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.

Region-Wise Interpretable Neonatal Brain Age Prediction via Pretrained Segmentation-Guided Attention

Dayeon Bak[1][0009−0002−4681−2592], Junghwa Kang[1][0000−0003−4923−8338], Hyun Gi Kim[2,3][0000−0002−0086−543X], and Yoonho Nam[1(✉)][0000−0003−2149−0072]

[1] Department of Biomedical Engineering, Hankuk University of Foreign Studies, Yongin-si, Gyeonggi-do, Korea
yoonhonam@hufs.ac.kr
[2] Department of Radiology and Research Institute of Radiology, Asan Medical Center, University of Ulsan College of Medicine, Seoul, Republic of Korea
[3] Department of Radiology, Eunpyeong St. Mary's Hospital, College of Medicine, The Catholic University of Korea, Seoul, Republic of Korea

Abstract. The neonatal brain undergoes rapid structural maturation during the perinatal period, yet accurate and anatomically interpretable assessment of brain maturity remains challenging. While deep learning–based brain age prediction models have shown strong predictive performance, most rely on post-hoc visualization techniques that provide limited region-level insight. In this study, we propose an interpretable neonatal brain age prediction framework guided by a pretrained model for neonatal brain segmentation. The framework extracts region-wise representations from a pretrained segmentation model and integrates them through a multi-region cross-attention module, enabling quantitative analysis of anatomical contributions to age prediction. Experiments on the Developing Human Connectome Project dataset and an independent institutional dataset demonstrate that the proposed model outperforms direct structural image-based brain age prediction baselines. Region-level analysis reveals age-dependent shifts in RIS patterns, with increasing values in gray matter and decreasing values in the basal ganglia and ventricles with advancing age. These findings suggest that segmentation-guided region-wise attention may improve interpretability in neonatal brain age prediction.

Keywords: Neonatal brain · Interpretability · Age prediction.

1 Introduction

The neonatal brain undergoes rapid and dynamic structural maturation during the perinatal period, with tissue-specific developmental trajectories [21]. These early structural changes are closely linked to later neurodevelopmental outcomes, highlighting the clinical importance of characterizing perinatal brain maturation [21]. With advances in MRI technology, these developmental processes can

be observed non-invasively in vivo, and brain age prediction may quantitatively estimate biological brain maturity [18,21]. The difference between predicted and chronological age, known as the brain age gap (BAG), reflects deviation from normative developmental trajectories and has been associated with neurodevelopmental outcomes in infancy, suggesting its potential as an early biomarker of perinatal brain development [5, 18].

However, BAG alone has limitations in comprehensively understanding neurodevelopmental status. As a single scalar value, BAG cannot fully explain the underlying biological mechanisms that drive atypical maturation. The same BAG value may arise from distinct neurodevelopmental processes [19]. Owing to these limitations, there has been increasing interest in interpretability approaches aimed at identifying the anatomical structures and developmental features that contribute to brain age prediction [3, 5–7, 11, 22].

Recently, deep learning-based approaches have been proposed for neonatal brain age prediction. CNN and transformer-based models can learn local and global image representations and are often interpreted using post-hoc analyses [6, 7, 22]. However, patch-level modeling may include tissue boundaries or partial tissue information, limiting the explicit use of tissue-specific developmental cues [12]. Mask-based approaches using white matter (WM) or gray matter (GM) images can emphasize developmental processes such as myelination or cortical folding, but they may discard original image information, depend on segmentation quality, and fail to preserve boundary information between brain regions [1, 6, 11]. Region-level features such as volume provide anatomical interpretability, but predefined features may constrain model expressiveness and limit the learning of high-dimensional developmental patterns [10, 16]. In addition, saliency maps and Grad-CAM are widely used, but their reliability can be affected by gradient noise, ReLU saturation, class score scaling, and low spatial resolution [5, 20, 22].

Although multi-task learning has recently been applied to adult brain studies to jointly improve structural segmentation and age prediction [3], these methods are mostly based on adult data and often rely on attention visualization rather than quantitative region-wise analysis. Therefore, tissue- and region-level interpretability remains limited in neonatal brain age prediction.

In this study, we propose an interpretable brain age prediction framework that provides region-wise interpretability by leveraging a segmentation model pretrained for neonatal brain segmentation. Through the use of region-wise representations derived from the pretrained BabySeg [8] model, the framework enables quantitative analysis of the relative importance of neonatal brain regions in age prediction. Overall, our study provides not only improved predictive performance but also a foundation for understanding neonatal brain development in a interpretable manner.

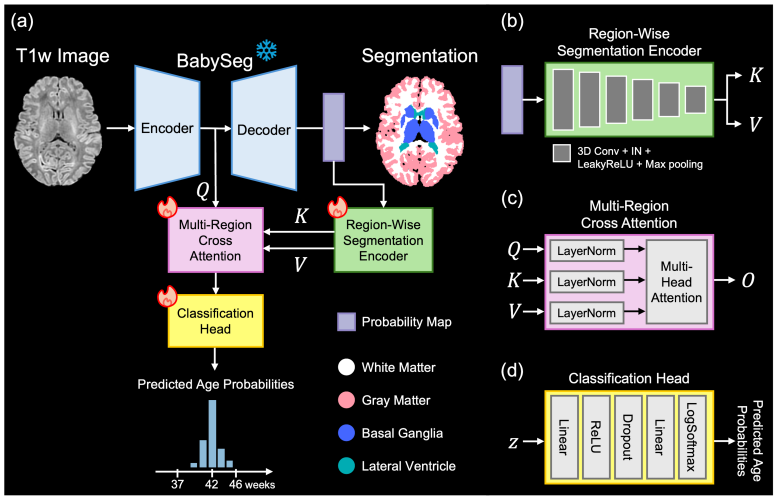

Fig. 1. (a) Overview of the proposed brain age prediction framework. Detailed architectures of (b) the region-wise segmentation encoder, (c) the multi-region cross-attention module, and (d) the classification head.

2 Methods

Since structural brain differences are minimal across narrow age ranges (e.g., 1–6 days), we discretized the age range and reformulated the regression task as a classification task. This design was intended to reduce the model's sensitivity to subtle within-bin variations and encourage it to learn more robust age-related imaging features. The original age range of 37–46 weeks was extended to 34–48 weeks by adding a 2.5-week margin to ensure complete Gaussian-shaped target distributions. This range was then divided into 14 one-week intervals. Consequently, the ground-truth age was converted into a normal distribution with a standard deviation of 8. An overview of the proposed brain age prediction framework is illustrated in Fig. 1(a). The architecture comprised four primary components: a pretrained segmentation network, a region-wise segmentation encoder, a multi-region cross-attention module, and a classification head.

2.1 Pretrained Segmentation Network

To obtain multi-region brain segmentations and the latent representation of the input T1-weighted (T1w) images, we utilized BabySeg, a UNet-based neonatal brain segmentation network [8]. BabySeg originally outputs a probability map of 22 regions. In this study, to focus on comparisons among major tissues and anatomical regions that show prominent structural changes during neonatal brain maturation, we reorganized the output into a 4-channel probability map consisting of white matter (WM), gray matter (GM), basal ganglia (BG), and lateral ventricles (LV), as shown in Fig. 1(a). The latent representation extracted

Table 1. Performance comparison on the in-dataset and external test sets. MAE is reported in weeks.

Model	dHCP T1 (In-dataset test set)						ESMH T1 (External test set)					
	Normal ($N = 40$)		Abnormal ($N = 158$)		Total ($N = 198$)		Normal ($N = 142$)		Abnormal ($N = 65$)		Total ($N = 207$)	
	MAE↓ (week)	$R\uparrow$	MAE↓ (week)	$R\uparrow$	MAE↓ (week)	$R\uparrow$	MAE↓ (week)	$R\uparrow$	MAE↓ (week)	$R\uparrow$	MAE↓ (week)	$R\uparrow$
SFCN	0.640 ±0.439	0.900	0.844 ±0.682	0.833	0.803 ±0.645	0.845	1.578 ±0.912	0.758	1.536 ±0.829	0.733	1.565 ±0.886	0.750
Swin-T	0.903 ±0.757	0.718	1.095 ±0.910	0.681	1.057 ±0.883	0.690	2.197 ±1.218	0.532	1.802 ±0.975	0.585	2.073 ±1.159	0.545
Proposed	**0.569 ±0.419**	**0.913**	**0.799 ±0.593**	**0.858**	**0.753 ±0.569**	**0.869**	**1.379 ±0.696**	**0.861**	**1.286 ±0.628**	**0.862**	**1.350 ±0.675**	**0.861**

from BabySeg was a 64-channel vector with a spatial resolution of 1/64 of the input T1w image.

2.2 Region-Wise Segmentation Encoder

The detailed architecture of the region-wise segmentation encoder is illustrated in Fig. 1(b). To extract representations from the probability maps, we utilized a convolutional block consisting of a 3D 3×3 convolution, instance normalization, LeakyReLU (slope = 0.2), and a max-pooling layer (kernel size and stride = 2). This block was repeated six times, increasing the channel dimension to 64. To maintain regional independence during feature extraction, the region dimension was transposed to the batch dimension.

2.3 Multi-Region Cross-Attention

To effectively fuse the global context of the latent representation with regional features, we propose a multi-region cross-attention module based on the cross-attention vision transformer [2]. This module utilizes the 3D global latent representation as the query, while regional features serve as the keys and values. This design allows the network to dynamically learn the relative contribution of each anatomical region to the final age prediction.

Feature Tokenization Let $X_{global} \in \mathbb{R}^{B\times E\times D\times H\times W}$ denote the input latent representation and $X_{region} \in \mathbb{R}^{B\times C\times E\times D\times H\times W}$ denote the regional feature tensor extracted from C partitioned regions. In our implementation, we set the embedding dimension $E = 64$, the spatial dimensions $D = H = W = 3$, and the number of regions $C = 4$ (WM, GM, BG, and LV).

For the attention operation, the 3D spatial dimensions were flattened into a 1D sequence. The global feature X_{global} generated the query tokens $Q \in \mathbb{R}^{B\times S\times E}$, where the sequence length $S = D \times H \times W = 27$. To ensure each region was represented by a robust summary vector, the regional features X_{region} underwent global average pooling across their spatial dimensions to produce the

key K and value V tokens:

$$K = V = \frac{1}{DHW} \sum_{d=1}^{D} \sum_{h=1}^{H} \sum_{w=1}^{W} X_{region}^{(d,h,w)} \in \mathbb{R}^{B \times C \times E}. \tag{1}$$

Consequently, each region was encapsulated into a single embedding vector, resulting in C total tokens.

Cross-Attention Mechanism The architecture of the cross-attention module is illustrated in Fig. 1(c). To maintain training stability, the Q, K, and V tokens were passed through layer normalization (LN) before being processed by the multi-head attention layer:

$$O, W_{attn} = \text{MultiHeadAttention}(\text{LN}(Q), \text{LN}(K), \text{LN}(V)) \tag{2}$$

The mechanism computed the dot-product between the global queries and regional keys to yield the attention weight matrix $W_{attn} \in \mathbb{R}^{B \times S \times C}$. The resulting output $O \in \mathbb{R}^{B \times S \times E}$ represented the updated global features, enriched by dynamically attending to the C regional features across all S spatial locations.

Region Importance Score The attention weights W_{attn} served as an interpretable metric indicating how strongly the model attended to each region during the decision-making process. The final region importance score (RIS), $S_{imp} \in \mathbb{R}^{B \times 1 \times C}$, was obtained by averaging the attention weights across all query tokens.

2.4 Classification Head

For the final prediction, global average pooling was applied to the updated representation O along the sequence dimension (S) to obtain a global vector $z \in \mathbb{R}^{B \times E}$. This vector was fed into a classification head consisting of a Multi-Layer Perceptron (MLP), as illustrated in Fig. 1(d). The prediction $\hat{y}$ was obtained via a LogSoftmax function:

$$\hat{y} = \text{LogSoftmax}(\text{MLP}(z)) \tag{3}$$

where $\hat{y}$ represented the predicted log-probabilities for each neonatal age class. Finally, the predicted age was calculated as the weighted average across all age bins:

$$\text{pred} = \sum_{i=1}^{14} \exp(\hat{y}_i) \cdot \text{age}_i, \tag{4}$$

where i denotes the class index and age_i represents the midpoints of the 14 discretized age bins, ranging from 34.5 to 47.5 weeks.

2.5 Loss Function

The network outputs log-probabilities, $\hat{y}$, for each age class. As previously described, the ground-truth age was represented as a normal distribution rather than a single-point label. The model was optimized using Kullback–Leibler divergence loss to minimize the discrepancy between the predicted distribution $\hat{y}$ and the target normal distribution derived from the actual age.

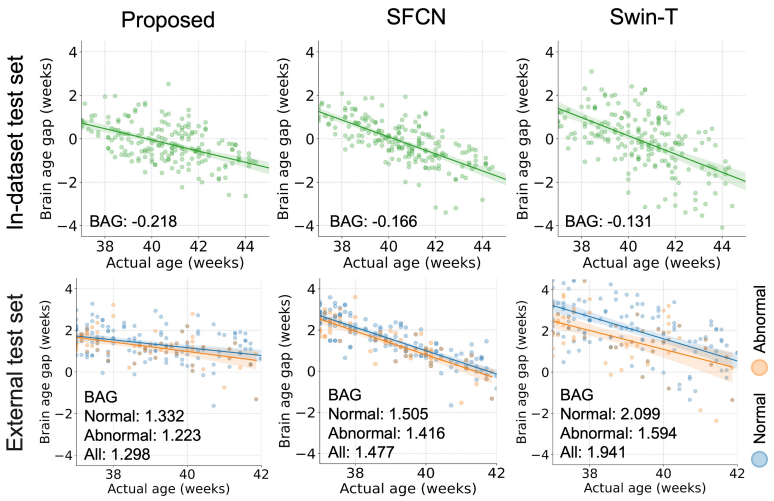

Fig. 2. Brain age gap comparison on the in-dataset and external test sets.

3 Experiments

3.1 Dataset

Although T2-weighted images provide clinically important information in neonatal brain MRI, high-resolution 3D T2-weighted images, unlike 2D T2-weighted images, may not always be acquired in routine clinical settings. Therefore, we used 3D T1w images to facilitate a fair comparison between the in-dataset and external datasets.

In-dataset In this study, 561 3D T1w images of neonates aged 37 to 46 weeks postmenstrual age (PMA) were obtained from the Developing Human Connectome Project (dHCP, third release) [4]. These images were previously preprocessed by the dHCP structural pipeline, which includes bias field correction [15]. An independent training set, validation set, and test set were used in this study. To train the model, data from 343 neonates with radiology scores of 1 or 2 were

used. For validation, data from 20 neonates were used, consisting of 10 neonates with a radiology score of 1 and 10 neonates with a radiology score of 2. Additionally, data from 198 neonates covering all radiology scores were used as the test set.

External dataset To evaluate the generalizability of the proposed model, we used an additional external test set consisting of 207 3D T1w images from from Eunpyeong St. Mary's Hospital (ESMH), acquired from neonates with PMA ranging from 37 to 42 weeks. This external test set consisted of 142 normal neonates and 65 neonates diagnosed with hypoxic-ischemic encephalopathy. All MRI data were acquired using a 3T MRI scanner (MAGNETOM Vida, Siemens Healthineers, Erlangen, Germany). 3D T1w MPRAGE images were acquired with an isotropic voxel size of 0.8 mm. The acquisition parameters for the T1w images were TR/TE = 2400/2.37 ms and flip angle = 15°.

Preprocessing Although additional preprocessing is not strictly required for using the T1w images as input to the BabySeg model, all T1w images underwent the same preprocessing pipeline to ensure a fair comparison under consistent conditions. First, the images were skull-stripped using the HD-BET brain extraction tool [9]. Next, the images were rigidly registered to the dHCP neonatal brain standard space [14]. Additionally, the images were resampled to a spatial shape of $192 \times 192 \times 192$. Finally, intensities were linearly scaled to the range $[0, 1]$.

3.2 Baseline Models

The proposed method was compared with both CNN-based and transformer-based architectures. A Simple Fully Convolutional Network (SFCN) [17], a lightweight 3D CNN widely used for brain age prediction, was selected as a baseline. Additionally, a 3D Swin Transformer (Swin-T) [13] with four hierarchical stages was implemented to compare our approach against a state-of-the-art transformer architecture. For a fair comparison, all baseline models were implemented as classification models that output probabilities across 14 classes, identical to the proposed model. All baseline models were retrained using the same training dataset as the proposed model.

3.3 Evaluation

The performance of the model was quantitatively evaluated using the mean absolute error (MAE) and the Pearson correlation coefficient (R). To evaluate the ability of the model to capture abnormal patterns, we calculated the BAG. To quantify region importance, the RIS was computed. The mean RIS for each region was calculated, and region-wise averages were further analyzed across age groups. Age groups were defined by dividing the samples into three groups with approximately balanced sample sizes. For the in-dataset test set, subjects were divided into three groups: 37–39 weeks (N=66), 39–41 weeks (N=68), and 41–45

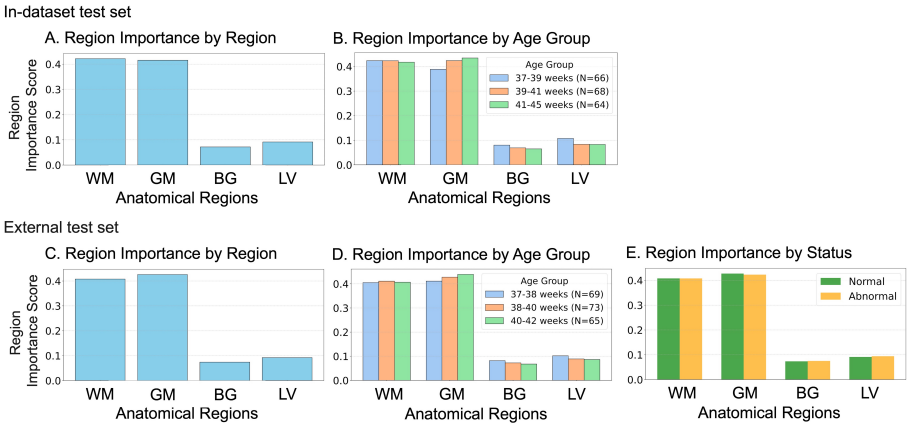

Fig. 3. Region-wise importance scores on the in-dataset and external test sets.

weeks (N=64). For the external test set, subjects were grouped into 37–38 weeks (N=69), 38–40 weeks (N=73), and 40–42 weeks (N=65). Additionally, in the external dataset, region importance was further analyzed according to status (normal vs. abnormal).

4 Results

As shown in Table 1, in the in-dataset test set, the proposed model achieved an MAE of 0.753 weeks, outperforming SFCN by 0.050 weeks and Swin-T by 0.304 weeks. The proposed model also achieved the highest R value of 0.869, which was higher than those of SFCN and Swin-T by 0.024 and 0.179, respectively. In the external test set, the proposed model achieved an MAE of 1.350 weeks, outperforming SFCN by 0.215 weeks and Swin-T by 0.723 weeks. The proposed model also maintained a high $R = 0.861$, which was comparable to its performance on the in-dataset test set and higher than those of SFCN and Swin-T by 0.111 and 0.316, respectively.

Figure 2 shows the BAG analysis. In the in-dataset test set, the BAG values were -0.218, -0.166, and -0.131 weeks for the proposed model, SFCN, and Swin-T, respectively. Swin-T showed the smallest absolute BAG, and the absolute difference between Swin-T and the proposed model was 0.087 weeks. All three models showed negative slopes between BAG and actual age. In the external test set, the BAG values were 1.298, 1.477, and 1.941 weeks for the proposed model, SFCN, and Swin-T, respectively. The proposed model showed the smallest absolute BAG, which was 0.643 weeks lower than that of Swin-T. In addition, all three models showed negative slopes, and the abnormal group tended to exhibit lower BAG values than the normal group.

WM and GM showed higher RIS values than BG and LV in both the in-dataset and external test sets (Fig. 3(A) and Fig. 3(C)). In both datasets, GM

Table 2. Ablation study of key components: MAE and correlation coefficient R.

Key components			**dHCP T1 (In-dataset test set)**						**Institutional Dataset T1 (External test set)**					
Latent	Prob. map	Cross Attn.	Normal ($N = 40$)		Abnormal ($N = 158$)		Total ($N = 198$)		Normal ($N = 142$)		Abnormal ($N = 65$)		Total ($N = 207$)	
			MAE↓ (week)	R↑	MAE↓ (week)	R↑	MAE↓ (week)	R↑	MAE↓ (week)	R↑	MAE↓ (week)	R↑	MAE↓ (week)	R↑
✓			0.889 ±0.598	0.771	1.014 ±0.792	0.741	0.988 ±0.757	0.750	**1.347 ±0.822**	0.671	**1.275 ±0.788**	0.631	**1.325 ±1.325**	0.659
	✓		0.717 ±0.537	0.857	0.844 ±0.617	0.841	0.819 ±0.603	0.845	1.476 ±0.766	0.831	1.343 ±0.761	0.826	1.434 ±0.765	0.827
✓	✓	✓	**0.569 ±0.419**	**0.913**	**0.799 ±0.593**	**0.858**	**0.753 ±0.569**	**0.869**	1.379 ±0.696	**0.861**	1.286 ±0.628	**0.862**	1.350 ±0.675	**0.861**

RIS increased with age, whereas BG and LV RIS decreased with age (Fig. 3(B) and Fig. 3(D)). As shown in Fig. 3(E), both the normal and abnormal groups showed higher RIS values in WM and GM than in BG and LV.

Table 2 presents the ablation study conducted to evaluate the contribution of each key component. In the in-dataset test set, the proposed method achieved the best performance, with an MAE of 0.753 weeks and an R value of 0.869. In the external test set, the latent-vector-only model showed a slightly lower MAE than the proposed method by 0.025 weeks. However, the proposed method achieved a substantially higher correlation coefficient, improving the R value by 0.202 compared with the latent-vector-only model.

5 Discussion

In this study, we proposed an interpretable brain age prediction framework that leverages a segmentation model pretrained for neonatal brain segmentation. The proposed method showed relatively stable predictive performance compared with the baseline models on the independent external dataset. Considering the variability encountered in real-world clinical settings, robust performance on external data may represent an important advantage in terms of model generalizability and potential clinical applicability. Stratified analysis between the normal and abnormal groups further showed distinct brain age prediction patterns, suggesting the potential of brain age prediction as a biomarker for neonatal brain development.

In addition, RIS analysis (Fig. 3) showed relatively higher RIS values for WM and GM. This pattern may be related to rapid WM myelination and cortical maturation during the term-equivalent period. In contrast, the decreasing RIS of BG and LV with age may indicate that deep gray matter and ventricular morphology are more informative at earlier ages, whereas cortical and WM maturation become increasingly dominant with advancing PMA. These findings suggest that combining deep learning-based brain age prediction with region-wise interpretability analysis may provide meaningful insights into neonatal brain developmental processes.

The ablation results showed that the model combining latent and probability map representations through cross-attention achieved better in-dataset prediction performance than the latent-only and probability map-only models. All

models used pretrained BabySeg representations. This suggests that the proposed cross-attention-based integration may have contributed to the observed performance gain. The pretrained BabySeg representations may further support generalizability, as the model was designed to learn features robust to dataset shifts. Future studies using the original 22 regions are needed to investigate the relationship between RIS and regional volume.

Acknowledgments. This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2026-25523975 and RS-2026-25477832).

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

References

1. Callaert, D.V., Ribbens, A., Maes, F., Swinnen, S.P., Wenderoth, N.: Assessing age-related gray matter decline with voxel-based morphometry depends significantly on segmentation and normalization procedures. Frontiers in aging neuroscience **6**, 124 (2014)
2. Chen, C.F.R., Fan, Q., Panda, R.: Crossvit: Cross-attention multi-scale vision transformer for image classification. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 357–366 (2021)
3. Chen, M., Wang, Y., Shi, Y., Feng, J., Feng, R., Guan, X., Xu, X., Zhang, Y., Jin, C., Wei, H.: Brain age prediction based on quantitative susceptibility mapping using the segmentation transformer. IEEE Journal of Biomedical and Health Informatics **28**(2), 1012–1021 (2023)
4. Edwards, A.D., Rueckert, D., Smith, S.M., Abo Seada, S., Alansary, A., Almalbis, J., Allsop, J., Andersson, J., Arichi, T., Arulkumaran, S., et al.: The developing human connectome project neonatal data release. Frontiers in neuroscience **16**, 886772 (2022)
5. Fang, Z., Pan, N., Liu, S., Li, H., Pan, M., Zhang, J., Li, Z., Liu, M., Ge, X.: Comparative analysis of brain age prediction using structural and diffusion mris in neonates. NeuroImage **299**, 120815 (2024)
6. Gan, S., Fang, C., Xu, X., Xu, R., Huang, J., Sun, D., He, Q.: A brain-age prediction model in 3d-cnn-vit deep learning network
7. He, S., Grant, P.E., Ou, Y.: Global-local transformer for brain age estimation. IEEE transactions on medical imaging **41**(1), 213–224 (2021)
8. Hoffmann, M., Zöllei, L., Dalca, A.V.: Deep infant brain segmentation from multi-contrast mri. arXiv preprint arXiv:2512.05114 (2025)
9. Isensee, F., Schell, M., Pflueger, I., Brugnara, G., Bonekamp, D., Neuberger, U., Wick, A., Schlemmer, H.P., Heiland, S., Wick, W., et al.: Automated brain extraction of multisequence mri using artificial neural networks. Human brain mapping **40**(17), 4952–4964 (2019)
10. Kanwal, M., Son, Y.: Deeply supervised multi-task autoencoder for biological brain age estimation using three dimensional t _1-weighted magnetic resonance imaging. arXiv preprint arXiv:2508.01565 (2025)

11. Li, M., Liu, J., Yang, M., Zhang, C., Zhao, N., Zhang, Z., Zheng, Q.: Myelination-attention-empowered deep learning model improved brain age prediction in children below 2 years of age. Pediatric Radiology pp. 1–11 (2025)
12. Lin, A., Chen, B., Xu, J., Zhang, Z., Lu, G., Zhang, D.: Ds-transunet: Dual swin transformer u-net for medical image segmentation. IEEE Transactions on Instrumentation and Measurement **71**, 1–15 (2022)
13. 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)
14. Ma, Q., Liang, K., Li, L., Masui, S., Guo, Y., Nosarti, C., Robinson, E.C., Kainz, B., Rueckert, D.: The developing human connectome project: A fast deep learning-based pipeline for neonatal cortical surface reconstruction. Medical Image Analysis **100**, 103394 (2025)
15. 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)
16. Park, S., Kim, H.G., Yang, H., Lee, M., Kim, R.E., Kim, S.H., Styner, M.A., Kim, J., Kim, J.R., Kim, D.: A regional brain volume–based age prediction model for neonates and the derived brain maturation index. European Radiology **34**(6), 3892–3902 (2024)
17. Peng, H., Gong, W., Beckmann, C.F., Vedaldi, A., Smith, S.M.: Accurate brain age prediction with lightweight deep neural networks. Medical image analysis **68**, 101871 (2021)
18. Sun, H., Mehta, S., Khaitova, M., Cheng, B., Hao, X., Spann, M., Scheinost, D.: Brain age prediction and deviations from normative trajectories in the neonatal connectome. Nature communications **15**(1), 10251 (2024)
19. Whitmore, L., Beck, D.: Current challenges and future directions for brain age prediction in children and adolescents. Nature Communications **16**(1), 7771 (2025)
20. Woerl, A.C., Disselhoff, J., Wand, M.: Initialization noise in image gradients and saliency maps. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1766–1775 (2023)
21. Wu, D., Zheng, W., Grant, P.E., Huang, H.: Imaging the developing connectome of perinatal brain (2023)
22. Zhao, H., Cai, H., Liu, M.: Transformer based multi-modal mri fusion for prediction of post-menstrual age and neonatal brain development analysis. Medical Image Analysis **94**, 103140 (2024)