure continuity in the resultant deformation maps between sections. This discontinuity often leads to a jarred

appearance of the images along axes other than the one in which the sections are sliced and registered. While images of coronal sections exhibit appropriate alignment post-registration as the slicing is done across the coronal axis, there is a marked emergence of artifacts in the axial and sagittal views. This highlights the challenge of maintaining spatial continuity across different views when utilizing 2D nonlinear registration. Some of these artifacts are shown in Figs. 9 and 10.

In contrast, our experiments with direct 3D nonlinear registration did not exhibit such artifacts, indicating better performance in maintaining spatial continuity. To take advantage of the strengths of both approaches, we have developed a combined methodology that involves stacking up all 2D points to generate 3D correspondences distributed throughout the 3D volume. These 3D correspondences are then employed to perform 3D Laplacian registration. This integration of 2D and 3D nonlinear registration methods in our approach addresses the limitations of 2D registration. A key factor contributing to this improvement is the comprehensive influence exerted on each voxel's deformation by all the correspondences in our 3D volume, extending beyond just the correspondences from the voxel's original section index. This results in a more regularized algorithm, introducing more constraints that must be satisfied in estimating the deformation. Using 2D correspondences, which are less noisy, and combining the precision inherent in 2D alignment with the spatial continuity derived from 3D registration allows for a more coherent and artifact-free alignment across all views and sections.

4.4. Section matching

The postprocessing step of RegBoost in Section 4.3 aims to refine closely aligned 3D volumes and boost the registration results. The better the initial alignment between I_m and I_f , the better the sections are matched. These matched sections serve as input to the postprocessing step. We describe some candidates for section matching algorithms and show that RegBoost's postprocessing improves the registration result in all of our cases. We present a simple keyframe matching algorithm that matches sections between the affine registered results. We also present LaplacianReg, Elastix (Klein et al., 2010), and ANTs (Avants et al., 2008) symmetric normalization (Avants et al., 2011) to do a firstpass nonlinear registration and reslice the sections. After reslicing, we get a new transformed moving image with sections that can be directly matched with the sections of I_f .

4.4.1. Keyframe matching: After retrieving the affine registration result, we can use an area-based keyframe matching method for section matching. In this method, the area represented by each section is computed by calculating the number of active pixels in that section. The active pixels are calculated as described in Section 4.1.2. When plotting the areas with the section indices, we get plots that are similar in shape. By matching the extrema in both plots and using them as keyframes, we can compute matched sections by interpolating between the matched extrema using 1D Laplacian registration. Table 2 shows applying RegBoost atop of keyframe matching improves the mutual information score, but keyframe matching performs poorly relative to the other methods because the keyframe matching assumes all regions of a section follow the same deformation, which is not accurate in reality.

4.4.2. LaplacianReg: In the postprocessing step of RegBoost, we apply 3D Laplacian registration on 2D correspondences found between 2D image slices. In contrast, within LaplacianReg, we directly identify the 3D correspondences and use them for Laplacian registration, as portrayed in Fig. 11. This algorithm takes the resampled and resliced affine registered result from the preprocessing step of RegBoost as the moving image and recomputes the point clouds. We collect the point clouds representing boundary and inner surfaces as described in Section 4.1.4. We then establish correspondences between the updated point clouds according to the process described in Section 4.1.7. We first use only the outer point clouds from both volumes to obtain the required correspondence. When these correspondences are used in Laplacian registration, it results in the transformed moving image as shown in Fig. 12 and Fig. 13. It can be observed that the sections align perfectly at the outer contours, but internal regions still suffer from misalignment.

We can expand our approach to include additional correspondences computed from internal surface contours to mitigate the artifacts observed within the internal alignments. However, this has proven to be particularly challenging, given the noisy nature of these internal structures. Notably, the point clouds representing the external surface are comparatively cleaner and facilitate easier correspondence identification. Directly including these internal point clouds in our current operation poses significant challenges in accurately computing normals, which can lead to suboptimal correspondence matching. An example of internal point clouds is depicted in Fig. 14. We can mitigate this noise by using matched sections to reduce the 3D search space into a more manageable 2D context. This motivates us to propose using matched sections within the postprocessing step of RegBoost.

The results from LaplacianReg indicate that we need good correspondences to get good results. It can be inferred from Table 2 that using RegBoost with LaplacianReg improves the mutual information score. For all results depicted in this table, LaplacianReg only uses external surface correspondences.

4.4.3. Other nonlinear registration methods: Despite imperfect section-matching, RegBoost can enhance the registration results outputted from keyframe matching and LaplacianReg methods. A good first-pass nonlinear registration can assist with better section matching, which RegBoost can further improve. We experiment with using RegBoost on ANTs Symmetric Normalization (SyN) (Avants et al., 2008, 2011) and affine + nonlinear B-spline registration from Elastix (Klein et al., 2010) for section matching. These methods are widely acclaimed and popularly used for nonlinear alignment tasks.

When using methods from ANTs and Elastix, registration results depend on the intensity profiles of I_m and I_f . To optimize for this factor, we rescaled the intensity profile of I_m to match the intensity profile of I_f . Our brain tissue data typically exhibits intensity ranges between 200 to 1000 a.u. However, due to viral injections, specific voxels representing activated cells exhibit exceptionally high intensities, reaching up to 30,000 a.u. To handle these outliers, we clipped the intensity values above the 99th percentile to the 99th percentile value itself and rescaled the values to span from 0 to 255. Rescaling assists in getting better alignments and registration results and standardizes the mutual information scores between different images to a standard scale. When using Elastix, we utilized a rigid + nonlinear

B-spline operation and extensively fine-tuned the parameters to optimize the registration outcomes. The initial parameters for our Elastix model were taken from Par0033 (Bink, 2016a) from the Elastix model zoo (Bink, 2016b). We ran a grid search on the set of parameters for this model to optimize its performance for our data further. These parameters have been provided in the code repository. For SyN, we used the default parameters from the ANTs toolkit.

We found that the results after using RegBoost were more pronounced when using Elastix than when using ANTs. Elastix provided decent internal alignment that RegBoost can improve upon. ANTs SyN, on the other hand, performs significantly better in the initial alignment, which makes the improvements from RegBoost more subtle and more challenging to spot. Nevertheless, we observed a general improvement in all methods by applying RegBoost.

We have also run experiments with the SynthMorph (Hoffmann et al., 2022) deep-learning framework with our STPT data. However, as discussed in Section 2, these methods operate on downsampled volumes and are primarily trained on MRI human brain scans. Our STPT mouse brain data differs significantly from the training images. Thus, these models have difficulty extrapolating to our data modality and produced suboptimal external and internal alignments, making them poor test cases for RegBoost improvements.

5. Results

We present the quantitative results of our study in the following tables. Table 1 shows the results of our rotation correction operation on the advanced Mattes mutual information score between different methods. Table 2 shows the results of our registration operations compared to Elastix and ANTs. Fig. 15 shows a 3D visualization of the Allen CCF annotations applied to a registered mouse brain volume. Figs. 16–19 shows the final results of using RegBoost with Elastix and ANTs. Our preprocessing and postprocessing steps have proven crucial in incrementally improving the alignment results, as seen from Tables 1 and 2.

We use the Mattes mutual information (Maes et al., 1996) for our quantitative analysis. Though other measures exist, such as the Hausdorff distance (Huttenlocher et al., 1993) and the Dice coefficient (Dice, 1945), these measures require some ground truth data for comparison. Our data modality of STPT mouse brains is high resolution and relatively niche; thus, creating manual, human-annotated labels for our data or finding similar data with extensive ground truth annotations is time-consuming. We note that the values from the mutual information score can only give a rough estimate of the global registration quality and may not capture tiny nuances, as it depends on intensity profiles, which other measures can potentially address (Heinrich et al., 2012). However, as our data contains cell activations with high-intensity values, these cells may interfere with descriptor-based methods and cross-correlation, as compared to mutual information's computation between similarity of intensity histograms.

The results from Elastix have some imperfections at the external boundary, as seen in Figs. 16 and 17. Elastix methods predominantly rely on parameterized B-spline models to estimate deformation fields. The reliance on parameters, specifically those part of the mathematical modeling, inherently limits the degree of freedom, leading to scenarios in which complex deformation fields are only approximated rather than precisely modeled. This is because these models can have limited numbers of parameters and hence limited degrees of freedom to capture complex deformation fields. This can allow for some error, as a modeled curve may miss some data points. In contrast, our interpolation of deformation at correspondence points computes a deformation vector for each point in space, making it a non-parameterized model similar to ANTs. LaplacianReg also uses a similar method and exhibits no artifacts at the external surface. By applying RegBoost on Elastix, we corrected these artifacts at the surface and fine-tuned the alignment of the internal regions. All of these factors contribute to a higher mutual information score.

We have generally observed very good results in registration when using ANTs. When applying RegBoost on top of it, we could correct for minor internal alignments as seen in Fig. 18. However, these are very minute corrections, and the improvements in registration score stem from the corrections in external alignment depicted in Fig. 19.

To analyze the invertibility of our output, we computed the discrete Jacobian determinant over the deformation field. We observe that 0.6% of voxels throughout the volume have a negative determinant. As seen in Fig. 20, these voxels with negative Jacobian determinants lie at or near our correspondences, and the deformation for the regions between them is smooth.

For our intended applications for segmentation and cell detection, we process cell counts that occur near detected cells and away from the boundaries of the brain tissue. So, the negative Jacobian determinant values observed in the boundary do not compromise our specific application. As demonstrated by our results, our deformations effectively register these regions. However, for applications requiring invertibility, the presence of negative Jacobians may be a limitation. Some approaches, such as those proposed in Pal et al. (2022), aim to mitigate negative Jacobians to ensure deformation field invertibility. Future work will focus on refining our deformation field to minimize these negative Jacobian determinants.

6. Discussion and conclusion

To summarize, our RegBoost method, which consists of affine preprocessing and nonlinear postprocessing steps, can improve the quality of existing registration algorithms, which we use for section matching, by using correspondences between boundary and edge features within input data.

Our preprocessing pipeline starts with rotation correction to symmetrically align the volume. Our experiments consistently reveal that aligning the axes of the moving image I_m with the coordinate axes as an initial step provided benefits in the registration pipeline outputs. This alignment consistently improved the mutual information score by at least 0.02 in all experiments. The rationale behind this improvement lies in establishing a symmetry

concerning the coordinate axes, which, in turn, serves as a solid starting point for methods employing gradient descent.

This alignment also initiates a well-structured coarse alignment, enhancing the establishment of reliable correspondences. The results of rotation correction are visually demonstrated in Fig. 2. In particular, the center groove is well vertically aligned, and the lower ends of the sections divided by this center groove align precisely on a horizontal line. A particularly significant outcome is the appearance of lobes, as depicted in the last column of the figure. These lobes should ideally emerge at the same section, but due to initial misalignment, they are slightly tilted, causing one lobe to appear ahead of the other. Following our correction process, both lobes align within the same section. This alignment significantly aids in establishing section-to-section correspondences, which is a crucial step in Laplacian interpolation.

Following rotation correction, we proceed with scaling and transforming the moving image I_m to perform affine registration, as outlined in Section 4.1. The objective at this stage is to coarsely align I_m with the fixed image I_f to facilitate the establishment of correspondences, which are used for Laplacian registration. We have shown the challenges in using 3D volume-to-volume correspondences and defined our postprocessing step in Section 4.3 to use 2D correspondences from matched sections.

When using only these 2D correspondences between the sections, registering corresponding sections independently will output discontinuities when these sections are stacked up and resampled from a different axis. To mitigate this issue, we aggregate all 2D correspondences throughout the volume to perform 3D interpolation, allowing the displacement from a correspondence in one section to affect the deformation in other sections. This approach reduces the degrees of freedom and serves as a regularizer for the registration process. As a result, we have produced improved registration outcomes as presented in Section 5.

However, this 3D approach produces more accurate results but has considerable memory requirements. Constructing the Laplacian matrix to represent the entire 3D space can be challenging due to its size. For an input volume represented by dimensions $n \times w \times h$, the corresponding 3D Laplacian matrix would have dimensions of $k_3 \times k_3$, where $k_3 = n \times w \times h$. In comparison, 2D sectional registration is similar to performing a piecewise Laplacian and requires a dimension of $k_2 \times k_2$, where $k_2 = w \times h$.

It is important to note that our method is sensitive to correspondences. Laplacian interpolation seeks to interpolate deformations at these correspondences and computes a deformation vector for each point in space. This leads to over-fitting because of the large degree of freedom for the deformation field. Therefore, establishing robust correspondence is paramount. Achieving such robustness requires an iterative refinement of correspondences. However, due to the lack of function-based modeling in Laplacian interpolation, finding a new set of correspondences can only occur after resampling the entire volume from the previous registration iteration. This approach can be time-consuming and suboptimal for an iterative algorithm.

Finally, we standardized the intensity range to a standard scale from 0 to 255. This adjustment proved to be beneficial, particularly for registration using Elastix. This improves mutual information metrics and simplifies the establishment of correspondences for the overall pipeline. As shown in Table 2, this preprocessing step yields positive results.

All our experiments were on a workstation with 256 GB of system memory (RAM) and a 24-core CPU with a 2.1 GHz clock speed. Elastix and ANTs SyN take approximately 2 min to complete. In comparison, our semi-optimized Python implementation of Laplacian interpolation in 3D space takes around 15 ± 2 min to finish. Approximately 8 min of runtime was spent in the Laplacian registration step; the rest was spent at I/O, indicating potential for further optimization.

6.1. Future work

Future work for our method can encompass multiple directions. Our registration framework currently uses an average intensity template volume as the fixed image. Since our method only requires retrieving boundary structures from our image data, a potential extension of this work can include applications that use segmented templates, such as the annotated template within the Allen CCF, as the fixed image instead of an average intensity template.

Another avenue for improvement is to reduce the computation times for our method. Better algebraic and geometric multi-grid linear solvers are needed to solve the linear system of equations. With faster solvers in place, we could iteratively refine the function, achieving better correspondences at each iteration, ultimately resulting in improved outcomes.

Future work can also explore the impact of different functions used in the registration process, such as MIND (Heinrich et al., 2012) and normalized cross-correlation.

Furthermore, although our work focuses on registering STPT-scanned mouse brains, our method is versatile enough to extend to other modalities outside mouse brain image data as a future application. One modality can include LightSheet-scanned cleared brain tissues, live mouse calcium imaging, and human brain data.

Lastly, a promising direction to explore would be in investigating the invertibility of our deformation fields. Currently, our Laplacian operation does not ensure bijection due to the existence of negative Jacobian determinants within the deformation field. As part of our future work, we plan to identify the source of these negative Jacobian determinant values and eliminate the causes to ensure this bijection.

Acknowledgments

The authors would like to thank Lydia Ng for her assistance with information regarding the Allen CCF. Research reported in this publication was supported by the National Institute on Aging of the National Institutes of Health under award number U01AG076791. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Data availability

Data used in this registration study will be made publicly available on Mendeley data at <https://data.mendeley.com/preview/y7wys7pbxv?a=de9ec347-0beb-4dbe-9aa5-9697a19f8feb>.

The Allen CCF annotations and averaged templates can be downloaded directly from the Allen Brain Map Community website using the AllenSDK at <https://community.brain-map.org/t/allen-mouse-ccf-accessing-and-using-related-data-and-tools/359>.

All code in this study is publicly available and can be reproduced from the GitHub repository at <https://github.com/naiveen/UCI-ALLEN-BrainRepositoryCode>.

References

- Agarwal Nitin, Xu Xiangmin, Gopi M, 2016. Robust registration of mouse brain slices with severe histological artifacts. In: Proceedings of the Tenth Indian Conference on Computer Vision, Graphics and Image Processing. ICVGIP '16, Association for Computing Machinery, New York, NY, USA, 10.1145/3009977.3010053.
- Agarwal Nitin, Xu Xiangmin, Gopi M, 2017. Geometry processing of conventionally produced mouse brain slice images. *J. Neurosci. Methods* 306, 10.1016/j.jneumeth.2018.04.008.
- Allen Institute for Brain Science, 2017. Common coordinate framework version 3. https://download.alleninstitute.org/informatics-archive/mouse_ccf_3_1_mm/annotation/ccf_2017/compacted/.
- Avants BB, Epstein CL, Grossman M, Gee JC, 2008. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* 12 (1), 26–41. 10.1016/j.media.2007.06.004, Special Issue on The Third International Workshop on Biomedical Image Registration – WBIR 2006. [PubMed: 17659998]
- Avants Brian B., Tustison Nicholas J., Song Gang, Cook Philip A., Klein Arno, Gee James C., 2011. A reproducible evaluation of ANTs similarity metric performance in brain image registration. *Neuroimage* 54 (3), 2033–2044. 10.1016/j.neuroimage.2010.09.025. [PubMed: 20851191]
- Baker AH, Jessup ER, Manteuffel T, 2005. A technique for accelerating the convergence of restarted GMRES. *SIAM J. Matrix Anal. Appl.* 26 (4), 962–984. 10.1137/S0895479803422014.
- Balakrishnan Guha, Zhao Amy, Sabuncu Mert Rory, Gutttag John V., Dalca Adrian V., 2018. VoxelMorph: A learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging* 38, 1788–1800, <https://api.semanticscholar.org/CorpusID:52281312>.
- Bashiri Fereshteh S., Baghaie Ahmadreza, Rostami Reihaneh, Yu Zeyun, D'Souza Roshan M., 2019. Multi-modal medical image registration with full or partial data: A manifold learning approach. *J. Imaging* 5 (1), 10.3390/jimaging5010005.
- Beg Mirza Faisal, Miller Michael I., Trounev Alain, Younes Laurent, 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.* 61, 139–157, <https://api.semanticscholar.org/CorpusID:17772076>.
- Besl Paul J., McKay Neil D., 1992. A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 239–256, <https://api.semanticscholar.org/CorpusID:21874346>.
- Bink Diewertje Ilse, 2016a. Neuropathological changes in mouse models of cardiovascular diseases. <https://api.semanticscholar.org/CorpusID:263460140>.
- Bink Diewertje Ilse, 2016b. Par0033. <https://web.archive.org/web/20221130100007/https://elastix.lumc.nl/modelzoo/par0033/>.
- Chandrashekhar Vikram, Tward Daniel J., Crowley Devin, Crow Ailey K., Wright Matthew A., Hsueh Brian Y., Gore Felicity, Machado Timothy A., Branch Audrey, Rosenblum Jared S., Deisseroth Karl, Vogelstein Joshua T., 2021. CloudReg: automatic terabyte-scale cross-modal brain volume registration. *Nat. Methods* 18 (8), 845–846. [PubMed: 34253927]

- Chen Junyu, Frey Eric C., He Yufan, Segars William P., Li Ye, Du Yong, 2022. TransMorph: Transformer for unsupervised medical image registration. *Med. Image Anal.* 82, 102615. 10.1016/j.media.2022.102615. [PubMed: 36156420]
- Dalca Adrian V., Bobu Andreea, Rost Natalia S., Golland Polina, 2016. Patch-based discrete registration of clinical brain images. In: Wu Guorong, Coupé Pierrick, Zhan Yiqiang, Munsell Brent C., Rueckert Daniel (Eds.), *Patch-Based Techniques in Medical Imaging*. Springer International Publishing, Cham, pp. 60–67.
- Dice Lee R., 1945. Measures of the amount of ecologic association between species. *Ecology* 26 (3), 297–302.
- Dupuis Paul, Grenander Ulf, 1998. Variational problems on flows of diffeomorphisms for image matching. *Quart. Appl. Math.* LVI (3), 587–600.
- Fawzi Ali, Achuthan Anusha, Belaton Bahari, 2021. Brain image segmentation in recent years: A narrative review. *Brain Sci.* 11 (8), 1055. 10.3390/brainsci11081055. [PubMed: 34439674]
- Fitzgibbon Andrew William, 2003. Robust registration of 2D and 3D point sets. *Image Vis. Comput.* 21, 1145–1153, <https://api.semanticscholar.org/CorpusID:7576794>.
- Gopi M, Krishnan S, Silva CT, 2000. Surface reconstruction based on lower dimensional localized delaunay triangulation. *Comput. Graph. Forum* 19 (3), 467–478. 10.1111/1467-8659.00439, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-8659.00439>.
- Heinrich Mattias P., Jenkinson Mark, Bhushan Manav, Matin Tahreema, Gleeson Fergus V., Brady Sir Michael, Schnabel Julia A., 2012. MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration. *Med. Image Anal.* 16 (7), 1423–1435. 10.1016/j.media.2012.05.008, Special Issue on the 2011 Conference on Medical Image Computing and Computer Assisted Intervention. [PubMed: 22722056]
- Hoffmann Malte, Billot Benjamin, Greve Douglas N., Iglesias Juan Eugenio, Fischl Bruce, Dalca Adrian V., 2022. SynthMorph: learning contrast-invariant registration without acquired images. *IEEE Trans. Med. Imaging* 41 (3), 543–558. [PubMed: 34587005]
- Huttenlocher DP, Klanderman GA, Rucklidge WJ, 1993. Comparing images using the hausdorff distance. *IEEE Trans. Pattern Anal. Mach. Intell.* 15 (9), 850–863. 10.1109/34.232073.
- Kim Yongsoo, Yang Guangyu Robert, Pradhan Kith, Venkataraju Kannan Umadevi, Bota Mihail, García del Molino Luis Carlos, Fitzgerald Greg, Ram Keerthi, He Miao, Levine Jesse Maurica, Mitra Partha, Huang Z. Josh, Wang Xiao-Jing, Osten Pavel, 2017. Brain-wide maps reveal stereotyped cell-type-based cortical architecture and subcortical sexual dimorphism. *Cell* 171 (2), 456–469.e22. 10.1016/j.cell.2017.09.020. [PubMed: 28985566]
- Klein Arno, Andersson Jesper, Ardekani Babak A., Ashburner John, Avants Brian, Chiang Ming-Chang, Christensen Gary E., Collins D. Louis, Gee James, Hellier Pierre, Song Joo Hyun, Jenkinson Mark, Lepage Claude, Rueckert Daniel, Thompson Paul, Vercauteren Tom, Woods Roger P., Mann J. John, Parsey Ramin V., 2009. Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *Neuroimage* 46 (3), 786–802. 10.1016/j.neuroimage.2008.12.037. [PubMed: 19195496]
- Klein Stefan, Staring Marius, Murphy Keelin, Viergever Max A., Pluim Josien P.W., 2010. Elastix: A toolbox for intensity-based medical image registration. *IEEE Trans. Med. Imaging* 29, 196–205, <https://api.semanticscholar.org/CorpusID:15850341>. [PubMed: 19923044]
- Krizhevsky Alex, Sutskever Ilya, Hinton Geoffrey E., 2012. ImageNet classification with deep convolutional neural networks. In: Pereira F, Burges CJ, Bottou L, Weinberger KQ (Eds.), In: *Advances in Neural Information Processing Systems*, vol. 25, Curran Associates, Inc., Arizona, https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf.
- Lein Ed S., Hawrylycz Michael J., Ao Nancy, Ayres Mikael, Bensinger Amy, Bernard Amy, Boe Andrew F., Boguski Mark S., Brockway Kevin S., Byrnes Emi J., Chen Lin, Chen Li, Chen Tsuey-Ming, Chin Chi, Mei Chong, Jimmy Crook, Brian E, Czaplinska Aneta, Dang Chinh N., Datta Suvro, Dee Nick R., Desaki Aimee L., Desta Tsega, Diep Ellen, Dolbeare Tim A., Donelan Matthew J., Dong Hong-Wei, Dougherty Jennifer G., Duncan Ben J., Ebbert Amanda J., Eichele Gregor, Estin Lili K., Faber Casey, Facer Benjamin A., Fields Rick, Fischer Shanna R., Fliss Tim P., Frensley Cliff, Gates Sabrina N., Glattfelder Katie J., Halverson Kevin R., Hart Matthew R., Hohmann John G., Howell Maureen P., Jeung Darren P., Johnson Rebecca A., Karr Patrick T.,

Kawal Reena, Kidney Jolene M., Knapik Rachel H., Kuan Chihchau L., Lake James H., Laramée Annabel R., Larsen Kirk D., Lau Christopher, Lemon Tracy A., Liang Agnes J., Liu Ying, Luong Lon T., Michaels Jesse, Morgan Judith J., Morgan Rebecca J., Mortrud Marty T., Mosqueda Nerick F., Ng Lydia L., Ng Randy, Orta Geralyn J., Overly Caroline C., Pak Tu H., Parry Sheana E., Pathak Sayan D., Pearson Owen C., Puchalski Ralph B., Riley Zackery L., Rockett Hannah R., Rowland Stephen A., Royall Joshua J., Ruiz Marcos J., Sarno Nadia R., Schaffnit Katherine, Shapovalova Nadiya V., Sivasay Taz, Slaughterbeck Clifford R., Smith Simon C., Smith Kimberly A., Smith Bryan I., Sodt Andy J., Stewart Nick N., Stumpf Kenda-Ruth, Sunkin Susan M., Sutram Madhavi, Tam Angelene, Teemer Carey D., Thaller Christina, Thompson Carol L., Varnam Lee R., Visel Axel, Whitlock Ray M., Wahnoutka Paul E., Wolkey Crissa K., Wong Victoria Y., Wood Matthew, Yaylaoglu Murat B., Young Rob C., Youngstrom Brian L., Yuan Feng, Xu Zhang, Bin Zwingman, Theresa A, Jones Allan R., 2007. Genome-wide atlas of gene expression in the adult mouse brain. *Nature* 445 (7124), 168–176. 10.1038/nature05453. [PubMed: 17151600]

Liebmann Thomas, Renier Nicolas, Bettayeb Karima, Greengard Paul, Tessier-Lavigne Marc, Flajolet Marc, 2016. Three-dimensional study of Alzheimer's disease hallmarks using the iDISCO clearing method. *Cell Rep.* 16 (4), 1138–1152. 10.1016/j.celrep.2016.06.060. [PubMed: 27425620]

Maes Frederik, Collignon Andre, Vandermeulen Dirk, Marchal Guy, Suetens Paul, 1996. Multi-modality image registration maximization of mutual information. In: *Proceedings of the 1996 Workshop on Mathematical Methods in Biomedical Image Analysis. MMBIA'96*, IEEE Computer Society, USA, p. 14.

Maintz J.B. Antoine, Viergever Max A., 1998. A survey of medical image registration. *Med. Image Anal.* 2 (1), 1–36. 10.1016/S1361-8415(01)80026-8. [PubMed: 10638851]

Myronenko Andriy, Song Xubo, 2010. Point set registration: Coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.* 32 (12), 2262–2275. 10.1109/TPAMI.2010.46. [PubMed: 20975122]

Oh Seung Wook, Harris Julie A., Ng Lydia, Winslow Brent, Cain Nicholas, Mihalas Stefan, Wang Quanxin, Lau Chris, Kuan Leonard, Henry Alex M., Mortrud Marty T., Ouellette Benjamin, Nguyen Thuc Nghi, Sorensen Staci A., Slaughterbeck Clifford R., Wakeman Wayne, Li Yang, Feng David, Ho Anh, Nicholas Eric, Hirokawa Karla E., Bohn Phillip, Joines Kevin M., Peng Hanchuan, Hawrylycz Michael J., Phillips John W., Hohmann John G., Wahnoutka Paul, Gerfen Charles R., Koch Christof, Bernard Amy, Dang Chinh, Jones Allan R., Zeng Hongkui, 2014. A mesoscale connectome of the mouse brain. *Nature* 508 (7495), 207–214. 10.1038/nature13186. [PubMed: 24695228]

Ourselin Sébastien, Roche Alexis, Prima Sylvain, Ayache Nicholas, 2000. Block matching: A general framework to improve robustness of rigid registration of medical images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. <https://api.semanticscholar.org/CorpusID:21483245>.

Pal Soumyadeep, Tennant Matthew, Ray Nilanjan, 2022. Towards Positive Jacobian: Learn to Postprocess for Diffeomorphic Image Registration with Matrix Exponential. In: *2022 26th International Conference on Pattern Recognition. ICPR*, IEEE Computer Society, Los Alamitos, CA, USA, pp. 3391–3398. 10.1109/ICPR56361.2022.9956119.

Qiu Zhiyu, Tang Huihui, Tian Dongsheng, 2009. Non-rigid medical image registration based on the thin-plate spline algorithm. In: *CSIE '09: Proceedings of the 2009 WRI World Congress on Computer Science and Information Engineering*, Vol. 2. 1730 Massachusetts Ave., NW Washington, DC United States, pp. 522–527. 10.1109/CSIE.2009.825,

Qu Lei, Li Yuanyuan, Xie Peng, Liu Lijuan, Wang Yimin, Wu Jun, Liu Yu, Wang Tao, Li Longfei, Guo Kaixuan, Wan Wan, Ouyang Lei, Xiong Feng, Kolstad Anna C., Wu Zhuohao, Xu Fang, Zheng Yefeng, Gong Hui, Luo Qingming, Bi Guoqiang, Dong Hongwei, Hawrylycz Michael, Zeng Hongkui, Peng Hanchuan, 2022. Cross-modal coherent registration of whole mouse brains. *Nature Methods* 19 (1), 111–118. 10.1038/s41592-021-01334-w. [PubMed: 34887551]

Ragan Timothy, Kadir Lolahon R, Venkataraju Kannan Umadevi, Bahlmann Karsten, Sutin Jason, Taranda Julian, Arganda-Carreras Ignacio, Kim Yongsoo, Seung H Sebastian, Osten Pavel, 2012. Serial two-photon tomography for automated ex vivo mouse brain imaging. *Nat. Methods* 9 (3), 255–258. [PubMed: 22245809]

- Reuter Martin, Rosas H. Diana, Fischl Bruce, 2010. Highly accurate inverse consistent registration: A robust approach. *Neuroimage* 53 (4), 1181–1196. 10.1016/j.neuroimage.2010.07.020. [PubMed: 20637289]
- Rueckert D, Sonoda LI, Hayes C, Hill DLG, Leach MO, Hawkes DJ, 1999. Non-rigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imaging* 18 (8), 712–721. 10.1109/42.796284. [PubMed: 10534053]
- Rusinkiewicz S, Levoy M, 2001. Efficient variants of the ICP algorithm. In: *Proceedings Third International Conference on 3-D Digital Imaging and Modeling*. IEEE, Quebec City, QC, Canada, pp. 145–152. 10.1109/IM.2001.924423.
- Shen Dinggang, Davatzikos Christos, 2002. HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imaging* 21, 1421–1439, <https://api.semanticscholar.org/CorpusID:5586563>. [PubMed: 12575879]
- Song Guoli, Han Jianda, Zhao Yiwen, Wang Zheng, Du Huibin, 2017. A review on medical image registration as an optimization problem. *Curr. Med. Imaging Rev.* 13, 10.2174/1573405612666160920123955.
- Sun Wenhong, Zhou Weidong, Yang Mingqiang, 2011. Medical image registration using thin-plate spline for automatically detecting and matching of point sets. In: *2011 5th International Conference on Bioinformatics and Biomedical Engineering*. IEEE, Wuhan, China, pp. 1–4. 10.1109/icbbe.2011.5780355.
- Trappolini Giovanni, Cosmo Luca, Moschella Luca, Marin Riccardo, Melzi Simone, Rodolà Emanuele, 2021. Shape registration in the time of transformers. *arXiv:2106.13679*.
- Tustison Nicholas, Avants Brian, Gee James, 2009. Directly manipulated free-form deformation image registration. *IEEE Trans. Image Process.* : Publ. IEEE Signal Process. Soc. 18, 624–635. 10.1109/TIP.2008.2010072. [PubMed: 19171516]
- Tward Daniel, Brown Timothy, Kageyama Yusuke, Patel Jaymin, Hou Zhipeng, Mori Susumu, Albert Marilyn, Troncoso Juan, Miller Michael, 2020. Diffeomorphic registration with intensity transformation and missing data: Application to 3D digital pathology of Alzheimer's disease. *Front. Neurosci.* 14, 10.3389/fnins.2020.00052.
- Vaswani Ashish, Shazeer Noam, Parmar Niki, Uszkoreit Jakob, Jones Llion, Gomez Aidan N, Kaiser Łukasz, Polosukhin Illia, 2017. Attention is all you need. In: Guyon I, Luxburg U. Von, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (Eds.), In: *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- Vercauteren Tom, Pennec Xavier, Perchant Aymeric, Ayache Nicholas, 2009. Diffeomorphic demons: Efficient non-parametric image registration. *Neuroimage* 45, S61–S72, <https://api.semanticscholar.org/CorpusID:220992908>. [PubMed: 19041946]
- Wachinger Christian, Navab Nassir, 2010. Manifold learning for multi-modal image registration. In: *BMVC*. pp. 1–12.
- Wang Quanxin, Ding Song-Lin, Li Yang, Royall Josh, Feng David, Lesnar Phil, Graddis Nile, Naeemi Maitham, Facer Benjamin, Ho Anh, Dolbeare Tim, Blanchard Brandon, Dee Nick, Wakeman Wayne, Hirokawa Karla E., Szafer Aaron, Sunkin Susan M., Oh Seung Wook, Bernard Amy, Phillips John W., Hawrylycz Michael, Koch Christof, Zeng Hongkui, Harris Julie A., Ng Lydia, 2020. The allen mouse brain common coordinate framework: A 3D reference atlas. *Cell* 181 (4), 936–953.e20. 10.1016/j.cell.2020.04.007. [PubMed: 32386544]
- Yeo B.T. Thomas, Sabuncu Mert, Vercauteren Tom, Ayache Nicholas, Fischl Bruce, Golland Polina, 2009. Spherical demons: Fast diffeomorphic landmark-free surface registration. *IEEE Trans. Med. Imaging* 29, 650–668. 10.1109/TMI.2009.2030797. [PubMed: 19709963]
- Zhang Zhengyou, 1994. Iterative point matching for registration of free-form curves and surfaces. *Int. J. Comput. Vis.* 13, 119–152, <https://api.semanticscholar.org/CorpusID:14673939>.
- Zitová Barbara, Flusser Jan, 2003. Image registration methods: a survey. *Image Vis. Comput.* 21 (11), 977–1000. 10.1016/S0262-8856(03)00137-9.




Fig. 1. Pipeline for affine registration:

Starting with our input data, we find features to compute the point clouds of I_m and I_f . We perform an initial rigid alignment on these point clouds, which consists of rotation, scaling, and translation, to align the point clouds together coarsely. We then compute matching correspondences between features in both point clouds to refine the transformation using ICP, computing a final affine transformation matrix.




Fig. 2. Example images showing before and after rotation correction:

A pair of coronal sections are shown in each row. The left side denotes the section before rotation correction, and the right side denotes the section after rotation correction and reslicing. In the top row, the approximated plane for the moving image is shown in red, while the approximated template plane is shown in green. Boxes denoting points of interest are marked where the alignment has improved. After re-slicing, the brain is less tilted to the right and more symmetric. Two additional points of interest are marked on the bottom row, indicating a set of lobes that now appear in the same section. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 3. Active pixels in example I_m sections:

On the left, active pixels in a coronal section are computed by applying a simple threshold. A lot of salt and pepper noise can be observed, and applying a min-max algorithm to estimate a bounding box will require filtering this noise. A sample noisy point is marked. Instead of min-max, we use the 3D volume stack to estimate the sagittal plane's active pixel count. The middle section shows a view of the sagittal axis at that specific noisy point. We can observe that the section corresponding to the noisy plane outputs a very low active pixel count, most of which is primarily from noise. On the right, a sagittal view section where we expect a good active pixel count is shown. All sections below this will have close to zero active pixels, as indicated in Fig. 4, which can be used to estimate the sagittal plane of the bounding box. This same process can be repeated in all axes.

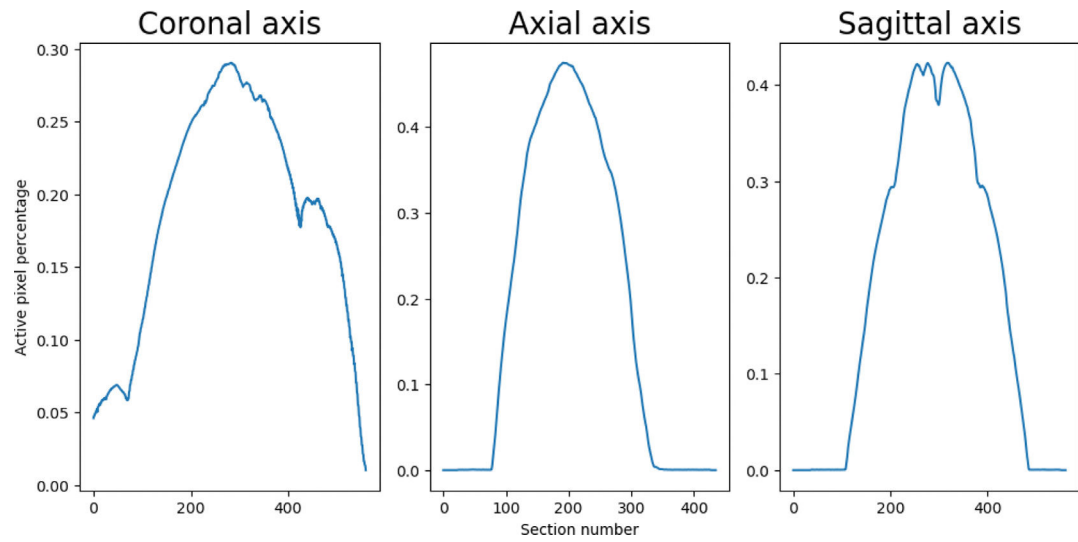


Fig. 4. Active pixels per section of I_m across the different axes:

Depicted is the percentage of active pixels, or pixels that represent the actual section information, out of all pixels for each section in I_m . The axes are defined as orthogonal to the planes in radiological convention.




Fig. 5. Example images across the coronal plane illustrating contour features:

Representative slices are shown with contours denoting the outer boundaries (green) and the internal features (red). These features are used later to establish correspondences for registration, as seen in Fig. 12. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

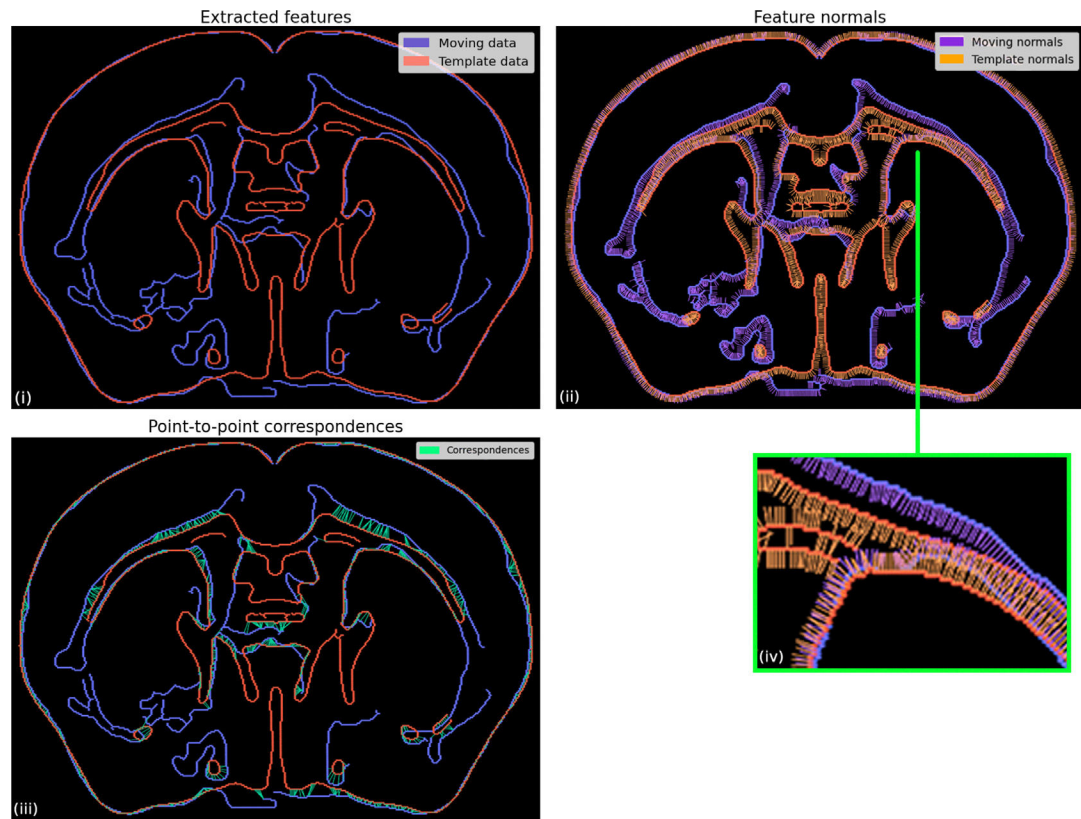


Fig. 6. Correspondence finding between representative mapped sections:

(i) depicts features from I_m (blue) and I_f (red). (ii) depicts the extracted normals from I_m (purple) and I_f (orange). (iii) depicts the established correspondences (green). (iv) depicts a zoomed-in view of a region in the extracted normals. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 7. Point cloud representation:

The point clouds I_m (blue) and I_f (red) are depicted (i) before any registration operations and (ii) after affine registration. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 8. RegBoost postprocessing:

We present the pipeline for the RegBoost postprocessing steps. After performing section matching, 3D correspondences from our data are matched together from section to section. We then perform Laplacian interpolation using a linear solver to compute a deformation map for registration.




Fig. 9. 2D Laplacian registration compared with 3D Laplacian registration across the axial plane:

The outputs of the 2D Laplacian (left column) and the 3D Laplacian registrations (right column) are depicted here. Some points of interest have been marked to highlight areas of note. In the 2D Laplacian output, discontinuities are apparent along the edges. The 3D Laplacian output is visually smoother without discontinuities.




Fig. 10. 2D Laplacian registration compared with 3D Laplacian registration across the coronal plane:

The outputs of the 2D Laplacian (left column) and the 3D Laplacian registrations (right column) are depicted here. Some points of interest have been marked to highlight areas of note. In the 2D Laplacian output, the deformations between the adjacent slices around the circled areas are discontinuous. The 3D Laplacian output maintains continuity across the sections.




Fig. 11. LaplacianReg:

We introduce LaplacianReg, a registration framework that uses 3D correspondences and the 3D Laplacian operator to align volumes together. The input data is first transformed according to the RegBoost preprocessing steps in Section 4.1. Correspondences are established between external boundary and internal features, then used in a 3D Laplacian operation to result in a final nonlinear registration.




Fig. 12. LaplacianReg, coronal comparison between registration using external boundaries and external boundaries with internal features:

In the left column, representative coronal slices are depicted, where LaplacianReg has been run using only the outer boundary features. In the right column, LaplacianReg has been run using the outer and internal boundary features. Some points of interest have been marked to highlight areas of note. Using external boundaries with internal features results in significantly better internal alignment, as demonstrated by the regions circled in green. The regions in the right column are aligned closer to the template structures outlined in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 13. LaplacianReg, axial comparison between registration using external boundaries and external boundaries with internal features:

In the left column, representative axial slices are depicted where LaplacianReg has been run using only the outer boundary features. In the right column, LaplacianReg has been run using the outer and internal boundary features. Some points of interest have been marked to highlight areas of note. Using external boundaries with internal features results in significantly better internal alignment, as demonstrated by the regions circled in green. The regions in the right column are aligned closer to the template structures outlined in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 14. 3D visualization of correspondences:

(i) external features of I_f (red) and their correspondences in I_m (green). (ii) internal features for both I_m (blue) and I_f (red). (iii) internal features of I_f and their correspondences in I_m (green). These internal features are noisy, but some smooth correspondences can be found where structures are neatly defined, such as around the center. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 15. 3D rendering of the final registered output:

- (i) A cross-section of the registered mouse brain volume indicating its segmented anatomy.
- (ii) The registered mouse brain overlays its corresponding annotated volume.




Fig. 16. Final results of registration and alignment over the coronal plane using Elastix:

The four columns depict registration using Elastix, Elastix + RegBoost, Elastix + RegBoost without the annotation template overlay, and the average template, respectively, for a set of example images. The baseline Elastix here is compared to highlight the noteworthy changes with our method. The images representing registration with 3D Laplacian refinement are more closely aligned to the fixed template image features (shown in red). Some points of interest have been marked to highlight areas of interest. The exterior and interior boundaries align significantly better than the baseline Elastix method. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

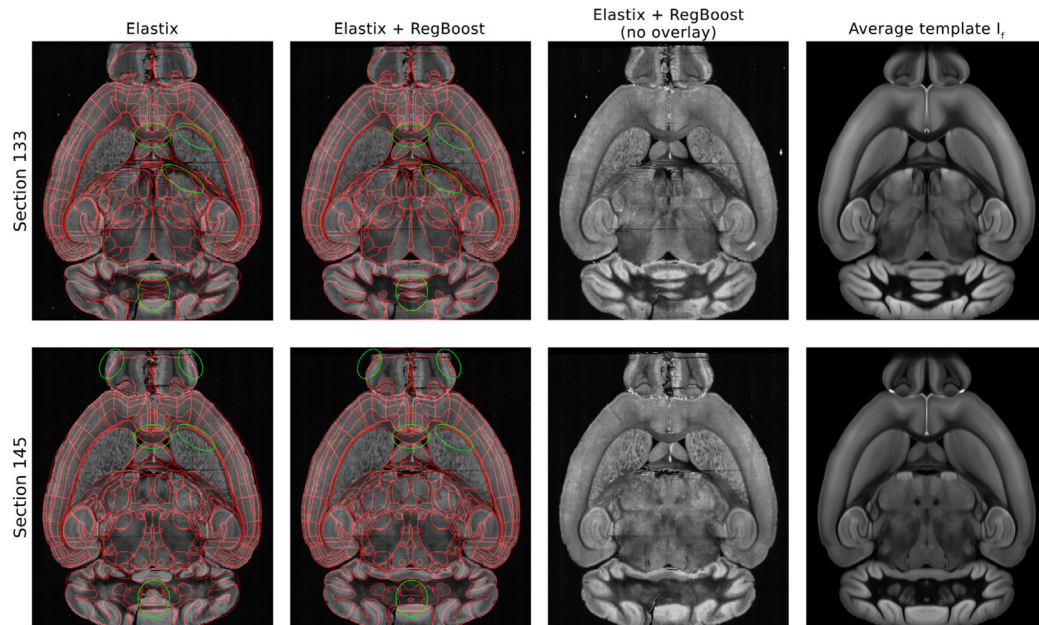


Fig. 17. Final results of registration and alignment over the axial plane using Elastix:

The four columns depict registration using Elastix, Elastix + RegBoost, Elastix + RegBoost without the annotation template overlay, and the average template, respectively, for a set of example images. The baseline Elastix here is compared to highlight the noteworthy changes with our method. The images representing registration with 3D Laplacian refinement are more closely aligned to the fixed template image features (shown in red). Some points of interest have been marked to highlight areas of interest. The internal regions, particularly the ventricles, align significantly better than the baseline Elastix method. Additionally, the tissues in the bottom-most circles for our method are present in section 133 and not in section 145, which are accurate to the average template I_f , where Elastix is not. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

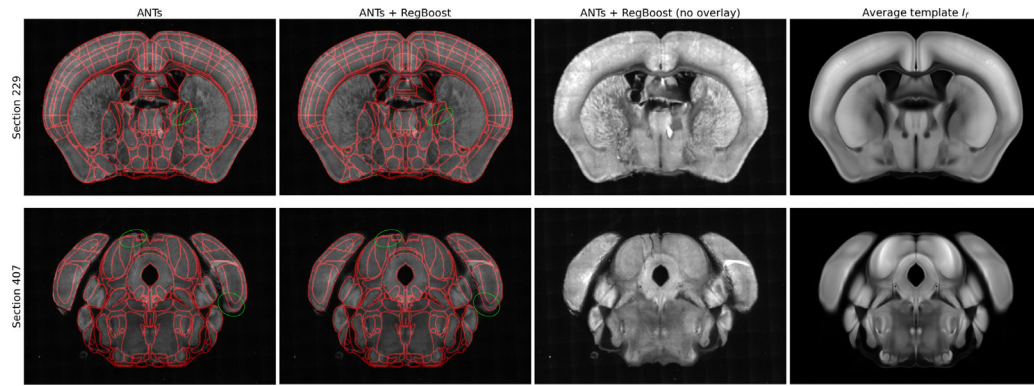


Fig. 18. Example coronal images depicting the final registration results using ANTs: ANTs already performs well in registering the input image. However, some areas in our implementation are visually better than the base ANTs result, particularly in the circled green regions. Our method can perform curve matching for these areas to align these features better. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

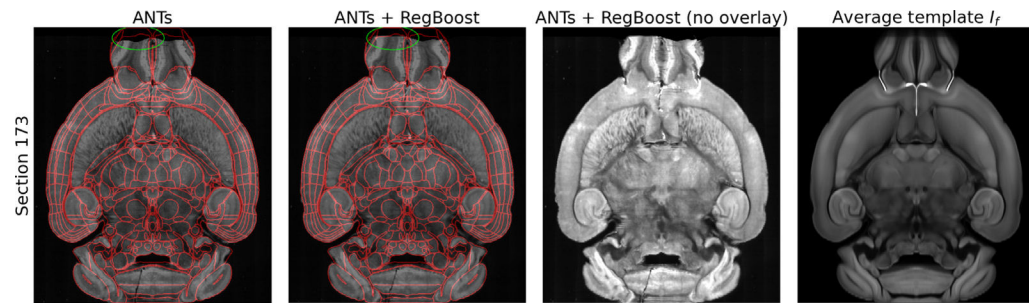


Fig. 19. Example axial images depicting the final registration results using ANTs:

ANTs already performs well in registering the input image. However, some areas in our implementation are visually better than the base ANTs result, particularly in the circled green regions. Our method can perform curve matching for these areas to align these features better. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 20. Jacobian determinants computed over example axial images:

Computing the Jacobian determinant over all the voxels, results in a negative factor in approximately 0.6% of the voxels. The affected voxels are highlighted in white in these three sections. These negative determinant pixels are mostly found at the boundaries where we have our correspondences, which indicates that the volume in between the correspondences is interpolated smoothly.

Table 1 Rotation correction results:

The computed advanced Mattes mutual Information metrics for each method with and without the rotation correction operation as defined in Section 4.1.1. Lower scores indicate better performance. No intensity scaling is used for these results. The results indicate that rotation correction in the initial rigid alignment improves mutual information scores across all methods.

Method	No rotation correction	Rotation correction
RegBoost + keyframe matching	-0.64	-0.66
Basic Elastix	-0.67	-0.69
LaplacianReg	-0.66	-0.72

Table 2 Quantitative results:

The final results are shown with each method's advanced Mattes mutual information scores, where lower scores indicate better outcomes. Each Laplacian operation utilizes our improvements (rotation correction, affine transformations, and intensity rescaling, if possible) and is compared with Elastix and ANTs out-of-the-box. LaplacianReg uses axis alignment, ICP affine transformations, and external 3D correspondences with the 3D Laplacian operator. RegBoost uses axis alignment, affine registration, and slice-to-slice 2D correspondences with the 3D Laplacian operator. The results indicate that our methods utilizing Laplacian operations outperform Elastix and ANTs, with the best-performing method being section-to-section correspondences with the 3D Laplacian operator.

Method	Adv. Mattes mutual info.
RegBoost + keyframe matching	-0.66
Basic Elastix	-0.67
LaplacianReg	-0.72
LaplacianReg + RegBoost	-0.74
Elastix + RegBoost	-0.82
ANTs SyN	-0.89
ANTs SyN + RegBoost	-0.92