SPHARM-Mamba: Rotation-Invariant Multiscale Modeling for Brain Age Prediction

Junho Choi[1][0009-0002-3144-2445], Mingyu Joo[2][0009-0003-7172-9253], Jiwon Son[1][0009-0005-8897-0945], Won Hwa Kim[⋆1,2][0000-0001-5393-0883], and Ilwoo Lyu[⋆1,2][0000-0001-5868-9603]

[1] Graduate School of Artificial Intelligence, POSTECH, Pohang, South Korea
[2] Department of Computer Science and Engineering, POSTECH, Pohang, South Korea
{wonhwa,ilwoolyu}@postech.ac.kr

Abstract. Accurate brain age prediction is essential for understanding neurodevelopmental trajectories and detecting abnormal aging patterns. Morphological features derived from cortical shape provide informative structural representations for brain age prediction. Recent deep learning approaches for surface analysis primarily rely on local aggregation mechanisms. Transformer- and Mamba-based architectures model long-range dependencies through patch partitioning combined with self-attention or selective scan. However, these strategies introduce high computational cost, uncertainty at patch boundaries, and sensitivity to surface rotations that necessitate anatomical registration or extensive data augmentation. To address these limitations, we propose SPHARM-Mamba, a spherical harmonics-based generic backbone for genus-zero surface data that eliminates patch partitioning and avoids surface registration. To this end, dense cortical signals are projected onto spherical harmonics to construct compact rotation-invariant descriptors naturally ordered from global to local scales. Mamba is then employed to model dependencies across harmonic degrees and capture multiscale interactions efficiently. Experiments on cortical age prediction demonstrate superior performance over existing geometric and patch-based models with substantially fewer parameters. The software is available at https://github.com/Shape-Lab/SPHARM-Mamba.

Keywords: Cortical surface analysis · Brain age prediction · Spherical harmonics · Mamba · Geometric deep learning

1 Introduction

Brain age prediction provides a valuable biomarker for quantifying neurobiological maturation across the lifespan and detecting abnormal developmental trajectories [1,4,5,17]. Estimating brain age from cortical surfaces employs morphological features to enable effective age prediction [19,28]. This often requires effec-

⋆ W. Kim and I. Lyu - co-corresponding authors

tively capturing geometric patterns on complex manifolds while preserving large-scale cortical organization and local morphological detail. Early convolution-based methods demonstrated strong capability in learning local cortical structures. To extend their receptive fields, later approaches introduced hierarchical graph sampling [20] and convolutional generalizations [3,26,29]. However, these methods lack explicit mechanisms for modeling long-range cortical interactions, which are essential for brain age prediction as age-related changes are spatially distributed across distant regions [6,23,27]. Accurate estimation thus requires multiscale modeling from global to local morphology.

Attention-based architectures have driven efforts to capture long-range dependencies on cortical surfaces. Surface Vision Transformer (SiT) [7] projects cortical surfaces onto the sphere and partitions them into patches that interact via global attention. This design enables interactions beyond local convolutional neighborhoods, but incurs substantial computational and memory overhead that limits the model to low-resolution representations. Later, Multiscale Surface Vision Transformer (MS-SiT) [8] replaces global attention with hierarchical local attention using shifted windows to reduce computational cost and support higher-resolution inputs for improved performance. In parallel, Surface Vision Mamba (SiM) [13] adopts the same patch-based framework and replaces the transformer with Mamba, achieving comparable predictive performance while improving memory efficiency and inference speed.

Despite these advances, patch-based architectures are still constrained by partitioning. This constraint introduces boundary inconsistencies across training samples under imperfect shape correspondence. Moreover, compressing high-dimensional structural relationships within each patch into a fixed-dimensional token embedding inevitably results in the loss of fine-grained geometric information [16]. Sequence-based formulations such as SiM [13] require explicit patch ordering which overlooks geometric continuity and causes a structural mismatch despite bidirectional modeling. As a result, patch-based transformer and Mamba models remain sensitive to rotations and typically require anatomical registration or extensive data augmentation to achieve competitive performance.

In this paper, we propose a spherical harmonics–based Mamba (SPHARM-Mamba), a generic backbone network for genus-zero surface analysis. The proposed method departs from patch-based representations and constructs compact rotation-invariant descriptors in the spherical harmonic domain. By performing convolution directly in the spectral domain, the framework preserves high-resolution surface information and enables rich channel-wise feature integration within each harmonic degree. Unlike patch representations, harmonic decomposition follows a natural progression from global structures to local details and thus forms a structure-aware ordered sequence intrinsically compatible with Mamba [11]. We further employ Mamba to model cross-degree dependencies and capture long-range relationships for multiscale representation learning. This design enables efficient and expressive surface modeling across spatial scales. The main contributions of this work are summarized as follows:

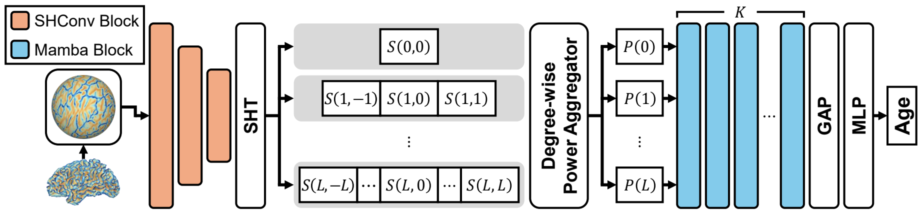

Fig. 1. Overview of the proposed SPHARM-Mamba framework. Input cortical surface features are projected onto the unit sphere to serve as input channels for SHConv blocks. Each SHConv block performs SHT followed by rotation-equivariant spectral convolution and ISHT. The resulting spherical harmonic coefficients are grouped by harmonic degree and converted into rotation-invariant descriptors through degree-wise power aggregation. The degree-ordered tokens are embedded and passed to K stacked Mamba blocks to model cross-degree multiscale dependencies. Global average pooling and an MLP head produce the final age prediction.

- We demonstrate superior cortical age regression performance without anatomical registration, outperforming existing geometric and patch-based models.
- We introduce a spherical harmonics framework that ensures rotation invariance and enables feature learning within and across harmonic degrees.
- We present a generic backbone for genus-zero surface analysis that achieves higher accuracy while using only 3.11% of the parameters of the best-performing baseline, MS-SiT [8].

2 Methods

2.1 Methodlogical Overview

Let $f : \mathbb{S}^2 \to \mathbb{R}^C$ denote C cortical geometric features defined on the unit sphere. Our goal is to learn a regression model that predicts a scalar brain age from the spherical signal f. In this study, we parameterize this regression model using neural networks. Specifically, we extract deep learning representations from f in the spectral domain to capture channel-wise interactions (Sec. 2.2). From these representations, we construct compact, rotation-invariant descriptors (Sec. 2.3). Finally, we employ a Mamba-based sequential model to capture harmonic-wise dependencies across spectral components within the resulting descriptors (Sec. 2.4). Figure 1 illustrates an overview of the proposed framework.

2.2 Spherical Harmonic Convolution

To extract informative features within our framework, we account for the noisy and heterogeneous nature of cortical geometry by employing spherical convolutions that respect the underlying spherical domain. Among the available formulations, we adopt spherical harmonics-based convolutions due to their natural preservation of rotational equivariance and computational efficiency in coefficient-domain operations [12]. Specifically, we perform harmonic decomposition by applying the spherical harmonic transform to the cortical surface signals f. At harmonic degree ℓ and order m, the i-th channel of f ($i = 1, \ldots, C$) can be encoded by its spectral coefficients $S_i(\ell, m)$. The spherical harmonic convolution (SHConv) is subsequently performed directly in the spectral domain. By constraining the convolutional filters to be zonal, equivariance to 3D rotations is inherently preserved [12]. In this manner, SHConv enables effective learning of feature representations by mixing spherical signals f in the spectral domain.

2.3 Rotation Invariant Spectral Descriptors

We now convert the learned rotation-equivariant representations into rotation-invariant descriptors, as equivariance alone does not guarantee invariance to arbitrary rotations. These invariant descriptors are then used as inputs to the Mamba model. From the rotational property of spherical harmonics [18], a 3D rotation preserves the distribution of spectral energy across harmonic degrees ℓ. Therefore, a 3D rotation only redistributes coefficients across different orders m within the same harmonic degree. More formally, we construct a rotation-invariant descriptor of the i-th channel by aggregating coefficient energy within each degree ℓ:

$$P_i(\ell) = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \left(S_i(\ell, m)\right)^2 . \tag{1}$$

After computing the degree-wise power descriptors $P(\ell)$, we obtain rotation-invariant representations for each channel up to the maximum degree $L \in \mathbb{Z}^+$. In the discrete setting, the descriptor dimensionality is bounded by L, independent of the mesh sampling density or the signal complexity. Consequently, memory usage remains stable even as the number of vertices increases. The computational cost of the transform scales linearly with the number of vertices $O(VL^2)$ while the memory requirement remains $O(L^2)$. In practice, L is a fixed constant (e.g., $L = 80$ in our experiments). This property enables the use of high-resolution surfaces without additional memory burden.

2.4 Cross-Degree Multiscale Modeling with Mamba

Within each harmonic degree, SHConv learns expressive feature representations through channel-wise interactions. However, it operates separately across degrees and does not model relationships between different harmonic degrees. Consequently, multiscale dependencies across the degree dimension remain unexplored. To address this limitation, we treat the degree-wise descriptors as an ordered sequence and employ a Mamba-based state-space model to capture cross-degree multiscale interactions.

Following the spherical harmonic decomposition, rotation-invariant descriptors $\{P_1(\ell), \ldots, P_C(\ell)\}$ are constructed via degree-wise power aggregation of the spectral coefficients. These descriptors exhibit an intrinsic coarse-to-fine ordering along the harmonic degree dimension. Lower degrees capture global structural patterns, whereas higher degrees encode progressively more localized structural details. Formally, let the rotation-invariant representation be denoted by degree-wise vectors $P(\ell) \in \mathbb{R}^C$ for $\ell = 0, 1, \dots, L$, where each degree serves as a token in a sequence ordered from low- to high-frequency components:

$$X_0 = \left[P(0), P(1), \dots, P(L) \right]. \tag{2}$$

Since these descriptors are derived from the degree-wise power spectrum, the initial sequence representation X_0 is invariant to rotations of the input spherical signal. The sequence is fed into stacked Mamba layers:

$$X_k = \mathrm{Mamba}(X_{k-1}), \quad k = 1, \dots, K, \tag{3}$$

resulting in the final sequence representation $X_K \in \mathbb{R}^{(L+1)\times C}$. This formulation enables explicit modeling of dependencies across harmonic degrees and facilitates multiscale integration between global and local spectral components. Finally, a global average pooling operation is applied along the degree dimension, and the pooled feature is passed to a multi-layer perceptron (MLP) regression head to predict the subject's age:

$$\mathrm{Age} = \mathrm{MLP}(\mathrm{Pool}(X_K)). \tag{4}$$

Unlike patch-based representations that rely on artificially imposed ordering, our sequence is naturally defined by the harmonic hierarchy. This structure preserves geometric consistency while compensating for the absence of cross-degree interactions in SHConv.

3 Experimental Setup

3.1 Dataset

We used T1 weighted scans from four publicly available datasets: HCP [10] (N=1,113), IXI [15] (N=563), OASIS-1 [25] (N=257), and OASIS-2 [24] (N=373). For OASIS-1 and OASIS-2, we included only healthy participants. They were combined into a large adult cohort consisting of 2,306 subjects spanning a wide age range from 19.98 to 98.00 years (mean ± stdev: 45.94 ± 21.53). The cortical surfaces were constructed via the standard FreeSurfer pipeline [9]. We used four FreeSurfer measures with icosahedral subdivision level up to 6: *curv*, *sulc* [9], cortical thickness, and surface area. We assessed robustness under three different alignment conditions: (1) the original dataset, (2) a version with cortical surface registration to *fsaverage* [9], and (3) a version transformed by uniformly sampled random $SO(3)$ rotations. All variants contained the same number of samples. Each hemisphere was used as an independent sample.

3.2 Implementation Details

Training was performed for 300 epochs using AdamW [22] with an initial learning rate of 1×10^{-4}, weight decay of 0.002, and a cosine learning rate schedule [21] with 20 warmup epochs. The batch size was set to 8, and Smooth L1 loss with $\beta = 10$ was used. Two SHConv layers were employed for spectral feature extraction. The first operated up to harmonic degree 80 with 128 output channels, and the second up to degree 40 with 256 channels. After degree-wise power aggregation, each channel was reduced to 41 scalar descriptors (*i.e.*, L=40). For cross-degree modeling, we used a Mamba module with model dimension 128, state dimension 64, convolution kernel size 4, expansion factor 2, and 4 layers. The regression head consisted of a single hidden layer with 128 units followed by ReLU activation.

3.3 Baseline Methods

We compared our method with representative convolution-based geometric deep learning models and patch-based models. For fair comparison across all methods, we ran parameter optimization and reported performance across all runs: SurfGNN [20] with input resolutions [5,124; 20,484; 81,924], and TSL depth [2, 4]; MoNet [26] with hidden dimensions [32, 64, 128, 256], and kernel sizes [3, 5, 7]; S2CNN [3] with bandwidth [30, 60, 120] and channel sizes [16, 32, 64, 128]; SU-Net [29] with C_1 [32, 64]; SiT [7] with transformer depth [8, 12, 16] and attention heads [4, 6, 8]; MS-SiT [8] with drop path rate [0.0, 0.05, 0.1] and window size [7, 9, 11]; and SiM [13] with model dimension [128, 192, 256], state dimension [8, 16, 32], and depth [12, 24]. For a fair evaluation, the dataset was divided into five folds with a 3:1:1 split for training, validation, and testing. All models were trained using Smooth L1 loss, and performance was evaluated using Mean Absolute Error (MAE, in years). Final results were reported as the mean and standard deviation of the test MAE across the five folds.

4 Results

4.1 Ablation Study

To validate the proposed framework, we performed ablation studies to examine the necessity of cross-degree modeling, degree ordering, spatial resolution, and harmonic degree truncation. We first evaluated the role of cross-degree modeling. The baseline model consists of SHConv layers followed by a multi-layer perceptron (MLP), as in [14], instead of Mamba. This configuration yields an MAE of 5.1. Replacing the MLP with the proposed cross-degree modeling using Mamba reduces the MAE to 4.006, corresponding to a 19.5% improvement. Notably, the MLP baseline uses a larger SHConv backbone with more parameters (0.998M vs. 0.856M) and a larger harmonic basis size (121^2 vs. 81^2). Its worse performance suggests that the gain comes from Mamba-based cross-degree modeling rather than increased model capacity. This result suggests the importance

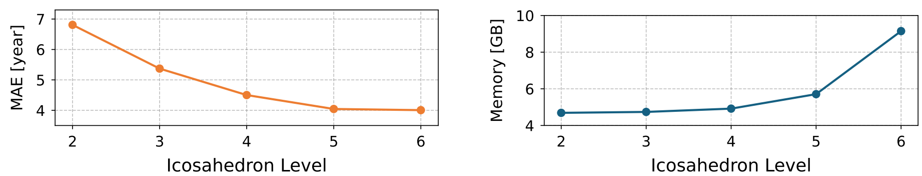

Fig. 2. Ablation study. Evaluation of icosahedral subdivision levels under the default configuration on an NVIDIA RTX 6000 Ada Generation GPU. The left panel shows MAE and the right panel shows GPU memory consumption for each subdivision level.

of explicitly modeling cross-degree multiscale dependencies for brain age prediction, where age-related cortical changes are spatially distributed across multiple structural scales.

We examined the effect of degree ordering on sequence modeling. Spherical harmonics follow a natural hierarchy from low-frequency global components to high-frequency local details. In our default setting, the proposed low-to-high degree ordering achieves an MAE of 4.006. Reversing the order increases the MAE to 4.634 and results in a 15.7% performance drop. This is likely because age-related cortical changes are spatially distributed and require a stable global representation that is progressively refined by local morphological variations. The proposed low-to-high degree ordering aligns with this principle and facilitates effective multiscale representation learning for accurate age estimation.

We evaluated performance across icosahedral subdivision levels ranging from level 2 to level 6. Fig. 2 plots the MAE and memory consumption across subdivision levels. Predictive accuracy improves as spatial resolution increases and indicates that higher-resolution surfaces provide more informative cortical representations. As described in Section 2, the number of vertices increases substantially while memory consumption remains nearly unchanged, which allows stable training at high resolution. Based on empirical performance and computational efficiency, we used a subdivision level of 6 for the final model.

We analyzed the impact of harmonic degree truncation in SHConv. The default configuration uses degree limits of (80, 40) for the first and second SHConv layers. Increasing the limits to (120, 80) reduces the MAE from 4.006 to 3.995. As studied in [2], harmonic degrees above 40 mainly reflect fine-scale geometric irregularities and segmentation noise rather than biologically meaningful structural patterns. Consistent with this observation, we adopt the (80, 40) configuration to balance efficiency and performance.

4.2 Comparison to Baselines

The proposed method achieves the lowest MAE of 4.006 years regardless of registration and improves performance by 32.6% over the best baseline on registered data with an MAE of 5.948, as shown in Table 1 and Figure 3. Under random $SO(3)$ rotations, most baselines show clear performance degradation. Patch-based methods are particularly unstable since random rotations al-

Table 1. Mean Absolute Error (MAE) on original, random rotation (Rotation), and registered (Registration) settings. Params: Number of learnable parameters (M).

Method	MAE Loss			Params
	Original	Rotation	Registration	(M)
S2CNN [3]	8.084 ± 0.102	9.112 ± 0.122	8.069 ± 0.279	0.846
SU-Net [29]	7.390 ± 0.375	9.641 ± 0.562	6.141 ± 0.094	3.735
MoNet [26]	6.105 ± 0.098	6.130 ± 0.111	6.078 ± 0.107	**0.643**
SurfGNN [20]	6.800 ± 0.227	7.776 ± 0.289	6.084 ± 0.216	20.892
SiT [7]	7.887 ± 0.398	8.559 ± 0.417	7.160 ± 0.217	5.613
MS-SiT [8]	6.500 ± 0.242	6.977 ± 0.138	5.948 ± 0.280	27.536
SiM [13]	13.236 ± 3.218	17.491 ± 0.917	7.580 ± 0.449	12.609
Ours	**4.006 ± 0.093**	**4.076 ± 0.161**	**4.100 ± 0.170**	0.856

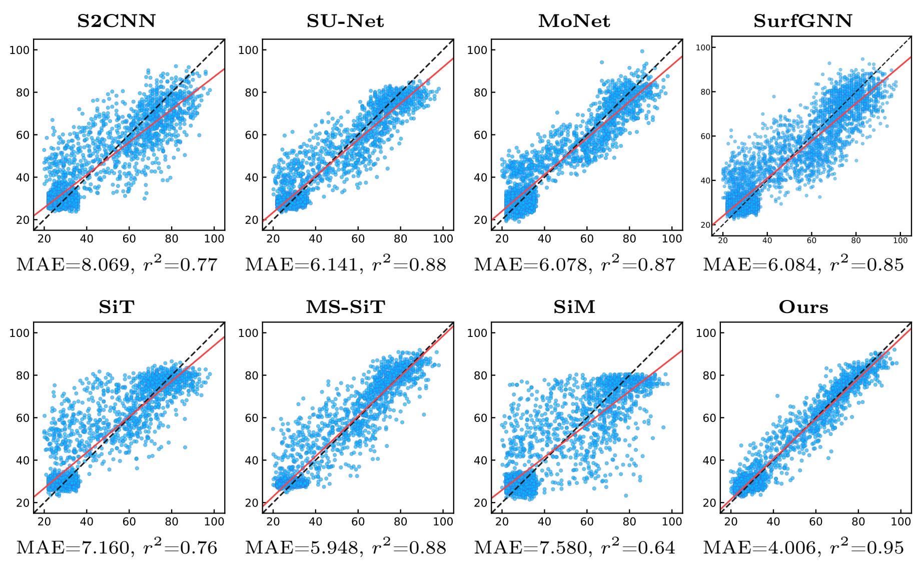

Fig. 3. Scatter plots of ground-truth age (x-axis) and predicted age (y-axis) on the registered dataset. The dashed black line denotes the identity line (y = x) and the solid red line indicates the fitted regression. MAE and the coefficient of determination ($\mathbf{r^2}$) are reported below each panel.

ter patch assignments and change the local content in each patch. In contrast, rotation-invariant methods including our method remain stable and maintain consistent accuracy. Because we construct rotation-invariant descriptors in the spherical harmonic domain, the model receives identical representations under arbitrary rotations. With the surface registration, all baselines show improved performance. As expected, patch-based models [7,8,13] exhibit a large gap since anatomical patterns become similar across corresponding patch indices. In contrast, convolution-based models (MoNet [26] and SurfGNN [20]) are less dependent on surface registration likely because local neighborhood aggregation preserves local structural information. Our model inherits this advantage in the spherical harmonic domain; the proposed spectral descriptors and cross-degree multiscale modeling effectively capture cortical structure without requiring anatomical alignment.

5 Conclusion

We introduced SPHARM-Mamba as a generic backbone for genus-zero surface analysis. The model replaced patch partitioning with spherical harmonic descriptors and removed anatomical registration. SHConv produced rotation-equivariant features that were converted into rotation-invariant tokens ordered from global to local scales and were robust to arbitrary $SO(3)$ rotations. Mamba was used to model cross-degree dependencies for efficient multiscale representation learning. On large-scale age prediction, the proposed method achieved superior MAE with fewer learnable parameters compared to competing approaches trained even with surface registration. Future work will extend our approach to diverse clinical prediction tasks with broad surface types.

Acknowledgments. This work was supported in part by the National Research Foundation of Korea (NRF) under RS-2024-00333931 and RS-2025-02216257, and in part by the Institute for Information & Communications Technology Planning & Evaluation (IITP) under RS-2019-II191906.

Disclosure of Interests. The authors have no competing interests in the paper.

References

1. Bermudez, C., Plassard, A.J., Chaganti, S., Huo, Y., Aboud, K.S., Cutting, L.E., Resnick, S.M., Landman, B.A.: Anatomical context improves deep learning on the brain age estimation task. Magnetic Resonance Imaging **62**, 70–77 (2019)
2. Chung, M.K., Worsley, K.J., Nacewicz, B.M., Dalton, K.M., Davidson, R.J.: General multivariate linear modeling of surface shapes using surfstat. Neuroimage **53**(2), 491–505 (2010)
3. Cohen, T.S., Geiger, M., Köhler, J., Welling, M.: Spherical cnns. arXiv preprint arXiv:1801.10130 (2018)
4. Cole, J.H., Franke, K.: Predicting age using neuroimaging: innovative brain ageing biomarkers. Trends in neurosciences **40**(12), 681–690 (2017)
5. Cole, J.H., Ritchie, S.J., Bastin, M.E., Valdés Hernández, M., Muñoz Maniega, S., Royle, N., Corley, J., Pattie, A., Harris, S.E., Zhang, Q., et al.: Brain age predicts mortality. Molecular psychiatry **23**(5), 1385–1392 (2018)
6. Corps, J., Rekik, I.: Morphological brain age prediction using multi-view brain networks derived from cortical morphology in healthy and disordered participants. Scientific reports **9**(1), 9676 (2019)
7. 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)

8. Dahan, S., Williams, L.Z., Rueckert, D., Robinson, E.C.: The multiscale surface vision transformer. ArXiv pp. arXiv–2303 (2024)
9. Fischl, B.: Freesurfer. Neuroimage **62**(2), 774–781 (2012)
10. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., et al.: The minimal preprocessing pipelines for the human connectome project. Neuroimage **80**, 105–124 (2013)
11. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. In: First conference on language modeling (2024)
12. Ha, S., Lyu, I.: Spharm-net: Spherical harmonics-based convolution for cortical parcellation. IEEE Transactions on Medical Imaging **41**(10), 2739–2751 (2022)
13. He, R., Zheng, W., Zhao, L., Wang, Y., Zhu, D., Wu, D., Hu, B.: Surface vision mamba: Leveraging bidirectional state space model for efficient spherical manifold representation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 599–608. Springer (2025)
14. Im, Y., Zhao, Y., Gutman, B.A., Thomopoulos, S.I., Haddad, E., Zhu, A.H., Jahanshad, N., Thompson, P.M., Ching, C.R.: A surface-based deep learning approach for cortical shape analysis. bioRxiv (2024)
15. IXI Dataset: Information eXtraction from Images. https://brain-development.org/ixi-dataset/
16. Kahatapitiya, K., Ryoo, M.S.: Swat: Spatial structure within and among tokens. arXiv preprint arXiv:2111.13677 (2021)
17. Kaufmann, T., van der Meer, D., Doan, N.T., Schwarz, E., Lund, M.J., Agartz, I., Alnæs, D., Barch, D.M., Baur-Streubel, R., Bertolino, A., et al.: Genetics of brain age suggest an overlap with common brain disorders. BioRxiv p. 303164 (2018)
18. Kazhdan, M., Funkhouser, T., Rusinkiewicz, S.: Rotation invariant spherical harmonic representation of 3 d shape descriptors. In: Symposium on geometry processing. vol. 6, pp. 156–164 (2003)
19. Li, G., Nie, J., Wang, L., Shi, F., Gilmore, J.H., Lin, W., Shen, D.: Measuring the dynamic longitudinal cortex development in infants by reconstruction of temporally consistent cortical surfaces. Neuroimage **90**, 266–279 (2014)
20. Li, Z., Zhang, J., Zeng, Y., et al.: Surfgnn: A robust surface-based prediction model with interpretability for coactivation maps of spatial and cortical features. Medical Image Analysis p. 103793 (2025)
21. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
22. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
23. Madan, C.R., Kensinger, E.A.: Predicting age from cortical structure across the lifespan. European Journal of Neuroscience **47**(5), 399–416 (2018)
24. Marcus, D.S., Fotenos, A.F., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Open access series of imaging studies: longitudinal mri data in nondemented and demented older adults. Journal of cognitive neuroscience **22**(12), 2677–2684 (2010)
25. Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. Journal of cognitive neuroscience **19**(9), 1498–1507 (2007)
26. Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., Bronstein, M.M.: Geometric deep learning on graphs and manifolds using mixture model cnns. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5115–5124 (2017)

27. Valizadeh, S., Hänggi, J., Mérillat, S., Jäncke, L.: Age prediction on the basis of brain anatomical measures. Human brain mapping **38**(2), 997–1008 (2017)
28. Wang, J., Li, W., Miao, W., Dai, D., Hua, J., He, H.: Age estimation using cortical surface pattern combining thickness with curvatures. Medical & biological engineering & computing **52**(4), 331–341 (2014)
29. 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: International Conference on Information Processing in Medical Imaging. pp. 855–866. Springer (2019)