Bridging EEG to fMRI: Semi-Supervised Functional Connectivity Translation via Latent Alignment

Xiaoye Li[1], Weiqi Li[1], Jia Mi[1], Shijie Huang[1], Yuanzhe He[1], Xin Lin[1], Zhiming Cui[1(✉)], and Dinggang Shen[1,2,3(✉)]

[1] School of Biomedical Engineering & State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai, China
[2] Shanghai United Imaging Intelligence Co., Ltd., Shanghai, China
[3] Shanghai Clinical Research and Trial Center, Shanghai, China
cuizm.neu.edu@gmail.com, Dinggang.Shen@gmail.com

Abstract. Functional magnetic resonance imaging (fMRI) and electroencephalography (EEG) are key tools for characterizing brain function, yet fMRI's high cost and limited portability hinder large-scale screening and longitudinal monitoring, motivating EEG-based inference of fMRI-derived representations. Existing EEG-to-fMRI methods predominantly perform signal-level regression and are sensitive to hemodynamic variability and EEG spatial ambiguity. We therefore reformulate EEG-to-fMRI translation as graph-to-graph mapping in functional connectivity (FC) space, directly targeting network organization with a time-aggregated representation. To exploit both scarce paired sessions and abundant unpaired data, we propose a two-stage semi-supervised framework: (i) within-modality FC reconstruction to learn robust representations for each modality and (ii) paired EEG-to-fMRI translation with asymmetric teacher–student latent alignment, which suppresses modality-dependent noise while preserving shared functional structure. We further use hierarchical fMRI FC, derived from connectivity-profile similarities, to capture higher-level network organization. Experiments on simultaneous EEG–fMRI datasets show consistent improvements over strong baselines in FC similarity and network-pattern agreement; the learned representations also improve downstream Alzheimer's disease classification.

Keywords: EEG-to-fMRI Translation · Functional Connectivity · Semi-supervised Learning · Latent Alignment

1 Introduction

Functional neuroimaging provides a critical window into the organization and dynamics of large-scale brain networks. Both functional magnetic resonance imaging (fMRI) and electroencephalography (EEG) capture neural activity, yet they differ markedly in cost, portability, and spatiotemporal resolution. fMRI measures blood-oxygen-level-dependent (BOLD) signals with high spatial resolution,

but its high cost and limited portability restrict large-scale screening and longitudinal monitoring. In contrast, EEG records electrical activity at millisecond resolution and is inexpensive and portable for diverse clinical and community settings [2]. Linked through neurovascular coupling [19,14], motivating to infer fMRI functional representations from widely available EEG to extend fMRI-level network assessment without the associated cost and logistical barriers.

Recent studies have explored EEG-to-fMRI translation using deep neural networks [14,12,13]. However, these methods formulate the task as signal-level regression, predicting fMRI time series at the regions of interest (ROI) or voxel level from EEG, which has two fundamental limitations. (1) **Signal-level** temporal correspondence is unstable because the hemodynamic response function (HRF) varies across brain regions and individuals, yet most methods assume a fixed HRF [12]. (2) **Signal-level** modeling is typically implemented in a localized, ROI-independent manner; e.g., NeuroBOLT [14] trains a separate predictor per ROI within a predefined temporal window, ignoring whole-brain spatial consistency and increasing computational cost for fine-grained parcellations. These limitations motivate a shift from **signal-level** to **network-level** characterization of brain function. Functional connectivity (FC) summarizes time-aggregated statistical dependencies between brain regions [6]. Because FC averages over extended temporal windows rather than relying on point-wise signal correspondence, it is less sensitive to HRF variability and instantaneous noise that destabilize signal-level methods [29]. Moreover, FC explicitly encodes inter-regional interactions and network structure, providing a more stable and biologically meaningful target space for cross-modal modeling. Additionally, simultaneous EEG-fMRI sessions are scarce, whereas large-scale EEG-only and fMRI-only datasets are readily available. Most existing translation methods rely solely on limited paired supervision and do not effectively exploit such unpaired data, restricting representation learning and generalization. This motivates a semi-supervised design that leverages abundant unpaired data for within-modality pretraining and reserves paired data for cross-modal alignment and translation.

In this paper, we reformulate EEG-to-fMRI translation as graph-to-graph mapping in FC space to focus on network-level functional structure, rather than enforcing signal-level correspondence. To address scarce simultaneous EEG-fMRI pairs, we develop a two-stage semi-supervised framework: Stage I leverages abundant unpaired EEG and fMRI data to learn modality-specific latent manifolds through self-reconstruction, and Stage II uses paired sessions for EEG-to-fMRI FC translation with a teacher–student scheme, where a frozen fMRI encoder guides the EEG encoder via asymmetric latent alignment. We follow standard EEG/fMRI preprocessing from [9] and construct hierarchical fMRI FC targets inspired by [28]; our contribution is integrating these components into a semi-supervised FC translation framework with unpaired representation learning and frozen fMRI-teacher-guided EEG alignment. Extensive experiments demonstrate superior preservation of FC structure and network patterns over signal-level and graph-based baselines, and the learned EEG representations further improve downstream Alzheimer's disease diagnosis.

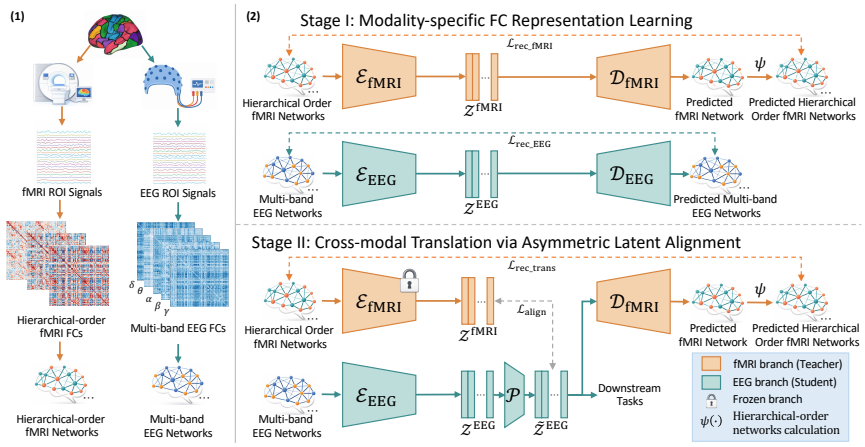

Fig. 1. Overview of our two-stage semi-supervised EEG-to-fMRI FC translation framework. (1) ROI time series are converted into multi-band EEG FCs and hierarchical-order fMRI FCs. (2) Stage I learns modality-specific representations via within-modality reconstruction, and Stage II aligns EEG embeddings to a frozen fMRI teacher to generate fMRI-like FCs.

2 Method

As illustrated in Fig. 1, our framework first constructs FC networks from EEG and fMRI recordings and then learns EEG-to-fMRI translation via two stages: (i) modality-specific FC representation learning using abundant unpaired data, and (ii) cross-modal translation via asymmetric latent alignment that maps EEG representations into the fMRI latent space using paired data. The learned EEG representations can be subsequently leveraged for downstream clinical tasks as a surrogate for fMRI.

2.1 Functional Connectivity Construction

We represent both EEG and fMRI at the network level using FCs, which summarize statistical dependencies between distributed brain regions and provide a time-aggregated description of large-scale functional organization (Fig. 1(1)). Throughout this work, we use the Schaefer-100 parcellation with $N = 100$ ROIs.
EEG FC. After preprocessing and source localization via a minimum-norm estimate (MNE) on the `fsaverage` template [9], signal-level EEG is averaged within each ROI of the Schaefer-100 atlas [20]. EEG FC is computed using power envelope connectivity (PEC), which reduces volume conduction artifacts and captures slow envelope fluctuations (< 1 Hz) related to fMRI hemodynamic dynamics [9]. Here $\delta/\theta/\alpha/\beta/\gamma$ denote carrier frequency bands, while the "$<$ 1 Hz" term refers to the slow envelope dynamics used by PEC. FCs are computed separately for five frequency bands: δ (0.5–4 Hz), θ (4–8 Hz), α (8–14 Hz), β (14–30 Hz), and γ (30–50 Hz), and stacked as $\mathbf{A}^{\mathrm{EEG}} \in \mathbb{R}^{5\times N\times N}$.

fMRI FC. Following preprocessing [9], BOLD signals are averaged within each ROI of the Schaefer-100 atlas [20] to obtain regional time series $\mathbf{T} \in \mathbb{R}^{N\times L}$. We compute first-order fMRI FC by pairwise Pearson correlation, yielding $\mathbf{C} = \mathbf{A}^{\text{fMRI}}_{\text{first}} \in \mathbb{R}^{N\times N}$, and apply Fisher's r-to-z transform for variance stabilization. Since first-order FC captures only pairwise interactions, we construct a hierarchical fMRI FC representation using a deterministic operator $\Psi(\cdot)$. Let $\mathbf{C}_{i,:}$ denote the ROI-wise connectivity profile of ROI i. We compute high-order FC as $\mathbf{H}_{ij} = \text{corr}(\mathbf{C}_{i,:}, \mathbf{C}_{j,:})$ and associated FC as $\mathbf{S}_{ij} = \text{corr}(\mathbf{C}_{i,:}, \mathbf{H}_{j,:})$. The resulting hierarchical representation is $\mathbf{A}^{\text{fMRI}}_{\text{hier}} = [\mathbf{C}, \mathbf{H}, \mathbf{S}] = \Psi(\mathbf{C})$. Here $\mathbf{H}$ captures similarity between whole-brain connection patterns, $\mathbf{S}$ captures cross-order similarity profiles, and Ψ has no learned parameters.

2.2 Two-stage Semi-supervised Learning Framework

Given limited paired EEG-fMRI data and abundant unpaired recordings, we adopt a two-stage semi-supervised design with within-modality reconstruction followed by cross-modal translation and alignment (Fig. 1(2)). For each modality $m \in \{\text{EEG}, \text{fMRI}\}$, the ROIs form a shared node set $\mathcal{V}$, and each ROI's row-wise connectivity profile serves as its node feature. Although FC matrices are symmetric, we retain the full $N \times N$ matrix to preserve these ROI-centered profiles; upper-triangular vectorization is used only for evaluation. With batch size B, the EEG and fMRI inputs have shapes $\mathbf{A}^{\text{EEG}} \in \mathbb{R}^{B\times 5\times N\times N}$ and $\mathbf{A}^{\text{fMRI}}_{\text{hier}} \in \mathbb{R}^{B\times 3\times N\times N}$, respectively.

Stage I: Modality-specific FC Representation Learning. Stage I learns a latent manifold for each modality from abundant unpaired data. An encoder–decoder pair $(\mathcal{E}_m, \mathcal{D}_m)$ reconstructs each modality's FC. A Transformer jointly models ROI-wise connectivity profiles across EEG frequency bands or fMRI hierarchical orders, followed by aggregation across bands or orders. The aggregated ROI features are then processed by Chebyshev graph convolutional (ChebGCN) encoders and decoders with adaptive learnable adjacency matrices. Each encoder outputs $\mathbf{Z}^m \in \mathbb{R}^{B\times N\times h}$. The EEG decoder reconstructs the five-band EEG FC, whereas the fMRI decoder reconstructs first-order FC and obtains its hierarchical target via $\Psi(\cdot)$. The Stage I objective is

$$\mathcal{L}_{\text{StageI}} = \|\hat{\mathbf{A}}^{\text{EEG}} - \mathbf{A}^{\text{EEG}}\|_F^2 + \|\hat{\mathbf{A}}^{\text{fMRI}}_{\text{hier}} - \mathbf{A}^{\text{fMRI}}_{\text{hier}}\|_F^2, \tag{1}$$

where $\|\cdot\|_F^2$ is the squared Frobenius norm over all dimensions.

Stage II: Cross-modal Translation via Asymmetric Latent Alignment. Stage I has learned robust modality-specific FC representations, while Stage II uses paired EEG-fMRI data for cross-modal alignment and EEG-to-fMRI FC translation.

Asymmetric Latent Alignment. We map EEG embeddings to the fMRI latent space with a teacher–student strategy, and decode them with the fMRI decoder to predict first-order fMRI FC. We treat the fMRI branch as the teacher since fMRI FC provides a more stable spatial organization, whereas EEG-derived FC is more affected by source localization ambiguity and measurement noise.

Thus, we freeze the pretrained fMRI encoder $\mathcal{E}^*_{\mathrm{fMRI}}$ as the teacher, while fine-tuning the EEG encoder $\mathcal{E}_{\mathrm{EEG}}$ and the fMRI decoder $\mathcal{D}_{\mathrm{fMRI}}$. Alignment is performed in latent space rather than raw FC space because the learned latent space compresses modality-specific variability and aligning graph-level summaries is more robust to per-edge noise. To bridge the two modalities, we introduce a learnable projector $\mathcal{P}$ to match embedding distributions and dimensionality, providing a lightweight adaptation layer without modifying the frozen teacher: $\tilde{\mathbf{Z}}^{\mathrm{EEG}} = \mathcal{P}(\mathbf{Z}^{\mathrm{EEG}})$. To obtain sample-level representations for alignment, we extract graph-level embeddings via mean pooling over nodes: $\mathbf{g}_b^m = \mathrm{Readout}(\mathbf{Z}_b^m) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{Z}_{b,i}^m$, and we define $\tilde{\mathbf{g}}_b^{\mathrm{EEG}} = \mathrm{Readout}(\tilde{\mathbf{Z}}_b^{\mathrm{EEG}})$ as the projected EEG graph embedding. Instance-level alignment uses cosine similarity,

$$\mathcal{L}_{\cos} = \frac{1}{B}\sum_{b=1}^{B}\left(1 - \cos\left(\mathrm{sg}(\mathbf{g}_b^{\mathrm{fMRI}}), \tilde{\mathbf{g}}_b^{\mathrm{EEG}}\right)\right), \tag{2}$$

where $\mathrm{sg}(\cdot)$ stops gradients through the teacher. We also minimize the Kullback–Leibler (KL) divergence between softened teacher and student distributions,

$$\mathcal{L}_{\mathrm{KL}} = \frac{1}{B}\sum_{b=1}^{B} D_{\mathrm{KL}}\left(\mathrm{softmax}\left(\mathrm{sg}(\mathbf{g}_b^{\mathrm{fMRI}})/\tau\right) \parallel \mathrm{softmax}\left(\tilde{\mathbf{g}}_b^{\mathrm{EEG}}/\tau\right)\right), \tag{3}$$

where τ is a temperature parameter. This KL term is an auxiliary teacher-student distillation regularizer: it softly encourages the EEG student embedding to follow the fMRI teacher's latent distribution, while FC reconstruction and cosine alignment remain the main constraints.

Cross-modal Translation. The projected EEG embeddings are decoded by the fMRI decoder to generate first-order fMRI FC, $\tilde{\mathbf{C}} = \mathcal{D}_{\mathrm{fMRI}}(\tilde{\mathbf{Z}}^{\mathrm{EEG}})$, and the translated hierarchical representation is obtained as $\tilde{\mathbf{A}}_{\mathrm{hier}}^{\mathrm{fMRI}} = \Psi(\tilde{\mathbf{C}})$. The Stage II objective is

$$\mathcal{L}_{\mathrm{StageII}} = \|\tilde{\mathbf{A}}_{\mathrm{hier}}^{\mathrm{fMRI}} - \mathbf{A}_{\mathrm{hier}}^{\mathrm{fMRI}}\|_F^2 + \lambda(\mathcal{L}_{\cos} + \beta\mathcal{L}_{\mathrm{KL}}), \tag{4}$$

where λ and β balance reconstruction and alignment terms.

3 Experiments and Results

3.1 Datasets

We use three data collections. **Stage I (unpaired).** TUEG [17] (10,874 subjects/16,986 EEG recordings) and UK Biobank [22] (40,000 rs-fMRI subjects). **Stage II (paired).** Simultaneous EEG-fMRI sessions are summarized in Table 1 and split at the subject level with a 7:1:2 train/val/test ratio. **Downstream AD EEG.** Seven datasets are aggregated (Table 1); only AD and NC subjects are retained. The final cohort includes 363 subjects/461 EEG recordings, evaluated by subject-level 5-fold cross-validation.

Table 1. Source-wise data summary. For Stage II, entries are subject/session counts; for AD EEG, entries are AD/NC recording counts after filtering.

Stage II paired EEG-fMRI			**Downstream AD EEG**		
Source	**Subj.**	**Sess.**	**Source**	**AD**	**NC**
NATVIEW [23]	22	392	ADFSU [25]	160	12
Oddball [26]	17	102	PearlNeuro [4]	0	78
NODDI [3]	15	15	ADFTD [16]	36	29
Rest & sleep [8]	33	255	ADSZ [1]	24	24
Eyes Open/Closed [15]	8	8	BrainLat_EEG [18]	20	28
Simult. EEG-fMRI [7]	20	59	AD-Auditory [11]	17	10
Total	**115**	**831**	APAVA [5]	12	11

3.2 Experimental Setup

State-of-the-Art Methods. We compare our method with signal-level and network-level methods to ensure a comprehensive evaluation. For signal-level, we adapt the representative method NeuroBOLT [14] to take EEG time series as input and predict ROI- or voxel-wise fMRI signals. The predicted fMRI signals are then converted into FC using the same pipeline as Sec. 2.1 to enable FC-level comparison. For network-level methods, representative graph-learning FC regression baselines with increasing modeling capacity are included: GCN [10], GAT [24], GFormer [27], and DGCNN [21]. All methods are trained and evaluated under identical subject-independent splits and preprocessing.

Evaluation metrics. We evaluate FC translation quality on **first-order** fMRI FC using: (i) **Mean Squared Error (MSE)** (↓) between predicted and ground-truth FC, measuring the element-wise accuracy, and (ii) **Pearson Correlation (Corr)** (↑), measuring pattern similarity.

Implementation details. All models are trained with AdamW (learning rate 1×10^{-4}) with batch size 64. For Stage II, we set the alignment weights $\lambda = 10.0$ and $\beta = 0.5$ in Eq.(4), and temperature $\tau = 0.2$ in Eq.(3), across all experiments, and set latent dimension $h = 128$ for all encoders.

3.3 Comparison with State-of-the-Art Methods

Quantitative Results. Table 2 summarizes EEG-to-fMRI FC translation under subject-independent evaluation. The results demonstrate that our method achieves the best overall performance, with an MSE of 0.0551 and a Corr of 0.6530. Notably, signal-level EEG-to-fMRI translation baseline NeuroBOLT performs substantially worse than network-level methods, indicating that direct signal-to-signal regression struggles to recover reliable inter-regional dependencies in the presence of hemodynamic delays and cross-modal temporal-scale mismatch. This limitation is particularly evident for NeuroBOLT, which learns ad hoc ROI-wise predictors independently for each region, ignoring global spatial context and thus failing to preserve coherent whole-brain network organization.

Table 2. Comparison to graph-level and signal-level baselines. Paired Wilcoxon signed-rank tests compare our method with each baseline: $^{\dagger}p < 0.05$ and $^{\ddagger}p < 0.001$.

Method	MSE ↓	Corr ↑
GCN [10]	0.1423 ± 0.0133	0.1659 ± 0.0444
GAT [24]	0.1428 ± 0.0067	0.2570 ± 0.0643
GFormer [27]	0.1206 ± 0.0065	0.3367 ± 0.0496
DGCNN [21]	0.1005 ± 0.0049	0.4495 ± 0.0480
NeuroBOLT [14]	0.2196 ± 0.0372	0.2176 ± 0.0595
Ours	$\mathbf{0.0551 \pm 0.0197}^{\dagger}$	$\mathbf{0.6530 \pm 0.1199}^{\ddagger}$

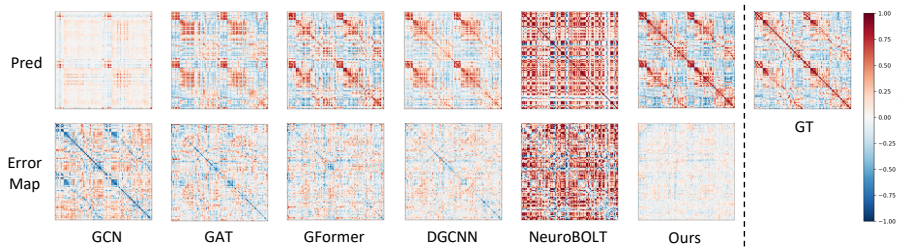

Fig. 2. Qualitative visualization of EEG-to-fMRI functional connectivity translation. Error maps visualize signed residuals $\Delta\mathbf{A} = \hat{\mathbf{A}} - \mathbf{A}$.

In contrast, network-level methods that explicitly model functional connectivity provide markedly more reliable results, demonstrating the advantage of performing cross-modal mapping directly in FC space. By integrating two-stage representation learning and hierarchical-order supervision, our method yields more accurate and structurally consistent fMRI-like network representations, and consistently outperforms existing network-level methods by a clear margin.

Qualitative Results. Fig. 2 shows representative first-order fMRI FC and predictions. Graph-based baselines vary in quality: GCN performs poorly, while GAT/GFormer/DGCNN recover reasonable but rough structure, with DGCNN being the strongest yet still leaving noticeable residual errors. The signal-level method largely fails to produce meaningful connectivity patterns, yielding noisy FCs without coherent network structure. In contrast, our method better matches the GT by preserving within-network blocks and key inter-network connections, with markedly reduced and less structured residuals.

3.4 Ablation Study

To evaluate each component, we conduct three ablations: (1) **w/o Stage I Pretraining** (train with paired data only), (2) **w/o Latent Alignment** (remove Stage II latent alignment), and (3) **w/o Hierarchical-order FC** (supervise first-order fMRI FC only). As shown in Table 3, removing Stage I causes the largest drop (MSE: $0.0551 \rightarrow 0.0648$; Corr: $0.6530 \rightarrow 0.6316$), confirming that

Table 3. Ablation study of our framework. We remove (i) Stage I within-modality reconstruction pretraining, (ii) the Stage II asymmetric latent alignment, and (iii) hierarchical-order FC supervision.

Method	MSE ↓	Corr ↑
w/o Stage I Pretraining	0.0648 ± 0.0237	0.6316 ± 0.1294
w/o Latent Alignment	0.0608 ± 0.0213	0.6496 ± 0.1256
w/o Hierarchical-order FC	0.0629 ± 0.0224	0.6427 ± 0.1259
Ours	$\mathbf{0.0551 \pm 0.0197}$	$\mathbf{0.6530 \pm 0.1199}$

Table 4. AD classification results with frozen EEG encoders and a strong GNN baseline (mean±std, % over 5 subject-independent folds). We report Balanced Accuracy (BA), Macro-F1 score, and Area Under the ROC Curve (AUC).

EEG Encoder Checkpoint	BA ↑	Macro-F1 ↑	AUC ↑
DGCNN [21]	76.2±4.4	79.0±6.9	88.9±4.4
Stage I only	77.4±2.4	76.1±4.4	88.6±1.7
Stage II only	78.1±5.3	79.3±7.8	88.9±3.7
Stage I + Stage II	**80.5±3.2**	**83.8±3.5**	**90.4±2.2**

within-modality reconstruction on abundant unpaired EEG and fMRI data is crucial for learning stable, modality-specific FC representations under scarce paired supervision. Removing hierarchical-order FC also degrades performance (MSE: $0.0551 \rightarrow 0.0629$; Corr: $0.6530 \rightarrow 0.6427$), indicating complementary structural and functional constraints beyond pairwise correlations that help preserve coherent topological organization. Disabling asymmetric alignment yields a moderate drop (MSE: $0.0551 \rightarrow 0.0608$; Corr: $0.6530 \rightarrow 0.6496$), suggesting Stage I provides a strong initialization while the latent-alignment terms further improve cross-modal consistency as a complementary regularizer.

3.5 Downstream Diagnosis for Alzheimer's Disease

To assess clinical relevance, we evaluate the learned EEG representations on Alzheimer's disease diagnosis. We freeze the EEG encoder and train a classifier (a two-layer MLP) on its embeddings to distinguish AD from NC. We report 5-fold results in Table 4, including a strong GNN baseline (DGCNN) and three training variants of our framework: **Stage I only**, **Stage II only**, and **Stage I+Stage II**. From Table 4, **Stage I only** already achieves competitive AD classification performance, indicating that within-modality reconstruction learns discriminative EEG representations. It slightly surpasses DGCNN in BA (77.4 ± 2.4 vs. 76.2 ± 4.4), while being lower in Macro-F1 and AUC. **Stage II only** improves Macro-F1 (79.3 ± 7.8) and matches DGCNN in AUC (88.9 ± 3.7 vs. 88.9 ± 4.4), suggesting that paired translation provides additional disease-relevant cues. Combining **Stage I+Stage II** achieves the best overall performance (BA/Macro-F1/AUC $= 80.5 \pm 3.2/83.8 \pm 3.5/90.4 \pm 2.2$), with the largest

gain in Macro-F1 (76.1→83.8), indicating that fMRI-guided knowledge transfer yields more clinically informative EEG representations.

4 Conclusion

We propose a two-stage semi-supervised EEG-to-fMRI translation framework in FC space, using unpaired data for within-modality representation learning and paired data for fMRI-guided asymmetric latent alignment. Hierarchical-order FC further constrains higher-level network patterns. Experiments show consistent improvements over strong baselines in FC similarity and network-pattern agreement, while downstream AD classification supports the utility of the learned EEG representations. Future work should evaluate leave-one-source/paradigm-out generalization and clinically controlled cohorts.

Acknowledgments. This work was supported in part by Brain Science and Brain-like Intelligence Technology-National Science and Technology Major Project (No. 2025ZD0215400), National Natural Science Foundation of China (grant numbers 62131015, U23A20295, 82441023, 82394432), the China Ministry of Science and Technology (STI2030-Major Projects-2022ZD0209000, STI2030-Major Projects-2022ZD0213100), Key-Area Research and Development Program of Guangdong Province (2023B0303040001), and HPC Platform of ShanghaiTech University.

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

References

1. Alves, C.L., Pineda, A.M., Roster, K., Thielemann, C., Rodrigues, F.A.: EEG functional connectivity and deep learning for automatic diagnosis of brain disorders: Alzheimer's disease and schizophrenia. Journal of Physics: Complexity **3**(2), 025001 (2022)
2. Cohen, M.X.: Where does EEG come from and what does it mean? Trends in Neurosciences **40**(4), 208–218 (2017)
3. Deligianni, F., Carmichael, D.W., Zhang, G.H., Clark, C.A., Clayden, J.D.: NODDI and tensor-based microstructural indices as predictors of functional connectivity. PLOS ONE **11**(4), e0153404 (2016)
4. Dzianok, P., Kublik, E.: PEARL-Neuro database: EEG, fMRI, health and lifestyle data of middle-aged people at risk of dementia. Scientific Data **11**(1), 276 (2024)
5. Escudero, J., Abásolo, D., Hornero, R., Espino, P., López, M.: Analysis of electroencephalograms in Alzheimer's disease patients with multiscale entropy. Physiological Measurement **27**(11), 1091–1106 (2006)
6. Friston, K.J.: Functional and effective connectivity: a review. Brain Connectivity **1**(1), 13–36 (2011)
7. Gallego-Rudolf, J., Corsi-Cabrera, M., Concha, L., Ricardo-Garcell, J., Pasaye-Alcaraz, E.: Simultaneous and independent electroencephalography and magnetic resonance imaging: A multimodal neuroimaging dataset. Data in Brief **51**, 109661 (2023)

8. Gu, Y., Sainburg, L.E., Han, F., Liu, X.: Simultaneous EEG and functional MRI data during rest and sleep from humans. Data in Brief **48**, 109059 (2023)
9. Jiao, Y., Zhao, K., Wei, X., et al.: Deep graph learning of multimodal brain networks defines treatment-predictive signatures in major depression. Molecular Psychiatry **30**, 3963–3974 (2025)
10. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: International Conference on Learning Representations (ICLR) (2017)
11. Lahijanian, M., Aghajan, H., Vahabi, Z.: Auditory gamma-band entrainment enhances default mode network connectivity in dementia patients. Scientific Reports **14**(1), 13153 (2024)
12. Lanzino, R., Fontana, F., Cinque, L., Scarcello, F., Maki, A.: NT-ViT: Neural transcoding vision transformers for EEG-to-fMRI synthesis. arXiv preprint arXiv:2409.11836 (2024)
13. Li, Y., Lou, A., Xu, Z., Wang, S., Chang, C.: Leveraging sinusoidal representation networks to predict fMRI signals from EEG. In: Medical Imaging 2024: Image Processing. Proceedings of SPIE, vol. 12926, p. 254. SPIE (2024)
14. Li, Y., Lou, A., Xu, Z., Zhang, S., Wang, S., Englot, D.J., Kolouri, S., Moyer, D., Bayrak, R.G., Chang, C.: NeuroBOLT: Resting-state EEG-to-fMRI synthesis with multi-dimensional feature mapping. In: The Thirty-eighth Annual Conference on Neural Information Processing Systems (2024)
15. van der Meer, J.N., Pampel, A., van Someren, E.J., Ramautar, J.R., van der Werf, Y.D., Gomez-Herrero, G., Lepsien, J., Hellrung, L., Hinrichs, H., Möller, H.E., Walter, M.: "Eyes Open – Eyes Closed" EEG/fMRI data set including dedicated "Carbon Wire Loop" motion detection channels. Data in Brief **7**, 990–994 (2016)
16. Miltiadous, A., Tzimourta, K.D., Afrantou, T., Ioannidis, P., Grigoriadis, N., Tsalikakis, D.G., Angelidis, P., Tsipouras, M.G., Glavas, E., Giannakeas, N., et al.: A dataset of scalp EEG recordings of Alzheimer's disease, frontotemporal dementia and healthy subjects from routine EEG. Data **8**(6), 95 (2023)
17. Obeid, I., Picone, J.: The Temple University Hospital EEG data corpus. Frontiers in Neuroscience **10**, 196 (2016)
18. Prado, P., Medel, V., Gonzalez-Gomez, R., Sainz-Ballesteros, A., Vidal, V., Santamaría-García, H., Moguilner, S., Mejia, J., Slachevsky, A., Behrens, M.I., et al.: The BrainLat project, a multimodal neuroimaging dataset of neurodegeneration from underrepresented backgrounds. Scientific Data **10**(1), 889 (2023)
19. Ritter, P., Villringer, A.: Simultaneous EEG–fMRI. Neuroscience & Biobehavioral Reviews **30**(6), 823–838 (2006)
20. Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.N., Holmes, A.J., et al.: Local-Global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. Cerebral Cortex **28**, 3095–3114 (2018)
21. Song, T., Zheng, W., Song, P., Cui, Z.: EEG emotion recognition using dynamical graph convolutional neural networks. IEEE Transactions on Affective Computing **11**(3), 532–541 (2020)
22. Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman, A., Young, A., Sprosen, T., Peakman, T., Collins, R.: UK Biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. PLOS Medicine **12**(3), e1001779 (2015)
23. Telesford, Q.K., Gonzalez-Moreira, E., Xu, T., Tian, Y., Colcombe, S.J., Cloud, J., Russ, B.E., Falchier, A., Nentwich, M., Madsen, J., Parra, L.C., Schroeder, C.E., Milham, M.P., Franco, A.R.: An open-access dataset of naturalistic viewing using simultaneous EEG-fMRI. Scientific Data **10**(1), 554 (2023)

24. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: International Conference on Learning Representations (ICLR) (2018)
25. Vicchietti, M.L., Ramos, F.M., Betting, L.E., Campanharo, A.S.: Computational methods of EEG signals analysis for Alzheimer's disease classification. Scientific Reports **13**(1), 8184 (2023)
26. Walz, J.M., Goldman, R.I., Muraskin, J., Brown, T.R., Sajda, P.: Simultaneous EEG-fMRI reveals temporal evolution of coupling between supramodal cortical attention networks and the brainstem. The Journal of Neuroscience **33**(49), 19212–19222 (2013)
27. Yun, S., Jeong, M., Kim, R., Kang, J., Kim, H.J.: Graph transformer networks. In: Advances in Neural Information Processing Systems. vol. 32 (2019)
28. Zhang, Y., Zhang, H., Chen, X., et al.: Hybrid high-order functional connectivity networks using resting-state functional MRI for mild cognitive impairment diagnosis. Scientific Reports **7**, 6530 (2017)
29. Zuo, X.N., Di Martino, A., Kelly, C., Shehzad, Z.E., Gee, D.G., Klein, D.F., Castellanos, F.X., Biswal, B.B., Milham, M.P.: The oscillating brain: complex and reliable. NeuroImage **49**(2), 1432–1445 (2010)