This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

CortexGen: A Geometric Generative Framework for Realistic Cortical Surface Generation Using Latent Flow Matching

Yuanzhuo Zhu[1,3], Kehan Li[1,3], Jianhua Ma[1,3,4], Chunfeng Lian[2,3,4(✉)], and Fan Wang[1,3(✉)]

[1] Key Laboratory of Biomedical Information Engineering of Ministry of Education, School of Life Science and Technology, Xi'an Jiaotong University, Xi'an, China
fan.wang@xjtu.edu.cn
[2] School of Mathematics and Statistics, Xi'an Jiaotong University, Xi'an, China
chunfeng.lian@xjtu.edu.cn
[3] Research Center for Intelligent Medical Equipment and Devices, Xi'an Jiaotong University, Xi'an China
[4] Pazhou Lab (Huangpu), Guangzhou, China

Abstract. Geometric deep learning has shown great potential for cortical surface analysis, but its performance often depends on a large-scale training set of cortical surfaces, which are traditionally derived from MRI scans through complex and time-consuming preprocessing pipelines. Although deep learning-based surface reconstruction methods have streamlined this process, they still rely on MRI data, limiting the availability of training data. To address this, we propose **CortexGen**, a geometric generative framework that synthesizes highly realistic cortical surfaces without requiring MRI scans. CortexGen employs geometric variational encoders to map cortical surfaces into a latent space, where latent flow matching models efficiently learn the true data distribution. This enables a two-stage cortical surface synthesis process: first, deforming an icosahedron-discretized sphere into a coarse cortical surface, and second, refining it into a high-resolution surface. Experiments show that CortexGen generates diverse, realistic cortical surfaces with 163,842 vertices in just 1.4 seconds per surface. Using these synthetic surfaces as augmented training data significantly improved learning-based cortical surface parcellation in few-shot settings. Our code and pretrained models are available at https://github.com/ladderlab-xjtu/CortexGen.

Keywords: Cortical surface synthesis · Latent flow matching · Cortical surface parcellation · Data augmentation.

1 Introduction

Recent years have witnessed significant advancements in geometric deep learning for cortical surface computing [24], particularly with promising outcomes in the fundamental task of parcellation. To address the inherent challenges posed by

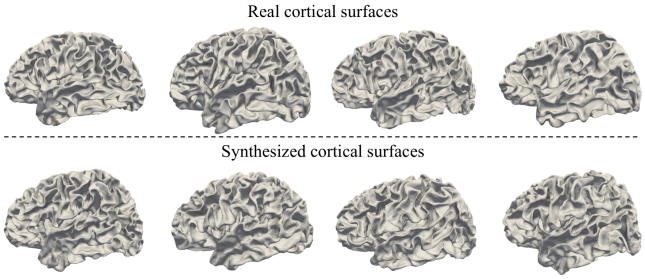

Fig. 1. We propose CortexGen, a geometric generative framework capable of synthesizing realistic high-resolution cortical surfaces without the need of MRI data. Shown here are representative examples of real (***top***) and synthesized cortical surfaces (***bottom***), respectively.

the irregularity data structure of cortical surfaces, substantial efforts [25,5,6,27] have focused on establishing dedicated network architectures for representation learning. However, beyond architecture designs, the availability of large-scale training data remains a critical factor in successfully training a deep model for cortical surface parcellation.

To get such training data, reconstructing the cortical surfaces from MRI scans remains the predominant approach. Traditional neuroimaging pipelines, such as FreeSurfer [3], involves a multi-step process, including image preprocessing, tissue segmentation, hemisphere separation, topology correction, and cortical surface reconstruction. This intricate process is time-intensive, typically requiring 2~3 hours per scan. Recent advancements have spurred the development of deep learning models for accelerated cortical surface reconstruction [9,18,13,26,1]. While have significantly streamlined the process, these methods have not addressed the issue of limited data source, as they still heavily depend on MRI scans as the primary input.

Generative learning may offer a promising alternative, i.e., synthesizing realistic cortical surfaces. Particularly, diffusion models [7,19] have emerged as a well-established paradigm in image synthesis [12,16,17], primarily due to their superior training stability compared to Generative Adversarial Networks (GANs) [4]. Although diffusion models have been applied to generate cortical curvature maps [22], their potential for synthesizing cortical surfaces remains unexplored.

In this paper, for the first time, we introduce a novel generative framework, termed **CortexGen**, for synthesizing highly realistic and fine-grained cortical surfaces (see Fig. 1). Our framework has two stages: 1) deforming an icosahedron-discretized sphere to generate a low-resolution cortical surface, and 2) refining it to obtain a high-resolution cortical surface. At each stage, we train a geometric variational autoencoder (GVAE) for cortical surface self-reconstruction, and employ flow matching [10,11], a powerful technique that enhances the training and inference efficiency of denoising models, to learn the synthesis of cortical surfaces in the latent space defined by the GVAE. These innovations collectively

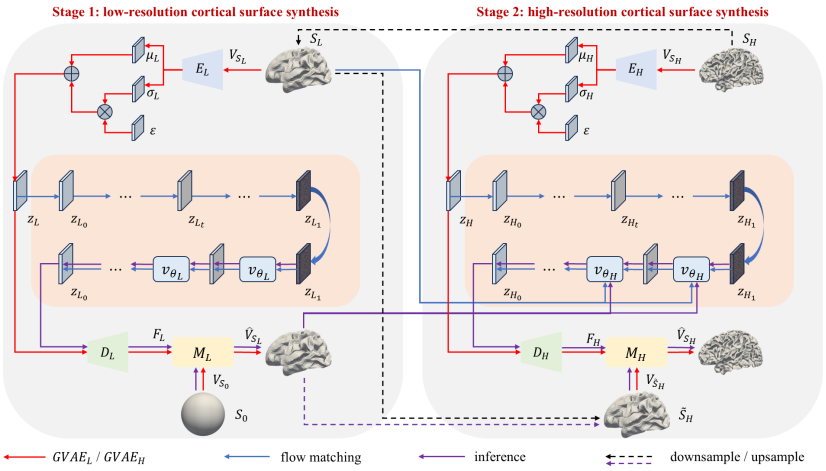

Fig. 2. The diagram of CortexGen, a two-stage generative framework based on latent flow matching for cortical surface synthesis. Arrows in different colors represent different workflows.

enable the *synthesis of a realistic cortical surface with 163,842 vertices in just 1.4 seconds* on an NVIDIA RTX3060 GPU. To further leverage the cortical surface synthesis capability of CortexGen, we apply it as a data augmentation method for few-shot cortical surface parcellation. Experimental results demonstrate that CortexGen led to substantial improvements of parcellation performance, which in turn justifies the high fidelity of synthesized cortical surfaces.

2 CortexGen

Our CortexGen is a geometric generative framework designed to synthesize realistic high-resolution cortical surfaces in the absence of MRI data. As illustrated in Fig. 2, CortexGen involves two stages: 1) synthesizing a low-resolution cortical surface by deforming an icosahedron-discretized sphere (a standard ico5 sphere in this study), and 2) synthesizing a high-resolution cortical surface (with the same vertical connectivity as a standard ico7 sphere) conditioned on the low-resolution initialization. Inspired by the powerful latent diffusion models [15], both stages are implemented in two steps: 1) training a GVAE for surface self-reconstruction, and 2) training a latent flow matching (LFM) model for high-quality generation in the GVAE hidden space.

2.1 Cortical Surface Self-Reconstruction

Given an icosahedron-reparameterized original cortical surface S_H, which shares the same number of vertices and connectivity as a standard ico7 sphere, its low-resolution counterpart S_L is derived by downsampling and smoothing, having

the same number of vertices and connectivity as a standard ico5 sphere. Two independent GVAEs are then trained for self-reconstruction of these respective surfaces. *Notably*, cortical surface self-reconstruction is performed by applying learnable deformation to an initial template surface, in line with common practices in current deep learning-based cortical surface reconstruction methods [9,18,13,26,1].

For the self-reconstruction of low-resolution cortical surface, we employ the convolution, pooling, and transposed convolution proposed by [25] to construct $GVAE_L$, which consists of an encoder E_L, a decoder D_L, and a deformation block M_L. The encoder E_L comprises three convolutional blocks, each incorporating a 1-ring convolution, batch normalization, and LeakyReLU activation, with pooling layers following the first two blocks. The decoder D_L, while maintaining the same three-block structure, differs in that the first two blocks are each followed by a transposed convolution.

The encoder E_L extracts a d-dimensional latent representation characterized by a distribution with mean $\mu_L \in \mathbb{R}^{642\times d}$ and variance $\sigma_L \in \mathbb{R}^{642\times d}$, such that $\mu_L, \sigma_L = E_L(V_{S_L})$, where $V_{S_L} \in \mathbb{R}^{10242\times 3}$ represents the 3D coordinates of all vertices on S_L. In the absence of MRI data and, consequently, the unavailability of the velocity field, the decoder D_L generates a N-channel feature map $F_L \in \mathbb{R}^{10242\times N}$, defined as $F_L = D_L(z_L)$, where $z_L = \mu_L + \sigma_L \cdot \varepsilon$, and $\varepsilon \sim \mathcal{N}(0,1)$. This feature map guides the deformation of the initial template spherical surface (a standard ico5 sphere, denoted as S_0) through the deformation block M_L.

Inspired by the successful application of neural ordinary differential equations (NODEs) [2] in cortical surface reconstruction [13,1], we model the deformation from S_0 to S_L as an ODE via a neural network, i.e., the deformation block M_L:

$$\frac{dV(t)}{dt} = M_L(V(t), F_L), V(0) = V_{S_0}, \tag{1}$$

where $V_{S_0} \in \mathbb{R}^{10242\times 3}$ contains the 3D coordinates of all vertices on S_0, $V(t)$ represents the coordinates at time $t \in [0,1]$, and the final coordinates $V(1)$ are expected to match V_{S_L}. To solve this ODE, we employ a fixed-step ODE solver, i.e., the forward Euler method, which can be expressed as:

$$V(x+1) = V(x) + h \cdot M_L(V(x), F_L), \tag{2}$$

where $x \in \{0, \cdots, X-1\}$, X is number of steps, $h = 1/X$ is the step size, and $M_L : \mathbb{R}^3 \times \mathbb{R}^N \to \mathbb{R}^3$ is implemented as a four-layer MLP.

We train $GVAE_L$ by minimizing a weighted sum of multiple loss function:

$$\mathcal{L}_{GVAE_L} = \mathcal{L}_1(\hat{V}_{S_L}, V_{S_L}) + \lambda_1\mathcal{L}_{lap}(\hat{V}_{S_L}) + \lambda_2\mathcal{L}_{nc}(\hat{V}_{S_L}) + \lambda_3\mathcal{L}_{KL}(\mu_L, \sigma_L), \tag{3}$$

where $\hat{V}_{S_L}$ represents the coordinates of the reconstructed low-resolution cortical surface, $\mathcal{L}_1$ is the mean absolute distance between corresponding vertices in $\hat{V}_{S_L}$ and V_{S_L}, $\mathcal{L}_{lap}$ is the mesh Laplacian loss that promotes smoothness of the reconstructed mesh, $\mathcal{L}_{nc}$ is the normal consistency loss that constrains the

cosine similarity between the normals of two adjacent faces, $\mathcal{L}_{KL}$ is the Kullback-Leibler divergence to regularize the discrepancy between the latent representation distribution and the standard normal distribution, and λ_1, λ_2, λ_3 are the corresponding weighting coefficients.

The $GVAE_H$ for high-resolution cortical surface self-reconstruction shares a similar architecture and optimization objective with $GVAE_L$, consisting of an encoder E_H, a decoder D_H, and a deformation block M_H, which are structurally identical to E_L, D_L, and M_L, respectively. However, $GVAE_H$ differs in two key aspects: 1) its input is $V_{S_H} \in \mathbb{R}^{163842\times 3}$, representing the 3D coordinates of all vertices on S_H, and 2) its initial surface, denoted as $\widetilde{S}_H$, is obtained by subdividing each triangular face on S_L twice, resulting in a mesh with the same vertex count and connectivity as S_H.

2.2 Latent Flow Matching

In the GVAE latent space, a deterministic continuous normalizing flow mapping the standard normal distribution to the latent distribution of cortical surfaces is trained via flow matching (FM) [10,11], a simulation-free method.

Based on the optimal transport theory [14,20], FM assumes that the transition from the source distribution (i.e., the standard normal distribution) to the target distribution (i.e., the latent representation distribution) follows a straight path between the two distributions, driven by a constant velocity field. Specifically, the training of denoising model for low-resolution cortical surface synthesis proceeds as follows. Given empirical observations of the target distribution $z_{L0} \sim \mathcal{N}(\mu_L, \sigma_L)$ and source distribution $z_{L1} \sim \mathcal{N}(0, 1)$, noise is added to z_{L0} according to:

$$z_{Lt} = tz_{L0} + (1-t)z_{L1}, \tag{4}$$

where z_{Lt} represents the linear interpolation between z_{L0} and z_{L1}, and time $t \in [0, 1]$ controls the degree of noise corruption. The velocity field is modeled by a neural network v_{θ_L} (i.e., the denoising model, implemented as a time-conditional Spherical U-Net [25] in this study) with parameter θ_L, which is optimized by minimizing:

$$\mathcal{L}_{v_{\theta_L}} = \mathbb{E}_{z_{Lt},t}\left[\|z_{L0} - z_{L1} - v_{\theta L}(z_{Lt}, t)\|_2^2\right]. \tag{5}$$

The training process for the denoising model in high-resolution cortical surface synthesis follows the same procedure as that for low-resolution cortical surface synthesis, with the exception that V_{S_L}, serving as conditional information, is passed to the denoising network v_{θ_H} along with z_{Ht} as:

$$\mathcal{L}_{v_{\theta_H}} = \mathbb{E}_{z_{Ht},t}\left[\|z_{H0} - z_{H1} - v_{\theta H}(z_{Ht}, V_{S_L}, t)\|_2^2\right]. \tag{6}$$

Notably, to bridge the gap between $\hat{V}_{S_L}$ and V_{S_L}, we refer to the practice in [21]. After training v_{θ_L} and v_{θ_H}, v_{θ_H} is finetuned on $\hat{V}_{S_L}$ as:

$$\mathcal{L}'_{v_{\theta_H}} = \mathbb{E}_{z_{Ht},t}\left[\left\|z_{H0} - z_{H1} - v_{\theta H}(z_{Ht}, \hat{V}_{S_L}, t)\right\|_2^2\right]. \tag{7}$$

2.3 Inference

Firstly, a low-resolution cortical surface is synthesized through the following steps: 1) the network v_{θ_L} denoises a sample drawn from the standard normal distribution to obtain a latent code; 2) this latent code is decoded by D_L to produce a feature map; 3) guided by this feature map, M_L deforms the sphere S_0 to generate a low-resolution cortical surface. *Next*, a high-resolution cortical surface is synthesized through the following steps: 1) conditioned on the synthesized low-resolution cortical surface, the finetuned v_{θ_H} denoises a sample drawn from the standard normal distribution to obtain a latent code; 2) this latent code is decoded by D_H to produce a feature map; 3) under the guidance of this feature map, M_H deforms the subdivided low-resolution cortical surface to produce the final high-resolution cortical surface.

3 Experiments

We first demonstrated the ability of CortexGen to synthesize realistic cortical surfaces by employing it as a data augmentation method for few-shot cortical surface parcellation. Then, we performed ablation studies on several key components of the framework to assess the effectiveness of our design.

3.1 Dataset

All experiments were conducted using the Baby Connectome Project (BCP) dataset [8], which contains 213 subjects and 417 icosahedron-reparameterized original cortical surface meshes. Each mesh corresponds to 36 regions of interest (ROIs) and has the same number of vertices and connectivity as a standard ico7 sphere. For preprocessing and data correction, we applied Laplace smoothing to all meshes by moving each vertex to the average coordinates of its adjacent vertices, and retained only those meshes without self-intersecting faces. The final dataset included 180 subjects and 327 meshes. The 3D vertex coordinates of all cortical surfaces were normalized to the range [-1, 1]. We focused only on the left hemisphere, given its similarity to the right hemisphere.

3.2 Experimental Setup

Following common practices in supervised deep learning, we split all cortical surfaces into training, validation, and test sets *at the subject level*, with the ratios of 60%, 20%, and 20%, respectively. We learned a deep learning-based cortical surface parcellation model on the training set, and evaluated its performance on the test set using the parameters that achieved the highest Dice coefficient on the validation set. This setup is referred to as "**Fully Supervised**".

In the context of *few-shot cortical surface parcellation*, where only a limited number of *labeled* cortical surfaces are available in the training set, we leveraged these labeled surfaces to generate pseudo-labels for the remaining unlabeled training surfaces, thus augmenting the dataset. Specifically, we first used

Table 1. Performance of Spherical U-Net and Cortex-Diffusion under different setups, quantified in terms of Dice Coefficient (%). The table also reports the mean improvement of "Real_aug" and "Gen_aug_N" over the baseline (i.e., "SD"), shown in parentheses after each item.

Setups	Baseline	Spherical U-Net	Cortex-Diffusion
SD	83.81 ± 1.93	-	-
Real_aug	-	83.85±1.80 (0.04±1.12)	83.66±1.76 (-0.15±1.15)
Gen_aug_200 (ours)	-	84.00±1.75 (0.19±0.81)	83.94±1.88 (0.13±1.01)
Gen_aug_500 (ours)	-	84.23±1.76 (0.42±0.59)	84.14±1.90 (0.33±0.74)
Gen_aug_1000 (ours)	-	84.33±1.83 (**0.52±0.56**)	84.35±1.92 (**0.54±0.64**)
Fully Supervised	-	89.06±2.95	88.96±2.95

FreeSurfer [3] to perform spherical mapping and calculate cortical attributes (i.e., mean curvature and sulcal depth) for all cortical surfaces in the training set. Then, we used Spherical Demons [23] to register each labeled spherical cortical surface to the unlabeled spherical cortical surfaces, obtaining the resulting warped label maps. For each vertex, the final pseudo-label is determined by the label with the maximum number of assigned labels. Finally, the combination of labeled and pseudo-labeled cortical surfaces formed the augmented training set. This setup is referred to as "**Real_aug**". In our implementation, the number of few-shot samples was set to 3 (i.e., 3-shot learning for parcellation).

A similar steup to "Real_aug", but with some differences, is defined by "**Gen_aug_N**", a data augmentation strategy based on our CortexGen. Specifically, we trained CortexGen on the training set and used it to synthesize **N** realistic cortical surfaces. The pseudo-labels for these surfaces were generated in the same manner as in "Real_aug". *Only the synthesized cortical surfaces* were used to form the augmented training set.

Finally, we defined a baseline termed "**SD**". In this setup, we registered the labeled cortical surfaces from the training set to all cortical surfaces in the test set. The pseudo-labels were then directly used as the parcellation predictions for the cortical surfaces in the test set.

We considered two deep learning methods for cortical surface parcellation: **1**) a spherical mapping-based model, i.e., **Spherical U-Net** [25], and **2**) an original cortical surface-based model, i.e., **Cortex-Diffusion** [27]. The performance of all models was quantitatively evaluated in terms of the **Dice** coefficient.

3.3 Results

Few-shot cortical surface parcellation. The results obtained by Spherical U-Net and Cortex-Diffusion under different setups, along with their improvements over the baseline setup, are summarized in Table 1, leading to two key observations. *First*, for both spherical surface-based and original surface-based models, although inferior to the "Fully Supervised" setup, "Gen_aug_N" consistently achieves greater performance gains than "Real_aug", demonstrating the broad applicability of "Gen_aug_N". In contrast, the performance gain from

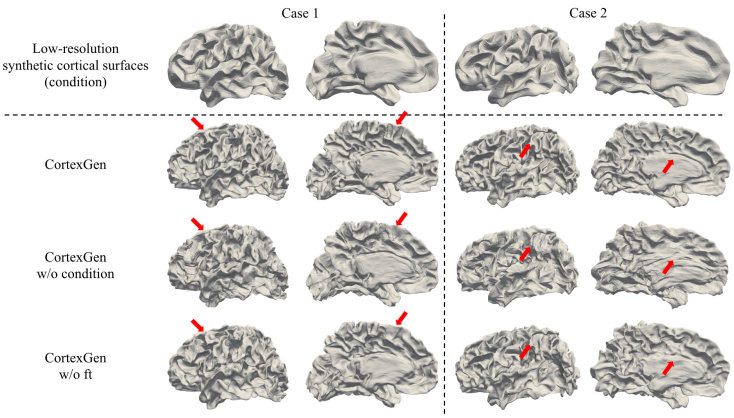

Fig. 3. Front and back views of high-resolution cortical surfaces (***bottom***) synthesized by different frameworks, conditioned on the same low-resolution synthetic cortical surfaces (***top***). Red arrows highlight the areas where CortexGen outperforms the other two variants in terms of quality.

"Real_aug" is sometimes even negative. *Second*, as N increases (i.e., as the number of available cortical surfaces for training increases), the performance gain from "Gen_aug_N" also increases. The superior performance of "Gen_aug_N" not only highlights the reliability of CortexGen as a data augmentation tool but also reflects the high quality of the synthesized cortical surfaces.

Ablation study on high-resolution cortical surface synthesis. In our framework, the high-resolution cortical surface is obtained by super-resolution of the synthesized low-resolution cortical surface, which involves two key factors: 1) the training of v_{θ_H} is conditioned on the low-resolution cortical surface; and 2) after the training of v_{θ_L} and v_{θ_H}, finetuning v_{θ_H} on the reconstructed low-resolution cortical surface is necessary to correct errors introduced during self-construction. To evaluate the effectiveness of these two components, we constructed two alternative frameworks: 1) **CortexGen w/o condition**, where v_{θ_H} only takes z_{Ht} and t as inputs to predict noise; and 2) **CortexGen w/o ft**, where v_{θ_H} directly predicts noise without fine-tuning.

As shown in Fig. 3, we visualized the representative high-resolution syntheses by these variants, given the same synthesized low-resolution conditions. A noticeable quality gap exists between the high-resolution cortical surfaces synthesized by the first two frameworks, highlighting the importance of using the low-resolution cortical surface as a condition when denoising with v_{θ_H}. Moreover, the visual advantage of CortexGen over CortexGen w/o ft indicates that finetuning significantly mitigates the reconstruction errors from $GVAE_L$, resulting in smoother cortical surfaces.

4 Conclusion

In this paper, we propose a novel generative framework, CortexGen, to synthesize realistic high-resolution cortical surface without the need for MRI data. By efficiently learning the data distribution of real cortical surfaces, CortexGen can generate diverse synthetic cortical surfaces, which can be used for data augmentation in the context of few-shot cortical surface parcellation.

Acknowledgments. This work was supported in part by Science and Technology Innovation 2030-Major Projects (No. 2022ZD0209000), NSFC Grant (Nos. 12326616), Natural Science Basic Research Program of Shaanxi (No. 2024JC-TBZC09), and Shaanxi Provincial Key Industrial Innovation Chain Project (No. 2024SF-ZDCYL-02-10).

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

References

1. Bongratz, F., Rickmann, A.M., Wachinger, C.: Neural deformation fields for template-based reconstruction of cortical surfaces from mri. Medical Image Analysis **93**, 103093 (2024)
2. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. Advances in neural information processing systems **31** (2018)
3. Fischl, B.: Freesurfer. Neuroimage **62**(2), 774–781 (2012)
4. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems **27** (2014)
5. Gopinath, K., Desrosiers, C., Lombaert, H.: Graph convolutions on spectral embeddings for cortical surface parcellation. Medical image analysis **54**, 297–305 (2019)
6. Ha, S., Lyu, I.: Spharm-net: Spherical harmonics-based convolution for cortical parcellation. IEEE Transactions on Medical Imaging **41**(10), 2739–2751 (2022)
7. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
8. 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)
9. Lebrat, L., Santa Cruz, R., de Gournay, F., Fu, D., Bourgeat, P., Fripp, J., Fookes, C., Salvado, O.: Corticalflow: a diffeomorphic mesh transformer network for cortical surface reconstruction. Advances in Neural Information Processing Systems **34**, 29491–29505 (2021)
10. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022)
11. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022)
12. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Repaint: Inpainting using denoising diffusion probabilistic models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11461–11471 (2022)

13. Ma, Q., Li, L., Robinson, E.C., Kainz, B., Rueckert, D., Alansary, A.: Cortexode: Learning cortical surface reconstruction by neural odes. IEEE Transactions on Medical Imaging **42**(2), 430–443 (2022)
14. McCann, R.J.: A convexity principle for interacting gases. Advances in mathematics **128**(1), 153–179 (1997)
15. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
16. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH 2022 conference proceedings. pp. 1–10 (2022)
17. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. IEEE transactions on pattern analysis and machine intelligence **45**(4), 4713–4726 (2022)
18. Santa Cruz, R., Lebrat, L., Fu, D., Bourgeat, P., Fripp, J., Fookes, C., Salvado, O.: Corticalflow++: boosting cortical surface reconstruction accuracy, regularity, and interoperability. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 496–505. Springer (2022)
19. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
20. Villani, C., et al.: Optimal transport: old and new, vol. 338. Springer (2008)
21. Weber, T., Ingrisch, M., Bischl, B., Rügamer, D.: Cascaded latent diffusion models for high-resolution chest x-ray synthesis. In: Pacific-Asia conference on knowledge discovery and data mining. pp. 180–191. Springer (2023)
22. Xie, Z., Dahan, S., Williams, L.Z., Cardoso, M.J., Robinson, E.C.: Cortical surface diffusion generative models. arXiv preprint arXiv:2402.04753 (2024)
23. Yeo, B.T., Sabuncu, M.R., Vercauteren, T., Ayache, N., Fischl, B., Golland, P.: Spherical demons: fast diffeomorphic landmark-free surface registration. IEEE transactions on medical imaging **29**(3), 650–668 (2009)
24. Zhao, F., Wu, Z., Li, G.: Deep learning in cortical surface-based neuroimage analysis: a systematic review. Intelligent Medicine **3**(1), 46–58 (2023)
25. Zhao, F., Xia, S., Wu, Z., Duan, D., Wang, L., Lin, W., Gilmore, J.H., Shen, D., Li, G.: Spherical u-net on cortical surfaces: methods and applications. In: Information Processing in Medical Imaging: 26th International Conference, IPMI 2019, Hong Kong, China, June 2–7, 2019, Proceedings 26. pp. 855–866. Springer (2019)
26. Zheng, H., Li, H., Fan, Y.: Coupled reconstruction of cortical surfaces by diffeomorphic mesh deformation. Advances in neural information processing systems **36**, 80608–80621 (2023)
27. Zhu, Y., Lian, C., Li, X., Wang, F., Ma, J.: Efficient cortical surface parcellation via full-band diffusion learning at individual space. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 162–172. Springer (2024)