

Non-Invasive Reconstruction of Intracranial EEG Across the Deep Temporal Lobe from Scalp EEG based on Conditional Normalizing Flow

Dongyi He^a, Bin Jiang^{a,*}, Kecheng Feng^a, Luyin Zhang^a, Ling Liu^a, Yuxuan Li^a, Yun Zhao^b, He Yan^a

^a*School of Artificial Intelligence, Chongqing University of Technology, Chongqing, 400054, China*

^b*School of Smart Health, Chongqing Polytechnic University of Electronic Technology, Chongqing, 401331, China*

Abstract

Although obtaining deep brain activity from non-invasive scalp electroencephalography (EEG) is crucial for neuroscience and clinical diagnosis, directly generating high-fidelity intracranial electroencephalography (iEEG) signals remains a largely unexplored field, limiting our understanding of deep brain dynamics. Current research primarily focuses on traditional signal processing or source localization methods, which struggle to capture the complex waveforms and random characteristics of iEEG. To address this critical challenge, this paper introduces NeuroFlowNet, a novel cross-modal generative framework that reconstructs iEEG signals from scalp EEG over multiple medial temporal lobe (MTL) subregions covered by depth microelectrodes in a public synchronized EEG-iEEG dataset. NeuroFlowNet is built on Conditional Normalizing Flow (CNF), which models complex conditional probability distributions through reversible transformations and latent-variable sampling. Additionally, the model integrates a multi-scale architecture and self-attention mechanisms to robustly capture fine-grained temporal details and long-range dependencies. Validation on a publicly available synchronized EEG-iEEG dataset under both cross-trial and cross-session within-subject protocols shows that NeuroFlowNet can generate band-limited iEEG signals that closely match the ground truth in the time domain, reproduce low-/mid-frequency spectral characteristics (0.5-50 Hz, including the alpha band 8-13 Hz), and preserve the inter-channel correlation structure associated with MTL functional connectivity within the retained bandwidth. Importantly, the cross-session results remain comparable to the cross-trial results, indicating that the model retains stable reconstruction performance on unseen recording sessions from the same subject. Moreover, compared with representative deterministic regression baselines under the same subject-specific protocol, NeuroFlowNet yields consistently lower errors in the inter-channel correlation matrix, indicating better preservation of band-limited cross-channel dependency structure. A conditional sampling and uncertainty analysis further demonstrates non-degenerate stochastic outputs under repeated latent draws for the same EEG input, with predictive variance strongly associated with reconstruction error across validation segments. These results support the feasibility of non-invasive iEEG reconstruction over the recorded MTL subregions, offering a practical route toward studying deep-brain dynamics from scalp EEG under a subject-specific setting. The code of this study is available in <https://github.com/hdy6438/NeuroFlowNet>

Keywords: Deep brain activity, Intracranial electroencephalography (iEEG), Scalp electroencephalography (EEG), Cross-modal generation, Generative models, Brain signal reconstruction,

1. Introduction

Electroencephalography (EEG) is a high-time-resolution brain activity monitoring technique that can track neuronal activation in real time and dynamically record brain activity. It plays a crucial role in brain-computer interfaces, clinical diagnosis, and cognitive neuroscience research [1, 2, 3, 4, 5]. Among these, scalp electroencephalography (EEG) measures the distribution of potentials on the scalp resulting from the propagation of postsynaptic potentials and action potentials from neuronal clusters in the brain, as detected by electrodes placed on the scalp [6, 4]. It offers the advantages of simplicity and high flexibility, making it widely used in clinical monitoring and diagnosis [7].

*Corresponding author. Email: jb20200132@cqut.edu.cn

However, due to the influence of cranial volume conduction effects during signal propagation, its spatial resolution is relatively low [8, 9, 3, 10, 11]. That is, the measured signals represent the superimposed potentials formed by the activities of numerous sources within the brain on the scalp surface. These signals are characterized by weak potentials, susceptibility to external noise interference, and low signal-to-noise ratios, making it difficult to accurately decode the complex dynamic relationships of deep brain neural activities [12, 13]. Recent studies have further explored neuroscience-informed dynamical EEG representations through spatio-spectral modeling and neural manifold decoding, demonstrating the value of structured representation learning for neural decoding and feature transformation analysis [14, 15]. This limitation restricts its further application in clinical settings. In contrast, intracranial electroencephalography (iEEG) can directly obtain neural activity signals from the cerebral cortex or deep nuclei, featuring high temporal-spatial resolution and high signal-to-noise ratio [16]. It is the gold standard for locating lesions in brain-related diseases and achieving high-precision neural decoding [17, 18]. However, its high surgical costs and the risks of immune reactions and infections following electrode implantation significantly limit its practical application [17].

Given the limited spatial resolution of EEG and the difficulty in obtaining iEEG signals, the brain electroencephalogram inverse problem has emerged as a research hotspot to achieve real-time monitoring and precise localization of neuronal activity [19, 20, 21]. Its core lies in how to infer the spatial distribution and temporal information of cortical and deep neuronal electrical activity from EEG signals recorded on the scalp [19, 6]. By solving the EEG inverse problem, it is possible to map non-invasive EEG signals to neural signals of approximately iEEG quality [22, 23]. This not only helps to “unlock” dynamic information from deep brain regions under non-invasive conditions [24] but also provides new technical support for precise localization of epileptic foci, analysis of brain functional mechanisms, and neural rehabilitation interventions [25].

However, significant differences in discharge patterns and signal characteristics between EEG and iEEG limit the effectiveness of traditional methods [26, 27]. Parameter-based localization methods based on equivalent current dipoles (ECD) are constrained by the difficulty in accurately estimating the number of dipoles [6, 28], while source reconstruction methods based on current distribution are hindered by discrepancies between prior constraints and actual neuronal activity [29], both of which struggle to accurately reconstruct complex neural dynamics [30, 31]. Additionally, traditional machine learning methods rely on manual feature extraction [32, 33, 34] and have limited nonlinear modeling capabilities [32], further restricting their representation capabilities in EEG inversion.

In recent years, deep learning techniques have significantly improved EEG inversion performance through their powerful nonlinear fitting capabilities [35, 36, 37]. For example, convolutional neural networks (CNNs) extract spatial correlations between multiple channels through spatial convolution, enhancing the accuracy of source space imaging [38, 39, 40]. Recurrent neural networks (RNNs) and their variants can model the temporal dynamic characteristics of signals, significantly improving the ability to distinguish between transient and sustained source activity [41, 42]. However, deep learning methods still face numerous challenges. Due to data limitations, model training often relies on large-scale synthetic data generated using finite element or boundary element models [43]. Synthetic data struggles to fully reflect the true complex mapping relationship between EEG and deep neuronal electrical activity, severely limiting the model’s performance in real clinical applications. Additionally, the outputs of this method are primarily continuous neural activity distribution maps, which can locate positions but struggle to capture neural response waveforms, thereby impacting further decoding analysis [44].

Generative models offer a new approach for high-fidelity mapping from EEG signals to iEEG signals, with the core mechanism being the use of generative networks to perform cross-modal conversion of easily obtainable non-invasive EEG signals into iEEG signals. Abdi-Sargezeh et al. were the first to introduce generative adversarial networks (GANs) to address this issue [45], successfully generating high-quality iEEG waveforms while demonstrating advantages in terms of efficiency and model complexity. Building on this, they proposed an EEG-to-EEG conversion model, named VAE-cGAN [46], combining variational autoencoders (VAE) and conditional generative adversarial networks (cGAN), and validated it on a dataset containing real iEEG signals recorded via Foramen Ovale electrodes. By encoding EEG into the latent space and regenerating iEEG signals, they further enhanced the fidelity and signal resolution of the mapped waveforms. The team achieved waveform reconstruction for the first time under conditions where real iEEG was used as a reference, providing a new technical pathway for non-invasive acquisition of high-fidelity iEEG signals. However, this research remains in its preliminary exploratory phase. Due to limitations such as the sampling range of Foramen Ovale channels and inherent defects like mode collapse in GAN models, current reconstruction is only feasible for localized regions of the medial temporal lobe (MTL), and expanding the number of reconstruction

channels may significantly degrade reconstruction quality and effectiveness. Additionally, expanding the number of reconstructed channels may significantly degrade reconstruction quality and effectiveness. Additionally, due to the complexity and randomness of neural activity in the brain, models like GAN and VAE, which use deterministic mappings, often fail to adequately reflect the spatiotemporal diversity and randomness of iEEG signals. In addition to waveform fidelity, a fundamental requirement for iEEG reconstruction is to preserve the multi-channel dependency structure across intracranial contacts, since functional connectivity analyses are typically derived from inter-channel covariance/correlation patterns [47, 48]. In an EEG-to-iEEG setting, the inverse mapping is underconstrained: multiple intracranial configurations can be compatible with similar scalp observations due to volume conduction and source mixing. Deterministic regressors trained with pointwise objectives may therefore converge to “average” solutions that reduce waveform error but attenuate network-level variability and distort the inter-channel correlation matrix.

To address these challenges, this study innovatively proposes a new cross-modal generative framework named NeuroFlowNet, which attempts for the first time to reconstruct iEEG signals across multiple MTL subregions using EEG. To capture the spatio-temporal diversity and randomness of iEEG signals, the model is based on Conditional Normalizing Flow (CNF), directly learning the complex conditional probability distribution $p(X_{iEEG}|X_{EEG})$ from EEG to iEEG. Compared to traditional deterministic mappings, the normalizing flow model maps complex data distributions to simple prior distributions through a series of reversible transformations, enabling repeated conditional sampling from latent Gaussian variables. This provides a direct mechanism for representing residual iEEG variability that is not uniquely determined by scalp EEG. Additionally, the model adopts a multi-scale architecture combined with self-attention mechanisms, further enhancing the fineness and robustness of iEEG waveform reconstruction. NeuroFlowNet has been validated for effectiveness on a publicly available dataset [49] of whole temporal lobe EEG-iEEG signals synchronously collected via Stereoelectroencephalography electrodes. Experimental results demonstrate that the model has advantages in signal fidelity, spectral feature reproduction, network structure recovery, and conditional sample uncertainty behavior, and that these properties remain stable when evaluation is shifted from unseen trials to unseen sessions within the same subject. We also assess computational latency under the final model configuration to examine whether the reversible generation process is compatible with real-time or near-real-time monitoring scenarios. This provides a more precise and reliable new paradigm for non-invasive analysis of deep brain dynamics.

2. Methods

The proposed model, NeuroFlowNet, demonstrated in Figure 1, is designed to generate iEEG signals from non-invasive EEG signals by learning the conditional probability distribution of iEEG given EEG, $p(X_t|X_c)$, where $X_t \in \mathbb{R}^{C_{out} \times L}$ represents the target iEEG signals and $X_c \in \mathbb{R}^{C_{in} \times L}$ represents the conditioning EEG signals, with C_{out} and C_{in} being the number of channels of iEEG and EEG, respectively, and L denoting the temporal length of the signals. NeuroFlowNet is architected as a multi-scale conditional normalizing flow, which transforms the complex distribution of iEEG data, conditioned on EEG data, into a simple, tractable base distribution (typically a standard multivariate Gaussian) via a sequence of invertible and differentiable transformations. The conditioning on EEG signals X_c is achieved through a series of coupling layers, where the transformation is conditioned on the EEG data. The model is trained using a maximum likelihood estimation approach, optimizing the parameters of the transformations to maximize the likelihood of the observed iEEG data given the EEG data.

2.1. Core Invertible Transformation: Conditional Affine Coupling Layer

The core of the NeuroFlowNet architecture is the conditional affine coupling layer, which allows for the transformation of the iEEG data while conditioning on the EEG data. The coupling layer operates by splitting the input data into two parts: one part is transformed while the other part through an identity transformation. Specifically, given an input vector $x \in \mathbb{R}^{C \times L}$ at any stage of the flow, it is first split along the channel dimension into two halves, $x_a \in \mathbb{R}^{C/2 \times L}$ and $x_b \in \mathbb{R}^{C/2 \times L}$. The first half, x_a , is passed through an identity transformation, while the second half, x_b , is transformed using a conditional affine transformation parameterized by a neural network. The transformation can be expressed as:

$$z_b = (x_b + \tau) \odot \exp(\sigma) \quad (1)$$

where $\odot$ denotes element-wise multiplication. The translation parameters $\tau \in \mathbb{R}^{C/2 \times L}$ and the log-scale parameters $\sigma \in \mathbb{R}^{C/2 \times L}$ are produced by a conditional neural network, denoted as g_{cond} , which takes x_a and the conditioning EEG

NeuroFlowNet




Figure 1: Overview of NeuroFlowNet architecture. The model consists of multiple scales, each containing a series of flow steps that transform the iEEG data conditioned on the EEG data. The final latent variable Z is obtained by concatenating the outputs from all scales.

data X_c as inputs, and outputs the parameters τ and σ :

$$\begin{bmatrix} \tau \\ \sigma \end{bmatrix} = g_{cond}(x_a, X_c) \quad (2)$$

The conditional neural network g_{cond} is designed to capture the complex relationships between the iEEG and EEG data, allowing for a more accurate transformation of the iEEG data based on the conditioning EEG data. The g_{cond} network is composed of an initial couple of convolutional layers with kernel size of 3 and 1, respectively, both followed by a ReLU activation function. Specifically, the first convolutional layers with kernel size of 3 captures local temporal features from the input data and projects the input to a higher-dimensional feature space $\mathbb{R}^{C_h \times L}$, and the second convolutional layer with kernel size of 1 serves as a channel-wise linear transformation, further processing the features while maintaining the temporal resolution. The output of the second convolutional layer is then fed into a multi-head self-attention mechanism to capture long-range dependencies in the data, which is crucial for modeling the complex temporal dynamics of iEEG signals.

The multi-head self-attention mechanism allows the model to capture multiple relationships between the input data and the conditioning data by using multiple sets of query, key, and value matrices, each with different learned parameters. The output of the attention mechanism is then passed through a final convolutional layer with kernel size of 3 to produce the final parameters τ and σ . For numerical stability, translation parameters τ is scaled by a hyperbolic tangent function ($\tanh$) to ensure that the values are bounded between -1 and 1, while the log-scale parameters σ are clipped to a range of $[-5, 5]$ to prevent numerical instability during the training process. The log-determinant of the




Figure 2: Flow step composition in NeuroFlowNet. Each flow step consists of an invertible 1×1 convolution followed by a conditional affine coupling layer, which transforms the input data while conditioning on the EEG data X_c .

Jacobian for this affine transformation is given by the sum of the log-scale factors:

$$\log \left| \det \frac{\partial z_b}{\partial x_b} \right| = \sum_{i,j} \sigma_{i,j} \quad (3)$$

The output of the coupling layer is $z = \text{concatenate}(x_a, z_b)$. The inverse transformation is obtained by applying the inverse of the affine transformation, which can be expressed as:

$$x_b = \exp(-\sigma) \odot (z_b - \tau) \quad (4)$$

where the inverse transformation is applied to the transformed data z_b to recover the original data x_b .

The inverse transformation is also computed in a similar manner, where the translation and log-scale parameters are obtained from the same conditional neural network g_{cond} , but with the input data being the transformed data z_a and the conditioning EEG data X_c , i.e.:

$$\begin{bmatrix} \tau \\ \sigma \end{bmatrix} = g_{cond}(z_a, X_c) \quad (5)$$

2.2. Flow Step Composition

As demonstrated in Figure 2, each flow step within NeuroFlowNet combines the conditional affine coupling layer with an invertible channel-mixing operation to enhance the expressive power of the model. Specifically, before the input x is fed into the affine coupling layer, it undergoes an invertible 1×1 convolutional operation, a channel-wise linear transformation, $y = Wx'$, where x' is the input to the flow step and $W \in \mathbb{R}^{c \times c}$ is a learnable weight matrix, which is initialized as an orthonormal matrix (via QR decomposition of a random Gaussian matrix) to ensure invertibility

from the start of training and promote stability. This 1×1 convolution allows for information to be mixed across channels at each temporal location independently. The log-determinant of the Jacobian for this transformation is:

$$\log \left| \det \frac{\partial y}{\partial x'} \right| = L \cdot \log |\det W| \quad (6)$$

The total log-determinant for a single flow step is the sum of the log-determinants from the 1×1 convolution and the subsequent conditional affine coupling layer. The inverse of the flow step involves applying the inverse affine coupling transformation followed by multiplication with W^{-1} .

2.3. Multi-Scale Architecture

The proposed NeuroFlowNet employs a multi-scale architecture to effectively model features and dependencies at varying resolutions within the iEEG signals. The model comprises N_S hierarchical scales. At each scale $k \in \{1, \dots, N_S\}$, the input, target iEEG signal X_t at the first scale or the output from the preceding finer scale ($k - 1$), is processed through a sequence of N_{steps} identical flow steps, as detailed in Section 2.2, with all transformations are conditioned on the EEG signal X_c .

Following the flow steps at scale k , the resulting tensor $y^{(k)} \in \mathbb{R}^{C_k \times L}$ is split along the channel dimension. A subset of channels, $z_{\text{part}}^{(k)} \in \mathbb{R}^{C_k/2 \times L}$, is extracted and designated as part of the final latent variable Z . The remaining channels, $x_{\text{in}}^{(k+1)} \in \mathbb{R}^{C_k/2 \times L}$, are propagated as input to the subsequent, typically coarser, scale ($k + 1$). The number of channels C_k is halved at each successive scale, facilitating a progressive abstraction of features.

This hierarchical processing strategy enables the model to decompose the learning task, capturing fine-grained details at the initial scales and progressively more abstract, global representations at deeper levels. The final output from the coarsest scale, $z_{\text{final}}^{(N_S)} \in \mathbb{R}^{C_{N_S} \times L}$, also contributes to the complete latent representation. The full latent variable Z is obtained by concatenating all extracted components and the final output:

$$Z = \text{concatenate}(z_{\text{part}}^{(1)}, z_{\text{part}}^{(2)}, \dots, z_{\text{part}}^{(N_S-1)}, z_{\text{final}}^{(N_S)}) \in \mathbb{R}^{C_{\text{out}} \times L} \quad (7)$$

where $C_{\text{out}} = \sum_{k=1}^{N_S} C_k/2 + C_{N_S}$ is the total number of channels in the latent variable Z

2.4. Training and Inference

2.4.1. Training

NeuroFlowNet is trained by maximizing the exact log-likelihood of the iEEG data X_t given the EEG context X_c . Algorithm 1 outlines the training procedure. Let f denote the overall transformation from the data space to the latent space, such that $Z = f(X_t, X_c)$. The log-likelihood of the iEEG data can be expressed as:

$$\log p(X_t | X_c) = \log p_Z(Z) + \sum_{i=1}^M \log \left| \det \frac{\partial z_i}{\partial z_{i-1}} \right| \quad (8)$$

where M is the total number of elementary invertible transformations that map the input X_t to the final latent variable Z (as defined in Equation 7), $z_0 = X_t$ and z_i is the output of the i -th transformation acting on its input z_{i-1} , then the final result of these M transformations is Z . The term $\frac{\partial z_i}{\partial z_{i-1}}$ is the Jacobian matrix of this i -th transformation. $p_Z(Z)$ is the log-probability density of Z under a standard multivariate Gaussian distribution $\mathcal{N}(0, I)$, which serves as the base distribution. The sum term $\sum_{i=1}^M \log \left| \det \frac{\partial z_i}{\partial z_{i-1}} \right|$ is the accumulated sum of the log-determinants of these Jacobians for all constituent invertible transformations (1×1 convolutions and affine coupling layers) across all flow steps and scales.

The model parameters are optimized by maximizing this log-likelihood objective:

$$\theta^* = \arg \max_{\theta} \log p(X_t | X_c; \theta) \quad (9)$$

where θ represents the parameters of the model, including the weights of the conditional neural network g_{cond} and the weights of the 1×1 convolutions.

Algorithm 1 Training Phase of NeuroFlowNet (Forward & Parameter Update)

Require: Batch $\{X_c, X_t\}$, model parameters θ , per-scale splits $\{S_k\}_{k=1}^{N_S}$, time length L

Ensure: Updated parameters θ by maximizing $\log p(X_t | X_c)$

```

1: Initialize  $\log\_det\_total \leftarrow \mathbf{0}_B$ ,  $o \leftarrow X_t$ ,  $\mathcal{Z} \leftarrow \emptyset$ 
2: for  $k = 1 \rightarrow N_S$  do
3:    $\log\_det_k \leftarrow \mathbf{0}_B$ 
4:   for  $i = 1 \rightarrow N_{steps}$  do ▷ FlowStep  $i$  in scale  $k$ 
5:      $u \leftarrow \text{1x1conv}_{W_i^{(k)}}(o)$ 
6:      $\Delta_{conv} \leftarrow L \cdot \text{slogdet}(W_i^{(k)})$  ▷ scalar
7:      $[u_a, u_b] \leftarrow \text{split\_half}(u)$ 
8:      $[\tau, \sigma] \leftarrow g_{cond}(u_a, X_c; \theta_{aff})$ 
9:      $\tau \leftarrow 10 \cdot \tanh(\tau)$ ,  $\sigma \leftarrow \text{clamp}(\sigma, -5, 5)$ 
10:     $z_b \leftarrow (u_b + \tau) \odot \exp(\sigma)$ 
11:     $o \leftarrow \text{concat}(u_a, z_b)$ 
12:     $\Delta_{affine} \leftarrow \sum_{c,t} \sigma_{c,t}$  ▷ per-sample vector of size  $B$ 
13:     $\log\_det_k \leftarrow \log\_det_k + (\Delta_{conv} + \Delta_{affine})$ 
14:  end for
15:   $z^{(k)} \leftarrow o[:, 1 : S_k, :]$  ▷ extract latent part
16:   $o \leftarrow o[:, S_k + 1 : C_k, :]$  ▷ remaining channels to next scale
17:  append  $z^{(k)}$  to  $\mathcal{Z}$ 
18:   $\log\_det\_total \leftarrow \log\_det\_total + \log\_det_k$ 
19: end for
20: append final  $o$  to  $\mathcal{Z}$  ▷ final leftover latent
21:  $Z \leftarrow \text{concat}(\mathcal{Z}, \text{channel})$ 
22:  $\log p_Z \leftarrow -\frac{1}{2} \sum_{c,t} (Z^2 + \log(2\pi))$  ▷ per-sample
23:  $\mathcal{L} \leftarrow -\frac{1}{B} \sum_{b=1}^B (\log p_Z^{(b)} + \log\_det\_total^{(b)})$ 
24: Backpropagate  $\mathcal{L}$  and update  $\theta$  (e.g., AdamW step)

```

2.4.2. Inference (iEEG Generation)

To generate iEEG samples $\hat{X}_t$ conditioned on a given EEG signal X_c , the inverse transformation f^{-1} is employed. Algorithm 2 outlines the inference procedure. First, a sample $\hat{Z}$ is drawn from the base Gaussian distribution $p_Z(Z)$. This $\hat{Z}$ is then partitioned into segments $\hat{z}_{part}^{(1)}, \dots, \hat{z}_{part}^{(N_S-1)}, \hat{z}_{final}^{(N_S)}$ corresponding to the multi-scale splitting performed during the forward pass. The generation process starts from the coarsest scale N_S . The segment $\hat{z}_{final}^{(N_S)}$ is transformed by applying the inverse of the flow steps within the N_S -th scale block. For each subsequent finer scale k (iterating from $N_S - 1$ down to 1), the corresponding latent segment $\hat{z}_{part}^{(k)}$ is concatenated with the output from the inverse transformation of the preceding coarser scale ($k + 1$). This combined tensor is then passed through the inverse flow steps of scale k . All inverse transformations are conditioned on the EEG context X_c . The final output of this sequential inverse process, after passing through the inverse of the first scale block, is the generated iEEG sample $\hat{X}_t = f^{-1}(\hat{Z}; X_c)$.

3. Experiments

3.1. Dataset

This study utilizes an electrophysiological dataset recorded from nine epilepsy patients during a verbal working memory task, as introduced by Boran et al. [49], which is publicly available at https://gin.g-node.org/USZ_NCH/Human_MTL_units_scalp_EEG_and_iEEG_verbal_WM. The dataset includes simultaneous EEG recorded using the 10-20 system and iEEG obtained via depth microelectrodes implanted in the medial temporal lobe (MTL). The EEG and iEEG signals are originally sampled at 256 Hz and 4 kHz, respectively, and later resampled to 200 Hz and 2 kHz. It also contains waveforms and spike times of 1526 single units, along with MNI coordinates in MNI152 space and anatomical labels for the intracranial electrodes.

Algorithm 2 Inference / Sampling Phase of NeuroFlowNet (Inverse Mapping)

Require: Conditioning X_c , model parameters θ , per-scale splits $\{S_k\}_{k=1}^{N_S}$, optional $\varepsilon \in \mathbb{R}^{B \times D \times L}$

Ensure: Generated iEEG samples $\hat{X}_t$

```
1: if no  $\varepsilon$  provided then
2:   sample  $\varepsilon \sim \mathcal{N}(0, I)$  ▷ shape matches total latent channels
3: end if
4: split  $\varepsilon$  into parts  $\{\tilde{z}^{(1)}, \dots, \tilde{z}^{(N_S)}, \tilde{z}^{(N_S+1)}\}$  according to  $\{S_k, \dots\}$ 
5:  $o \leftarrow \tilde{z}^{(N_S+1)}$  ▷ start from coarsest latent
6: for  $k = N_S \rightarrow 1$  do ▷ reverse through scales
7:    $o \leftarrow \text{concat}(\tilde{z}^{(k)}, o)$  ▷ reconstruct full channel tensor at scale  $k$ 
8:   for each FlowStep  $i = N_{\text{steps}}, \dots, 1$  (reversed order) do
9:      $[z_a, z_b] \leftarrow \text{split\_half}(o)$ 
10:     $[\tau, \sigma] \leftarrow g_{\text{cond}}(z_a, X_c; \theta_{\text{aff}})$ 
11:     $\tau \leftarrow 10 \cdot \tanh(\tau)$ ,  $\sigma \leftarrow \text{clamp}(\sigma, -5, 5)$ 
12:     $u_b \leftarrow \exp(-\sigma) \odot (z_b - \tau)$  ▷ affine inverse
13:     $u \leftarrow \text{concat}(z_a, u_b)$ 
14:     $o \leftarrow \text{1x1conv}_{W^{-1}}(u)$  ▷ apply inverse 1x1 conv
15:   end for
16: end for
17: Output:  $\hat{X}_t \leftarrow o$ 
```

The nine subjects (five female, four male) are all right-handed and aged between 18 and 56. Their pathologies include hippocampal sclerosis, gliosis, xanthoastrocytoma, brain contusion, and focal cortical dysplasia. Each participant has a varying number and location of implanted depth electrodes within the MTL. Due to significant inter-subject variability in electrode implantation (both in terms of anatomical coverage and number of recorded units), only subjects S1, S6, and S9 were included in this study. These three subjects were selected because (i) they have comprehensive electrode coverage across key MTL subregions of interest, (ii) they exhibit a sufficient number of well-isolated single units for reliable statistical analysis, and (iii) their recording sessions contained minimal artifacts and consistent task performance. Due to clinical considerations and individualized implantation plans, the exact anatomical locations of intracranial electrodes differ across subjects [49]. To illustrate this variability, Figure 3 presents a 3D schematic of depth electrode trajectories in the MTL for three subjects (S1, S6 and S9) used in this study.

Each subject completes between two and seven recording sessions. Table 1 summarizes the characteristics of the subjects, including their personal information, pathologies, implanted scalp and intracranial electrodes, and number of sessions. The EEG signals are recorded using the international 10-20 system, with electrodes placed at standard locations such as Fp1, Fp2, F7, F3, Cz, Pz, O1, and O2. The iEEG signals are recorded from depth microelectrodes targeting multiple MTL subregions: anterior hippocampus (AHL/AHR), amygdala (AL/AR), entorhinal cortex (ECL/ECR), and parahippocampal gyrus (PHL/PHR), as well as lateral rhinal area (LR). Each trajectory consisted of eight microelectrode contacts arranged along a single stereotactic path.

The dataset provides a rich source of electrophysiological data, allowing for the exploration of cross-modal generation of iEEG signals from EEG signals. The iEEG signals are used as the target signals X_t for model training and evaluation, while the EEG signals are used as the conditioning input X_c . The dataset is suitable for training and evaluating the NeuroFlowNet model, as it contains a diverse set of subjects with varying pathologies and electrode placements. The model aims to learn the complex relationships between the non-invasive EEG signals and the invasive iEEG signals, enabling the generation of iEEG signals from EEG data.

3.2. Data Preprocessing and Division

The iEEG and EEG signals are preprocessed to remove artifacts and noise, ensuring high-quality data for model training. The preprocessing steps include wavelet denoising, downsampling, data division, and normalization.

The wavelet denoising is performed using the discrete wavelet transform (DWT) with Symlet wavelets of order 4 (sym4), which helps suppress broadband noise while retaining the main morphological characteristics of the signals.




Figure 3: Medial temporal lobe (MTL) depth electrode locations for subjects used in this study (S1, S6 and S9). Each panel shows coronal, sagittal, and axial views of depth electrode trajectories within MTL subregions. Colors indicate anatomical targets: anterior hippocampus (AHL/AHR, red/light red), amygdala (AL/AR, orange/yellow), entorhinal cortex (ECL/ECR, green/light green), parahippocampal gyrus (PHL/PHR, purple/magenta), and lateral rhinal area (LR, blue). Each trajectory consists of eight electrode contacts positioned along a single stereotactic path.

A three-level decomposition is applied to each channel independently, producing the coefficient set A_3, D_3, D_2, D_1 , where A_3 denotes the level-3 approximation coefficients and D_ℓ denotes the detail coefficients at level ℓ . The approximation coefficients A_3 are retained without modification, and soft-thresholding is applied to all detail sub-bands D_1, D_2, D_3 . The noise level is estimated using the median absolute deviation (MAD) computed over the concatenation of detail coefficients across levels, $\hat{\sigma} = \text{median}(|\text{concat}(D_1, D_2, D_3)|)/0.6745$. A universal threshold is then determined based on the segment length (n) (number of time samples) as $T = \hat{\sigma} \sqrt{2 \ln(n)}$, and each detail coefficient (w) is shrunk via soft thresholding $S_T(w) = \text{sign}(w) \max(|w| - T, 0)$. Finally, the denoised signal is reconstructed by inverse DWT (IDWT) with symmetric boundary extension, and the reconstructed sequence is cropped to the original length to ensure temporal alignment with the original signal. We note that the effect of this denoising procedure on transient waveform morphology was not directly evaluated in the present study. Therefore, the subsequent waveform and spectral evaluations should be interpreted as reconstruction performance relative to the denoised and band-limited target signals, rather than as evidence that all transient morphological features in the original broadband iEEG recordings were preserved.

In this work, to align the temporal grid between modalities for supervised cross-modal learning, we further down-sampled iEEG to 200 Hz using an explicit anti-aliasing procedure, which constrains the analyzable bandwidth to ≤ 100 Hz (and our evaluations to 0.5-50 Hz / 8-13 Hz). Specifically, to prevent aliasing, we apply an explicit anti-aliasing low-pass filter prior to decimation: iEEG is first low-pass filtered with a cutoff below the target Nyquist frequency (100 Hz) and a transition band (implemented in a zero-phase manner via forward-backward filtering), and then resampled to 200 Hz using a polyphase/decimation routine. As a consequence, the maximum representable frequency after downsampling is 100 Hz, and all frequency-domain evaluations and claims in this paper are restricted to the preserved low-/mid-frequency range (PSD: 0.5-50 Hz; alpha: 8-13 Hz; see Tables 2 and 3).

Each subject participates in multiple recording sessions, and each session consists of multiple trials of the verbal working memory task. Due to the individualized clinical implantation schemes (i.e., subject-specific electrode locations, trajectories, and the number of available contacts/units), the intracranial targets and coverage are not directly comparable across subjects; therefore, a controlled cross-subject generalization setting is not feasible with this dataset. Accordingly, we adopt subject-specific (within-subject) evaluation protocols, where models are trained and validated using data from the same subject. To assess robustness at different granularities, we report two within-subject settings. In the cross-trial setting, within each session, 90% of trials are randomly assigned to the training set and the remaining 10% to the evaluation set, so the model is tested on unseen trials from sessions observed during training. In the cross-session setting, entire recording sessions are held out from model fitting and used only for evaluation, so

Table 1: Subject Characteristics

No.	Age	Sex	Pathology	Implanted electrodes	Sessions
1	24	F	Xanthoastrocytoma WHO II	F3, F4, C3, C4, P3, P4, O1, O2, F7, F8, T3, T4, T5, T6, Fz, Cz, Pz, A1, A2 AHL, AL, ECL, LR, PHL, PHR	4
2	39	M	Gliosis	F3, F4, C3, C4, O1, O2, A1, A2 AHL, AHR, AL, AR, ECL, ECR, PHL, PHR	7
3	18	F	Hippocampal sclerosis	F3, F4, C3, C4, O1, O2, A1, A2 AHL, AHR, AL, ECL, PHL	3
4	28	M	Brain contusion	F3, F4, C3, C4, P3, P4, O1, O2, F7, F8, T3, T4, T5, T6, Fz, Cz, Pz, A1, A2 AHL, AHR, AL, AR, ECL, ECR, PHL, PHR	2
5	20	F	Focal cortical dysplasia	Fp1, Fp2, F3, F4, C3, C4, P3, P4, O1, O2, F7, F8, T4, T5, T6, Fz, Cz, Pz, A1, A2 AHL, AL, DRR, PHR	3
6	31	M	Hippocampal sclerosis	Fp1, Fp2, F3, F4, C3, C4, O1, O2, A1, A2 AHL, AHR, AL, AR, ECL, ECR, PHL, PHR	7
7	47	M	Hippocampal sclerosis	F3, F4, C3, C4, O1, O2, A1, A2 AHL, AHR, AL, AR, ECL, ECR, PHL, PHR	4
8	56	F	Hippocampal sclerosis	F3, F4, C3, C4, P3, P4, O1, O2, F7, F8, T3, T4, T5, T6, Fz, Cz, Pz, A1, A2 AHL, AHR, AL, AR, ECL, ECR, PHL, PHR	5
9	19	F	Hippocampal sclerosis	F3, F4, C3, C4, O1, O2, A1, A2 AHL, AHR, AL, AR, ECL, ECR, PHL, PHR	2

¹ Each iEEG electrode label corresponds to eight depth microelectrode contacts positioned along a single trajectory.

² Abbreviations for MTL Regions: AHL/AHR - Anterior Hippocampus (Left/Right), AL/AR - Amygdala (Left/Right), ECL/ECR - Entorhinal Cortex (Left/Right), PHL/PHR - Parahippocampal Gyrus (Left/Right), LR - Lateral Rhinal Area, DRR - Dorsal Rhinal Region.

the model is tested on unseen sessions from the same subject. The cross-trial and cross-session results are reported in Tables 2 and 3, respectively. After splitting, the recordings are further divided into smaller segments of 200 samples (1 second) using a sliding-window approach with a 90% overlap for the training set, while the evaluation set is divided into segments of 200 samples with no overlap. Z-score normalization is applied to each segment, ensuring that the mean and standard deviation of the signals are zero and one, respectively. This normalization step is crucial for stabilizing the training process and improving the model’s performance.

3.2.1. Regression baseline comparison setup

To evaluate whether NeuroFlowNet preserves inter-channel correlation structure more effectively than deterministic alternatives, we additionally compare it with four aligned sequence-to-sequence regression baselines: (i) a point-wise linear mapping implemented as a 1×1 convolution (Linear), (ii) a shallow 1D CNN regressor (Shallow CNN), (iii) a non-causal 1D U-Net regressor (1D U-Net), and (iv) a tiny Transformer regressor (Tiny Transformer). All models are trained and evaluated using the same subject-specific setting (S1, S6, S9), the same trial-level split (90%/10%), and the same segmentation strategy (1-second windows of 200 samples; training windows with overlap and validation windows without overlap). This ensures that any performance differences are attributable to modeling choices rather than data partitioning or preprocessing.

For each 1-second segment, we compute the Pearson inter-channel correlation matrix $\mathbf{R} \in \mathbb{R}^{C \times C}$ from the iEEG channels over the time axis. Let $\mathbf{R}^{\text{pred}}$ and $\mathbf{R}^{\text{true}}$ denote the predicted and ground-truth correlation matrices, respectively. We report two matrix-level errors:

$$\mathcal{E}_{\text{MAE}}(\mathbf{R}) = \frac{1}{C^2} \|\mathbf{R}^{\text{pred}} - \mathbf{R}^{\text{true}}\|_1, \quad \mathcal{E}_F(\mathbf{R}) = \|\mathbf{R}^{\text{pred}} - \mathbf{R}^{\text{true}}\|_F, \quad (10)$$

where $\|\cdot\|_1$ is the element-wise ℓ_1 norm and $\|\cdot\|_F$ is the Frobenius norm. Metrics are first averaged within each subject and then across subjects (S1, S6, S9), reported as mean $\pm$ std. Lower values indicate better preservation of inter-channel correlation structure in the reconstructed band-limited signals.

3.3. Hyperparameter Selection and Ablation Study

The performance of the NeuroFlowNet model is influenced by several key hyperparameters, including the number of scales N_S , the number of flow steps per scale N_{steps} , and the number of channels in the hidden layers of the conditional neural network C_h . To determine the optimal values for these hyperparameters, we conducted a systematic




Figure 4: Hyperparameter selection results for NeuroFlowNet. The performance of the model is evaluated on the validation set and averaged across all subjects and electrodes. The hyperparameters include N_S , the number of scales (a); N_{steps} , the number of flow steps per scale (b); and C_h , the number of channels in the hidden layers of the conditional neural network (c). The results indicate that the model achieves optimal performance with $N_S = 2$, $N_{steps} = 6$, and $C_h = 128$.

evaluation using a one-factor-at-a-time (OFAT) approach. This method involves varying one hyperparameter while keeping the others fixed at their default values, allowing us to isolate the effect of each hyperparameter on model performance. The candidate values for each hyperparameter were selected based on preliminary experiments and domain knowledge. The performance of the model was evaluated on the validation set and averaged across all subjects to ensure that the selected hyperparameters generalize well across different individuals. Specifically, as shown in Figure 4(a), varying the number of scales $N_S \in \{1, 2, 3\}$ leads to a small improvement when increasing from $N_S = 1$ to $N_S = 2$ (higher correlation with lower MAE). Further increasing to $N_S = 3$ brings no consistent gain, with performance remaining essentially unchanged within the error bars. As illustrated in Figure 4(b), we evaluated the number of flow steps per scale $N_{steps} \in \{2, 4, 6, 8\}$. The correlation improves as N_{steps} increases up to 6, while MAE shows only minor variation; increasing to $N_{steps} = 8$ yields no clear additional benefit. Finally, as shown in Figure 4(c), we assessed the effect of the hidden channels $C_h \in \{32, 64, 128, 256\}$ in the conditional network. The model achieves the best (or near-best) validation performance at $C_h = 128$, whereas smaller widths underperform slightly and a larger width ($C_h = 256$) does not improve results. Based on these observations and considering computational efficiency, we select $N_S = 2$, $N_{steps} = 6$, and $C_h = 128$ as the hyperparameter setting for training NeuroFlowNet.

To validate the architectural choice of the conditional network g_{cond} in the affine coupling layers, we conduct a controlled ablation by removing the MHSA module from g_{cond} while keeping all other components and training/evaluation settings unchanged. Figure 5 summarizes the results on the validation set. Overall, incorporating MHSA yields better performance when averaged across subjects and electrodes. The mean correlation increases from 0.479 ± 0.077 (No MHSA) to 0.497 ± 0.060 (With MHSA), and the mean MAE decreases from $42.74 \pm 9.66 \mu V$ to $42.09 \pm 8.35 \mu V$ (mean $\pm$ std across subjects and electrodes). These results suggest that MHSA strengthens g_{cond} in capturing cross-modal dependencies between EEG and iEEG, improving reconstruction consistency (higher correlation) and slightly reducing amplitude error (lower MAE) on average. This ablation provides component-level evidence that the attention mechanism is a beneficial design element in g_{cond} for cross-modal generation.

3.4. Implementation Details

The NeuroFlowNet model is implemented using PyTorch, a popular deep learning framework. The model architecture consists of multiple scales, each containing a series of flow steps that transform the iEEG data conditioned on the EEG data. The model is trained using the AdamW optimizer with a learning rate of 1×10^{-3} and a weight decay of 1×10^{-2} . The model is trained on a single NVIDIA RTX 4080 GPU with 16 GB of memory. The batch size is set to 128, and the model is trained for 100 epochs. The N_S (number of scales), N_{steps} (number of flow steps per scale) and C_h (number of channels in the hidden layers of the conditional neural network) are treated as hyperparameters. Candidate values are chosen empirically and then evaluated on the validation set using OFAT procedure, where one hyperparameter is varied while the others are fixed, as illustrated in Section 3.3. The final model architecture is set to have $N_S = 2$ scales, with $N_{steps} = 6$ flow steps per scale, and $C_h = 128$ channels in the hidden layers of the conditional neural network. The model is trained separately for each subject, and the best-performing model on the validation set is selected. Early stopping is implemented based on the performance on the validation set, with a patience of 10 epochs.




Figure 5: Ablation of the MHSA module in the conditional network g_{cond} . Correlation (higher is better) and MAE (lower is better) are reported on the validation set. Bars indicate the mean across subjects and electrodes, and error bars denote the standard deviation across subjects

4. Results and Discussion

The performance of NeuroFlowNet is evaluated based on its ability to generate iEEG signals from EEG signals in terms of temporal waveform, power spectral density (PSD) waveform, alpha PSD waveform and functional connection mode. The results reveal a clear pattern of region- and subject-specific variability, reflecting both the strengths and current limitations of the model.

4.1. Temporal Waveform Comparison

A qualitative assessment of the model’s performance is presented in Figure 6, which displays exemplar generated iEEG signals (orange) overlaid with the corresponding ground truth iEEG signals (blue) for a 200ms segment from Subject S6. The figure provides a visual comparison across various MTL subregions, including the anterior hippocampus (AHL/AHR), amygdala (AL/AR), entorhinal cortex (ECL/ECR), and parahippocampal gyrus (PHL/PHR). As showed in Figure 6, the generated waveforms demonstrate a high degree of fidelity to the ground truth signals, successfully capturing the principal morphological characteristics and temporal dynamics. The model effectively replicates both the slower, underlying oscillations and the relatively faster fluctuations that remain observable under the 200 Hz bandwidth. The alignment is particularly strong in regions such as the amygdala and the parahippocampal gyrus, where the generated signals closely trace the contours of the ground truth waveforms in terms of phase and relative amplitude.

In more structurally complex or deeper regions, such as the anterior hippocampus and entorhinal cortex, the model still captures the general trend of the signal but occasionally shows minor discrepancies in peak amplitude or precise temporal alignment. For instance, some of the fast transients in the ground truth signal are slightly smoothed or have a marginally lower amplitude in the generated counterpart. Nevertheless, the overall correspondence remains robust, underscoring NeuroFlowNet’s capability to learn the complex conditional mapping from non-invasive EEG to invasive iEEG signals across multiple anatomical targets. This visual evidence corroborates that the model generates temporally coherent and physiologically plausible iEEG waveforms.

To quantify the temporal waveform performance of the model, Tables 2 and 3 report the mean absolute error (MAE), Pearson correlation coefficient (Corr), and cosine similarity across multiple MTL subregions for subjects S1, S6, and S9 under the cross-trial and cross-session settings, respectively. Averaged across all listed regions, the cross-trial setting yields an MAE of $39.93 \mu\text{V}$, a correlation of 0.50, and a cosine similarity of 0.58, while the cross-session setting yields $47.22 \mu\text{V}$, 0.54, and 0.61. Therefore, moving from unseen trials to unseen sessions does not cause a systematic collapse of waveform similarity; instead, the temporal metrics remain in the same range, indicating that NeuroFlowNet captures subject-specific EEG-to-iEEG mappings that are not restricted to a single recording session.

Table 2: Cross-trial performance of NeuroFlowNet on iEEG signal generation.

Subject	Region	Temporal waveform			PSD waveform (0.5-50Hz)		Alpha PSD waveform (8-13Hz)		
		MAE (μ V)	Corr	Cosine	Corr	Cosine	Corr	Cosine	Power MAPE (%)
S1	AHL	40.66 $\pm$ 21.53	0.35 $\pm$ 0.28	0.47 $\pm$ 0.24	0.66 $\pm$ 0.22	0.70 $\pm$ 0.19	0.25 $\pm$ 0.53	0.70 $\pm$ 0.21	2.65 $\pm$ 4.40
	AL	42.82 $\pm$ 28.98	0.39 $\pm$ 0.30	0.50 $\pm$ 0.26	0.67 $\pm$ 0.23	0.71 $\pm$ 0.20	0.29 $\pm$ 0.52	0.71 $\pm$ 0.21	3.33 $\pm$ 7.35
	ECL	49.15 $\pm$ 27.35	0.40 $\pm$ 0.25	0.50 $\pm$ 0.21	0.66 $\pm$ 0.24	0.70 $\pm$ 0.21	0.32 $\pm$ 0.52	0.72 $\pm$ 0.21	6.49 $\pm$ 16.65
	LR	33.86 $\pm$ 20.08	0.60 $\pm$ 0.28	0.65 $\pm$ 0.25	0.76 $\pm$ 0.23	0.79 $\pm$ 0.20	0.45 $\pm$ 0.48	0.78 $\pm$ 0.18	1.24 $\pm$ 2.18
	PHL	35.68 $\pm$ 15.08	0.56 $\pm$ 0.33	0.62 $\pm$ 0.28	0.74 $\pm$ 0.23	0.77 $\pm$ 0.20	0.44 $\pm$ 0.49	0.77 $\pm$ 0.20	1.23 $\pm$ 2.87
	PHR	82.42 $\pm$ 72.44	0.24 $\pm$ 0.29	0.39 $\pm$ 0.27	0.53 $\pm$ 0.25	0.59 $\pm$ 0.22	0.23 $\pm$ 0.53	0.70 $\pm$ 0.20	2.85 $\pm$ 5.36
S6	AHL	44.08 $\pm$ 27.47	0.44 $\pm$ 0.23	0.51 $\pm$ 0.21	0.61 $\pm$ 0.21	0.67 $\pm$ 0.18	0.33 $\pm$ 0.51	0.74 $\pm$ 0.19	1.58 $\pm$ 3.63
	AHR	40.80 $\pm$ 21.73	0.58 $\pm$ 0.21	0.63 $\pm$ 0.19	0.68 $\pm$ 0.20	0.73 $\pm$ 0.17	0.48 $\pm$ 0.50	0.80 $\pm$ 0.19	0.95 $\pm$ 1.88
	AL	45.78 $\pm$ 19.99	0.46 $\pm$ 0.23	0.52 $\pm$ 0.21	0.63 $\pm$ 0.21	0.67 $\pm$ 0.18	0.39 $\pm$ 0.48	0.76 $\pm$ 0.19	1.75 $\pm$ 2.97
	AR	26.77 $\pm$ 9.21	0.74 $\pm$ 0.16	0.77 $\pm$ 0.14	0.78 $\pm$ 0.17	0.81 $\pm$ 0.15	0.67 $\pm$ 0.38	0.87 $\pm$ 0.14	0.61 $\pm$ 1.01
	ECL	53.71 $\pm$ 21.22	0.49 $\pm$ 0.23	0.54 $\pm$ 0.21	0.67 $\pm$ 0.21	0.71 $\pm$ 0.18	0.37 $\pm$ 0.51	0.75 $\pm$ 0.19	1.75 $\pm$ 3.41
	ECR	34.92 $\pm$ 14.32	0.64 $\pm$ 0.17	0.68 $\pm$ 0.16	0.71 $\pm$ 0.20	0.75 $\pm$ 0.18	0.59 $\pm$ 0.43	0.84 $\pm$ 0.15	1.05 $\pm$ 2.48
	PHL	25.73 $\pm$ 15.79	0.39 $\pm$ 0.29	0.49 $\pm$ 0.25	0.60 $\pm$ 0.21	0.66 $\pm$ 0.18	0.32 $\pm$ 0.51	0.74 $\pm$ 0.20	1.38 $\pm$ 2.51
	PHR	36.89 $\pm$ 16.07	0.60 $\pm$ 0.20	0.64 $\pm$ 0.18	0.70 $\pm$ 0.19	0.74 $\pm$ 0.17	0.51 $\pm$ 0.45	0.79 $\pm$ 0.18	0.84 $\pm$ 1.53
S9	AHL	29.92 $\pm$ 10.79	0.67 $\pm$ 0.19	0.73 $\pm$ 0.15	0.81 $\pm$ 0.17	0.82 $\pm$ 0.16	0.36 $\pm$ 0.49	0.73 $\pm$ 0.20	0.88 $\pm$ 1.24
	AHR	43.49 $\pm$ 26.44	0.38 $\pm$ 0.28	0.50 $\pm$ 0.24	0.69 $\pm$ 0.23	0.71 $\pm$ 0.21	0.20 $\pm$ 0.55	0.67 $\pm$ 0.22	7.20 $\pm$ 14.62
	AL	35.30 $\pm$ 11.46	0.61 $\pm$ 0.19	0.66 $\pm$ 0.16	0.78 $\pm$ 0.19	0.80 $\pm$ 0.17	0.34 $\pm$ 0.50	0.72 $\pm$ 0.21	1.45 $\pm$ 3.01
	AR	29.48 $\pm$ 21.12	0.34 $\pm$ 0.34	0.46 $\pm$ 0.30	0.64 $\pm$ 0.25	0.68 $\pm$ 0.22	0.22 $\pm$ 0.54	0.67 $\pm$ 0.23	4.82 $\pm$ 8.81
	ECL	38.73 $\pm$ 17.66	0.55 $\pm$ 0.25	0.61 $\pm$ 0.21	0.75 $\pm$ 0.22	0.77 $\pm$ 0.20	0.34 $\pm$ 0.52	0.72 $\pm$ 0.22	1.62 $\pm$ 2.73
	ECR	43.11 $\pm$ 18.69	0.44 $\pm$ 0.26	0.51 $\pm$ 0.23	0.71 $\pm$ 0.22	0.74 $\pm$ 0.20	0.24 $\pm$ 0.52	0.70 $\pm$ 0.21	3.46 $\pm$ 5.78
	PHL	26.60 $\pm$ 14.10	0.73 $\pm$ 0.20	0.78 $\pm$ 0.17	0.84 $\pm$ 0.16	0.85 $\pm$ 0.15	0.46 $\pm$ 0.46	0.77 $\pm$ 0.18	0.93 $\pm$ 1.39
	PHR	38.48 $\pm$ 17.46	0.49 $\pm$ 0.26	0.55 $\pm$ 0.25	0.74 $\pm$ 0.22	0.76 $\pm$ 0.20	0.33 $\pm$ 0.51	0.71 $\pm$ 0.21	5.35 $\pm$ 10.49

¹ The results of each region are averaged across all electrodes within that region.

² Abbreviations for MTL Regions: AHL/AHR - Anterior Hippocampus (Left/Right), AL/AR - Amygdala (Left/Right), ECL/ECR - Entorhinal Cortex (Left/Right), PHL/PHR - Parahippocampal Gyrus (Left/Right), LR - Lateral Rhinal Area.

At the regional level, the same broad performance pattern is preserved in both settings. More stable reconstructions are typically observed in regions such as S6-AR and S9-PHL in the cross-trial setting and S6-AR, S6-ECR, S1-LR, and S1-PHL in the cross-session setting, all of which show temporal correlations above 0.69 and, in some cases, above 0.80. In contrast, deeper or more difficult targets, such as S1-AHL and especially S1-PHR, remain challenging across both settings. This consistency suggests that the dominant source of variability is region- and subject-dependent observability rather than a simple overfitting to particular trials.

Notably, MAE is more sensitive to session-level amplitude shifts than correlation-based metrics. For example, in subject S1, the average MAE increases from 47.43 μ V in the cross-trial setting to 70.21 μ V in the cross-session setting, whereas the average temporal correlation simultaneously increases from 0.42 to 0.51. This divergence indicates that session-wise changes in signal scale or baseline can affect amplitude fidelity even when waveform morphology remains aligned. Taken together, the temporal results support a robustness claim that is appropriately limited: NeuroFlowNet generalizes to unseen sessions within a subject, but amplitude calibration is still more fragile than waveform-shape recovery.

4.2. Power Spectral Density Comparison

To evaluate the frequency reproduction capability within the preserved bandwidth of NeuroFlowNet, Figure 7 shows a comparison of the PSD waveforms of the generated signals (orange) and the actual signals (blue) for eight MTL subregions of subject S6. The PSD waveform for each region is the average of all electrodes in that region. The analysis reveals that NeuroFlowNet successfully captures the salient spectral characteristics of the native iEEG signals. Across all depicted regions, the generated PSDs closely mirror the ground truth within the preserved bandwidth, particularly in replicating the prominent power peak in the alpha band (8-13 Hz) and a secondary peak in the theta band (4-8 Hz) [50, 51]. These frequency bands are physiologically critical in the MTL; theta oscillations are integral to memory encoding and retrieval, while alpha rhythms are associated with functional inhibition and attentional modulation. The model’s accurate generation of these peaks indicates that it learns to infer the state of these underlying cognitive processes from the conditioning EEG data. Furthermore, the model correctly reproduces the aperiodic,

Table 3: Cross-session performance of NeuroFlowNet on iEEG signal generation.

Subject	Region	Temporal waveform			PSD waveform (0.5-50Hz)		Alpha PSD waveform (8-13Hz)		
		MAE (μ V)	Corr	Cosine	Corr	Cosine	Corr	Cosine	Power MAPE (%)
S1	AHL	52.48 $\pm$ 25.26	0.32 $\pm$ 0.27	0.43 $\pm$ 0.24	0.71 $\pm$ 0.22	0.74 $\pm$ 0.20	0.17 $\pm$ 0.50	0.69 $\pm$ 0.19	3.00 $\pm$ 6.60
	AL	67.56 $\pm$ 44.54	0.56 $\pm$ 0.23	0.63 $\pm$ 0.19	0.81 $\pm$ 0.19	0.83 $\pm$ 0.17	0.21 $\pm$ 0.52	0.69 $\pm$ 0.20	5.43 $\pm$ 16.05
	ECL	67.11 $\pm$ 32.60	0.46 $\pm$ 0.24	0.55 $\pm$ 0.19	0.73 $\pm$ 0.23	0.76 $\pm$ 0.21	0.18 $\pm$ 0.50	0.68 $\pm$ 0.20	10.32 $\pm$ 21.51
	LR	46.66 $\pm$ 23.64	0.83 $\pm$ 0.16	0.84 $\pm$ 0.14	0.93 $\pm$ 0.12	0.94 $\pm$ 0.11	0.40 $\pm$ 0.49	0.78 $\pm$ 0.18	1.15 $\pm$ 2.07
	PHL	38.74 $\pm$ 17.30	0.84 $\pm$ 0.15	0.84 $\pm$ 0.13	0.93 $\pm$ 0.12	0.93 $\pm$ 0.11	0.41 $\pm$ 0.50	0.77 $\pm$ 0.19	0.71 $\pm$ 1.00
	PHR	148.69 $\pm$ 70.33	0.05 $\pm$ 0.21	0.29 $\pm$ 0.27	0.46 $\pm$ 0.23	0.53 $\pm$ 0.20	0.08 $\pm$ 0.52	0.67 $\pm$ 0.20	5.69 $\pm$ 8.00
S6	AHL	48.31 $\pm$ 30.98	0.49 $\pm$ 0.23	0.56 $\pm$ 0.20	0.63 $\pm$ 0.21	0.68 $\pm$ 0.18	0.38 $\pm$ 0.50	0.76 $\pm$ 0.19	1.32 $\pm$ 2.46
	AHR	42.83 $\pm$ 19.44	0.62 $\pm$ 0.19	0.66 $\pm$ 0.17	0.70 $\pm$ 0.20	0.74 $\pm$ 0.17	0.53 $\pm$ 0.47	0.82 $\pm$ 0.17	0.76 $\pm$ 1.25
	AL	48.63 $\pm$ 21.44	0.52 $\pm$ 0.22	0.57 $\pm$ 0.20	0.65 $\pm$ 0.21	0.69 $\pm$ 0.19	0.44 $\pm$ 0.48	0.77 $\pm$ 0.19	1.51 $\pm$ 2.70
	AR	28.86 $\pm$ 10.62	0.76 $\pm$ 0.14	0.79 $\pm$ 0.12	0.79 $\pm$ 0.17	0.82 $\pm$ 0.15	0.71 $\pm$ 0.38	0.89 $\pm$ 0.14	0.55 $\pm$ 1.03
	ECL	56.35 $\pm$ 21.61	0.52 $\pm$ 0.21	0.57 $\pm$ 0.19	0.66 $\pm$ 0.21	0.71 $\pm$ 0.18	0.43 $\pm$ 0.48	0.77 $\pm$ 0.18	1.55 $\pm$ 2.61
	ECR	35.72 $\pm$ 11.22	0.69 $\pm$ 0.16	0.72 $\pm$ 0.15	0.74 $\pm$ 0.18	0.77 $\pm$ 0.16	0.67 $\pm$ 0.39	0.87 $\pm$ 0.14	0.76 $\pm$ 1.18
	PHL	29.04 $\pm$ 16.25	0.44 $\pm$ 0.30	0.55 $\pm$ 0.24	0.61 $\pm$ 0.22	0.67 $\pm$ 0.19	0.34 $\pm$ 0.51	0.75 $\pm$ 0.19	1.10 $\pm$ 1.86
	PHR	38.91 $\pm$ 15.87	0.64 $\pm$ 0.18	0.68 $\pm$ 0.16	0.73 $\pm$ 0.19	0.76 $\pm$ 0.16	0.53 $\pm$ 0.45	0.81 $\pm$ 0.18	0.75 $\pm$ 1.43
S9	AHL	30.48 $\pm$ 10.93	0.66 $\pm$ 0.19	0.71 $\pm$ 0.16	0.80 $\pm$ 0.18	0.81 $\pm$ 0.16	0.39 $\pm$ 0.51	0.74 $\pm$ 0.20	1.22 $\pm$ 2.68
	AHR	42.39 $\pm$ 26.01	0.37 $\pm$ 0.29	0.47 $\pm$ 0.26	0.69 $\pm$ 0.22	0.71 $\pm$ 0.20	0.24 $\pm$ 0.52	0.69 $\pm$ 0.21	7.02 $\pm$ 15.96
	AL	37.40 $\pm$ 13.91	0.58 $\pm$ 0.21	0.63 $\pm$ 0.19	0.76 $\pm$ 0.19	0.78 $\pm$ 0.17	0.35 $\pm$ 0.51	0.72 $\pm$ 0.21	1.58 $\pm$ 3.21
	AR	29.71 $\pm$ 22.06	0.34 $\pm$ 0.35	0.45 $\pm$ 0.31	0.65 $\pm$ 0.24	0.69 $\pm$ 0.21	0.30 $\pm$ 0.54	0.70 $\pm$ 0.23	4.51 $\pm$ 8.39
	ECL	40.97 $\pm$ 20.06	0.53 $\pm$ 0.28	0.58 $\pm$ 0.26	0.73 $\pm$ 0.21	0.76 $\pm$ 0.19	0.38 $\pm$ 0.52	0.74 $\pm$ 0.21	1.71 $\pm$ 3.39
	ECR	43.92 $\pm$ 20.59	0.41 $\pm$ 0.27	0.48 $\pm$ 0.23	0.67 $\pm$ 0.23	0.70 $\pm$ 0.20	0.28 $\pm$ 0.52	0.70 $\pm$ 0.21	3.86 $\pm$ 8.00
	PHL	26.67 $\pm$ 13.45	0.72 $\pm$ 0.19	0.76 $\pm$ 0.17	0.83 $\pm$ 0.17	0.84 $\pm$ 0.15	0.47 $\pm$ 0.49	0.77 $\pm$ 0.20	1.07 $\pm$ 1.94
	PHR	37.42 $\pm$ 16.85	0.49 $\pm$ 0.27	0.56 $\pm$ 0.24	0.75 $\pm$ 0.20	0.77 $\pm$ 0.19	0.38 $\pm$ 0.50	0.74 $\pm$ 0.20	5.57 $\pm$ 11.69

¹ The results of each region are averaged across all electrodes within that region.

² Abbreviations for MTL Regions: AHL/AHR - Anterior Hippocampus (Left/Right), AL/AR - Amygdala (Left/Right), ECL/ECR - Entorhinal Cortex (Left/Right), PHL/PHR - Parahippocampal Gyrus (Left/Right), LR - Lateral Rhinal Area.

1/f-like decay in power at higher frequencies, a hallmark of background neural activity [52, 53]. This demonstrates a robust generalization beyond just the dominant oscillatory components.

To further scrutinize the model’s performance on the functionally significant alpha band, Figure 8 provides a detailed analysis of generated versus true alpha power. The scatter plot (Figure 8a) shows a strong positive correlation, with the linear fit line ($y=0.67x+0.13$) closely following the ideal line ($y=x$). This indicates that the model reliably captures the relative changes in alpha power. The slope of less than one suggests a slight regression to the mean, where the model may temper the most extreme high or low power values, a common characteristic when inferring high-resolution signals from lower signal-to-noise data. The Bland-Altman plot (Figure 8b) reinforces this finding by showing that the mean difference between true and generated alpha power is exceptionally low (0.01), indicating no systematic bias in the model’s predictions. The data points are evenly scattered within the 95% limits of agreement, confirming that while individual predictions have some variance, the model’s estimations are, on average, highly accurate. Collectively, the spectral analysis demonstrates that NeuroFlowNet generates iEEG signals that are not only temporally plausible but also spectrally realistic, preserving the key neurophysiological frequency fingerprints essential for downstream analysis.

To quantify the frequency-domain reproduction capability of the model, Tables 2 and 3 summarize the 0.5-50 Hz PSD waveform metrics (Corr and cosine similarity) and the 8-13 Hz α PSD waveform metrics (Corr, cosine similarity, and mean absolute percentage error (MAPE)) for multiple MTL subregions of subjects S1, S6, and S9. Across all listed regions, the average PSD correlation/cosine pair changes only slightly from 0.698/0.733 in the cross-trial setting to 0.725/0.756 in the cross-session setting. Likewise, the average alpha-band cosine similarity remains nearly unchanged (0.744 vs. 0.751), and the average alpha power MAPE stays low in both settings (2.43% vs. 2.78%). These results indicate that the low-/mid-frequency spectral fingerprints learned by NeuroFlowNet are stable not only across unseen trials but also across unseen sessions from the same subject.

At the region level, the spectral performance is again heterogeneous but consistent with the temporal observations. In the broad frequency range (0.5-50 Hz), PSD correlations remain high in many regions, from 0.53 ± 0.25 (PHR, S1) to 0.84 ± 0.16 (PHL, S9) in the cross-trial setting and up to 0.93 ± 0.12 in S1-LR and S1-PHL in the cross-session




Figure 6: Exemplar generated iEEG signals (orange) overlaid with ground truth iEEG signals (blue) for various MTL subregions, specifically showing a 200ms segment from Subject S6. The figure illustrates the performance of NeuroFlowNet in generating iEEG signals from EEG data across different anatomical targets, including anterior hippocampus (AHL/AHR), amygdala (AL/AR), entorhinal cortex (ECL/ECR), and parahippocampal gyrus (PHL/PHR).




Figure 7: Comparison of power spectral density (PSD) between ground truth iEEG signals (blue) and generated iEEG signals (orange) for a 200 ms segment from Subject S6. The figure shows the average PSD across all electrodes within each MTL subregion, including anterior hippocampus (AHL/AHR), amygdala (AL/AR), entorhinal cortex (ECL/ECR), and parahippocampal gyrus (PHL/PHR).




Figure 8: (a) Scatter plot of normalized true alpha power versus normalized generated alpha power. The red dashed line represents the ideal line ($y = x$), while the green solid line indicates the linear fit ($y = 0.67x + 0.13$). (b) Bland-Altman plot showing the difference in alpha power against the mean alpha power. The red dashed line indicates the mean difference (0.01), and the green and blue dashed lines represent the upper (95% CI, 0.28) and lower (95% CI, -0.27) limits of agreement, respectively.

setting. Most PSD cosine similarities exceed 0.70, suggesting that the overall spectral profile is preserved even when the exact waveform amplitude is imperfect. In the alpha band, cosine similarities remain stable, whereas correlations fluctuate more strongly, indicating that the model is better at preserving relative alpha-band shape and power than exact segment-wise alpha fluctuations.

The main degradation under the session-level split is concentrated in a small number of difficult regions, especially S1-PHR and S1-ECL, where both temporal and alpha-band errors are larger. Even so, the overall session-level spectral metrics remain comparable to the trial-level metrics. This is important for the robustness discussion: the model is not merely memorizing trial-specific spectral statistics from sessions seen during training, but instead preserves meaningful band-limited frequency structure when transferred to held-out sessions. Such behavior is particularly relevant for downstream neurophysiological analyses that rely more on stable band-power organization than on exact pointwise waveform matching [54, 55, 56].

4.3. Inter-Channel Correlation Comparison

Beyond generating realistic single-channel waveforms, a crucial test for the model is its ability to capture the complex functional relationships between different neural populations. These relationships are reflected in the inter-channel correlation structure of the iEEG signals. Figure 9 presents a qualitative comparison of these correlation patterns, displaying the 50 strongest pairwise Pearson correlations for both the ground truth (a) and generated (b) iEEG signals from a 200ms segment of Subject S6.

The circular plots reveal a close correspondence between the ground truth and the generated data. The model reproduces salient features of the correlation topology visible in this band-limited representation. Key patterns, such as strong correlations between homologous regions in the left and right hemispheres (e.g., between channels in AHL and AHR), are clearly preserved.

Furthermore, the model captures strong correlations between anatomically and functionally related regions within the same hemisphere, such as the entorhinal cortex and hippocampus. The visual similarity of the two plots—from the dense connections within the parahippocampal gyrus to the sparser but strong links spanning distant regions—suggests that NeuroFlowNet is not merely generating a collection of independent, realistic time series, but is also retaining cross-channel covariance patterns observable in the reconstructed signals. Since Pearson correlation is primarily sensitive to shared low-/mid-frequency fluctuations, this comparison is most appropriately interpreted as evidence




Figure 9: Comparison of inter-channel correlation. Circular graphs show the connectivity patterns of (a) the original iEEG signals and (b) the generated iEEG signals. Both plots illustrate a 200 ms segment from Subject S6. Pairwise Pearson correlation coefficients are computed between all channels, and only the top 50 strongest absolute correlations are visualized for clarity.

that the model preserves one important aspect of MTL dependency structure within the retained bandwidth, rather than fully recovering the underlying neural or network dynamics.

4.4. Conditional Sampling and Uncertainty Analysis

To directly examine whether NeuroFlowNet behaves as a stochastic conditional generator rather than a deterministic reconstructor, we performed a conditional sampling and uncertainty analysis on the validation set. For each validation EEG segment, we repeatedly drew $K = 20$ independent latent samples from the base Gaussian distribution and generated 20 corresponding iEEG reconstructions using the same trained subject-specific model and the same preprocessing/postprocessing pipeline. This protocol keeps the conditioning EEG fixed and changes only the latent variable, so any variation among the generated outputs reflects the learned conditional distribution $p(X_{iEEG} | X_{EEG})$.

For each segment, we computed the sample mean, sample variance, and empirical 5th–95th percentile interval across the 20 conditional samples. Conditional sample diversity was quantified in two complementary ways: (i) the average pairwise RMSE between generated waveforms in the time domain, and (ii) the average pairwise Frobenius distance between inter-channel correlation matrices computed from different samples. We further evaluated the 90% predictive interval coverage probability (PICP), mean prediction interval width (MPIW), MAE and temporal correlation of the sample mean, and uncertainty-error consistency. The latter was measured as the Spearman correlation between segment-level predictive uncertainty (mean sample variance across channels and time) and segment-level reconstruction error (MAE of the sample mean).

The analysis provides direct empirical evidence that NeuroFlowNet generates non-degenerate conditional samples. Across subjects, the average pairwise time-domain diversity ranged from 23.48 to 35.59 μV , and the pairwise diversity of inter-channel correlation matrices ranged from 7.55 to 8.64 in Frobenius norm (Table 4). Thus, repeated sampling under an identical EEG condition does not collapse to a single deterministic trace; instead, the model produces a family of EEG-consistent iEEG trajectories with measurable waveform-level and connectivity-level variation. The representative example in Figure 10(a) illustrates this behavior visually: individual stochastic reconstructions fluctuate around the sample mean while retaining the main morphology of the ground-truth signal.

The uncertainty estimates also showed a consistent relationship with reconstruction difficulty. Segment-level predictive variance was strongly and positively correlated with segment-level MAE for all three subjects (S1: $\rho =$




Figure 10: Representative conditional sampling and uncertainty analysis for Subject S6. (a) For a fixed scalp EEG segment, $K = 20$ independent latent samples generate multiple iEEG trajectories for channel mAHL1. The solid line denotes the sample mean, the shaded area denotes the empirical 90% predictive interval, and the dashed line denotes the ground-truth iEEG. (b) Across all validation segments of S6, each point represents one segment. Predictive uncertainty is defined as the mean sample variance across channels and time, and reconstruction error is defined as the MAE between the sample mean and the ground truth. Higher predictive uncertainty is strongly associated with larger reconstruction error ($\rho = 0.858$, $p = 4.52 \times 10^{-68}$).

Table 4: Conditional sampling and uncertainty metrics on the validation set.

Subject	Segments	Channels	Diversity RMSE (μV)	Corr. diversity (Frobenius)	PICP ₉₀	MPIW (μV)	Mean variance (μV^2)	Mean MAE (μV)	Mean Corr.	Unc.-error ρ
S1	112	48	35.59	7.55	0.381	62.00	841.51	49.37	0.490	0.958 ($p = 2.03 \times 10^{-61}$)
S6	231	64	28.56	8.64	0.432	51.57	434.66	35.78	0.599	0.858 ($p = 4.52 \times 10^{-68}$)
S9	70	64	23.48	7.99	0.382	42.94	286.27	34.01	0.572	0.878 ($p = 1.92 \times 10^{-23}$)

$K = 20$ conditional samples were generated for each validation EEG segment. Diversity RMSE is the average pairwise RMSE among conditional samples. Correlation diversity is the average pairwise Frobenius distance between inter-channel correlation matrices. PICP₉₀ and MPIW are computed from the empirical 5th–95th percentile interval. Mean Corr. denotes the temporal correlation between the sample mean reconstruction and ground-truth iEEG.

0.958, $p = 2.03 \times 10^{-61}$; S6: $\rho = 0.858$, $p = 4.52 \times 10^{-68}$; S9: $\rho = 0.878$, $p = 1.92 \times 10^{-23}$). This indicates that the stochastic spread of conditional samples is not arbitrary noise; segments for which the model is more uncertain also tend to be segments with larger reconstruction error. Region-level results showed the same pattern. For example, the S1 PHR region exhibited both the largest uncertainty (mean variance: $2580.25 \mu V^2$) and the largest sample-mean error ($84.46 \mu V$), whereas lower-error regions such as S1 PHL and S6/S9 PHL showed substantially smaller uncertainty.

At the same time, the empirical predictive intervals were not fully calibrated: PICP₉₀ ranged from 0.381 to 0.432, well below the nominal 0.90 level. Therefore, these results should not be interpreted as evidence of a calibrated Bayesian posterior interval. Rather, they support the more specific conclusion that NeuroFlowNet exhibits genuine conditional sample diversity and that its sample variance provides an informative, although under-dispersed, uncertainty signal. This directly strengthens the generative interpretation of the model while keeping the scope of the claim limited to stochastic behavior and uncertainty-error consistency under the present subject-specific reconstruction protocol.

4.5. Functional Connectivity Fidelity Comparison with Regression Baselines

As emphasized above, an important requirement for non-invasive iEEG reconstruction is not only waveform plausibility at each channel, but also the preservation of inter-channel correlation structure that reflects MTL functional connectivity. In practice, many downstream analyses in cognitive and clinical neuroscience rely on network-level signatures (e.g., co-fluctuations across contacts/regions), and a model that produces visually plausible traces yet distorts cross-channel dependencies may be of limited utility for connectivity-driven interpretation. Therefore, we com-

Table 5: Functional connectivity fidelity (inter-channel correlation structure) comparison between NeuroFlowNet and regression baselines.

Matrix	Linear (1×1)	Shallow CNN	1D U-Net	Tiny Transformer	NeuroFlowNet (Ours)
$\mathcal{E}_{\text{MAE}}(\mathbf{R})$	0.426 ± 0.064	0.318 ± 0.032	0.302 ± 0.013	0.360 ± 0.080	0.288 ± 0.042
$\mathcal{E}_F(\mathbf{R})$	30.45 ± 6.45	23.37 ± 3.99	22.08 ± 2.57	26.04 ± 6.86	21.61 ± 4.69

Note: $\mathbf{R}$ denotes the inter-channel Pearson correlation matrix computed over time for each 1-second segment. $\mathcal{E}_{\text{MAE}}(\mathbf{R}) = \frac{1}{C^2} \|\mathbf{R}^{\text{pred}} - \mathbf{R}^{\text{true}}\|_1$ and $\mathcal{E}_F(\mathbf{R}) = \|\mathbf{R}^{\text{pred}} - \mathbf{R}^{\text{true}}\|_F$, averaged within each subject and then across subjects, reported as mean $\pm$ std.

plement our within-model connectivity demonstration with a controlled comparison to representative deterministic regression baselines under the protocol defined in Section 3.2.1.

Table 5 summarizes the correlation-structure fidelity across models. NeuroFlowNet achieves the lowest errors among the compared methods, with $\mathcal{E}_{\text{MAE}}(\mathbf{R}) = 0.288 \pm 0.042$ and $\mathcal{E}_F(\mathbf{R}) = 21.61 \pm 4.69$. Compared with the pointwise linear baseline, NeuroFlowNet reduces correlation-matrix MAE from 0.426 ± 0.064 to 0.288 ± 0.042 (a relative reduction of $\sim 32\%$), and reduces Frobenius error from 30.45 ± 6.45 to 21.61 ± 4.69 (a relative reduction of $\sim 29\%$). This indicates that a simple per-time linear mapping is insufficient to recover the cross-channel dependency patterns expressed in these band-limited iEEG signals, whereas NeuroFlowNet better preserves the corresponding correlation structure under the same training data regime.

Beyond the linear baseline, NeuroFlowNet also consistently outperforms the stronger nonlinear regressors included in this comparison (Shallow CNN, 1D U-Net, and Tiny Transformer). While these deterministic models can learn nonlinear mappings and often generate temporally smooth outputs, they still tend to emphasize pointwise signal reconstruction objectives, which can inadvertently *attenuate* or *homogenize* cross-channel relationships—especially when the inverse mapping from scalp EEG to deep iEEG is underconstrained. In such ill-posed settings, a deterministic regressor is incentivized to output a “central” solution that reduces average prediction error across trials, which may appear reasonable at the level of individual traces but can shrink trial-to-trial covariance and distort the correlation matrix. The observed gap in $\mathcal{E}_{\text{MAE}}(\mathbf{R})$ and $\mathcal{E}_F(\mathbf{R})$ therefore suggests that NeuroFlowNet better preserves the correlation structure most visible to the present Pearson-correlation analysis, aligning with the qualitative similarity of connectivity graphs shown in Figure 9. Because this metric is dominated by shared low-/mid-frequency fluctuations and by the bandwidth retained after preprocessing, the result is best understood as improved preservation of correlation structure within that reconstructed frequency range.

4.6. Architectural Interpretation and Scope of Connectivity Preservation

One conceptual explanation for this advantage lies in the inductive biases of conditional normalizing flows and the specific architectural design of NeuroFlowNet. First, CNF optimizes exact conditional likelihood, enabling the model to represent the *conditional distribution* of iEEG given EEG rather than collapsing all variability into a single deterministic output. Conceptually, this is consistent with the non-uniqueness of EEG inverse problems: multiple intracranial configurations can give rise to similar scalp observations due to volume conduction and source mixing. By introducing a latent variable and learning an invertible mapping, NeuroFlowNet can in principle retain residual degrees of freedom that are not fully determined by EEG, instead of forcing them into a mean prediction. This provides a plausible account of why the model may better preserve trial-level co-variations across channels, which then manifest as more faithful correlation matrices, although the specific physiological meaning of the latent representation is not established by the present experiments.

The invertible 1×1 convolution in each flow step provides an explicit channel-mixing mechanism. Unlike standard convolutions that may learn channel interactions implicitly (and potentially inconsistently across layers), the invertible mixing operation encourages structured information exchange across channels at every temporal location while maintaining reversibility. This is particularly relevant in multi-channel iEEG, where functional connectivity is encoded in coordinated activity patterns across contacts/regions. By repeatedly mixing channels through invertible transforms, the model can more effectively capture and propagate global cross-channel dependencies throughout the flow, which can improve the fidelity of reconstructed correlation structure.

In addition, the multi-scale architecture further supports structure preservation by decomposing the learning problem across resolutions. Fine scales focus on local temporal detail, whereas coarser scales capture more global, slowly

varying patterns that often drive connectivity estimates. Since correlation matrices integrate information over time, errors in low-frequency or large-scale components can disproportionately affect connectivity fidelity. The hierarchical split-and-propagate design allows NeuroFlowNet to allocate capacity to both local waveform characteristics and global coordination patterns, which may explain its stronger performance in $\mathcal{E}_{\text{MAE}}(\mathbf{R})$ and $\mathcal{E}_F(\mathbf{R})$.

The role of the multi-head self-attention (MHSA) module should also be interpreted in this context. MHSA is not introduced primarily as a smoothing operator; rather, its main function is to model long-range temporal dependencies that are difficult to capture with local convolutions alone. In EEG-to-iEEG reconstruction, such dependencies can link temporally separated events, phase relationships, and slowly evolving contextual states that shape the instantaneous waveform. Any apparent smoothing in the generated traces is therefore better understood as a by-product of learning temporally coherent conditional structure under a band-limited setting, not as the principal purpose of attention itself.

Relatedly, the different architectural components likely contribute unevenly to waveform reconstruction and connectivity preservation. The finer scales and local convolutional operations are expected to contribute more directly to pointwise waveform morphology, including peak shape, local oscillatory detail, and short-lag temporal alignment. By contrast, the coarser scales, MHSA, and repeated channel-mixing operations are more naturally aligned with preserving slowly varying coordination patterns across time and channels, which are the main drivers of the inter-channel correlation analysis used here. Because the present experiments were not designed as a strict component-wise ablation, this interpretation should be viewed as a mechanistic explanation of the observed performance profile rather than a definitive quantitative attribution. Nevertheless, it provides a coherent account of why NeuroFlowNet can simultaneously improve single-channel realism and better retain band-limited connectivity structure.

These findings should be interpreted within the subject-specific scope of this dataset. The present comparison demonstrates that, under identical within-subject training conditions, NeuroFlowNet yields improved fidelity of band-limited inter-channel correlation structure relative to representative deterministic regressors. Together with the conditional sampling results in Section 4.4, this supports a more specific and empirically grounded claim: within the current protocol, the CNF formulation improves structure-preserving reconstruction over deterministic regressors while also producing observable stochastic conditional samples with informative uncertainty behavior. A protocol-matched comparison with other generative families, such as diffusion models or GANs, remains an important direction for future work. Future studies will also extend these connectivity-driven evaluations to broader cohorts and more diverse brain states, and will assess whether this correlation-structure advantage translates into improved performance on downstream tasks that explicitly rely on network-level features (e.g., connectivity-based decoding or seizure network characterization).

4.7. Computational latency and real-time feasibility

Since NeuroFlowNet relies on reversible transformations, we further evaluated the inference latency to assess its feasibility for real-time monitoring. The latency was measured using the final model configuration selected in this study, i.e., $N_S = 2$, $N_{\text{steps}} = 6$, and $C_h = 128$. For each test, a 1-s EEG segment sampled at 200 Hz was used as the conditioning input, and the model generated the corresponding iEEG segment under batch size 1. All measurements were performed in evaluation mode with gradient computation disabled, and the first 50 runs were discarded as warm-up. On a single NVIDIA RTX 4080 GPU, the average inference latency was 9.536 ± 2.449 ms per 1-s segment, with a median latency of 8.724 ms and a 95th-percentile latency of 14.253 ms. The corresponding real-time factor was 0.00954, indicating that the inverse reversible transformations required less than 1% of the input window duration. Even the maximum observed latency, 26.609 ms, remained far below the 1000-ms window length. These results suggest that NeuroFlowNet has practical potential for real-time or near-real-time monitoring.

The computational cost of the reversible transformations is mainly determined by the repeated evaluations of the conditional network g_{cond} in the affine coupling layers and the inverse 1×1 convolutions. Importantly, during inference, the model does not require iterative optimization; iEEG generation is achieved by a single inverse pass through the learned flow. Therefore, although the reversible architecture introduces additional transformation steps compared with simple deterministic regression models, its inference process remains deterministic in computational schedule and can be efficiently accelerated on modern GPUs. Future work will further optimize the model for deployment on low-power edge devices and streaming EEG systems.

4.8. Discussion of Physiological Determinants in Regional Variability

The observed performance heterogeneity across MTL subregions can be partially attributed to neuroanatomical and physiological factors. In terms of structural complexity, the anterior hippocampus (AHL/AHR) is characterized by intricate laminar organization and high neuronal density, resulting in complex and non-linear iEEG signal patterns [57]. This complexity likely exceeds the current representational capacity of NeuroFlowNet’s conditional normalizing flow, contributing to reduced fidelity in this region [58]. Besides, deeply situated structures like AHL experience greater signal attenuation in EEG recordings due to intervening tissues (cerebrospinal fluid, skull), reducing the signal-to-noise ratio of the conditioning data and complicating accurate inference of iEEG dynamics [59]. In contrast, more superficial regions such as the parahippocampal gyrus benefit from proximity to the cortical surface, aiding model performance [60]. Additionally, regions with strong cortico-MTL connectivity, such as the entorhinal cortex (ECL/ECR), may exhibit iEEG dynamics that are more inferable from scalp-recorded potentials, aligning with their relatively high temporal and spectral performance [61].

4.9. Discussion of Bandwidth and Denoising Limitations

A key preprocessing step in this work is downsampling iEEG from 2 kHz to 200 Hz to match scalp EEG and to enforce an aligned temporal grid for conditional modeling. While this alignment is important for stable supervised training and fair cross-modal comparison, it imposes an explicit bandwidth constraint: after downsampling, the effective Nyquist frequency is 100 Hz, and spectral content above 100 Hz cannot be represented. In addition, the required anti-aliasing low-pass filter applied before decimation/resampling can attenuate components near the cutoff, which may further reduce the sharpness of rapid fluctuations even below 100 Hz.

Another related preprocessing limitation is the use of wavelet denoising before model training. Although denoising was applied to suppress broadband noise and improve signal quality, its effect on transient waveform morphology was not directly evaluated in the present study. In particular, the denoising procedure may smooth sharp transient components or attenuate part of the faster fluctuations. Therefore, the waveform comparisons reported in this work should be interpreted as reconstruction performance relative to the denoised and band-limited target iEEG signals, rather than as evidence that all transient morphological features in the original broadband iEEG recordings were preserved.

This limitation is particularly relevant for intracranial biomarkers that are typically studied in higher frequency ranges, such as high-gamma activity and high-frequency oscillations, as well as very fast transient events. These phenomena are not the target of the present reconstruction setting and cannot be faithfully assessed under a 200 Hz sampling pipeline. Accordingly, the term “high-fidelity” in this study refers to fidelity within the preserved bandwidth of the preprocessing chain and relative to the denoised target signals, rather than broadband fidelity across the full spectrum available in the original 2 kHz iEEG.

Consistent with this scope, our quantitative frequency-domain evaluations are restricted to low-/mid-frequency bands, specifically PSD similarity in 0.5–50 Hz and alpha-band (8–13 Hz) characteristics (Tables 2 and 3). Moreover, our functional-connectivity analyses rely on inter-channel correlation matrices computed from these preprocessed time series, and our conclusions are therefore confined to (i) temporal waveform plausibility at 200 Hz resolution, (ii) preservation of low-/mid-frequency spectral fingerprints, and (iii) recovery of inter-channel dependency structure within this bandwidth.

Future work will extend the framework toward higher-frequency intracranial signatures by (i) training and evaluating at higher target sampling rates, such as 500 Hz or 1 kHz, when synchronized EEG-iEEG data of sufficient quality are available, and/or (ii) adopting multi-resolution modeling strategies in which low-frequency components are reconstructed at 200 Hz for cross-modal alignment while higher-frequency components are modeled with an additional dedicated branch or a hierarchical refinement stage. Future studies should also include a systematic raw-versus-denoised comparison to quantify how preprocessing affects transient morphology before model training. These directions would enable a more comprehensive assessment of whether non-invasive EEG can support reconstruction of clinically and cognitively relevant high-frequency iEEG biomarkers beyond the present denoised and band-limited setting.

4.10. Discussion of Physiological Interpretation of the Latent Space Representation Z

To further strengthen the interpretability of NeuroFlowNet, we discuss the potential physiological meaning of the learned latent representation Z . At the present stage, the physiological interpretation of Z is best viewed as a

source of working hypotheses about the residual intracranial dynamics captured by the model, rather than as a direct identification of specific physiological processes. In conditional normalizing flows, Z is obtained through an invertible transformation $Z = f(X_i, X_c)$, where X_i and X_c denote the target iEEG and conditioning scalp EEG, respectively. Due to the bijective nature of the flow mapping, Z should not be interpreted as a single physiological “source” or a direct anatomical generator in the classical sense. Instead, we view Z as a compact parametrization of the residual degrees of freedom in intracranial activity that are not uniquely constrained by scalp EEG, i.e., a principled representation of the conditional uncertainty in $p(X_{iEEG} | X_{EEG})$ [62].

From a neurophysiological perspective, this residual component can plausibly reflect multiple factors. First, it may capture trial-to-trial variability of local circuit dynamics within the MTL that is only weakly expressed on the scalp due to volume conduction, source mixing, and limited signal-to-noise ratio [63, 64, 65]. Second, it may encode micro-scale and transient activity (e.g., sharp fluctuations and other local population events) that substantially shapes iEEG waveforms but is attenuated or spatially smeared in scalp recordings [65, 66, 67]. Third, it may represent state-dependent stochastic modulations of ongoing rhythms, where fluctuations in cognitive state or vigilance alter intracranial oscillatory patterns without being fully predictable from scalp EEG alone [68]. Importantly, this interpretation is consistent with the well-known non-uniqueness of EEG inverse problems: multiple intracranial configurations can produce similar scalp observations, and introducing a latent variable provides a natural probabilistic mechanism to represent this ambiguity rather than forcing a single deterministic reconstruction [63, 64].

This perspective also helps contextualize the regional variability discussed above. For deeper or more structurally complex targets (e.g., the anterior hippocampus), scalp observability is expected to be reduced, and thus a larger portion of iEEG variability may remain underconstrained by EEG [65, 66, 69]. In NeuroFlowNet, such underconstrained components may be represented through the latent pathway Z , enabling the model to avoid collapsing to an “average” waveform and instead generate diverse yet EEG-consistent iEEG samples [62]. We emphasize that the above points constitute a physiologically motivated interpretation of Z rather than an established one-to-one identification of latent dimensions with specific neural generators or physiological processes. Future work should test these hypotheses using dedicated interpretability analyses (e.g., linking latent factors to band-limited power modulation, event-related signatures, or unit/spike-informed markers) on larger cohorts with broader brain-state coverage [70].

4.11. Discussion of Clinical Translation and Robustness Considerations

NeuroFlowNet demonstrates the feasibility of reconstructing MTL iEEG patterns from scalp EEG using a public synchronized EEG-iEEG dataset. The cross-session evaluation further provides a more stringent within-subject robustness test than random trial splitting, as the model is required to generalize to entire recording sessions excluded from training. The broadly comparable performance between cross-session and cross-trial evaluations suggests that NeuroFlowNet is not merely exploiting trial-level redundancy within individual sessions, but instead captures a relatively stable, subject-specific EEG-to-iEEG mapping.

Nevertheless, the present evidence should be interpreted within a subject-specific and task-dependent scope [70]. The dataset consists of recordings from epilepsy patients with individualized implantation schemes, and intracranial electrode coverage is not standardized across subjects [71]. Consequently, direct cross-subject generalization cannot be rigorously established using the current data. Thus, although our findings support the methodological potential of non-invasive iEEG reconstruction, they do not yet warrant strong clinical claims, particularly for high-stakes applications such as non-invasive pre-surgical localization [71].

Clinical translation of EEG-to-iEEG reconstruction ultimately depends on robust performance across both subjects and brain states. Inter-subject variability arises from differences in anatomy, skull conductivity, pathology, medication, and electrode implantation targets, all of which can alter the scalp-to-intracranial mapping [69, 63, 64, 72]. Cross-state variability is also expected, as intracranial dynamics change substantially across vigilance and cognitive conditions (e.g., rest vs. task engagement), and across clinically relevant contexts (e.g., interictal vs. ictal periods) [68, 71]. These factors may introduce distribution shifts that challenge purely subject-specific modeling and motivate broader validation [65, 66].

To move toward clinical readiness, several steps are necessary. First, larger multi-subject datasets with more comparable intracranial coverage (or standardized ROI definitions) are required to enable systematic evaluation of cross-subject generalization and the development of adaptation strategies that account for anatomical and physiological differences [71, 69]. Second, robustness should be assessed across heterogeneous brain states, ideally including

resting state, sleep, multiple cognitive tasks, and interictal/ictal segments, to establish state generalization [68, 71]. Third, the clinical value of reconstructed signals should be validated in downstream endpoints where performance can be directly linked to decision-making, such as interictal spike detection, seizure onset zone (SOZ) localization, or surgical outcome prediction [71]. Taken together, we position NeuroFlowNet as a methodological step toward non-invasive deep-brain signal inference, while recognizing that comprehensive cross-subject and cross-state validation is an essential prerequisite for clinical translation [71].

5. Conclusion

In this study, we introduced NeuroFlowNet, a novel generative framework based on conditional normalizing flows, to address the enduring challenge of generating high-fidelity (band-limited) iEEG from non-invasive EEG within the preserved bandwidth. Our work makes a pivotal contribution by demonstrating, for the first time, the feasibility of reconstructing iEEG signals across multiple MTL subregions. By directly modeling the conditional probability distribution, NeuroFlowNet provides a latent-variable mechanism for representing residual iEEG variability that is not fully determined by scalp EEG. Our results confirm that the generated signals not only achieve high fidelity in temporal waveforms and spectral characteristics but also successfully preserve the complex inter-channel correlation structure, reflecting the underlying functional connectivity of the MTL network. The cross-session experiments further show that these properties remain broadly stable when evaluation is shifted from unseen trials to unseen recording sessions from the same subject, strengthening the robustness claim of the method. In addition, a baseline comparison against representative deterministic regressors shows that NeuroFlowNet yields lower errors in the inter-channel correlation matrix under the same within-subject protocol, supporting its advantage in structure-preserving reconstruction rather than purely pointwise fitting. Conditional sampling with 20 latent draws per EEG segment further demonstrates non-degenerate stochastic reconstructions and a strong positive association between predictive variance and reconstruction error, providing direct evidence of generative behavior under the present protocol. The latency analysis further shows that the final model can generate a 1-s iEEG segment in only a few milliseconds on a modern GPU, supporting its practical potential for real-time or near-real-time monitoring. These findings suggest that probabilistic iEEG reconstruction may provide a useful non-invasive window into deep brain dynamics, particularly for analyses that require both waveform plausibility and multi-channel dependency structure.

However, we acknowledge certain limitations. The model’s performance exhibits variability across different MTL subregions, with deeper and more structurally complex areas like the anterior hippocampus proving more challenging to reconstruct accurately. This suggests that the signal-to-noise ratio of the conditioning EEG and the inherent complexity of the target signal are key performance determinants. Additionally, since iEEG is downsampled to 200 Hz with anti-aliasing, the present results and spectral claims are limited to the preserved low-/mid-frequency range (≤ 100 Hz; evaluated mainly in 0.5-50 Hz and 8-13 Hz), and do not address high-frequency intracranial biomarkers such as high-gamma/HFO. The conditional sampling analysis also showed under-coverage of empirical 90% prediction intervals, indicating that the current uncertainty estimates are informative but not fully calibrated. Moreover, the present study does not yet establish cross-subject or cross-state generalization, and it does not include a protocol-matched comparison with alternative generative models such as diffusion models or GANs. Clinical deployment should therefore be considered premature until the framework is validated on larger cohorts, heterogeneous brain states, and broader generative baselines.

Future work will focus on four main directions: 1) Enhancing the model architecture by incorporating more sophisticated conditioning mechanisms that can better leverage spatial information from the EEG montage. 2) Expanding the framework to a subject-independent model by collecting larger-scale datasets with more standardized and comparable electrode implantation coverage across subjects, which will enable a systematic evaluation of cross-subject generalization and support the development of advanced domain adaptation techniques to account for inter-subject anatomical and physiological variability [73, 72, 74]. 3) Improving uncertainty calibration and comparing NeuroFlowNet with protocol-matched generative baselines, including diffusion- or GAN-based conditional generators. 4) Validating the clinical utility of the generated iEEG in downstream tasks, such as automated seizure detection and cognitive state decoding, to bridge the gap between technical innovation and practical application.

CRediT authorship contribution statement

Dongyi He: Conceptualization, Methodology, Software, Data curation, Writing - original draft, Writing - review & editing, Visualization. **Bin Jiang:** Conceptualization, Supervision, Writing - review & editing, Funding acquisition, Writing - original draft, Formal analysis. **Kecheng Feng:** Writing - review & editing, Visualization, Formal analysis. **Luyin Zhang:** Writing - review & editing. **Ling Liu:** Writing - review & editing. **Yuxuan Li:** Writing - review & editing. **Yun Zhao:** Funding acquisition, Formal analysis, Writing - review & editing, Supervision, Resources. **He Yan:** Supervision, Resources. **All authors have read and approved the final manuscript.**

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethics statement

This study was conducted in publicly available datasets, thus ethical approval was not required.

Declaration of generative AI and AI-assisted technologies in the writing process.

During the preparation of this work the author(s) used ChatGPT in order to check the grammar and spelling of the manuscript. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

Funding Declaration

This research was funded by the Scientific and Technological Research Program of the Chongqing Education Commission (KJZD-K202303103, KJQN202501104) and Chongqing Municipal Key Project for Technology Innovation and Application Development (CSTB2024TIAD-KPX0042).

References

- [1] A. Craik, Y. He, J. L. Contreras-Vidal, Deep learning for electroencephalogram (eeg) classification tasks: a review, *Journal of neural engineering* 16 (3) (2019) 031001.
- [2] A. Kumari, D. R. Edla, A study on brain–computer interface: Methods and applications, *SN Computer Science* 4 (2) (2022) 98.
- [3] B. Jiang, X. Wang, D. He, S. Cheng, M. Liao, D. Miao, Q. Xia, Y. Zhao, G. Li, A comprehensive review of deep learning for motor imagery eeg: From healthy subjects to patients, *Biomedical Signal Processing and Control* 117 (2026) 109680.
- [4] Y. Li, W. Zeng, W. Dong, D. Han, L. Chen, H. Chen, Z. Kang, S. Gong, H. Yan, W. T. Siok, et al., A tale of single-channel electroencephalogram: Devices, datasets, signal processing, applications, and future directions, *IEEE Transactions on Instrumentation and Measurement* (2025).
- [5] D. He, W. T. Siok, N. Wang, Toward robust, reproducible, and widely accessible intracranial language brain-computer interfaces: A comprehensive review of neural mechanisms, hardware, algorithms, evaluation, clinical pathways and future directions, *arXiv preprint arXiv:2603.12279* (2026).
- [6] C. M. Michel, D. Brunet, Eeg source imaging: a practical review of the analysis steps, *Frontiers in neurology* 10 (2019) 325.

- [7] A. Biasucci, R. Leeb, I. Iturrate, S. Perdikis, A. Al-Khodairy, T. Corbet, A. Schnider, T. Schmidlin, H. Zhang, M. Bassolino, et al., Brain-actuated functional electrical stimulation elicits lasting arm motor recovery after stroke, *Nature communications* 9 (1) (2018) 2421.
- [8] M. Hassan, F. Wendling, Electroencephalography source connectivity: aiming for high resolution of brain networks in time and space, *IEEE Signal Processing Magazine* 35 (3) (2018) 81–96.
- [9] B. He, L. Astolfi, P. A. Valdés-Sosa, D. Marinazzo, S. O. Palva, C.-G. Bénar, C. M. Michel, T. Koenig, Electrophysiological brain connectivity: theory and implementation, *IEEE transactions on biomedical engineering* 66 (7) (2019) 2115–2137.
- [10] Y. Zhao, B. Jiang, Y. Yan, D. He, X. Zhang, S. Zou, Y. Zhu, X. Hu, G. Xie, Rgftslanet: A cross-subject classification model for decoding eeg-based motor imagery tasks, in: 2025 Cross Strait Radio Science and Wireless Technology Conference (CSRSWTC), IEEE, 2025, pp. 1–4.
- [11] D. He, S. Li, B. Jiang, H. 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, *Journal of Neural Engineering* 22 (5) (2025) 056042.
- [12] X. Ma, W. Chen, Z. Pei, Y. Zhang, J. Chen, Attention-based convolutional neural network with multi-modal temporal information fusion for motor imagery eeg decoding, *Computers in Biology and Medicine* 175 (2024) 108504.
- [13] N. Padfield, J. Zabalza, H. Zhao, V. Masero, J. Ren, Eeg-based brain-computer interfaces using motor-imagery: Techniques and challenges, *Sensors* 19 (6) (2019) 1423.
- [14] H. Yu, Z. Lin, F. Li, J. Liu, C. Liu, J. Wang, Spatospectral representation and neural decoding of somatic perception of acupuncture stimulations, *IEEE Journal of Biomedical and Health Informatics* (2025).
- [15] H. Yu, F. Zeng, D. Liu, J. Wang, J. Liu, Neural manifold decoder for acupuncture stimulations with representation learning: An acupuncture-brain interface, *IEEE Journal of Biomedical and Health Informatics* 29 (6) (2025) 4147–4160.
- [16] J. Parvizi, S. Kastner, Human intracranial eeg: promises and limitations, *Nature neuroscience* 21 (4) (2018) 474.
- [17] J. Parvizi, S. Kastner, Promises and limitations of human intracranial electroencephalography, *Nature neuroscience* 21 (4) (2018) 474–483.
- [18] R. Mukamel, I. Fried, Human intracranial recordings and cognitive neuroscience, *Annual review of psychology* 63 (1) (2012) 511–537.
- [19] R. Grech, T. Cassar, J. Muscat, K. P. Camilleri, S. G. Fabri, M. Zervakis, P. Xanthopoulos, V. Sakkalis, B. Vanrumste, Review on solving the inverse problem in eeg source analysis, *Journal of neuroengineering and rehabilitation* 5 (1) (2008) 25.
- [20] S. M. M. Basri, H. Al-Nashash, H. Mir, Inverse problem for m/eeg source localization: a review, *IEEE Sensors Journal* (2024).
- [21] S. Baillet, Magnetoencephalography for brain electrophysiology and imaging, *Nature neuroscience* 20 (3) (2017) 327–339.
- [22] A. Sohrabpour, Y. Lu, P. Kankirawatana, J. Blount, H. Kim, B. He, Effect of eeg electrode number on epileptic source localization in pediatric patients, *Clinical Neurophysiology* 126 (3) (2015) 472–480.
- [23] T. Hedrich, G. Pellegrino, E. Kobayashi, J.-M. Lina, C. Grova, Comparison of the spatial resolution of source imaging techniques in high-density eeg and meg, *Neuroimage* 157 (2017) 531–544.

- [24] B. He, A. Sohrabpour, E. Brown, Z. Liu, Electrophysiological source imaging: a noninvasive window to brain dynamics, *Annual review of biomedical engineering* 20 (1) (2018) 171–196.
- [25] M. Seeber, L.-M. Cantonas, M. Hoevels, T. Sesia, V. Visser-Vandewalle, C. M. Michel, Subcortical electrophysiological activity is detectable with high-density eeg source imaging, *Nature communications* 10 (1) (2019) 753.
- [26] P. Ryvlin, Seeg in 2025: progress and pending challenges in stereotaxy methods, biomarkers and radiofrequency thermocoagulation, *Current opinion in neurology* 38 (2) (2025) 111–120.
- [27] C. Herff, D. J. Krusienski, P. Kubben, The potential of stereotactic-eeg for brain-computer interfaces: current progress and future directions, *Frontiers in neuroscience* 14 (2020) 123.
- [28] S. K. Veeramalla, V. H. R. Talari, Multiple dipole source localization of eeg measurements using particle filter with partial stratified resampling, *Biomedical engineering letters* 10 (2) (2020) 205–215.
- [29] T. R. Knösche, M. Gräser, A. Anwander, Prior knowledge on cortex organization in the reconstruction of source current densities from eeg, *Neuroimage* 67 (2013) 7–24.
- [30] R. R. Ramírez, D. Wipf, S. Baillet, Neuroelectromagnetic source imaging of brain dynamics, in: *Computational Neuroscience*, Springer, 2010, pp. 127–155.
- [31] S. Cui, L. Duan, B. Gong, Y. Qiao, F. Xu, J. Chen, C. Wang, Eeg source localization using spatio-temporal neural network, *China Communications* 16 (7) (2019) 131–143.
- [32] M. Saeidi, W. Karwowski, F. V. Farahani, K. Fiok, R. Taiar, P. A. Hancock, A. Al-Juaid, Neural decoding of eeg signals with machine learning: a systematic review, *Brain sciences* 11 (11) (2021) 1525.
- [33] F. Hassan, S. F. Hussain, Review of eeg signals classification using machine learning and deep-learning techniques, in: *Advances in non-invasive biomedical signal sensing and processing with machine learning*, Springer, 2023, pp. 159–183.
- [34] H. Yu, X. Lei, Z. Song, C. Liu, J. Wang, Supervised network-based fuzzy learning of eeg signals for alzheimer’s disease identification, *IEEE Transactions on Fuzzy Systems* 28 (1) (2019) 60–71.
- [35] J. Rong, R. Sun, B. Joseph, G. Worrell, B. He, Deep learning-based eeg source imaging is robust under varying electrode configurations, *Clinical Neurophysiology* (2025).
- [36] L. Bozhkov, P. Georgieva, Overview of deep learning architectures for eeg-based brain imaging, in: *2018 International joint conference on neural networks (IJCNN)*, IEEE, 2018, pp. 1–7.
- [37] Z. Hui, L. Ying, W. Lingyue, Y. Ning, Y. Shuo, Eeg dipole source localization using deep neural network, *Int. J. Circuits Syst. Signal Process* 16 (2022) 132–140.
- [38] L. Hecker, R. Rupprecht, L. Tebartz Van Elst, J. Kornmeier, Convdip: A convolutional neural network for better eeg source imaging, *Frontiers in Neuroscience* 15 (2021) 569918.
- [39] G. Huang, K. Liu, J. Liang, C. Cai, Z. H. Gu, F. Qi, Y. Li, Z. L. Yu, W. Wu, Electromagnetic source imaging via a data-synthesis-based convolutional encoder–decoder network, *IEEE Transactions on Neural Networks and Learning Systems* 35 (5) (2022) 6423–6437.
- [40] S. M. Kaviri, R. Vinjamuri, Integrating electroencephalography source localization and residual convolutional neural network for advanced stroke rehabilitation, *Bioengineering* 11 (10) (2024) 967.
- [41] J. C. Bore, P. Li, L. Jiang, W. M. Ayedh, C. Chen, D. J. Harmah, D. Yao, Z. Cao, P. Xu, A long short-term memory network for sparse spatiotemporal eeg source imaging, *IEEE Transactions on Medical Imaging* 40 (12) (2021) 3787–3800.

- [42] A. Sikka, H. Jamalabadi, M. Krylova, S. Alizadeh, J. N. Van Der Meer, L. Danyeli, M. Deliano, P. Vicheva, T. Hahn, T. Koenig, et al., Investigating the temporal dynamics of electroencephalogram (eeg) microstates using recurrent neural networks, *Human brain mapping* 41 (9) (2020) 2334–2346.
- [43] M. Doan, S. Plis, Scaling synthetic brain data generation, *IEEE Journal of Biomedical and Health Informatics* (2024).
- [44] G. Cisotto, A. Zancanaro, I. F. Zoppis, S. L. Manzoni, hveegnet: a novel deep learning model for high-fidelity eeg reconstruction, *Frontiers in Neuroinformatics* 18 (2024) 1459970.
- [45] B. Abdi-Sargezeh, A. Oswal, S. Sanei, Mapping scalp to intracranial eeg using generative adversarial networks for automatically detecting interictal epileptiform discharges, in: *2023 IEEE Statistical Signal Processing Workshop (SSP)*, IEEE, 2023, pp. 710–714.
- [46] B. Abdi-Sargezeh, S. Shirani, A. Valentin, G. Alarcon, S. Sanei, Eeg-to-eeg: Scalp-to-intracranial eeg translation using a combination of variational autoencoder and generative adversarial networks, *Sensors* 25 (2) (2025) 494.
- [47] G. Chiarion, L. Sparacino, Y. Antonacci, L. Faes, L. Mesin, Connectivity analysis in eeg data: a tutorial review of the state of the art and emerging trends, *Bioengineering* 10 (3) (2023) 372.
- [48] L. Wei, T. Qiu, Z. Wang, W. Mei, A functional connectivity metric method for eeg time series via nonlinear symbolization, *Biomedical Signal Processing and Control* 103 (2025) 107498.
- [49] E. Boran, T. Fedele, P. Klaver, P. Hilfiker, L. Stieglitz, T. Grunwald, J. Sarnthein, Persistent hippocampal neural firing and hippocampal-cortical coupling predict verbal working memory load, *Science advances* 5 (3) (2019) eaav3687.
- [50] N. A. Herweg, E. A. Solomon, M. J. Kahana, Theta oscillations in human memory, *Trends in cognitive sciences* 24 (3) (2020) 208–227.
- [51] O. Jensen, A. Mazaheri, Shaping functional architecture by oscillatory alpha activity: gating by inhibition, *Frontiers in human neuroscience* 4 (2010) 186.
- [52] B. J. He, Scale-free brain activity: past, present, and future, *Trends in cognitive sciences* 18 (9) (2014) 480–487.
- [53] M. A. Kramer, C. J. Chu, The 1/f-like behavior of neural field spectra are a natural consequence of noise driven brain dynamics, *bioRxiv* (2023).
- [54] G. Assenza, F. Capone, L. Di Biase, F. Ferreri, L. Florio, A. Guerra, M. Marano, M. Paolucci, F. Ranieri, G. Salomone, et al., Oscillatory activities in neurological disorders of elderly: biomarkers to target for neuro-modulation, *Frontiers in Aging Neuroscience* 9 (2017) 189.
- [55] P. J. Uhlhaas, W. Singer, Neural synchrony in brain disorders: relevance for cognitive dysfunctions and pathophysiology, *neuron* 52 (1) (2006) 155–168.
- [56] Q. Zhou, J. Lin, L. Yao, Y. Wang, Y. Han, K. Xu, Relative power correlates with the decoding performance of motor imagery both across time and subjects, *Frontiers in Human Neuroscience* 15 (2021) 701091.
- [57] G. Li, S. Jiang, S. E. Paraskevopoulou, M. Wang, Y. Xu, Z. Wu, L. Chen, D. Zhang, G. Schalk, Optimal referencing for stereo-electroencephalographic (seeg) recordings, *NeuroImage* 183 (2018) 327–335.
- [58] C. Winkler, D. Worrall, E. Hoogeboom, M. Welling, Learning likelihoods with conditional normalizing flows, *arXiv preprint arXiv:1912.00042* (2019).
- [59] M. F. Hnazaee, B. Wittevrongel, E. Khachatryan, A. Libert, E. Carrette, I. Dauwe, A. Meurs, P. Boon, D. Van Roost, M. M. Van Hulle, Localization of deep brain activity with scalp and subdural eeg, *NeuroImage* 223 (2020) 117344.

- [60] K. Vlcek, I. Fajnerova, T. Nekovarova, L. Hejtmanek, R. Janca, P. Jezdik, A. Kalina, M. Tomasek, P. Krsek, J. Hammer, et al., Mapping the scene and object processing networks by intracranial eeg, *Frontiers in human neuroscience* 14 (2020) 561399.
- [61] J. Yang, D. Cao, C. Guo, L. Stieglitz, D. Ledergerber, J. Sarnthein, J. Li, Enhanced role of the entorhinal cortex in adapting to increased working memory load, *Nature Communications* 16 (1) (2025) 5798.
- [62] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, B. Lakshminarayanan, Normalizing flows for probabilistic modeling and inference, *Journal of Machine Learning Research* 22 (57) (2021) 1–64.
- [63] S. Asadzadeh, T. Y. Rezaei, S. Beheshti, A. Delpak, S. Meshgini, A systematic review of eeg source localization techniques and their applications on diagnosis of brain abnormalities, *Journal of neuroscience methods* 339 (2020) 108740.
- [64] I. Zorzos, I. Kakkos, E. M. Ventouras, G. K. Matsopoulos, Advances in electrical source imaging: A review of the current approaches, applications and challenges, *Signals* 2 (3) (2021) 378–391.
- [65] M. C. Piastra, A. Nüßing, J. Vorwerk, M. Clerc, C. Engwer, C. H. Wolters, A comprehensive study on electroencephalography and magnetoencephalography sensitivity to cortical and subcortical sources, *Human Brain Mapping* 42 (4) (2021) 978–992.
- [66] J. Afnan, Z. Cai, J.-M. Lina, C. Abdallah, E. Delaire, T. Avigdor, V. Ros, T. Hedrich, N. von Ellenrieder, E. Kobayashi, et al., Eeg/meg source imaging of deep brain activity within the maximum entropy on the mean framework: Simulations and validation in epilepsy, *Human Brain Mapping* 45 (10) (2024) e26720.
- [67] A. A. Liu, S. Henin, S. Abbaspoor, A. Bragin, E. A. Buffalo, J. S. Farrell, D. J. Foster, L. M. Frank, T. Gedankien, J. Gotman, et al., A consensus statement on detection of hippocampal sharp wave ripples and differentiation from other fast oscillations, *Nature communications* 13 (1) (2022) 6000.
- [68] I. Pershin, G. Candrian, M. Münger, G.-M. Baschera, M. Rostami, D. Eich, A. Müller, Vigilance described by the time-on-task effect in eeg activity during a cued go/nogo task, *International Journal of Psychophysiology* 183 (2023) 92–102.
- [69] J. Vorwerk, C. H. Wolters, D. Baumgarten, Global sensitivity of eeg source analysis to tissue conductivity uncertainties, *Frontiers in Human Neuroscience* 18 (2024) 1335212.
- [70] E. Boran, T. Fedele, A. Steiner, P. Hilfiker, L. Stieglitz, T. Grunwald, J. Sarnthein, Dataset of human medial temporal lobe neurons, scalp and intracranial eeg during a verbal working memory task, *Scientific data* 7 (1) (2020) 30.
- [71] B. C. Jobst, F. Bartolomei, B. Diehl, B. Frauscher, P. Kahane, L. Minotti, A. Sharan, N. Tardy, G. Worrell, J. Gotman, Intracranial eeg in the 21st century, *Epilepsy currents* 20 (4) (2020) 180–188.
- [72] X. Wang, Y. Zhao, D. He, Q. Xia, G. Li, N. Wang, N. Peng, B. Jiang, Pa-tcnet: Pathology-aware temporal calibration with physiology-guided target refinement for cross-subject motor imagery eeg decoding in stroke patients, *arXiv preprint arXiv:2604.16554* (2026).
- [73] Y. Zhao, D. He, F. Ren, Q. Xia, L. Xu, G. Xie, X. Zhang, R. Yang, S. Zou, B. Jiang, Rmetnet: A cross-subject motor imagery eeg signal classification model based on tslanet and riemannian geometry features, *PloS one* 21 (4) (2026) e0347671.
- [74] X. Wang, Y. Zhao, D. He, Q. Xia, G. Li, X. Xing, Y. Pan, B. Jiang, Cfspmnet: Cross-subject fourier-guided spatial-patch mamba network for eeg motor imagery decoding in stroke patients, *arXiv preprint arXiv:2605.10111* (2026).