



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.

Network-Aware Bilinear Tokenization for Brain Functional Connectivity Representation Learning

Leo Milecki¹, Qingyu Hu^{1,2}, Bahram Jafarsteh¹, Mert R. Sabuncu^{1,2}, and Qingyu Zhao¹

¹ Department of Radiology, Weill Cornell Medicine, New York, NY, USA.

² School of Electrical and Computer Engineering, Cornell University and Cornell Tech, New York, NY, USA. lem4012@med.cornell.edu

Abstract. Masked autoencoders (MAEs) have recently shown promise for self-supervised representation learning of resting-state brain functional connectivity (FC). However, a fundamental question remains unresolved: how should FC matrices be tokenized to align with the intrinsic modular organization of large-scale brain networks? Existing approaches typically adopt region-centric or graph-based schemes that treat FC as structurally homogeneous elements and overlook the large-scale network brain organization. We introduce NERVE (Network-Aware Representations of Brain Functional Connectivity via Bilinear Tokenization), a self-supervised learning framework that redefines FC tokenization by partitioning FC matrices into patches of intra- and inter-network connectivity blocks. Unlike image-based MAE, where fixed-size patches share a common tokenizer, FC patches defined by network pairs are heterogeneous in size and correspond to distinct functional roles. To resolve this problem, NERVE embeds FC patches through a novel structured bilinear factorization. This formulation preserves network identity and reduces parameter complexity from quadratic to linear scaling in the number of networks. We evaluate NERVE across three large-scale developmental cohorts (ABCD, PNC, and CCNP) for behavior and psychopathology prediction. Compared to structurally agnostic MAE variants and graph-based self-supervised baselines, the proposed network-aware formulation yields more stable and transferable representations, particularly in cross-cohort evaluation. Ablation studies confirm that the proposed bilinear network embedding and anatomically grounded parcellation are critical for performance. These findings highlight the importance of incorporating domain-specific structural priors into self-supervised learning for functional connectomics. Code is available at: <https://github.com/leomlck/NERVE>

Keywords: rs-fMRI · Brain Functional Connectivity · Deep Learning.

1 Introduction

Resting-state functional magnetic resonance imaging (rs-fMRI) enables the estimation of functional connectivity (FC), defined as the temporal correlation be-

tween spatially distributed brain regions. FC has become a central tool for studying individual differences in large-scale brain organization and their association with cognition, behavior, and mental health [15,17,14,29]. However, extracting compact predictive representations from FC remains challenging due to its high dimensionality, low signal-to-noise ratio, and substantial inter-subject variability [31,2,28]. Indeed, large-scale studies have reported that increasing model complexity does not reliably improve performance over classical approaches [9,23,27,10], suggesting that more appropriate inductive biases may be required.

Masked autoencoders (MAE) [8] provide a principled framework for representation learning by partitioning the input into units, embedding each unit into a token representation, masking a subset of tokens, and reconstructing the masked content from the visible ones. In computer vision, these units correspond to spatial image patches that naturally align with local structure. When adapted to FC, however, the notion of a “patch” lacks a canonical definition, as the arrangement of regions in an FC matrix is largely arbitrary and does not necessarily reflect spatial or functional locality. Defining appropriate patch units and their corresponding token embeddings thus becomes central for applying MAE to FC. Existing approaches for FC adopt heuristic tokenization schemes prior to masking [4,32,3,20,6]. BrainMass [32], for example, treats individual regions as units and masks randomly selected rows of the FC matrix, while RS-MAE [20] masks grouped regions inspired by spatiotemporal strategies. Graph-based methods similarly operate on node-level embeddings, where masking is applied to node tokens [5,12,22,30]. Despite promising results, a fundamental question remains: how should FC matrices be partitioned into patches and embedded into tokens to align the learned representations with the modular, network-level organization of brain FC?

Our central insight is that the conceptual analog of spatially neighboring pixels in an image is groups of brain regions that share similar functional dynamics, e.g., regions organized into large-scale functional networks [33]. Under this view, the natural counterpart of an image patch is a connectivity block defined by intra- or inter-network interactions. However, a key challenge arises: FC patches defined by pairs of networks vary in dimensionality, precluding the use of a shared patch encoder. Moreover, each network carries distinct functional roles, suggesting that tokenization should preserve network identity rather than collapse all patches into a homogeneous representation space. To address this, we introduce a novel and parameter-efficient *bilinear tokenization* scheme. Instead of learning independent embeddings for each network-pair patch, we learn network-specific region embeddings and model inter-network connectivity via bilinear interactions. This factorization replaces quadratic growth in patch-specific parameters with linear scaling in the number of networks, while explicitly encoding network identity and structured intra- and inter-network interactions.

Integrating this design into an MAE framework, we introduce **NERVE** (Network-Aware Representations of Brain Functional Connectivity via Bilinear Tokenization), a self-supervised approach tailored to the modular organization of brain FC. We evaluate NERVE across three large-scale adolescent neuroimag-




Fig. 1. Overview of NERVE. **A.** The functional connectivity (FC) matrix is partitioned into patches defined by pairs of functional brain networks. **B. Network-aware Bilinear Tokenization.** Each functional network is assigned learnable network-specific weights at initialization, and patch tokens are computed through structured bilinear interactions between network weights during forward. **C. MAE Framework.** We apply a standard MAE framework to the proposed network-aware tokens, thereby introducing a functionally informed inductive bias over connectivity structure.

ing cohorts on the challenging task of predicting behavioral and psychopathology scores. Our results demonstrate that NERVE learns more informative and transferable FC representations, outperforming alternative tokenization strategies and existing self-supervised learning (SSL) approaches.

2 Methods

Let $\mathcal{D} = \{X^{(i)}\}_{i=1}^N$ denote FC matrices of N participants, where each $X^{(i)} \in \mathbb{R}^{R \times R}$ is a correlation matrix constructed from functional time series across R brain regions. We aim to learn structured and transferable representations of X in a self-supervised manner. To this end, we adopt a transformer-based MAE framework, which partitions X into patches, encodes each patch into a token, randomly masks a subset of tokens, and reconstructs the masked content from the visible tokens.

Network-based Patching. The way data are partitioned into patches determines what structure the model can exploit through masking and reconstruction. While image-based MAE relies on a natural decomposition into fixed-size image patches, FC matrices lack a canonical patching scheme, making tokenization a non-trivial design choice. Rather than treating rows of X as patches [4, 32, 3], we observe that the conceptual analog of neighboring pixels in images is groups of brain regions that share similar functional dynamics. This organization is captured by large-scale functional networks [33]. Therefore, we propose to reorganize X by grouping the R regions into N_n established functional networks $\{\mathcal{N}_1, \dots, \mathcal{N}_{N_n}\}$ (e.g., Visual, Default, Dorsal Attention). For each network pair (l, m) with $l \leq m$, we define a connectivity block: $x_{l,m} \in \mathbb{R}^{|\mathcal{N}_l| \times |\mathcal{N}_m|}$ representing

intra- ($l = m$) or inter-network ($l < m$) connectivity (Fig. 1A). These connectivity patches are then treated as the basic units for masking and reconstruction. The total number of patches is: $N_{\text{patch}} = \frac{N_n(N_n+1)}{2}$.

Shared vs. Patch-specific Tokenization. In image-based MAE, patches are fixed-size and interchangeable, enabling a shared projection for tokenization. Here, network-defined patches vary in size and correspond to specific network interactions, making them structurally and semantically distinct. Embedding these irregular and network-specific patches, therefore, requires a dedicated network-aware tokenization strategy. A straightforward adaptation of image-based MAE to FC would flatten each patch $x_{l,m}$ and project it through a linear transformation (denoted as *shared* linear): $t_{l,m} = W^\top \text{vec}(x_{l,m})$, where $W \in \mathbb{R}^{S_{\max} \times d_E}$ with d_E the encoder embedding dimension, and $S_{\max} = \max_{l,m} |\mathcal{N}_l| \cdot |\mathcal{N}_m|$ the maximum flattened patch size. Because patches have different sizes, zero-padding to the largest patch size $S_{\max}$ is required. While parameter-efficient, this shared projection enforces a common representation across semantically distinct network interactions and ignores the structural heterogeneity of FC. An alternative is to assign a distinct projection layer to each network pair (denoted as *patch-specific* linear): $t_{l,m} = W_{l,m}^\top \text{vec}(x_{l,m})$, where $W_{l,m} \in \mathbb{R}^{(|\mathcal{N}_l||\mathcal{N}_m|) \times d_E}$. Although this allows patch-specific (network pairs) modeling, it introduces a quadratic growth in parameters with respect to the number of networks N_n , which quickly becomes impractical and risks overfitting.

Bilinear Tokenization. To preserve network specificity while maintaining parameter efficiency, we propose a bilinear network-aware tokenization. Each functional network $\mathcal{N}_l$ is assigned a learnable matrix $U_l \in \mathbb{R}^{|\mathcal{N}_l| \times d_E}$, where each column represents a network-specific embedding dimension. For an FC patch $x_{l,m} \in \mathbb{R}^{|\mathcal{N}_l| \times |\mathcal{N}_m|}$ between networks l and m , we construct the corresponding tokenizer via a column-wise Kronecker (Khatri–Rao) [16] product (Fig. 1B):

$$W_{l,m} = U_l \odot U_m \in \mathbb{R}^{(|\mathcal{N}_l||\mathcal{N}_m|) \times d_E}, \quad t_{l,m} = W_{l,m}^\top \text{vec}(x_{l,m}) \in \mathbb{R}^{d_E}.$$

where $\odot$ denotes the Khatri–Rao product defined elementwise as $[W_{l,m}]_{(i,j),k} = [U_l]_{i,k} [U_m]_{j,k}$. Conceptually, instead of learning an independent projection for each network pair, we learn network-level region embeddings and model inter-network connectivity as bilinear interactions between them. The contribution of a connection between region $i \in \mathcal{N}_l$ and region $j \in \mathcal{N}_m$ to embedding dimension k is given by the product $[U_l]_{i,k} [U_m]_{j,k}$, capturing structured interactions while sharing parameters across networks. This low-rank factorization of patch-specific projections replaces quadratic growth in patch-specific parameters with a linear scaling in the number of networks, while explicitly encoding network identity and structured intra- and inter-network interactions. Specifically, if we assume the R brain regions are approximately evenly distributed across the N_n networks, we can approximate $S_{\max} \approx (R/N_n)^2$. Under this assumption, the parameter complexities for *shared* linear, *patch-specific* linear, and *bilinear* embedding are: $\mathcal{O}\left(\frac{R^2}{N_n^2} \times d_E\right)$, $\mathcal{O}(R^2 \times d_E)$, and $\mathcal{O}(R \times d_E)$, respectively, highlighting the trade-off between parameter efficiency and expressivity by shifting complexity from universal, patch-specific to network-centric representations.

MAE for FC. Given the network-aware token sequence $T \in \mathbb{R}^{N_{\text{patch}} \times d_E}$, a learnable CLS token is prepended, and learnable positional embeddings are added to encode the identity of each network pair. For each subject, a fixed proportion of patch tokens is randomly masked, and only the visible tokens are processed by a transformer encoder, producing contextualized representations. To reconstruct masked connectivity, mask tokens are first inserted at masked positions, and the full token sequence is then processed by a lightweight transformer decoder. The decoded tokens are then projected back to the FC patch space using a bilinear decoding layer consistent with the tokenization scheme, yielding reconstructed patches $\hat{x}_{l,m}$ (Fig. 1C). The reconstruction objective is:
$$\mathcal{L}_{\text{recon}} = \frac{1}{|\mathcal{M}|} \sum_{(l,m) \in \mathcal{M}} \|\hat{x}_{l,m} - x_{l,m}\|_2^2.$$

Implementation Details. We use Schaefer’s 17-network parcellation ($R = 400$, $N_n = 17$, $N_{\text{patch}} = 153$) [26] with a masking ratio of 0.5. The encoder uses 4 layers, 4 heads, and $d_E = 256$. The decoder has 1 layer, 2 heads with $d_D = 64$. Models are trained for 4,000 epochs using AdamW with a learning rate of $1e^{-2}$, weight decay of $1e^{-2}$, cosine scheduling with linear warmup for 400 epochs, batch size of 1,024, and mixed precision using Pytorch on a single NVIDIA L40 GPU. We perform preliminary hyperparameter tuning on the training data for the masking ratio, transformer encoder, and decoder length, width, and learning rate. For brevity, we report results using the best-performing configuration and focus ablations on methodological design choices central to our contribution.

3 Experiments and Results

Datasets & Downstream Evaluation. We evaluate our framework on three large-scale developmental rs-fMRI cohorts: ABCD (N=1,791, age 9–10, 53.8% F) [7], PNC (N=1,416, age 8–23, 53.6% F) [25], and CCNP (N=178, age 6–17, 52.3% F) [19]. For ABCD, we directly use the FC matrices preprocessed and released by CBIG [21], which contain a subset of the full ABCD cohort with the highest image quality. We preprocess PNC and CCNP data using the publicly available DeepPrep pipeline [24], which includes motion correction, nuisance regression, and temporal and spatial filtering. All FC matrices are defined by the $R = 400$ regions Schaefer atlas [26]. We train the MAE on the FC matrices from the combined ABCD+PNC datasets, treating CCNP as out-of-domain samples, and apply the trained model to extract FC representations for all three datasets. Then, within each dataset, we use the representations to predict behavioral phenotypes. For ABCD, we predict the *Internalizing*, *Externalizing*, and *Total* scores from the Child Behavior Checklist (CBCL) [1]. For PNC and CCNP, we predict the Reproducible Brain Charts (RBC) [11] harmonized *Internalizing*, *Externalizing*, and general psychopathology (*p-factor*) scores. Kernel Ridge Regression (KRR) is adopted as the downstream regression model following prior large-scale evaluations demonstrating its strong and stable performance for behavioral prediction compared to more complex neural network-based probes [9]. KRR is assessed using stratified 10-fold cross-validation within each dataset, with age and sex effects removed via linear regressions fit on the training folds. Finally,

Table 1. Behavioral prediction performance across developmental cohorts. We report Pearson correlation $r_{\pm\Delta}$ with bootstrap 95% CI half-width Δ from 10-fold cross-validation. n.: negative values, OOD: Out-of-domain.

Dataset →	ABCD			PNC			CCNP (OOD)		
Method ↓	Int.	Ext.	Total	Int.	Ext.	p-factor	Int.	Ext.	p-factor
BrainNetCNN	.05 $\pm$.04	.08 $\pm$.04	.11 $\pm$.04	.07 $\pm$.03	.08 $\pm$.04	.05 $\pm$.03	.15 $\pm$.12	.10 $\pm$.13	.01 $\pm$.12
BrainGNN	.05 $\pm$.04	.08 $\pm$.05	.06 $\pm$.04	.08 $\pm$.04	.04 $\pm$.04	.06 $\pm$.04	.17 $\pm$.13	.16 $\pm$.16	.04 $\pm$.10
BrainNetTF	—	—	—	.08 $\pm$.04	.07 $\pm$.05	.05 $\pm$.04	.01 $\pm$.09	.21 $\pm$.14	.05 $\pm$.09
GraphMAE	.06 $\pm$.05	.06 $\pm$.04	.06 $\pm$.04	.09 $\pm$.05	.08 $\pm$.05	.09 $\pm$.05	.07 $\pm$.16	.22 $\pm$.16	.12 $\pm$.13
GATE	—	—	—	.07 $\pm$.06	.08 $\pm$.05	.09 $\pm$.04	.04 $\pm$.09	.14 $\pm$.14	.11 $\pm$.08
BrainGSLs	.04 $\pm$.05	.05 $\pm$.05	.05 $\pm$.05	.05 $\pm$.05	.08 $\pm$.06	.04 $\pm$.05	n.	.22 $\pm$.15	n.
BrainMass	.06 $\pm$.05	.08 $\pm$.05	.01 $\pm$.04	.03 $\pm$.06	.09 $\pm$.06	.10 $\pm$.05	.04 $\pm$.13	.14 $\pm$.14	.13 $\pm$.16
NERVE	.11 $\pm$.06	.09 $\pm$.05	.13 $\pm$.04	.14 $\pm$.05	.13 $\pm$.05	.12 $\pm$.05	.08 $\pm$.14	.33 $\pm$.14	.13 $\pm$.15

Gray shading: $p < 0.05$, two-sided paired bootstrap test between top-1 and 2 methods.

we report the Pearson correlation between the true and predicted behavioral scores concatenated across testing folds. Uncertainty is quantified via bootstrap resampling over subjects (1,000 iterations) to estimate 95% confidence intervals (CI). The statistical significance of performance differences between the top two methods was evaluated using two-sided paired bootstrap tests on subject-level out-of-fold predictions (1,000 iterations) for each behavioral variable.

Comparison Methods. We benchmark NERVE against established methods for FC encoding, restricting comparisons to approaches operating on FC matrices to ensure fair and interpretable evaluation. Compared encoders include BrainMass [32], a foundation-style model combining temporal masking and contrastive learning, and graph-based SSL methods: GraphMAE [12], BrainGSLs [30], and GATE [22], which apply node-masked graph encoders with SSL objectives. All models are assessed using the same protocol: training encoders on ABCD+PNC and evaluating downstream KRR within each dataset via 10-fold cross-validation. To put the prediction performance in the context of fully supervised, task-specific learning capacity, we also evaluate state-of-the-art supervised FC-based prediction models, which include BrainNetCNN [15], BrainGNN [17], and BrainNetTF [14]. These supervised models are trained and evaluated end-to-end within each dataset by 10-fold cross-validation, as behavioral targets derive from different assessment instruments and are not strictly interchangeable across cohorts. GATE and BrainNetTF require BOLD time series preprocessing and are therefore evaluated only on PNC and CCNP, as only FC matrices are available for ABCD in the CBIG release.

Main Results. Past rs-fMRI studies have consistently shown that continuous behavioral prediction from FC is intrinsically challenging, with state-of-the-art models typically achieving modest effect sizes (often $r < 0.15$) even in large cohorts [21,9]. In this context, the representations learned by NERVE demonstrate strong and stable predictive power across datasets (Table 1). Specifically, in in-domain evaluations (ABCD, PNC), NERVE achieves the highest or tied-

Table 2. Tokenization ablation. We report Pearson correlation $r_{\pm\Delta}$ with bootstrap 95% CI half-width Δ from 10-fold cross-validation. n. indicates negative values.

Dataset $\rightarrow$	ABCD			PNC			CCNP		
Method $\downarrow$	Int.	Ext.	Total	Int.	Ext.	p-factor	Int.	Ext.	p-factor
<i>shared</i> linear	.06 $\pm$.06	.06 $\pm$.05	.02 $\pm$.05	.10 $\pm$.05	.11 $\pm$.05	.10 $\pm$.06	.06 $\pm$.10	.31 $\pm$.15	.12 $\pm$.15
<i>specific</i> linear	.03 $\pm$.08	.08 $\pm$.05	.06 $\pm$.05	.14 $\pm$.06	.12 $\pm$.05	.09 $\pm$.06	.01 $\pm$.11	.03 $\pm$.13	n.
<i>bilinear</i>	.11 $\pm$.09	.09 $\pm$.05	.13 $\pm$.04	.14 $\pm$.05	.13 $\pm$.05	.12 $\pm$.05	.08 $\pm$.14	.33 $\pm$.14	.13 $\pm$.15

Table 3. Parcellation and network structure ablation. We report Pearson r on 10-fold cross-validation. For the *Permutation* baselines, we report mean $\pm$ std r over 100 random permutations of region-to-network assignment.

Dataset $\rightarrow$	ABCD			PNC			CCNP		
Parcellation $\downarrow$	Int.	Ext.	Total	Int.	Ext.	p-factor	Int.	Ext.	p-factor
Permutation	.04 $\pm$.02	.07 $\pm$.02	.07 $\pm$.03	.09 $\pm$.02	.10 $\pm$.03	.09 $\pm$.01	-.04 $\pm$.09	.19 $\pm$.05	.00 $\pm$.06
16 $\times$ 16 <i>shared</i>	.02	.02	.05	.05	.09	.11 †	-.03	.11	.01
16 $\times$ 16 <i>specific</i>	.02	.10 †	.06	.04	.09	.08	-.02	.25 †	.11 †
16 $\times$ 16 <i>bilinear</i>	.01	.08	.06	.05	.10	.08	.00	.11	-.05
7-network	.05	.09 †	.09 †	.07	.10	.07	-.01	.22	.07 †
17-network	.11 †	.09 †	.13 †	.14 †	.13 †	.12 †	.08 †	.33 †	.13 †

Bold † : $p < 0.05$ by one-sided permutation test compared to *Permutation* baselines.

highest performance across all behavioral targets, supporting the advantage of incorporating network-level structure over structurally agnostic SSL baselines and supervised FC architectures trained within a single dataset. Among self-supervised baselines, GraphMAE and BrainMass show competitive performance on specific targets but exhibit higher variability across targets. In the out-of-domain CCNP cohort, NERVE exhibits strong cross-cohort generalization, achieving the highest observed correlation for externalizing symptoms ($r = .33$) and competitive performance for the general psychopathology (p-factor) score. These results further support the role of network-aware tokenization in enhancing representation stability beyond the training distribution. The only target where NERVE does not achieve top performance is the internalizing phenotype in the CCNP cohort, where supervised models trained specifically on CCNP obtain higher correlations.

Tokenization Ablation. We compare three tokenization strategies within the NERVE framework: i) *shared* linear, ii) patch-*specific* linear, and iii) the proposed *bilinear* network-aware formulation (Table 2). The bilinear formulation consistently achieves the strongest overall performance, particularly on in-domain datasets, while the patch-specific variant performs second-best in most settings. The shared linear formulation remains competitive on the out-of-domain CCNP samples, consistent with the idea that reduced parameterization benefits generalization. Overall, the proposed bilinear strategy achieves

the strongest predictive performance, supporting the effectiveness of combining constrained parameterization with an explicit network-level inductive bias.

Network Structure and Parcellation. To evaluate the interaction between network-based patching and the proposed network-specific bilinear tokenization, we design controlled ablations that modify patch definitions while keeping all other architectural and training components fixed, isolating the effect of anatomically grounded tokenization in NERVE. First, we perform 100 permutation baselines by randomly shuffling region indices prior to patch assignment within the 17-network parcellation [33]. This serves as a null model without meaningful network structure. Second, we evaluate a vanilla MAE configuration that partitions the FC matrix into fixed continuous 16×16 square patches following image-based practice [8]. We then encode these patches using the *shared*, *patch-specific*, and proposed *bilinear* network-aware formulation, thereby removing explicit functional network grounding. Third, we test a coarser 7-network parcellation [33] to assess sensitivity to network granularity. Table 3 reports performance and one-sided permutation test p-values computed against 100 region-permuted baselines. NERVE with the 17-network parcellation achieves the highest performance and significantly outperforms the baselines, supporting the relevance of anatomically informed patch definitions. Notably, when network structure is disrupted, patch-specific linear tokenization tends to outperform the bilinear formulation, suggesting that the proposed bilinear factorization is most effective when aligned with meaningful functional organization.

4 Discussion & Conclusion

Our results consistently demonstrate that incorporating network-level structure improves stability and cross-cohort generalizability of FC representations. Nevertheless, several limitations should be acknowledged. First, evaluation was conducted within developmental cohorts and psychopathology measures. While this setting reflects a clinically relevant and challenging prediction task, further evaluation across additional populations and phenotypic domains will help assess the generalizability of the learned representations. Second, our framework currently operates on static FC derived from a predefined parcellation. Extensions incorporating dynamic FC patterns [13] or multimodal integration with structural connectivity derived from diffusion MRI [18] represent promising directions for extending the proposed framework. Future work will focus on scaling to more heterogeneous datasets, exploring alternative downstream tasks, and improving interpretability of the learned representations. In particular, future analyses will examine the learned network-specific weights in the bilinear embedding and characterize attention patterns between network-aware tokens to better understand how large-scale network connectivity patterns across parcellations contribute to representation learning.

In this work, we introduced a self-supervised framework that incorporates large-scale functional network structure into brain connectivity representation learning through a bilinear tokenization strategy. Our results demonstrate that

network-aware modeling provides a principled inductive bias for functional connectomics and supports scalable cross-cohort brain-behavior mapping.

Acknowledgments. This work was supported in part by NIH Grant AA028840 (QZ), a BBRF Young Investigator Grant (QZ), and a NIARR Pilot Resource Grant (QZ).

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

References

1. Achenbach, T.M., Edelbrock, C.S.: The classification of child psychopathology: A review and analysis of empirical efforts. *Psychological Bulletin* **85**(6), 1275–1301 (1978)
2. Button, K.S., Ioannidis, J.P., Mokrysz, C., Nosek, B.A., Flint, J., Robinson, E.S., et al.: Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience* **14**(5), 365–376 (2013)
3. Caro, J.O., Fonseca, A.H.d.O., Averill, C., Rizvi, S.A., Rosati, M., Cross, J.L., et al.: BrainLM: A foundation model for brain activity recordings. In: ICLR. Curran Associates, Inc. (2024)
4. Dong, Z., Li, R., Wu, Y., Nguyen, T., Su, J., Chong, et al.: Brain-JEPA: Brain Dynamics Foundation Model with Gradient Positioning and Spatiotemporal Masking. In: NeurIPS. vol. 37, pp. 86048–86073. Curran Associates, Inc. (2024)
5. Farahani, F.V., Karwowski, W., Lighthall, N.R.: Application of graph theory for identifying connectivity patterns in human brain networks: A systematic review. *Frontiers in Neuroscience* **13** (2019)
6. Gao, J., Ge, B., Qiang, N., Zhao, S.: 3D masked autoencoder with spatiotemporal transformer for modeling of 4D fMRI data. *Medical Image Analysis* **107**(Pt B), 103861 (2026)
7. Garavan, H., Bartsch, H., Conway, K., Decastro, A., Goldstein, R.Z., Heeringa, S., et al.: Recruiting the ABCD sample: design considerations and procedures. *Developmental Cognitive Neuroscience* **32**, 16–22 (2018)
8. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked Autoencoders Are Scalable Vision Learners. In: CVPR. IEEE, Inc. (2021)
9. He, T., Kong, R., Holmes, A.J., Nguyen, M., Sabuncu, M.R., et al.: Deep neural networks and kernel regression achieve comparable accuracies for functional connectivity prediction of behavior and demographics. *NeuroImage* **206** (2020)
10. He, T., Kong, R., Holmes, A.J., Sabuncu, M.R., Eickhoff, S.B., Bzdok, et al.: Is deep learning better than kernel regression for functional connectivity prediction of fluid intelligence? In: PRNI. IEEE, Inc. (2018)
11. Hoffmann, M.S., Moore, T.M., Axelrud, L.K., Tottenham, N., Pan, P.M., Miguel, et al.: An Evaluation of Item Harmonization Strategies Between Assessment Tools of Psychopathology in Children and Adolescents. *Assessment* **31**(2), 502–517 (2024)
12. Hou, Z., Liu, X., Cen, Y., Dong, Y., Yang, H., Wang, C., et al.: GraphMAE: Self-Supervised Masked Graph Autoencoders. In: SIGKDD. pp. 594–604. Association for Computing Machinery (2022)
13. Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, et al.: Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage* **80**, 360–378 (2013)

14. Kan, X., Dai, W., Cui, H., Zhang, Z., Guo, Y., Yang, C.: Brain Network Transformer. In: NeurIPS. vol. 35. Curran Associates, Inc. (2022)
15. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., et al.: BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* **146**, 1038–1049 (2017)
16. Khatri, C.G., Radhakrishna Rao, C.: Solutions to Some Functional Equations and Their Applications to Characterization of Probability Distributions. *The Indian Journal of Statistics* **30**(2), 167–180 (1968)
17. Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, et al.: BrainGNN: Interpretable Brain Graph Neural Network for fMRI Analysis. *Medical Image Analysis* **74**, 102233 (2021)
18. Litwińczuk, M.C., Muhlert, N., Cloutman, L., Trujillo-Barreto, N., Woollams, A.: Combination of structural and functional connectivity explains unique variation in specific domains of cognitive function. *NeuroImage* **262**(3), 119531 (2022)
19. Liu, S., Wang, Y.S., Zhang, Q., Zhou, Q., Cao, L.Z., Jiang, C., et al.: Chinese Color Nest Project : An accelerated longitudinal brain-mind cohort. *Developmental Cognitive Neuroscience* **52**, 101020 (2021)
20. Ma, H., Xu, Y., Tian, L.: RS-MAE: Region-State Masked Autoencoder for Neuropsychiatric Disorder Classifications Based on Resting-State fMRI. *IEEE transactions on neural networks and learning systems* **36**(6), 10707–10720 (2025)
21. Ooi, L.Q.R., Chen, J., Zhang, S., Kong, R., Tam, A., Li, J., et al.: Comparison of individualized behavioral predictions across anatomical, diffusion and functional connectivity MRI. *NeuroImage* **263**, 119636 (2022)
22. Peng, L., Wang, N., Xu, J., Zhu, X., Li, X.: GATE: Graph CCA for Temporal Self-Supervised Learning for Label-Efficient fMRI Analysis. *IEEE Transactions on Medical Imaging* **42**(2), 391–402 (2023)
23. Pervaiz, U., Vidaurre, D., Woolrich, M.W., Smith, S.M.: Optimising network modelling methods for fMRI. *NeuroImage* **211**, 116604 (2020)
24. Ren, J., An, N., Lin, C., Zhang, Y., Sun, Z., Zhang, et al.: DeepPrep: an accelerated, scalable and robust pipeline for neuroimaging preprocessing empowered by deep learning. *Nature Methods* **22**(3), 473–476 (2025)
25. Satterthwaite, T.D., Elliott, M.A., Ruparel, K., Loughhead, J., Prabhakaran, K., Calkins, et al.: Neuroimaging of the Philadelphia Neurodevelopmental Cohort. *NeuroImage* **86**, 544–553 (2014)
26. Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.N., Holmes, et al.: Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral cortex* **28**(9), 3095–3114 (2018)
27. Schulz, M.A., Yeo, B.T., Vogelstein, J.T., Mourao-Miranada, J., Kather, J.N., Kording, K., et al.: Different scaling of linear models and deep learning in UKBiobank brain images versus machine-learning datasets. *Nature Communications* **11**(1), 1–15 (2020)
28. Tiego, J., Martin, E.A., DeYoung, C.G., Hagan, K., Cooper, S.E., Pasion, et al.: Precision behavioral phenotyping as a strategy for uncovering the biological correlates of psychopathology. *Nature Mental Health* **1**(5), 304–315 (2023)
29. Wei, W., Zhang, K., Chang, J., Zhang, S., Ma, L., Wang, H., et al.: Analyzing 20 years of Resting-State fMRI Research: Trends and collaborative networks revealed. *Brain Research* **1822**, 148634 (2024)
30. Wen, G., Cao, P., Liu, L., Yang, J., Zhang, X., Wang, F., et al.: Graph Self-Supervised Learning With Application to Brain Networks Analysis. *IEEE Journal of Biomedical and Health Informatics* **27**(8), 4154–4165 (2023)

31. Woo, C.W., Chang, L.J., Lindquist, M.A., Wager, T.D.: Building better biomarkers: brain models in translational neuroimaging. *Nature Neuroscience* **20**(3), 365–377 (2017)
32. Yang, Y., Ye, C., Su, G., Zhang, Z., Chang, Z., Chen, H., et al.: BrainMass: Advancing Brain Network Analysis for Diagnosis with Large-scale Self-Supervised Learning. *IEEE Transactions on Medical Imaging* **43**(11), 4004–4016 (2024)
33. Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., et al.: The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology* **106**(3), 1125–1165 (2011)