LBTAR: LEARNABLE BRAIN-TOPOLOGY-GUIDED AUTOREGRESSIVE CHANNEL RECOVERY FOR EEG SPATIAL SUPER-RESOLUTION

Anonymous authors
Paper under double-blind review

ABSTRACT

High-density electroencephalography (EEG) captures rich spatial information, but its acquisition and deployment costs are relatively high. Low-density EEG is easier to acquire, but its limited spatial coverage makes it difficult to characterize regional brain activity at a fine spatial scale. Therefore, reconstructing high-density EEG from low-density observations has become an important problem in EEG spatial super-resolution. Existing methods mostly adopt one-shot reconstruction or fixed-order autoregressive recovery, which limits their ability to adapt to sample-specific channel dependencies and dynamically evolving recovery states. To address this limitation, we propose LBTAR, which replaces predefined recovery paths with a state-dependent sequential recovery process. LBTAR treats EEG channels as nodes and integrates channel-to-region assignments, inter-regional adjacency relationships, and electrode spatial information to construct a brain-region-structure-guided channel topology. It further jointly considers the currently visible context, the remaining candidate set, the dynamic channel topology, and channel-level prediction reliability to adaptively select the next group of channels for recovery. At each step, the mask-conditioned reconstruction network produces full-channel predictions, while only the selected channel signals are written back to progressively update the recovery state. Channel-level uncertainty is further used to adjust subsequent topology and recovery decisions, thereby mitigating error propagation. Experiments on three public datasets, namely SEED, SEED-IV, and Localize-MI, show that LBTAR achieves favorable time-domain reconstruction performance and frequency-domain consistency across multiple spatial super-resolution ratios, while maintaining relatively low computational overhead and stable performance at high missing-channel ratios. Ablation experiments further validate the effectiveness of dynamic recovery, brain-region structural priors, and uncertainty-aware modeling.

1 INTRODUCTION

Electroencephalography (EEG) offers high temporal resolution and has been widely used in applications such as cognitive state analysis Jiang et al. (2025); Roy et al. (2019), brain–computer interfaces Yin et al. (2024), and computer-assisted diagnosis of neurological disorders Shang et al. (2024). High-density EEG provides rich spatial information through denser electrode sampling. However, it involves high acquisition costs, complex setup and deployment procedures, and stringent requirements for electrode placement and contact quality, thereby limiting its use in portable brain–computer interfaces, mobile health monitoring, and naturalistic settings. In contrast, low-density EEG is less costly, easier to set up, and more flexible to deploy. Nevertheless, its limited spatial coverage results in spatial undersampling, making it difficult to adequately characterize the spatial distribution of regional brain activity and inter-channel dependencies Wang et al. (2025). Therefore, reconstructing high-density EEG from sparse channel observations has become an important problem in EEG spatial super-resolution.

As illustrated in Fig. 2(a)–(b) in Appendix A, existing EEG spatial super-resolution methods mainly follow two paradigms: one-shot reconstruction and fixed-order autoregressive recovery. The former directly predicts all missing channels through spatial interpolation (Perrin et al., 1989; Dong et al.,

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

2021), end-to-end mapping (Li et al., 2025), or generative modeling (Wang et al., 2025; Liu et al., 2025), but makes limited use of the sequential dependencies among missing channels. The latter progressively recovers predefined channel groups and uses previously recovered signals as additional context for subsequent predictions (Liu et al., 2026). However, its channel grouping and recovery order are typically determined by fixed geometric rules, which limits its ability to adapt to sample-specific channel dependencies and dynamically evolving recovery states.

EEG channel dependencies are influenced not only by the geometric distances between electrodes, but also by channel-to-region assignments, inter-regional adjacency relationships, and sample-specific patterns of brain activity Agarwal et al. (2022); Zhou et al. (2023). As recovery proceeds, the visible context expands and the effective channel topology evolves, while the prediction reliability of previously recovered channels also affects subsequent predictions. Therefore, brain-region structural priors, the current recovery state, and prediction reliability should be jointly considered when determining the channel recovery path.

Based on the above analysis, we propose LBTAR (Learnable Brain-Topology-Guided Autoregressive Channel Recovery), a brain-topology-guided learnable autoregressive channel recovery framework for EEG spatial super-resolution, as shown in Fig. 2(c). LBTAR formulates EEG spatial super-resolution as a state-dependent sequential recovery process and integrates channel-to-region assignments, inter-regional adjacency relationships, and electrode spatial information to construct a channel-level topology. At each recovery stage, the model adaptively selects the next group of channels for recovery by jointly considering the currently visible context, the remaining candidate channels, the dynamic channel topology, and channel-level prediction reliability. The mask-conditioned reconstruction network then produces full-channel predictions, while only the selected channel signals are written back to progressively update the recovery state. Channel-level uncertainty is further fed back to adjust the dynamic topology and guide channel selection at subsequent steps, thereby mitigating error propagation. The main contributions of this work are summarized as follows:

- We propose LBTAR, a brain-region-topology-guided learnable autoregressive framework for EEG spatial super-resolution, which formulates missing-channel reconstruction as a state-dependent sequential recovery process guided by brain-region structural priors and progressively expands the visible context through iterative write-back.
- We design a state-aware learnable channel recovery policy that jointly considers the currently visible context, remaining candidate channels, dynamic brain-region topology, and channel-level prediction reliability to dynamically score candidate channels and adaptively select the next group of channels to be recovered, without requiring predefined channel groups or recovery orders.
- We introduce channel-level uncertainty modeling to explicitly estimate prediction reliability and feed it back into subsequent topology adjustment and channel selection, thereby reducing error propagation caused by using low-reliability predictions as subsequent context.
- We conduct experiments under multiple spatial super-resolution ratios on three public datasets, including SEED, SEED-IV, and Localize-MI. The effectiveness of LBTAR and its stability under high missing-channel ratios are evaluated in terms of time-domain reconstruction, frequency-domain consistency, and computational efficiency, while ablation experiments further analyze the contributions of the core modules.

2 Related Work

Existing EEG spatial super-resolution methods mainly include spatial interpolation, end-to-end deep mapping, generative reconstruction, and autoregressive recovery. Spatial interpolation methods estimate missing-channel signals based on electrode locations or spatial relationships among channels Perrin et al. (1989); Dong et al. (2021). End-to-end deep mapping methods employ deep networks to learn nonlinear mappings between low-density and high-density EEG, as well as cross-channel dependencies Li et al. (2025). Generative methods recover high-density signals from low-density observations through adversarial generation or diffusion modeling Wang et al. (2025); Liu et al. (2025); Corley & Huang (2018). The first three categories mainly adopt one-shot holistic reconstruction, directly predicting all missing channels from low-density EEG.

Unlike one-shot reconstruction, CAFE formulates EEG spatial super-resolution as a channel-group-level autoregressive process. It groups missing channels according to their geometric distances from the observed channels and progressively recovers the groups in a predefined order Liu et al. (2026). Recovered channels are added to the visible set at subsequent stages, thereby progressively expanding the prediction context. This method establishes sequential recovery relationships among missing channels. However, its channel grouping and recovery order are predetermined by geometric rules.

In EEG analysis tasks such as emotion recognition and brain–computer interfaces, previous studies have shown that dependencies among EEG channels are influenced not only by the spatial distances between electrodes, but also by channel-to-region assignments, inter-regional adjacency relationships, and the current signal state Agarwal et al. (2022); Zhou et al. (2023). However, the application of such structural priors to EEG spatial super-resolution remains relatively limited. As the recovery process proceeds, the visible context continues to expand, while the prediction reliability of recovered channels also affects subsequent recovery. Based on these observations, this study jointly considers brain-region structural priors, dynamic recovery states, and prediction reliability to learn sample-specific channel recovery paths.

3 METHOD

Given a high-density EEG signal $\mathbf{X}_H \in \mathbb{R}^{C\times T}$, where C and T denote the number of channels and the number of temporal samples, respectively, let $\mathcal{L}$ denote the set of observed low-density channels. The missing-channel set and the low-density input are defined as $\mathcal{U} = \{1, 2, \ldots, C\} \setminus \mathcal{L}$ and $\mathbf{X}_L = \mathbf{X}_H[\mathcal{L}, :]$, respectively. The model estimates the missing-channel signals from $\mathbf{X}_L$ and merges them with the observed channels to obtain the complete high-density EEG reconstruction. Unlike one-shot reconstruction or sequential recovery based on predefined channel groups, we formulate EEG spatial super-resolution as a state-dependent autoregressive channel recovery process and adaptively determine the recovery path according to the current recovery state, dynamic brain-region topology, and channel-level prediction reliability:

$$\mathcal{L} \to \mathcal{S}_1 \to \mathcal{S}_2 \to \cdots \to \mathcal{S}_G \tag{1}$$

Here, G denotes the total number of recovery steps, and $\mathcal{S}_t \subseteq \mathcal{U}$ denotes the channel set selected at the t-th step. The selected channel sets satisfy

$$\mathcal{S}_i \cap \mathcal{S}_j = \emptyset, \quad i \neq j, \qquad \bigcup_{t=1}^{G} \mathcal{S}_t = \mathcal{U} \tag{2}$$

3.1 OVERVIEW OF LBTAR

This paper proposes LBTAR (Learnable Brain-Topology-Guided Autoregressive Channel Recovery), a brain-topology-guided learnable autoregressive channel recovery method, whose overall architecture is illustrated in Fig. 1. LBTAR consists of four modules: the Brain Topology Builder, Learnable Recovery Policy, Mask-Conditioned Reconstruction Backbone, and Channel Uncertainty Head. The model first maps the low-density observations to the full channel space and constructs the initial assembled EEG and visible-channel mask through zero padding. At each recovery step, the model constructs a dynamic channel topology based on brain-region structural priors and the current prediction reliability. Subsequently, the mask-conditioned reconstruction network encodes the current recovery state and generates channel-node representations and full-channel predictions. The learnable recovery policy scores the remaining candidate channels and selects the next group of channels to be recovered. Finally, the predictions of the selected channels are written back to the current state, and channel-level uncertainty is used to update the reliability information for the next step. This process is performed iteratively until all missing channels have been recovered.

3.2 BRAIN TOPOLOGY BUILDER

To incorporate spatial and brain-region structural priors, we design a Brain Topology Builder that treats EEG channels as nodes and integrates electrode spatial locations, channel-to-region assignments, and brain-region adjacency information to construct a channel-level topology guided by

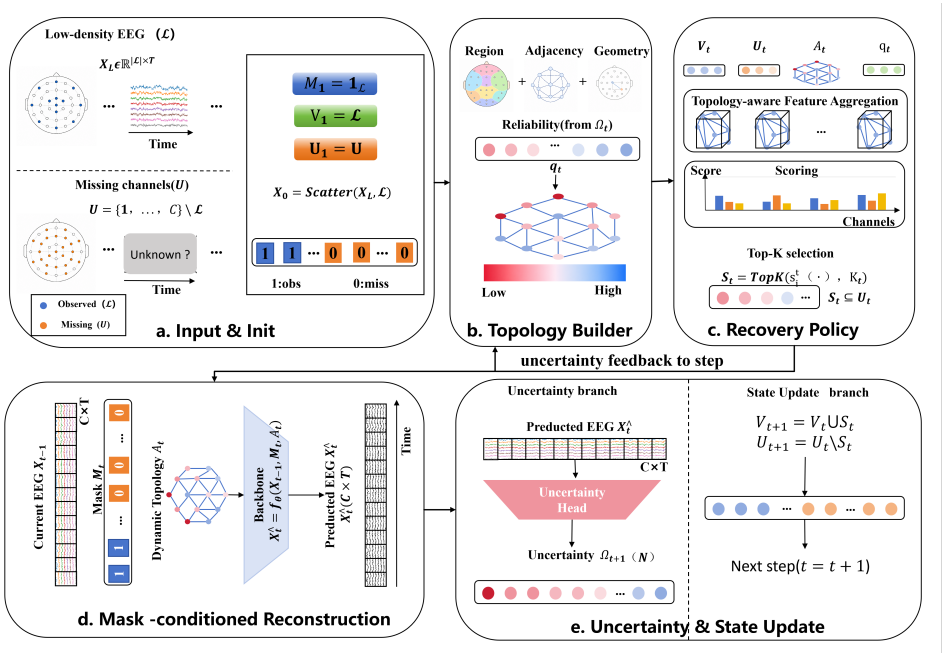


Figure 1: Overall architecture of LBTAR.

brain-region structure. Let the high-density EEG contain C channels. The spatial coordinate and brain-region label of the i-th channel are denoted by $\mathbf{p}_i$ and r_i, respectively. The electrode spatial relationship, same-region relationship, and brain-region adjacency relationship are defined as

$$A_{\text{spa}}(i,j) = \exp\left(-\frac{\|\mathbf{p}_i - \mathbf{p}_j\|_2^2}{\tau}\right) \tag{3}$$

$$A_{\text{reg}}(i,j) = \mathbb{I}\left(r_i = r_j\right) \tag{4}$$

$$A_{\text{adj}}(i,j) = B_{\text{adj}}\left(r_i, r_j\right) \tag{5}$$

The base channel topology is defined as

$$A_{\text{base}} = \text{Norm}\left(\lambda_1 A_{\text{spa}} + \lambda_2 A_{\text{reg}} + \lambda_3 A_{\text{adj}}\right) \tag{6}$$

Here, τ is the distance scale parameter, $\mathbb{I}(\cdot)$ denotes the indicator function, and B_{adj} is a predefined brain-region-level adjacency matrix. The coefficients λ_1, λ_2, and λ_3 control the contributions of the electrode spatial relationship, same-region relationship, and brain-region adjacency relationship, respectively. $\text{Norm}(\cdot)$ denotes the adjacency matrix normalization operation.

To reduce the influence of low-reliability predictions on subsequent recovery, the model converts the channel-level uncertainty $\Omega_t(i)$ available before the t-th recovery step into reliability:

$$q_t(i) = \begin{cases} 1, & t = 1 \\ \exp[-\,\text{softplus}(\Omega_t(i))]\,, & t > 1 \end{cases} \tag{7}$$

The reliability of all channels is initialized to 1 at $t = 1$ and is subsequently updated according to the predicted channel-level uncertainty.

$$A_t(i,j) = A_{\text{base}}(i,j)\, q_t(i)\, q_t(j) \tag{8}$$

This mechanism preserves the EEG spatial and brain-region structural priors while dynamically adjusting the topological connection strengths according to channel prediction reliability, thereby providing reliability-aware structural information for subsequent channel selection and signal reconstruction.

3.3 Learnable Recovery Policy

Fixed recovery orders are difficult to adapt to sample-specific channel dependencies and dynamically changing recovery states. To address this issue, we design a Learnable Recovery Policy that adaptively selects the channels to be recovered at each recovery step.

At the t-th step, for each candidate channel $i \in \mathcal{U}_t$, the model integrates the channel node representation, topological relationships, and prediction reliability to construct a state feature:

$$\mathbf{z}_i^{(t)} = \left[\mathbf{h}_i^{(t)} \,\|\, \sum_{j \in \mathcal{V}_t} A_t(i,j) \quad \| \sum_{j \in \mathcal{U}_t \setminus \{i\}} A_t(i,j) \,\|\, q_t(i) \right] \tag{9}$$

Here, $\mathbf{h}_i^{(t)}$ denotes the node representation of channel i at the t-th step. The two summation terms characterize the topological relationships of the channel with the currently visible channel set $\mathcal{V}_t$ and the other remaining candidate channels, respectively. A_t denotes the dynamic channel topology, $q_t(i)$ denotes the channel prediction reliability, and $\|$ denotes feature concatenation.

The scoring network generates recovery scores from the state features and selects the K_t channels with the highest scores from the remaining candidate channel set $\mathcal{U}_t$ as the recovery set at the current step:

$$\mathcal{S}_t = \mathrm{TopK}\left(\left\{ \mathrm{MLP}_\phi\left(\mathbf{z}_i^{(t)}\right) \mid i \in \mathcal{U}_t \right\}, K_t \right) \tag{10}$$

Here, ϕ denotes the learnable parameters of the scoring network, and K_t denotes the number of channels selected at the t-th step. By jointly considering the current recovery state, dynamic channel topology, and prediction reliability, this policy adaptively determines the channels to be recovered at each step without requiring predefined channel groups or a fixed recovery order.

3.4 Mask-conditioned Reconstruction Backbone

To explicitly distinguish visible channels from missing channels, we design a Mask-conditioned Reconstruction Backbone that jointly uses the currently assembled EEG and the visible-channel mask as reconstruction conditions. The model first maps the low-density EEG into the complete channel space according to the positions of the observed channels and initializes the missing channels to zero. At the t-th step, the binary mask $\mathbf{M}_t$ is constructed according to the current visible-channel set $\mathcal{V}_t$, where visible channels are assigned a value of 1 and the remaining channels are assigned a value of 0. After being expanded along the temporal dimension, the mask is concatenated with the currently assembled EEG and fed into the reconstruction backbone shared across all recovery steps:

$$\mathbf{Z}_t = \mathrm{Concat}\left(\mathbf{X}^{(t-1)}, \mathrm{Expand}(\mathbf{M}_t) \right) \tag{11}$$

$$\mathbf{H}_t, \widehat{\mathbf{X}}_t = f_\theta\left(\mathbf{Z}_t, A_t\right) \tag{12}$$

Here, $\mathbf{Z}_t \in \mathbb{R}^{2C \times T}$ denotes the mask-conditioned input, f_θ denotes the shared reconstruction backbone, and $\mathbf{H}_t \in \mathbb{R}^{C \times d}$ denotes the channel node representations, where $\mathbf{h}_i^{(t)} = \mathbf{H}_t(i,:)$. $\widehat{\mathbf{X}}_t \in \mathbb{R}^{C \times T}$ denotes the complete channel prediction at the t-th step.

The network employs dilated temporal convolutions to extract EEG temporal features and uses the dynamic brain-region topology A_t for cross-channel information interaction. Although the network predicts all channels at each step, state write-back is applied only to the currently selected channel set $\mathcal{S}_t$. After completing G recovery steps, the final high-density EEG reconstruction is obtained.

3.5 Channel Uncertainty Head

To evaluate the prediction reliability of the recovered channels and mitigate error accumulation, we design a Channel Uncertainty Head to perform channel-level uncertainty estimation.

After channel prediction is completed at the t-th step, the uncertainty head generates the channel-level uncertainty used in the next step based on the current prediction results, the assembled EEG before state write-back, and the visible-channel mask:

$$\mathbf{\Omega}_{t+1} = u_\psi\left(\widehat{\mathbf{X}}_t, \mathbf{X}^{(t-1)}, \mathbf{M}_t\right) \in \mathbb{R}^C \tag{13}$$

Here, $\Omega_{t+1}(i)$ denotes the predicted log-variance of channel i, with a larger value indicating higher uncertainty. Based on statistical features of the predicted signals, assembled signals, and visibility states, the model generates channel-level uncertainty through a multilayer perceptron and uses it for the uncertainty loss, dynamic topology adjustment, and subsequent channel selection, thereby mitigating the influence of low-reliability predictions on subsequent recovery.

3.6 Optimization and Rollout Strategy

To jointly optimize signal reconstruction, channel-level uncertainty estimation, and the dynamic recovery policy, LBTAR adopts the following joint training objective.

$$\mathcal{J} = \mathcal{L}_{\text{re}} + \lambda_{\text{un}}\mathcal{L}_{\text{un}} + \lambda_{\text{sp}}\mathcal{L}_{\text{sp}} + \lambda_{\text{policy}}\mathcal{L}_{\text{policy}} \tag{14}$$

where λ_{un}, λ_{sp}, and λ_{policy} are the weights of the corresponding loss terms.

For each candidate channel $i \in \mathcal{U}_t$, the time-domain reconstruction error is defined as

$$e_i^{(t)} = \frac{1}{T}\left\|\widehat{\mathbf{X}}_t(i,:) - \mathbf{X}_H(i,:)\right\|_2^2 \tag{15}$$

and the log-power spectrum is defined as

$$\mathcal{P}(x) = \log\left(1 + |\text{STFT}(x)|^2\right) \tag{16}$$

where T denotes the number of temporal samples. The reconstruction loss, uncertainty loss, and frequency-domain consistency loss are respectively defined as

$$\mathcal{L}_{\text{re}} = \frac{1}{G}\sum_{t=1}^{G}\frac{1}{|\mathcal{S}_t|}\sum_{i\in\mathcal{S}_t} e_i^{(t)} \tag{17}$$

$$\mathcal{L}_{\text{un}} = \frac{1}{G}\sum_{t=1}^{G}\frac{1}{|\mathcal{S}_t|}\sum_{i\in\mathcal{S}_t}\left[e_i^{(t)}\exp\left(-\Omega_{t+1}(i)\right) + \Omega_{t+1}(i)\right] \tag{18}$$

$$\mathcal{L}_{\text{sp}} = \frac{1}{G}\sum_{t=1}^{G}\frac{1}{|\mathcal{S}_t|}\sum_{i\in\mathcal{S}_t}\left\|\mathcal{P}\left(\widehat{\mathbf{X}}_t(i,:)\right) - \mathcal{P}\left(\mathbf{X}_H(i,:)\right)\right\|_1 \tag{19}$$

Here, $\Omega_{t+1}(i)$ denotes the predicted log-variance of channel i estimated from the prediction at the t-th recovery step. The above losses constrain the time-domain reconstruction error, prediction uncertainty, and spectral consistency, respectively.

Since the Learnable Recovery Policy employs a non-differentiable hard Top-K operation, we define the soft selection probability and policy auxiliary loss over the candidate set $\mathcal{U}_t$ as

$$\pi_i^{(t)} = \text{Softmax}_{i\in\mathcal{U}_t}\left(\frac{s_i^{(t)}}{\tau_p}\right) \tag{20}$$

$$\mathcal{L}_{\text{policy}} = \frac{1}{G}\sum_{t=1}^{G}\sum_{i\in\mathcal{U}_t}\pi_i^{(t)}\,\text{stopgrad}\left(e_i^{(t)}\right) \tag{21}$$

Here, τ_p is the temperature coefficient, and $\text{stopgrad}(\cdot)$ denotes the stop-gradient operation. This loss encourages the scoring network to prioritize channels with lower current prediction errors,

thereby mitigating error propagation. The stop-gradient operation prevents the prediction-error term from backpropagating into the reconstruction network.

To alleviate the inconsistency in historical context between training and inference, for each channel $i \in \mathcal{S}_t$ selected at the t-th step, state write-back is uniformly expressed as

$$\mathbf{X}^{(t)}(i,:) = \alpha_t \mathbf{X}_H(i,:) + (1-\alpha_t)\,\widehat{\mathbf{X}}_t(i,:) \tag{22}$$

$$\alpha_t = \begin{cases} 1, & \text{teacher forcing} \\ 0, & \text{self rollout} \\ \mathrm{Bernoulli}(\rho_e), & \text{scheduled sampling} \end{cases} \tag{23}$$

Here, ρ_e denotes the probability of using the ground-truth signal for write-back at the e-th training epoch. The originally observed channels always retain their ground-truth inputs, while the unselected missing channels maintain their states from the previous recovery step. During inference, $\alpha_t = 0$ is used uniformly, corresponding to self rollout. Channel reliability is used only for subsequent topology adjustment and channel selection, and does not directly affect the predicted EEG waveforms.

4 Experiments

4.1 Experimental Setup

4.1.1 Datasets

We evaluate LBTAR on three publicly available EEG datasets: SEED, SEED-IV, and Localize-MI. SEED Zheng & Lu (2015) and SEED-IV Zheng et al. (2019) are emotion EEG datasets with 62 channels, while Localize-MI Mikulan et al. (2020) is a high-density EEG dataset with 256 channels. For each dataset, the complete-channel signals are treated as the high-density reference, a predefined subset of channels is used as the low-density input, and the remaining channels are regarded as recovery targets. Further details regarding the datasets and preprocessing procedures are provided in Appendix E.

4.1.2 Experimental Settings

Following the evaluation protocols of ESTformer and CAFE, we conduct experiments under predefined low-density (LD) channel layouts. For SEED and SEED-IV, three spatial super-resolution ratios of $2\times$, $4\times$, and $8\times$ are considered, retaining 31, 16, and 8 observed channels, respectively; for Localize-MI, four ratios of $2\times$, $4\times$, $8\times$, and $16\times$ are considered, retaining 128, 64, 32, and 16 observed channels, respectively. Four LD layouts are adopted for each ratio to cover different scalp regions. All methods use consistent data splits, temporal windows, normalization procedures, channel layouts, and training and testing conditions.

The evaluation metrics include normalized mean squared error (NMSE), Pearson correlation coefficient (PCC), signal-to-noise ratio (SNR), and spectral mean absolute error (Spec-MAE), all of which are computed only over the channels to be recovered. We also report the number of parameters, giga floating-point operations (GFLOPs), and the average inference time per sample to measure the storage, computational, and deployment overheads of the models. Unless otherwise specified, all results are averaged over the four LD layouts and all test subjects.

Further details regarding the implementation, stage-wise recovery settings, training configuration, and evaluation metrics are provided in Appendix F.

4.2 Overall Reconstruction Performance

To provide a comprehensive comparison across different technical paradigms, we select five representative EEG spatial super-resolution methods as baselines: the spatial interpolation method SIPerrin et al. (1989), the generative method EEGSR-GANCorley & Huang (2018), the Transformer-based end-to-end method ESTformerLi et al. (2025), the diffusion-based method

SRGDiffLiu et al. (2025), and the fixed-order autoregressive method CAFELiu et al. (2026). Table 1 reports the reconstruction results on SEED, SEED-IV, and Localize-MI, with all methods evaluated under identical data splits, low-density channel layouts, input window lengths, normalization procedures, and evaluation metrics. As the spatial super-resolution factor increases, the number of observed channels decreases and the performance of all methods generally degrades, indicating the increased difficulty caused by higher channel-missing ratios. LBTAR achieves the best overall performance across the three datasets and different super-resolution factors, while still maintaining a low NMSE of 0.07 and a high PCC of 0.96 under the $16\times$ setting on Localize-MI. These results demonstrate that, through dynamic channel selection, iterative state updating, and guidance from brain topology and prediction reliability, LBTAR can more effectively exploit progressively expanded contextual information and maintain stable reconstruction performance across different channel scales and missing ratios.

To evaluate the frequency-domain consistency of the reconstructed EEG, we adopt Spec-MAE to measure the mean absolute error between the log-power spectra of the reconstructed signals and the ground-truth high-density EEG, where a lower value indicates more accurate spectral reconstruction. As reported in Table 3 in Appendix C, LBTAR achieves the lowest Spec-MAE across all super-resolution ratios on the three datasets and maintains a value of 0.16 under the $16\times$ setting on Localize-MI, indicating that it effectively preserves spectral consistency even under a high proportion of missing channels.

Additional qualitative waveform comparisons are provided in Fig. 3 in Appendix B. Specifically, Fig. 3 compares LBTAR and CAFE, the two methods with the best overall quantitative performance, on a representative missing channel. Both methods recover the overall temporal trend of the ground-truth signal, while LBTAR more closely matches the ground-truth EEG in terms of amplitude variations, local peaks, and dynamic fluctuation patterns. These qualitative observations are consistent with the quantitative results in Table 1, further supporting the temporal reconstruction capability of LBTAR.

Table 1: Reconstruction performance of different EEG spatial super-resolution methods on the SEED, SEED-IV, and Localize-MI datasets. SNR is reported in dB. The best results are shown in bold.

Model	Metric	SEED			SEED-IV			Localize-MI			
		$2\times$	$4\times$	$8\times$	$2\times$	$4\times$	$8\times$	$2\times$	$4\times$	$8\times$	$16\times$
LBTAR	NMSE ↓	**0.11**	**0.20**	**0.33**	**0.09**	**0.16**	**0.25**	**0.02**	**0.04**	**0.06**	**0.07**
	PCC ↑	**0.94**	**0.89**	**0.81**	**0.95**	**0.91**	**0.86**	**0.99**	**0.98**	**0.97**	**0.96**
	SNR ↑	**9.25**	**6.80**	**4.72**	**10.31**	**7.85**	**5.95**	**15.89**	**13.42**	**12.35**	**11.39**
CAFE	NMSE ↓	0.22	0.29	0.42	0.22	0.27	0.40	0.07	0.08	0.10	0.13
	PCC ↑	0.88	0.84	0.76	0.88	0.85	0.78	0.96	0.95	0.94	0.93
	SNR ↑	6.40	5.43	3.75	6.50	5.62	4.05	11.52	10.81	9.82	8.81
SRGDiff	NMSE ↓	0.31	0.46	0.67	0.32	0.43	0.62	0.19	0.20	0.27	0.36
	PCC ↑	0.84	0.79	0.72	0.83	0.77	0.71	0.90	0.89	0.86	0.80
	SNR ↑	5.10	3.38	1.70	4.98	3.62	2.09	7.27	6.93	5.65	4.41
ESTformer	NMSE ↓	0.42	0.46	0.54	0.38	0.45	0.54	0.13	0.15	0.17	0.23
	PCC ↑	0.76	0.74	0.68	0.79	0.75	0.69	0.96	0.94	0.92	0.88
	SNR ↑	3.77	3.38	2.63	4.22	3.46	2.66	8.81	8.16	7.59	6.37
EEGSR-GAN	NMSE ↓	0.50	0.56	0.63	0.48	0.64	0.74	0.15	0.34	0.41	0.46
	PCC ↑	0.72	0.66	0.60	0.75	0.61	0.57	0.93	0.81	0.78	0.74
	SNR ↑	3.01	2.51	1.95	3.18	1.91	1.31	8.25	4.62	3.79	3.35
SI	NMSE ↓	0.35	0.51	0.69	0.33	0.48	0.69	0.09	0.10	0.13	0.22
	PCC ↑	0.81	0.74	0.67	0.82	0.76	0.68	0.95	0.95	0.93	0.89
	SNR ↑	4.60	2.93	1.61	4.86	3.18	1.59	10.70	9.92	8.86	6.66

4.3 Ablation Study

To evaluate the contributions of LBTAR's core components and different topology configurations, we conduct ablation experiments on the SEED dataset under the $2\times$, $4\times$, and $8\times$ settings, as shown

in Table 2. The analysis is conducted from two perspectives: recovery mechanisms, and structural priors with reliability modeling. The complete model is denoted as Full LBTAR, while detailed configurations of all ablation variants are provided in Appendix G.

In terms of recovery mechanisms, both Random Fixed Order and Fixed Geo-Distance outperform One-shot Recovery, yet remain inferior to Full LBTAR. This result suggests that sequential recovery benefits from the progressively expanded visible context, while state-dependent dynamic channel selection provides further advantages over random or predefined fixed recovery orders. Regarding structural priors and reliability modeling, removing either Brain Topology or uncertainty modeling consistently degrades reconstruction performance. Spatial-only Topology outperforms the variant without brain topology but remains overall inferior to Full LBTAR, indicating that electrode spatial relationships provide useful structural priors, while channel-to-region assignments and inter-regional adjacency contribute additional complementary information. Shuffled Region Topology also performs worse than the full model, suggesting that the observed gains are associated with meaningful brain-region structural information rather than merely introducing additional graph connections. The performance degradation of w/o Uncertainty further demonstrates the effectiveness of incorporating prediction reliability into sequential recovery.

Table 2: Ablation study of LBTAR on the SEED dataset under different spatial super-resolution factors. SNR is reported in dB. The best results are shown in bold.

Variant	2×			4×			8×		
	NMSE ↓	PCC ↑	SNR ↑	NMSE ↓	PCC ↑	SNR ↑	NMSE ↓	PCC ↑	SNR ↑
One-shot Recovery	0.26	0.86	5.81	0.37	0.79	4.33	0.47	0.72	3.24
Random Fixed Order	0.23	0.88	6.62	0.34	0.81	4.57	0.45	0.74	3.46
w/o Brain Topology	0.16	0.92	8.05	0.26	0.86	5.90	0.39	0.78	4.08
Fixed Geo-Distance	0.17	0.91	7.73	0.28	0.85	5.55	0.40	0.78	4.01
w/o Uncertainty	0.15	0.91	8.09	0.27	0.85	5.60	0.37	0.79	4.32
Spatial-only Topology	0.14	0.93	8.32	0.24	0.87	6.13	0.34	0.81	4.65
Shuffled Region Topology	0.16	0.91	7.80	0.27	0.86	5.70	0.37	0.79	4.33
Full LBTAR	**0.11**	**0.94**	**9.25**	**0.20**	**0.89**	**6.80**	**0.33**	**0.81**	**4.72**

4.4 Efficiency Comparison

As reported in Table 4 in Appendix D, LBTAR contains only 0.701 M parameters, requires 1.021 GFLOPs, and achieves a single-sample inference time of 2.013 ms under the same hardware configuration and batch size of 1. These values are lower than those of CAFE, SRGDiff, and ESTformer, showing that LBTAR maintains strong reconstruction performance with a compact model size and low computational overhead.

5 Conclusion

We propose LBTAR, a brain-topology-guided learnable autoregressive channel recovery method that formulates EEG spatial super-resolution as a state-dependent dynamic sequential recovery process. By jointly incorporating brain-region structural priors, the recovery state, and channel-level prediction reliability, LBTAR adaptively selects channels for recovery and progressively reconstructs high-density EEG through iterative write-back. Experiments on three public datasets show that LBTAR achieves strong temporal and spectral reconstruction performance across different super-resolution ratios and remains robust under severe channel sparsity. Ablation studies further validate the effectiveness of sequential recovery, state-dependent dynamic channel selection, brain-topology priors, and uncertainty modeling, while LBTAR also maintains a low parameter count, computational cost, and inference overhead.

AI use statement

In this work, we used generative AI tools to assist with language polishing, LaTeX formatting, and improving the clarity and organization of the manuscript. Generative AI tools were not used to

generate or modify experimental data or results, nor were they used to independently determine the scientific conclusions of this work. All AI-assisted content was reviewed and revised by the authors, and all experimental results, analyses, and technical claims were independently verified by the authors. The authors take full responsibility for the final content of this work, including any text or other artifacts produced with the assistance of generative AI.

REFERENCES

Rupal Agarwal, Marvin Andujar, and Shaun Canavan. Classification of emotions using eeg activity associated with different areas of the brain. *Pattern Recognition Letters*, 162:71–80, 2022.

IA Corley and Y Huang. Deep eeg super-resolution: Upsampling eeg spatial resolution with generative adversarial networks, ieee embs international conference on biomedical and health informatics (bhi), 2018.

Li Dong, Lingling Zhao, Yufan Zhang, Xue Yu, Fali Li, Jianfu Li, Yongxiu Lai, Tiejun Liu, and Dezhong Yao. Reference electrode standardization interpolation technique (resit): a novel interpolation method for scalp eeg. *Brain Topography*, 34(4):403–414, 2021.

Wei-Bang Jiang, Yansen Wang, Bao-Liang Lu, and Dongsheng Li. Neurolm: A universal multitask foundation model for bridging the gap between language and eeg signals. In *International conference on learning representations*, volume 2025, pp. 55436–55457, 2025.

Dongdong Li, Zhongliang Zeng, Zhe Wang, and Hai Yang. Estformer: Transformer utilising spatiotemporal dependencies for electroencephalogram super-resolution. *Knowledge-Based Systems*, 317:113345, 2025.

Hongjun Liu, Leyu Zhou, Zijianghao Yang, and Chao Yao. Step-aware residual-guided diffusion for eeg spatial super-resolution. *arXiv preprint arXiv:2510.19166*, 2025.

Hongjun Liu, Leyu Zhou, Zijianghao Yang, Rujun Han, Shitong Duan, Kuanjian Tang, and Chao Yao. Cafe: Channel-autoregressive factorized encoding for robust biosignal spatial super-resolution. *arXiv preprint arXiv:2602.17011*, 2026.

Ezequiel Mikulan, Simone Russo, Sara Parmigiani, Simone Sarasso, Flavia Maria Zauli, Annalisa Rubino, Pietro Avanzini, Anna Cattani, Alberto Sorrentino, Steve Gibbs, et al. Simultaneous human intracerebral stimulation and hd-eeg, ground-truth for source localization methods. *Scientific data*, 7(1):127, 2020.

François Perrin, Jacques Pernier, Olivier Bertrand, and Jean Francois Echallier. Spherical splines for scalp potential and current density mapping. *Electroencephalography and clinical neurophysiology*, 72(2):184–187, 1989.

Yannick Roy, Hubert Banville, Isabela Albuquerque, Alexandre Gramfort, Tiago H Falk, and Jocelyn Faubert. Deep learning-based electroencephalography analysis: a systematic review. *Journal of neural engineering*, 16(5):051001, 2019.

Shunuo Shang, Yingqian Shi, Yajie Zhang, Mengxue Liu, Hong Zhang, Ping Wang, and Liujing Zhuang. Artificial intelligence for brain disease diagnosis using electroencephalogram signals. *Journal of Zhejiang University-SCIENCE B*, 25(10):914–940, 2024.

Shuqiang Wang, Tong Zhou, Yanyan Shen, Ye Li, Guoheng Huang, and Yong Hu. Generative ai enables eeg super-resolution via spatio-temporal adaptive diffusion learning. *IEEE Transactions on Consumer Electronics*, 71(1):1034–1045, 2025.

Kang Yin, Elissa Yanting Lim, and Seong-Whan Lee. Gitgan: Generative inter-subject transfer for eeg motor imagery analysis. *Pattern Recognition*, 146:110015, 2024.

Wei-Long Zheng and Bao-Liang Lu. Investigating critical frequency bands and channels for eeg-based emotion recognition with deep neural networks. *IEEE Transactions on autonomous mental development*, 7(3):162–175, 2015.

Wei-Long Zheng, Wei Liu, Yifei Lu, Bao-Liang Lu, Andrzej Cichocki, et al. Emotionmeter: A multimodal framework for recognizing human emotions. *IEEE Trans. Cybern.*, 49(3):1110–1122, 2019.

Yijin Zhou, Fu Li, Yang Li, Youshuo Ji, Guangming Shi, Wenming Zheng, Lijian Zhang, Yuanfang Chen, and Rui Cheng. Progressive graph convolution network for eeg emotion recognition. *Neurocomputing*, 544:126262, 2023.

A Comparison of EEG Spatial Super-Resolution Paradigms

Fig. 2 illustrates the differences among one-shot reconstruction, fixed-order autoregressive recovery, and the proposed state-dependent dynamic recovery paradigm.

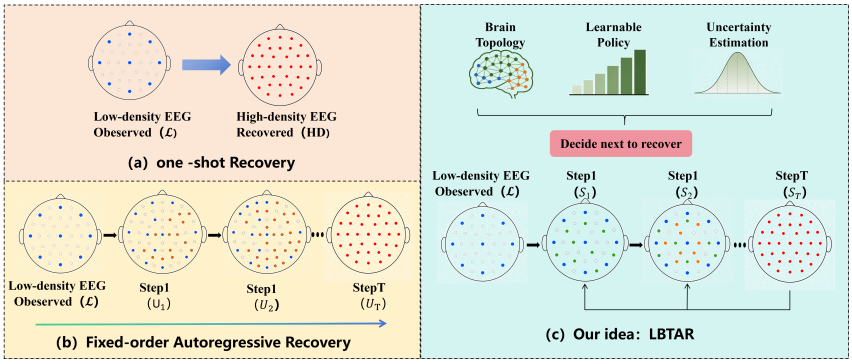


Figure 2: Comparison of EEG spatial super-resolution paradigms: (a) one-shot recovery, (b) fixed-order autoregressive recovery, and (c) LBTAR dynamic recovery.

B Qualitative EEG Reconstruction Results

Fig. 3 presents a qualitative comparison between LBTAR and CAFE on a representative missing channel under the 2× spatial super-resolution setting.

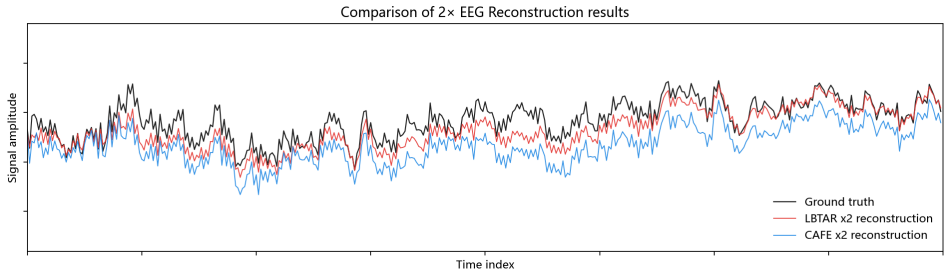


Figure 3: Qualitative comparison of EEG waveform reconstruction results between LBTAR and CAFE under the 2× spatial super-resolution setting.

C Spectral Reconstruction Performance

Table 3 reports the frequency-domain reconstruction performance of different methods under different spatial super-resolution ratios.

Table 3: Comparison of Spec-MAE under different spatial super-resolution factors on the SEED, SEED-IV, and Localize-MI datasets. Lower values indicate better performance.

Model	SEED			SEED-IV			Localize-MI			
	2×	4×	8×	2×	4×	8×	2×	4×	8×	16×
LBTAR (Ours)	**0.20**	**0.32**	**0.45**	**0.17**	**0.28**	**0.39**	**0.05**	**0.11**	**0.13**	**0.16**
CAFE	0.39	0.44	0.53	0.40	0.45	0.54	0.17	0.20	0.24	0.27
SRGDiff	0.47	0.75	0.93	0.50	0.59	0.85	0.31	0.33	0.39	0.47
ESTformer	0.53	0.57	0.63	0.54	0.61	0.60	0.22	0.24	0.29	0.35
EEGSR-GAN	0.60	0.88	1.12	0.62	0.85	1.06	0.38	0.53	0.62	0.66
SI	0.47	0.59	0.71	0.50	0.75	0.82	0.20	0.21	0.24	0.29

D Computational Efficiency

Table 4 compares the computational efficiency of LBTAR with representative EEG spatial super-resolution methods in terms of parameter count, GFLOPs, and inference time.

Table 4: Computational efficiency comparison of different EEG spatial super-resolution methods. Lower values indicate better efficiency.

Method	Params (M)	GFLOPs	Runtime (ms)
LBTAR (Ours)	**0.701**	**1.021**	**2.013**
CAFE	2.342	3.238	8.204
SRGDiff	20.421	5.439	92.541
ESTformer	12.112	4.302	5.039

E Dataset and Preprocessing Details

E.1 SEED

SEED is a publicly available affective EEG dataset released by the BCMI Laboratory at Shanghai Jiao Tong University. It contains EEG recordings from 15 participants. Positive, neutral, and negative emotional states were elicited using emotional film clips, and brain activity was recorded using a 62-channel EEG system. In this study, the complete 62-channel signals are treated as the high-density EEG reference. A subset of channels is retained according to predefined low-density channel layouts as model input, while the remaining unobserved channels are used as recovery targets.

Dataset homepage: `https://bcmi.sjtu.edu.cn/home/seed/seed.html`

Data application portal: `https://bcmi.sjtu.edu.cn/ApplicationForm/apply_form/`

Access to the SEED dataset requires applicants to complete and submit a data-use agreement. After approval, the dataset maintainers provide the download link and access password.

E.2 SEED-IV

SEED-IV is an extension of the SEED dataset and contains data from 15 participants. Each participant completed three experimental sessions on different days, with 24 trials in each session. A 62-channel EEG system was used to record brain activity elicited by four emotional categories: happiness, sadness, fear, and neutrality. In this study, the complete 62-channel signals are treated as the high-density EEG reference. Low-density inputs under different spatial super-resolution ratios are constructed according to predefined channel layouts, and the unobserved channels are used as recovery targets.

Dataset homepage: `https://bcmi.sjtu.edu.cn/home/seed/seed-iv.html`

Data application portal: https://bcmi.sjtu.edu.cn/ApplicationForm/apply_form/

SEED-IV follows the same data-access procedure as SEED. After approval, the dataset maintainers provide the corresponding download link and access password.

E.3 LOCALIZE-MI

Localize-MI is a publicly available dataset containing synchronized intracranial electrical stimulation and high-density EEG recordings. The data were collected during the presurgical evaluation of patients with drug-resistant epilepsy. Single-pulse electrical stimulation was delivered through implanted electrodes while 256-channel high-density EEG signals were recorded simultaneously. Owing to its dense spatial sampling, the dataset provides a suitable reference for EEG spatial super-resolution under high channel-density conditions. In this study, the complete 256-channel signals are treated as the high-density EEG reference and are used to evaluate reconstruction performance in a large channel space and under high missing-channel ratios.

Dataset access: https://doi.gin.g-node.org/10.12751/g-node.1cc1ae/

The repository provides online data browsing and a complete archive download. The dataset is distributed under the CC BY-NC-SA 4.0 license.

E.4 DATA PREPROCESSING AND SAMPLE CONSTRUCTION

For SEED and SEED-IV, the EEG signals were first downsampled from the original sampling rate of 1000 Hz to 200 Hz and band-pass filtered between 0.5 and 75 Hz. Trials containing sensor failures, corrupted signals, or invalid recordings were subsequently removed through manual inspection. The processed continuous EEG signals were then divided into non-overlapping 4-s segments. At a sampling rate of 200 Hz, each sample contains 800 temporal points and has a size of 62×800.

For Localize-MI, the original 256-channel EEG signals were recorded at 8000 Hz. Preprocessing included 0.1-Hz high-pass filtering and notch filtering at 50, 100, 150, and 200 Hz to suppress power-line interference and its harmonics. Abnormal channels and trials were then removed, and the trials were temporally aligned according to the peak of the stimulation artifact. Finally, segments from 250 ms before stimulation to 10 ms after stimulation were extracted. Each segment spans 260 ms and has a size of 256×2081.

For all three datasets, the processed samples were randomly divided into training and test sets at a ratio of 80% to 20%, and 10% of the training set was further reserved as the validation set. The validation set was used for hyperparameter selection and early stopping, whereas the test set was used only for final performance evaluation after model training.

In the spatial super-resolution task, the complete-channel signals are treated as the high-density EEG reference. According to the target spatial super-resolution ratio, only the channels specified by a predefined layout are retained as low-density observations, while the remaining channels serve as recovery targets, thereby forming temporally aligned low-density/high-density EEG pairs. For SEED and SEED-IV, experiments are conducted at $2\times$, $4\times$, and $8\times$ spatial super-resolution ratios. For Localize-MI, experiments are conducted at $2\times$, $4\times$, $8\times$, and $16\times$.

Table 5: Summary of datasets and preprocessing settings.

Dataset	Recording setup	Main preprocessing	Window and sample size	SR ratios
SEED	62 channels; 1000 Hz	Downsampling to 200 Hz; 0.5–75 Hz band-pass filtering; abnormal-trial removal	4 s; 62×800	$2\times$, $4\times$, $8\times$
SEED-IV	62 channels; 1000 Hz	Downsampling to 200 Hz; 0.5–75 Hz band-pass filtering; abnormal-trial removal	4 s; 62×800	$2\times$, $4\times$, $8\times$
Localize-MI	256 channels; 8000 Hz	0.1 Hz high-pass filtering; 50/100/150/200 Hz notch filtering; abnormal-channel and trial removal; stimulus alignment	260 ms; 256×2081	$2\times$, $4\times$, $8\times$, $16\times$

F Additional Implementation Details

F.1 Stage-Wise Recovery Settings

At recovery stage t, LBTAR scores the unrecovered channels according to the currently visible channels, the remaining channels, brain-topology relations, channel-level uncertainty, and current signal representations, and then selects the channels with the highest scores for reconstruction.

The stage-wise recovery schedules for SEED and SEED-IV under different spatial super-resolution ratios are summarized in Table 6.

Table 6: Stage-wise recovery settings for SEED and SEED-IV.

SR Ratio	Observed	Missing	Recovered per Stage
$2\times$	31	31	5, 10, 16
$4\times$	16	46	7, 16, 23
$8\times$	8	54	15, 18, 21

The numbers of channels recovered in the three stages sum to the total number of missing channels under the corresponding ratio.

F.2 Optimization, Training, and Reproducibility Settings

The model is trained using the AdamW optimizer with an initial learning rate of 2×10^{-3}, a weight decay of 1×10^{-4}, and a batch size of 256. A cosine annealing learning-rate schedule is adopted, and the gradient clipping threshold is set to 1.0. The random seed is fixed to 42.

All experiments are conducted on a workstation equipped with an NVIDIA GeForce RTX 5090 GPU with 32 GB of memory, an AMD Ryzen 9 9950X CPU with 16 cores and 32 threads, and 128 GB of system memory. The models are implemented in Python 3.10 and PyTorch 2.11 with CUDA 12.8 and cuDNN 9.19, and all experiments use FP32 numerical precision.

Table 7: Optimization and training settings of LBTAR.

Configuration	Setting
Optimizer	AdamW
Initial learning rate	2×10^{-3}
Weight decay	1×10^{-4}
Batch size	256
Learning-rate scheduler	Cosine annealing
Gradient clipping threshold	1.0
Random seed	42
Training device	RTX 5090 (32 GB)

F.3 Evaluation Metrics

Let $\mathcal{U}$ denote the set of missing channels, $|\mathcal{U}|$ its cardinality, and T the signal length. The ground-truth and reconstructed signals over the missing channels are denoted by $X_{\mathcal{U}}$ and $\hat{X}_{\mathcal{U}}$, respectively. All reconstruction metrics are calculated only over missing channels.

F.3.1 Normalized Mean Squared Error

$$\mathrm{NMSE} = \frac{\left\| X_{\mathcal{U}} - \hat{X}_{\mathcal{U}} \right\|_F^2}{\|X_{\mathcal{U}}\|_F^2 + \epsilon} \tag{24}$$

A lower NMSE indicates a smaller relative reconstruction error.

F.3.2 Pearson Correlation Coefficient

For the i-th missing channel,

$$\mathrm{PCC}_i = \frac{\sum_{t=1}^{T}(X_{i,t} - \bar{X}_i)(\hat{X}_{i,t} - \bar{\hat{X}}_i)}{\sqrt{\sum_{t=1}^{T}(X_{i,t} - \bar{X}_i)^2}\sqrt{\sum_{t=1}^{T}(\hat{X}_{i,t} - \bar{\hat{X}}_i)^2} + \epsilon} \tag{25}$$

The overall PCC is

$$\mathrm{PCC} = \frac{1}{|\mathcal{U}|}\sum_{i \in \mathcal{U}} \mathrm{PCC}_i \tag{26}$$

F.3.3 Signal-to-Noise Ratio

$$\mathrm{SNR} = 10 \log_{10} \frac{\|X_{\mathcal{U}}\|_F^2}{\left\| X_{\mathcal{U}} - \hat{X}_{\mathcal{U}} \right\|_F^2 + \epsilon} \tag{27}$$

A higher SNR indicates better reconstruction quality.

F.3.4 Spectral Mean Absolute Error

Spec-MAE measures the mean absolute discrepancy between the log-power spectra of the reconstructed and ground-truth EEG signals over the missing channels. A lower Spec-MAE indicates better frequency-domain reconstruction.

G Detailed Ablation Settings

To facilitate the interpretation of the ablation results in Table 2, we provide detailed descriptions of the evaluated variants below.

One-shot Recovery. This variant removes the sequential recovery process and reconstructs all missing channels in a single forward pass. Intermediate predictions are not progressively written back into the current observation, and therefore previously recovered channels cannot be exploited as additional context for subsequent reconstruction.

Random Fixed Order. This variant retains the sequential recovery framework but replaces the learnable channel-selection policy with a random recovery strategy. Channels are progressively recovered and written back into the current visible state.

Fixed Geo-Distance. This variant replaces the state-dependent learnable channel-selection mechanism with a predefined recovery order determined by electrode geometric distances.

w/o Brain Topology. This variant removes the brain-topology prior constructed from electrode spatial relationships, channel-to-region assignments, and inter-regional adjacency, while retaining sequential recovery, learnable channel selection, and the remaining state information.

w/o Uncertainty. This variant removes channel-level uncertainty estimation and its feedback to subsequent topology adjustment and channel selection.

Spatial-only Topology. This variant retains only the spatial geometric relationships among electrodes while removing channel-to-region assignments and inter-regional adjacency information.

Shuffled Region Topology. This variant preserves the original electrode spatial locations but randomly shuffles channel-to-region assignments, thereby disrupting meaningful brain-region correspondence while preserving the overall topology-construction procedure.

Full LBTAR. The complete model jointly incorporates sequential recovery, state-dependent learnable channel selection, brain-topology guidance, and channel-level uncertainty feedback.