Imaging the Topology of Dynamic Brain Connectivity

Peilin He[1][0009−0002−7382−330X] and Tananun Songdechakraiwut[2][0009−0008−5637−2811] ⋆

[1] Division of Natural and Applied Sciences, Duke Kunshan University, China
[2] Department of Computer Science, Duke University, USA
t.song@duke.edu

Abstract. Functional brain connectivity changes dynamically over time, making its representation challenging for learning on non-Euclidean data. We present a framework that encodes dynamic functional connectivity as an image representation of evolving network topology. Persistent graph homology summarizes global organization across scales, yielding Wasserstein distance-preserving embeddings stable under resolution changes. Stacking these embeddings forms a topological image that captures temporal reconfiguration of brain networks. This design enables convolutional architectures and transfer learning from pretrained foundational models to operate effectively under limited and imbalanced data. Applied to early-stage cognitive impairment assessment on OASIS-3, the approach achieves strong CDR–SB prediction performance, supporting dynamic topological imaging as a promising framework for capturing functional alterations relevant to Alzheimer's disease.

Keywords: Persistent homology · Geometric deep learning · Transfer learning · Alzheimer's disease.

1 Introduction

Learning effective representations of functional brain connectivity remains a central challenge in computational neuroimaging. Functional networks exhibit complex, time-varying organization on non-Euclidean domains, creating a mismatch with models and features designed for static, grid-structured data. Conventional correlation- and graph-based summaries can fit the dominant normal population well, but they collapse temporal evolution and higher-order organization into low-order statistics, reducing sensitivity to subtle, early-stage impairments under severe class imbalance.

This work presents a framework that converts time-resolved functional connectivity graphs into compact image representations of dynamic brain-network topology. Persistent graph homology [1] is used to summarize global network organization across connectivity scales, yielding Wasserstein distance-preserving

⋆ Corresponding author: t.song@duke.edu

embeddings that are stable to resolution changes [2]. Stacking these embeddings over time produces a two-dimensional topological image that represents the temporal evolution of brain-network organization.

This transformation links non-Euclidean brain-network analysis with modern computer vision, enabling convolutional architectures and transfer learning from pretrained medical-imaging models [3] to capture functional dynamics under limited and imbalanced data. Applied to early-stage cognitive impairment assessment on OASIS-3 [4], the framework attains strong prediction accuracy on CDR–SB scores, suggesting sensitivity to subtle functional alterations relevant to early Alzheimer's disease.

2 Topological Image Representation

Functional magnetic resonance imaging (fMRI) measures spontaneous fluctuations in the blood-oxygen-level-dependent (BOLD) signal, which reflect hemodynamic responses coupled to neural activity across the brain. We parcellate the brain into 116 regions of interest (ROIs) using the Automated Anatomical Labeling (AAL) atlas and extract the mean BOLD signal within each ROI. Following the dynamic-connectivity preprocessing in [7], we compute pairwise Pearson correlations using overlapping windows of 5 time points (15,s) with a step size of 2 time points (6,s), yielding a sequence of time-resolved functional connectivity matrices.

Each connectivity matrix is represented as an undirected, weighted graph $G = (V, W)$, where V denotes brain regions and $W = [w_{ij}] \in \mathbb{R}^{|V|\times|V|}$ is a symmetric weight matrix encoding the strength of pairwise functional connections. To extract global and interpretable structure from these graphs, we apply persistent graph homology (PGH) [1], which summarizes the evolution of connected components (0-homology) and cycles (1-homology) across all edge-weight thresholds via closed-form computation. For each graph G we consider a filtration of thresholded subgraphs $G_{\epsilon_0} \supseteq G_{\epsilon_1} \supseteq \cdots \supseteq G_{\epsilon_k}$ obtained by retaining edges with weight greater than $\epsilon_0 \leq \epsilon_1 \leq \cdots \leq \epsilon_k$. We use this superlevel filtration over the connectivity edge weights throughout all experiments. PGH tracks the birth and death of topological features across this filtration, and because the number of connected components (β_0) increases while the number of independent cycles (β_1) decreases monotonically with ϵ [1], the topology of G can be equivalently summarized by the multiset of birth times of components $B(G)$ and death times of cycles $D(G)$.

To obtain a smooth and resolution-flexible representation, we follow [2, 6] and apply inverse transform sampling to the empirical distributions of $B(G)$ and $D(G)$. Concretely, we treat $B(G)$ and $D(G)$ as empirical distributions and sample their quantile functions at fixed grids, yielding two vectors $\mathbf{v}_B \in \mathbb{R}^m$ and $\mathbf{v}_D \in \mathbb{R}^n$ consisting of uniformly spaced birth and death quantiles. This interpolation allows the topological summaries to be resampled at arbitrary resolutions (by varying m and n) while approximately preserving Wasserstein distances between the underlying persistence diagrams, ensuring stability under changes in

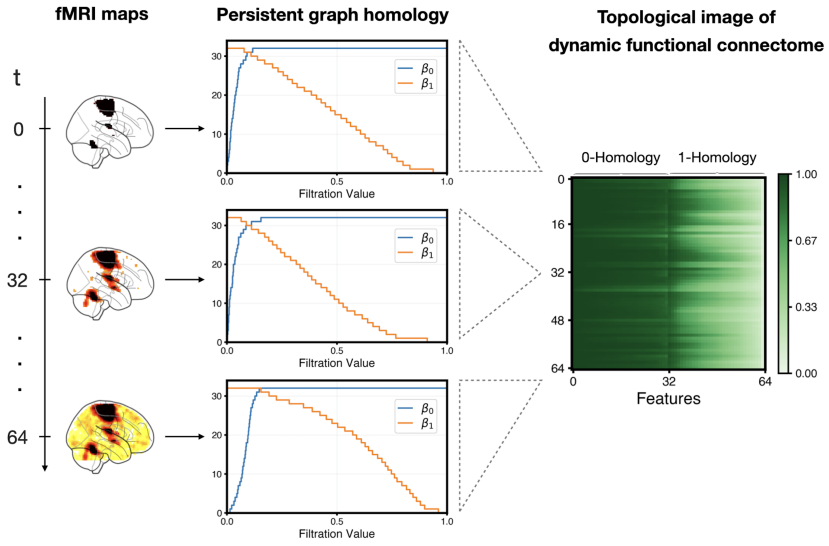

Fig. 1. Topological representation of a dynamic functional connectome. At each time point, fMRI activation maps are converted into functional connectivity graphs and summarized via persistent graph homology, yielding topological curves for 0- and 1-dimensional homology (β_0, β_1) over filtration values. Stacking these summaries across time produces a 2D topological image, where rows correspond to time windows and columns encode features from 0- and 1-homology.

sampling resolution $[6, 2]$. We then concatenate the two embeddings to form a unified topological feature vector $F(G) = [\mathbf{v}_B \,||\, \mathbf{v}_D] \in \mathbb{R}^{m+n}$ that jointly encodes 0- and 1-dimensional homological features. In our implementation, we use $m = 36$ birth quantiles and $n = 37$ death quantiles, yielding 73 topological features per time window.

For a sequence of sliding windows $\{G_t\}_{t=1}^{T}$, we compute $F(G_t)$ for each t and stack the resulting vectors along the temporal axis to form an image-like representation $\mathcal{I} = [F(G_1)^\top; F(G_2)^\top; \ldots; F(G_T)^\top] \in \mathbb{R}^{T\times(m+n)}$. Each row of $\mathcal{I}$ corresponds to the topological summary of one time window, and each column traces the evolution of a sampled topological feature over time. Viewed as a two-dimensional map, the vertical axis encodes time and the horizontal axis encodes topological features; smooth regions indicate stable network topology, whereas localized changes highlight dynamic reconfigurations of connected components and cycles. An overview of the pipeline is given in Fig. 1.

3 Application to Early-Stage Cognitive Impairment Assessment

3.1 Datasets

We evaluate our representation framework using the OASIS-3 dataset [4], which provides longitudinal, multimodal neuroimaging and detailed clinical assessments. The cohort comprises 839 participants with resting-state fMRI and corresponding diagnostic labels: 772 with normal cognition (N), 34 with mild cognitive impairment (MCI), and 33 with cognitive impairment not meeting MCI criteria (IMP). All fMRI data were preprocessed using fMRIPrep with a standardized and reproducible workflow, following preprocessing practices commonly used in prior OASIS-3 dynamic-connectivity studies [7]. Subjects with dementia were excluded to focus on early-stage cognitive decline relevant to Alzheimer's disease. Among the 34 MCI cases, one subject with an extreme Clinical Dementia Rating Sum of Boxes (CDR-SB) score (9.0) was treated as an outlier and excluded. To capture this early impairment spectrum, the MCI and IMP groups were merged into a single category (IMP+MCI) and contrasted with the normal cohort. Data partitioning was performed at the participant level using OASISID, with each participant assigned exclusively to one training, validation, or test partition, thereby preventing longitudinal subject leakage.

3.2 Model configurations

Zero-inflated log-normal (ZILN) loss. To respect the non-negative, zero-inflated, right-skewed distribution of CDR–SB, the proposed CNNs and ZILN-based baseline variants are trained with a ZILN likelihood. For subject i, the network predicts a zero-inflation probability $\pi_i \in (0,1)$ and log-normal parameters (μ_i, σ_i) for $y_i > 0$, with conditional density $p(y_i \mid \pi_i, \mu_i, \sigma_i) = \pi_i \mathbf{1}_{\{y_i=0\}} + (1-\pi_i)\mathbf{1}_{\{y_i>0\}} f_{\mathrm{LN}}(y_i; \mu_i, \sigma_i)$, where $f_{\mathrm{LN}}(y_i; \mu_i, \sigma_i) = (y_i \sigma_i \sqrt{2\pi})^{-1} \exp\big(-(\ln y_i - \mu_i)^2/(2\sigma_i^2)\big)$ for $y_i > 0$. The loss is the average negative log-likelihood

$$\mathcal{L}_{\mathrm{ZILN}} = -\frac{1}{N}\sum_{i=1}^{N}\left[\mathbf{1}_{\{y_i=0\}} \log \pi_i + \mathbf{1}_{\{y_i>0\}} \left(\log(1-\pi_i) + \log f_{\mathrm{LN}}(y_i; \mu_i, \sigma_i)\right)\right].$$

.

Ridge loss for feature-based baselines. Feature-based baselines, including static FC, graph statistics, and TDA features, use linear ridge regression with squared-error loss and ℓ_2 regularization. Given features $\mathbf{x}_i \in \mathbb{R}^d$ and predictions $\hat{y}_i = \mathbf{x}_i^\top \boldsymbol{\beta}$, we minimize

$$\mathcal{L}_{\mathrm{Ridge}} = \frac{1}{N}\sum_{i=1}^{N}\left(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}\right)^2 + \lambda \|\boldsymbol{\beta}\|_2^2.$$

Mean-squared-error (MSE) loss for deep connectome baselines. For MSE-based neural baselines, including the MSE variants of BrainNetCNN [13] and the BNT-style Transformer [14] , the model directly predicts a scalar CDR–SB score $\hat{y}_i$. We train these models by minimizing

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2 .$$

Architecture and optimization We use five ImageNet-1K-pretrained convolutional backbones on the proposed topological images: DenseNet-121, EfficientNet-B0, ShuffleNetV2-0.5, ResNet-18, and ResNet-50. For single-channel inputs, the pretrained first-layer RGB weights are averaged across the channel dimension. Topological images are resized to 160×160 pixels for EfficientNet-B0, ResNet-18, DenseNet-121, and ResNet-50, and to 224×224 pixels for ShuffleNetV2-0.5. For ResNet-18 and ResNet-50, we use small-stem variants, replacing the initial 7×7 stride-2 convolution and first max-pooling operation with a 3×3 stride-1 stem to preserve spatial detail on low-resolution inputs. Models are trained on single-channel topological images for 60 epochs with AdamW and a warmup–cosine learning-rate schedule, using mild in-plane rotations and translations during training. We adopt partial fine-tuning by freezing early backbone stages with a stage-wise freeze fraction of 0.90 and updating the remaining late stages and prediction head. All CNNs are optimized with the ZILN loss described above.

Vector-based baselines We compare the proposed image-based topological models with non-image baselines derived from the same fMRI data. Static functional connectivity (FC) is represented by vectorizing the upper triangular entries of each subject-level FC matrix, following standard resting-state FC analysis [8]. Graph-statistics features are extracted from thresholded FC graphs and include low-order network descriptors commonly used in brain network analysis [10]. PGH-vector features are constructed by flattening classical persistent-homology summaries [11, 12]. For each vector representation, we evaluate ridge regression [9] as a conventional linear baseline and an MLP trained with the ZILN objective as a distribution-aware neural baseline. Hyperparameters are selected using the validation split only, and the selected model is evaluated once on the held-out test split using the same CDR–SB target and MAE-based reporting protocol as the CNN models.

Deep connectome baselines To compare against stronger neural baselines operating directly on connectomes, we further evaluate BrainNetCNN [13] and a Brain Network Transformer (BNT)-style model [14] on static FC matrices. For BrainNetCNN, we include both an MSE regression variant and a matched ZILN variant. For the BNT-style Transformer, each ROI is represented by its static-FC connection profile, and the original classification formulation is adapted to CDR–SB regression by replacing the output head with either an MSE regression

Table 1. CDR–SB regression on the test set. Mean absolute error (MAE) with 95% confidence intervals for impaired (IMP+MCI) and normal (N) subjects.

Type	Method	MAE (95% CI)	
		IMP+MCI	N
Our	EfficientNet-B0	0.7775 (0.7370, 0.8179)	0.3677 (0.3391, 0.3964)
	ResNet-18	0.7251 (0.6602, 0.7899)	0.4451 (0.4149, 0.4753)
	ShuffleNetV2-0.5	0.8573 (0.7637, 0.9509)	0.3806 (0.3380, 0.4232)
	DenseNet-121	0.6667 (0.6035, 0.7299)	0.4981 (0.4175, 0.5786)
	ResNet-50	0.6998 (0.6580, 0.7416)	0.4422 (0.4161, 0.4682)
Baselines	Static FC + Ridge	1.0106 (0.9120, 1.1093)	0.1943 (0.1645, 0.2240)
	Static FC + ZILN	1.0993 (1.0374, 1.1613)	0.0681 (0.0506, 0.0855)
	Graph statistics + Ridge	1.0430 (0.9901, 1.0960)	0.1216 (0.1076, 0.1357)
	Graph statistics + ZILN	1.0666 (1.0302, 1.1031)	0.0985 (0.0944, 0.1026)
	PGH vector + Ridge	0.9969 (0.9055, 1.0884)	0.3001 (0.2528, 0.3474)
	PGH vector + ZILN	1.1063 (1.0684, 1.1441)	0.0728 (0.0600, 0.0857)
	BrainNetCNN + MSE	1.0439 (1.0097, 1.0782)	0.1247 (0.1091, 0.1403)
	BrainNetCNN + ZILN	1.0243 (0.9810, 1.0676)	0.1143 (0.1045, 0.1241)
	BNT-style Transformer + MSE	1.1058 (1.0398, 1.1718)	0.0588 (0.0525, 0.0651)
	BNT-style Transformer + ZILN	1.1324 (1.1036, 1.1613)	0.0309 (0.0234, 0.0384)

head or a three-parameter ZILN head. These deep connectome baselines use the same train/validation/test splits and the same test-set evaluation metrics as the proposed CNN models.

3.3 Results and clinical relevance

Our target is the CDR-SB score. We evaluate CDR–SB regression using a subject-level stratified five-fold protocol. In each split, one fold is reserved for testing; from the remaining development folds, 60% of the full cohort is used for training and the rest for validation. We report mean absolute error (MAE) across the five held-out folds separately for normal (N) and impaired (IMP+MCI) groups. The 95% confidence intervals (CIs) are obtained from 5,000 bootstrap resamples. Across the full imaging cohort, the CDR-SB distribution is highly skewed: among N subjects, 95.5% have CDR-SB = 0 and 4.5% have CDR-SB = 0.5; among IMP+MCI subjects, scores range from 0.5 to 4.0, with 35.4% at 0.5, 38.5% at 1.0, 13.8% at 1.5, and 12.3% at larger than or equal to 2.0. This extreme concentration at 0 in the N group makes its MAE highly sensitive to small deviations from zero, whereas IMP+MCI errors are evaluated over a broader, clinically relevant range.

Table 1 summarizes mean MAEs and 95% CIs for all methods aggregated across the five held-out test folds. On IMP+MCI subjects, all proposed topological CNNs achieve lower point-estimate MAE than all vector-based and deep connectome baselines, with 95% CIs lying entirely below 1.0 CDR–SB point. In contrast, 9 of 10 baseline point estimates exceed 1.0, and the upper bounds of all baseline 95% CIs exceed 1.0. For N subjects, the baselines obtain much smaller MAEs, consistent with the strong concentration of CDR–SB scores near zero in this group, whereas the proposed CNNs exhibit higher N-group MAE. This pattern suggests an observed trade-off between near-zero calibration in normal subjects and severity prediction among impaired subjects. Reported minimum

Table 2. Effect of freezing fraction on ResNet-18 under subject-level 5-fold test evaluation. CIs are computed across fold-level test MAEs.

Frozen fraction %	MAE (95% CI)	
	IMP+MCI	N
0%	0.9043 (0.8046, 1.0041)	0.2188 (0.1639, 0.2737)
17%	0.8981 (0.7835, 1.0128)	0.2321 (0.1647, 0.2996)
33%	0.8559 (0.7140, 0.9979)	0.2699 (0.1847, 0.3552)
50%	0.8542 (0.7496, 0.9588)	0.2605 (0.2080, 0.3129)
67%	0.7967 (0.6900, 0.9034)	0.3086 (0.2300, 0.3871)
90%	0.6966 (0.5768, 0.8163)	0.4127 (0.3106, 0.5147)
100%	0.5997 (0.4906, 0.7088)	0.6440 (0.5465, 0.7414)

clinically important differences (MCIDs) for CDR–SB in mildly impaired populations are approximately 1.0 to 2.0 points [5]; on this cohort, both the point estimates and 95% CIs for IMP+MCI predictions with our models remain below the lower end of this range. These errors are numerically below reported MCID ranges for CDR–SB, providing a useful reference for the scale of prediction error, although MCID and prediction error quantify different constructs.

3.4 Ablation analyses

In this study, we assess how two design choices affect performance, measured by MAE with 95% bootstrap CIs: (i) the fraction of the backbone frozen during training and (ii) the proportion of subjects used for training.

Effect of backbone frozen depth. Table 2 shows that the point-estimate MAE remains below 1.0 for both diagnostic groups across all frozen fractions, although the IMP+MCI confidence intervals slightly exceed 1.0 under the 0% and 17% freezing settings. Increasing the frozen fraction shifts error allocation toward the impaired group: IMP+MCI MAE decreases from 0.9043 at 0% freezing to 0.5997 at 100% freezing, whereas N-group MAE increases from 0.2188 to 0.6440. This illustrates a trade-off between severity prediction in impaired subjects and near-zero calibration in normal subjects. Settings up to 90% freezing preserve the desired ordering MAE for N is less than MAE for IMP+MCI, whereas full freezing reverses this ordering despite achieving the lowest IMP+MCI MAE.

Effect of training set size. Table 3 shows that training-set size has a limited effect on IMP+MCI prediction under the current setup. Across 1%–60% training fractions, IMP+MCI MAE remains in a narrow range (from 0.6529 to 0.6869), with all 95% CIs below 0.74. In contrast, N group MAE gradually improves with more training data, reaching its lowest value at 60%. Thus, additional training subjects mainly improve near-zero calibration for normal subjects, while impaired-subject prediction is relatively stable across training fractions.

Table 3. Training-set size ablation for CDR–SB regression. Mean absolute error (MAE) with 95% confidence intervals for impaired (IMP+MCI) and normal (N) subjects under subject-level stratified five-fold evaluation.

Train %	Train IMP+MCI / N	MAE (95% CI)	
		IMP+MCI	N
1%	1 / 7	0.6529 (0.6131, 0.6926)	0.5006 (0.4277, 0.5736)
5%	3 / 39	0.6538 (0.6121, 0.6955)	0.4777 (0.4700, 0.4853)
10%	7 / 77	0.6723 (0.6285, 0.7160)	0.4576 (0.4337, 0.4815)
20%	13 / 155	0.6625 (0.6194, 0.7055)	0.4681 (0.4564, 0.4797)
30%	19–20 / 231–232	0.6785 (0.6605, 0.6965)	0.4534 (0.4351, 0.4718)
40%	26 / 309	0.6709 (0.6118, 0.7299)	0.4576 (0.4246, 0.4905)
50%	33 / 386	0.6755 (0.6557, 0.6953)	0.4569 (0.4262, 0.4877)
60%	39–40 / 463–464	0.6869 (0.6420, 0.7318)	0.4458 (0.4195, 0.4721)

4 Discussion

Encoding the topology of dynamic functional connectivity into image-like representations and training CNNs to regress CDR–SB offers a potentially low-cost, scalable tool for early screening, progression monitoring, and clinical-trial enrichment in Alzheimer's disease, with natural extensions to other neuropsychiatric disorders and longitudinal neuroimaging studies. However, the reasoning chain from time×scale topological image features to clinical conclusions remains insufficiently transparent: learned convolutional features are not readily attributable to specific region-to-region pathways or transient network events, which may limit uptake in clinical settings that require interpretable evidence. Therefore, improving transparency should be a priority. In addition, it is important to involve clinicians in the evaluation and refinement of explanation methods to build trust and facilitate practical deployment. The current comparison does not fully disentangle the contribution of topological encoding from that of temporal information, since the proposed models use dynamic connectivity whereas the non-topological baselines use static connectivity. A matched dynamic non-topological baseline would help isolate these effects in future work.

5 Compliance with ethical standards

This research study was conducted retrospectively using human subject data made available in open access by OASIS-3 project [4]. Ethical approval was not required as confirmed by the license attached with the open access data.

References

1. Songdechakraiwut, T., Chung, M.K.: Topological learning for brain networks. *The Annals of Applied Statistics* **17**(1), 403–433 (2023)
2. Skraba, P., Turner, K.: Wasserstein stability for persistence diagrams. arXiv preprint arXiv:2006.16824 (2020)
3. Mienye, I.D., Swart, T.G., Obaido, G., Jordan, M., Ilono, P.: Deep convolutional neural networks in medical image analysis: A review. *Information* **16**(3), 195 (2025)

4. LaMontagne, P.J., Benzinger, T.L.S., Morris, J.C., et al.: OASIS-3: Longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer disease. *medRxiv* (2019). https://doi.org/10.1101/2019.12.13.19014902
5. Muir, R.T., Hill, M.D., Black, S.E., Smith, E.E.: Minimal clinically important difference in Alzheimer's disease: Rapid review. *Alzheimer's & Dementia* **20**(5), 3352–3363 (2024)
6. Songdechakraiwut, T., Krause, B.M., Banks, M.I., Nourski, K.V., Van Veen, B.D.: Wasserstein distance-preserving vector space of persistent homology. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2023, pp. 277–286 (2023)
7. Yi, L., Lutz, M.W., Wu, Y., Li, Y., Songdechakraiwut, T.: Topological and geometric signatures of brain network dynamics in Alzheimer's disease. *Alzheimer's & Dementia* **21**(8), e70545 (2025). https://doi.org/10.1002/alz.70545
8. Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S.: Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine* **34**(4), 537–541 (1995). https://doi.org/10.1002/mrm.1910340409
9. Hoerl, A.E., Kennard, R.W.: Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **12**(1), 55–67 (1970). https://doi.org/10.1080/00401706.1970.10488634
10. Rubinov, M., Sporns, O.: Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage* **52**(3), 1059–1069 (2010). https://doi.org/10.1016/j.neuroimage.2009.10.003
11. Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. *Discrete & Computational Geometry* **28**, 511–533 (2002). https://doi.org/10.1007/s00454-002-2885-2
12. Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discrete & Computational Geometry* **33**(2), 249–274 (2005). https://doi.org/10.1007/s00454-004-1146-y
13. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* **146**, 1038–1049 (2017). https://doi.org/10.1016/j.neuroimage.2016.09.046
14. Kan, X., Dai, W., Cui, H., Zhang, Z., Guo, Y., Yang, C.: Brain network transformer. In: *Advances in Neural Information Processing Systems*, vol. 35, pp. 25586–25599 (2022).