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.

CortexAdapt3D: Parameter-Efficient Fine-Tuning of General 3D Foundation Models for Cortical Surface Analysis

Kehan Li[1,3], Chen Shen[3], Shuang Hao[1,3], Fan Wang[1,3], Jianhua Ma[1,3*], and Chunfeng Lian[2,3*]

[1] Key Laboratory of Biomedical Information Engineering of Ministry of Education, School of Life Science and Technology, Xi'an Jiaotong University, Xi'an, China
[2] School of Mathematics and Statistics, Xi'an Jiaotong University, Xi'an, China
[3] Research Center for Intelligent Medical Equipment and Devices (IMED), Xi'an Jiaotong University, Xi'an, China
`{jhma, chunfeng.lian}@xjtu.edu.cn`

Abstract. Cortical surface analysis is a specialized 3D shape modeling paradigm of fundamental importance in neuroscience. Despite the rapid advancement of pre-trained generic 3D models, their direct application to cortical data is challenged by the highly convoluted topology of the cortex and the presence of fine-scale anatomical features absent in generic 3D objects. To bridge this domain gap, we propose CortexAdapt3D, a parameter-efficient fine-tuning framework that adapts geometric priors from large-scale 3D pre-training to cortical surface analysis. First, to respect intrinsic manifold topology, we adopt a geodesic-aware patch sampling strategy on the cortical icosahedral (ICO) sphere. Second, we inject relative positional bias to capture the fine-grained local geometric dependencies inherent in highly folded cortical morphology. This enables effective fusion of geometric priors from large-scale generic 3D data with domain-specific cortical features. Finally, recognizing that cortical surfaces encode complementary neuroanatomical signals beyond spatial coordinates, we propose an Attribute-aware Spectral Modulation Adapter (ASMA). This spectral-domain adapter leverages attention to dynamically fuse generic 3D priors with domain-specific attributes like sulcal depth and curvature. Extensive experiments across classification, regression, and segmentation tasks demonstrate that CortexAdapt3D consistently outperforms both representative 3D fine-tuning strategies and dedicated cortex-specific architectures. Our findings demonstrate the effectiveness of cortical-specific adaptation for transferring generic 3D priors to the tested infant cortical analysis tasks, and suggest a promising direction for broader cortical representation learning. The code will be available at `https://github.com/ladderlab-xjtu/CortexAdapt3D`.

Keywords: Cortical Surface Analysis · 3D Pre-training · Parameter-Efficient Fine-tuning · Geometric Deep Learning

* Corresponding authors.

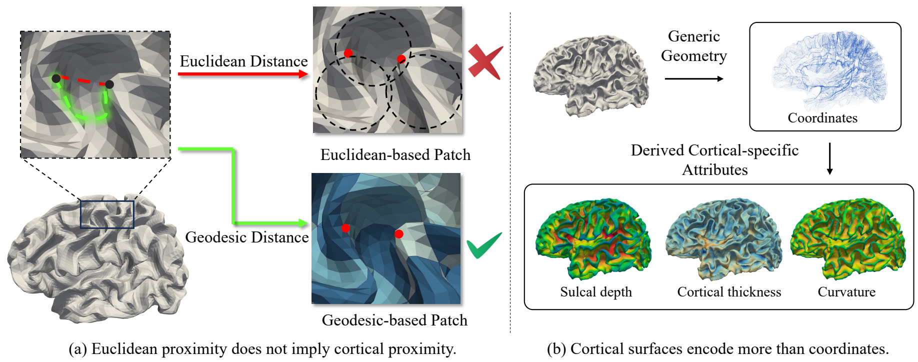

Fig. 1. (a) Euclidean proximity does not imply cortical proximity: Euclidean KNN can connect points across folds, whereas geodesic-based patches respect intrinsic cortical topology. (b) Cortical surfaces extend beyond generic geometry, encoding rich derived attributes (sulcal depth, thickness, curvature) that provide anatomically meaningful signals beyond coordinates.

1 Introduction

The cerebral cortex serves as the primary structural substrate for essential human functions, ranging from sensory perception and language to higher-level cognition [10]. In structural neuroimaging, the cortex is commonly modeled as a 2D manifold embedded in 3D space. Compared to traditional voxel-based representations, surface-based models retain detailed geometric information [18] that is intrinsically linked to cortical parcellation, neurodevelopmental trajectories, and neurodegenerative processes [15, 19, 4].

Current cortical surface analysis approaches [15, 19, 4, 22] are largely restricted by specialized architectures and limited dataset scales, which hinders their ability to learn truly generalizable representations. In contrast, the field of 3D self-supervised representation learning [11, 20, 1] has seen a surge in models pre-trained on large-scale, generic 3D datasets, yielding robust geometric priors with strong transferability. Recent advances in parameter-efficient fine-tuning (PEFT) [16, 21, 6], such as adapters and prompt-based tuning, allow these generic priors to be transferred to downstream tasks without the computational burden of full-parameter updates [8, 13].

Despite the potential of 3D foundation models, a substantial domain gap exists between generic 3D object modeling and cortical surface analysis. This gap stems from the cortex's highly convoluted folding patterns, where Euclidean proximity misrepresents intrinsic connectivity, and its rich set of neuroanatomical attributes (e.g., thickness, sulcal depth) that carry complementary information to spatial coordinates. Generic PEFT methods fail to address these unique characteristics, limiting the effective transfer of pre-trained knowledge.

To tackle these challenges, we introduce **CortexAdapt3D**, a fine-tuning framework designed to adapt large-scale 3D geometric priors to the intricacies of cortical morphology. Our approach is built upon three core strategies: *First*, following the topology-preserving design philosophy of cortical Transformer models [3, 2], we construct triangular patches directly on an icosahedral (ICO) sphere to address the discrepancy between Euclidean and intrinsic geometry on highly folded surfaces (Fig. 1(a)). This geodesic-aware patch construction preserves intrinsic cortical topology, ensuring neighborhoods are defined by cortical proximity rather than misleading Euclidean closeness [12]. *Second*, recognizing that absolute positional encodings [11, 20, 1, 17] are suboptimal for capturing fine-scale structural dependencies, we inject learnable relative positional biases into the pre-trained Transformer's attention mechanism [7]. This enhances the model's capacity to represent local geometric relationships critical to cortical folding patterns. *Most critically*, as shown in Fig. 1(b), cortical surfaces encode rich neuroanatomical attributes (e.g., sulcal depth, curvature, thickness) that provide complementary information to spatial coordinates. To jointly model geometry and these attributes, we propose an Attribute-aware Spectral Modulation Adapter (ASMA). Building upon spectral-domain adaptation [6], ASMA employs cross-attention to modulate geometric features in the spectral domain using attribute-derived signals, allowing anatomically meaningful attributes to guide the refinement of geometric representations.

Extensive evaluations on tasks including preterm classification, cortical parcellation, and age regression demonstrate that CortexAdapt3D consistently outperforms both specialized cortical models and state-of-the-art 3D fine-tuning strategies. Our results suggest that properly adapted 3D foundation models provide an effective path for the evaluated infant cortical tasks and a promising direction for broader cortical representation learning.

2 Methods

2.1 Overview

Given a cortical surface, we represent it as a triangular mesh embedded in $\mathbb{R}^3$. Unlike generic 3D object modeling, cortical surfaces exhibit highly convoluted folding patterns and rich neuroanatomical attributes (Fig. 1). To effectively adapt geometric priors learned from large-scale generic 3D pre-training, we propose **CortexAdapt3D**, a parameter-efficient fine-tuning framework tailored to cortical surface analysis (Fig. 2).

The framework first partitions the cortical surface into triangular patches defined on an ICO sphere. Patch coordinates are processed by a frozen patch encoder. The resulting patch embeddings, together with cortical-specific attributes and the computed eigen basis, are then separately fed into the Transformer Backbone. Within each Transformer layer, we inject relative positional bias into the pretrained attention module to better capture local geometric dependencies, and insert an ASMA to incorporate cortical attributes through spectral-domain attention. During fine-tuning, only the relative positional bias and ASMA modules, along with the task head, are trainable.

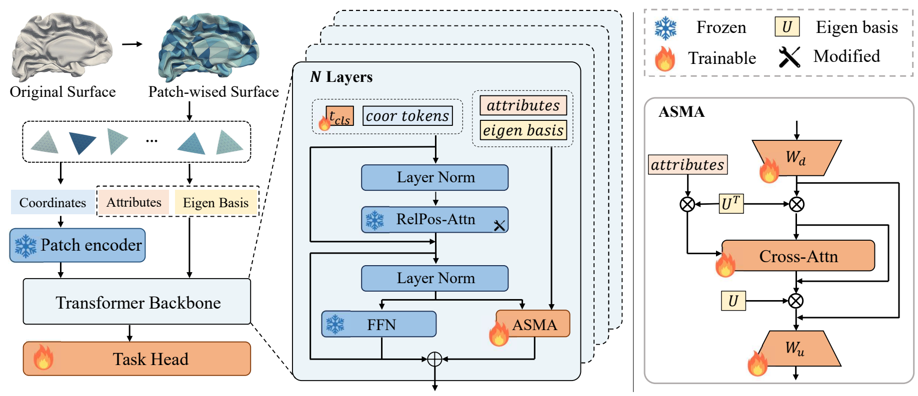

Fig. 2. Overview of the Adaptation Framework. The cortical surface is partitioned into triangular patches on the ICO sphere and processed by a frozen patch encoder and Transformer layers. Within each layer, we inject relative positional bias into the pretrained attention and incorporate an Attribute-aware Spectral Modulation Adapter (ASMA) for efficient fine-tuning.

2.2 Geodesic-aware Patch Construction

Standard point-based 3D models typically construct local neighborhoods using Farthest Point Sampling (FPS) followed by Euclidean k-Nearest Neighbor (kNN) grouping [12]. However, as illustrated in Fig. 1(a), Euclidean proximity does not necessarily correspond to cortical proximity on highly folded surfaces. Points that are close in $\mathbb{R}^3$ may be geodesically distant across sulcal folds.

To preserve intrinsic topology, we construct triangular patches directly on an ICO sphere. Specifically, all cortical surfaces are remeshed to the ICO5 topology with 10,242 vertices and 20,480 faces. We construct 320 geodesic triangular patches on the ICO sphere, where each patch contains 45 vertices. The fixed connectivity is used for topology-preserving patch assignment. Patch representations are constructed from vertex coordinates and cortical attributes, serving as Transformer inputs.

2.3 Relative Positional Bias Injection

Most 3D pre-trained models use absolute positional encodings derived from patch centers [11, 20, 1, 17], which are effective for masked reconstruction but less appropriate for highly folded cortical surfaces requiring fine-grained local geometric modeling. To enhance local structural modeling, we inject a learnable relative positional bias (RPB) into the pretrained self-attention by augmenting the attention logits with a term B_{ij}:

$$Attention(Q, K, V) = Softmax\left(\frac{QK^\top}{\sqrt{d}} + B\right)V$$

Table 1. Preterm classification results on dHCP dataset (%). Classification results are reported as point estimates on the predefined test split.

Setting	Method	ACC	Precision	Recall	AUC	SPE
Fully Supervised (from scratch)	SphericalCNN	71.03	62.49	65.74	63.17	75.20
	AGConv	87.59	84.21	77.48	91.21	95.58
Full Fine-tuning	PointMAE	87.59	84.21	77.48	85.79	95.58
	PointDif	83.45	75.93	77.06	87.14	88.50
PEFT (PointDif)	IDPT	88.28	85.00	79.04	89.49	95.58
	DAPT	81.38	73.12	74.61	85.09	86.73
	PointGST	84.82	77.85	79.07	89.82	89.38
	Ours	**91.03**	**89.83**	**83.05**	**91.76**	**97.35**
PEFT (PointMAE)	IDPT	88.28	83.62	81.28	**92.06**	93.81
	DAPT	81.38	73.49	76.85	83.55	84.96
	PointGST	86.90	82.58	77.03	89.82	94.69
	Ours	**91.72**	**89.44**	**85.73**	91.87	**96.46**

Here, B_{ij} encodes the relative positional relationship between patch i and patch j. By explicitly incorporating pairwise relative geometry alongside absolute positional information, the attention mechanism better captures local structural dependencies intrinsic to cortical topology.

2.4 Attribute-aware Spectral Modulation Adapter (ASMA)

As shown in Fig. 1(b), cortical surfaces encode substantially more than geometric coordinates. In addition to spatial shape, they provide derived attributes that reflect morphological organization and structural variation. To effectively integrate these cortical-specific signals, we introduce the ASMA module. ASMA follows the spectral adaptation paradigm of PointGST [6], but tailors it to cortical surfaces by removing the Euclidean local-basis branch and introducing cortical-attribute-conditioned cross-attention.

Let $H \in \mathbb{R}^{N\times d}$ denote the token representations from a Transformer layer. The features are first reduced in dimensionality through a learnable linear transformation: $H_d = HW_d$, where $W_d \in \mathbb{R}^{d\times d^{'}}$ is a trainable linear projection with $d^{'} < d$. We then project features into the spectral domain using the eigen basis of the graph Laplacian: $\tilde{H} = U^{\top} H_d$, where $U \in \mathbb{R}^{N\times N}$ is the orthonormal eigenvector matrix, satisfying $L = U\Lambda U^{\top}$. This spectral projection enhances feature discriminability by providing a frequency-domain representation[6]. Cortical attributes $A \in \mathbb{R}^{N\times d_a}$ are similarly projected and linearly transformed, where d_a is the adapter embedding dimension. In the spectral domain, we apply cross-attention to modulate geometric features:

$$Z = CrossAttn(\tilde{H}, \tilde{A}),$$

where geometry features serve as queries and attribute features as keys and values. This mechanism allows anatomically meaningful attributes to dynamically

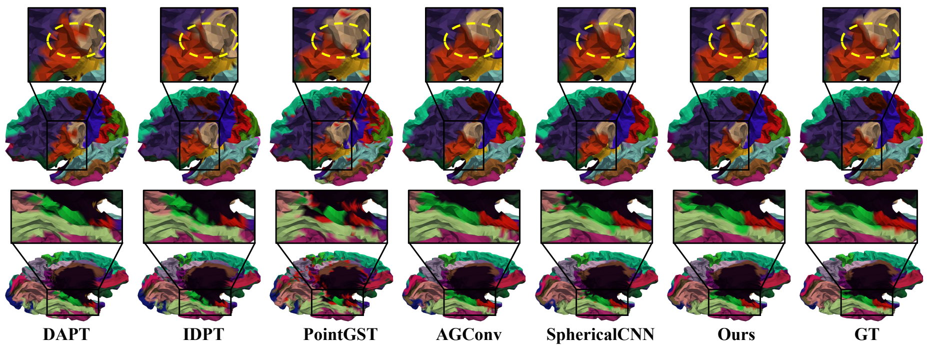

Fig. 3. Qualitative comparison of cortical parcellation results on the BCP dataset. Highlighted regions (yellow dashed circles and zoom-in views) illustrate that our method produces segmentations that are more consistent with the ground truth, particularly along intricate sulcal boundaries and highly folded regions, demonstrating improved structural continuity and boundary delineation.

refine geometric representations. Before transforming back to the spatial domain, a residual connection is applied in the spectral space to preserve the original frequency components. The modulated features are subsequently projected back to the spatial domain: $H^{'} = UZ$, followed by another residual connection with the original spatial features. Finally, the refined representation is up-projected by W_u to restore the feature dimension and produce the adapter output.

3 Experiments

3.1 Dataset and comparison methods

Dataset. We evaluate CortexAdapt3D on two publicly available infant cortical surface datasets: dHCP[9] and BCP[5]. The Developing Human Connectome Project (dHCP) provides cortical surfaces across a wide range of postmenstrual ages (PMA). Following prior works, we use dHCP for two tasks: preterm classification and PMA age regression. The Baby Connectome Project (BCP) provides infant cortical surfaces with anatomical parcellation labels, which we use for cortical parcellation (segmentation). All surfaces are remeshed onto a standardized icosahedral (ICO) sphere, and cortical attributes (sulcal depth, thickness, curvature) are computed for downstream modeling.
Comparison Methods. We compare CortexAdapt3D against three categories of baselines. **(1) Fully supervised training from scratch.** We include representative models trained entirely on target datasets without pre-training: SphericalCNN[19], a widely adopted spherical convolutional network for cortical surfaces, and AGConv[14], a representative geometry-aware convolutional model for point clouds. **(2) Parameter-efficient fine-tuning (PEFT).** We evaluate state-of-the-art 3D point cloud adaptation methods, including IDPT[16],

Table 2. Cortical parcellation on BCP (%) and PMA age regression on dHCP (week). Results are reported as mean ± standard deviation across test samples.

Setting	Method	Parcellation (BCP)		Regression (dHCP)
		ACC ↑	Dice ↑	MAE ↓
Fully Supervised (from scratch)	SphericalCNN	91.63 ± 4.12	87.26 ± 4.12	1.16 ± 0.92
	AGConv	90.88 ± 3.50	88.45 ± 2.94	0.78 ± 0.58
Full Fine-tuning	PointMAE	79.84 ± 2.71	75.01 ± 2.34	0.86 ± 0.60
	PointDif	80.33 ± 2.46	75.99 ± 2.19	0.85 ± 0.61
PEFT (PointDif)	IDPT	79.57 ± 2.37	75.31 ± 2.16	0.76 ± 0.60
	DAPT	79.27 ± 2.55	74.82 ± 2.24	0.75 ± 0.56
	PointGST	71.14 ± 2.42	68.52 ± 2.17	0.87 ± 0.64
	Ours	$\mathbf{91.91 \pm 3.96}$	$\mathbf{89.90 \pm 3.65}$	$\mathbf{0.74 \pm 0.55}$
PEFT (PointMAE)	IDPT	80.23 ± 2.63	75.91 ± 2.29	0.77 ± 0.56
	DAPT	79.65 ± 2.15	74.83 ± 1.97	0.76 ± 0.57
	PointGST	70.06 ± 2.57	65.16 ± 2.20	0.78 ± 0.67
	Ours	$\mathbf{91.92 \pm 3.90}$	$\mathbf{89.96 \pm 3.79}$	$\mathbf{0.76 \pm 0.57}$

DAPT[21], and PointGST[6], built upon pretrained backbones (i.e., Point-MAE[11] and PointDif[20]). **(3) Full fine-tuning of pretrained models.** For completeness, we also report results of full-parameter fine-tuning, where all backbone weights are updated. Our method is implemented on the same pretrained backbones for fair comparison.

3.2 Results

Preterm Classification on dHCP. As shown in Table 1, CortexAdapt3D consistently achieves the best performance across most evaluation metrics. It can be observed that both full fine-tuning and parameter-efficient adaptation of 3D pre-trained models yield competitive performance on this coarse-grained classification task, highlighting the potential of transferring generic 3D knowledge to cortical surface analysis. Building upon this foundation, our cortical-specific adaptation design further enhances representation learning and surpasses all competing methods, demonstrating the advantage of targeted fine-tuning for cortical geometry and attributes.

Cortical Parcellation on BCP. Table 2 shows that CortexAdapt3D consistently outperforms both generic 3D pre-training baselines and existing parameter-efficient tuning methods. Notably, generic PEFT strategies exhibit limited effectiveness on fine-grained cortical parcellation, where accurate modeling of highly folded sulcal boundaries is required. With our cortical-specific adaptation design, CortexAdapt3D not only substantially improves over these generic approaches but also outperforms specialized cortical architectures (e.g., SphericalCNN), highlighting the strong transferability of generic 3D priors when properly adapted. Figure 3 further provides qualitative comparisons. Our method produces segmentations that are more consistent with the ground truth, particularly along complex sulcal boundaries and highly folded cortical regions, demonstrating improved structural continuity and boundary delineation.

Table 3. Ablation study on dHCP preterm classification (%). The base model is a PointGST-like unconditioned spectral adapter. We incrementally enable geodesic patch construction (*G-patch*), relative positional bias (*RPB*), and attribute conditioning within the spectral adapter (*Attr*).

G-patch	RPB	Attr	ACC	Precision	Recall	AUC	SPE
			84.83	77.82	80.19	88.30	88.50
	✓	✓	88.97	85.01	81.72	91.68	94.69
✓		✓	85.52	79.31	77.27	89.55	92.04
✓	✓		86.21	82.54	74.35	87.53	95.58
✓	✓	✓	**91.03**	**89.83**	**83.05**	**91.76**	**97.35**

PMA Age Regression on dHCP. For age regression (Table 2), CortexAdapt3D achieves the lowest MAE among all compared methods. Notably, while SphericalCNN achieves strong performance on the surface-specific parcellation task, its regression performance is substantially weaker. This observation highlights a key limitation of highly specialized surface architectures: although effective for task-specific segmentation, they may struggle to transfer across different predictive settings.

Overall, the consistent improvement of CortexAdapt3D across classification, parcellation, and regression tasks demonstrates that attribute-aware spectral modulation enhances geometric representations in a task-agnostic manner. These results indicate that properly adapted generic 3D priors provide stronger cross-task generalization than purely surface-specific designs.

3.3 Ablation Studies

We conduct ablation experiments on the dHCP preterm classification task to quantify the contribution of each proposed component: geodesic-aware triangular patching (*G-patch*), relative positional bias injected into the pretrained attention (*RPB*), and *attribute conditioning* within our spectral adapter (*Attr*). Results are reported in Table 3.

Our *Base* setting is not the vanilla PointDif backbone; instead, it corresponds to the spectral-domain adapter fine-tuning framework **without** incorporating cortical attributes (i.e., an unconditioned spectral adapter). Starting from this *Base*, introducing *RPB* together with *Attr* yields a consistent improvement, validating the importance of enhancing local geometric dependency modeling and leveraging cortical-specific signals for feature refinement. Using *G-patch* construction alone provides limited improvement when combined with only one cue, indicating that topology-preserving neighborhoods are most effective when paired with complementary cortical-specific modeling. Notably, the full model achieves the best overall performance, demonstrating that these designs are complementary and jointly bridge the domain gap between generic 3D pre-training and cortical surface analysis.

4 Conclusion

This paper introduces CortexAdapt3D, a parameter-efficient framework for adapting generic 3D pre-trained models to cortical surface analysis. By integrating topology-aware ICO patching, relative positional bias, and ASMA, CortexAdapt3D reduces the domain gap between generic 3D object modeling and highly convoluted cortical manifolds. Experiments on classification, parcellation, and regression show consistent gains over training from scratch, full fine-tuning, and existing PEFT methods, highlighting the value of cortical-specific adaptation for infant cortical representation learning. Extending this framework to adult and pathological cortical analysis remains an important direction for future work.

Acknowledgments. This work was supported in part by Brain Science and Brain-like Intelligence Technology - National Science and Technology Major Project (2022ZD0209000), National Natural Science Foundation of China (Nos. T2522028 and 12326616), Natural Science Basic Research Program of Shaanxi (No. 2024JC-TBZC-09), 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. Chen, G., Wang, M., Yang, Y., Yu, K., Yuan, L., Yue, Y.: Pointgpt: Autoregressively generative pre-training from point clouds. Advances in Neural Information Processing Systems **36**, 29667–29679 (2023)
2. Cheng, J., Zhao, F., Wu, Z., Yuan, X., Wang, L., Gilmore, J.H., Lin, W., Zhang, X., Li, G.: Stf: A spherical transformer for versatile cortical surfaces applications. NeuroImage p. 121370 (2025)
3. Dahan, S., Fawaz, A., Williams, L.Z., Yang, C., Coalson, T.S., Glasser, M.F., Edwards, A.D., Rueckert, D., Robinson, E.C.: Surface vision transformers: Attention-based modelling applied to cortical analysis. In: International Conference on Medical Imaging with Deep Learning. pp. 282–303. PMLR (2022)
4. Ha, S., Lyu, I.: Spharm-net: Spherical harmonics-based convolution for cortical parcellation. IEEE Transactions on Medical Imaging **41**(10), 2739–2751 (2022)
5. 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)
6. Liang, D., Feng, T., Zhou, X., Zhang, Y., Zou, Z., Bai, X.: Parameter-efficient fine-tuning in spectral domain for point cloud learning. IEEE Transactions on Pattern Analysis and Machine Intelligence (2025)
7. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)

8. Lv, B., Zha, Y., Dai, T., Yuerong, X., Chen, K., Xia, S.T.: Adapting pre-trained 3d models for point cloud video understanding via cross-frame spatio-temporal perception. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12413–12422 (2025)
9. Makropoulos, A., Robinson, E.C., Schuh, A., Wright, R., Fitzgibbon, S., Bozek, J., Counsell, S.J., Steinweg, J., Vecchiato, K., Passerat-Palmbach, J., et al.: The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. Neuroimage **173**, 88–112 (2018)
10. Molnár, Z., Clowry, G.J., Šestan, N., Alzu'bi, A., Bakken, T., Hevner, R.F., Hüppi, P.S., Kostović, I., Rakic, P., Anton, E., et al.: New insights into the development of the human cerebral cortex. Journal of anatomy **235**(3), 432–451 (2019)
11. Pang, Y., Tay, E.H.F., Yuan, L., Chen, Z.: Masked autoencoders for 3d point cloud self-supervised learning. World Scientific Annual Review of Artificial Intelligence **1**, 2440001 (2023)
12. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. Advances in neural information processing systems **30** (2017)
13. Wang, M., Wang, H., Li, Y., Kong, X., Du, J., Shen, G., Xia, F.: Trackany3d: Transferring pretrained 3d models for category-unified 3d point cloud tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 28249–28259 (2025)
14. Wei, M., Wei, Z., Zhou, H., Hu, F., Si, H., Chen, Z., Zhu, Z., Qiu, J., Yan, X., Guo, Y., et al.: Agconv: Adaptive graph convolution on 3d point clouds. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(8), 9374–9392 (2023)
15. Xue, C., Wang, F., Zhu, Y., Li, H., Meng, D., Shen, D., Lian, C.: Neuro explainer: Fine-grained attention decoding to uncover cortical development patterns of preterm infants. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 202–211. Springer (2023)
16. Zha, Y., Wang, J., Dai, T., Chen, B., Wang, Z., Xia, S.T.: Instance-aware dynamic prompt tuning for pre-trained point cloud models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14161–14170 (2023)
17. Zhang, R., Guo, Z., Gao, P., Fang, R., Zhao, B., Wang, D., Qiao, Y., Li, H.: Point-m2ae: multi-scale masked autoencoders for hierarchical point cloud pre-training. Advances in neural information processing systems **35**, 27061–27074 (2022)
18. Zhao, F., Wu, Z., Li, G.: Deep learning in cortical surface-based neuroimage analysis: a systematic review. Intelligent Medicine **3**(1), 46–58 (2023)
19. Zhao, F., Wu, Z., Wang, L., Lin, W., Gilmore, J.H., Xia, S., Shen, D., Li, G.: Spherical deformable u-net: Application to cortical surface parcellation and development prediction. IEEE transactions on medical imaging **40**(4), 1217–1228 (2021)
20. Zheng, X., Huang, X., Mei, G., Hou, Y., Lyu, Z., Dai, B., Ouyang, W., Gong, Y.: Point cloud pre-training with diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22935–22945 (2024)
21. Zhou, X., Liang, D., Xu, W., Zhu, X., Xu, Y., Zou, Z., Bai, X.: Dynamic adapter meets prompt tuning: Parameter-efficient transfer learning for point cloud analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14707–14717 (2024)
22. 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)