

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 NEUROCYCLE: PHYSIOLOGICALLY CONSTRAINED CYCLING FOR GENERATING NEURAL INFORMATION- RICH fMRI FROM EEG

Anonymous authors

Paper under double-blind review

ABSTRACT

Electroencephalography (EEG) provides millisecond-level temporal resolution but suffers from poor spatial precision, whereas functional magnetic resonance imaging (fMRI) offers fine-grained spatial detail at the expense of cost and latency. Leveraging their complementarity, an emerging direction is to synthesize fMRI from EEG, enriching EEG with spatial information while retaining its efficiency. However, existing EEG to fMRI generation methods often lack designs to preserve information completeness and neglect neurophysiological priors, leading to reconstructions that may appear plausible but fail to ensure neuroscientific validity. We introduce NeuroCycle, a cyclic EEG–fMRI generation framework that enforces information completeness and neuroscientific plausibility. It incorporates two neurophysiological priors: (i) a Cross-Modal ROI-wise Structural Module that aligns fMRI embeddings with EEG-derived correlation patterns to preserve regional organization, and (ii) an R2E Physiological Connectivity Guidance Module that supervises covariance matrices via Riemannian-to-Euclidean mapping to maintain functional connectivity. The bidirectional cycle (EEG→fMRI→EEG) further enforces information completeness and cross-modal alignment, ensuring that synthesized fMRI retains key neural information. Experiments on NODDI and Oddball datasets show consistent improvements over state-of-the-art baselines, producing sharper voxel-wise fMRI with richer neural information, preserved connectivity, and stronger cross-modal alignment.

1 INTRODUCTION

Electroencephalography (EEG) is one of the most widely used modalities for non-invasive brain signal acquisition due to its high temporal resolution, low cost, and accessibility Gevins et al. (1995). It has been extensively applied in clinical monitoring and neuromodulation. However, EEG suffers from inherently limited spatial precision, which restricts its effectiveness in tasks requiring fine-grained localization, such as precise diagnosis and prediction of neurological disorders Srinivasan (1999). In contrast, functional magnetic resonance imaging (fMRI) provides stable observations of brain states with high spatial resolution, making it particularly powerful for brain region localization and disease prediction. Yet, the acquisition of fMRI is expensive, relies on bulky equipment, and is constrained by low temporal resolution, preventing its deployment in real-time monitoring and large-scale clinical scenarios Menon & Goodyear (2001). In this context, EEG and fMRI provide complementary advantages across temporal and spatial dimensions Huster et al. (2012); Mulert et al. (2004). Both originate from the same underlying neural activity: EEG captures neuronal electrical discharges, while fMRI reflects the hemodynamic response. Based on this complementarity, an important research question naturally arises: Can we synthesize spatially fine-grained fMRI representations solely from EEG data, thereby retaining the efficiency and affordability of EEG while overcoming its spatial limitations?

This question has attracted increasing attention, and several pioneering studies follow a common pipeline that encodes EEG signals into embeddings and subsequently generates fMRI representations through transformer- or GAN-based frameworks (e.g., NT-ViT Lanzino et al. (2024), E2fGAN Roos et al. (2025)). However, existing methods face critical limitations (Figure 1(a)). First, existing methods generally lack explicit modeling and supervision of neurophysiological pri-

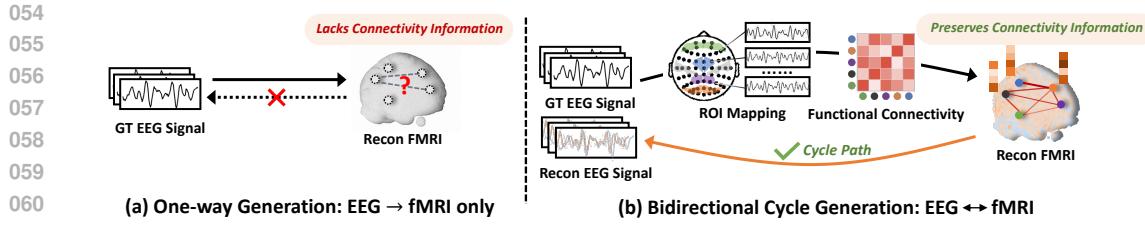


Figure 1: Traditional one-way EEG→fMRI generation (a) overlooks physiological priors and loses information completeness, whereas our cyclic framework (b) integrates connectivity priors and cycle consistency to preserve neural information and enhance biological plausibility.

ors, such as brain ROI-level organization and functional connectivity, making it difficult to ensure that the generated fMRI representations preserve complete and physiologically consistent information aligned with EEG data, thereby limiting their neuroscientific interpretability. Second, most approaches adopt a unidirectional EEG → fMRI mapping, neglecting the preservation of complete neurophysiological information. As a result, the synthesized fMRI may appear visually plausible but fail to guarantee that it faithfully retains the essential statistical characteristics embedded in EEG signals, which undermines its clinical utility and downstream applicability. Moreover, current methods fall short in capturing cross-modal consistency, limiting their ability to uncover and exploit the deep physiological coupling between EEG and fMRI.

To address these challenges, we propose a cyclic cross-modal EEG–fMRI generation framework (Figure 1(b)) that integrates neurophysiological priors to ensure both information completeness and neuroscientific plausibility. First, we incorporate neurophysiological priors through two dedicated modules: the Cross-Modal ROI-wise Structural Module, which aligns EEG-derived correlation patterns with reconstructed fMRI ROI embeddings to preserve regional organization, and the R2E Physiological Connectivity Guidance Module, which supervises EEG covariance matrices via Riemannian-to-Euclidean mapping to maintain functional connectivity consistency. Second, unlike unidirectional EEG → fMRI mappings that risk losing essential neural information, we introduce a pathway that regenerates EEG from the synthesized fMRI, thereby forming a cyclic bidirectional framework (EEG → fMRI → EEG) based on flow matching to guarantee that essential neural information is fully preserved. This cyclic constraint strengthens semantic alignment between modalities, ensuring that the generated fMRI captures both plausible spatial patterns and key neural signal statistics. Together, these innovations provide a principled framework for cross-modal brain signal generation, enhancing both neuroscientific interpretability and potential clinical utility.

In summary, our main contributions are threefold: (1) We propose the first cycle-based EEG–fMRI bidirectional generation framework via Flow Matching, which preserves the full and reversible neurophysiological details embedded in fMRI, thereby ensuring information completeness and modality alignment. (2) We incorporate neurophysiological priors through two modules: A Cross-Modal ROI-wise Structural Module that enforces consistency between EEG-derived connectivity patterns and reconstructed fMRI ROI embeddings, preserving brain regional organization; An R2E Physiological Connectivity Guidance Module that aligns EEG-derived and reconstructed EEG covariance matrices by mapping them from the Riemannian manifold to the Euclidean tangent space, enhancing functional connectivity interpretability and indirectly improving the physiological information integrity preserved in the reconstructed fMRI. (3) We validate the framework on NODDI Deligianni et al. (2016; 2014) and Oddball Walz et al. (2013; 2014; 2015) datasets, achieving superior fMRI reconstruction performance in preserving neural information and cross-modal consistency.

2 RELATED WORK

Early EEG-to-fMRI synthesis studies applied convolutional, graph-attentional, transformer, and adversarial designs to directly generate voxel-wise fMRI. CNN-TC Liu & Sajda (2019) used CNN-based transcribing but struggled with long-range dependencies. CNN-TAG Calhas & Henriques (2022) introduced graph attention over electrodes. NT-ViT Lanzino et al. (2024) employed self-attention for temporal–spectral relations. E2fGAN/E2fNet Roos et al. (2025) enhanced robustness under noise and shifts. Despite architectural diversity, these approaches are largely data-driven map-

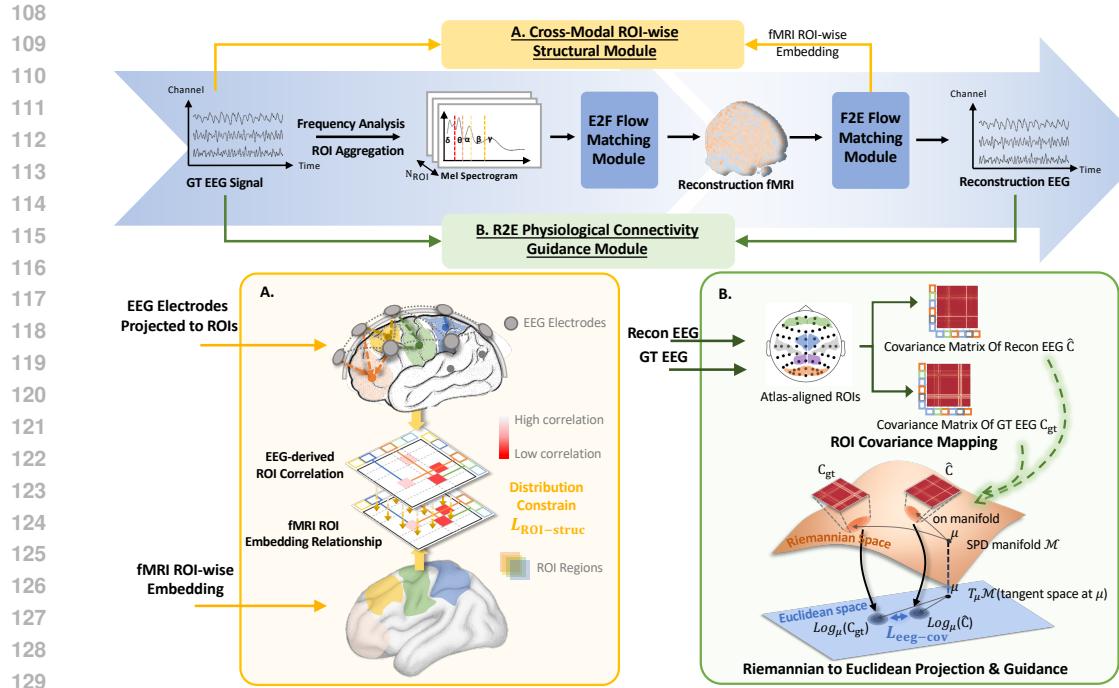


Figure 2: Overview of the proposed cyclic EEG-fMRI generation framework. EEG signals are mapped to fMRI via Flow Matching and back-translated to EEG for cycle consistency. Two guidance modules enhance neurophysiological plausibility: (A) Cross-Modal ROI-wise Structural Module for preserving ROI-level organization across modalities; (B) R2E Physiological Connectivity Guidance for functional connectivity alignment.

pings that do not enforce information integrity or neurophysiological constraints (e.g., connectivity priors), which may limit information completeness and biological plausibility in the synthesized volumes.

Distinct from the above voxel-wise setting, **NeuroBOLT** Li et al. (2024) performs ROI-level time-series synthesis in resting state via multi-dimensional feature mapping that fuses spatiotemporal and multi-scale spectral embeddings. While effective for parcellation-based dynamics and network analysis, its outputs are parcellation signals rather than 3D volumes, thereby discarding intra-ROI heterogeneity and precluding fine-grained spatial localization. In contrast, our **NeuroCycle** introduces a cyclic pathway (EEG \rightarrow fMRI \rightarrow EEG) to guarantee information completeness, while integrating neurophysiological priors to produce biologically constrained voxel-wise reconstructions that move beyond coarse ROI dynamics reconstruction.

3 METHOD

3.1 OVERVIEW

The overall framework of the EEG-fMRI Cyclic Generation Model is illustrated in Figure 2. For the EEG modality, the data of a subject with C electrodes is denoted as $X^{EEG} = (x_1, x_2, \dots, x_C) \in \mathbb{R}^{C \times T}$, where $x_i \in \mathbb{R}^T$ represents the temporal signal of the i -th electrode and T is the length of the time window. For the fMRI modality, the volumetric data at a single time point is represented as $X^{fMRI} \in \mathbb{R}^{V_x \times V_y \times V_z}$, where V_x, V_y, V_z denote the voxel dimensions along the three spatial axes. Our objective is to reconstruct fMRI signals \hat{X}^{fMRI} from EEG inputs X^{EEG} that not only appear realistic, but also faithfully preserve physiological structures and contain complete neural information compared to ground-truth fMRI X^{fMRI} . To achieve this, we design a cyclic mapping of EEG \rightarrow fMRI \rightarrow EEG. In this cycle, the reconstructed EEG \hat{X}^{EEG} is compared against the ground-truth EEG X^{EEG} to provide additional supervision. This cyclic constraint ensures that the

generated fMRI retains sufficient information to regenerate the original EEG, thereby guaranteeing both the information integrity of the reconstructed fMRI and the cross-modal consistency between EEG and fMRI.

To extract spectral information, the input raw EEG signal X^{EEG} is transformed into a Mel-spectrogram representation, following the same preprocessing strategy adopted in Lanzino et al. (2024). This produces $X_{spec}^{EEG} \in \mathbb{R}^{C \times F \times T}$, with F denoting the number of Mel frequency bands. Subsequently, electrodes are projected to atlas-defined cortical regions, and ROI-level spectral features are aggregated as $X_{roi}^{EEG} \in \mathbb{R}^{N_{roi} \times F \times T}$. We then input X_{roi}^{EEG} into the EEG ROI encoder E_{EEG} to obtain ROI-level latent embeddings $Z_{roi}^{EEG} \in \mathbb{R}^{N_{roi} \times d}$, where d is the embedding dimension. These ROI embeddings are concatenated and projected through a projection layer to derive a compact global representation Z_g^{EEG} . The global embedding Z_g^{EEG} is then fed into the EEG-to-fMRI (E2F) flow matching module $F_{E \rightarrow F}$ to predict the corresponding fMRI global embedding Z_g^{fMRI} . This embedding is further decoded by decoder D_{fMRI} to reconstruct the volumetric fMRI data \hat{X}^{fMRI} . Next, the reconstructed fMRI \hat{X}^{fMRI} is encoded by the fMRI encoder E_{fMRI} into its ROI-level embedding $\hat{Z}_{roi}^{fMRI} \in \mathbb{R}^{N_{roi} \times d}$. This embedding is then fed into the proposed **Cross-Modal ROI-wise Structural Module**, which enforces that the fMRI ROI embeddings preserve the connectivity relationship inferred from EEG, thereby maintaining functional connectivity patterns that reflect structural organization and ensuring consistency across modalities. Then, following the procedure described above, a compact global representation obtained \hat{Z}_g^{fMRI} is mapped to the EEG embedding space through the fMRI-to-EEG (F2E) flow matching module $F_{F \rightarrow E}$. The resulting representation is then passed into the EEG decoder D_{EEG} to generate the reconstructed raw EEG signal $\hat{X}^{EEG} \in \mathbb{R}^{C \times T}$. The reconstructed EEG \hat{X}^{EEG} together with the ground-truth EEG X^{EEG} are fed into the proposed **R2E Physiological Connectivity Guidance Module**, which maps EEG-derived ROI-level functional connectivity matrices from the Riemannian manifold to the Euclidean tangent space and enforces agreement between the reconstructed and ground-truth EEG connectivity by minimizing their tangent-space discrepancy, thereby ensuring that the generated EEG preserves physiological connectivity relationships and, in turn, enhancing the information integrity of the generated fMRI.

3.2 ROI ENCODER AND BIDIRECTIONAL CROSS-MODAL FLOW MATCHING MODULE

To bridge the heterogeneous representations of EEG and fMRI, we design ROI-based encoders that project both modalities into a shared latent space while preserving global and regional brain dynamics, followed by bidirectional flow matching modules learning consistent cross-modal transformations. The module design is illustrated in Figure 3.

The raw EEG signal $X^{EEG} \in \mathbb{R}^{C \times T}$, is first transformed into a Mel-spectrogram representation $X_{spec}^{EEG} \in \mathbb{R}^{C \times F \times T}$, to incorporate spectral information. To aggregate electrode-level features into anatomically meaningful regions, we apply a predefined electrode-to-ROI mapping $\mathcal{M}_{EEG \rightarrow ROI}$, where each electrode contributes to a cortical ROI according to a distance-based weighting between the electrode location and the ROI center in the atlas Babiloni et al. (2004). Let $p_c \in \mathbb{R}^3$ denote the spatial coordinate of the c -th electrode, and $q_r \in \mathbb{R}^3$ denote the center coordinate of the r -th ROI. The distance between electrode c and ROI r is $d_{c,r} = \|p_c - q_r\|_2$. We define the unnormalized weight as the inverse distance $w_{c,r} = \frac{1}{d_{c,r} + \varepsilon}$, where ε is a small constant to avoid division by zero. After normalization across all electrodes, the final weight is $\alpha_{c,r} = \frac{w_{c,r}}{\sum_{c'=1}^C w_{c',r}}$. Given the Mel-spectrogram representation of EEG signals X_{spec}^{EEG} , the ROI-level EEG spectral features for ROI r are obtained as a weighted combination of electrode features: $X_{ROI}^{EEG}[r, f] = \sum_{c=1}^C \alpha_{c,r} \frac{1}{T} \sum_{t=1}^T X_{spec}^{EEG}[c, f, t]$, where $f = 1, \dots, F$ indexes Mel frequency bands.

These sequences, together with fMRI ROI representations, are encoded by modality-specific VAE-style encoders. Each encoder E_{EEG} and E_{fMRI} outputs ROI-level Gaussian parameters: $(\mu_{roi}^{EEG}, \log \sigma_{roi}^{2,EEG})$, $(\mu_{roi}^{fMRI}, \log \sigma_{roi}^{2,fMRI})$. ROI-level latent variables are sampled via the reparameterization trick: $z_{roi}^m = \mu_{roi}^m + \sigma_{roi}^m \odot \epsilon$, $\epsilon \sim \mathcal{N}(0, I)$, where $m \in \{\text{EEG}, \text{fMRI}\}$ denotes the modality.

To obtain a compact whole-brain representation, all ROI embeddings are first concatenated and then projected through a transformation function f_{proj} . This projected feature vector parameterizes

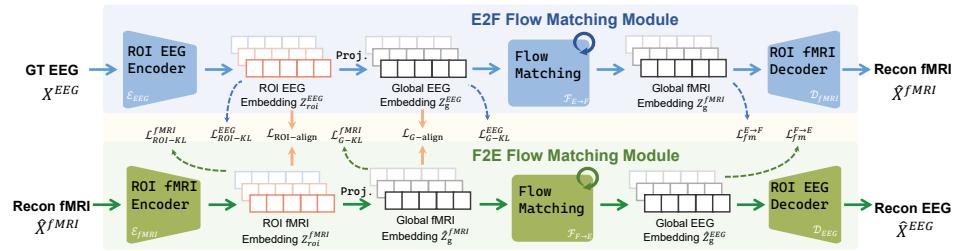


Figure 3: Illustration of the ROI Encoder and Bidirectional E2F/F2E Flowing Matching Module

a global Gaussian distribution: $(\mu_g^m, \log \sigma_g^{2,m}) = f_{\text{proj}}(\text{Concat}(z_{\text{roi}}^m))$, from which the global latent variable is sampled via the reparameterization trick: $z_g^m = \mu_g^m + \sigma_g^m \odot \epsilon$, $\epsilon \sim \mathcal{N}(0, I)$, where $m \in \{\text{EEG, fMRI}\}$ denotes the modality.

On top of these latent encodings, two flow matching modules are introduced: E2F flow matching module $F_{E \rightarrow F}$ maps EEG latent variables into the fMRI space, while F2E flow matching module $F_{F \rightarrow E}$ maps fMRI latents into the EEG space. Given a source latent z_s and a target latent z_t , the flow predicts $\hat{z}_t = F(z_s)$ and is trained with the mean-squared error loss $\mathcal{L}_{\text{fm}}(z_s, z_t) = \|F(z_s) - z_t\|^2$. The bidirectional flow losses are:

$$\mathcal{L}_{\text{fm}}^{E \rightarrow F} = \|F_{E \rightarrow F}(z_g^{\text{EEG}}) - z_g^{\text{fMRI}}\|_2^2, \quad \mathcal{L}_{\text{fm}}^{F \rightarrow E} = \|F_{F \rightarrow E}(z_g^{\text{fMRI}}) - z_g^{\text{EEG}}\|_2^2. \quad (1)$$

To regularize the latent space and stabilize training, KL divergence regularization is applied to both global and ROI-level distributions:

$$\mathcal{L}_{\text{KL}}^m = \mathcal{L}_{\text{G-KL}}^m + \mathcal{L}_{\text{ROI-KL}}^m = D_{\text{KL}}(q(z_g^m | X^m) \| \mathcal{N}(0, I)) + D_{\text{KL}}(q(z_{\text{roi}}^m | X^m) \| \mathcal{N}(0, I)). \quad (2)$$

The overall KL loss is then: $\mathcal{L}_{\text{KL}} = \mathcal{L}_{\text{KL}}^{\text{EEG}} + \mathcal{L}_{\text{KL}}^{\text{fMRI}}$.

Moreover, to explicitly align regional representations across modalities, we introduce an ROI-level correspondence loss $\mathcal{L}_{\text{ROI-align}}$, which enforces regional consistency between EEG and fMRI ROI embeddings. This loss minimizes the distance between paired EEG-fMRI ROI embeddings while contrasting against embeddings from other ROIs based on a contrastive InfoNCE objective. Similarly, global alignment between EEG and fMRI pairs latents is also encouraged. The two losses are defined as:

$$\mathcal{L}_{\text{ROI-align}} = -\frac{1}{N_{\text{ROI}}} \sum_{r=1}^{N_{\text{ROI}}} \log \frac{\exp(\text{sim}(z_{\text{roi},r}^{\text{EEG}}, z_{\text{roi},r}^{\text{fMRI}}) / \tau)}{\sum_{r'} \exp(\text{sim}(z_{\text{roi},r}^{\text{EEG}}, z_{\text{roi},r'}^{\text{fMRI}}) / \tau)}, \quad \mathcal{L}_{\text{align}} = -\log \frac{\exp(\text{sim}(z_g^{\text{EEG}}, z_g^{\text{fMRI}}) / \tau)}{\sum_{z'} \exp(\text{sim}(z_g^{\text{EEG}}, z') / \tau)}. \quad (3)$$

where $z_{\text{roi},r}^{\text{EEG}}$ and $z_{\text{roi},r}^{\text{fMRI}}$ denote the d -dimensional embeddings of ROI r , $\text{sim}(\cdot, \cdot)$ is the cosine similarity, and τ is a temperature parameter.

Finally, the modality-specific decoders take the EEG and fMRI latents as input to reconstruct the original signals, $\hat{X}^{\text{fMRI}} = D_{\text{fMRI}}(F_{E \rightarrow F}(z_g^{\text{EEG}}))$, $\hat{X}^{\text{EEG}} = D_{\text{EEG}}(F_{F \rightarrow E}(z_g^{\text{fMRI}}))$, with reconstruction losses is defined as the mean squared error (MSE) between the predicted and ground-truth signals:

$$\mathcal{L}_{\text{recon}}^m = \frac{1}{|X^m|} \|\hat{X}^m - X^m\|_2^2, \quad m \in \{\text{EEG, fMRI}\}. \quad (4)$$

The overall training objective of this module is :

$$\mathcal{L} = \lambda_{\text{fm}}(\mathcal{L}_{\text{fm}}^{E \rightarrow F} + \mathcal{L}_{\text{fm}}^{F \rightarrow E}) + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}} + \lambda_{\text{align}} \mathcal{L}_{\text{align}} + \lambda_{\text{ROI-align}} \mathcal{L}_{\text{ROI-align}} + \lambda_{\text{recon}}(\mathcal{L}_{\text{recon}}^{\text{fMRI}} + \mathcal{L}_{\text{recon}}^{\text{EEG}}). \quad (5)$$

This formulation links EEG and fMRI latent spaces through reversible flow mappings, while reconstruction and regularization terms preserve both global and regional fidelity, thereby enhancing cross-modal consistency.

3.3 CROSS-MODAL ROI-WISE STRUCTURAL MODULE

To ensure that the reconstructed fMRI preserves realistic regional structural organization, we introduce an ROI-level structural constraint for fMRI data in Cross-Modal ROI-wise Structural Module,

denoted as $\mathcal{L}_{\text{ROI-struct}}$. This constraint leverages functional connectivity (FC) estimated from ground-truth EEG to regulate the pairwise relationships among fMRI ROI embeddings. Since FC reflects the strength of statistical dependencies between brain regions Rojas et al. (2018); Bijsterbosch et al. (2018), ROI embeddings with strong EEG-derived connectivity are expected to exhibit stronger associations and thus should lie closer in the fMRI embedding space, whereas those with weak connectivity should remain farther apart. In this way, the constraint enforces cross-modal consistency while grounding the embedding geometry in neurophysiological plausibility.

Formally, given the ROI-level EEG signals obtained from the electrode-to-ROI mapping described in Section 3.2, we compute the functional connectivity matrix $W \in \mathbb{R}^{N_{\text{ROI}} \times N_{\text{ROI}}}$ by first estimating the covariance across ROI time series and then normalizing each entry to obtain the Pearson correlation between ROI r and ROI r' . The resulting normalized connectivity matrix is defined as: $\widetilde{W}_{r,r'} = \frac{W_{r,r'}}{\sum_{i < j} W_{i,j}}$, which serves as a structural prior over brain ROI interactions.

For the fMRI side, let $Z_{\text{ROI}}^{\text{fMRI}} = \{z_r^{\text{fMRI}}\}_{r=1}^{N_{\text{ROI}}}$, $z_r^{\text{fMRI}} \in \mathbb{R}^d$ denote the ROI-level embeddings produced by the fMRI encoder. We then define the pairwise squared distance between ROI embeddings: $D_{r,r'} = \|z_r^{\text{fMRI}} - z_{r'}^{\text{fMRI}}\|_2^2$. The ROI-wise structural constraint is formulated as the weighted average of these distances:

$$\mathcal{L}_{\text{ROI-struct}} = \frac{1}{N_{\text{pairs}}} \sum_{r < r'} \widetilde{W}_{r,r'} D_{r,r'}, \quad (6)$$

where $N_{\text{pairs}} = N_{\text{ROI}}(N_{\text{ROI}} - 1)/2$ is the number of ROI pairs.

By constraining fMRI ROI-wise relationships based on EEG-derived connectivity, this module transfers brain structural information across modalities and encourages fMRI reconstructions that are more consistent with underlying neurophysiology.

3.4 R2E PHYSIOLOGICAL CONNECTIVITY GUIDANCE MODULE

To ensure that the generated fMRI retains information sufficient for reconstructing neurophysiologically plausible EEG, we constrain the *EEG covariance connectivity* of the EEG reconstructed from generated fMRI to match that of ground-truth EEG on the SPD manifold S_d^{++} . Let $\Sigma_{\text{GT}}^{\text{EEG}}$ and $\Sigma_{\text{Gen}}^{\text{EEG}}$ denote the ROI-level EEG *covariance* matrices computed from ground-truth EEG and from the EEG reconstructed from generated fMRI, respectively (both estimated with an OAS-type shrinkage estimator; the electrode-to-ROI operator M is given in Sec. 3.3). A technical difficulty is that such covariance-based connectivity matrices live on the non-Euclidean manifold of symmetric positive-definite (SPD) matrices, where naïve Euclidean losses are geometrically inappropriate and often numerically unstable. Following the recent Riemannian-to-Euclidean (R2E) strategy Mellot et al. (2024); Collas et al. (2025), we enforce positive-definiteness and stability via symmetrization, shrinkage, and eigenvalue clamping:

$$\Sigma \leftarrow \frac{1}{2}(\Sigma + \Sigma^\top), \quad \Sigma \leftarrow (1 - \lambda)\Sigma + \lambda I, \quad \lambda \in (0, 1), \quad \Lambda \leftarrow \max(\Lambda, \epsilon).$$

We then adopt the Log–Euclidean (LE) framework, which diffeomorphically maps S_d^{++} to the vector space of symmetric matrices $\text{Sym}(d)$ via the matrix logarithm; using the Frobenius norm in the log-domain yields the LE distance: $d_{\text{LE}}(\Sigma_1, \Sigma_2) = \|\log(\Sigma_1) - \log(\Sigma_2)\|_F$. Our EEG covariance consistency loss is therefore:

$$\mathcal{L}_{\text{eeg-cov}} = \|\log(\Sigma_{\text{Gen}}^{\text{EEG}}) - \log(\Sigma_{\text{GT}}^{\text{EEG}})\|_F^2. \quad (7)$$

This geometry-aware loss is applied end-to-end through the fMRI \rightarrow EEG path, thereby indirectly supervising the generated fMRI to preserve essential covariance connectivity structure.

3.5 TRAINING STRATEGY

To achieve stable cross-modal generation, we adopt a two-stage training strategy. In the first stage, we independently train the EEG \rightarrow fMRI and fMRI \rightarrow EEG pathways with reconstruction and alignment losses, ensuring modality-faithful encoders/decoders ($E_{\text{EEG}}, E_{\text{fMRI}}, D_{\text{EEG}}, D_{\text{fMRI}}$) and stable bidirectional flow modules ($F_{E \rightarrow F}, F_{F \rightarrow E}$) without cycle constraints and $\mathcal{L}_{\text{ROI-struct}}$ and $\mathcal{L}_{\text{eeg-cov}}$.

In the second stage, we introduce cycle pathway (EEG \rightarrow fMRI \rightarrow EEG) to enforce reversibility of cross-modal mappings. By requiring the reconstructed EEG signals to remain consistent with their

324 ground truth, the cycle loss preserves shared neural information and indirectly enhances the integrity
 325 of reconstructed fMRI. The overall training objective at stage two combines all losses:
 326

$$327 \mathcal{L}_{\text{total}} = \lambda_{\text{fm}} (\mathcal{L}_{\text{fm}}^{E \rightarrow F} + \mathcal{L}_{\text{fm}}^{F \rightarrow E}) + \lambda_{\text{KL}} (\mathcal{L}_{\text{KL}} + \mathcal{L}_{\text{ROI-KL}}) + \lambda_{\text{align}} (\mathcal{L}_{\text{align}} + \mathcal{L}_{\text{ROI-align}}) + \lambda_{\text{recon}} (\mathcal{L}_{\text{recon}}^{\text{fMRI}} + \mathcal{L}_{\text{recon}}^{\text{EEG}}) \\ 328 + \lambda_{\text{eeg-cov}} \mathcal{L}_{\text{eeg-cov}} + \lambda_{\text{ROI-struct}} \mathcal{L}_{\text{ROI-struct}}. \quad (8)$$

329
 330 Here, \mathcal{L}_{fm} denotes the flow matching loss, \mathcal{L}_{KL} and $\mathcal{L}_{\text{ROI-KL}}$ are global and ROI-level KL loss, $\mathcal{L}_{\text{align}}$
 331 and $\mathcal{L}_{\text{ROI-align}}$ enforce cross-modal alignment at global and ROI levels, $\mathcal{L}_{\text{recon}}$ terms regularize the
 332 fidelity of fMRI/EEG reconstructions, $\mathcal{L}_{\text{eeg-cov}}$ enforces EEG functional connectivity consistency in
 333 the Riemannian space, $\mathcal{L}_{\text{ROI-struct}}$ transfers EEG-derived structural priors to fMRI ROI embeddings.
 334 The coefficients λ balance the contributions of different objectives.
 335

336 This staged training strategy first stabilizes the cross-modal mappings by learning modality-
 337 preserving encoders, decoders, and flow modules, and then progressively incorporates structural
 338 priors and cycle-consistency constraints, thereby enabling the generation of fMRI that are not only
 339 neurobiologically faithful, but also preserve more complete neural information from EEG.
 340

4 EXPERIMENT

4.1 SETTINGS

Datasets

346 We conduct experiments on two public EEG–fMRI datasets, NODDI Deligianni et al. (2016; 2014)
 347 and Oddball Walz et al. (2013; 2014; 2015), both providing synchronized multimodal recordings
 348 from healthy adults. The NODDI dataset contains simultaneous resting-state EEG–fMRI recordings
 349 from 15 subjects (originally 17). EEG signals were acquired using a 64-channel cap to capture
 350 millisecond-level neural dynamics, while the corresponding fMRI data were collected with a clinical
 351 1.5T scanner. The Oddball dataset consists of recordings from 17 subjects performing auditory and
 352 visual oddball tasks. Here, EEG signals were recorded across 43 channels with high-frequency
 353 sampling to capture attentional modulations, and fMRI data were obtained using a 3T scanner to
 354 measure task-related hemodynamic responses.
 355

Comparison Methods

356 We compare our method against six representative baselines: CNN-TC Liu & Sajda (2019),
 357 CNN-TAG Calhas & Henriques (2022), NT-VIT Lanzino et al. (2024), E2fNet Roos et al. (2025),
 358 E2fGAN Roos et al. (2025), and Spec2VolCAMU He et al. (2025), providing a diverse set of ap-
 359 proaches for EEG-to-fMRI generation.
 360

Experimental Setups

361 We conduct experiments on two public EEG–fMRI datasets, NODDI and Oddball. NODDI Zhang
 362 et al. (2012) comprises 15 subjects with 63-channel EEG recorded at 512 Hz over 2.16 s (1106 time
 363 points) paired with fMRI volumes of size $30 \times 64 \times 64$, yielding about 4,500 samples. Oddball Calhas
 364 & Henriques (2022) includes 17 subjects with 34-channel EEG (512 Hz, 2.0 s; 1024 time points)
 365 paired with fMRI volumes of size $32 \times 64 \times 64$, totaling around 17,340 samples. Raw EEG signals are
 366 transformed into Mel spectrograms using $n_{\text{mels}}=16$ and hop length of 16 samples, resulting in input
 367 sizes of [63, 16, 70] for NODDI and [34, 16, 65] for Oddball. To establish cross-modal correspon-
 368 dence, all fMRI volumes are parcellated using the Harvard–Oxford cortical atlas Craddock et al.
 369 (2012), and EEG electrodes are aligned to cortical regions via a fixed Biosemi64-to-ROI mapping
 370 matrix.
 371

We report results under two validation settings: fixed split and leave-one-subject-out (LOSO) in
 372 Table 1: (1) Fixed split. Following prior work Lanzino et al. (2024), we adopt the split where
 373 a subset of subjects is used as the test set (NODDI: 2 subjects; Oddball: 4 subjects) and each
 374 experiment is repeated three times with different random seeds, and we report the mean \pm standard
 375 deviation. (2) LOSO. To assess cross-subject generalization, we further conduct leave-one-subject-
 376 out (LOSO) cross-validation. In each fold, one subject is held out for testing while the remaining
 377 subjects are used for training. The final performance is obtained by averaging across all folds and
 reporting the corresponding standard deviations.
 378

Dataset		NODDI			Oddball		
Method	Validation	RMSE ↓	SSIM ↑	PSNR ↑	RMSE ↓	SSIM ↑	PSNR ↑
CNN-TC Liu & Sajda (2019)	Fixed	0.46 ± 0.08	0.449 ± 0.060	\	0.86 ± 0.03	0.189 ± 0.038	\
CNN-TAG Calhas & Henriques (2022)	Fixed	0.40 ± 0.02	0.462 ± 0.020	\	0.70 ± 0.09	0.200 ± 0.017	\
NT-ViT Lanzino et al. (2024)	Fixed	0.07 ± 0.01	0.602 ± 0.005	23.10 ± 1.24	0.08 ± 0.012	0.651 ± 0.012	21.94 ± 1.217
E2Net Roos et al. (2025)	Fixed	0.17 ± 0.003	0.617 ± 0.003	15.39 ± 0.153	0.12 ± 0.003	0.64 ± 0.042	18.42 ± 0.220
E2iGAN Roos et al. (2025)	Fixed	0.15 ± 0.021	0.549 ± 0.084	16.48 ± 0.963	0.09 ± 0.01	0.575 ± 0.056	20.92 ± 0.972
Spec2VolCAMU He et al. (2025)	Fixed	0.096 ± 0.01	0.686 ± 0.051	20.35 ± 0.836	0.09 ± 0.03	0.712 ± 0.083	21.21 ± 3.00
NeuroCycle (ours)	Fixed	0.069 ± 0.004	0.716 ± 0.015	23.22 ± 0.473	0.060 ± 0.008	0.896 ± 0.038	24.43 ± 0.934
NT-ViT Lanzino et al. (2024)	LOSO	0.080 ± 0.01	0.581 ± 0.048	21.94 ± 1.056	0.070 ± 0.01	0.633 ± 0.064	23.10 ± 1.351
E2Net Roos et al. (2025)	LOSO	0.102 ± 0.02	0.609 ± 0.062	19.82 ± 1.451	0.13 ± 0.03	0.626 ± 0.053	17.73 ± 1.825
Spec2VolCAMU He et al. (2025)	LOSO	0.088 ± 0.02	0.698 ± 0.068	21.11 ± 1.956	0.074 ± 0.01	0.742 ± 0.052	22.61 ± 1.373
NeuroCycle (ours)	LOSO	0.065 ± 0.009	0.715 ± 0.046	23.74 ± 1.31	0.061 ± 0.01	0.894 ± 0.053	24.29 ± 1.182

Table 1: Quantitative comparison on NODDI and Oddball datasets using RMSE, SSIM, and PSNR. Bold numbers indicate the best performance, and underlined numbers denote the second best.

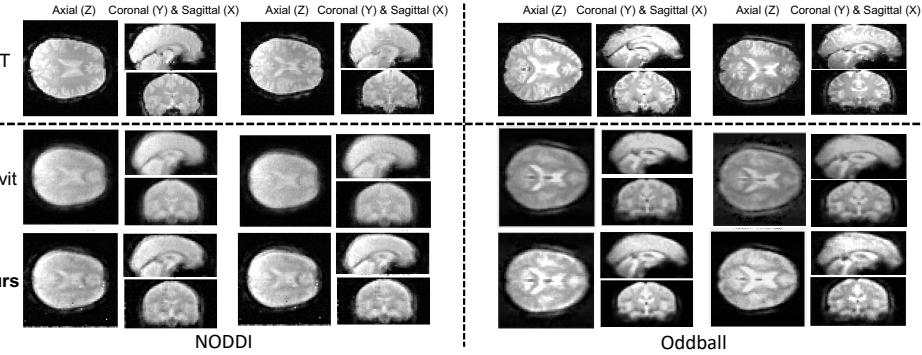


Figure 4: Three-view slices (axial, coronal, sagittal) visualization of reconstructed fMRI on NODDI and Oddball datasets.

We train the model in two stages. In the first stage, the EEG→fMRI and fMRI→EEG pathways are optimized independently for 50 epochs using Adam ($lr = 0.0001$, batch size = 64). In the second stage, we introduce cycle consistency and connectivity constraints, and continue training for 150 epochs with the same learning rate 0.0001. All experiments are implemented in PyTorch and run on a single NVIDIA RTX 4090 GPU.

Training and Evaluation Metrics

To assess the quality of the reconstructed fMRI volumes, we employ three standard image-level evaluation metrics: Root Mean Squared Error (RMSE), Structural Similarity Index (SSIM), and Peak Signal-to-Noise Ratio (PSNR). All metrics are computed in the voxel space on a per-sample basis and averaged over the test set.

4.2 RESULTS AND ANALYSIS

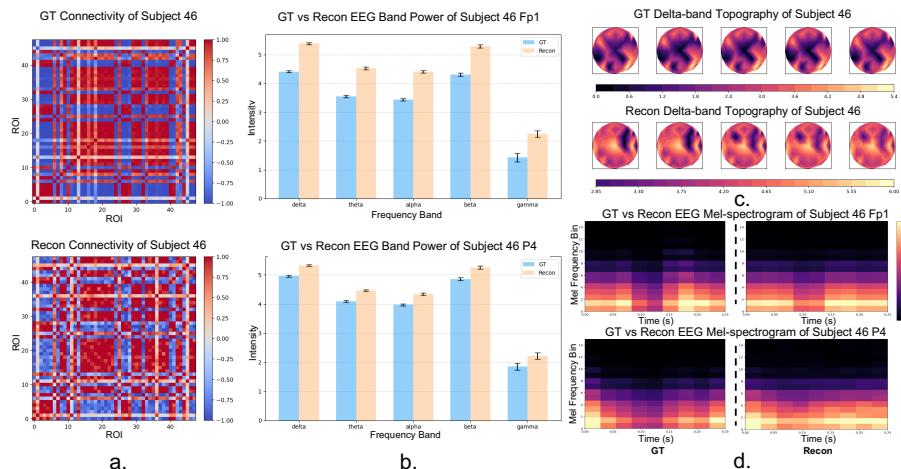
As shown in Table 1, our proposed framework (NeuroCycle) achieves the best performance on both NODDI and Oddball datasets across all metrics (RMSE, SSIM, and PSNR), consistently outperforming previous baselines. In particular, NeuroCycle reduces reconstruction error while substantially improving structural similarity and signal-to-noise ratio. Furthermore, as shown in the three-view slices (axial, coronal, sagittal) in Figure 4, our method preserves more structural and physiological details, yielding clearer and sharper reconstructions that are closer to the ground truth compared with competing approaches.

4.3 ABLATION STUDY

We design three ablation settings to examine the role of each component: (i) w/o Cycle, removing fMRI→EEG pathway and cycle loss, leaving only a unidirectional EEG→fMRI mapping; (ii) w/o R2E-PCG, excluding the R2E functional connectivity guidance; and (iii) w/o CM-RS, removing

432	Dataset	NODDI			Oddball		
433	Method \ Metric	RMSE \downarrow	SSIM \uparrow	PSNR \uparrow	RMSE \downarrow	SSIM \uparrow	PSNR \uparrow
434	w/o Cycle	0.13	0.51	19.56	0.12	0.769	23.14
435	w/o R2E-PCG	0.097	0.64	20.22	0.074	0.836	24.28
436	w/o CM-RS	0.082	0.68	21.62	0.083	0.818	24.67
437	NeuroCycle (ours)	0.068	0.715	23.31	0.061	0.894	25.28

438
439 Table 2: Ablation study on NODDI and Oddball datasets. Results show that removing any component
440 leads to performance degradation, confirming the complementary contributions of all three.
441



458 Figure 5: Visualization of reconstructed EEG versus ground truth (GT) EEG within the proposed
459 EEG→fMRI→EEG cyclic generation framework. (a) ROI-wise functional connectivity matrices
460 of GT and reconstructed EEG for Subject 46. (b) Band power comparison (δ , θ , α , β , γ) across
461 channels 3 and 7 for Subject 46, showing consistency between GT and reconstructed EEG. (c) Topo-
462 graphic maps of two major frequency bands (beta and delta) for Subject 46, where reconstructed
463 patterns closely resemble GT spatial distributions. (d) Mel-spectrogram comparison.

464
465
466 the cross-modal ROI-wise structural constraint. Table 2 summarizes the overall ablation results by
467 sequentially removing each of the three proposed components. We observe consistent performance
468 drops across fMRI quality metrics, demonstrating that every module contributes to the overall effec-
469 tiveness of our framework.

470 To intuitively demonstrate the effect of the cycle supervision, we examine the reconstructed EEG
471 from four complementary perspectives—structural connectivity, spectral fidelity, spatial topogra-
472 phy, and the Mel-spectrogram Nuwer (1988); Deligianni et al. (2014). Representative results for
473 Subject 46 in the NODDI dataset are visualized in Fig. 5. First, we assessed structural consistency
474 at the connectivity level. As shown in Fig. 5(a), the correlation matrix of the reconstructed EEG
475 preserves ROI-wise connectivity patterns comparable to those of the ground truth, indicating that
476 the cyclic pathway enforces meaningful structural organization. Second, we examined frequency
477 fidelity. As shown in Fig. 5(b), the band power distributions on channels Fp1 and P4 demonstrate
478 that the reconstructed EEG maintains the characteristic spectral profile across δ – γ bands, supporting
479 the preservation of frequency-specific information. Third, we analyzed spatial distributions in terms
480 of topographic activity. As shown in Fig. 5(c), the reconstructed signals recover coherent patterns
481 in the major frequency band: β and δ bands, which align with ground-truth spatial maps. And as
482 shown in Fig. 5(d), the reconstructed EEG Mel-spectrogram preserves the dominant low-frequency
483 components as well as the overall temporal modulation patterns.

484 These analyses together demonstrate that the EEG reconstructed from the generated fMRI preserves
485 the characteristic features of the original EEG across connectivity, spectral, and spatial dimensions.
486 This indicates that cycle supervision indirectly enforces the generated fMRI to retain complete neu-

ral information, thereby providing strong evidence for the effectiveness and rationality of the proposed framework in improving fMRI generation quality.

5 CONCLUSION

In this work, we proposed a cyclic EEG-fMRI generation framework that integrates neurophysiological priors to ensure information completeness and neuroscientific plausibility. The Cross-Modal ROI-wise Structural Module preserves ROI-level organization, and the R2E Physiological Connectivity Guidance Module enforces functional connectivity consistency. The cyclic bidirectional design further guarantees preservation of essential neural information. Experiments on two public datasets demonstrate superior fMRI reconstruction with richer neural information and improved cross-modal consistency, providing a principled approach for interpretable and clinically relevant brain signal generation.

REFERENCES

Claudio Babiloni, Giuliano Binetti, Emanuele Cassetta, Daniele Cerboneschi, Gloria Dal Forno, Claudio Del Percio, Florinda Ferreri, Raffaele Ferri, Bartolo Lanuzza, Carlo Miniussi, et al. Mapping distributed sources of cortical rhythms in mild alzheimer's disease. a multicentric eeg study. *Neuroimage*, 22(1):57–67, 2004.

Janine Diane Bijsterbosch, Mark W Woolrich, Matthew F Glasser, Emma C Robinson, Christian F Beckmann, David C Van Essen, Samuel J Harrison, and Stephen M Smith. The relationship between spatial configuration and functional connectivity of brain regions. *elife*, 7:e32992, 2018.

David Calhas and Rui Henriques. Eeg to fmri synthesis benefits from attentional graphs of electrode relationships. *arXiv preprint arXiv:2203.03481*, 2022.

Antoine Collas, Ce Ju, Nicolas Salvy, and Bertrand Thirion. Riemannian flow matching for brain connectivity matrices via pullback geometry. *arXiv preprint arXiv:2505.18193*, 2025.

R Cameron Craddock, G Andrew James, Paul E Holtzheimer III, Xiaoping P Hu, and Helen S Mayberg. A whole brain fmri atlas generated via spatially constrained spectral clustering. *Human brain mapping*, 33(8):1914–1928, 2012.

Fani Deligianni, Maria Centeno, David W Carmichael, and Jonathan D Clayden. Relating resting-state fmri and eeg whole-brain connectomes across frequency bands. *Frontiers in neuroscience*, 8:258, 2014.

Fani Deligianni, David W Carmichael, Gary H Zhang, Chris A Clark, and Jonathan D Clayden. Noddi and tensor-based microstructural indices as predictors of functional connectivity. *Plos one*, 11(4):e0153404, 2016.

Alan Gevins, Harrison Leong, Michael E Smith, Jian Le, and Robert Du. Mapping cognitive brain function with modern high-resolution electroencephalography. *Trends in neurosciences*, 18(10):429–436, 1995.

Dongyi He, Shiyang Li, Bin Jiang, and He Yan. Spec2volcamu-net: A spectrogram-to-volume model for eeg-to-fmri reconstruction based on multi-directional time-frequency convolutional attention encoder and vision-mamba u-net. *arXiv preprint arXiv:2505.09521*, 2025.

René J Huster, Stefan Debener, Tom Eichele, and Christoph S Herrmann. Methods for simultaneous eeg-fmri: an introductory review. *Journal of Neuroscience*, 32(18):6053–6060, 2012.

Romeo Lanzino, Federico Fontana, Luigi Cinque, Francesco Scarcello, and Atsuto Maki. Neural transcoding vision transformers for eeg-to-fmri synthesis. In *European Conference on Computer Vision*, pp. 53–70. Springer, 2024.

Yamin Li, Ange Lou, Ziyuan Xu, Shengchao Zhang, Shiyu Wang, Dario Englot, Soheil Kolouri, Daniel Moyer, Roza Bayrak, and Catie Chang. Neurobolt: Resting-state eeg-to-fmri synthesis with multi-dimensional feature mapping. *Advances in neural information processing systems*, 37:23378–23405, 2024.

540 Xueqing Liu and Paul Sajda. A convolutional neural network for transcoding simultaneously ac-
 541 quired eeg-fmri data. In *2019 9th International IEEE/EMBS Conference on Neural Engineering*
 542 (*NER*), pp. 477–482. IEEE, 2019.

543 Apolline Mellot, Antoine Collas, Sylvain Chevallier, Alex Gramfort, and Denis A Engemann.
 544 Geodesic optimization for predictive shift adaptation on eeg data. *Advances in Neural Infor-
 545 mation Processing Systems*, 37:32828–32855, 2024.

546 Ravi S Menon and Bradley G Goodyear. Spatial and temporal resolution in fmri. *Functional MRI:
 547 an introduction to methods*, pp. 145–158, 2001.

548 Christoph Mulert, Lorenz Jäger, Robert Schmitt, Patrick Bussfeld, Oliver Pogarell, Hans-Jürgen
 549 Möller, Georg Juckel, and Ulrich Hegerl. Integration of fmri and simultaneous eeg: towards a
 550 comprehensive understanding of localization and time-course of brain activity in target detection.
 551 *Neuroimage*, 22(1):83–94, 2004.

552 Marc R Nuwer. Quantitative eeg: I. techniques and problems of frequency analysis and topographic
 553 mapping. *Journal of Clinical Neurophysiology*, 5(1):1–44, 1988.

554 Gonzalo M Rojas, Carolina Alvarez, Carlos E Montoya, MariA De la Iglesia-Vaya, Jaime E Cis-
 555 ternas, and Marcelo Gálvez. Study of resting-state functional connectivity networks using eeg
 556 electrodes position as seed. *Frontiers in neuroscience*, 12:235, 2018.

557 Kristofer Grover Roos, Atsushi Fukuda, and Quan Huu Cap. From brainwaves to brain scans: A
 558 robust neural network for eeg-to-fmri synthesis. *arXiv preprint arXiv:2502.08025*, 2025.

559 Ramesh Srinivasan. Methods to improve the spatial resolution of eeg. *International journal of
 560 bioelectromagnetism*, 1(1):102–111, 1999.

561 Jennifer M Walz, Robin I Goldman, Michael Carapezza, Jordan Muraskin, Truman R Brown, and
 562 Paul Sajda. Simultaneous eeg-fmri reveals temporal evolution of coupling between supramodal
 563 cortical attention networks and the brainstem. *Journal of Neuroscience*, 33(49):19212–19222,
 564 2013.

565 Jennifer M Walz, Robin I Goldman, Michael Carapezza, Jordan Muraskin, Truman R Brown, and
 566 Paul Sajda. Simultaneous eeg–fmri reveals a temporal cascade of task-related and default-mode
 567 activations during a simple target detection task. *Neuroimage*, 102:229–239, 2014.

568 Jennifer M Walz, Robin I Goldman, Michael Carapezza, Jordan Muraskin, Truman R Brown, and
 569 Paul Sajda. Prestimulus eeg alpha oscillations modulate task-related fmri bold responses to audi-
 570 tory stimuli. *NeuroImage*, 113:153–163, 2015.

571 Hui Zhang, Torben Schneider, Claudia A Wheeler-Kingshott, and Daniel C Alexander. Noddi: prac-
 572 tical in vivo neurite orientation dispersion and density imaging of the human brain. *Neuroimage*,
 573 61(4):1000–1016, 2012.

574 **A APPENDIX**

575 **A.1 LLM USAGE**

576 Large Language Models (LLMs) were used to support the preparation of this manuscript. In par-
 577 ticular, an LLM was employed to assist with refining the paper, focusing on improving readabil-
 578 ity, polishing grammar, and ensuring clarity and conciseness. The model also provided alternative
 579 phrasings and stylistic adjustments to better align with academic writing conventions.

580 The LLM was not involved in research ideation, methodological design, experimental implemen-
 581 tation, or result analysis. All scientific contributions including conceptual development, data process-
 582 ing, model design, and evaluation were carried out exclusively by the authors. The role of the LLM
 583 was limited to linguistic refinement and presentation quality.

584 The authors take full responsibility for the manuscript’s scientific content, including any portions of
 585 text that were edited or polished using the LLM. Care has been taken to ensure that the use of the
 586 LLM adheres to ethical standards and does not introduce plagiarism or scientific misconduct.