

Lifespan Cortical Surface Reconstruction from Thick-Slice Clinical MRI

Xiuyu Dong^{1,2}, Kaibo Tang¹, Dan Hu¹, Zhengwang Wu¹, Li Wang¹, Weili Lin¹ and Gang Li¹(✉)

¹ Department of Radiology and Biomedical Research Imaging Center, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

² Joint Department of Biomedical Engineering, University of North Carolina at Chapel Hill and North Carolina State University, Chapel Hill, NC, USA
gang_li@med.unc.edu

Abstract. Accurately characterizing brain morphological changes throughout human lifespan is crucial for understanding brain development, aging, and disorders. At the core of this endeavor lies cortical surface reconstruction (CSR), which underpins the computation of essential brain morphological features. However, existing CSR methods face two major limitations. First, cortical surfaces are typically reconstructed from 3D MRI data with high isotropic resolution, which is confined to research settings. In contrast, clinical MRI scans are collected with high in-plane but low through-plane resolution. Second, most CSR pipelines are designed either for adult or pediatric populations, restricting their applicability across the lifespan. To this end, we develop a deep learning framework that harnesses MRI super-resolution (SR) as a bridging mechanism, leveraging the complementary information SR provides to jointly perform SR and CSR with a coarse-to-fine strategy. Specifically, we introduce a dual-decoder age-conditioned temporal attention network (DATAN) with a shared encoder, which simultaneously performs CSR and SR from thick-slice clinical MRI. By jointly training on the SR task, the shared encoder captures richer cortical features, thereby enhancing CSR performance. Through a two-stage coarse-to-fine approach, incremental refinements in the SR output progressively restore fine-scale details otherwise lost in low-resolution scans, ultimately improving CSR fidelity. Furthermore, to facilitate accurate CSR across the lifespan, we exploit the age-conditioning module of our framework and train our model on a large, diverse MRI dataset spanning ages from 1 to 100 years. Experimental results demonstrate that our method, despite requiring only thick-slice clinical MRI scans, achieves consistently improved CSR performance across the entire human lifespan.

Keywords: Cortical Surface Reconstruction, Lifespan, Clinical MRI.

1 Introduction

Brain development is a lifelong, intricate process that can be characterized by morphological changes in the cerebral cortex, including variations in cortical thickness, sulcal

depth, surface area, etc. [1-4]. Reconstructing cortical surfaces from brain magnetic resonance imaging (MRI) scans is essential for the quantitative analysis of these morphological features and their changes over time. Existing large-scale neuroimaging datasets not only offer new opportunities to better understand cortical development over the entire human lifespan but also demand fast and robust cortical surface reconstruction (CSR) method that can generalize well across different age groups and acquisition protocols. Available neuroimaging pipelines, such as FreeSurfer [5], BrainSuite [6], the dHCP pipeline [7], the HCP pipeline [8], and iBEAT v2.0 [9], have achieved great success in cortical surface reconstruction within the research community. However, these pipelines typically involve multiple processing steps, which incur high computational costs and lengthy processing times, thereby limiting scalability to the rapidly growing volume of neuroimaging data.

Recent CSR methods based on geometric deep learning [10-20] have significantly reduced processing time from several hours to a few seconds for each subject, leveraging either implicit surface representations [10-12] or explicit mesh deformations [14-20]. However, two major challenges still remain: a) most of these methods are trained on datasets from specific age groups and are thus unlikely to generalize well to unseen data across the human lifespan; b) they generally rely on high-resolution (HR) 3D isotropic MRI scans, making them inapplicable for clinical MRI scans, which are generally acquired at low-resolution with large slice thickness.

To address these limitations, we propose a novel approach for lifespan CSR using low-resolution (LR) thick-slice brain MRIs commonly encountered in clinical practice. Our framework introduces a dual-decoder age-conditioned temporal attention network (DATAN) with a shared encoder. In particular, MRI super-resolution (SR) is treated as a bridging mechanism, which enables DATAN to learn complementary information relevant to both SR and CSR. Unlike CorticalFlow++ [16], which refines surface reconstruction through a cascade of deformation blocks, our two-stage design jointly performs SR and CSR, where the output SR image of the first stage is used as input of the second stage. This allows the second stage to leverage enriched spatial details, leading to more accurate surface predictions. To ensure accurate CSR across the entire human lifespan and facilitate its application in large-scale neuroimaging studies, we train and evaluate our method on a diverse dataset consisting of 8,987 T1-weighted (T1w) scans with ages spanning from 1 to 100 years.

2 Method

As illustrated in **Fig. 1(a)**, we developed a two-stage framework, which employs dual-decoder age-conditioned temporal attention network (DATAN) as backbone, for lifespan cortical surface reconstruction from thick-slice clinical brain MRI. In Stage 1, given a low-resolution thick-slice MRI (LR) as input, the DATAN jointly reconstructs a super-resolution MRI (SR-MRI 1) and a coarse cortical surface (Surface 1). Then, in Stage 2, DATAN takes SR-MRI 1 as input and reconstructs a second super-resolution MRI (SR-MRI 2), along with a fine cortical surface (Surface 2). Notably, both stages employ the same network and share the same initial template surface.

2.1 Dual-decoder Age-conditional Temporal Attention Network

We introduce DATAN, which is designed to reconstruct cortical surfaces across the entire human lifespan while overcoming the difficulties posed by low-resolution, anisotropic imaging data in clinical settings. As shown in **Fig. 1(b)**, DATAN takes a LR MRI (linearly interpolated to isotropic resolution) as input. With DATAN’s dual-decoder design, we are able to jointly optimize DATAN on the SR and CSR tasks.

Instead of treating SR and CSR as sequential and independent processes, we perform them jointly to leverage their shared reliance on rich spatial representation, particularly in the cortical regions where LR images typically appear blurred. To this end, DATAN employs a shared encoder that extracts rich spatial representations, followed by SR and CSR decoders, enabling concurrent training on both tasks.

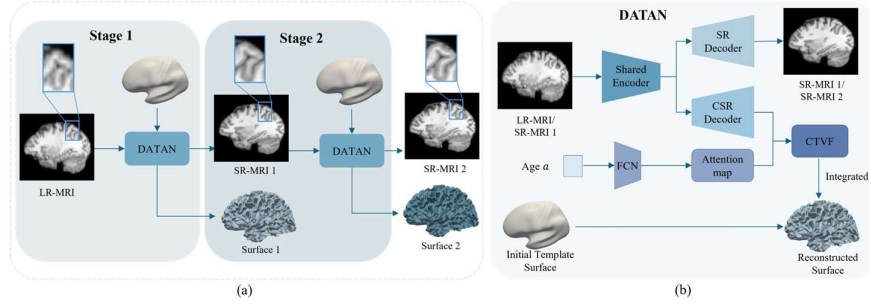


Fig. 1. (a) Coarse-to-fine two-stage lifespan cortical surface reconstruction from thick-slice clinical MRI. In Stage 1, given LR-MRI as input, DATAN jointly predicts SR-MRI 1 and the coarse cortical surface (Surface 1). In Stage 2, DATAN takes SR-MRI 1 as input and predicts the fine cortical surface (Surface 2) along with SR-MRI 2. LR-MRI: low-resolution (thick-slice) MRI; SR-MRI 1: coarse super-resolution MRI predicted in Stage 1; SR-MRI 2: fine super-resolution MRI predicted in Stage 2; Surface 1: coarse cortical surface reconstructed in Stage 1; Surface 2: fine cortical surface reconstructed in Stage 2. (b) The architecture of the dual-decoder age-conditioned temporal attention network (DATAN). With the dual-decoder design, DATAN is able to jointly perform MRI super-resolution (SR) and cortical surface reconstruction (CSR). Training on both tasks simultaneously enables the shared encoder to learn richer and complementary features, which improves the model’s performance on both tasks. Finally, the model is conditioned on age and trained on a large and diverse MRI dataset spanning from 1 to 100 years of age, which facilitates the model’s generalizability across the entire human lifespan. FCN: fully connected network; SVFs: multiple stationary velocity fields; CTVF: conditional time-varying velocity field.

The SR decoder generates super-resolution MRI from low-resolution input, serving two key purposes. On the one hand, inferring spatial details from low-resolution image compels the shared encoder to learn richer spatial representation, which directly benefits the CSR decoder. On the other hand, the SR process enhances spatial details of the input data that might otherwise be lost due to the anisotropic nature of clinical MRI, which subsequently improves CSR performance in Stage 2.

Following CoTAN [19], we model the diffeomorphic deformation $\phi_t: \mathbb{R}_x^3 \rightarrow \mathbb{R}_x^3$ through a time-varying velocity field (TVF) $v_t: \mathbb{R}_x^3 \rightarrow \mathbb{R}_x^3$, which, compared to stationary velocity field (SVF), allows us to capture more complex, large-scale deformations while still preserving smoothness and invertibility. The relationship between ϕ_t and v_t is given by the following ordinary differential equation (ODE) [10, 15, 16, 19–24]:

$$\frac{\partial \phi_t}{\partial t} = v_t(\phi_t), \phi_0 = Id, t \in [0, T] \quad (1)$$

where $Id: \mathbb{R}_x^3 \rightarrow \mathbb{R}_x^3$ is the identity mapping. The equivalent integral equation (IE) allows one to numerically solve the ODE in (1), i.e.,

$$\phi_T = Id + \int_0^T v_t(\phi_t) dt. \quad (2)$$

Let $S_0 \subset \mathbb{R}^3$ denotes the initial surface template with points $x_0 \in S_0$ and define $x_t := \phi_t(x_0)$. Evaluating both sides of (2) at x_0 , we get the equation as follows:

$$x_T = \phi_T(x_0) = x_0 + \int_0^T v_t(\phi_t(x_0)) dt = x_0 + \int_0^T v_t(x_t) dt. \quad (3)$$

Hence, in DATAN, we let the CSR decoder predict a conditional time-varying velocity field (CTVF) v_t , and obtain the predicted surface $S_T \subset \mathbb{R}^3$ with points $x_T \in S_T$ by numerically solving the IE in (3).

Finally, to handle large variation in cortical features and eventually facilitate accurate CSR across lifespan, we employ a fully connected network (FCN) that learns a conditional temporal attention map given time t and scan age a as input. The learned attention maps are then used to weigh each velocity field during integration.

2.2 Loss Function

In our approach, we design a hybrid loss function that enables joint optimization for both the SR and CSR tasks. Specifically, during training, we adaptively adjust the weights for SR loss (\mathcal{L}_{SR}) and CSR loss (\mathcal{L}_{CSR}), similar to the approach in [25]. The overall per-subject loss (\mathcal{L}_D) is defined as follows:

$$\mathcal{L}_D = \frac{1}{2\alpha_1^2} \mathcal{L}_{\text{SR}} + \frac{1}{2\alpha_2^2} \mathcal{L}_{\text{CSR}} + \log \alpha_1 \alpha_2, \quad (4)$$

where α_1 and α_2 are trainable hyper-parameters that are initialized to 1 and updated during training. The SR and CSR losses are described in the following sections.

Super-Resolution Loss. Given a pair of low- and high-resolution images (L, H), the SR task predicts the output super-resolution image \hat{H} . We use the mean squared error (MSE) as the loss function for the SR task, i.e., $\mathcal{L}_{\text{SR}} = \frac{1}{|\Omega|} \|\hat{H} - H\|_2^2$, where $|\Omega|$ represents the total number of voxels in the MRI.

Cortical Surface Reconstruction Loss. For white matter surface reconstruction, we use Chamfer distance loss \mathcal{L}_{cd} to quantify the deviation of the predicted surface from the ground-truth surface. In addition, to encourage the prediction of more plausible surfaces, we incorporate two regularization terms that are popular in literatures [14,19], namely, Laplacian loss \mathcal{L}_{lap} to encourage smoothness of the predicted surface, and normal consistency loss \mathcal{L}_{nc} to encourage consistency between the normal vectors of adjacent faces. The final CSR loss is given by:

$$\mathcal{L}_{CSR} = \mathcal{L}_{cd} + \lambda_{lap}\mathcal{L}_{lap} + \lambda_{nc}\mathcal{L}_{nc} \quad (5)$$

with weights λ_{lap} and λ_{nc} . For pial surface reconstruction, we follow [11,19] and use the ground-truth (GT) white matter surfaces as the input for training. Then MSE loss can be computed between the vertices of predicted and ground-truth pial surfaces.

3 Experiments

3.1 Dataset and Implementation Details

As illustrated in **Table 1**, we extensively evaluated our lifespan DATAN framework on a large-scale T1-weighted (T1w) MRI dataset comprising 8,987 scans from seven publicly available datasets. These datasets include diverse acquisition protocols and demographics, spanning infants, adolescents, adults, and elderly. Except ADNI, the remaining datasets were split into training (60%), validation (10%), and testing (30%). ADNI was used exclusively for external validation, with no involvement in training. All images were bias-corrected to start with, and treated as high-resolution (HR) MRIs. To simulate clinically acquired thick-slice data, we followed prior super-resolution studies [26, 27] and performed downsampling in the original physical space of each image. Specifically, consistent with the typical anisotropic spacing encountered in clinical MRI scans, we downsampled the high-resolution (HR) MRIs by a factor of 4 in the axial direction. All HR-LR MRI pairs were affinely aligned to the MNI152 template and cropped to a fixed size of $176 \times 224 \times 160$. During training, LR images were first linearly interpolated to match HR resolution before entering the network. GT cortical surfaces were reconstructed using the FreeSurfer pipeline [5], except for BCP and NDAR infant datasets, which were processed with iBEAT v2.0 [9]. We trained our model using the Adam optimizer with an initial learning rate of 1×10^{-5} , and empirically set the loss weights to $\lambda_{lap} = 0.5$ and $\lambda_{nc} = 5 \times 10^{-4}$.

3.2 Comparative Results

To demonstrate the advantage of our approach, we compared it against two most recent template-based cortical surface reconstruction approaches: CoTAN [19] and CoSeg [20]. To evaluate the geometric accuracy of the reconstructed cortical surfaces, we employ two widely used metrics: average symmetric surface distance (ASSD) and Hausdorff distance (HD). The ASSD quantifies the average discrepancy between two

surfaces. It is computed by sampling a set of points from one surface and determining the mean distance between these points and their corresponding points on the other surface. This process is then repeated in the reverse direction, and the final ASSD is obtained as the average of these two mean distances. Meanwhile, the HD metric measures the maximum surface discrepancy. To enhance robustness, we adopt the 90th percentile distance (HD90) rather than the absolute maximum, following the recommendations in [15] and [19]. Additionally, we assess the topological quality of the reconstructed cortical surfaces using the ratio of self-intersecting faces (SIF), which quantifies the extent of topological errors in the reconstructed surface.

Table 2 demonstrates that our method achieves superior performance in both white matter (WM) and pial surface reconstruction compared to other approaches. For white matter surface reconstruction, our method yields an ASSD of 0.261 mm for the left hemisphere, surpassing CoTAN and CoSeg by 39.44% and 63.34%, respectively. Furthermore, our method also improves HD accuracy, reducing errors by 38.20% and 62% compared to CoTAN and CoSeg. For pial surface reconstruction, our method produces 0.007% self-intersections, outperforming both CoTAN and CoSeg. Notably, CoSeg leverages cortical ribbon segmentations for weak supervision during training instead of GT surfaces but uses GT surfaces for evaluation. Hence, its poor performance can be attributed to the quality of segmentation in training data. Age-specific results are illustrated in **Figs. 2** and **3**. We evaluate our method on a large-scale T1w MRI dataset comprising seven public datasets, each covering a distinct age range from infancy to late adulthood (**Table 1**). To demonstrate our method’s generalizability across age groups, we visualize representative reconstructed cortical surfaces and their corresponding error distance maps for each dataset (**Figs. 2** and **3**), rather than reporting summarized quantitative metrics. Our method performs especially well in younger age groups, where baseline methods often struggle to capture fine cortical structures due to significant developmental variability. Despite the challenges posed by low-resolution lifespan MRI data, our method consistently produces more accurate and smoother surfaces with minimal self-intersections.

Table 1. Summary of datasets used across different age groups in our study. *: ADNI is used solely for validation, not training.

Dataset	Modality	Resolution (mm ³)	Matrix Size	Age range (Years)	Number of Scans
BCP [28]	T1w	0.8×0.8×0.8	208×300×320	1-6	707
NDAR [29]	T1w	1×1×1	195×233×150	1-2	941
HCPD [30]	T1w	0.8×0.8×0.8	320×320×320	6-22	652
HCP [31]	T1w	0.7×0.7×0.7	260×311×260	22-37	1111
HCPA [32]	T1w	0.8×0.8×0.8	320×320×320	36-100	721
NKI [33]	T1w	1×1×1	256×256×256	6-85	1613
ADNI* [34]	T1w	1×1×1	256×256×256	55-97	3242

3.3 Ablation Study

For validation, we conducted ablation studies to assess the effectiveness of each component of our method. We adopt CoTAN as our baseline, which consists only of the CSR decoder and takes low-resolution (LR) images as inputs. Next, we introduced

the SR decoder into our DATAN framework using LR inputs. The quantitative results for different settings are presented in **Table 3**, demonstrating that our SR tasks significantly enhance CSR performance on LR lifespan MRIs. By reconstructing spatial details from LR images, the SR decoder force the shared encoder to learn richer spatial representations, thereby improving the quality of CSR. Furthermore, the incorporation of the coarse-to-fine two-stage refinement learning strategy provides additional improvements, highlighting its effectiveness in deformation learning. Our DATAN at Stage 2 exhibits fewer errors and artifacts compared to other settings. These observed enhancements underscore the effectiveness of our proposed method.

Table 2. Comparison of different CSR methods in reconstructing the white matter and pial surfaces of the left hemisphere. Values are mean \pm std over all subjects.

Method	White Matter Surface			Pial Surface		
	ASSD (mm)	HD90 (mm)	SIF (%)	ASSD (mm)	HD90 (mm)	SIF (%)
CoTAN	0.431 \pm 0.103	0.953 \pm 0.262	0.007 \pm 0.030	0.442 \pm 0.098	0.989 \pm 0.263	0.057 \pm 0.067
CoSeg	0.712 \pm 0.092	1.550 \pm 0.417	0.188 \pm 0.121	0.735 \pm 0.157	1.783 \pm 0.697	0.395 \pm 0.225
Ours	0.261 \pm 0.057	0.589 \pm 0.145	0.001 \pm 0.010	0.400 \pm 0.083	0.880 \pm 0.216	0.007 \pm 0.023

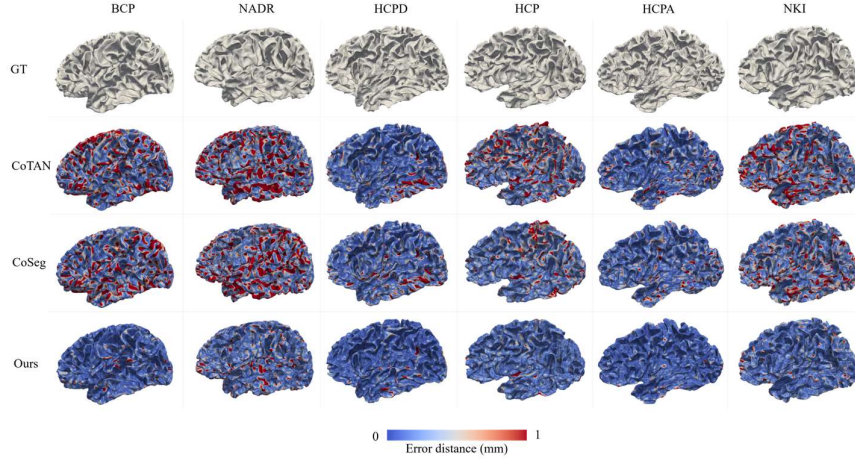


Fig. 2. Visualization of the reconstructed WM cortical surfaces for lifespan datasets. The geometric error map depicts the chamfer distance (from 0 to 1 mm) between the predicted surfaces and the ground truth.

Table 3. Ablation study for cortical surface reconstruction performance of our framework. Values are mean \pm std over all subjects. Baseline: CoTAN with LR input; S1: Stage 1; S2: Stage 2.

Method	White Matter Surface			Pial Surface		
	ASSD (mm)	HD90 (mm)	SIF (%)	ASSD (mm)	HD90 (mm)	SIF (%)
Baseline (LR)	0.431 \pm 0.10	0.953 \pm 0.26	0.007 \pm 0.03	0.442 \pm 0.10	0.989 \pm 0.26	0.057 \pm 0.07
DATAN (S1)	0.302 \pm 0.06	0.690 \pm 0.15	0.002 \pm 0.01	0.396 \pm 0.09	0.887 \pm 0.23	0.026 \pm 0.03
DATAN (S2)	0.261 \pm 0.06	0.589 \pm 0.14	0.001 \pm 0.01	0.400 \pm 0.08	0.880 \pm 0.22	0.007 \pm 0.02




Fig. 3. Visualization of the reconstructed pial cortical surfaces for lifespan datasets. The geometric error map depicts the chamfer distance (from 0 to 1 mm) between the predicted surfaces and the ground truth.

3.4 External Validation

To evaluate the robustness and generalizability of our method, we performed additional experiments on the ADNI dataset, which serves as an external dataset with 3,242 scans. The results in terms of ASSD, HD90, and SIF (**Table 4**) is comparable to the results on the testing dataset, which highlights our framework’s generalizability to unseen data.

Table 4. External validation for cortical surface reconstruction performance of our framework on ADNI dataset. Values are mean \pm std over all subjects.

	ASSD (mm)	HD90 (mm)	SIF (%)
White Matter Surface	0.281 ± 0.087	0.637 ± 0.394	0.001 ± 0.059
Pial Surface	0.445 ± 0.088	0.993 ± 0.336	0.005 ± 0.076

4 Conclusion

In this work, we introduce a novel two-stage framework for coarse-to-fine lifespan cortical surface reconstruction from thick-slice clinical brain MRI scans, which facilitates neuroimage analysis of lifespan cortical development in clinical settings. To tackle the challenges of the low-resolution (LR), anisotropic MRI in clinical practice, we present DATAN that employs a dual-decoder structure for jointly performing SR and CSR tasks. The two decoders share a common encoder to learn rich representation of spatial features to improve the CSR from LR input. Furthermore, we train our framework in a coarse-to-fine manner for more precise reconstruction. Experiments on the demanding task of cortical surface reconstruction using lifespan LR MRI data demonstrate the superiority of our approach. In future work, we will explore more extensive

data augmentation strategies, to accommodate for diverse MRI data in clinical settings and validate our method on more real clinical data.

Acknowledgments. The work of Gang Li was supported in part by NIH grants (AG075582, NS128534, EB037388, ES033518, MH123202, and NS135574).

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

References

1. Zhao, F., Wu, Z., Li, G.: Deep learning in cortical surface-based neuroimage analysis: a systematic review. *Intelligent Medicine* **3**(1), 46-58 (2023)
2. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis: I. segmentation and surface reconstruction. *Neuroimage* **9**(2), 179-194 (1999)
3. Li, G., Nie, J., Wang, L., et al: Measuring the dynamic longitudinal cortex development in infants by reconstruction of temporally consistent cortical surfaces. *Neuroimage* **90**, 266-279 (2014)
4. Fischl, B., Dale, A.M.: Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proceedings of the National Academy of Sciences* **97**(20), 11050-11055 (2000)
5. Fischl, B.: FreeSurfer. *Neuroimage* **62**(2), 774-781 (2012)
6. Shattuck, D.W., Leahy, R.M.: BrainSuite: an automated cortical surface identification tool. *Medical Image Analysis* **6**(2), 129-142 (2002)
7. Makropoulos, A., Robinson, E.C., Schuh, A., et al.: The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. *Neuroimage* **173**, 88-112 (2018)
8. Glasser, M.F., Sotiropoulos, S.N., et al.: The minimal preprocessing pipelines for the human connectome project. *Neuroimage* **80**, 105-124 (2013)
9. Wang, L., Wu, Z., Chen, L., et al.: iBEAT V2. 0: a multisite-applicable, deep learning-based pipeline for infant cerebral cortical surface reconstruction. *Nature protocols* **18**(5), 1488-1509 (2023)
10. Cruz, R.S., Lebrat, L., Bourgeat, P., et al.: DeepCSR: A 3D deep learning approach for cortical surface reconstruction. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 806-815 (2021)
11. Ma, Q., Li, L., Robinson, E.C., et al.: CortexODE: Learning cortical surface reconstruction by neural ODEs. *IEEE Transactions on Medical Imaging* (2022)
12. Hoopes, A., Iglesias, J.E., et al.: TopoFit: rapid reconstruction of topologicallycorrect cortical surfaces. *Proceedings of machine learning research* **172**, 508 (2022)
13. Zheng, H., Li, H., Fan, Y.: SurfNN: Joint Reconstruction of Multiple Cortical Surfaces from Magnetic Resonance Images. *arXiv preprint. arXiv:2303.02922* (2023)
14. Bongratz, F., Rickmann, A.M., Pölsterl, S., Wachinger, C.: Vox2Cortex: fast explicit reconstruction of cortical surfaces from 3D MRI scans with geometric deep neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20773-20783 (2022)
15. Lebrat, L., Santa Cruz, R., de Gournay, F., et al.: CorticalFlow: a diffeomorphic mesh transformer network for cortical surface reconstruction. *Advances in Neural Information Processing Systems* **34**, 29491-29505 (2021)

16. Santa Cruz, R., Lebrat, L., Fu, D., et al.: CorticalFlow++: Boosting cortical surface reconstruction accuracy, regularity, and interoperability. In: Medical Image Computing and Computer Assisted Intervention. pp. 496–505. Springer (2022)
17. Chen, X., Zhao, J., Liu, S., et al.: SurfFlow: A flow-based approach for rapid and accurate cortical surface reconstruction from infant brain MRI. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 380–388. Springer (2023)
18. Dong, X., Wu, Z., Ma, L., et al.: Cycle-consistent learning for fetal cortical surface reconstruction. In: Medical Image Computing and Computer Assisted Intervention. pp. 212–222. Springer (2024)
19. Ma, Q., Li, L., Kyriakopoulou, V., et al.: Conditional temporal attention networks for neonatal cortical surface reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 312–322. Springer (2023)
20. Ma, Q., Li, L., Robinson, E. C., et al.: Weakly Supervised Learning of Cortical Surface Reconstruction from Segmentations. In: Medical Image Computing and Computer Assisted Intervention. pp. 766–777. Springer (2024)
21. Gupta, K., Chandraker, M.: Neural mesh flow: 3D manifold mesh generation via diffeomorphic flows. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. pp. 1747–1758 (2020)
22. Ashburner, J.: A fast diffeomorphic image registration algorithm. *Neuroimage* **38**(1), 95–113 (2007)
23. Balakrishnan, G., Zhao, A., Sabuncu, M.R., et al.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
24. Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision* **61**, 139–157 (2005)
25. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7482–7491 (2018)
26. Zhao, C., Dewey, B.E., Pham, D.L., et al.: SMORE: a self-supervised anti-aliasing and super-resolution algorithm for MRI using deep learning. *IEEE Transactions on Medical Imaging* **40**(3), 805–817 (2020)
27. Wu, W., Qu, R., Shi, D., et al.: Cortical Surface Reconstruction from 2D MRI with Segmentation-Constrained Super-Resolution and Representation Learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 99–108. Springer (2024)
28. Howell, B.R., Styner, M.A., Gao, W., et al.: The UNC/UMN Baby Connectome Project (BCP): An overview of the study design and protocol development. *Neuroimage* **185**, 891–905 (2019)
29. Girault, J.B., Cornea, E., Goldman, B.D., et al.: Cortical structure and cognition in infants and toddlers. *Cerebral cortex* **30**(2), 786–800 (2020)
30. Somerville, L.H., Bookheimer, S.Y., Buckner, R.L., et al.: The Lifespan Human Connectome Project in Development: A large-scale study of brain connectivity development in 5–21 year olds. *Neuroimage* **183**, 456–468 (2018)
31. Van Essen, D.C., Smith, S.M., Barch, D.M., et al.: The WU-Minn human connectome project: an overview. *Neuroimage* **80**, 62–79 (2013)
32. Bookheimer, S.Y., Salat, D.H., Terpstra, M., et al.: The lifespan human connectome project in aging: an overview. *Neuroimage* **185**, 338–348 (2013)

33. Tobe, R.H, MacKay-Brandt, A., Lim, R., et al.: A longitudinal resource for studying connectome development and its psychiatric associations during childhood. *Scientific data* **9**(1), 300 (2022)
34. Jack Jr, C.R., Bernstein, M.A., Fox, N.C., et al.: The Alzheimer’s disease neuroimaging initiative (ADNI): MRI methods. *Journal of Magnetic Resonance Imaging* **27**(4), 685–691 (2008)