Dynamic Spatial-Temporal Mapping for EEG-to-fMRI Synthesis via Gaussian-Guided Heterogeneous Graph Neural Network

Jialun Zheng[1 [0000−0002−4957−596X]], Jie Liu[2 [0000−0002−1327−1315]], Jiannong Cao[1 [0000−0002−2725−2529]], and Divya Saxena[3 [0000−0002−6847−585X]]

[1] Department of Computing, The Hong Kong Polytechnic University, Hong Kong, HKSAR, China
[2] Department of Computer Science, City University of Hong Kong, Hong Kong, HKSAR, China
[3] School of Artificial Intelligence and Data Science, Indian Institute of Technology (IIT) Jodhpur, Rajasthan, India
{22069255r}@connect.polyu.hk

Abstract. EEG-to-fMRI synthesis aims to generate high spatial-temporal resolution fMRI, pivotal for diagnosing brain diseases, using accessible EEG signals. However, current static regression methods fail to account for regional variations in EEG-fMRI time delays and cannot disentangle spatially mixed EEG signals. To address this, we propose DSTP, a framework that explicitly parameterizes this dynamic spatial-temporal mapping. We introduce a Dynamic Heterogeneous Graph Neural Network (DHGNN) to model the evolving mapping between scalp electrodes and brain regions of interest (ROIs). Moreover, we embed a learnable Gaussian Mask into the message-passing mechanism. This acts as an adaptive temporal gate, effectively filtering the EEG history to identify the specific time window that aligns with each ROI's hemodynamic response. Additionally, a Transformer module integrates the historical evolution of these states, while a multi-objective loss preserves temporal dynamics and brain topological consistency alongside reconstruction accuracy [4]. Extensive experiments on multiple real-world datasets demonstrate that DSTP achieves state-of-the-art performance. Notably, the model exhibits superior zero-shot generalization across tasks and spatial resolutions, proving its ability to learn intrinsic physiological mapping rules rather than overfitting to specific distributions.

Keywords: Graph Neural Network · EEG-to-fMRI synthesis · Cross-modal · Brain Network.

1 Introduction

EEG-to-fMRI synthesis has emerged as a promising technique to generate high quality nueroimages from accessible electroencephalography (EEG) [17, 2]. This technology is essential for enabling affordable, high spatial-temporal resolution

brain activity monitoring. However, accurate synthesis is constrained by the complex, non-linear relationship, known as neurovascular coupling between EEG and fMRI. This ambiguity makes it difficult to achieve synthesis that is both spatially precise and temporally consistent [5], limiting the clinical reliability of current applications.

Existing approaches typically attempt to bridge this modality gap by learning robust EEG representations [15, 14, 18] and projecting them onto fMRI space via regression. While early methods assumed a fixed Hemodynamic Response Function (HRF), recent works have leveraged Generative Diffusion models [17] and Transformer-based architectures [11, 9] to capture varying temporal dependencies. However, they learn an implicit end-to-end mapping without explicitly parameterizing the underlying physiological process. Consequently, they face two major limitations. ***1) Temporal Oversimplification:*** They fail to explicitly model the region-specific variance of the HRF. In reality, hemodynamic delays are not uniform but vary significantly across different brain regions and subjects. ***2) Spatial Ambiguity:*** They struggle to resolve the inverse problem caused by volume conduction [6], where electrical signals measured at the scalp are spatially smeared and deflected from their true cortical sources. Thus, the cross-modal correspondence is not a static regression task but a Dynamic Spatial-Temporal Mapping that evolves with the underlying neural states.

To address these limitations, we propose **DSTP**, a Dynamic Spatial-Temporal Mapping framework designed to explicitly model the evolving physiology between EEG and fMRI. To resolve spatial ambiguity, we design a Dynamic Heterogeneous Graph Neural Network (DHGNN), inspired by prior graph learning approaches for brain network modeling [13]. Unlike static methods, the heterogeneous graph contains scalp electrodes and Regions of Interest (ROIs) as distinct type of nodes connected via evolving directional edges. To move beyond implicit temporal assumptions, we introduce a Gaussian Mask-Guided Message Passing mechanism. This module utilizes a learnable Gaussian distribution to generate adaptive temporal masks. These masks act as explicit physiological filters during message passing, effectively simulating the ROI-specific HRF and capturing time-varying delays. Finally, to ensure synthesis quality, we integrate a multi-objective loss function enforcing reconstruction accuracy, as well as temporal and topological consistency. Our contributions are listed as follows:

1. We propose a Dynamic Heterogeneous Graph Neural Network to explicitly model the evolving mapping between scalp electrodes and brain regions. This dynamic graph-based formulation effectively mitigates spatial ambiguity caused by volume conduction.
2. We introduce a Learnable Gaussian Mask mechanism that acts as an adaptive temporal gate. It captures region-specific hemodynamic delays and dispersions, aligning EEG electrical dynamics with fMRI BOLD signals in a biologically plausible manner.
3. Extensive experiments demonstrate that DSTP achieves state-of-the-art performance on both resting-state and task-based datasets, with superior zero-shot generalization capabilities across subjects, tasks, and spatial resolutions.

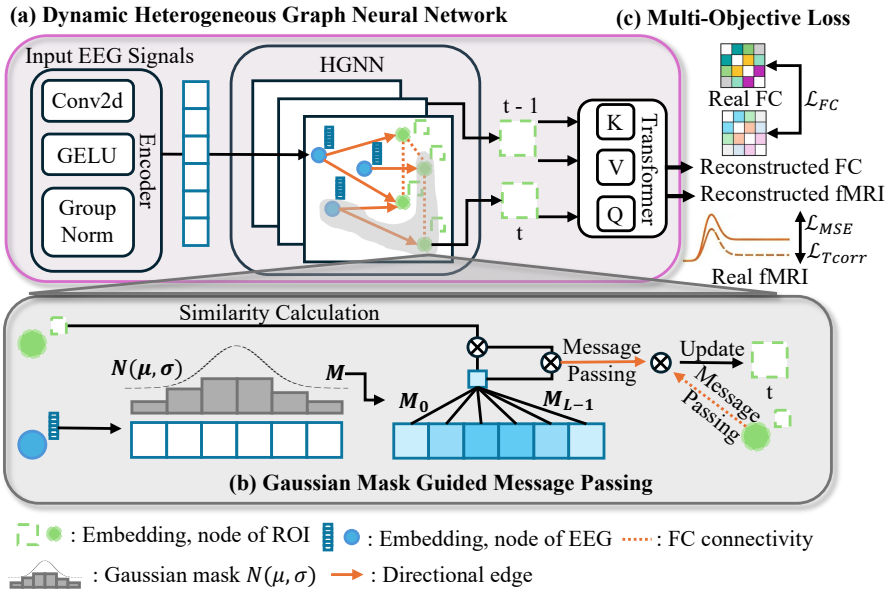

Fig. 1. Overview of the proposed DSTP framework. (a) DHGNN models evolving spatial-temporal relationships between scalp electrodes and brain ROIs. (b) Then, a Gaussian mask act as an adaptive temporal gate to filter region-specific hemodynamic delays. (c) Multi-Objective Loss optimizes reconstruction accuracy, temporal consistency, and global functional connectivity.

2 Method

2.1 Dynamic Heterogeneous Graph Construction

The mapping from scalp EEG electrodes to ROIs involves a complex, non-linear inverse problem. To model this, we construct a Dynamic Heterogeneous Graph. Unlike methods relying on static network structure [2], our graph structure dynamically evolves [20, 19], allowing the model to capture time-varying dependencies between the electrode and ROIs. We construct a heterogeneous graph $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$ at timestamp t with N_E Electrode Nodes and N_R ROI Nodes, connected by directed Electrode-to-ROI edges $\mathcal{E}_{ER}$ and undirected Functional Connectivity edges $\mathcal{E}_{FC}$. Here, E_{ER} denotes the directed cross-modal edges from EEG electrode nodes to ROI nodes, which model the information flow from scalp electrical activity to regional BOLD responses. In contrast, E_{FC} denotes the undirected ROI–ROI functional connectivity edges, which encode pairwise dependencies among brain regions and serve as a topological prior during message passing. The j-th electrode node is encoded from input EEG signals as $\mathbf{H}_{E,j}^{(t)} \in \mathbb{R}^{L \times d}$, where L is the window size, d is the feature dimension. The r-th ROI node is initialized with a learnable embedding $\mathbf{h}_{R,r}^{(0)} \in \mathbb{R}^d$.

2.2 Dynamic Spatial-Temporal Mapping

Since hemodynamic delays vary significantly across brain regions and subjects, traditional fixed-HRF models fail to capture this variability, while purely data-driven approaches risk overfitting to spurious correlations. To address this, we propose a three-stage strategy. Start with a Learnable Gaussian Mask explicitly models region-specific temporal delays. Followed by a Guided Message Passing mechanism resolves spatial ambiguities from volume conduction. Finally, a Temporal Evolution module integrates historical states to reconstruct the accumulative BOLD response.

Learnable Gaussian Mask As HRF follows a distinct unimodal shape, rising to a peak before dispersing08, we adopt a Gaussian function to approximate this active response window. For the r-th brain region (ROI), we generate a temporal weight gaussian mask $\mathbf{M}_r \in \mathbb{R}^L$. The mask value at the i-th time step for $i \in \{0, \ldots, L-1\}$ is defined as:

$$\mathbf{M}_r[i] = \exp\left(-\frac{(i-\mu_r)^2}{2\sigma_r^2}\right), \tag{1}$$

where μ_r and σ_r represent the learnable delay and dispersion parameters for region r. They are initialized to align with canonical physiological knowledge (e.g., $\mu \approx 6$s) but are optimized end-to-end.

Gaussian Mask Guided Message Passing To effectively integrate gaussian mask from previous section for modeling of ROI specific HRF, we propose a Gaussian Mask Guided Message Passing mechanism. This module uses the region-specific hemodynamic delay not as a hard filter, but as a soft guidance. For each electrode-ROI pair (j, r), we extract a delay-aware feature vector. We apply the Gaussian mask $\mathbf{M}_r$ to the electrode's history $\mathbf{H}_{E,j}$ to highlight the temporal window most relevant to ROI r:

$$\mathbf{u}_{j\to r} = \sum_{i=0}^{L-1} \mathbf{M}_r[i] \cdot \mathbf{H}_{E,j}[i,:], \tag{2}$$

where $\mathbf{u}_{j\to r} \in \mathbb{R}^d$ represents the electrode's signal conditioned on the hemodynamic properties of ROI r. Then, we determine the importance of electrode j to ROI r by calculating the similarity between the Gaussian-weighted feature $\mathbf{u}_{j\to r}$ and the previous ROI state $\mathbf{h}_r^{(t-1)}$. This similarity serves as the edge weight for message passing:

$$e_{jr} = \text{score}(\mathbf{h}_r^{(t-1)}, \mathbf{u}_{j\to r}) = (\mathbf{h}_r^{(t-1)}\mathbf{W}_Q)(\mathbf{u}_{j\to r}\mathbf{W}_K)^\top, \tag{3}$$

where $\mathbf{W}_Q$ and $\mathbf{W}_K$ are learnable weights. We normalize these scores across all electrodes to obtain the attention weights α_{jr}:

$$\alpha_{jr} = \frac{\exp(e_{jr})}{\sum_{u\in\mathcal{V}_E}\exp(e_{ur})}. \tag{4}$$

Finally, the ROI node aggregates the information based on these guided attention weights. We also incorporate functional connectivity from neighboring regions:

$$\tilde{\mathbf{h}}_r^{(t)} = \mathbf{h}_r^{(t-1)} + \underbrace{\sum_{j \in \mathcal{V}_E} \alpha_{jr} \mathbf{u}_{j \to r}}_{\text{Gaussian Mask Guided}} + \underbrace{\sum_{k \in \mathcal{N}_r} \tilde{A}_{rk}^{fc} \mathbf{W}_{fc} \mathbf{h}_k^{(t-1)}}_{\text{FC}}, \tag{5}$$

where $\tilde{A}_{rk}^{fc}$ is the normalized functional connectivity weight between ROIs r and k, computed from the training set, and fixed during training/inference. We use this fixed FC term as a stable global prior to preserve biologically meaningful inter-regional topology, while the time-varying electrode-to-ROI interactions are modeled dynamically through the Gaussian-guided attention mechanism. $\mathbf{W}_{fc}$ is a learnable linear transformation. The aggregated embedding is then updated via activation function:

$$\mathbf{h}_r^{(t)} = \sigma\left(\text{LayerNorm}(\tilde{\mathbf{h}}_r^{(t)})\right). \tag{6}$$

Temporal Evolution and Readout The DHGNN produces a spatial embedding $\mathbf{h}_r^{(t)}$ at each timestep, but fMRI signal generation is an accumulative process. We therefore apply a Transformer-based Temporal Module that attends over the historical states $[\mathbf{h}_r^{(t-T)}, \ldots, \mathbf{h}_r^{(t-1)}]$ to retrieve the past context most relevant to the current hemodynamic response. Specifically, the current state $\mathbf{h}_r^{(t)}$ serves as the Query, while the historical states are Keys $\mathbf{K}_r$ and Values $\mathbf{V}_r$:

$$\mathbf{z}_r^{(t)} = \text{Softmax}\left(\frac{\mathbf{h}_r^{(t)} \mathbf{W}_Q \left(\mathbf{K}_r\right)^\top}{\sqrt{d}}\right) \mathbf{V}_r. \tag{7}$$

The predicted fMRI value is then produced by a readout MLP that fuses the current spatial state with the retrieved temporal context:

$$\hat{y}_r^{(t)} = \text{MLP}_{out}\left(\mathbf{z}_r^{(t)} + \mathbf{h}_r^{(t)}\right). \tag{8}$$

2.3 Multi-Objective Loss

To ensure signal fidelity and biologically plausible connectivity, we optimize a composite loss combining three terms. The MSE loss $\mathcal{L}_{MSE}$ penalizes pointwise errors. Since MSE alone cannot capture temporal dynamics, we add a correlation loss that maximizes the Pearson correlation between predicted and ground-truth time series:

$$\mathcal{L}_{Tcorr} = 1 - \frac{1}{R} \sum_{r=1}^{R} \rho(\hat{\mathbf{y}}_{r,:}, \mathbf{y}_{r,:}). \tag{9}$$

To further preserve global brain topology, we impose a connectivity loss $\mathcal{L}_{FC}$ that penalizes the mean squared difference between the predicted and ground-truth functional connectivity matrices. The final objective is:

$$\mathcal{L}_{total} = \lambda_\alpha \mathcal{L}_{MSE} + \lambda_\beta \mathcal{L}_{Tcorr} + \lambda_\gamma \mathcal{L}_{FC}. \tag{10}$$

Table 1. Experimental Results of the Comparison Methods.

Setting	Method	Tcorr ↑	Pixcorr ↑	ConnMSE ↓
Intra-subject	BIOT [16]	$0.339_{\pm 0.169}$	$0.369_{\pm 0.206}$	$0.224_{\pm 0.134}$
	LaBraM [7]	$0.184_{\pm 0.164}$	$0.248_{\pm 0.226}$	$0.381_{\pm 0.229}$
	Li et al. [10]	$0.391_{\pm 0.195}$	$0.451_{\pm 0.224}$	$0.154_{\pm 0.101}$
	NeuroBOLT [11]	$0.369_{\pm 0.164}$	$0.402_{\pm 0.213}$	$0.259_{\pm 0.182}$
	UnEBOLT [9]	$0.416_{\pm 0.166}$	$0.490_{\pm 0.176}$	$0.201_{\pm 0.148}$
	DSTP (Ours)	$\mathbf{0.458_{\pm 0.173}}$	$\mathbf{0.537_{\pm 0.186}}$	$\mathbf{0.127_{\pm 0.153}}$
Inter-subject	BIOT [16]	$0.413_{\pm 0.035}$	$0.486_{\pm 0.135}$	$0.098_{\pm 0.055}$
	LaBraM [7]	$0.214_{\pm 0.058}$	$0.393_{\pm 0.089}$	$0.122_{\pm 0.068}$
	Li et al. [10]	$0.370_{\pm 0.038}$	$0.483_{\pm 0.109}$	$0.080_{\pm 0.047}$
	NeuroBOLT [11]	$0.413_{\pm 0.052}$	$0.428_{\pm 0.145}$	$0.126_{\pm 0.080}$
	UnEBOLT [9]	$0.418_{\pm 0.036}$	$0.549_{\pm 0.120}$	$0.092_{\pm 0.066}$
	DSTP (Ours)	$\mathbf{0.470_{\pm 0.055}}$	$\mathbf{0.559_{\pm 0.090}}$	$\mathbf{0.077_{\pm 0.073}}$

3 Experiments

3.1 Experimental Settings

Datasets We conduct extensive experiments on multiple EEG-fMRI datasets provided by [11] that are acquired from different sites, scanners, covering both resting-state and task-evoked paradigms. 1) Resting-state dataset consists of 29 simultaneous EEG-fMRI scans collected from 22 healthy volunteers (7 subjects with re-test scans). Participants were instructed to rest passively with their eyes closed for a duration of 20 minutes. We conduct main experiments, intra-subject and cross-subject evaluation on this dataset. 2) Auditory Task Dataset contains 16 scans from 10 subjects performing an auditory oddball task. Subjects kept their eyes closed and responded to random binaural tones via button presses.

Baselines We compare DSTP against multiple state-of-the-art methods, categorized into general EEG representation learners and EEG-to-fMRI translators.

Implementation Details DSTP is implemented in PyTorch using AdamW (20 epochs, batch 16, LR 5×10^{-4} with cosine annealing). Gaussian masks are initialized at $\mu = 10s$ and $\sigma = 8s$, with loss weights set to $\lambda_\alpha = 2.5, \lambda_\beta = 0.2, \lambda_\gamma = 0.5$. Following [9], we adopt the same data preprocessing protocol and experimental setup, including the atlas definition for ROI extraction and the general acquisition setting used for synchronized EEG-fMRI modeling. We evaluate performance via: (1) TCorr: ROI-wise temporal Pearson correlation; (2) PixCorr: Pearson correlation between flattened upper-triangular functional connectivity (FC) elements; and (3) ConnMSE: MSE between FC matrices. All results are reported as mean ± standard deviation across evaluation folds.

Table 2. Ablation Study Results

Setting	Method	Tcorr ↑	Pixcorr ↑	ConnMSE ↓
Intra-subject	DSTP w/o Gaussian	$0.412_{\pm 0.182}$	$0.528_{\pm 0.120}$	$\mathbf{0.095}_{\pm 0.158}$
	DSTP w/o DHGNN	$0.385_{\pm 0.195}$	$0.482_{\pm 0.135}$	$0.164_{\pm 0.188}$
	DSTP w/o Multi-loss	$0.425_{\pm 0.118}$	$0.530_{\pm 0.110}$	$0.142_{\pm 0.165}$
	DSTP (Ours)	$\mathbf{0.458}_{\pm \mathbf{0.173}}$	$\mathbf{0.537}_{\pm \mathbf{0.186}}$	$0.127_{\pm \mathbf{0.153}}$
Inter-subject	DSTP w/o Gaussian	$0.428_{\pm 0.090}$	$0.545_{\pm 0.095}$	$0.085_{\pm 0.038}$
	DSTP w/o DHGNN	$0.392_{\pm 0.110}$	$0.498_{\pm 0.115}$	$0.158_{\pm 0.072}$
	DSTP w/o Multi-loss	$0.441_{\pm 0.095}$	$0.550_{\pm 0.092}$	$0.135_{\pm 0.045}$
	DSTP (Ours)	$\mathbf{0.470}_{\pm \mathbf{0.055}}$	$\mathbf{0.559}_{\pm \mathbf{0.090}}$	$\mathbf{0.077}_{\pm \mathbf{0.073}}$

3.2 Experimental Results Comparison and Analysis

Table 1 presents the quantitative comparison of our proposed DSTP against state-of-the-art methods. Here, we conduct experiments under two settings: *intra-subject*, where training and testing data are derived from temporally split segments of the same scan, and *inter-subject*, where the model is trained on a group of subjects and evaluated on strictly unseen subjects. As shown, DSTP achieves overall best performance across all metrics in both settings. Specifically, we observe: DSTP significantly outperforms the strong baseline UnEBOLT in terms of ConnMSE. While UnEBOLT proves effective in pixel-level regression, it exhibits limitations in preserving global functional connectivity. DSTP overcomes this by explicitly modeling the complex brain topology via the DHGNN module, which captures not just static constraints but the dynamic evolution of topological structures.

3.3 Ablation Study

We evaluate the contribution of each component under both intra-subject and inter-subject settings in Table 2. (1) w/o Gaussian: Removing Gaussian-guided message passing leads to a significant drop in Tcorr, confirming its necessity for modeling regional HRF temporal variance. (2) w/o DHGNN: Replacing the dynamic graph with a standard regression backbone causes the sharpest performance decline across all metrics, particularly ConnMSE. This underscores that DHGNN is vital for capturing the brain's non-Euclidean topology and its dynamic evolution. (3) w/o Multi-loss: Training solely with MSE loss degrades ConnMSE despite stable Pixcorr, proving that multi-objective optimization is essential for preserving biological functional connectivity beyond simple pointwise accuracy. Notably, in the intra-subject setting, the full model does not achieve the lowest ConnMSE, indicating a trade-off between connectivity preservation and other reconstruction objectives.

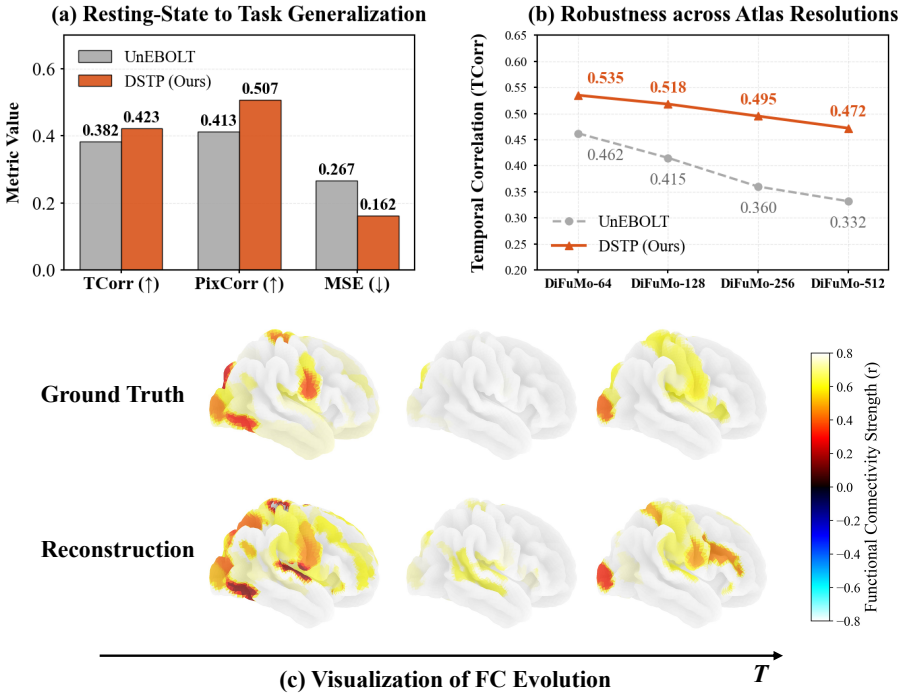

Fig. 2. Generalization and Visualization Analysis

3.4 Visualization and Generalization Study

As shown in Fig. 2a and Fig. 2b, DSTP exhibits superior zero-shot generalization compared to UnEBOLT, achieving higher accuracy when transferring from resting-state pretraining to unseen auditory tasks [8]. Notably, DSTP maintains stable performance across varying spatial resolutions, from DiFuMo 64 to 512 [3, 1], whereas UnEBOLT's performance degrades at higher granularities. We attribute this robustness to the DHGNN's ability to learn intrinsic, resolution-agnostic neurovascular coupling rules. By treating ROIs as graph nodes rather than fixed output dimensions, DSTP captures a continuous topological mapping from scalp potentials to cortical sources that remains effective regardless of parcellation granularity. Moreover, Fig. 2c visualizes the evolution of brain activity patterns. The reconstructed maps demonstrate high temporal fidelity with the ground truth, accurately tracking the propagation of neural activation. The synchronized shifting of hotspots across the time axis T confirms that our framework effectively captures the dynamic physiological states and the underlying spatial-temporal mapping [12].

4 Conclusion

In this paper, we presented DSTP, a novel framework for EEG-to-fMRI synthesis that explicitly models the dynamic spatial-temporal relationship between scalp signals and brain regions. By integrating a Dynamic Heterogeneous Graph Neural Network (DHGNN) and a learnable Gaussian Mask mechanism, our model effectively addresses the challenges of spatial volume conduction and regional variability in hemodynamic delays. Extensive experiments demonstrate that DSTP achieves state-of-the-art performance in both intra-subject and inter-subject settings. Furthermore, its superior zero-shot generalization across different tasks and spatial resolutions suggests that DSTP captures intrinsic physiological mapping rules rather than merely overfitting to specific datasets, offering a robust and scalable solution for high-resolution brain activity monitoring. A promising direction for future work is to further investigate the physiological plausibility of the learned temporal parameters and their relationship to region-specific hemodynamic responses.

Acknowledgments. This work is supported in part by Hong Kong RGC General Research Fund (No.: PolyU-15205924); and in part by the Research Institute for Artificial Intelligence of Things, The Hong Kong Polytechnic University. We would also like to acknowledge datasets acquired in the Advanced MRI Section at the National Institute of Neurological Disorders and Stroke (NINDS), National Institutes of Health (NIH).

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

References

1. Bryce, N.V., Flournoy, J.C., Moreira, J.F.G., Rosen, M.L., Sambook, K.A., Mair, P., McLaughlin, K.A.: Brain parcellation selection: An overlooked decision point with meaningful effects on individual differences in resting-state functional connectivity. NeuroImage **243**, 118487 (2021)
2. Calhas, D., Henriques, R.: Eeg to fmri synthesis benefits from attentional graphs of electrode relationships. arXiv preprint arXiv:2203.03481 (2022)
3. Dadi, K., Varoquaux, G., Machlouzarides-Shalit, A., Gorgolewski, K.J., Wassermann, D., Thirion, B., Mensch, A.: Fine-grain atlases of functional modes for fmri analysis. NeuroImage **221**, 117126 (2020)
4. Han, K., Su, Y., He, L., Zhan, L., Plis, S., Calhoun, V., Yang, C.: Rethinking functional brain connectome analysis: do graph deep learning models help. npj Artificial Intelligence **2**(1), 19 (2026)
5. He, B., Ji, E., Zong, X., Liang, Z., Huang, G., Zhang, L.: Grasti-acl: Graph spatial-temporal infomax with adversarial contrastive learning for brain disorders diagnosis based on resting-state fmri. Medical Image Analysis p. 103815 (2025)
6. He, B., Sohrabpour, A., Brown, E., Liu, Z.: Electrophysiological source imaging: a noninvasive window to brain dynamics. Annual review of biomedical engineering **20**(1), 171–196 (2018)

7. Jiang, W.B., Zhao, L., Lu, B.L.: Large brain model for learning generic representations with tremendous eeg data in bci. In: International Conference on Learning Representations. vol. 2024, pp. 16405–16426 (2024)
8. Kieliba, P., Madugula, S., Filippini, N., Duff, E.P., Makin, T.R.: Large-scale intrinsic connectivity is consistent across varying task demands. PloS one **14**(4), e0213861 (2019)
9. Li, Y., Lou, A., Li, C., Wang, S., Pourmotabbed, H., Xu, Z., Zhang, S., Englot, D.J., Kolouri, S., Moyer, D., Bayrak, R.G., Chang, C.: Unebolt: A unified model for eeg-to-bold translation and functional connectivity reconstruction. In: Proceedings of The 9th International Conference on Medical Imaging with Deep Learning. Proceedings of Machine Learning Research, vol. 315, pp. 2338–2351 (2026)
10. 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. vol. 12926, pp. 795–800. SPIE (2024)
11. 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. Advances in neural information processing systems **37**, 23378–23405 (2024)
12. Mitra, A., Snyder, A.Z., Blazey, T., Raichle, M.E.: Lag threads organize the brain's intrinsic activity. Proceedings of the National Academy of Sciences **112**(17), E2235–E2244 (2015)
13. Noman, F., Phan, R.C.W., Ombao, H., Ting, C.M.: Adaptive graph learning with multi-graph convolutions for brain disorder classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 56–65. Springer (2025)
14. Wagh, N., Wei, J., Rawal, S., Berry, B.M., Varatharajah, Y.: Evaluating latent space robustness and uncertainty of eeg-ml models under realistic distribution shifts. Advances in Neural Information Processing Systems **35**, 21142–21156 (2022)
15. Wang, G., Liu, W., He, Y., Xu, C., Ma, L., Li, H.: Eegpt: Pretrained transformer for universal and reliable representation of eeg signals. Advances in Neural Information Processing Systems **37**, 39249–39280 (2024)
16. Yang, C., Westover, M., Sun, J.: Biot: Biosignal transformer for cross-data learning in the wild. Advances in Neural Information Processing Systems **36**, 78240–78260 (2023)
17. Yao, W., Lyu, Z., Mahmud, M., Zhong, N., Lei, B., Wang, S.: Catd: Unified representation learning for eeg-to-fmri cross-modal generation. IEEE Transactions on Medical Imaging (2025)
18. Zhang, Y., Yu, Y., Li, H., Wu, A., Chen, X., Liu, J., Zeng, L.L., Hu, D.: Dmae-eeg: a pretraining framework for eeg spatiotemporal representation learning. IEEE Transactions on Neural Networks and Learning Systems (2025)
19. Zheng, J., Liu, J., Cao, J., Wang, X., Yang, H., Chen, Y., Yu, P.S.: Dp-dgad: A generalist dynamic graph anomaly detector with dynamic prototypes. arXiv preprint arXiv:2508.00664 (2025)
20. Zheng, J., Saxena, D., Cao, J.: Coin-gnn: Inductive spatial-temporal prediction for continuous distribution shifts via graph neural networks. IEEE Transactions on Knowledge and Data Engineering (2025)