Cortical Surface Reconstruction from 2D MRI with Segmentation-Constrained Super-Resolution and Representation Learning

Wenxuan Wu[1,2], Ruowen Qu[1], Dongzi Shi[1], Tong Xiong[1], Xiangmin Xu[1,3], Xiaofen Xing[1(✉)], and Xin Zhang[1(✉)]

[1] The School of Electronic and Information Engineering, South China University of Technology, Guangzhou 510640, China
eexinzhang@scut.edu.cn
xfxing@scut.edu.cn
[2] BGI Research, Hangzhou 310030, China
[3] Institute of Artificial Intelligence, Hefei Comprehensive National Science Center, Hefei 230088, China

Abstract. Cortical surface reconstruction typically relies on high-quality 3D brain MRI to establish the structure of cortex, playing a pivotal role in unveiling neurodevelopmental patterns. However, clinical challenges emerge due to elevated costs and prolonged acquisition times, often resulting in low-quality 2D brain MRI. To optimize the utilization of clinical data for cerebral cortex analysis, we propose a two-stage method for cortical surface reconstruction from 2D brain MRI images. The first stage employs segmentation-constrained MRI super-resolution, concatenating the super-resolution (SR) model and cortical ribbon segmentation model to emphasize cortical regions in the 3D images generated from 2D inputs. In the second stage, two encoders extract features from the original and super-resolution images. Through a shared decoder and the mask-swap module with multi-process training strategy, cortical surface reconstruction is achieved by mapping features from both the original and super-resolution images to a unified latent space. Experiments on the developing Human Connectome Project (dHCP) dataset demonstrate a significant improvement in geometric accuracy over the leading-SR based cortical surface reconstruction methods, facilitating precise cortical surface reconstruction from 2D images. The code is open-sourced at: https://github.com/SCUT-Xinlab/CSR-from-2D-MRI.

Keywords: Cortical surface reconstruction · Multi-modal representation learning · MRI super-resolution.

1 Introduction

Cortical surface analysis plays a pivotal role in unraveling the intricacies of brain structure, providing crucial insights into neuroanatomy, connectivity, and the neural basis of cognition and disorders. The cortical surface is typically

reconstructed from high-quality 3D brain MRI data. However, obtaining high-resolution images in clinical settings proves challenging, especially when dealing with newborns who frequently exhibit head movements. Consequently, clinical data often consist of high in-plane resolution but low-resolution images with significant slice gaps in the through-plane direction due to these constraints.

To enhance the analysis of clinical data, the process of reconstructing cortical surfaces from LR clinical images typically involves two key steps. The first step is to downsample real high-resolution (HR) brain MRI, creating paired low-resolution (LR) and HR images. These LR-HR image pairs are then used to train MRI Super-Resolution methods [10, 11, 20] to reduce slice gaps and as a result enhancing the quality of the generated 3D images. For the second step, the predicted SR MRI can be used to reconstruct the cortical surface through traditional neuroimage processing pipelines, such as FreeSurfer [3–5], for adults or the dHCP pipeline [1, 6, 16] for newborns. However, processing large volumes of clinical data is time-intensive and challenging for these pipelines. Alternatively, learning-based approaches [2, 7, 12–14, 18] transform SR brain MRI from the first step into implicit surface representations or explicit deformation fields, significantly reducing the time required for surface reconstruction to seconds per subject.

Nevertheless, even though learning-based cortical surface reconstruction methods address the issue of inference speed, they still require high-quality SR 3D brain MRI as input to achieve accurate reconstruction. Yet, when facing significant slice gaps in LR brain MRI, state-of-the-art (SOTA) SR methods struggle to generate SR images with details matching real 3D brain MRI. This limitation hinders the model's ability to accurately capture cerebral cortex information.

To address this problem, we aim to further enhance the performance of cortical surface reconstruction from both image-wise and feature-wise perspectives. In addressing the image-wise aspect, our objective is to emphasize on cerebral cortex information in the generated SR brain MRI. To achieve this, we employ the segmentation mask as a topology guidance, directing the SR model to prioritize the cortical region. In addressing the feature-wise aspect, inspired by the benefits of multi-modal learning approaches [8, 9, 15, 17, 19] in achieving consistency across diverse data modalities in a high-dimensional feature space, we treat SR brain MRI and real HR brain MRI as distinct modalities. Employing a feature alignment strategy, our goal is to enable a cortical surface reconstruction model trained on HR brain MRI to effectively adapt to the features of SR brain MRI. This adaptation facilitates the accurate reconstruction of the cerebral cortex.

In this paper, we introduce a two-stage model for precise cortical surface reconstruction from 2D MRI, which consists of Segmentation Constrained MRI Super-Resolution (SCSR) and Feature Alignment Guided Cortical Surface Reconstruction (FAGCSR). The main contributions of this paper can be summarize as follows: 1) The proposed SCSR, based on the advanced SR method FSTNet, incorporates a segmentation model for the cortical ribbon as a constraint, enriching the predicted SR brain MRI with pertinent cortical information for improved surface reconstruction. 2) In the FAGCSR, we propose a mask-swap module with

a novel multi-process training strategy, gradually aligning SR brain MRI features with HR brain MRI in a unified latent space. During inference phase, the sole use of predicted SR suffices for generating high-quality cortical surface, enabling efficient leveraging of abundant LR clinical data and broadening the scope of cerebral cortex analyses. 3) Experiments on the publicly available dHCP dataset demonstrate the effectiveness of achieving precise cortical surface reconstruction from LR 2D brain MRI.

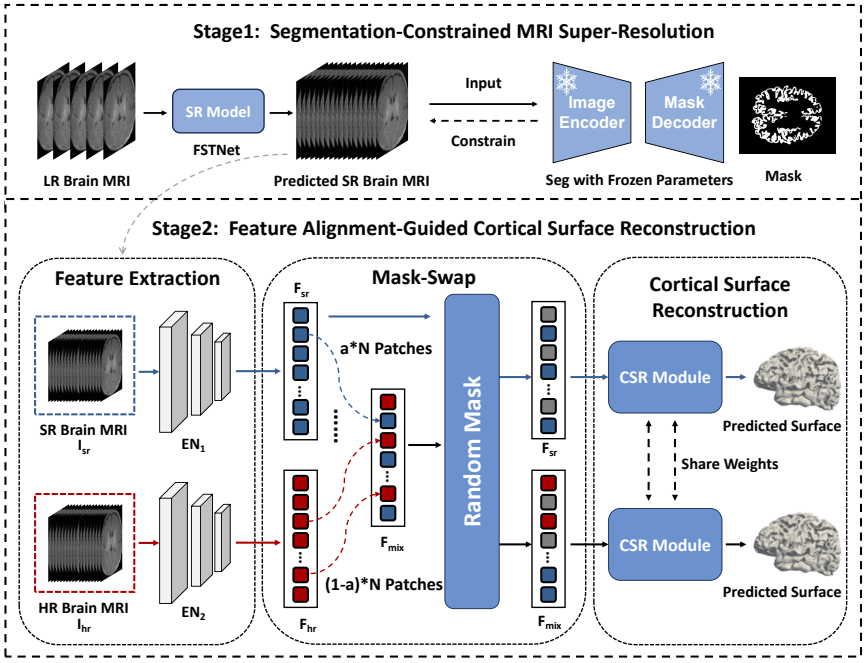

Fig. 1. The Architecture of the proposed method.

2 Method

2.1 Model Architecture

As illustrated in Fig. 1, we propose a two-stage model for cortical surface reconstruction from 2D LR brain MRI images. In stage 1, we aim to enhance surface reconstruction performance at the image level. Building upon a state-of-the-art MRI super-resolution model, we incorporate a cortical segmentation model as a constraint to guide the focus of the SR model on cortical regions. This ensures the generation of HR brain MRI with a more complete cortical surface topology. The objective of Stage 2 is to enhance the model performance at the feature

level. We introduce a dual-path cortical surface reconstruction model with a mask-swap module. It serves as a bridge to align features from paired SR 3D brain MRI and real 3D brain MRI into a shared latent space. This enables the decoder to accurately generate the surface from the SR features.

2.2 Segmentation Constrained MRI Super-Resolution

In this section, we present the details of Segmentation-Constrained MRI Super-Resolution (SCSR) employed for the generation of high-resolution cortical topology enhanced MRIs. Initially, two models are trained separately: an SR model, FSTNet, and a cortical ribbon segmentation model, Seg, with FSTNet recognized as one of the premier SR models. For FSTNet, given an input LR MRI, $\boldsymbol{I_{lr}} \in \mathbb{R}^{X\times Y\times Z}$, exhibiting high resolution in axial slices (XY) and lower resolution in the through-plane direction (Z), its role is to generate HR MRI, $\boldsymbol{I_{sr}} \in \mathbb{R}^{X\times Y\times kZ}$, with k denoting the up-sampling rate. Simultaneously, a 2D Unet is trained for the cortical ribbon segmentation using high-resolution axial slices and corresponding masks. Subsequently, we sequentially combine FSTNet and Seg with frozen parameters. The continued training of FSTNet under the constraint of Seg ensures a focused attention on the cortical regions, resulting in the generation of HR MRIs $\boldsymbol{I_{hr}}$ with a more comprehensive cortical topology.

During the isolated training of FSTNet, we use their default loss function [10,11], including MSE loss L_{MSE} and GAN loss L_G. The total loss is represented as $L_{SR} = L_{MSE} + \lambda_G L_G$. For the exclusive training of Seg, we only employ Dice loss L_{Dice}. In the joint training of FSTNet and Seg, we freeze the parameters of Seg, and the total loss is denoted as $L_1 = L_{SR} + \mu L_{Dice}$.

2.3 Feature Alignment Guided Cortical Surface Reconstruction

Mask-Swap Module To align the features of SR 3D brain MRI with those of real 3D brain MRI in a common latent space, we employ a dual-path design and extend the mask-swap module to brain data. Initially, the predicted SR MRI $\boldsymbol{I_{sr}}$ and real HR MRI are partitioned into non-overlapping patches of dimensions $N \times 8$, where $N = \frac{X}{2} \times \frac{Y}{2} \times \frac{Z}{2}$ represents the number of patches, and 8 signifies the embedding dimension of each patch. Following this, both $\boldsymbol{I_{sr}}$ and $\boldsymbol{I_{hr}}$ undergo encoding into the feature space using two Transformer encoders, EN_1 and EN_2, resulting in corresponding features of the same dimensions, namely $\boldsymbol{F_{sr}} = EN_1(\boldsymbol{I_{sr}})$ and $\boldsymbol{F_{hr}} = EN_2(\boldsymbol{I_{hr}})$. Utilizing the mask-swap module, we randomly select $N_1 = a * N$ tokens from $\boldsymbol{F_{sr}}$ and complementarily select $N_2 = (1-a)*N$ tokens from $\boldsymbol{F_{hr}}$. These tokens are then combined to form a mixed feature, $\boldsymbol{F_{mix}}$, where $a \in [0, 1]$ represents the weight corresponding to the number of tokens selected from $\boldsymbol{F_{sr}}$. Subsequently, $\boldsymbol{F_{mix}}$ is input into the decoder for surface reconstruction, effectively aligning $\boldsymbol{F_{sr}}$ and $\boldsymbol{F_{hr}}$ within the same feature space.

Multi-Process Training Strategy Next, we will present the details of aligning $\boldsymbol{F_{sr}}$ to $\boldsymbol{F_{hr}}$ through a multi-process training strategy with the mask-swap

module. Initially, with a set to 0, the model is trained solely on the features corresponding to the real HR image $\boldsymbol{F_{hr}}$ until convergence, enabling the reconstruction of high-quality cortical surfaces. Subsequently, the model undergoes iterative training, starting with the already converged state at $a = 0$, and proceeds with additional training cycles at $a = 0.25$, $a = 0.5$, $a = 0.75$, achieving convergence at each state. During this progression, the encoder gradually aligns $\boldsymbol{F_{sr}}$ with $\boldsymbol{F_{hr}}$, bringing them into a coherent latent space conducive to the generation of accurate surfaces. To prevent catastrophic forgetting, a dual-path training strategy is adopted instead of directly transitioning a from 0.75 to 1. Leveraging the model trained at $a = 0.75$, a is set to 0.5 to generate $\boldsymbol{F_{mix}}$, while retaining the features of $\boldsymbol{F_{sr}}$. Both sets of features are simultaneously input into the shared-decoder for cortical surface reconstruction. Throughout this process, the $\boldsymbol{F_{mix}}$ branch ensures alignment between the features of $\boldsymbol{F_{hr}}$ and $\boldsymbol{F_{sr}}$, while the $\boldsymbol{F_{sr}}$ branch ensures direct fitting of the predicted SR brain MRI. This multi-process training approach allows the proposed method, during the testing phase, to directly reconstruct accurate cortical surfaces from a single SR image (Fig. 1).

Cortical Surface Reconstruction This subsection introduces cortical surface reconstruction based on diffeomorphic transformation. The approach utilizes a Neural Ordinary Differential Equation (NODE) to model the reconstruction process. Following the CoTAN [12], we generate a time-varying velocity field (TVF) from 3D brain MRI to deform a template mesh to the cortical surface. Taking HR brain MRI $\boldsymbol{I_{hr}}$ as an example, the encoder initially maps $\boldsymbol{I_{hr}}$ to features $\boldsymbol{F_{hr}}$. Subsequently, the decoder transforms the feature into a flow field $U \in \mathbb{R}^{X \times Y \times Z \times 3}$, driving the deformation process. Given a template surface $S_0 \in N_v \times 3$, the definition of the deformation ODE and its solution over time $t \in [0, T]$ are computed as:

$$\frac{\partial}{\partial t}\phi_t = U,\ \phi_0 = \boldsymbol{S}_0, \tag{1}$$

$$\boldsymbol{S}_T = \phi_T\left(S_0\right) = \boldsymbol{S}_0 + \int_0^T \boldsymbol{U}\mathrm{d}s, \tag{2}$$

where ϕ_t represents the deformation trajectory of S_0 and $S_T = \phi_T(S_0)$ represents the final surface after deformation, which is the predicted cortical surface.

In practical applications, we employ the numerical solution to solve the ODE. For the corresponding time $t \in [0, T]$ and solution $\phi_T\left(S_0\right)$, we discretize the ODE using a fixed-step method. Taking Euler method as an example, the definition and solution are represented as $S_{n+1} = S_n + h\Delta S_n$, where $n \in [0, N]$. N is the number of total steps of iteration, S_N is the predicted cortical surface, ΔS_n denotes the flow field U representing the variation at each step, and $h = T/N$ is the step size.

For the training of surface reconstruction, following prior work [12], we incorporate Chamfer distance loss $\mathcal{L}_c$ to minimize the disparity between predicted and ground-truth surfaces, Laplacian loss $\mathcal{L}_{lap}$ for mesh smoothness and

normal consistency loss $\mathcal{L}_n$ to ensure consistent normals between the generated surface and the ground-truth. The total loss function is represented as $\mathcal{L} = \mathcal{L}_c + \lambda_{lap}\mathcal{L}_{lap} + \lambda_n\mathcal{L}_n$, where λ_{lap} and λ_n represent the weights assigned to the Laplacian loss and normal consistency loss, respectively.

3 Experiments

3.1 Setup

Dataset and Evaluation Metrics In our study, we employed 877 T2-weighted brain MRI images from the third release of the developing Human Connectome Project (dHCP). These scans are acquired from newborns with a PMA ranging from 26 to 45 weeks. For generating LR brain MRI, we down sample the real MRI by a rate of 4. We partitioned this dataset into training, validation, and test sets, with respective allocations of 70%, 10%, and 20%. We assess geometric accuracy using Chamfer Distance (CD), average symmetric surface distance (ASSD), and Hausdorff Distance (HD) as evaluation metrics. Calculation of these distances are performed on point clouds, each containing 100k uniformly sampled points derived from both the predicted and pseudo ground-truth surfaces.

Implementation Details Experiments were conducted using PyTorch 1.13.1 and an NVIDIA RTX 4090 GPU. For the SCSR training stage, we employ a λ_G value of 0.001 for independent SR training, conducting 20,000 iterations. Independent Seg training spans 200 epochs, while joint training involves setting μ to 0.5 and training for 20,000 iterations. In the FAGCSR stage, we set λ_{lap} to 0.5 and λ_n to 0.0005. Training includes 100 epochs for each a value (0, 0.25, 0.5, 0.75), followed by an additional 200 epochs in a dual-path manner with a set to 0.5. Both stages use the Adam optimizer with a learning rate of 0.0001. Throughout the experiments, focus was solely directed towards the left-hemi white matter surface.

3.2 Comparative Results

The conventional approach for reconstructing the cortical surface from LR 2D brain MRI initially employs SR methods to transform 2D images into 3D counterparts. Subsequently, the predicted SR 3D brain MRI undergoes processing using surface reconstruction methods designed for HR 3D brain MRI. To assess the effectiveness of our proposed method, we now compare it with several cortical surface extraction approaches, including Corticalflow [7], CFPP [18], and CoTAN [12]. We utilize the SR 3D brain MRI output by FSTNet as the input for retraining the surface reconstruction models with default settings. Table 1 demonstrates the superior performance of our proposed method on CD, surpassing the conventional approach by 16.38%, 13.10%, and 8.75% for Corticalflow, CFPP, and CoTAN, respectively. Additionally, our method outperforms in terms of ASSD by 13.83%, 9.86%, and 9.27%. Furthermore, the proposed method excels

Table 1. Comparative results of neonatal cortical surface reconstruction on the dHCP dataset. Methods employing SCSR demonstrate statistically significant improvements over their respective baseline versions (paired t-test $p < 0.05$).

Method	Input MRI	Left White Matter Surface Reconstruction		
		CD (mm)	ASSD (mm)	HD (mm)
Corticalflow	FSTNet	0.299 ± 0.068	0.159 ± 0.025	0.341 ± 0.073
	FSTNet-SCSR	0.292 ± 0.062	0.152 ± 0.023	0.324 ± 0.061
CFPP	FSTNet	0.290 ± 0.063	0.152 ± 0.024	0.324 ± 0.065
	FSTNet-SCSR	0.285 ± 0.062	0.147 ± 0.023	0.312 ± 0.060
CoTAN	FSTNet	0.274 ± 0.055	0.151 ± 0.022	0.324 ± 0.055
	FSTNet-SCSR	0.269 ± 0.066	0.144 ± 0.022	0.306 ± 0.056
Proposed	FSTNet	0.252 ± 0.045	0.139 ± 0.020	0.291 ± 0.046
	FSTNet-SCSR	$\mathbf{0.250 \pm 0.049}$	$\mathbf{0.137 \pm 0.020}$	$\mathbf{0.288 \pm 0.049}$
CFPP	Original 3D	0.238 ± 0.061	0.119 ± 0.021	0.246 ± 0.064

in HD with improvements of 15.54%, 11.11%, and 11.11%. The Fig. 2 illustrates the predicted white matter surfaces generated by these methods, accompanied by their respective error maps computed against the ground-truth. Our method demonstrates fewer errors and artifacts compared to alternative approaches. This observed enhancements indicate that our proposed method not only builds upon but further enhances the performance of one of the leading SR methods FSTNet in downstream cortical surface reconstruction tasks.

Table 2. Ablation Study on FAGCSR Stage. Methods employing FAGCSR demonstrate statistically significant improvements over their respective baseline versions (paired t-test $p < 0.05$).

SR Method	Mask	Swap	Left White Matter Surface Reconstruction		
			CD (mm)	ASSD (mm)	HD (mm)
Bicubic	✗	✗	0.400 ± 0.098	0.229 ± 0.034	0.523 ± 0.092
	✗	✓	0.366 ± 0.081	0.213 ± 0.030	0.484 ± 0.081
	✓	✓	0.364 ± 0.081	0.213 ± 0.031	0.486 ± 0.083
FSTNet	✗	✗	0.277 ± 0.065	0.150 ± 0.023	0.316 ± 0.054
	✗	✓	0.258 ± 0.051	0.142 ± 0.297	0.252 ± 0.045
	✓	✓	0.252 ± 0.045	0.139 ± 0.020	0.291 ± 0.046
FSTNet-SCSR	✗	✗	0.258 ± 0.046	0.142 ± 0.020	0.298 ± 0.049
	✗	✓	0.257 ± 0.049	0.139 ± 0.020	0.292 ± 0.050
	✓	✓	$\mathbf{0.255 \pm 0.049}$	$\mathbf{0.137 \pm 0.020}$	$\mathbf{0.288 \pm 0.049}$

3.3 Ablation Study

To further substantiate the effectiveness of our proposed approach, we conducted ablation experiments on the SCSR and feature alignment module. Using FSTNet

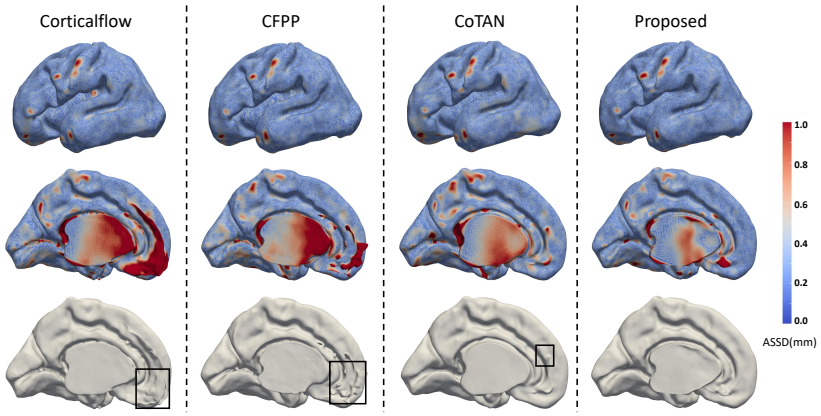

Fig. 2. Visualization of the error map computed from the predicted white matter surfaces and pseudo ground-truth of a single subject.

as the baseline for SR, we investigated the impact of the SCSR module on the cortical surface reconstruction model's performance. The models were retrained using SR 3D brains generated by FSTNet and FSTNet with SCSR as inputs for four surface reconstruction methods (Corticalflow [7], CFPP [18], CoTAN [12], and our proposed method). As detailed in Table 2, SCSR significantly enhances the geometric accuracy of the predicted surface, yielding average improvements of 1.7% on CD, 3.49% on ASSD, and 3.90% on HD compared to FSTNet without SCSR. The quantitative results demonstrate that SCSR effectively directs the attention of the SR model towards cortical ribbon, leading to improved performance in downstream surface reconstruction tasks.

Subsequently, we evaluate the impact of the FAGCSR stage on 3D brain images generated by three SR methods (Bicubic, FSTNet, FSTNet-SCSR) with real HR 3D brain images on model performance. Feature alignment is achieved via mask-swap module, utilizing swap for feature alignment and the mask for mapping features to a more generalized latent space. When employing both mask and swap, the geometric accuracy of the predicted cortical surface experiences significant improvement by 6.84%, 6.14%, and 6.33% on CD, ASSD, and HD, respectively. Upon removing the mask component, a slight dip in geometric accuracy was observed.

4 Conclusion

This study introduces a two-stage model for accurate cortical surface reconstruction from 2D MRI. The Segmentation-Constrained MRI Super-Resolution (SCSR) stage employs the topological structure of the cortical ribbon for guidance, generating SR brain MRI with crucial cerebral cortex information. The Feature Alignment-Guided Cortical Surface Reconstruction (FAGCSR) stage

aligns the SR brain MRI to the feature space of the corresponding real image. Building upon a leading SR model, proposed method further enhance the geometric accuracy at both image-wise and feature-wise levels, showcasing SOTA performance in the task of reconstructing cortical surfaces from 2D low-resolution brain MRI.

Acknowledgments. This work is supported by Key-Area Research and Development Program of Guangdong Province (2023B0303040001), Guangdong Basic and Applied Basic Research Foundation (2024A1515010180) and Guangdong Provincial Key Laboratory of Human Digital Twin (2022B1212010004).

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

References

1. Jelena Bozek, Antonios Makropoulos, Andreas Schuh, Sean Fitzgibbon, Robert Wright, Matthew F Glasser, Timothy S Coalson, Jonathan O'Muircheartaigh, Jana Hutter, Anthony N Price, et al. Construction of a neonatal cortical surface atlas using multimodal surface matching in the developing human connectome project. *NeuroImage*, 179:11–29, 2018.
2. Rodrigo Santa Cruz, Leo Lebrat, Pierrick Bourgeat, Clinton Fookes, Jurgen Fripp, and Olivier Salvado. Deepcsr: A 3d deep learning approach for cortical surface reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 806–815, 2021.
3. Anders M Dale, Bruce Fischl, and Martin I Sereno. Cortical surface-based analysis: I. segmentation and surface reconstruction. *Neuroimage*, 9(2):179–194, 1999.
4. Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.
5. Bruce Fischl, Martin I Sereno, and Anders M Dale. Cortical surface-based analysis: Ii: inflation, flattening, and a surface-based coordinate system. *Neuroimage*, 9(2):195–207, 1999.
6. Emer J Hughes, Tobias Winchman, Francesco Padormo, Rui Teixeira, Julia Wurie, Maryanne Sharma, Matthew Fox, Jana Hutter, Lucilio Cordero-Grande, Anthony N Price, et al. A dedicated neonatal brain imaging system. *Magnetic resonance in medicine*, 78(2):794–804, 2017.
7. Leo Lebrat, Rodrigo Santa Cruz, Frederic de Gournay, Darren Fu, Pierrick Bourgeat, Jurgen Fripp, Clinton Fookes, and Olivier Salvado. Corticalflow: a diffeomorphic mesh transformer network for cortical surface reconstruction. *Advances in Neural Information Processing Systems*, 34:29491–29505, 2021.
8. Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
9. Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022.
10. Zhiyang Lu, Zheng Li, Jun Wang, Jun Shi, and Dinggang Shen. Two-stage self-supervised cycle-consistency network for reconstruction of thin-slice mr images.

In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24*, pages 3–12. Springer, 2021.
11. Zhiyang Lu, Jian Wang, Zheng Li, Shihui Ying, Jun Wang, Jun Shi, and Dinggang Shen. Two-stage self-supervised cycle-consistency transformer network for reducing slice gap in mr images. *IEEE Journal of Biomedical and Health Informatics*, 2023.
12. Qiang Ma, Liu Li, Vanessa Kyriakopoulou, Joseph V Hajnal, Emma C Robinson, Bernhard Kainz, and Daniel Rueckert. Conditional temporal attention networks for neonatal cortical surface reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 312–322. Springer, 2023.
13. Qiang Ma, Liu Li, Emma C Robinson, Bernhard Kainz, Daniel Rueckert, and Amir Alansary. Cortexode: Learning cortical surface reconstruction by neural odes. *IEEE Transactions on Medical Imaging*, 42(2):430–443, 2022.
14. Qiang Ma, Emma C Robinson, Bernhard Kainz, Daniel Rueckert, and Amir Alansary. Pialnn: a fast deep learning framework for cortical pial surface reconstruction. In *Machine Learning in Clinical Neuroimaging: 4th International Workshop, MLCN 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 4*, pages 73–81. Springer, 2021.
15. Soumi Maiti, Yifan Peng, Takaaki Saeki, and Shinji Watanabe. Speechlmscore: Evaluating speech generation using speech language model. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
16. Antonios Makropoulos, Emma C Robinson, Andreas Schuh, Robert Wright, Sean Fitzgibbon, Jelena Bozek, Serena J Counsell, Johannes Steinweg, Katy Vecchiato, Jonathan Passerat-Palmbach, et al. The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. *Neuroimage*, 173:88–112, 2018.
17. Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
18. Rodrigo Santa Cruz, Léo Lebrat, Darren Fu, Pierrick Bourgeat, Jurgen Fripp, Clinton Fookes, and Olivier Salvado. Corticalflow++: Boosting cortical surface reconstruction accuracy, regularity, and interoperability. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 496–505. Springer, 2022.
19. Ziqiang Zhang, Sanyuan Chen, Long Zhou, Yu Wu, Shuo Ren, Shujie Liu, Zhuoyuan Yao, Xun Gong, Lirong Dai, Jinyu Li, et al. Speechlm: Enhanced speech pre-training with unpaired textual data. *arXiv preprint arXiv:2209.15329*, 2022.
20. Can Zhao, Blake E Dewey, Dzung L Pham, Peter A Calabresi, Daniel S Reich, and Jerry L Prince. 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.