
Adapting HFMCA to Graph Data: Self-Supervised Learning for Generalizable fMRI Representations

Jakub Frać
Vrije Universiteit Amsterdam

Qiang Li
TReNDS Center

Guido Van Wingen
Amsterdam University Medical Center

Shujian Yu
Vrije Universiteit Amsterdam

Abstract

1 Functional magnetic resonance imaging (fMRI) analysis faces major challenges due
2 to limited data and variability across studies. Existing self-supervised methods from
3 computer vision often rely on positive-negative pairs, which are difficult to define
4 for neuroimaging data. We adapt the Hierarchical Functional Maximal Correlation
5 Algorithm (HFMCA) to graph-structured fMRI, providing a principled framework
6 that measures statistical dependence and enables robust self-supervised pretraining.
7 Across five neuroimaging datasets, our method yields competitive embeddings for
8 multiple classification tasks and transfers effectively to unseen domains. Code and
9 supplementary material: <https://github.com/fr30/mri-eigenencoder>

10 1 Introduction

11 Functional magnetic resonance imaging (fMRI) reveals brain dynamics, with resting-state functional
12 connectivity serving as a key biomarker for neurological and psychiatric disorders [15, 7, 13]. Deep
13 learning, however, struggles with limited data, heterogeneous preprocessing, and domain shifts.

14 Self-supervised learning (SSL) methods adapted from computer vision [17, 20, 3, 14] address these
15 issues. Graph-based approaches, modeling functional connectivity matrices, offer interpretable,
16 low-dimensional representations and preserve network topology through augmentations.

17 The **Hierarchical Functional Maximal Correlation Algorithm (HFMCA)** [11] extends this by measur-
18 ing statistical dependence across multiple feature hierarchies without contrastive pairs [4, 10, 21, 1].
19 Operating on graph-structured data, it captures richer dependencies and yields more generalizable
20 neuroimaging representations.

21 **Contributions:** (1) We adapt HFMCA to graph-structured fMRI data, representing the first appli-
22 cation and extension of this framework to brain connectivity networks. (2) We demonstrate that
23 HFMCA-pretrained encoders produce competitive embeddings for neuroimaging classification tasks
24 across diverse datasets. (3) We show effective transfer learning capabilities, particularly in scenarios
25 where limited labelled data is available. 4) We evaluate neural scaling laws in the context of fMRI
26 graph encoders, showing that naive pretraining data scaling may induce negative transfer.

27 2 Background and Methods

28 2.1 Functional Maximal Correlation Algorithm

29 The Functional Maximal Correlation Algorithm (FMCA) [11] measures statistical dependence
30 between two random variables X and Y by maximizing the correlation between their nonlinear




Figure 1: HFMCA learns graph representations by maximising dependence between low- and high-level features from multiple augmentations.

Model	BSNIP	AOMIC (Sex)	HCP (Sex)
Baseline _F	47.7 ± 0.9	56.9 ± 1.3	63.0 ± 2.4
VICReg _F	46.4 ± 0.9	60.0 ± 1.5	65.7 ± 1.5
BT _F	46.2 ± 1.2	60.9 ± 0.8	66.0 ± 1.5
SimCLR _F	<u>47.9 ± 0.5</u>	59.2 ± 1.5	66.6 ± 1.5
HFMCA _F	48.7 ± 0.8	59.4 ± 1.1	<u>66.0 ± 1.9</u>

Table 1: Frozen encoder accuracy (%) on unseen datasets.

Model	BSNIP	AOMIC (Sex)	HCP (Sex)
Baseline	47.8 ± 0.8	62.5 ± 1.8	71.1 ± 2.2
VICReg	47.1 ± 1.0	60.5 ± 1.7	66.0 ± 3.2
BT	47.2 ± 1.2	60.3 ± 1.4	66.1 ± 1.1
SimCLR	48.7 ± 0.6	<u>63.7 ± 1.8</u>	68.2 ± 3.1
HFMCA	<u>48.4 ± 0.8</u>	64.6 ± 1.4	<u>70.2 ± 0.6</u>

Table 2: Unfrozen encoder accuracy (%) on unseen datasets.

transformations. It decomposes their joint density $\rho(X, Y) = p(X, Y)/(p(X)p(Y))$ into orthogonal components, whose neural approximations $f_\theta(X)$ and $g_\omega(Y)$ learn maximally dependent features. The objective $\mathcal{L}_{FMCA} = \log \det R_{FG} - \log \det R_F - \log \det R_G$ encourages orthogonal features within each view (via R_F, R_G) while aligning them across views (via R_{FG}).

2.2 Problem Setup and Graph Construction

Given a small labelled clinical dataset $D_c = \{(X_i, Y_i)\}$ and a large unlabelled population dataset $D_p = \{X_j\}$, each subject is represented by a functional connectivity matrix $X \in \mathbb{R}^{|V| \times |V|}$ encoding pairwise correlations between brain regions. To mitigate overfitting on D_c , we pretrain an encoder f_θ on D_p with a self-supervised FMCA objective and fine-tune it for diagnosis prediction.

2.3 Hierarchical FMCA on Graphs

We adapt FMCA to hierarchical graph representations by contrasting local and global feature dependencies. For a graph X and its augmentations $Y_i = \mathcal{T}_i(X)$, a shared encoder f_θ^1 produces embeddings Z_i . These are projected into low-level features $Z_L = [f_\omega^2(Z_1), \dots, f_\omega^2(Z_T)]$ and a high-level aggregate $Z_H = \sum_i f_{\omega_i}^3(Z_i)$. Low- and high-level features Z_L and Z_H are used to construct correlation matrices R_L, R_H and R_{LH} . The objective $\mathcal{L}_{HFMCA} = \log \det R_{LH} - \log \det R_L - \log \det R_H$ maximises statistical dependence between Z_L and Z_H , enforcing multi-view consistency across graph augmentations such as random walk sampling, node dropping, and edge perturbation. After pretraining, projection heads are removed, and only the encoder is used for downstream tasks.

We employ a Graph Transformer based on the GPS architecture [16], combining local message passing with global attention. Random Walk Positional Encodings [8] preserve brain region topology, and graph-level embeddings are obtained via global mean pooling.




Figure 2: HFMCA pretrained with increasing data volumes. No monotonic scaling trend is observed, consistent with recent findings on negative transfer in graph models.

3 Experiments

3.1 Experimental Setup

We evaluate our approach on five fMRI datasets covering diverse diagnostic and demographic tasks: **REST** [5]: MDD and sex classification (1642 subjects; 51.6% MDD; 61.0% Male) **ABIDE** [6]: ASD classification (866 subjects; 53.6% ASD) **BSNIP** [19]: Schizophrenia/Bipolar (1464 subjects; 43.7/34.2/22.1% Healthy/SZ/BP) **AOMIC** [18]: Sex (881 subjects; 51.9% Male) **HCP** [9]: Sex (443 subjects; 55.5% Male). Models are pretrained on REST and ABIDE (2000+ subjects) and evaluated on all datasets, including unseen ones (BSNIP, AOMIC, HCP).

The encoder with two projection heads is trained for 200 epochs using Adam (10^{-3} lr, 10^{-6} weight decay, batch size 256). After pretraining, projection heads are discarded and the backbone is fine-tuned. We compare HFMCA against SimCLR [4], Barlow Twins [21], VICReg [1], and a randomly initialized baseline. Evaluation follows nested 5-fold cross-validation with 10 runs, using both frozen and unfrozen encoder variants.

3.2 Results

Linear classifiers trained on frozen and unfrozen encoders assess transferability. The results (Tables 1 and 2) show that HFMCA consistently outperforms random initialization (*Baseline*) and remains competitive with other methods. Notably, it achieves more stable performance on average, exhibiting lower variance across experimental runs.

To investigate scaling laws, we pretrained HFMCA on increasing dataset sizes: **REST**, **REST+ABIDE**, **REST+ABIDE+HCP**, and **REST+ABIDE+HCP+BSNIP**. As shown in Figure 2, performance peaks for REST+ABIDE and drops when adding more data, consistent with reports of negative transfer in large graph models [12, 2].

4 Conclusion

We successfully extended HFMCA to graph-structured fMRI data, providing a theoretically principled approach to self-supervised representation learning. Our method achieves competitive performance across five neuroimaging datasets. The demonstrated transfer learning capabilities and stable training make HFMCA particularly suitable for neuroimaging applications. Future work should explore larger-scale datasets and investigate HFMCA as a component of foundational models for brain imaging. Even though the initial scaling law analysis suggests greater complexity compared to text and vision domains, the framework and transferability indicate important contributions toward generalizable computational models of brain function.

83 5 Potential Negative Societal Impact

84 This work could enable unintended uses such as inferring cognitive or mental health traits without
85 consent, raising privacy and fairness concerns. Moreover, pretrained models may inherit demographic
86 or dataset biases, leading to inequitable outcomes if applied clinically.

87 References

- 88 [1] Adrien Bardes, Jean Ponce, and Yann Lecun. Vicreg: Variance-invariance-covariance regular-
89 ization for self-supervised learning. In *ICLR*, 2022.
- 90 [2] Yuxuan Cao et al. When to pre-train graph neural networks? from data generation perspective!
91 In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data*
92 *Mining*, pages 142–153, 2023.
- 93 [3] J. O. Caro et al. BrainLM: A foundation model for brain activity recordings. In *ICLR*, 2024.
- 94 [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework
95 for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020.
- 96 [5] X. Chen et al. The direct consortium and the rest-meta-mdd project: towards neuroimaging
97 biomarkers of major depressive disorder. *Psychoradiology*, 2(1):32–42, 2022. doi: 10.1093/
98 psyrad/kkac005.
- 99 [6] C. Craddock, Y. Benhajali, C. Chu, F. Chouinard, A. Evans, A. Jakab, B. Khundrakpam, J. D.
100 Lewis, Q. Li, M. Milham, C. Yan, and P. Bellec. The neuro bureau preprocessing initiative:
101 open sharing of preprocessed neuroimaging data and derivatives. *Neuroinformatics*, 2013.
- 102 [7] B. de Kwaasteniet, E. Ruhe, M. Caan, M. Rive, S. Olabbariaga, M. Groefsema, L. Heesink,
103 G. van Wingen, and D. Denys. Relation between structural and functional connectivity in
104 major depressive disorder. *Biological Psychiatry*, 74(1):40–47, 2013. ISSN 0006-3223. doi:
105 <https://doi.org/10.1016/j.biopsych.2012.12.024>. URL <https://www.sciencedirect.com/science/article/pii/S0006322313000401>.
- 106 [8] Vijay Prakash Dwivedi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson.
107 Graph neural networks with learnable structural and positional representations. In *International*
108 *Conference on Learning Representations*, 2022.
- 109 [9] D. C. Van Essen et al. The human connectome project: A data acquisition perspective.
110 *NeuroImage*, 62(4):2222–2231, 2012. ISSN 1053-8119. doi: [https://doi.org/10.1016/j.](https://doi.org/10.1016/j.neuroimage.2012.02.018)
111 [neuroimage.2012.02.018](https://doi.org/10.1016/j.neuroimage.2012.02.018). URL [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S1053811912001954)
112 [pii/S1053811912001954](https://www.sciencedirect.com/science/article/pii/S1053811912001954). Connectivity.
- 113 [10] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for
114 unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020.
- 115 [11] Bo Hu, Yuheng Bu, and José C Príncipe. Learning orthonormal features in self-supervised
116 learning using functional maximal correlation. In *2024 IEEE International Conference on*
117 *Image Processing (ICIP)*, pages 472–478, 2024.
- 118 [12] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure
119 Leskovec. Strategies for pre-training graph neural networks. In *ICLR*, 2020.
- 120 [13] M. Liang, Y. Zhou, T. Jiang, Z. Liu, L. Tian, H. Liu, and Y. Hao. Widespread functional
121 disconnectivity in schizophrenia with resting-state functional magnetic resonance imaging.
122 *Neuroreport*, 17(2):209–213, February 2006. doi: 10.1097/01.wnr.0000198434.06518.b8.
- 123 [14] I. Malkiel, G. Rosenman, L. Wolf, and T. Hendler. Self-supervised transformers for fmri
124 representation. In *Proceedings of the 5th International Conference on Medical Imaging with*
125 *Deep Learning*, pages 895–913, 2022.
- 126

- 127 [15] C. S. Monk, S. J. Peltier, J. L. Wiggins, S.-J. Weng, M. Carrasco, S. Risi, and C. Lord.
128 Abnormalities of intrinsic functional connectivity in autism spectrum disorders,. *NeuroImage*,
129 47(2):764–772, 2009. ISSN 1053-8119. doi: <https://doi.org/10.1016/j.neuroimage.2009.04.069>.
130 URL <https://www.sciencedirect.com/science/article/pii/S1053811909004327>.
- 131 [16] Ladislav Rampášek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and
132 Dominique Beaini. Recipe for a general, powerful, scalable graph transformer. In *NeurIPS*,
133 volume 35, pages 14501–14515, 2022.
- 134 [17] C. Shi, Y. Wang, Y. Wu, S. Chen, R. Hu, M. Zhang, B. Qiu, and X. Wang. Self-supervised
135 pretraining improves the performance of classification of task functional magnetic resonance
136 imaging. *Frontiers in Neuroscience*, 17:1199312, June 2023. doi: 10.3389/fnins.2023.1199312.
- 137 [18] L. Snoek, M. M. van der Miesen, T. Beemsterboer, A. Van Der Leij, A. Eigenhuis, and H. S.
138 Scholte. The amsterdam open mri collection, a set of multimodal mri datasets for individual
139 difference analyses. *Scientific Data*, 8(1):1–23, 2021. doi: 10.1038/s41597-021-00986-6.
- 140 [19] C. A. Tamminga, G. Pearlson, M. Keshavan, J. Sweeney, B. Clementz, and G. Thaker.
141 Bipolar and schizophrenia network for intermediate phenotypes: Outcomes across the psy-
142 chosis continuum. *Schizophrenia Bulletin*, 40:S131–S137, 02 2014. ISSN 0586-7614. doi:
143 10.1093/schbul/sbt179. URL <https://doi.org/10.1093/schbul/sbt179>.
- 144 [20] X. Wang, Y. Chu, Q. Wang, L. Cao, L. Qiao, L. Zhang, and M. Liu. Unsupervised contrastive
145 graph learning for resting-state functional mri analysis and brain disorder detection. *Human*
146 *Brain Mapping*, 44(17):5672–5692, Dec 2023. doi: 10.1002/hbm.26469. Epub 2023 Sep 5.
- 147 [21] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-
148 supervised learning via redundancy reduction. In *ICML*, pages 12310–12320, 2021.