Cycle-consistent Learning for Fetal Cortical Surface Reconstruction

Xiuyu Dong[1], Zhengwang Wu[1], Laifa Ma[1], Ya Wang[1], Kaibo Tang[1], He Zhang[2], Weili Lin[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
[2] Department of Radiology, Obstetrics and Gynecology Hospital, Fudan University, Shanghai, China
gang_li@med.unc.edu

Abstract. Fetal cortical surface reconstruction is crucial for quantitative analysis of normal and abnormal prenatal brain development. While there are many cortical surface reconstruction methods available for adults and infants, there remains a notable scarcity of dedicated techniques for fetal cortical surface reconstruction. Of note, fetal brain MR images present unique challenges, characterized by nonuniform low tissue contrast associated with extremely rapid brain development and folding during the prenatal stages and low imaging resolution, as well as susceptibility to severe motion artifacts. Moreover, the smaller size of fetal brains results in much narrower cortical ribbons and sulci. Consequently, the fetal cortical surfaces are more prone to be influenced by partial volume effects and tissue boundary ambiguities. In this work, we develop a multi-task, priori-knowledge supervised fetal cortical surface reconstruction method based on deep learning. Our method incorporates a cycle-consistent strategy, utilizing prior knowledge and multiple stationary velocity fields to enhance its representation capabilities, enabling effective learning of diffeomorphic deformations from the template surface mesh to the inner and outer surfaces. Specifically, our framework involves iteratively refining both inner and outer surfaces in a cyclical manner by mutually guiding each other, thus improving accuracy especially for ambiguous and challenging cortical regions. Evaluation on a fetal MRI dataset with 83 subjects shows the superiority of our method with a geometric error of 0.229 ± 0.047 mm and 0.023 ± 0.058% self-intersecting faces, indicating promising surface geometric and topological accuracy. These results demonstrate a great advancement over state-of-the-art deep learning methods, while maintaining high computational efficiency.

Keywords: Fetal Cortical Surface Reconstruction, Cycle-consistent, Multi-task.

1 Introduction

Cortical surface reconstruction (CSR) aims to generate the mesh-represented surfaces of the thin, convoluted cerebral cortex from brain magnetic resonance images (MRI), including inner (white matter) and outer (pial) cortical surfaces [1]. The inner cortical surface represents the interface between the cortical gray matter (GM) and white matter (WM), while the outer surface represents the interface between the cerebrospinal fluid

(CSF) and GM. CSR is thus of vital importance for surface-based neuroimaging analyses. Traditional CSR approaches usually employ empirically defined computational pipelines, which typically consist of two main steps [2]. The first step is to reconstruct the inner cortical surface, involving the correction of topological errors in the segmented WM and generating the inner surface using classic tessellation methods, such as the Marching Cubes algorithm [3] and its variants [4]. The second step is to reconstruct the outer surface based on the inner surface. This is achieved by deforming the inner surface to the GM/CSF boundary with an effective deformation strategy, which needs to overcome challenges that opposing sulcal banks could be closer than the MRI resolution, while avoiding surface mesh self-intersection. Various CSR pipelines for neuroimaging analysis have been developed and widely utilized, such as FreeSurfer [5], BrainSuite [6], dHCP pipeline [7], and iBEAT v2.0 [8]. However, these pipelines share a common limitation in terms of computational efficiency.

Recently, several deep learning-based methods have emerged for cortical surface reconstruction with high efficiency. Cruz et al. [9] introduced the DeepCSR framework, which predicts implicit surface representations utilizing occupancy fields and signed distance functions (SDFs). Other methods such as CortexODE [10] and SurfNN [11] deform an initial surface reconstructed from a tissue segmentation map following subcortical filling and topology correction. Recent approaches [12-14] directly initiate deformation either from a convex hull or an inflated white matter/inner surface, e.g., CorticalFlow [12], CorticalFlow++ [13], and CoTAN [14]. However, these methods primarily target adults or neonates and are not fully suitable for fetuses for the following reasons. 1) Deep learning-based approaches heavily rely on training data, while fetal brain MRI scans are very challenging to acquire and are often of lower quantity and quality compared to those of adults and neonates. As shown in **Fig. 1**, fetal brain MRI exhibits lower image resolution and severer partial volume effects compared to adults and infants. 2) Noise in fetal MRI is often apparent, stemming from unavoidable motion artifacts. Therefore, purely relying on local features for deformation prediction is more challenging in fetal MRI. 3) During the prenatal period, the cerebral cortex undergoes exceptionally rapid development and folding, transitioning from a smooth shape to a highly convoluted structure. This leads to significant variations in cortical shape and size across ages and subjects, resulting in difficulties in learning deformation prediction. 4) The smaller size of fetal brains results in much narrower cortical ribbons and sulci. Consequently, the fetal cortical surfaces are more prone to be influenced by partial volume effects and ambiguities in image appearances.

To overcome the challenges mentioned above, we propose a novel approach for fetal cortical surface reconstruction, leveraging large neonatal data for pretraining to improving the network training, multi-task learning to fully explore the contextual information, and prior knowledge supervision to alleviate the ambiguity in representation learning (see **Fig. 2**). Specifically, our framework incorporates a cycle-consistent strategy, which involves stepwise deformation of an initial smooth template mesh to sequentially align inner and outer surfaces. Then, it iteratively refines both surfaces in a cyclical manner by mutually guiding each other, thus improving accuracy especially for ambiguous and challenging cortical regions. Furthermore, our method simultaneously predicts the signed distance function (SDF) and tissue map. Integrating these tasks into a

unified framework greatly improves the accuracy and robustness of fetal cortical surface reconstruction, even with limited and lower-quality fetal brain MRI data.

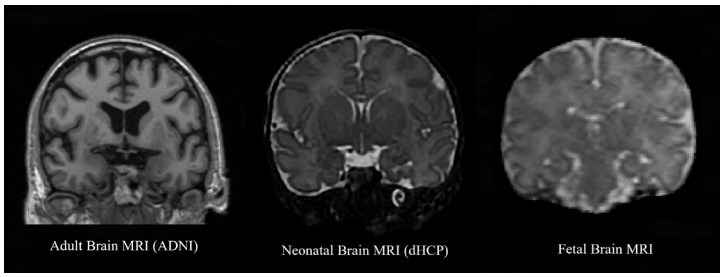

Fig. 1. Left: adult brain T1w MRI from the ADNI dataset [15]. Middle: neonatal brain T2w MRI from the dHCP dataset [16]. Right: fetal brain T2w MRI.

2 Method

Overview. We develop a multi-task, priori-knowledge supervised, and cycle-consistent framework for fetal cortical surface reconstruction. Specifically, this framework incorporates a cycle-consistent strategy, utilizing prior knowledge and multiple stationary velocity fields (SVFs) to enhance its representation capabilities, enabling precise learning of deformations from the template surface to the inner and outer surfaces, as illustrated in **Fig. 2**. The framework takes a 3D reconstructed fetal brain MRI scan and a surface mesh template as the initial input, predicting diffeomorphic deformation fields from the template to the inner surface and from the inner surface to the outer surface. Simultaneously, it estimates the corresponding signed distance function (SDF) for brain tissue segmentation. Based on the initial prediction results, we further employ an iterative computational loop between the inner surface and outer surface to improve the overall accuracy of the reconstruction. Due to the GPU memory limitation, we limit the number of iterations to 2 in this work.

Pretraining Strategy. To address the challenge of limited availability of fetal MRI data, we adopt a pretraining strategy to mitigate this scarcity. During the pretraining phase, we pretrained our deformation network using the public dHCP preterm and term neonatal 3D MRI dataset (877 subjects) [7]. This neonatal dataset includes early developing brains from 27 to 45 postmenstrual weeks, thus enabling our model to effectively learn dynamic cortical folding patterns during the perinatal stage. This foundational training facilitated capturing crucial structural characteristics and variation patterns essential for accurate fetal cortical surface modeling. Furthermore, the pretraining phase equipped the model with generalizable feature representations, pivotal for robust inference and adaptation to various fetal scans. This strategy not only accelerated conver-

gence during model refinement, but also enhanced reconstruction accuracy, outperforming conventional methods by mitigating the risk of overfitting to limited, low-quality fetal data.

Cycle-consistent Refinement. To tackle the intrinsic difficulty arising from the low contrast in fetal images and the ambiguous tissue boundaries, we introduce a cycle-consistent learning strategy, as shown in **Fig. 2(a)**. This strategy involves iteratively refining predicted deformation fields by considering the intricate relationship between the inner surface and outer surface. Instead of predicting the large deformation in a single step, we break down the substantial deformation into a series of smaller adjustments. By leveraging the inherent relationship between the inner surface and outer surface, we effectively guide the network's deformation field converging to a more precise position, so as to compensate for ambiguities associated with poor image quality. The diffeomorphic deformation flow is estimated at each iteration, which is parameterized by several conditional time-varying velocity fields (CTVF) [14] v_t. The diffeomorphic deformation flow can then be computed by solving the following ordinary differential equation (ODE) [10,12,13,17,18,19,20]:

$$\frac{\partial}{\partial t}\phi_t = v_t(\phi_t) \tag{1}$$

where ϕ_t and v_t denote the diffeomorphic deformation and CTVF at time $t \in [0, T]$. For each iteration i, given the initial surface S_0^i with points $x_0^i \in S_0^i$, and the estimated diffeomorphic flow ϕ_T^i, we compute the deformed points $x_T^i = \phi_T^i(x_0^i)$ on the deformed surface $S_T^i = \phi_T^i(S_0^i)$. It is important to note that in the first iteration, only the T2-weighted MRI serves as input. In subsequent iterations, the model input expands to include the Signed Distance Function (SDF) output from the previous iteration. Formally, $SDF^i, \phi_T^i = \mathcal{F}^i(I,\ SDF^{i-1})$, with $SDF^0 = 0$, where $\mathcal{F}^i$ represents the i-th deformation network and I denotes the input T2-weighted image.

Multi-task Priori-knowledge based Deformation Network. Each component of our deformation network consists of two branches: one dedicated to predicting the flow field and the other dedicated to predicting the segmentation mask. The incorporation of brain tissue segmentation as an auxiliary task endows the highly correlated segmentation and deformation prediction tasks to interact with each other to improve the network learning for the models with anatomy-awareness, facilitating superior capturing of cortical boundaries and preventing the prediction of surfaces with artifacts. Notably, our tissue map is computed by a signed distance function, which represents an implicit surface. The SDF encodes rich anatomical information (e.g., the relative position and shape information), thus aiding in the reconstruction of high-quality surfaces. Our deformation network adopts a 3D UNet-like structure, simultaneously predicting a signed distance function feature map and multiple stationary velocity fields (SVFs). To address the challenge of accommodating large deformations, while maintaining architectural simplicity, we extract feature maps at K resolution levels. Subsequently, we upsample multi-scale feature maps and learn N volumetric SVFs for each resolution. Considering

the high variation between fetal brains, we encode the gestational age (GA) as a conditioning variable in our network, as in CoTAN [14]. This allows learning knowledge-conditioned attention map to adaptively adjust the weight of each SVF [14]. Thus, the integration of these weighted SVFs results in a refined and precise learning of the deformation field. To elaborate, as illustrated in **Fig. 2(b)**, a Fully Connected Network (FCN) is employed to encode GA a into one $(K \times N) \times 1$ feature vector. Following reshaping and SoftMax activation, the FCN produces an attention map. Subsequently, a conditional deformation field is predicted by calculating the weighted sum of all SVFs and the attention map. This mechanism thus facilitates individualized deformation for different subjects.

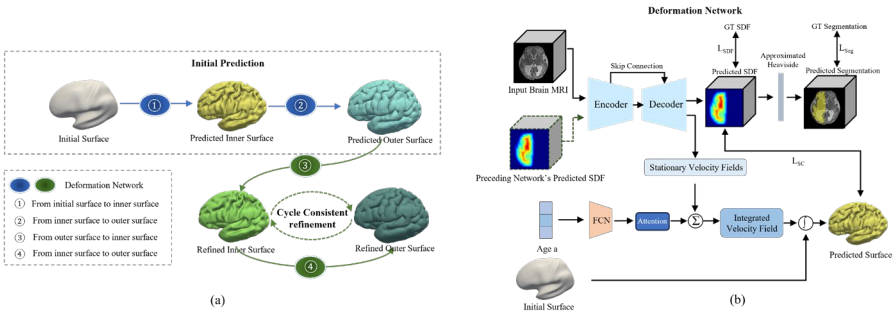

Fig. 2. Overview of our framework. (a) Cycle-consistent fetal cortical surface reconstruction pipeline. The deformation network first deforms from a template surface to the inner surface and then to the outer surface as initial prediction. Subsequently, the cycle improves the deformation in iterative steps to refine both by mutually guiding each other. (b) The architecture of the deformation network in our framework. Given a fetal brain MRI volume and a template surface, the deformation network learns to predict signed distance function (SDF) and then uses Heaviside function to compute the segmentation mask. Simultaneously, a priori knowledge-conditioned attention map, encoding by age a, guides the weighting of multiple stationary velocity fields (SVFs). By integrating these SVFs, the network produces a robust deformation field, leading to accurate fetal cortical surfaces.

Loss Function. In our approach, we design a comprehensive loss function, encompassing a surface space loss $\mathcal{L}_S$ and an image space loss $\mathcal{L}_I$, balanced by a hyperparameter λ:

$$\mathcal{L}_{total} = \mathcal{L}_S + \lambda \mathcal{L}_I \tag{2}$$

Image Space Loss. The image space loss $\mathcal{L}_I$ is a combination of signed distance loss and segmentation loss, weighted by α, which serves as an auxiliary task:

$$\mathcal{L}_I = \mathcal{L}_{SDF} + \alpha \mathcal{L}_{Seg} \tag{3}$$

The component of $\mathcal{L}_{SDF} = \mathcal{L}_1 + \mathcal{L}_{SC}$ includes the $\mathcal{L}_1$ term ensuring the similarity between the predicted implicit surface and the ground truth implicit surface, while the $\mathcal{L}_{SC}$ term acts as a double constraint—a smooth constraint for the reconstructed surface and

a structure constraint for the predicted implicit surface. Specifically, we extract the explicit surface from the predicted SDF (implicit surface) of the first branch and then compute the Chamfer distance with the reconstructed surface of the second branch. Additionally, the $\mathcal{L}_{Seg}$ supervises the overlapping between the predicted segmentation (obtained from SDF using the Heaviside function) and the ground truth segmentation, defined through the Dice loss.

Surface Space Loss. The $\mathcal{L}_S$ term is dedicated to surface reconstruction. In the initial iteration ($i = 1$) for inner surface reconstruction, $\mathcal{L}_S = \mathcal{L}_{cd} + \gamma\mathcal{L}_{lap} + \delta\mathcal{L}_{nc}$. Here, $\mathcal{L}_{cd}$ computes the distance between two surfaces, using the Chamfer distance loss [21-24], the Laplacian loss $\mathcal{L}_{lap}$ regularizes the smoothness of the surface mesh, and the normal consistency loss $\mathcal{L}_{nc}$ constrains the cosine similarity between the normals of two adjacent faces, as in CoTAN [14]. For subsequent iterations ($i > 1$), $\mathcal{L}_S$ transitions to only using the Chamfer distance loss. The outer surface reconstruction employs the Chamfer distance loss throughout iterations.

3 Experiments

Dataset and Implementation Details. Our framework was evaluated on a 3D reconstructed fetal brain MRI dataset consisting of 83 subjects ranging from 21 to 35 weeks of gestation. We divided the dataset into training, validation, and testing sets with a ratio of 6:1:3. The MRI volumes were affinely aligned to the MNI152 template and clipped to a size of 176 × 224 × 160 at 1 mm³ isotropic resolution. To generate pseudo-ground truth cortical surface meshes, we first generated initial tissue segmentation maps [25] and further extensively manually corrected them and then reconstructed the inner and outer surfaces using a deformable surface model [26]. During our training phase, the hyperparameters λ and α were set to 1×10^{-4} and 1, respectively. Our network was trained sequentially, with parameters frozen after the first iteration and subsequent iterations initiated. We employed the Adam optimizer with an initial learning rate setting to 1×10^{-5}, and the weights $\gamma = 0.5$ and $\delta = 5 \times 10^{-4}$.

Comparative Results. For validation, we first performed ablation studies to ensure the effectiveness of each part of our method. To demonstrate the advantage of our approach, we compared with the most recent learning-based cortical surface reconstruction approaches, including CortexODE [10] and CoTAN [14], as these methods can well reconstruct neonatal cortical surfaces. To measure geometric accuracy of reconstructed cortical surfaces, we employ two metrics [9]: average symmetric surface distance (ASSD) and Hausdorff distance (HD). The ASSD provides an assessment of the average distance between two surfaces. This involves sampling a batch of points on one surface and calculating the mean distance between these points and their corresponding points on the other surface. The process is then reciprocated, resulting in the final ASSD as the average of these mean distances. Simultaneously, the HD metric computes the maximum distance between two surfaces, with an emphasis on robustness achieved by

computing the 90th percentile distance rather than the maximum value, as suggested by [9] and [12]. In addition, the topology quality of the reconstructed cortical surfaces will be evaluated using the ratio of self-intersecting faces (SIF).

Experimental results in **Table 1** indicate that our method yields an ASSD of 0.2275 mm for the left hemisphere and 0.2285 mm for the right hemisphere for the pial/outer surface. For the inner surface, our method yields an ASSD of 0.2581 mm for the left hemisphere and 0.2627 mm for the right hemisphere. In contrast, CortexODE exhibits an average ASSD of 0.2821 mm for the left hemisphere and 0.3126 mm for the right hemisphere for the pial surface. For the inner surface, CortexODE yields an average ASSD of 0.3027 mm for the left hemisphere and 0.3512 mm for the right hemisphere. These results indicate a substantial improvement in ASSD, comparing our method with CortexODE, while CoTAN shows clear improvements in ASSD and HD for both inner and pial surfaces compared with CortexODE, it still falls short of our method. **Fig. 3** shows representative cases for a visual comparison at various ages. Our method consistently produces more accurate and smoother surfaces with fewer self-intersecting faces, despite the influence of poor quality and low contrast in fetal MRI scans.

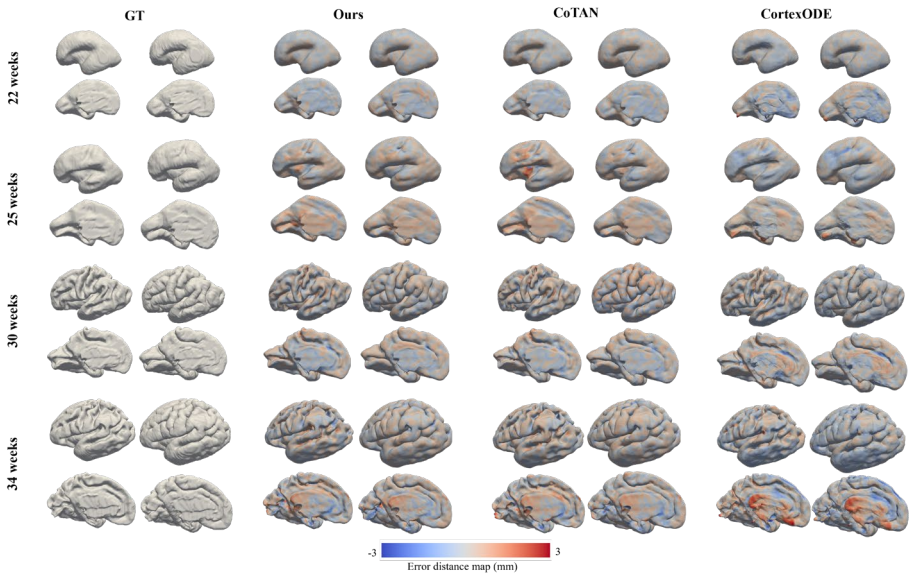

Fig. 3. Visual comparison of cortical inner and outer surfaces reconstructed by various methods across ages. Blue color indicates under-deformation error distance, while red color indicates over-deformation error distance.

Ablation Study. We conducted an ablation study to assess the effects of the signed distance function branch and the use of a cycle-consistent strategy on surface reconstruction performance. Quantitative results for different settings are presented in **Table 2**. It is evident that our cycle-consistent strategy reduces errors in CSR from fetal brain

MRIs. The signed distance function branch, while contributing only slightly to quantitative results, has a more substantial impact on the visual quality of the generated surfaces, notably reducing topology errors in predictions when visually examined in the testing data. Furthermore, the incorporation of the cycle-consistent refinement learning strategy yields noticeable additional improvements, underscoring its effectiveness in deformation learning. These error distance maps, computed by measuring the distance of each vertex on the predicted surface to its nearest counterpart on the ground truth surface, provide valuable insights into geometric errors.

Table 1. Comparison of different CSR methods in reconstructing the inner and pial surfaces of the left (L) and right (R) hemispheres. Values are mean ± std over all subjects.

	Inner Surface					
	ASSD (mm)		HD (mm)		SIF (%)	
	L	R	L	R	L	R
CortexODE	0.3027±0.1252	0.3512±0.1164	0.7413±0.4125	0.8276±0.3350	0.6720±1.0608	0.0104±0.0247
CoTAN	0.2673±0.0712	0.2738±0.0609	0.6145±0.2511	0.6085±0.1462	0.0389±0.1451	0.0135±0.0164
Ours	**0.2581±0.0705**	**0.2627±0.0602**	**0.5939±0.2528**	**0.5875±0.1511**	**0.0251±0.0973**	**0.0063±0.0169**
	Pial Surface					
	ASSD (mm)		HD (mm)		SIF (%)	
	L	R	L	R	L	R
CortexODE	0.2821±0.1200	0.3126±0.1278	0.6786±0.4124	0.7503±0.3958	0.4111±0.5247	0.2284±0.2166
CoTAN	0.2326±0.0622	0.2431±0.0590	0.4917±0.1519	0.5079±0.1200	**0.0492±0.1655**	0.0297±0.0333
Ours	**0.2275±0.0638**	**0.2285±0.0470**	**0.4820±0.1545**	**0.4768±0.1055**	0.0566±0.1817	**0.0228±0.0584**

Table 2. Cortical surface reconstruction performance with respect to signed distance function branch (SDF and Seg) and cycle-consistent strategy. Left (L) and right (R) hemispheres of the inner and pial surfaces are compared separately. Values are mean ± std over all subjects.

		SDF		√	√	√
		Seg			√	√
		Cycle-consistent				√
ASSD (mm)	Inner	L	0.2707±0.0716	0.2675±0.0704	0.2679±0.0698	0.2581±0.0705
		R	0.2819±0.0662	0.2671±0.0556	0.2663±0.0538	0.2627±0.0602
	Pial	L	0.2280±0.0630	0.2262±0.0613	0.2244±0.0611	0.2275±0.0638
		R	0.2325±0.0489	0.2336±0.0494	0.2306±0.0455	0.2285±0.0470
HD90 (mm)	Inner	L	0.6250±0.2398	0.6198±0.2557	0.6045±0.2272	0.5939±0.2528
		R	0.6494±0.1898	0.5894±0.1399	0.5867±0.1319	0.5875±0.1511
	Pial	L	0.4815±0.1481	0.4723±0.1380	0.4673±0.1260	0.4820±0.1545
		R	0.4957±0.1245	0.4874±0.1089	0.4819±0.0986	0.4768±0.1055
SIF (%)	Inner	L	0.0125±0.0483	0.0424±0.1670	0.0125±0.0483	0.0251±0.0973
		R	0.0053±0.0096	0.0038±0.0105	0.0053±0.0096	0.0063±0.0169
	Pial	L	0.0381±0.1311	0.0700±0.2180	0.0381±0.1311	0.0566±0.1817
		R	0.0441±0.0621	0.0277±0.0635	0.0441±0.0621	0.0228±0.0584

4 Conclusion

In this study, we introduce a novel cycle-consistent, diffeomorphic flow-based method for reconstruction of fetal cortical surfaces, which facilitates precise analysis of fetal cortical development and early, better detection of neurodevelopmental anomalies. To tackle the challenges stemming from the rapid changes during the prenatal period and the low contrast of MRI scans, we present a refinement learning strategy that ensures cycle consistency. This strategy involves gradually deforming a surface mesh template towards the inner surface and subsequently towards the pial surface, followed by iterative cyclical refinement. Experiments on the demanding task of cortical surface reconstruction using fetal brain MRI data demonstrate the superiority of our approach. Our method achieves significant improvements in mesh quality and reduction of geometric errors, compared to state-of-the-art deep learning methods, while maintaining high computational efficiency, i.e., 0.3 seconds for each hemisphere. Further improvements can be realized through the incorporation of advanced post-processing methods and novel loss functions to further reduce the SIF ratio and enhance surface quality. In future, we will also comprehensively evaluate our method on larger cohorts of fetal scans.

Acknowledgments. The work of Gang Li was supported in part by NIH grants (MH123202, ES033518, 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. Li, G., Nie, J., Wu, G., Wang, Y., Shen, D.: Consistent reconstruction of cortical surfaces from longitudinal brain MR images. Neuroimage **59**(4), 3805-3820 (2012)
3. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. ACM siggraph Computer Graphics **21**(4), 163–169 (1987)
4. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis: I. Segmentation and surface reconstruction. Neuroimage **9**(2), 179–194 (1999)
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. 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)
9. 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)

10. Ma, Q., Li, L., Robinson, E.C., et al.: CortexODE: Learning cortical surface reconstruction by neural ODEs. IEEE Transactions on Medical Imaging (2022)
11. Zheng, H., Li, H., Fan, Y.: SurfNN: Joint Reconstruction of Multiple Cortical Surfaces from Magnetic Resonance Images. arXiv preprint. arXiv:2303.02922 (2023)
12. 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)
13. 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)
14. 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)
15. 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)
16. Edwards, A.D., Rueckert, D., Smith, S.M., et al.: The developing human connectome project neonatal data release. Frontiers in neuroscience **16** (2022)
17. 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)
18. Ashburner, J.: A fast diffeomorphic image registration algorithm. Neuroimage **38**(1), 95–113 (2007)
19. 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)
20. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International journal of computer vision **61**, 139–157 (2005)
21. 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)
22. Wang, N., Zhang, Y., Li, Z., et al.: Pixel2Mesh: Generating 3D mesh models from single RGB images. In: Proceedings of the European conference on computer vision (ECCV). pp. 52–67 (2018)
23. Wickramasinghe, U., Remelli, E., Knott, G., Fua, P.: Voxel2Mesh: 3D mesh model generation from volumetric data. In: Medical Image Computing and Computer Assisted Intervention. pp. 299–308. Springer (2020)
24. Butt, M.A., Maragos, P.: Optimum design of chamfer distance transforms. IEEE Transactions on Image Processing **7**(10), 1477–1484 (1998)
25. Fidon, L., Aertsen, M., Kofler, F., et al.: A Dempster-Shafer approach to trustworthy AI with application to fetal brain MRI segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2024)
26. 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)