GRFBRAIN: GRAPH-STRUCTURED RECTIFIED FLOWS FOR EEG DYNAMIC MODELING

Anonymous authors
Paper under double-blind review

ABSTRACT

Forecasting time-varying functional connectivity from electroencephalography (EEG) requires modeling both history-dependent trends and structured variability across channels. Conditional flow matching provides a framework for distributional forecasting, yet it remains unclear whether graph-informed source distributions offer practical advantages over isotropic noise and strong deterministic predictors. We introduce a graph-structured residual flow framework that separates conditional mean prediction from stochastic residual transport. A history-only predictor estimates the future connectivity graph, while a graph Gaussian source encodes dependencies derived from past connectivity through a Laplacian-based covariance. A conditional velocity field transports source samples to future graph residuals, with transport time explicitly distinguished from physical EEG time. Our study identifies the conditions and controls needed to distinguish useful residual transport from improvements attributable to deterministic prediction, learned representations, and sampling effects.

1 INTRODUCTION

Learning informative representations of electroencephalography (EEG) is central to clinical settings such as seizure detection (Pradeepkumar et al., 2026). What makes this difficult is that the relevant information is not located in a single channel. A seizure, for instance, is characterized less by abnormal activity at one electrode than by the emergence, spread, and dissolution of synchrony across electrodes. Representations must therefore encode two coupled aspects of brain dynamics: how channel-wise activity evolves, and how the functional relationships among channels over time.

Predictive learning offers a route to EEG representation learning by requiring a model to infer future activity from its history. Existing graph-based methods pursue this idea through next-window prediction, evolving graph convolutions, or continuous-time spectral dynamics (Tang et al., 2022; Kotoge et al., 2025; Jia et al., 2026). However, these objectives are deterministic: regression maps each history to a single future, encouraging representations of the conditional mean while suppressing structured uncertainty. This is particularly limiting near transitions, where similar histories may lead to different futures and discriminative cues can lie in deviations from the average trajectory.

Distributional objectives model multiple plausible futures rather than a single point estimate. Flow matching (FM) realizes this by learning a history-conditioned velocity field that transports a simple source distribution toward future states (Lipman et al., 2023), encouraging representations to capture both predictable trends and residual variability. In EEG, however, this variability is structured across channels and reflected in recent functional connectivity, whereas standard FM starts from isotropic Gaussian noise that treats channels and channel pairs as exchangeable, leaving the velocity field to recover these dependencies during transport. The central question is therefore not whether isotropic noise is sufficiently expressive, but whether a history-informed source defines a more informative predictive task for representation learning.

This observation suggests that the source is not merely a sampling convenience but part of the learning task, and hence an inductive bias on the learned representation, as shown in Figure 1. We therefore ask: *can the channel dependencies observed in the history define the stochastic prediction task itself, rather than serving only as conditioning input to the encoder?*

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Table 1: Comparison of representative methods for modeling EEG dynamics.

Model	Formulation	Continuous	Graphical	Stochastic
CNN-LSTM	Spatial-temporal encoding	✗	✗	✗
BIOT (Yang et al., 2023)	Biosignal Transformer	✗	✗	✗
EvolveGCN (Pareja et al., 2020)	Evolving graph convolution	✗	✓	✗
DCRNN	Diffusion-graph convolutional RNN	✗	✓	✗
neural ODE	latent dynamics	✓	✗	✗
neural SDE	Stochastic latent dynamics	✓	✗	✓
ODEBrain (Jia et al., 2026)	Temporal-spatial latent dynamics	✓	✓	✗
Flow Matching	Isotropic latent transport	✓	✗	✓
GRFBrain	**Conditioned joint graph flow**	✓	✓	✓

✓ denotes that the property is implemented; ✗ denotes that it is absent from the evaluated formulation.

In this paper, we answer this question with GRFBrain, a graph-structured residual flow framework for predictive EEG representation learning. GRFBrain first decomposes the future EEG state into a history-derived reference and a residual, allowing the flow objective to focus on the variation that remains beyond the predictable trend. It then constructs a history-conditioned Gaussian source whose covariance is derived from the recent functional graph, so that source perturbations reflect the channel dependencies observed in each sample. The source energy is normalized to match that of an IID Gaussian control, so that their comparison isolates the role of relational structure from that of noise magnitude. Finally, a joint node–edge velocity network transports node and edge residuals in a shared state space, so that their bidirectional interactions shape the learned representation. For downstream prediction, the velocity network is evaluated once at the history-derived reference state, which corresponds to zero in residual coordinates. This readout requires neither source sampling nor numerical integration. As summarized in Table 1, existing methods are typically either graph-aware but deterministic, or stochastic without adaptive graph structure. Our contributions are threefold:

- We formulate predictive EEG representation learning as conditional transport in a joint node–edge residual space, separating the history-predictable component of the future from its remaining structured variability.
- We introduce a history-conditioned graph Gaussian source with energy matched to an IID baseline, together with a coupled node–edge velocity network that models interactions between channel activity and functional connectivity.
- We evaluate GSRF on seizure detection and conduct controlled studies of source geometry, node–edge coupling, transport targets, and forecasting horizon. GSRF improves F1 over deterministic graph-dynamics and standard flow-matching baselines, while the ablations isolate the benefit of history-adaptive source structure.

2 Related Works

Temporal graph methods for modeling EEG dynamics. Multichannel EEG can be represented as a graph, with channels as nodes and inter-channel relations as edges, motivating models that jointly capture temporal activity and spatial dependencies. Early work introduced temporal graph convolutions for seizure detection (Covert et al., 2019), followed by diffusion-convolutional recurrent networks with future-signal pretraining on geometric or correlation-based graphs (Tang et al., 2022). Later methods extended this framework through state-space modeling and dynamic graph learning in GraphS4mer (Tang et al., 2023), efficient residual recurrent updates in REST (Afzal et al., 2024), and explicit time-varying node-edge states in EvoBrain (Kotoge et al., 2025). Continuous-time extensions include BrainODE for reconstructing irregularly sampled fMRI with graph-aided neural ODEs (Han et al., 2024) and ODEBrain for learning continuous latent EEG graph trajectories under future-graph supervision (Jia et al., 2026). Despite this progression from temporal encoding to dynamic and continuous trajectories, these methods produce discriminative representations, point forecasts, or single trajectories rather than conditional distributions over future node-edge states.

Generative methods for brain modeling. Generative graph models learn joint distributions over node and edge variables. GDSS evolves continuous graph states through coupled stochastic differen-

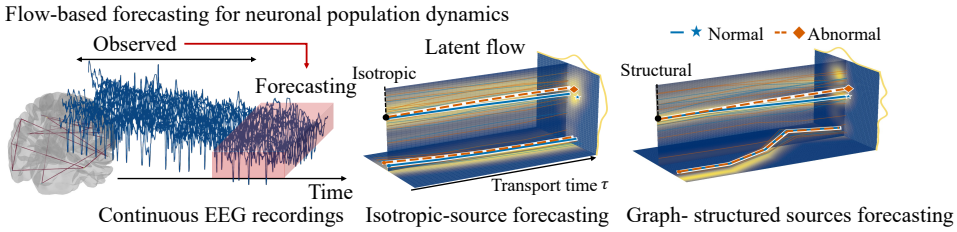


Figure 1: (Left) Continuous EEG real-time neuronal activity recordings. (Mid) An isotropic gaussian source-based method learns dynamic representations. (Right) Ours provides a graph conditioned source-based approach for learning the neuronal population dynamics.

tial equations, while DiGress denoises categorical graphs in discrete spaces (Jo et al., 2022; Vignac et al., 2023); both target graph synthesis rather than EEG-conditioned forecasting. Flow Matching (FM) learns continuous transport along prescribed probability paths (Lipman et al., 2023), with extensions addressing graph generation and geometry (Eijkelboom et al., 2024; Jiang et al., 2026), source–target coupling and conditional sources (Tong et al., 2024; Chen et al., 2025; Issachar et al., 2025; Kim et al., 2026), and topology-aware priors or objectives (Borovitskiy et al., 2021; Wyrwal et al., 2026). Closest to our setting, GiFlow uses graph-informed priors for spatiotemporal imputation but does not jointly transport evolving edges (Zhang et al., 2026), whereas DiffeoCFM models brain-connectivity distributions without conditioning their evolution on EEG history (Collas et al., 2025). `GRFBrain` addresses this gap through a history-conditioned graph Gaussian source, reference-centered future residuals, and a shared velocity field for joint node-edge transport.

3 Preliminary and Problem Formulation

Conditional Flow Matching. Given condition $\mathcal{C}$, we sample $\boldsymbol{z}_0 \sim p_0(\cdot \mid \mathcal{C})$, $\boldsymbol{z}_1 \sim p_1(\cdot \mid \mathcal{C})$, and $\tau \sim \mathcal{U}[0,1]$. The linear path $\boldsymbol{z}_\tau = (1-\tau)\boldsymbol{z}_0 + \tau \boldsymbol{z}_1$ has target velocity $\boldsymbol{u}_\tau = \boldsymbol{z}_1 - \boldsymbol{z}_0$. Conditional Flow Matching (CFM) learns

$$\mathcal{L}_{\mathrm{CFM}}(\theta) = \mathbb{E}_{\mathcal{C},\boldsymbol{z}_0,\boldsymbol{z}_1,\tau}\left[\|v_\theta(\boldsymbol{z}_\tau, \tau \mid \mathcal{C}) - \boldsymbol{u}_\tau\|_2^2\right]. \tag{1}$$

Although $\boldsymbol{u}_\tau$ is defined per sampled pair, the optimal regressor is the marginal velocity

$$v^\star(\boldsymbol{z}, \tau \mid \mathcal{C}) = \mathbb{E}[\boldsymbol{z}_1 - \boldsymbol{z}_0 \mid \boldsymbol{z}_\tau = \boldsymbol{z}, \mathcal{C}], \tag{2}$$

whose flow transports $p_0(\cdot \mid \mathcal{C})$ to $p_1(\cdot \mid \mathcal{C})$ (Lipman et al., 2023). Training requires no numerical integration, while generation solves $d\boldsymbol{z}_\tau/d\tau = v_\theta(\boldsymbol{z}_\tau, \tau \mid \mathcal{C})$. Here, τ is an artificial transport coordinate, distinct from the physical EEG window time s; the learned velocity therefore describes distributional transport rather than physiological dynamics.

Source-Target Mismatch in EEG Graph Transport. Standard CFM uses the isotropic source $p_0(\boldsymbol{z} \mid \mathcal{C}) = \mathcal{N}(\mathbf{0}, \boldsymbol{I})$, as an independent identically distributed (IID). Although expressive in principle, this source ignores the relational geometry available in the past EEGs. Future EEG graphs couple channel-wise node features with pairwise edges, and their unpredictable component retains history-dependent cross-channel structure. An isotropic source needs a finite-capacity velocity field to learn the conditional shift, and node-edge dependence. This mismatch raises two questions.

1. How should the state and source be defined? The future state should be centered by its temporal component, with source and target defined in the same residual node-edge space. The source should use observed history, encode the graph topology to control the total variance, so that a comparison with an IID source isolates the covariance structure rather than the noise scale.

2. How should nodes and edges be transported jointly? Because node activity and functional connectivity are coupled, the velocity field should support bidirectional node-to-edge and edge-to-node interactions instead of fixing the adjacency matrix or learning independent fields.

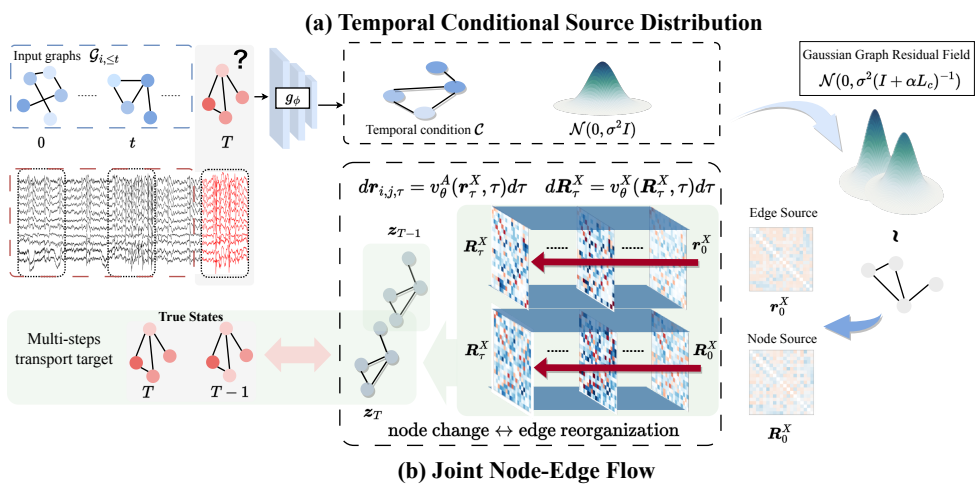


Figure 2: Overview of the proposed GRFBrain. (a) Temporal EEG features and functional association graphs define state-dependent states and a graph Gaussian residual source, whose covariance is adapted to the observed graph states. (b) A joint node-edge velocity network transports node and edge residuals along a shared flow coordinate, coupling channel-wise state variations.

These requirements motivate our predictor-centered residual state, temporal-conditioned graph source, and shared node-edge velocity field.

4 METHODOLOGY

Figure 2 shows the system overview of GRFBrain. First, it decomposes the temporal EEG graph embedding with sequential encoding as condition states to construct a graph-structured Gaussian source. Second, a Joint Node-Edge velocity field between source and target state.

4.1 HISTORY-CONDITIONED GRAPH GAUSSIAN SOURCE

Let $\mathcal{C} = \{(\boldsymbol{X}_s, \boldsymbol{A}_s)\}_{s=t-H+1}^{t}$ denote an EEG history of H physical windows, where $\boldsymbol{X}_s \in \mathbb{R}^{N\times F}$ contains channel-wise spectral features and $\boldsymbol{A}_s \in \mathbb{R}^{N\times N}$ is a directed top-k graph. We represent its undirected continuous edge state by $\boldsymbol{a}_s = \mathcal{U}(\boldsymbol{S}_s) \in \mathbb{R}^m$, where $\boldsymbol{S}_s = (\boldsymbol{A}_s + \boldsymbol{A}_s^\top)/2$ has zero diagonal, $\mathcal{U}$ extracts its strict upper triangle, and $m = N(N-1)/2$.

We model the future graph relative to a history-dependent reference. Let $\boldsymbol{M}_X \in \mathbb{R}^{N\times F}$ be the mean historical node state and $\boldsymbol{m}_A \in \mathbb{R}^m$ the edge reference produced by a frozen graph forecaster g_ϕ. The target residuals are

$$\boldsymbol{M}_X = \frac{1}{H}\sum_{s=t-H+1}^{t} \boldsymbol{X}_s, \qquad \boldsymbol{R}_1^X = \boldsymbol{X}_{t+1} - \boldsymbol{M}_X, \qquad \boldsymbol{r}_1^A = \boldsymbol{a}_{t+1} - \boldsymbol{m}_A. \tag{3}$$

This decomposition assigns the predictable component to the reference model and leaves the conditional variation to the flow. Unlike an independent and identically distributed (i.i.d.) Gaussian source, our source uses the relational geometry of the observed history. Define $\overline{\boldsymbol{S}}_{\mathcal{C}} = H^{-1}\sum_{s=t-H+1}^{t} \boldsymbol{S}_s$ and construct

$$\boldsymbol{L}_{\mathcal{C}} = \boldsymbol{I}_N - \boldsymbol{D}_{\mathcal{C}}^{-1/2}\overline{\boldsymbol{S}}_{\mathcal{C}}\boldsymbol{D}_{\mathcal{C}}^{-1/2}, \qquad \boldsymbol{K}_{\mathcal{C}} = \frac{N(\boldsymbol{I}_N + \alpha\boldsymbol{L}_{\mathcal{C}})^{-\nu}}{\mathrm{tr}[(\boldsymbol{I}_N + \alpha\boldsymbol{L}_{\mathcal{C}})^{-\nu}]}, \qquad \boldsymbol{G}_{\mathcal{C}} = \boldsymbol{K}_{\mathcal{C}}^{1/2}. \tag{4}$$

Here, $\boldsymbol{D}_{\mathcal{C}}$ is the stabilized degree matrix, $\alpha \ge 0$, and $\nu > 0$. Trace normalization gives $\mathrm{tr}(\boldsymbol{K}_{\mathcal{C}}) = N$, while $(1+\alpha\lambda_i)^{-\nu}$ suppresses high graph-frequency modes. Node and edge residuals are sampled as

$$\boldsymbol{R}_0^X = \sigma_X \boldsymbol{G}_{\mathcal{C}}\boldsymbol{\epsilon}_X, \qquad \boldsymbol{r}_0^A = \sigma_A s_{\mathcal{C}}\,\mathcal{U}(\boldsymbol{G}_{\mathcal{C}}\boldsymbol{\Xi}_A\boldsymbol{G}_{\mathcal{C}}), \qquad \boldsymbol{\Xi}_A = \mathcal{U}^{-1}(\boldsymbol{\epsilon}_A), \tag{5}$$

where $\boldsymbol{\epsilon}_X$ and $\boldsymbol{\epsilon}_A$ are independent standard Gaussian innovations, and $s_{\mathcal{C}}$ normalizes the edge energy. The two sources are conditionally independent but share $\boldsymbol{G}_{\mathcal{C}}$, so their covariance follows the historical graph while their total energies remain matched to the i.i.d. control. In the original coordinates, the source is centered at $(\boldsymbol{M}_X, \boldsymbol{m}_A)$. The graph is used only as an empirical conditioning structure and is not interpreted as an anatomical or causal brain network. Appendix B provides the forecaster objective, numerical stabilization, and covariance derivations.

4.2 Joint Node–Edge Residual Flow

We model node and edge residuals with a shared velocity field

$$v_\theta : \mathbb{R}^{N\times F} \times \mathbb{R}^m \times [0,1] \times \mathcal{C} \to \mathbb{R}^{N\times F} \times \mathbb{R}^m,$$

implemented by a joint node–edge Transformer (Jo et al., 2022; Vignac et al., 2023). It maintains one node token per channel and one symmetric edge token per unordered channel pair, including pairs with zero observed weight.

Let $\boldsymbol{c}_i$ denote the temporal encoding of channel i, and let $\boldsymbol{c}_A$ summarize the edge history. The global condition $\boldsymbol{q}_\tau = [N^{-1}\sum_i \boldsymbol{c}_i; \boldsymbol{c}_A; \mathrm{Emb}(\tau)]$ modulates each block through feature-wise linear modulation (FiLM) (Perez et al., 2018). Initial tokens are

$$\begin{aligned} \boldsymbol{h}_i^{(0)} &= P_X([\boldsymbol{R}^X_{\tau,i}; \boldsymbol{M}_{X,i}]) + P_N(\boldsymbol{c}_i), \\ \boldsymbol{e}_{ij}^{(0)} &= P_A([r^A_{\tau,ij}; m_{A,ij}]) + P_P(\chi(\boldsymbol{c}_i, \boldsymbol{c}_j)) + P_O([L_{\mathcal{C},ij}; K_{\mathcal{C},ij}]), \qquad i<j, \end{aligned} \tag{6}$$

where P_X, P_N, P_A, P_P, P_O are learned projections and $\chi(\boldsymbol{u},\boldsymbol{v}) = [\boldsymbol{u}+\boldsymbol{v}; |\boldsymbol{u}-\boldsymbol{v}|; \boldsymbol{u}\odot\boldsymbol{v}]$ is symmetric in its arguments.

Each interaction block uses edge tokens as attention biases and incident-edge messages for node updates. The updated endpoint tokens then revise their edge through the same symmetric descriptor:

$$\begin{aligned} \boldsymbol{h}_i^{(\ell+1)} &= \mathrm{NodeBlock}_\ell\left(\boldsymbol{h}_i^{(\ell)}, \{\boldsymbol{e}_{ij}^{(\ell)}\}_j; \boldsymbol{q}_\tau\right), \\ \boldsymbol{e}_{ij}^{(\ell+1)} &= \mathrm{EdgeBlock}_\ell\left(\boldsymbol{e}_{ij}^{(\ell)}, \chi(\boldsymbol{h}_i^{(\ell+1)}, \boldsymbol{h}_j^{(\ell+1)}); \boldsymbol{q}_\tau\right). \end{aligned} \tag{7}$$

After J blocks, separate affine heads produce $\boldsymbol{v}_\theta^X \in \mathbb{R}^{N\times F}$ and $\boldsymbol{v}_\theta^A \in \mathbb{R}^m$. Thus, edges guide node updates while node states jointly determine edge updates.

4.3 Conditional Flow-Matching Objective

For each future target, we sample $(\boldsymbol{R}_0^X, \boldsymbol{r}_0^A)$ from the conditional source and draw a shared $\tau \sim \mathcal{U}(0,1)$. Node and edge residuals follow

$$\begin{aligned} \boldsymbol{R}_\tau^X &= (1-\tau)\boldsymbol{R}_0^X + \tau \boldsymbol{R}_1^X, \quad & \boldsymbol{u}^X &= \boldsymbol{R}_1^X - \boldsymbol{R}_0^X, \\ \boldsymbol{r}_\tau^A &= (1-\tau)\boldsymbol{r}_0^A + \tau \boldsymbol{r}_1^A, & \boldsymbol{u}^A &= \boldsymbol{r}_1^A - \boldsymbol{r}_0^A. \end{aligned} \tag{8}$$

Source and target residuals are paired independently given $\mathcal{C}$; hence, this is conditional linear Flow Matching rather than an optimal-transport coupling.

To balance the two state spaces, let q_X and q_A be the fixed per-coordinate target-velocity energies estimated from the training set: $q_X = \max\{\widehat{\mathbb{E}}_{\mathrm{train}}[\|\boldsymbol{u}^X\|_F^2/(NF)], \epsilon_q\}$ and $q_A = \max\{\widehat{\mathbb{E}}_{\mathrm{train}}[\|\boldsymbol{u}^A\|_2^2/m], \epsilon_q\}$. The normalized conditional Flow Matching (CFM) objective is

$$\mathcal{L}_{\mathrm{CFM}}(\theta) = \mathbb{E}\left[\frac{\|\boldsymbol{v}_\theta^X - \boldsymbol{u}^X\|_F^2}{NFq_X} + \frac{\|\boldsymbol{v}_\theta^A - \boldsymbol{u}^A\|_2^2}{mq_A}\right], \tag{9}$$

where the expectation covers training histories and targets, conditional source samples, and τ. The objective remains standard CFM (Lipman et al., 2023); the proposed design changes the residual target, source covariance, and joint node–edge state space.

5 EXPERIMENTS

In this section, we conduct experiments to answer the following research questions:
- RQ1. Does GRFBrain strengthen seizure detection capacity through temporal-conditioned flow matching?
- RQ2. How does source geometry affect the development of latent transport?
- RQ3. Does the source-target objective of $\mathcal{L}_{\text{CFM}}$ facilitate dynamic optimization?
More detailed experiment settings can be found in the Appendix A.

5.1 EXPERIMENTAL SETUP

Tasks. We evaluate GRFBrain on two downstream tasks: *seizure detection* and *abnormal EEG classification*. These complementary tasks assess the utility of the learned EEG representations in identifying seizure-related activity and distinguishing abnormal from EEG patterns. Detailed descriptions of the datasets are provided in Section C.4.

Baseline methods. We organize the baselines with the following learning paradigms. First, we consider conventional sequence-based EEG models, including CNN-LSTM (Ahmedt-Aristizabal et al., 2020) and the Transformer-based biosignal model BIOT (Yang et al., 2023). Second, we compare with discrete time graph dynamics models, including DCRNN (Li et al., 2017) and EvolveGCN (Pareja et al., 2020), which model temporal graph evolution through recurrent mechanisms. We also include continuous time dynamics models. ODE-RNN (Rubanova et al., 2019), a neural stochastic differential equation (SDE) model (Liu et al., 2019), graph differential equations (GDEs) (Poli et al., 2019), and ODEBrain (Jia et al., 2026), which explicitly models continuous EEG graph dynamics. Finally, we also compare with flow-based predictive baselines, including standard Flow matching (Lipman et al., 2023) and Rectified flow (Liu et al., 2022).

Table 2: Results (AUROC↑, F1↑) on TUSZ (12s and 60s seizure detection). The upper block compares GRFBrain with discrete and continuous baselines (bold = best); the lower block ablates the node-edge interaction routes within GRFBrain.

Model	Method	T(s)	AUROC	F1
Flow-based	BIOT	12	0.772±0.006	0.294±0.006
		60	0.642±0.009	0.256±0.003
	DCRNN	12	0.825±0.002	0.416±0.009
		60	0.802±0.003	0.413±0.005
	Flow matching	12	0.813±0.002	0.372±0.013
		60	0.725±0.006	0.311±0.028
	GRFBrain	12	**0.877±0.003**	**0.523±0.014**
		60	**0.831±0.004**	**0.488±0.032**
GRFBrain	Edge-to-Node	12	0.867±0.004	0.488±0.007
		60	0.821±0.034	0.424±0.003
	Node-to-Edge	12	0.848±0.017	0.462±0.013
		60	0.817±0.029	0.414±0.047
	No-cross	12	0.673±0.007	0.374±0.033
		60	0.519±0.006	0.334±0.017

Metrics. To answer **RQ1**, we evaluate the model using the Area Under the Receiver Operating Characteristic Curve (AUROC) and the F1 score. The AUROC measures the ability of the models across varying thresholds, while the F1 score highlights the balance between precision and recall at its optimal threshold for classification. For **RQ2**, we measure the structural similarity of the predicted graph using the Global Jaccard Index (GJI) $\texttt{GJI}(\mathcal{E}_{true}, \mathcal{E}_{Pred}) = \frac{|\mathcal{E}_{true} \cap \mathcal{E}_{Pred}|}{|\mathcal{E}_{true} \cup \mathcal{E}_{Pred}|}$ (Castrillo et al., 2018). For **RQ3**, We compute the cosine similarity of predicted node embeddings.

5.2 RESULTS

5.2.1 MAIN RESULT

RQ1 concerns the forecasting-based transport capability on EEG. Table 3 compares GRFBrain with conventional EEG models, discrete graph-dynamics methods, continuous-time models, and generic flow-based predictors on TUSZ and TUAB. On TUSZ, the multi-step variant achieves the highest Accuracy of 0.883 ± 0.004, improving over the corresponding ODEBrain variant across all three metrics. The single-step variant obtains the best F1 score of 0.523 ± 0.014, an absolute gain of 2.7 percentage points over ODEBrain (0.496 ± 0.017), while remaining competitive in AUROC (0.877 ± 0.003 versus 0.881 ± 0.006). On TUAB, both variants consistently outperform their

Table 3: Main results on TUSZ (12s seizure detection) and TUAB abnormal EEG classification, reported as mean ± standard deviation over three seeds. **Bold** and underline indicate the best and second-best results for each metric, respectively. † denotes models trained with a multi-step predictive objective, and ‡ denotes models trained with a single-step predictive objective.

Method	TUSZ			TUAB		
	Acc	F1	AUROC	Acc	F1	AUROC
CNN-LSTM	0.735 ± 0.003	0.347 ± 0.012	0.757 ± 0.003	0.741 ± 0.002	0.736 ± 0.007	0.813 ± 0.003
BIOT	0.702 ± 0.003	0.294 ± 0.006	0.772 ± 0.006	0.717 ± 0.002	0.713 ± 0.004	0.788 ± 0.002
EvolveGCN	0.769 ± 0.002	0.385 ± 0.005	0.791 ± 0.004	0.708 ± 0.003	0.707 ± 0.002	0.777 ± 0.003
DCRNN	0.816 ± 0.002	0.416 ± 0.009	0.825 ± 0.002	0.768 ± 0.004	0.769 ± 0.002	0.848 ± 0.002
latent-ODE	0.827 ± 0.004	0.470 ± 0.005	0.849 ± 0.004	0.749 ± 0.003	0.745 ± 0.002	0.829 ± 0.004
latent-ODE (RK4)	0.821 ± 0.003	0.465 ± 0.001	0.845 ± 0.004	0.746 ± 0.002	0.739 ± 0.002	0.823 ± 0.003
ODE-RNN	0.802 ± 0.002	0.455 ± 0.007	0.855 ± 0.003	0.751 ± 0.003	0.744 ± 0.004	0.838 ± 0.005
neural SDE	0.857 ± 0.002	0.467 ± 0.003	0.851 ± 0.002	0.768 ± 0.003	0.751 ± 0.003	0.834 ± 0.002
GDEs	0.849 ± 0.003	0.475 ± 0.005	0.841 ± 0.003	0.757 ± 0.003	0.737 ± 0.006	0.823 ± 0.004
Flow matching	0.874 ± 0.002	0.372 ± 0.013	0.813 ± 0.002	0.741 ± 0.011	0.719 ± 0.007	0.803 ± 0.006
Rectified flow	0.721 ± 0.001	0.195 ± 0.008	0.566 ± 0.002	0.753 ± 0.004	0.721 ± 0.004	0.811 ± 0.003
ODEBrain†	0.869 ± 0.003	0.488 ± 0.015	0.875 ± 0.005	0.771 ± 0.005	0.770 ± 0.005	0.849 ± 0.003
ODEBrain‡	0.877 ± 0.004	0.496 ± 0.017	$\mathbf{0.881 \pm 0.006}$	0.778 ± 0.003	0.774 ± 0.005	0.857 ± 0.005
GRFBrain†	$\mathbf{0.883 \pm 0.004}$	0.503 ± 0.012	0.876 ± 0.001	$\mathbf{0.789 \pm 0.006}$	$\mathbf{0.786 \pm 0.004}$	$\mathbf{0.878 \pm 0.004}$
GRFBrain‡	0.878 ± 0.001	$\mathbf{0.523 \pm 0.014}$	0.877 ± 0.003	0.787 ± 0.004	0.783 ± 0.004	0.877 ± 0.006

– denotes a configuration that has not been evaluated on the corresponding corpus.

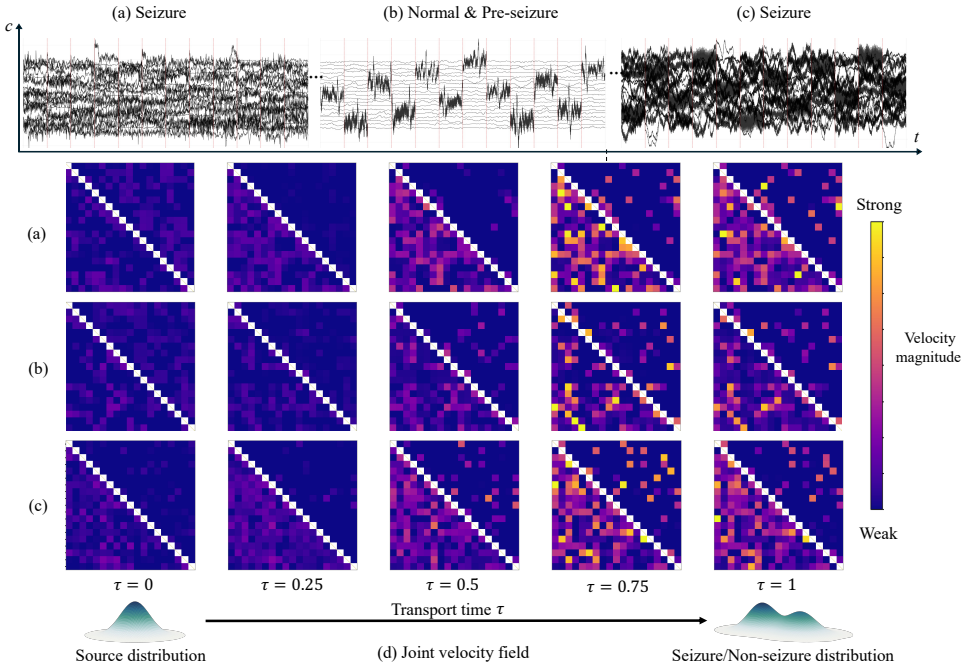


Figure 3: Visualization of the proposed joint node-edge velocity field v_θ (middle) obtained by GRFBrain. The learned velocity exists the distinct patterns of graph reconfiguration, which seizure represents higher and dense magnitude and non-seizure represents lower and sparse magnitude.

ODEBrain counterparts, with the multi-step variant achieving the best Accuracy, F1, and AUROC of 0.789 ± 0.006, 0.786 ± 0.004, and 0.878 ± 0.004, respectively. Generic flow learning alone does not reproduce these gains. On TUSZ, standard flow matching attains a relatively high Accuracy of 0.874 ± 0.002, but its F1 and AUROC decrease to 0.372 ± 0.013 and 0.813 ± 0.006; Rectified Flow performs worse across all three metrics. Taken together, these results answer RQ1: distributional forecasting improves EEG representation when flow matching is coupled with history-conditioned graph structure and joint node-edge transport, rather than applied as a generic predictive objective.

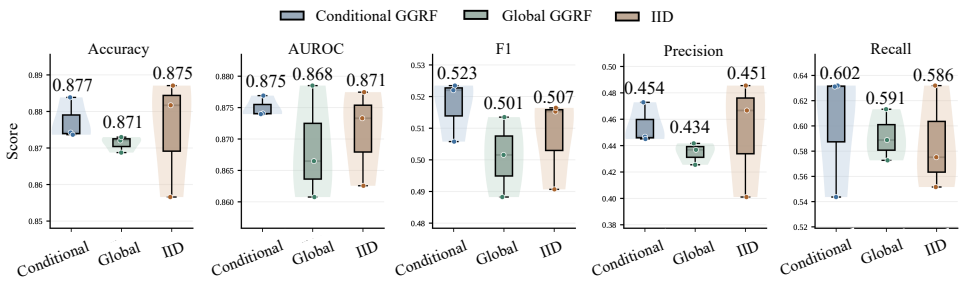
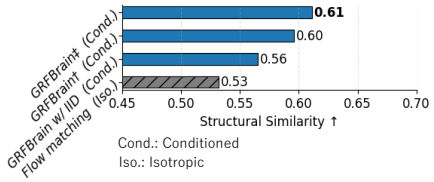
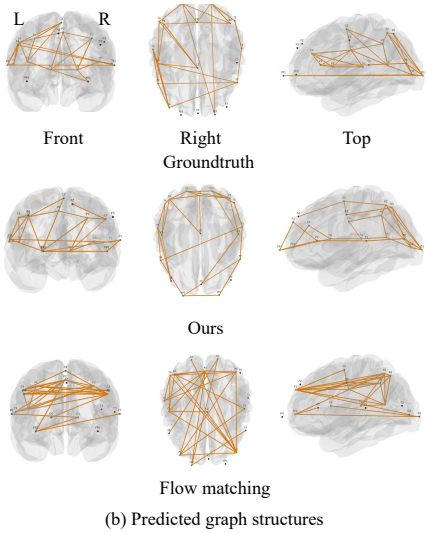


Figure 5: Effect of source geometry in `GRFBrain`. IID Gaussian, global GGRF, and our conditioned GGRF sources are compared across downstream classification metrics. Our conditional GGRF consistently improves the learned representation quality.

Beyond downstream performance, we further investigate whether the learned joint field v_θ captures state-dependent reorganization of channel associations. We analyze the edge-velocity field v_θ^A across the transport time τ. Positive and negative velocities respectively indicate tendencies to strengthen and weaken pairwise associations. As shown in Figure 3, edge-velocity magnitudes are relatively weak near the source side and become increasingly structured toward larger τ. Rather than exhibiting a uniform global change, the field concentrates on subsets of channel pairs, suggesting that the learned representation is associated with selective graph reconfiguration. These results suggest that the joint velocity encourages representations that encode not only channel-wise spectral states but also how their relational organization may be reconfigured. This provides a complementary explanation for the downstream gains by interpreting the learned flow as an estimation of latent EEG dynamics. Taken together, the quantitative and qualitative results answer **RQ1** affirmatively. `GRFBrain` improves seizure detection by coupling conditional flow matching with joint node-edge transport, thereby transforming temporal-dependent graph reconfiguration into class representations.



(a) Predicted graph structural similarity scores result



(b) Predicted graph structures

Figure 4: Results on (a) graph similarity and (b) functional connections.

RQ2 concerns the geometry of noise. Figure 5 compares the conditional GGRF with a global GGRF and an IID Gaussian source. The conditional GGRF performs best across all five metrics, achieving an Accuracy of 0.877, an AUROC of 0.875, and an F1 score of 0.523. Compared with the global GGRF, it improves F1 and AUROC by 2.2 and 0.7, respectively. It also outperforms the IID Gaussian in F1, AUROC, and Recall despite their similar Accuracy and Precision. The weaker performance of the global GGRF indicates that graph correlation alone is insufficient. The IID Gaussian discards inter-channel dependencies, whereas the global GGRF imposes the same relational geometry on each sample. In contrast, the conditional GGRF adapts its covariance to the observed EEG history. Because the source determines both the interpolated states and their velocity targets, this sample-specific covariance produces transport paths aligned with the current graph structure. These results answer **RQ2**: latent transport benefits from a history-conditioned source geometry rather than an isotropic or globally shared source, yielding more discriminative node-edge representations.

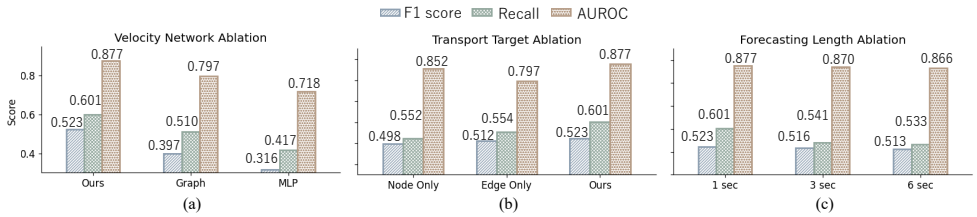


Figure 6: Summary of ablation study. (a) Velocity-network parameterization, (b) node and edge transport targets, and (c) forecasting horizon. The proposed joint node–edge velocity network, joint transport supervision, and short-horizon prediction provide the strongest overall performance.

Necessity of bidirectional node-edge interaction. Table 2 compares velocity-field designs on 12 s and 60 s EEG clips. The full bidirectional field achieves the highest AUROC/F1 of $0.877/0.523$ and $0.831/0.488$, respectively. Its advantage over standard FM shows that the generic objective alone is insufficient, while the within-`GRFBrain` ablations isolate the interaction mechanism. Removing cross-stream messages reduces AUROC/F1 to $0.673/0.374$ and $0.519/0.334$, despite retaining both node and edge states. One-way interaction partially recovers performance, with Edge-to-Node consistently outperforming Node-to-Edge, likely because it directly updates the node representations consumed by the classifier. Neither direction alone matches bidirectional interaction. Relative to Edge-to-Node, the full field improves F1 by 0.035 and 0.064 at 12 s and 60 s, with an AUROC gain of 0.010 in both settings. Thus, within this architecture, representing both state types is insufficient; reciprocal node-edge communication is required for the strongest downstream performance.

RQ3 concerns consistency in the graphs with source-target objective of $\mathcal{L}_{\text{CFM}}$. Figure 4 compares the predicted functional graphs with their target structures. `GRFBrain` increases the graph-similarity score from 0.53 for the FM to 0.61. Its similarity matrix also better preserves local correlations and blockwise dependency patterns, whereas the baseline exhibits larger structural discrepancies. These results support the role of $\mathcal{L}_{\text{CFM}}$ in coordinating node and edge evolution along the transport path. We therefore answer RQ3 affirmatively: source-target supervision improves the structural fidelity of future-graph prediction by guiding the two-stream velocity field to learn consistent node-edge dynamics.

5.3 Ablation Study

Effects of velocity architecture and forecasting options. Figure 6 evaluates the velocity architecture, transport targets, and forecasting horizon. The joint velocity network achieves an F1 of 0.523 and an AUROC of 0.877, outperforming the graph-only ($0.397/0.797$) and MLP ($0.316/0.718$) alternatives. This result attributes the gain to explicit node–edge coupling rather than the Flow Matching objective alone. Joint node–edge supervision also performs best: removing edge supervision reduces F1/Recall from $0.523/0.601$ to $0.498/0.552$, whereas edge-only supervision retains a competitive F1 of 0.512 but lowers AUROC to 0.797. Node dynamics and relational reconfiguration therefore provide complementary predictive information. Finally, extending the horizon from 1 s to 3 s and 6 s progressively decreases F1 ($0.523 \to 0.516 \to 0.513$) and Recall ($0.601 \to 0.541 \to 0.533$), suggesting that the 1 s target better captures rapidly evolving EEG dynamics, while longer horizons introduce greater predictive uncertainty.

6 Conclusion

In this work, we introduced `GRFBrain`, a temporal-conditioned graph Flow Matching framework for predictive EEG representation learning. `GRFBrain` decomposes node-edge sequences into past-derived states and residuals, and learns their joint transport through a bidirectional node-edge velocity field. Experiments on TUSZ and TUAB demonstrate improvements over graph-dynamics and generic flow-based baselines, while the ablations support the roles of conditional source geometry and joint node-edge transport. **Limitation**: Current evaluation is limited to epoched EEG and short-horizon supervision. Future work will investigate probabilistic calibration, longer forecasting horizons, and transfer to other neurological tasks.

AI USE STATEMENT

We used OpenAI Codex to assist with manuscript drafting and English-language editing, literature searches, LaTeX formatting and bibliography preparation, and consistency checks between mathematical descriptions and the implementation. AI outputs were used as suggestions for writing and technical review. The authors remain responsible for the accuracy of the text, references, mathematical statements,experimental claims, and associated research artifacts, including AI-assisted material.

ETHICS STATEMENT

This study uses existing, de-identified clinical EEG recordings from TUAB, TUSZ. TUAB and TUSZ are derived from the Temple University Hospital EEG Corpus. Access to and use of each dataset remain subject to its respective license and data-use requirements. Our experiments evaluate EEG classification on these benchmarks. Performance on these datasets does not establish clinical utility or generalization to other patient populations, hospitals, or acquisition settings; these require separate evaluation.

REFERENCES

Arshia Afzal, Grigorios Chrysos, Volkan Cevher, and Mahsa Shoaran. REST: Efficient and accelerated EEG seizure analysis through residual state updates. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 271–290. PMLR, 2024. URL https://proceedings.mlr.press/v235/afzal24a.html.

David Ahmedt-Aristizabal, Tharindu Fernando, Simon Denman, Lars Petersson, Matthew J. Aburn, and Clinton Fookes. Neural memory networks for seizure type classification. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 569–575. IEEE, 2020. doi: 10.1109/EMBC44109.2020.9175641. URL https://doi.org/10.1109/EMBC44109.2020.9175641.

Viacheslav Borovitskiy, Iskander Azangulov, Alexander Terenin, Peter Mostowsky, Marc Deisenroth, and Nicolas Durrande. Matérn Gaussian Processes on Graphs. In Arindam Banerjee and Kenji Fukumizu (eds.), *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 2593–2601. PMLR, 2021. URL https://proceedings.mlr.press/v130/borovitskiy21a.html.

Eduar Castrillo, Elizabeth León, and Jonatan Gómez. Dynamic structural similarity on graphs. *arXiv preprint arXiv:1805.01419*, 2018.

Chen Chen, Pengsheng Guo, Liangchen Song, Jiasen Lu, Rui Qian, Tsu-Jui Fu, Xinze Wang, Wei Liu, Yinfei Yang, and Alex Schwing. CAR-Flow: Condition-aware reparameterization aligns source and target for better flow matching. In *Advances in Neural Information Processing Systems*, volume 38, pp. 94919–94945. Curran Associates, Inc., 2025. doi: 10.52202/085713-3173. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/8838314582eb031f55aee65a703e76c5-Abstract-Conference.html.

Zheng Chen, Lingwei Zhu, Haohui Jia, and Takashi Matsubara. A two-view eeg representation for brain cognition by composite temporal-spatial contrastive learning. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, pp. 334–342. SIAM, 2023.

Antoine Collas, Ce Ju, Nicolas Salvy, and Bertrand Thirion. Riemannian flow matching for brain connectivity matrices via pullback geometry. In *Advances in Neural Information Processing Systems*, volume 38, pp. 59589–59639. Curran Associates, Inc., 2025. doi: 10.52202/085713-1993. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/5616112a0120c15bf7d47a6bccc21bc3-Abstract-Conference.html.

Ian C. Covert, Balu Krishnan, Imad Najm, Jiening Zhan, Matthew Shore, John Hixson, and Ming Jack Po. Temporal graph convolutional networks for automatic seizure detection. In Finale Doshi-Velez, Jim Fackler, Ken Jung, David Kale, Rajesh Ranganath, Byron Wallace, and Jenna Wiens (eds.), *Proceedings of the 4th Machine Learning for Healthcare Conference*, volume 106 of *Proceedings of Machine Learning Research*, pp. 160–180. PMLR, 2019. URL https://proceedings.mlr.press/v106/covert19a.html.

Floor Eijkelboom, Grigory Bartosh, Christian A. Naesseth, Max Welling, and Jan-Willem van de Meent. Variational flow matching for graph generation. In *Advances in Neural Information Processing Systems*, volume 37, pp. 11735–11764. Curran Associates, Inc., 2024. doi: 10.52202/079017-0374. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/15b780350b302a1bf9a3bd273f5c15a4-Abstract-Conference.html.

Kaiqiao Han, Yi Yang, Zijie Huang, Xuan Kan, Ying Guo, Yang Yang, Lifang He, Liang Zhan, Yizhou Sun, Wei Wang, and Carl Yang. BrainODE: Dynamic brain signal analysis via graph-aided neural ordinary differential equations. In *2024 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI)*, pp. 1–8. IEEE, 2024. doi: 10.1109/BHI62660.2024.10913768. URL https://doi.org/10.1109/BHI62660.2024.10913768.

Noam Issachar, Mohammad Salama, Raanan Fattal, and Sagie Benaim. Designing a conditional prior distribution for flow-based generative models. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL https://openreview.net/forum?id=Teh9Bq4giF.

Haohui Jia, Zheng Chen, Lingwei Zhu, Rikuto Kotoge, Jathurshan Pradeepkumar, Yasuko Matsubara, Jimeng Sun, Yasushi Sakurai, and Takashi Matsubara. ODEBrain: Continuous-time EEG graph for modeling dynamic brain networks. In C. Vondrick, B. Hariharan, C. Raffel, L. Pinto, D. Yang, and A. Faust (eds.), *International Conference on Learning Representations*, volume 2026, pp. 106978–106996, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/aebec8058f23a445353c83ede0e1ec48-Abstract-Conference.html.

Keyue Jiang, Jiahao Cui, Xiaowen Dong, and Laura Toni. Bures-wasserstein flow matching for graph generation. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/e3b1291b7529063172d927588d2b03a1-Abstract-Conference.html.

Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs via the system of stochastic differential equations. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 10362–10383. PMLR, 2022. URL https://proceedings.mlr.press/v162/jo22a.html.

Junwan Kim, Jiho Park, Seonghu Jeon, and Seungryong Kim. Better source, better flow: Learning condition-dependent source distribution for flow matching, 2026. URL https://arxiv.org/abs/2602.05951. Preprint.

Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Rikuto Kotoge, Zheng Chen, Tasuku Kimura, Yasuko Matsubara, Takufumi Yanagisawa, Haruhiko Kishima, and Yasushi Sakurai. EvoBrain: Dynamic multi-channel EEG graph modeling for time-evolving brain networks. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, pp. 144514–144543. Curran Associates, Inc., 2025. doi: 10.52202/085713-4837. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/d51ce6040fe3dafd260411593f05a1fa-Abstract-Conference.html.

Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*, 2017.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

Xuanqing Liu, Tesi Xiao, Si Si, Qin Cao, Sanjiv Kumar, and Cho-Jui Hsieh. Neural sde: Stabilizing neural ode networks with stochastic noise. *arXiv preprint arXiv:1906.02355*, 2019.

Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, and Charles Leiserson. Evolvegcn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 5363–5370, 2020.

Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, volume 32, pp. 3942–3951. AAAI Press, 2018. doi: 10.1609/aaai.v32i1.11671. URL https://ojs.aaai.org/index.php/AAAI/article/view/11671.

Michael Poli, Stefano Massaroli, Junyoung Park, Atsushi Yamashita, Hajime Asama, and Jinkyoo Park. Graph neural ordinary differential equations, 2019. URL https://arxiv.org/abs/1911.07532.

Jathurshan Pradeepkumar, Xihao Piao, Zheng Chen, and Jimeng Sun. Tokenizing single-channel EEG with time-frequency motif learning. In *The Fourteenth International Conference on Learning Representations*, 2026.

Yulia Rubanova, Ricky T. Q. Chen, and David K. Duvenaud. Latent ordinary differential equations for irregularly-sampled time series. In Hanna Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché Buc, Emily Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32, pp. 5320–5330. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/hash/42a6845a557bef704ad8ac9cb4461d43-Abstract.html.

Vinit Shah, Eva von Weltin, Silvia Lopez, James Riley McHugh, Lillian Veloso, Meysam Golmohammadi, Iyad Obeid, and Joseph Picone. The Temple University Hospital seizure detection corpus. *Frontiers in Neuroinformatics*, 12:83, 2018. doi: 10.3389/fninf.2018.00083. URL https://www.frontiersin.org/journals/neuroinformatics/articles/10.3389/fninf.2018.00083/full.

Siyi Tang, Jared A. Dunnmon, Khaled Kamal Saab, Xuan Zhang, Qianying Huang, Florian Dubost, Daniel L. Rubin, and Christopher Lee-Messer. Self-supervised graph neural networks for improved electroencephalographic seizure analysis. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=k9bx1EfHl_.

Siyi Tang, Jared A. Dunnmon, Liangqiong Qu, Khaled K. Saab, Tina Baykaner, Christopher Lee-Messer, and Daniel L. Rubin. Modeling multivariate biosignals with graph neural networks and structured state space models. In Bobak J. Mortazavi, Tasmie Sarker, Andrew Beam, and Joyce C. Ho (eds.), *Proceedings of the Conference on Health, Inference, and Learning*, volume 209 of *Proceedings of Machine Learning Research*, pp. 50–71. PMLR, 2023. URL https://proceedings.mlr.press/v209/tang23a.html.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=CD9Snc73AW.

Clément Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. DiGress: Discrete denoising diffusion for graph generation. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=UaAD-Nu86WX.

Kacper Wyrwal, İsmail İlkan Ceylan, and Alexander Tong. Topological flow matching. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/238e6167319ad5da33c7b38594a1edb1-Abstract-Conference.html.

Chaoqi Yang, M. Brandon Westover, and Jimeng Sun. BIOT: Biosignal transformer for cross-data learning in the wild. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 78240–78260. Curran Associates, Inc., 2023. doi: 10.52202/075280-3420. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/f6b30f3e2dd9cb53bbf2024402d02295-Abstract-Conference.html.

Ling Yang and Shenda Hong. Unsupervised time-series representation learning with iterative bilinear temporal-spectral fusion. In *International conference on machine learning*, pp. 25038–25054. PMLR, 2022.

Zepeng Zhang, Aref Einizade, Jhony H. Giraldo, and Olga Fink. Spatiotemporal imputation with graph-informed flow matching, 2026. URL https://arxiv.org/abs/2606.06682. Accepted at the 43rd International Conference on Machine Learning (ICML 2026); proceedings metadata not yet available.