SCALE: Structure-preserving Cycle-consistent Atlas Learning for 4D Infant Cerebellum Atlas Construction

Qingrui Sha[1,2], Jiameng Liu[1], Dan Hu[1], Jiale Cheng[1,2], Yue Sun[1,2], Zhengwang Wu[1], Weili Lin[1], Li Wang[1], and Gang Li[1(✉)]

[1] Department of Radiology and Biomedical Research Imaging Center, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA
gang_li@med.unc.edu

[2] Lampe Joint Department of Biomedical Engineering, University of North Carolina at Chapel Hill and North Carolina State University, Chapel Hill, NC, USA

Abstract. Constructing a four-dimensional (4D) infant cerebellum atlas is critical for modeling spatiotemporal brain development and identifying early cerebellar abnormalities. However, this task is particularly challenging due to extremely complex folia folding patterns, dynamic growth, rapid myelination, and pronounced variations in MR image appearance across developmental stages. These factors compromise population-level anatomical correspondences and introduce longitudinal inconsistencies, causing existing atlas construction methods to produce blurred tissue boundaries and fail to preserve fine-grained cerebellar structures. To address these challenges, we propose a structure-preserving and cycle-consistent atlas learning (SCALE) framework for 4D infant cerebellum atlas construction. SCALE jointly optimizes atlas generation and deformable registration through a dual-network architecture. Specifically, it incorporates a Hessian-based enhanced structural representation that effectively captures the accordion-like laminar sheets of cerebellar folia, enabling more accurate alignment. To promote longitudinal consistency, SCALE introduces a cross-temporal cycle-consistency constraint that forms closed deformation loops. This constraint enforces agreement between direct within-subject longitudinal transformation paths and atlas-mediated transformation paths, thereby regularizing atlas evolution to be consistent with individual growth trajectories. Extensive experiments on the UNC/UMN Baby Connectome Project dataset demonstrate that SCALE generates a high-fidelity 4D infant cerebellum atlas with sharper structural details, more stable population-level alignment, and greater longitudinal consistency across development. These improvements may facilitate early detection of atypical developmental deviations of the cerebellum.

Keywords: 4D infant cerebellum atlas · Longitudinal registration.

1 Introduction

The cerebellum, also known as the "little brain" , plays a crucial role in sensorimotor coordination, cognitive development, and the emergence of higher-order functions [16,26]. Accurately modeling its spatiotemporal developmental trajectories during infancy is essential for understanding early brain organization and identifying deviations associated with neurodevelopmental disorders [20,9]. A four-dimensional (4D) infant cerebellum atlas provides a powerful representation for capturing population-level commonality, variability, and longitudinal changes [24]. However, the cerebellum exhibits extremely fine-grained and highly folded folia structures, whose rapid growth, dynamic myelination, and substantial variations in MR imaging appearance make the estimation of anatomical correspondences and the maintenance of longitudinal consistency particularly challenging [11,8,18].

To address the structural complexity and appearance variations, many studies have explored the use of auxiliary anatomical priors. For example, tissue segmentation maps have been incorporated to guide atlas building [19,22]. However, such priors alone are insufficient for the cerebellum, as tissue segmentation maps provide only limited representation of the fine-grained geometry of cerebellar folia, which form thin, accordion-like laminar sheets with complex folding and branching patterns that vary substantially across subjects [23]. Given the uniquely folded structure of the cerebellum, a Hessian-based enhanced structural representation, which inherently captures the accordion-like laminar sheets of cerebellar folia, offers a more direct descriptor of cerebellar geometry. Instead of being treated as independent output channels during atlas construction, Hessian-based features should be derived from the atlas intensity itself to provide meaningful geometric supervision. Beyond structural priors, longitudinal consistency has not been adequately explored. Existing efforts have largely been limited to cross-temporal atlas alignment [8] and do not exploit the fact that within-subject longitudinal registration is inherently more accurate and reliable than atlas-to-subject alignment, particularly given the highly folded and subject-specific patterns of cerebellar folia. To this end, we leverage direct within-subject longitudinal transformation to supervise atlas-mediated transformation paths, thereby encouraging atlas evolution to remain consistent with individual developmental trajectories.

Overall, we propose a structure-preserving and cycle-consistent atlas learning (SCALE) framework for constructing a longitudinally consistent infant cerebellum atlas. The main contributions of this work are as follows:

1. We integrate a Hessian-based enhanced structural representation derived directly from the atlas intensity volume to preserve cerebellar folia geometry, enabling robust population-level alignment despite the highly intricate and rapidly evolving folia patterns in infancy.
2. We introduce a cross-temporal cycle-consistency constraint that constructs closed deformation loops. By enforcing agreement between within-subject longitudinal registration and atlas-mediated transformation paths, this con-

straint effectively guides and regularizes atlas evolution, yielding a sharper, longitudinally more consistent, and anatomically more meaningful atlas.
3. Extensive experiments on the UNC/UMN Baby Connectome Project dataset demonstrate that SCALE produces a high-fidelity 4D infant cerebellar atlas with sharper structures, clearer folia delineation, and more stable population-level alignment and longitudinal consistency across development.

2 Method

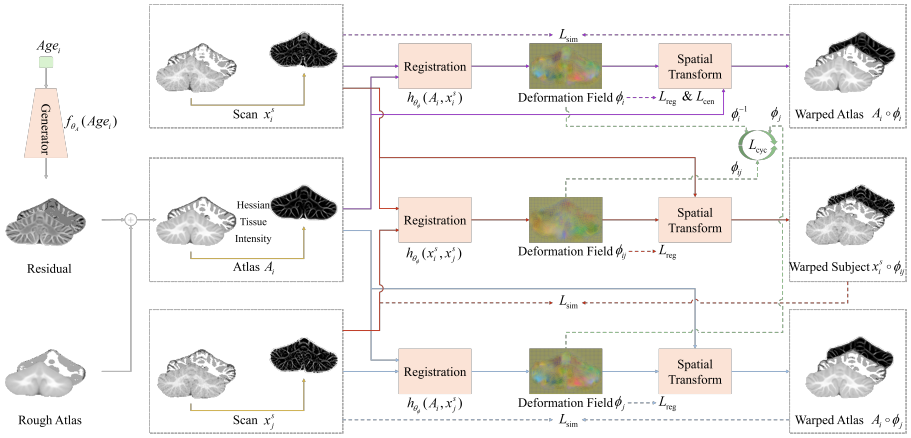

Fig. 1. Overview of the proposed SCALE framework. For a subject s, $\boldsymbol{x}_i^s$ and $\boldsymbol{x}_j^s$ denote scans acquired at two different time points. The atlas-generation network predicts an Age_i-conditioned residual to refine a rough atlas, and a Hessian-based map is derived from the generated intensity for structural guidance. In the cycle-registration network, the purple, blue, and red paths represent atlas-to-$\boldsymbol{x}_i^s$, atlas-to-$\boldsymbol{x}_j^s$, and $\boldsymbol{x}_i^s$-to-$\boldsymbol{x}_j^s$ registrations, respectively, while the green path forms a cycle that enforces cross-temporal cycle-consistency.

2.1 Network Overview

Fig. 1 illustrates the architecture of our structure-preserving and cycle-consistent atlas learning (SCALE) framework. SCALE utilizes a dual-network architecture consisting of an atlas generation network (AG-Net) and a cycle-registration network (CR-Net), which are jointly optimized. The model takes as inputs $\boldsymbol{x}_i^s = (\boldsymbol{I}_i^s, \boldsymbol{S}_i^s, \boldsymbol{H}_i^s)$ together with the associated age Age_i, and optionally a paired longitudinal scan $\boldsymbol{x}_j^s$ at time j. Here, $\boldsymbol{I}_i^s, \boldsymbol{S}_i^s$, and $\boldsymbol{H}_i^s$ denote the intensity volume, tissue segmentation map, and Hessian eigenvalue map, respectively. SCALE consists of an atlas-generation module $\boldsymbol{A}_i = f_{\boldsymbol{\theta_A}}(Age_i)$ and a

registration module $h_{\boldsymbol{\theta}_{\boldsymbol{\phi}}}(\cdot,\cdot)$. The atlas generator synthesizes the age-specific atlas $\boldsymbol{A}_i = (\boldsymbol{I}_{A_i}, \boldsymbol{S}_{A_i}, \boldsymbol{H}_{A_i})$, while the Hessian eigenvalue map $\boldsymbol{H}_{A_i}$ is computed from the intensity volume $\boldsymbol{I}_{A_i}$. The registration module is a single UNet-like network invoked three times with shared parameters to estimate the deformation fields required for atlas-to-subject and within-subject longitudinal registration. For example, $h_{\boldsymbol{\theta}_{\boldsymbol{\phi}}}(\boldsymbol{A}_i, \boldsymbol{x}_i^s)$ predicts the stationary velocity field $\boldsymbol{\phi}_i$ that warps the atlas to the subject scan $\boldsymbol{x}_i^s$. The remaining registration steps used for cycle construction are described in the following subsections.

2.2 Atlas-generation Network (AG-Net)

AG-Net follows the conditional atlas-generation design used in AtlasMorph [19], where the atlas is synthesized from an age-conditioned latent representation. Given the input age Age_i, the network generates multi-channel atlas outputs, including the intensity volume and the tissue segmentation map.

To further preserve fine-scale cerebellar folia geometry, AG-Net additionally incorporates a Hessian-based enhanced structural representation that is analytically derived from the generated intensity volume. Specifically, the voxel-wise Hessian matrix is computed from the second-order spatial derivatives of the Gaussian-smoothed intensity volume at multiple scales. Following eigen-decomposition, at each voxel and scale, we select the eigenvalue with the largest absolute magnitude while retaining its sign, thereby capturing strong second-order responses associated with thin and highly folded cerebellar structures. The responses across scales are subsequently fused by selecting, at each voxel, the signed response with the largest absolute magnitude, yielding the final Hessian representation $\boldsymbol{H}_{A_i}$ [12]. This Hessian-based representation provides robust geometric cues for downstream losses, enabling sharper folia delineation and improved anatomical consistency between the atlas and the subject.

2.3 Cycle-registration Network (CR-Net)

The registration component, CR-Net, is a UNet-based convolutional network following the VoxelMorph architecture [7]. Given a pair of 3D volumes, CR-Net predicts a stationary velocity field, which is integrated via the scaling-and-squaring method [2] to obtain forward and inverse deformation fields.

As shown in Fig. 1, CR-Net is invoked three times with shared parameters $\boldsymbol{\theta}_{\boldsymbol{\phi}}$ to estimate the deformation fields required for cycle-consistent learning. The purple and blue paths compute the atlas-to-subject fields $\boldsymbol{\phi}_i$ and $\boldsymbol{\phi}_j$, mapping $\boldsymbol{A}_i$ to $\boldsymbol{x}_i^s$ and $\boldsymbol{x}_j^s$, respectively. The red path yields the within-subject deformation field $\boldsymbol{\phi}_{ij}$, directly registering the two time points of the same subject. These fields are coupled through the green path to form a closed deformation loop. When longitudinal pairs are unavailable, only the purple path is used.

Due to subject-specific cerebellar folding, direct registration between two time points of the same subject is typically more reliable than aligning the atlas to a subject. We therefore use $\boldsymbol{\phi}_{ij}$ as a supervisory signal to regularize the atlas-mediated fields $\boldsymbol{\phi}_i$ and $\boldsymbol{\phi}_j$, enforcing consistency between direct longitudinal

transformations and atlas-based deformation paths to improve atlas sharpness. By jointly aligning the atlas to longitudinal scans at different ages, this constraint guides and stabilizes atlas evolution across developmental stages and yields a more longitudinally consistent atlas.

2.4 Loss Functions

We jointly optimize atlas-to-subject and within-subject longitudinal registrations by maximizing multi-channel alignment while enforcing smooth and unbiased deformations. Specifically, we encourage agreement between each subject and the warped atlas across the intensity volume, tissue segmentation map, and Hessian eigenvalue map, while simultaneously maximizing alignment between subject pairs to provide more reliable cross-temporal correspondences that guide atlas-mediated transformations.

The overall loss is defined as

$$L = L_{\text{sim}} + \lambda_1 L_{\text{reg}} + \lambda_2 L_{\text{cen}} + \lambda_3 L_{\text{cyc}}. \tag{1}$$

The similarity term measures multi-channel agreement for both atlas-to-subject and within-subject registrations. Since the inputs are multi-channel, we adopt channel-specific similarity metrics: normalized cross-correlation [3] for intensity volume, Dice loss [5] for tissue segmentation map, and mean squared error [17] for Hessian eigenvalue map. For example, the similarity between the atlas $\boldsymbol{A}_i$ and subject scan $\boldsymbol{x}_i^s$ is $L_{\text{sim}} = -L_{\text{lncc}}(\boldsymbol{I}_{\boldsymbol{A}_i} \circ \boldsymbol{\phi}_i, \boldsymbol{I}_i^s) - \gamma_1 L_{\text{dice}}(\boldsymbol{S}_{\boldsymbol{A}_i} \circ \boldsymbol{\phi}_i, \boldsymbol{S}_i^s) + \gamma_2 L_{\text{mse}}(\boldsymbol{H}_{\boldsymbol{A}_i} \circ \boldsymbol{\phi}_i, \boldsymbol{H}_i^s)$, where $\boldsymbol{I}$, $\boldsymbol{S}$, and $\boldsymbol{H}$ denote the intensity volume, tissue segmentation map, and Hessian eigenvalue map, respectively, and $\boldsymbol{\phi}_i$ warps the atlas into the space of subject $\boldsymbol{x}_i^s$. The full similarity loss sums analogous terms over all atlas-to-subject and within-subject registrations.

To ensure physically plausible transformations, we penalize spatial gradients of the displacement fields [6,10], $L_{reg} = ||\nabla \boldsymbol{u}_i||_2^2 + ||\nabla \boldsymbol{u}_j||_2^2 + ||\nabla \boldsymbol{u}_{ij}||_2^2$, where $\boldsymbol{u} = \boldsymbol{\phi} - \text{Id}$ is the displacement field and Id denotes the identity transform.

To maintain an unbiased atlas located near the geometric center of the population, we use the scaling-and-squaring method [2] to obtain inverse deformation fields in the atlas space and encourage the average inverse displacement to be close to zero [19]: $L_{\text{cen}} = ||\frac{1}{K}\sum_{k=1}^{K} \boldsymbol{u}_{i,k}^{-1}||_2^2$, where $\boldsymbol{u}_i^{-1} = \boldsymbol{\phi}_i^{-1} - \text{Id}$. The K deformation fields correspond to subjects from adjacent developmental stages, preventing the atlas from drifting toward any single subject.

To enforce agreement between direct within-subject registration and atlas-mediated deformation paths, we introduce a cycle-consistency loss:

$$L_{\text{cyc}} = L_{\text{mse}}(\boldsymbol{\phi}_{ij}, \boldsymbol{\phi}_i^{-1} \circ \boldsymbol{\phi}_j), \tag{2}$$

where $\boldsymbol{\phi}_{ij}$ denotes the direct deformation from subject scan $\boldsymbol{x}_i^s$ to $\boldsymbol{x}_j^s$, while $\boldsymbol{\phi}_i^{-1} \circ \boldsymbol{\phi}_j$ represents the composite transformation through the atlas. Leveraging reliable within-subject longitudinal registration, this constraint promotes consistent mapping and regularizes atlas evolution, yielding sharper and more longitudinally consistent atlases.

3 Experiments

3.1 Dataset and Preprocessing

Our 4D infant cerebellum atlas is constructed from MRI scans provided by the publicly available UNC/UMN Baby Connectome Project (BCP) [14]. The dataset contains 519 T1-weighted images from 255 infants, including 145 infants (56.9%) with longitudinal follow-up scans. All scans were acquired at an isotropic resolution of $0.8 \times 0.8 \times 0.8\,\mathrm{mm}^3$ and preprocessed using iBEAT V2.0 [25,21]. Each volume was then affinely aligned to the averaged 24-month image space using ANTs [4]. A rough initial atlas was obtained by voxel-wise averaging across all training scans. For quantitative evaluation, the dataset was split into training and testing sets in a 9:1 ratio at the subject level, with no infant appearing in both sets to prevent longitudinal data leakage.

3.2 Experimental Settings

Training Details. We optimized the model using the Adam [15] optimizer with a learning rate of 0.0001 and trained it for 1,000 epochs until convergence. Loss-weight hyperparameters were selected via grid search, yielding $\lambda_1 = 1$, $\lambda_2 = 0.1$, $\lambda_3 = 0.1$, $\gamma_1 = 0.5$, and $\gamma_2 = 0.8$. The cycle-consistency loss is applied only when longitudinal scans are available, and longitudinal pairs with sufficiently large anatomical differences are selected to avoid degenerated identity mappings. To prevent conditional atlases from being biased, we follow prior work [19] and compute the centering loss only within local temporal neighborhoods. With a batch size of 8, the centering term is evaluated within each mini-batch. Additional implementation details will be provided in the publicly released code.

Comparison Methods. We compare SCALE with several state-of-the-art atlas construction frameworks: (1) ANTs [4]: a widely used symmetric groupwise normalization (SyGN) framework implemented in the ANTs package. (2) AtlasMorph [19]: a conditional atlas construction method that introduces a temporal neighborhood-based centrality loss to improve atlas centrality. (3) DACN [8]: an atlas-building framework that integrates longitudinal deformable registration to produce anatomically plausible atlas. (4) MultiMorph [1]: a fast feedforward framework that generates population-specific anatomical atlas in a single pass. To ensure fairness, all comparison methods use the same multi-channel inputs as our approach.

Evaluation. To quantitatively assess atlas quality, we performed a segmentation-transfer experiment, in which atlas labels were warped to each subject image using a separately trained VoxelMorph model. Anatomical alignment was evaluated using Dice scores [5] and the average symmetric surface distance (ASSD) [13]. Field regularity and topological correctness were assessed by computing the Jacobian determinant, $|J_\phi|$, of the deformation field ϕ , where voxels with $|J_\phi| \leq 0$ indicate local folding errors. We further quantified atlas centrality by measuring the mean squared displacement magnitude of the mean displacement field $\|\overline{u}\|_2^2$.

Table 1. Comparison of registration accuracy using different atlases.

Methods	Dice ↑	ASSD (mm) ↓	Folds ↓	Centrality ↓
ANTs	0.730 ± 0.024	0.616 ± 0.064	**1.015 ± 0.171**	0.102 ± 0.035
AtlasMorph	0.773 ± 0.018	0.544 ± 0.055	68.422 ± 31.114	0.087 ± 0.021
DACN	0.780 ± 0.020	0.424 ± 0.053	6.640 ± 1.935	0.091 ± 0.037
MultiMorph	0.781 ± 0.018	0.403 ± 0.050	1.241 ± 0.153	0.050 ± 0.008
SCALE	**0.803 ± 0.019**	**0.232 ± 0.046**	1.141 ±0.213	**0.048 ± 0.009**

3.3 Results

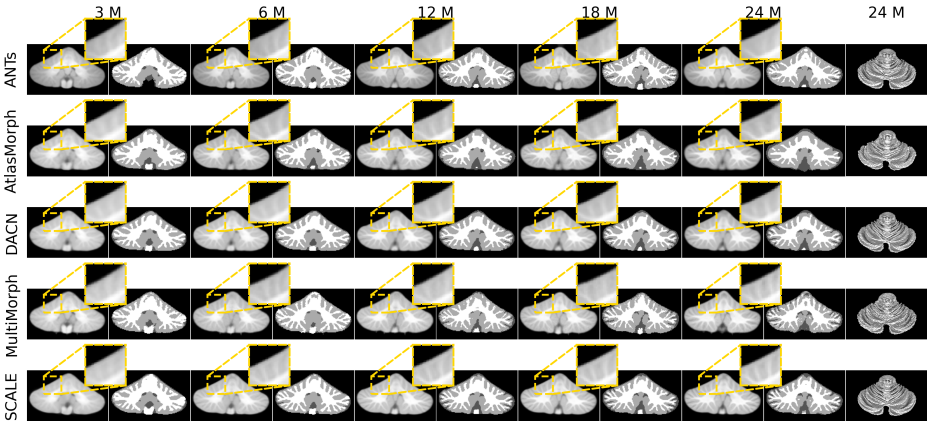

Fig. 2. Visual comparison of atlases, tissue segmentation maps, and the 24-month (24 M) white-matter surface (rightmost column) generated by different methods.

Table 1 summarizes the quantitative results on the BCP dataset [14]. Atlases with clearer anatomical structures generally yield higher registration accuracy. Although all competing methods use the same multi-channel inputs, our approach achieves the best Dice scores and ASSD performance, while maintaining deformation-field plausibility comparable to other state-of-the-art methods. In terms of atlas centrality, our method surpasses ANTs and performs comparably to AtlasMorph and MultiMorph. Fig. 2 shows the atlases produced by different methods. The atlas generated by our model exhibits sharper and more coherent anatomical structures.

By leveraging longitudinal registration as a reliable supervisory signal, the cycle-consistency term regularizes atlas evolution and helps produce clearer and more developmentally consistent atlases. As shown in Fig. 3, registering subjects to their age-matched atlases reveals a clear upward trend in Dice scores with increasing age. This indicates that the atlases generated by our method become progressively sharper and more structurally coherent over developmental stages. A similar trend can be observed in DACN, although our method consistently

Table 2. Quantitative ablation results of the proposed components.

Methods	Dice ↑	ASSD (mm) ↓	Folds ↓	Centrality ↓
Baseline	0.773 ± 0.018	0.544 ± 0.055	68.422 ± 31.114	0.087 ± 0.021
w/ H-Direct	0.781 ± 0.021	0.446 ± 0.049	1.045 ± 0.203	0.055 ± 0.010
w/ H-Derived	0.796 ± 0.018	0.236 ± 0.041	**0.912 ± 0.210**	0.057 ± 0.009
w/ Cycle	0.786 ± 0.019	0.336 ± 0.049	1.432 ± 0.203	**0.047 ± 0.005**
SCALE	**0.803 ± 0.019**	**0.232 ± 0.046**	1.141 ±0.213	0.048 ± 0.009

achieves higher accuracy across all ages. In contrast, MultiMorph, AtlasMorph, and ANTs exhibit a noticeable decline in Dice scores around 12–18 months. This degradation suggests that the atlases produced by these methods become less anatomically reliable at certain stages, likely due to subgroup-induced structural inconsistencies that blur fine-grained folia patterns. To further evaluate robustness and generalizability, we perform external validation on a cohort of 40 atypical infant scans without tissue segmentation. After aligning each scan to its age-matched atlas, we compute the Pearson correlation coefficient between each pair of warped images. Our method achieves 0.953 ± 0.006, which is significantly higher than MultiMorph (0.949 ± 0.007), DACN (0.950 ± 0.008), AtlasMorph (0.946 ± 0.010), and ANTs (0.944 ± 0.010) ($p<0.05$), indicating sharper structural detail and more reliable population-level alignment across atypical subjects.

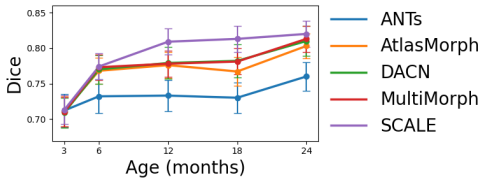

Fig. 3. Dice trajectories across five ages for different methods.

Table 2 reports the ablation results for each model component. Using the Hessian eigenvalue map as an additional output channel (w/ H-Direct, row 2) yields only marginal improvement. In contrast, constraining learning with a Hessian eigenvalue map derived from the generated atlas intensity volume (w/ H-Derived, row 3) leads to a clear performance gain, indicating more effective supervision. We attribute this to the fact that deriving the Hessian map from the intensity volume allows gradients to propagate more directly to the T1w generation branch. Rows 3 and 4 correspond to the two configurations of our method, each contributing to sharper and more anatomically coherent atlases.

4 Conclusion

In this study, we introduce a structure-preserving and cycle-consistent framework for 4D infant cerebellum atlas construction, leveraging unique anatomical and geometric characteristics of the cerebellum and informative longitudinal scans. By integrating a Hessian-based enhanced structural representation and cross-temporal cycle-consistency, our method produces atlases with sharper anatomical structures and markedly improved longitudinal consistency. Experiments on the BCP dataset show that SCALE achieves higher registration accuracy than competing approaches while maintaining plausible deformations and strong atlas centrality.

Acknowledgments. This work was supported in part by NIH grants (EB037388, AG075582, NS128534, NS135574, and HD119560).

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

References

1. Abulnaga, S.M., Hoopes, A., Dey, N., Hoffmann, M., Fischl, B., Guttag, J., Dalca, A.: Multimorph: On-demand atlas construction. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 30906–30917 (2025)
2. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 924–931. Springer (2006)
3. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Medical Image Analysis **12**(1), 26–41 (2008)
4. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ants similarity metric performance in brain image registration. Neuroimage **54**(3), 2033–2044 (2011)
5. Bajcsy, R., Kovačič, S.: Multiresolution elastic matching. Computer Vision, Graphics, and Image Processing **46**(1), 1–21 (1989)
6. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: An unsupervised learning model for deformable medical image registration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9252–9260 (2018)
7. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. IEEE Transactions on Medical Imaging **38**(8), 1788–1800 (2019)
8. Chen, L., Wu, Z., Hu, D., Pei, Y., Zhao, F., Sun, Y., Wang, Y., Lin, W., Wang, L., Li, G., et al.: Construction of longitudinally consistent 4d infant cerebellum atlases based on deep learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 139–149. Springer (2021)
9. Chen, L., Wu, Z., Hu, D., Wang, Y., Zhao, F., Zhong, T., Lin, W., Wang, L., Li, G.: A 4d infant brain volumetric atlas based on the unc/umn baby connectome project (bcp) cohort. NeuroImage **253**, 119097 (2022)

10. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning for fast probabilistic diffeomorphic registration. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I. pp. 729–738. Springer (2018)
11. Diedrichsen, J., Balsters, J.H., Flavell, J., Cussans, E., Ramnani, N.: A probabilistic mr atlas of the human cerebellum. neuroimage **46**(1), 39–46 (2009)
12. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: International conference on medical image computing and computer-assisted intervention. pp. 130–137. Springer (1998)
13. Heimann, T., Van Ginneken, B., Styner, M.A., Arzhaeva, Y., Aurich, V., Bauer, C., Beck, A., Becker, C., Beichel, R., Bekes, G., et al.: Comparison and evaluation of methods for liver segmentation from ct datasets. IEEE transactions on medical imaging **28**(8), 1251–1265 (2009)
14. Howell, B.R., Styner, M.A., Gao, W., Yap, P.T., Wang, L., Baluyot, K., Yacoub, E., Chen, G., Potts, T., Salzwedel, A., et al.: The unc/umn baby connectome project (bcp): An overview of the study design and protocol development. NeuroImage **185**, 891–905 (2019)
15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. CoRR **abs/1412.6980** (2014)
16. Lyu, W., Thung, K.H., Huynh, K.M., Wang, L., Lin, W., Ahmad, S., Yap, P.T.: Functional development of the human cerebellum from birth to age five. Nature Communications **16**(1), 6350 (2025)
17. Marmolin, H.: Subjective mse measures. IEEE Transactions on Systems, Man, and Cybernetics **16**(3), 486–489 (1986)
18. Nettekoven, C., Zhi, D., Shahshahani, L., Pinho, A.L., Saadon-Grosman, N., Buckner, R.L., Diedrichsen, J.: A hierarchical atlas of the human cerebellum for functional precision mapping. Nature Communications **15**(1), 8376 (2024)
19. Rakic, M., Hoopes, A., Mazdak, S., Sabuncu, M.R., Guttag, J.V., Dalca, A.V.: Atlasmorph: Learning conditional deformable templates for brain mri. Medical Image Analysis p. 103893 (2025)
20. Sathyanesan, A., Zhou, J., Scafidi, J., Heck, D.H., Sillitoe, R.V., Gallo, V.: Emerging connections between cerebellar development, behaviour and complex brain disorders. Nature Reviews Neuroscience **20**(5), 298–313 (2019)
21. Sun, Y., Wang, L., Gao, K., Ying, S., Lin, W., Humphreys, K.L., Li, G., Niu, S., Liu, M., Wang, L.: Self-supervised learning with application for infant cerebellum segmentation and analysis. Nature Communications **14**(1), 4717 (2023)
22. Tang, K., Chen, L., Wu, Z., Zhao, F., Wang, Y., Lin, W., Wang, L., Li, G.: Generation of anatomy-realistic 4d infant brain atlases with tissue maps using generative adversarial networks. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2024)
23. Van Essen, D.C.: A 2020 view of tension-based cortical morphogenesis. Proceedings of the National Academy of Sciences **117**(52), 32868–32879 (2020)
24. Wang, B., LeBel, A., D' Mello, A.M.: Ignoring the cerebellum is hindering progress in neuroscience. Trends in Cognitive Sciences **29**(4), 318–330 (2025)
25. Wang, L., Wu, Z., Chen, L., Sun, Y., Lin, W., Li, G.: ibeat v2. 0: a multisite-applicable, deep learning-based pipeline for infant cerebral cortical surface reconstruction. Nature protocols **18**(5), 1488–1509 (2023)
26. Wang, Y., Chen, L., Wu, Z., Li, T., Sun, Y., Cheng, J., Zhu, H., Lin, W., Wang, L., Huang, W., et al.: Longitudinal development of the cerebellum in human infants during the first 800 days. Cell reports **42**(4) (2023)