SUPPORT-CONDITIONED NEURAL FIELDS FOR EEG SPATIAL SUPER-RESOLUTION UNDER VARIABLE ELECTRODE TOPOLOGIES

Anonymous authors
Paper under double-blind review

ABSTRACT

EEG spatial super-resolution must reconstruct dense scalp signals from electrode sets whose number and spatial arrangement may vary across recordings. Existing approaches either rely on prescribed spatial interpolation, assume predefined channel layouts, or achieve coordinate flexibility without explicitly adapting the reconstruction function to the temporal and spectral characteristics of the current recording. We propose ScalpINR, an amortized support-conditioned neural field that directly instantiates a recording-specific coordinate-to-signal function from an unordered set of observed EEG signals and electrode coordinates. The observed support generates recording-dependent temporal and spectral controls for a shared coordinate decoder, while a dual-support objective encourages compatible fields to be inferred from different electrode subsets of the same recording. Once conditioned on the available electrodes, ScalpINR can be queried at missing or training-unseen electrode coordinates without test-time optimization. Experiments on SEED, ESAA, and BCI2000 demonstrate strong performance under both random missing-channel reconstruction and strict held-out-electrode evaluation. On ESAA, ScalpINR reduces NMSE by 37.5% relative to the strongest baseline when all non-held-out electrodes are available as support.

1 INTRODUCTION

EEG spatial super-resolution (EEGSR) reconstructs dense scalp signals from sparse electrode measurements, enabling high-density spatial analysis with substantially fewer sensors (Lawhern et al., 2018; Nejedly et al., 2019; Jaramillo-Jimenez et al., 2023; Wang et al., 2025). Unlike conventional super-resolution, however, reducing EEG electrodes does not merely lower the number of measurements; it changes where the underlying neural activity is observed. In practical acquisition, contact failures, channel-quality fluctuations, and device configurations further cause the number and spatial arrangement of reliable electrodes to vary across recordings (Leach et al., 2020; Moumane et al., 2024). EEGSR must therefore recover spatially dense signals from changing observation topologies, rather than from a single predefined low-density montage.

Existing approaches to EEG spatial reconstruction provide different forms of flexibility, but do not simultaneously address variable support topology, coordinate-queryable generation, and recording-specific signal dynamics. Classical interpolation, such as spherical splines, can operate on changing electrode locations, but its predictions are governed by prescribed spatial smoothness rather than learned EEG structure. Data-driven EEG super-resolution models, including Transformer- and diffusion-based approaches, learn richer spatiotemporal dependencies but are commonly optimized over predefined channel inventories or target layouts (Li et al., 2025a; Wang et al., 2025; Liu et al., 2026). More general set-conditioned predictors and conditional implicit neural representations can instead map an unordered observation set to arbitrary queries (Garnelo et al., 2018; Kim et al., 2019; Li et al., 2025b), showing that set conditioning and coordinate querying alone are not sufficient to characterize an EEG-specific solution. Recent EEG models further narrow this gap: Neural Brain Fields represents an individual recording as a continuous spatiotemporal field but fits the field separately for each recording, whereas ZUNA supports flexible reconstruction across arbitrary channel subsets and scalp locations through a large diffusion model (Kedem et al., 2025; Warner et al., 2026). What remains less explored is how a shared model can directly infer a recording-specific EEG field

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.

from the currently observed electrodes while adapting the field's temporal and spectral computation to those observations, rather than merely conditioning a fixed predictor on a support embedding.

These observations lead to two requirements for flexible EEG neural-field reconstruction. First, the available electrodes should determine not only *what* information conditions the prediction, but also *how* the inferred field processes the current recording. Because EEG recordings differ in temporal dynamics and spectral composition, a fixed decoder conditioned only through a latent vector may be insufficient when the support set changes substantially. We therefore seek a field whose temporal receptive-field mixture, spectral response, and temporal–spectral balance can be adapted from the observed EEG itself. Second, this recording-specific field should be stable with respect to the particular subset through which the recording is observed. Two support sets sampled from the same underlying EEG segment provide different amounts and spatial configurations of evidence, but should induce compatible predictions at common query locations. We term this requirement *support-topology consistency*, where support topology denotes both the number and spatial arrangement of the observed electrodes.

We realize these principles with ScalpINR, a support-conditioned neural field that directly maps an arbitrary set of observed EEG signals and electrode coordinates to a recording-specific coordinate-to-signal function. ScalpINR generates support-dependent temporal and spectral controls for a shared coordinate decoder, allowing its computation to adapt to the current recording. A dual-support training objective further encourages different electrode subsets of the same recording to induce compatible fields. At inference, the available electrodes are encoded once, after which complete waveforms can be queried at desired electrode coordinates, including training-held-out locations, without test-time optimization.

Our contributions are summarized as follows:

- We identify two requirements for EEG spatial super-resolution under variable electrode availability: the inferred field should adapt its functional behavior to the observed EEG and remain consistent across different support topologies of the same recording.
- We propose ScalpINR, an amortized support-conditioned EEG neural field that generates recording-dependent temporal and spectral controls for a shared coordinate decoder, together with a dual-support objective that promotes support-topology consistency without requiring test-time optimization for each recording.
- We evaluate ScalpINR under both random missing-channel reconstruction and a strict unseen-electrode protocol in which held-out electrodes are excluded from training entirely. Experiments across three EEG datasets demonstrate robust reconstruction under sparse and variable support and direct waveform generation at electrode locations never used during training.

2 METHOD

ScalpINR addresses EEG spatial super-resolution when the available electrodes vary in number and spatial arrangement. Given the currently observed support set, it infers a recording-specific neural field whose computation is adapted by the observed EEG itself. The support information enters the decoder at two complementary levels: time-resolved modulation injects local waveform context, while support-generated controls adapt its temporal-scale mixing, spectral response, and temporal–spectral fusion.

Given an observed support set $\mathcal{S} = \{(p_i, x_i)\}_{i=1}^{N_s}$, where p_i is the 3D electrode coordinate and $x_i \in \mathbb{R}^T$ is its observed EEG waveform, ScalpINR instantiates a recording-specific function $f_{\mathcal{S}} \equiv f_\theta(\cdot \mid \mathcal{S})$. Here, all recordings share the model parameters θ; different support sets instantiate different conditioned fields through a forward pass rather than through separately optimized networks. A target coordinate p_q can then query the conditioned field to generate its complete waveform, $f_{\mathcal{S}}(p_q) = \hat{\mathbf{x}}_q \in \mathbb{R}^T$. Neither the support cardinality nor the queried electrode locations are fixed by the model architecture.

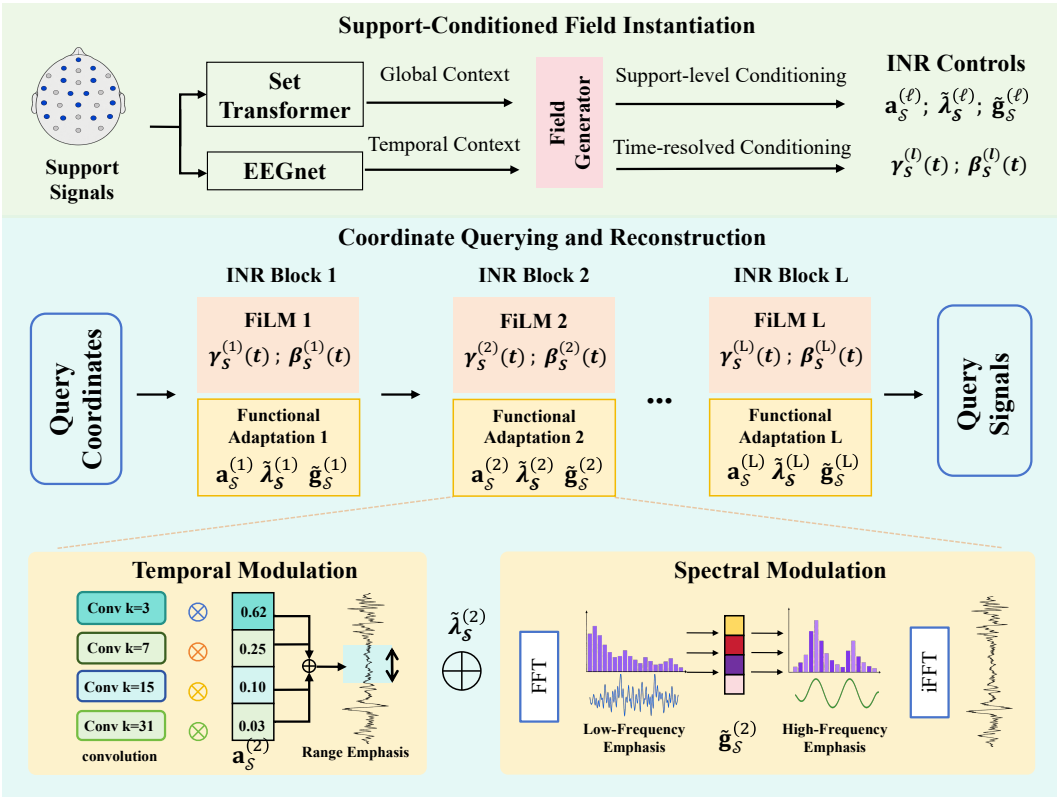


Figure 1: **Overview of ScalpINR.** The framework consists of two stages. **Top: support-conditioned field instantiation. Bottom: coordinate querying and reconstruction.**

2.1 Position-guided support-set encoder

The support-set encoder transforms the available electrodes into the conditioning information required to instantiate a recording-specific EEG field. It produces a global representation $z_{\mathcal{S}}$ that summarizes the observed support set, together with a time-resolved context $c_{\mathcal{S}}(1:T)$ that retains support-derived information at each sample time. The unpooled context $c_{\mathcal{S}}(t)$, together with $z_{\mathcal{S}}$, generates time-dependent FiLM parameters for every decoder layer, whereas its temporal mean is used to generate compact layer-wise controls over temporal-scale mixing, spectral response, and temporal–spectral fusion.

Specifically, shared waveform and coordinate encoders combine each observed electrode's signal and spatial position into a token u_i. A Set Transformer (Lee et al., 2019) exchanges information across the support set and produces an updated feature h_i for each electrode. The global representation $z_{\mathcal{S}}$ is then obtained by permutation-invariant mean pooling over these features. In parallel, the observed waveforms are lifted onto learned spatial anchors using distance-based weights and processed by a temporal encoder to produce $c_{\mathcal{S}}(1:T)$. These two representations feed complementary conditioning branches: a time-dependent FiLM modulator consumes the unpooled temporal context, whereas the field generator consumes $z_{\mathcal{S}}$ together with the temporally pooled context.

2.2 Support-generated EEG neural field

Generic conditional INRs (Li et al., 2025b; Naour et al., 2024) typically inject observation-level information through feature concatenation or feature-wise modulation while retaining a shared operator structure across recordings. Such conditioning changes the hidden representation but does not explicitly adapt which temporal scales or spectral responses are emphasized for the current EEG segment. ScalpINR instead conditions the decoder at two complementary resolutions. Fine-grained, time-dependent FiLM parameters are generated from $z_{\mathcal{S}}$ and the unpooled context $c_{\mathcal{S}}(t)$, while

compact support-level parameters control temporal-scale mixing, spectral response, and temporal–spectral fusion in each decoder block.

For fine-grained temporal conditioning, a shared modulator maps the global context and the context at each time index to layer-wise FiLM parameters:

$$\left\{\gamma_{\mathcal{S}}^{(\ell)}(t), \beta_{\mathcal{S}}^{(\ell)}(t)\right\}_{\ell=1}^{L} = H_{\mathrm{FiLM}}\left([z_{\mathcal{S}}; c_{\mathcal{S}}(t)]\right), \qquad t = 1, \ldots, T. \tag{1}$$

For the complementary support-level controls, we temporally average the time-resolved context and concatenate it with the global representation. A field generator then maps this compact support representation to the block-wise temporal, spectral, and fusion control parameters:

$$\left\{\tilde{\mathbf{a}}_{\mathcal{S}}^{(\ell)}, \tilde{\mathbf{g}}_{\mathcal{S}}^{(\ell)}, \tilde{\lambda}_{\mathcal{S}}^{(\ell)}\right\}_{\ell=1}^{L} = H_{\theta}\left(\left[z_{\mathcal{S}}; \frac{1}{T}\sum_{t=1}^{T} c_{\mathcal{S}}(t)\right]\right). \tag{2}$$

We apply a softmax to $\tilde{\mathbf{a}}_{\mathcal{S}}^{(\ell)}$ to obtain the temporal-mixture weights $\mathbf{a}_{\mathcal{S}}^{(\ell)}$, and a sigmoid to $\tilde{\lambda}_{\mathcal{S}}^{(\ell)}$ to obtain the fusion gate $\lambda_{\mathcal{S}}^{(\ell)} \in (0, 1)$. The spectral response in Eq. equation 6 uses $\tanh\big(\tilde{\mathbf{g}}_{\mathcal{S}}^{(\ell)}\big)$ as its bounded coefficients. These support-level controls are shared across time, query locations, and hidden feature channels.

Given these support-conditioned parameters, a target coordinate p_q is first converted into a temporal feature sequence. Let γ_p and γ_t denote fixed Fourier-feature mappings for the 3D electrode coordinate and normalized time, respectively, and let E_p and E_t denote their learned projections into the decoder feature space. The spatial embedding is shared across the temporal grid, whereas the temporal embedding varies with t:

$$V_q^{(0)}(t) = \mathrm{LN}\left(E_p(\gamma_p(p_q)) + E_t(\gamma_t(\bar{t}))\right), \qquad t = 1, \ldots, T, \tag{3}$$

where $\bar{t}$ denotes the normalized temporal coordinate. The resulting sequence $V_q^{(0)} \in \mathbb{R}^{T\times D}$ is subsequently processed by L support-conditioned decoder blocks. At decoder block ℓ, the time-dependent FiLM parameters first modulate the hidden representation:

$$\bar{V}_q^{(\ell)}(t) = \mathrm{LN}\left(\mathrm{GELU}\left(\left[1 + s_{\mathrm{mod}} \tanh\left(\gamma_{\mathcal{S}}^{(\ell)}(t)\right)\right] \odot W^{(\ell)} V_q^{(\ell)}(t) + s_{\mathrm{mod}} \beta_{\mathcal{S}}^{(\ell)}(t)\right) + V_q^{(\ell)}(t)\right), \tag{4}$$

where s_{mod} controls the magnitude of the FiLM modulation.

The temporal branch then processes $\bar{V}_q^{(\ell)}$ using a bank of shared depthwise temporal operators $\{\mathcal{T}_m\}_{m=1}^{M}$ with different receptive fields:

$$V_{\mathrm{temp}}^{(\ell)} = \sum_{m=1}^{M} a_{\mathcal{S},m}^{(\ell)} \mathcal{T}_m\left(\bar{V}_q^{(\ell)}\right). \tag{5}$$

In parallel, we construct a recording-specific spectral response from $\tilde{\mathbf{g}}_{\mathcal{S}}^{(\ell)}$. Let $\{B_b(f)\}_{b=1}^{B}$ denote shared, normalized Gaussian bases spanning the normalized rFFT frequency axis. The spectral response is

$$R_{\mathcal{S}}^{(\ell)}(f) = 1 + \eta_{\mathrm{spec}} \sum_{b=1}^{B} \tanh\left(\tilde{g}_{\mathcal{S},b}^{(\ell)}\right) B_b(f), \tag{6}$$

where η_{spec} bounds the magnitude of spectral modulation. We then apply this response to the temporal-frequency representation of the FiLM-modulated feature:

$$V_{\mathrm{spec}}^{(\ell)} = \mathrm{IRFFT}_t\left(R_{\mathcal{S}}^{(\ell)} \odot \mathrm{RFFT}_t\left(\bar{V}_q^{(\ell)}\right)\right). \tag{7}$$

Finally, the support-generated fusion gate combines the temporal and spectral responses, and the fused feature is added to the FiLM-modulated representation through a residual update:

$$V_q^{(\ell+1)} = \mathrm{LN}_{\mathrm{func}}^{(\ell)}\left(\bar{V}_q^{(\ell)} + \eta^{(\ell)}\left[\lambda_{\mathcal{S}}^{(\ell)} V_{\mathrm{temp}}^{(\ell)} + \left(1 - \lambda_{\mathcal{S}}^{(\ell)}\right) V_{\mathrm{spec}}^{(\ell)}\right]\right). \tag{8}$$

where $\eta^{(\ell)}$ is a shared, support-independent residual gain.

Applying the L support-conditioned decoder blocks recursively yields

$$V_q^{(L)} = D_{\mathcal{S}}\left(V_q^{(0)}\right), \tag{9}$$

where $D_{\mathcal{S}}$ denotes the shared coordinate decoder instantiated by the time-resolved FiLM parameters and support-generated functional controls derived from $\mathcal{S}$. A shared output projection then produces the reconstructed waveform:

$$\hat{x}_q(t) = w_{\text{out}}^{\top} V_q^{(L)}(t) + b_{\text{out}}, \qquad \hat{\mathbf{x}}_q = f_{\mathcal{S}}(p_q). \tag{10}$$

2.3 Support-topology consistency

The neural field inferred from a recording should remain stable when the same underlying EEG is observed through different electrode subsets. During training, we independently sample two channel masks from the same complete EEG segment, producing support views $\mathcal{S}^{(a)}$ and $\mathcal{S}^{(b)}$ with independently sampled support ratios r_a and r_b. Let $\mathcal{S}^{(\text{hi})}$ and $\mathcal{S}^{(\text{lo})}$ denote the views with the higher and lower support ratios, respectively, with ties broken arbitrarily. Both views are processed by the same ScalpINR with fully shared parameters; they are two support-conditioned forward passes of one model and differ only in the electrodes supplied as input. Because the masks are sampled independently, $\mathcal{S}^{(\text{hi})}$ is not required to contain $\mathcal{S}^{(\text{lo})}$.

Both support-conditioned fields are evaluated at the training-eligible electrode locations. Let $\mathcal{Q}_{\text{miss}}^{(\text{lo})}$ denote the electrode locations unavailable in the lower-support view among these locations. Under the strict unseen-electrode protocol, permanently held-out electrodes are excluded from this set and therefore do not participate in the topology-consistency objective. We impose an asymmetric cross-view consistency loss at $\mathcal{Q}_{\text{miss}}^{(\text{lo})}$ and stop gradients through the prediction conditioned on the higher-support view:

$$\mathcal{L}_{\text{topo}} = \frac{1}{|\mathcal{Q}_{\text{miss}}^{(\text{lo})}|} \sum_{p_q \in \mathcal{Q}_{\text{miss}}^{(\text{lo})}} \ell\left(f_{\mathcal{S}^{(\text{lo})}}(p_q), \text{sg}\left[f_{\mathcal{S}^{(\text{hi})}}(p_q)\right]\right). \tag{11}$$

Here, $\text{sg}[\cdot]$ only prevents gradients from this loss from flowing through the higher-support forward pass; it does not introduce a separately parameterized or pretrained model.

Because the two masks are sampled independently, a location missing from the lower-support view may be either observed or unobserved in the higher-support view. The objective directly penalizes changes in the generated waveform caused by varying the support topology for the same underlying EEG segment. Unlike $\mathcal{L}_{\text{rec}}$, which independently anchors each view to the measured waveform, $\mathcal{L}_{\text{topo}}$ explicitly couples the predictions obtained from the two support configurations.

2.4 Training objective and inference

Irrespective of the asymmetric gradient assignment in $\mathcal{L}_{\text{topo}}$, both support views are supervised symmetrically by the reconstruction and fidelity objectives. For each support view, missing electrodes form the reconstruction set $\mathcal{R}^{(v)}$. We additionally select a subset of observed electrodes $\mathcal{P}^{(v)}$ as visible-but-predicted locations, requiring the instantiated field to remain consistent with the measurements used to construct it. The two supervision terms are

$$\mathcal{L}_{\text{rec}} = \frac{1}{2} \sum_{v \in \{a,b\}} \frac{1}{|\mathcal{R}^{(v)}|} \sum_{p_q \in \mathcal{R}^{(v)}} \ell\left(f_{\mathcal{S}^{(v)}}(p_q), \mathbf{x}(p_q)\right), \tag{12}$$

$$\mathcal{L}_{\text{fid}} = \frac{1}{2} \sum_{v \in \{a,b\}} \frac{1}{|\mathcal{P}^{(v)}|} \sum_{p_i \in \mathcal{P}^{(v)}} \ell\left(f_{\mathcal{S}^{(v)}}(p_i), \mathbf{x}_i\right). \tag{13}$$

The complete objective is

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \lambda_{\text{fid}} \mathcal{L}_{\text{fid}} + \lambda_{\text{topo}} \mathcal{L}_{\text{topo}}. \tag{14}$$

We use an L1 distance for waveform reconstruction and fidelity supervision. The precise protocol-specific loss weights and the topology-consistency distance are provided in Appendix C.5. At inference, the available EEG channels are encoded once to generate the time-resolved FiLM parameters

together with the support-dependent temporal-scale, spectral-response, and fusion controls. The resulting field can then be repeatedly queried at desired electrode coordinates without retraining or test-time optimization.

3 EXPERIMENTS

We organize the experiments around four questions. First, how well does ScalpINR reconstruct missing EEG channels when the target electrode locations are represented during training? Second, can it generalize to electrode coordinates that are completely excluded from training? Third, which components of the support-conditioned functional adaptation and support-topology consistency contribute to the reconstruction gains? Finally, do the reconstructed signals preserve task-relevant information for independently trained downstream models?

3.1 EXPERIMENT SETUP

Datasets. We evaluate ScalpINR on three publicly available EEG datasets, including **SEED** (Zheng & Lu, 2015), **ESAA** (Li et al., 2022) (denoted as ESAA), and the **EEG Motor Movement/Imagery Dataset (EEGMMIDB)** collected using BCI2000 (Schalk, 2009; Schalk et al., 2004) (denoted as BCI2000). These datasets cover emotion-related EEG, auditory-attention EEG, and motor-imagery EEG, providing diverse electrode layouts, signal statistics, and downstream decoding tasks for evaluating both reconstruction fidelity and task utility. For each dataset, we construct subject-disjoint training, validation, and test partitions before window extraction. No trial or overlapping windows are shared across partitions. For each unseen-electrode split, the held-out electrodes are selected before any data-dependent preprocessing and are excluded from all training and validation inputs, reconstruction targets, losses, normalization statistics, and model-selection criteria. Detailed dataset descriptions and preprocessing pipelines are provided in Appendix C.1.

Baselines. We compare ScalpINR with four categories of reconstruction approaches. Spherical spline interpolation serves as a classical geometry-based baseline. CNP and ANP provide generic set-to-target regression controls (Garnelo et al., 2018; Kim et al., 2019), while Conditional INR-Concat and Conditional INR-FiLM are used as parameter-matched controls for generic implicit conditioning through concatenation and feature-wise modulation, respectively (Perez et al., 2018). For coordinate-flexible EEG reconstruction, we compare against ZUNA and Neural Brain Fields (Warner et al., 2026; Kedem et al., 2025); NBF fits a separate field for each evaluated recording, whereas ZUNA performs coordinate-aware diffusion-based reconstruction. Fixed-layout baselines include ImputeINR, CSDI (Tashiro et al., 2021), ESTformer (Li et al., 2025a), SRGDiff (Liu et al., 2026), and MCMA-EEG, our adaptation of MCMA (Chen et al., 2024) to fixed EEG montages. Fixed-layout methods are evaluated only when the target electrode locations have appeared during training. All methods use identical data splits, support masks, query targets, and evaluation code wherever their model assumptions permit.

Evaluation protocol. For each support ratio, we draw 50 test-time support masks and use the same paired masks for all applicable methods. Signal-level metrics are computed only on complementary query electrodes. NMSE and PCC measure waveform reconstruction. We separately evaluate frequency-domain reconstruction error as complementary spectral measures. Their exact definitions are provided in Appendix C.4. For unseen-electrode evaluation, each trainable model is retrained for every held-out-electrode split, and results are aggregated across the predefined splits. Downstream evaluation uses task models trained only on real high-density EEG and subsequently frozen for all reconstruction methods.

3.2 RECONSTRUCTION AT SEEN ELECTRODE LOCATIONS

We first consider the standard missing-channel setting, in which the target electrode locations appear during training but their signals are excluded from the test-time support set. Both the number and spatial configuration of the support electrodes vary across samples, and all metrics are computed only on the masked query electrodes.

Table 1: Random missing-channel reconstruction results at $r = 0.125$, the sparsest support ratio shared by the three datasets. Lower NMSE and frequency error indicate better reconstruction, while higher PCC indicates better signal correlation. Best results are in bold and second-best results are underlined.

Method	SEED			ESAA			BCI2000		
	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓
Classical and coordinate-queryable baselines									
Spherical Spline	2.06 ± 0.07	0.40 ± 0.05	0.60 ± 0.06	0.55 ± 0.03	0.70 ± 0.04	0.51 ± 0.06	0.68 ± 0.03	0.71 ± 0.04	0.51 ± 0.06
ZUNA	0.99 ± 0.05	0.02 ± 0.02	0.67 ± 0.06	0.75 ± 0.05	0.50 ± 0.05	0.57 ± 0.06	0.99 ± 0.05	0.01 ± 0.02	0.39 ± 0.05
Neural Brain Fields	1.34 ± 0.07	0.51 ± 0.05	0.50 ± 0.06	0.48 ± 0.05	$\underline{0.76} \pm 0.04$	0.46 ± 0.06	0.90 ± 0.05	0.32 ± 0.05	0.32 ± 0.05
Generic set-conditioned controls									
CNP	1.95 ± 0.08	0.42 ± 0.05	0.37 ± 0.05	0.54 ± 0.05	0.66 ± 0.04	0.25 ± 0.05	0.89 ± 0.05	0.40 ± 0.02	0.28 ± 0.05
ANP	2.51 ± 0.08	0.34 ± 0.05	0.39 ± 0.05	0.72 ± 0.05	0.46 ± 0.05	0.31 ± 0.05	0.82 ± 0.05	0.46 ± 0.02	0.24 ± 0.06
Fixed-layout baselines									
ImputeINR	1.54 ± 0.08	0.47 ± 0.05	0.21 ± 0.04	0.57 ± 0.05	0.67 ± 0.04	0.13 ± 0.04	0.68 ± 0.05	0.71 ± 0.03	0.13 ± 0.04
CSDI	2.38 ± 0.08	0.35 ± 0.05	0.42 ± 0.06	0.68 ± 0.05	0.57 ± 0.05	0.28 ± 0.05	0.82 ± 0.05	0.67 ± 0.04	0.14 ± 0.05
ESTFormer	0.69 ± 0.05	0.57 ± 0.05	0.12 ± 0.04	0.58 ± 0.05	0.66 ± 0.04	0.12 ± 0.04	0.70 ± 0.05	0.73 ± 0.04	0.11 ± 0.04
SRGDiff	1.87 ± 0.08	0.42 ± 0.05	0.33 ± 0.05	0.59 ± 0.05	0.68 ± 0.04	$\underline{0.07} \pm 0.02$	0.71 ± 0.05	0.69 ± 0.04	$\underline{0.08} \pm 0.02$
MCMA-EEG	$\underline{0.64} \pm 0.05$	$\underline{0.60} \pm 0.04$	$\underline{0.11} \pm 0.02$	$\underline{0.45} \pm 0.05$	0.75 ± 0.04	0.09 ± 0.05	$\underline{0.31} \pm 0.05$	$\underline{0.83} \pm 0.04$	$\underline{0.08} \pm 0.02$
ScalpINR	$\mathbf{0.54} \pm 0.04$	$\mathbf{0.67} \pm 0.08$	$\mathbf{0.09} \pm 0.02$	$\mathbf{0.39} \pm 0.04$	$\mathbf{0.78} \pm 0.07$	$\mathbf{0.05} \pm 0.01$	$\mathbf{0.25} \pm 0.04$	$\mathbf{0.86} \pm 0.06$	$\mathbf{0.06} \pm 0.01$

Table 2: Strict unseen-electrode generation results under different support ratios r, evaluated only on permanently held-out electrodes. Here, r denotes the fraction of non-held-out electrodes used as visible support, and $r = 1$ means that all non-held-out electrodes are available while the held-out electrodes are used only as query targets. We report mean ± standard deviation across multiple random hold-out splits. Lower NMSE and frequency error and higher PCC indicate better performance. Best results are in bold and second-best results are underlined.

r	Method	SEED (held-out only)			ESAA (held-out only)			BCI2000 (held-out only)		
		NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓
1	Spherical Spline	$\underline{0.29} \pm 0.03$	$\underline{0.84} \pm 0.04$	$\underline{0.23} \pm 0.02$	$\underline{0.16} \pm 0.04$	$\underline{0.94} \pm 0.05$	$\underline{0.22} \pm 0.02$	$\underline{0.18} \pm 0.03$	$\underline{0.90} \pm 0.04$	$\underline{0.24} \pm 0.02$
	CNP	1.02 ± 0.07	0.12 ± 0.02	0.40 ± 0.06	0.78 ± 0.05	0.16 ± 0.02	0.35 ± 0.05	0.98 ± 0.05	0.14 ± 0.03	0.36 ± 0.05
	ANP	1.06 ± 0.08	0.04 ± 0.02	0.35 ± 0.05	0.83 ± 0.05	0.06 ± 0.02	0.39 ± 0.05	1.00 ± 0.05	0.05 ± 0.02	0.28 ± 0.05
	ZUNA	0.85 ± 0.05	0.32 ± 0.05	0.37 ± 0.05	0.70 ± 0.05	0.51 ± 0.06	0.26 ± 0.05	0.90 ± 0.05	0.24 ± 0.02	0.32 ± 0.05
	Neural Brain Fields	0.90 ± 0.05	0.29 ± 0.02	0.42 ± 0.06	0.67 ± 0.04	0.34 ± 0.05	0.32 ± 0.05	0.90 ± 0.05	0.33 ± 0.05	0.39 ± 0.05
	ScalpINR	$\mathbf{0.26} \pm 0.04$	$\mathbf{0.86} \pm 0.06$	$\mathbf{0.09} \pm 0.01$	$\mathbf{0.10} \pm 0.03$	$\mathbf{0.95} \pm 0.05$	$\mathbf{0.12} \pm 0.01$	$\mathbf{0.16} \pm 0.04$	$\mathbf{0.95} \pm 0.05$	$\mathbf{0.06} \pm 0.01$
0.5	Spherical Spline	$\underline{0.73} \pm 0.03$	0.63 ± 0.04	0.39 ± 0.05	$\underline{0.42} \pm 0.03$	$\underline{0.80} \pm 0.04$	0.41 ± 0.06	$\underline{0.28} \pm 0.03$	$\underline{0.87} \pm 0.04$	$\underline{0.26} \pm 0.02$
	CNP	1.06 ± 0.08	0.11 ± 0.02	0.37 ± 0.05	0.82 ± 0.05	0.15 ± 0.02	0.38 ± 0.04	0.98 ± 0.05	0.14 ± 0.02	0.34 ± 0.05
	ANP	1.10 ± 0.08	0.04 ± 0.02	$\underline{0.32} \pm 0.02$	0.88 ± 0.05	0.06 ± 0.02	0.41 ± 0.06	1.00 ± 0.05	0.05 ± 0.02	0.41 ± 0.06
	ZUNA	0.93 ± 0.05	0.12 ± 0.02	0.43 ± 0.07	0.77 ± 0.05	0.47 ± 0.05	0.43 ± 0.06	0.95 ± 0.05	0.21 ± 0.02	0.32 ± 0.05
	Neural Brain Fields	0.94 ± 0.05	0.26 ± 0.02	0.36 ± 0.05	0.72 ± 0.05	0.33 ± 0.05	$\underline{0.25} \pm 0.02$	0.90 ± 0.05	0.33 ± 0.05	$\underline{0.26} \pm 0.02$
	ScalpINR	$\mathbf{0.35} \pm 0.03$	$\mathbf{0.81} \pm 0.06$	$\mathbf{0.11} \pm 0.02$	$\mathbf{0.22} \pm 0.04$	$\mathbf{0.94} \pm 0.05$	$\mathbf{0.08} \pm 0.01$	$\mathbf{0.16} \pm 0.04$	$\mathbf{0.94} \pm 0.05$	$\mathbf{0.06} \pm 0.01$
0.25	Spherical Spline	1.30 ± 0.07	$\underline{0.51} \pm 0.05$	$\underline{0.26} \pm 0.02$	$\underline{0.68} \pm 0.03$	$\underline{0.75} \pm 0.04$	0.30 ± 0.05	$\underline{0.49} \pm 0.03$	$\underline{0.78} \pm 0.04$	0.38 ± 0.05
	CNP	1.11 ± 0.09	0.10 ± 0.02	0.45 ± 0.06	0.86 ± 0.05	0.15 ± 0.02	0.32 ± 0.05	0.98 ± 0.05	0.14 ± 0.02	0.36 ± 0.05
	ANP	1.15 ± 0.08	0.03 ± 0.02	0.37 ± 0.05	0.92 ± 0.05	0.05 ± 0.03	0.36 ± 0.05	1.00 ± 0.05	0.05 ± 0.02	0.43 ± 0.06
	ZUNA	1.05 ± 0.08	0.06 ± 0.02	0.44 ± 0.06	0.81 ± 0.05	0.45 ± 0.05	$\underline{0.21} \pm 0.02$	0.92 ± 0.04	0.04 ± 0.02	0.50 ± 0.06
	Neural Brain Fields	$\underline{0.98} \pm 0.05$	0.23 ± 0.02	0.27 ± 0.05	0.78 ± 0.05	0.32 ± 0.06	0.40 ± 0.06	0.89 ± 0.05	0.33 ± 0.05	$\underline{0.35} \pm 0.02$
	ScalpINR	$\mathbf{0.58} \pm 0.04$	$\mathbf{0.75} \pm 0.07$	$\mathbf{0.13} \pm 0.01$	$\mathbf{0.35} \pm 0.04$	$\mathbf{0.90} \pm 0.05$	$\mathbf{0.06} \pm 0.01$	$\mathbf{0.43} \pm 0.04$	$\mathbf{0.86} \pm 0.06$	$\mathbf{0.19} \pm 0.01$
0.125	Spherical Spline	2.41 ± 0.07	$\underline{0.37} \pm 0.05$	0.51 ± 0.06	0.87 ± 0.03	$\underline{0.69} \pm 0.04$	$\underline{0.31} \pm 0.02$	$\underline{0.72} \pm 0.03$	$\underline{0.69} \pm 0.04$	$\underline{0.35} \pm 0.02$
	CNP	1.14 ± 0.08	0.08 ± 0.02	0.46 ± 0.06	0.89 ± 0.06	0.13 ± 0.02	0.41 ± 0.07	0.98 ± 0.05	0.13 ± 0.02	0.37 ± 0.05
	ANP	1.18 ± 0.07	0.02 ± 0.03	0.40 ± 0.06	0.95 ± 0.04	0.05 ± 0.02	0.45 ± 0.05	1.00 ± 0.05	0.05 ± 0.02	0.42 ± 0.06
	ZUNA	1.21 ± 0.08	0.01 ± 0.02	$\underline{0.27} \pm 0.02$	$\underline{0.79} \pm 0.05$	0.43 ± 0.05	0.43 ± 0.06	0.95 ± 0.05	0.02 ± 0.02	0.36 ± 0.05
	Neural Brain Fields	$\underline{1.02} \pm 0.08$	0.19 ± 0.02	0.43 ± 0.06	0.84 ± 0.06	0.31 ± 0.05	0.33 ± 0.05	0.89 ± 0.05	0.33 ± 0.04	0.44 ± 0.06
	ScalpINR	$\mathbf{0.70} \pm 0.04$	$\mathbf{0.69} \pm 0.08$	$\mathbf{0.24} \pm 0.01$	$\mathbf{0.47} \pm 0.04$	$\mathbf{0.83} \pm 0.06$	$\mathbf{0.21} \pm 0.01$	$\mathbf{0.62} \pm 0.04$	$\mathbf{0.86} \pm 0.06$	$\mathbf{0.29} \pm 0.01$

Table 1 compares all methods at $r = 0.125$, the sparsest support ratio shared by the three datasets. ScalpINR achieves the best result on all nine dataset–metric combinations. Relative to the best competing method in NMSE, ScalpINR reduces the error by 15.6%, 13.3%, and 19.4% on SEED, ESAA, and BCI2000, respectively. The improvement is consistent across complementary measures: PCC increases from 0.60 to 0.67, from 0.76 to 0.78, and from 0.83 to 0.86, while frequency-domain error decreases from 0.11 to 0.09, from 0.07 to 0.05, and from 0.08 to 0.06. These results show that ScalpINR preserves both waveform-level and spectral structure when only a small fraction of the electrode montage is observed. Complete results across support ratios are provided in Appendix D.

3.3 Querying Unseen Electrode Locations

We next evaluate whether ScalpINR can query electrode locations excluded from training. For each held-out split, approximately 10% of the electrodes are permanently removed from both the training support sets and the training reconstruction targets, such that no training loss is computed on their signals. At test time, these electrodes are provided only through their coordinates and evaluated as query targets. The support ratio r is defined over the remaining electrodes; $r = 1$

Table 3: Unified ablation on SEED under the unseen-electrode protocol at $r = 0.5$. Lower NMSE and frequency-domain MAE and higher PCC indicate better performance. Results are reported as mean $\pm$ standard deviation over the random hold-out splits.

Variant	Params (M)	NMSE ↓	PCC ↑	Freq. MAE ↓
Generic set-conditioned models				
Conditional INR-Concat	0.2762	0.52 ± 0.03	0.74 ± 0.04	0.14 ± 0.02
Conditional INR-FiLM	0.3967	0.45 ± 0.02	0.72 ± 0.05	0.14 ± 0.01
ScalpINR ablations				
Fixed field parameters	0.402	0.39 ± 0.02	0.80 ± 0.03	0.12 ± 0.01
w/o support-generated temporal weights	0.409	0.38 ± 0.02	0.80 ± 0.03	0.12 ± 0.01
w/o support-generated spectral response	0.409	0.37 ± 0.01	0.80 ± 0.02	0.12 ± 0.01
w/o support-generated fusion gate	0.410	0.40 ± 0.03	0.77 ± 0.05	0.14 ± 0.02
w/o topology consistency	0.438	0.39 ± 0.02	0.79 ± 0.04	0.13 ± 0.01
ScalpINR	0.438	$\mathbf{0.35} \pm 0.01$	$\mathbf{0.81} \pm 0.03$	$\mathbf{0.11} \pm 0.01$

uses all non-held-out electrodes as support, whereas $r < 1$ additionally introduces support sparsity. As shown in Table 2, ScalpINR achieves the best NMSE, PCC, and frequency-domain error across every evaluated dataset and support ratio among methods that can query previously unseen electrode coordinates. In the densest-support setting ($r = 1$), where all non-held-out electrodes are observed, ScalpINR reduces NMSE from 0.29 to 0.26 on SEED, from 0.16 to 0.10 on ESAA, and from 0.18 to 0.16 on BCI2000 relative to the strongest competing method. The 37.5% relative reduction on ESAA is particularly notable because the evaluated target electrodes are completely absent from training.

The advantage remains when coordinate generalization is combined with sparse support. At $r = 0.5$, ScalpINR reduces NMSE from 0.73 to 0.35 on SEED, from 0.42 to 0.22 on ESAA, and from 0.28 to 0.16 on BCI2000, corresponding to relative reductions of 52.1%, 47.6%, and 42.9%, respectively. The gains are accompanied by consistently higher PCC and lower frequency-domain error across all support ratios. Since the held-out electrodes are excluded from training inputs, reconstruction targets, normalization statistics, and model selection, these results demonstrate generalization to genuinely unseen electrode coordinates rather than interpolation over a memorized channel inventory.

3.4 Ablation Studies

We conduct controlled ablations to examine whether the gains of ScalpINR arise from generic support conditioning, support-generated functional adaptation, or the topology-consistency objective. We compare against two conditional INR baselines using the same support encoder and coordinate-query interface: Conditional INR-Concat injects the support representation by feature concatenation, whereas Conditional INR-FiLM applies support-dependent feature-wise modulation. We further replace all support-generated temporal weights, spectral responses, and fusion gates with recording-independent parameters, remove each type of functional control individually, and remove the topology-consistency objective while keeping the remaining architecture unchanged.

Table 3 shows that ScalpINR consistently outperforms both generic conditional INR baselines and all component ablations. Compared with Conditional INR-FiLM, ScalpINR reduces NMSE from 0.45 to 0.35, while improving PCC from 0.72 to 0.81 and reducing frequency-domain error from 0.14 to 0.11. Replacing all support-generated controls with fixed parameters increases NMSE to 0.39, while removing the temporal weights, spectral response, fusion gate, or topology consistency increases NMSE to 0.38, 0.37, 0.40, and 0.39, respectively. These results show that support-dependent temporal–spectral adaptation and cross-support consistency provide complementary contributions to reconstruction quality.

3.5 Downstream Decoding with Reconstructed EEG

Pointwise reconstruction metrics do not directly indicate whether the recovered EEG preserves information useful for downstream analysis. We therefore evaluate reconstructed EEG at $r = 0.25$

Table 4: Task-oriented evaluation at $r = 0.25$. We report classification accuracy (%) using a task-specific classifier and the frozen EEGPT encoder. The task classifier and the linear probe associated with EEGPT are trained only on GT-HD training EEG and are shared across all reconstruction methods. Observed electrodes retain their measured signals, and each method fills only the missing electrodes. Best and second-best reconstruction results are highlighted in bold and underlined, respectively.

Input / Reconstruction	Task classifier			EEGPT		
	SEED	ESAA	BCI2000	SEED	ESAA	BCI2000
GT-HD (reference)	40.37 ± 5.24	53.12 ± 3.06	70.84 ± 3.42	37.18 ± 3.84	47.08 ± 2.12	48.72 ± 2.74
Sparse input	34.26 ± 5.17	50.73 ± 3.24	51.19 ± 4.06	32.03 ± 4.11	42.97 ± 2.48	41.15 ± 3.02
Spherical Spline	33.58 ± 5.04	50.41 ± 3.11	53.92 ± 3.87	34.26 ± 3.95	45.13 ± 2.21	44.08 ± 2.86
ImputeINR	34.11 ± 5.15	50.86 ± 3.18	66.27 ± 3.55	33.17 ± 3.88	43.68 ± 2.35	42.31 ± 2.95
CSDI	33.69 ± 5.09	50.34 ± 3.29	50.75 ± 4.12	31.58 ± 4.16	42.15 ± 2.56	40.26 ± 3.11
ESTformer	$\underline{37.42} \pm 4.83$	51.28 ± 2.87	67.53 ± 3.61	$\underline{35.42} \pm 3.72$	$\underline{46.27} \pm 2.04$	$\underline{45.54} \pm 2.68$
SRGDiff	33.17 ± 5.31	50.95 ± 3.21	50.48 ± 4.23	32.84 ± 4.03	43.92 ± 2.29	41.76 ± 3.05
MCMA-EEG	34.63 ± 4.92	$\underline{51.37} \pm 2.94$	$\underline{68.21} \pm 3.48$	34.73 ± 3.81	44.86 ± 2.17	43.29 ± 2.81
ZUNA	33.84 ± 4.86	50.52 ± 3.02	50.66 ± 3.98	31.96 ± 4.08	41.88 ± 2.61	39.74 ± 3.16
Neural Brain Fields	31.29 ± 5.41	49.74 ± 3.38	45.38 ± 4.47	33.51 ± 3.99	45.64 ± 2.14	44.37 ± 2.79
ScalpINR	$\mathbf{40.58} \pm 5.28$	$\mathbf{53.46} \pm 3.12$	$\mathbf{69.13} \pm 3.36$	$\mathbf{36.08} \pm 3.76$	$\mathbf{46.75} \pm 1.98$	$\mathbf{48.12} \pm 2.71$

using two downstream consumers: a task-specific classifier and a linear probe built on the frozen EEGPT encoder. For each dataset, the task-specific classifier is trained and selected only on ground-truth high-density training and validation EEG. Similarly, the EEGPT encoder is kept frozen and its linear probe is trained only on ground-truth high-density EEG. Once trained, both downstream models are fixed and shared across all reconstruction methods, preventing them from adapting to method-specific reconstruction artifacts. At test time, all methods complete the same sparse recordings: observed electrodes retain their measured signals, while only missing electrodes are replaced by reconstructed signals. GT-HD provides the full-density reference, whereas Sparse input evaluates the incomplete recordings without reconstruction.

Table 4 shows that ScalpINR achieves the highest downstream accuracy among all reconstruction methods for all six dataset–consumer combinations. With the task-specific classifiers, ScalpINR reaches 40.58%, 53.46%, and 69.13% accuracy on SEED, ESAA, and BCI2000, outperforming the strongest competing reconstruction method by 3.16, 2.09, and 0.92 percentage points, respectively. Its performance also closely matches the GT-HD reference, with absolute gaps of only 0.21, 0.34, and 1.71 percentage points.

The same pattern holds with the independently frozen EEGPT representation. ScalpINR obtains 36.08%, 46.75%, and 48.12% accuracy, outperforming the strongest competing reconstruction by 0.66, 0.48, and 2.58 percentage points. These values are within 1.10, 0.33, and 0.60 percentage points of the GT-HD reference. Because neither downstream consumer is trained on reconstructed EEG, the consistent gains indicate that ScalpINR preserves task-relevant signal structure rather than merely optimizing pointwise reconstruction metrics.

4 Conclusion

We presented ScalpINR, an amortized support-conditioned neural field for EEG spatial super-resolution under variable electrode availability. ScalpINR infers a recording-specific coordinate-to-signal field using support-dependent temporal–spectral controls and a dual-support objective for topology consistency. Experiments on SEED, ESAA, and BCI2000 show strong reconstruction under random missing-channel and strict unseen-electrode settings, including waveform generation at electrode coordinates excluded from training without test-time optimization. Ablation and downstream results further support the proposed components and show preservation of task-relevant EEG information. Future work will consider broader electrode configurations and acquisition conditions.

REFERENCES

Jiarong Chen, Wanqing Wu, Tong Liu, and Shenda Hong. Multi-channel masked autoencoder and comprehensive evaluations for reconstructing 12-lead ecg from arbitrary single-lead ecg. *NPJ Cardiovascular Health*, 1(1):34, 2024.

Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *International conference on machine learning*, pp. 1704–1713. PMLR, 2018.

David Ha, Andrew M. Dai, and Quoc V. Le. Hypernetworks. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=rkpACe1lx.

Alberto Jaramillo-Jimenez, Diego A Tovar-Rios, Johann Alexis Ospina, Yorguin-Jose Mantilla-Ramos, Daniel Loaiza-López, Verónica Henao Isaza, Luisa María Zapata Saldarriaga, Valeria Cadavid Castro, Jazmin Ximena Suarez-Revelo, Yamile Bocanegra, et al. Spectral features of resting-state EEG in parkinson's disease: A multicenter study using functional data analysis. *Clinical Neurophysiology*, 151:28–40, 2023.

Shahar Ain Kedem, Itamar Zimerman, and Eliya Nachmani. Neural brain fields: A nerf-inspired approach for generating nonexistent eeg electrodes. *arXiv preprint arXiv:2601.00012*, 2025.

Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. *arXiv preprint arXiv:1901.05761*, 2019.

Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and Brent J Lance. EEGNet: a compact convolutional neural network for EEG-based brain–computer interfaces. *Journal of neural engineering*, 15(5):056013, 2018.

Etienne Le Naour, Louis Serrano, Léon Migus, Yuan Yin, Ghislain Agoua, Nicolas Baskiotis, Patrick Gallinari, and Vincent Guigue. Time series continuous modeling for imputation and forecasting with implicit neural representations. *Transactions on Machine Learning Research Journal*, 2024.

Sven Leach, Ku-young Chung, Laura Tüshaus, Reto Huber, and Walter Karlen. A protocol for comparing dry and wet eeg electrodes during sleep. *Frontiers in neuroscience*, 14:586, 2020.

Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International conference on machine learning*, pp. 3744–3753. PMLR, 2019.

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

Mengxuan Li, Ke Liu, Jialong Guo, Jiajun Bu, Hongwei Wang, and Haishuai Wang. Imputeinr: time series imputation via implicit neural representations for disease diagnosis with missing data. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pp. 9241–9249, 2025b.

Peiwen Li, Enze Su, Jia Li, Siqi Cai, Longhan Xie, and Haizhou Li. ESAA: An EEG-speech auditory attention detection database. In *2022 25th Conference of the Oriental COCOSDA International Committee for the Co-ordination and Standardisation of Speech Databases and Assessment Techniques (O-COCOSDA)*, pp. 1–6, 2022. doi: 10.1109/O-COCOSDA202257103.2022.9997944.

Hongjun Liu, Leyu Zhou, Zijianghao Yang, and Chao Yao. Step-aware residual-guided diffusion for eeg spatial super-resolution. *The Fourteenth International Conference on Learning Representations*, 2026.

Hanane Moumane, Jérémy Pazuelo, Mérie Nassar, Jose Yesith Juez, Mario Valderrama, and Michel Le Van Quyen. Signal quality evaluation of an in-ear eeg device in comparison to a conventional cap system. *Frontiers in Neuroscience*, 18:1441897, 2024.

Etienne Le Naour, Louis Serrano, Léon Migus, Yuan Yin, Ghislain Agoua, Nicolas Baskiotis, Patrick Gallinari, and Vincent Guigue. Time Series Continuous Modeling for Imputation and Forecasting with Implicit Neural Representations. *Transactions on Machine Learning Research Journal*, April 2024. doi: 10.48550/arXiv.2306.05880. URL https://hal.science/hal-04759780.

Petr Nejedly, Vaclav Kremen, Vladimir Sladky, Mona Nasseri, Hari Guragain, Petr Klimes, Jan Cimbalnik, Yogatheesan Varatharajah, Benjamin H Brinkmann, and Gregory A Worrell. Deep-learning for seizure forecasting in canines with epilepsy. *Journal of neural engineering*, 16(3): 036031, 2019.

Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

Gerwin Schalk. EEG Motor Movement/Imagery Dataset. *PhysioNet*, 2009. doi: 10.13026/C28G6P. URL [https://doi.org/10.13026/C28G6P]. Version 1.0.0.

Gerwin Schalk, Dennis J McFarland, Thilo Hinterberger, Niels Birbaumer, and Jonathan R Wolpaw. Bci2000: a general-purpose brain-computer interface (bci) system. *IEEE Transactions on biomedical engineering*, 51(6):1034–1043, 2004.

Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csdi: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in neural information processing systems*, 34:24804–24816, 2021.

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. doi: 10.1109/TCE.2025.3528438.

Christopher Warner, Jonas Mago, JR Huml, Mohamed Osman, and Beren Millidge. Zuna: Flexible eeg superresolution with position-aware diffusion autoencoders. *arXiv preprint arXiv:2602.18478*, 2026.

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. doi: 10.1109/TAMD.2015.2431497.

A LLM USAGE

Large Language Models (LLMs) were used to aid in the writing and polishing of the manuscript. Specifically, we used an LLM to assist in refining the language, improving readability, and ensuring clarity in various sections of the paper. The model helped with tasks such as sentence rephrasing, grammar checking, and enhancing the overall flow of the text.

It is important to note that the LLM was not involved in the ideation, research methodology, or experimental design. All research concepts, ideas, and analyses were developed and conducted by the authors. The contributions of the LLM were solely focused on improving the linguistic quality of the paper, with no involvement in the scientific content or data analysis.

The authors take full responsibility for the content of the manuscript, including any text generated or polished by the LLM. We have ensured that the LLM-generated text adheres to ethical guidelines and does not contribute to plagiarism or scientific misconduct.

B ADDITIONAL RELATED WORK

B.1 EEG SPATIAL SUPER-RESOLUTION AND FLEXIBLE ELECTRODE RECONSTRUCTION

Classical spatial interpolation methods, such as spherical splines, support variable electrode locations but rely on prescribed spatial smoothness. Learning-based EEG spatial super-resolution methods instead exploit data-driven spatial and temporal dependencies through convolutional, Transformer-based, and generative architectures (Li et al., 2025a; Wang et al., 2025; Liu et al., 2026). Most of these methods assume a predefined channel inventory, making it difficult to accommodate electrode locations that are absent from the training montage.

Position-aware generative models relax this constraint. ZUNA use spatiotemporal positional encoding and masked diffusion to reconstruct variable electrode subsets and locations (Warner et al., 2026). Coordinate flexibility is therefore not unique to our work. ScalpINR instead amortizes a support-conditioned neural field whose temporal-scale gates, spectral gates, and fusion parameters are generated from the observed support set and reused across coordinate queries. This replaces iterative diffusion sampling with a single support-conditioned field instantiation followed by deterministic query evaluation.

B.2 SET-CONDITIONED NEURAL FIELDS AND IMPLICIT REPRESENTATIONS

Set-conditioned function learning is closely related to neural processes, which encode an unordered context set and predict values at arbitrary target coordinates; attentive variants additionally model query–context interactions (Garnelo et al., 2018; Kim et al., 2019). Conditional implicit neural representations and hypernetworks provide another mechanism for amortizing instance-specific functions through latent modulation or generated parameters (Ha et al., 2017). Related approaches such as TimeFlow and ImputeINR apply implicit representations to temporal imputation and forecasting (Le Naour et al., 2024; Li et al., 2025b), but do not directly address EEG reconstruction under spatially varying electrode montages.

Neural Brain Fields is more directly related: it fits a continuous EEG field separately to each recording and queries unseen spatial or temporal coordinates (Kedem et al., 2025). ScalpINR instead learns a shared mapping from an observed electrode set to a recording-specific field, allowing the field to be instantiated in a single forward pass. The key distinction is therefore not continuous coordinate querying itself, but amortized support-conditioned field generation, in which the observed montage controls temporal-scale gates, spectral gates, and their fusion, together with consistency across different support views of the same underlying recording.

C Experimental Details

C.1 Datasets and Preprocessing

We evaluate ScalpINR on three EEG datasets spanning distinct recording paradigms and downstream tasks: SEED for emotion recognition, ESAA for auditory-attention decoding, and the EEG Motor Movement/Imagery Dataset collected using BCI2000 for motor-imagery classification. For each dataset, the complete canonical montage is treated as the high-density reference. Sparse observations are constructed by retaining a subset of electrodes as the support set, while the remaining target electrodes are reconstructed from their spatial coordinates.

SEED. SEED is used to evaluate spatial reconstruction of emotion-related EEG. For random missing-channel reconstruction, we evaluate support ratios $r \in \{0.5, 0.25, 0.125\}$.

ESAA. ESAA is used to evaluate spatial reconstruction of auditory-attention EEG. Its electrode configuration additionally allows evaluation under more severe spatial sparsity, and we consider support ratios $r \in \{0.5, 0.25, 0.125, 0.0625\}$.

BCI2000. The EEG Motor Movement/Imagery Dataset collected using BCI2000 is used to evaluate spatial reconstruction of motor-imagery EEG. We consider support ratios $r \in \{0.5, 0.25, 0.125\}$.

Partitioning and leakage control. For all datasets, subject-disjoint training, validation, and test partitions are established before window extraction, ensuring that neither trials nor overlapping windows are shared across partitions. For the strict unseen-electrode protocol, held-out electrodes are determined before data-dependent preprocessing and are excluded from all training and validation inputs, reconstruction targets, losses, normalization statistics, and model-selection criteria. Consequently, signals from these electrodes are never observed during model development.

C.2 Support-Mask and Held-Out-Electrode Protocols

We consider two complementary reconstruction protocols that distinguish between sparse observations at known electrode locations and generalization to electrode coordinates that are entirely unseen during training.

Random missing-channel reconstruction. For each sample, a support set is drawn from the canonical montage according to support ratio r, and the complementary electrodes serve as query targets. Both the number and spatial configuration of observed electrodes therefore vary with the sampled support mask. For each support ratio, we evaluate 50 test-time support masks and use identical masks across all applicable methods to ensure paired comparisons. Reconstruction metrics are computed only on the complementary query electrodes.

Strict unseen-electrode reconstruction. To explicitly evaluate spatial generalization beyond the electrode inventory available during training, approximately 10% of the electrodes are permanently held out from both training and validation. These electrodes never appear as support observations or reconstruction targets during model development and are introduced only through their spatial coordinates as query locations at test time.

Under this protocol, the support ratio r is defined with respect to the remaining non-held-out electrodes. Thus, $r = 1$ denotes the setting in which all available non-held-out electrodes are observed, whereas smaller values of r additionally subsample this available support set. Every applicable trainable model is trained independently for each held-out-electrode split. This protocol therefore separates interpolation among electrode locations represented during training from genuine coordinate generalization to previously unseen locations.

C.3 Baselines and Protocol Compatibility

Our comparisons cover three complementary classes of baselines. Spherical spline interpolation provides a classical geometry-based reference. CNP and ANP represent generic set-conditioned

predictors capable of mapping an observed context set to arbitrary query locations, while ZUNA and Neural Brain Fields provide stronger coordinate-flexible EEG reconstruction baselines. We additionally compare against ImputeINR, CSDI, ESTformer, SRGDiff, and MCMA-EEG as representative learning-based reconstruction methods designed around a predefined channel layout.

Wherever their model assumptions permit, all methods use the same preprocessing, subject partitions, support masks, query targets, and evaluation pipeline. In the random missing-channel setting, all methods are evaluated on identical support–query configurations. In the strict unseen-electrode setting, fixed-layout methods are not evaluated because the permanently held-out target electrodes are absent from their training output inventory. We therefore restrict this setting to methods that can produce predictions at explicitly specified electrode coordinates.

C.4 Metrics and Aggregation

We evaluate reconstruction quality using normalized mean squared error (NMSE), Pearson correlation coefficient (PCC), and frequency-domain reconstruction error. All reconstruction metrics are computed only on query electrodes rather than on the observed support electrodes. For a fixed support mask, let $\mathcal{I}$ be the set of all test-window–query-electrode pairs, and let $\mathbf{x}_{nq}, \hat{\mathbf{x}}_{nq} \in \mathbb{R}^T$ denote the ground-truth and reconstructed waveforms for $(n, q) \in \mathcal{I}$. We compute NMSE directly in the time domain as

$$\mathrm{NMSE} = \frac{\sum_{(n,q)\in\mathcal{I}} \|\hat{\mathbf{x}}_{nq} - \mathbf{x}_{nq}\|_2^2}{\sum_{(n,q)\in\mathcal{I}} \|\mathbf{x}_{nq}\|_2^2}. \tag{15}$$

For the frequency-domain metric, we apply a forward-normalized one-sided DFT along the time axis. Specifically, for $f = 0, \ldots, F-1$ with $F = \lfloor T/2 \rfloor + 1$,

$$X_{nq}(f) = \frac{1}{T}\sum_{t=0}^{T-1} x_{nq}(t)\exp\left(-\mathrm{j}\frac{2\pi f t}{T}\right), \qquad \hat{X}_{nq}(f) = \frac{1}{T}\sum_{t=0}^{T-1} \hat{x}_{nq}(t)\exp\left(-\mathrm{j}\frac{2\pi f t}{T}\right). \tag{16}$$

The reported frequency error is the mean absolute error between these complex Fourier coefficients:

$$\mathrm{FreqMAE} = \frac{1}{|\mathcal{I}|F}\sum_{(n,q)\in\mathcal{I}}\sum_{f=0}^{F-1}\left|\hat{X}_{nq}(f) - X_{nq}(f)\right|. \tag{17}$$

Here, $|\cdot|$ denotes the complex modulus. Thus, this metric is neither a magnitude-spectrum nor a log-spectrum error; it uses every nonnegative-frequency bin with equal weight and applies no additional spectral normalization. It therefore has the same units as the preprocessed waveform. Lower NMSE and frequency error indicate better reconstruction, whereas higher PCC indicates stronger agreement with the ground-truth waveform. PCC is computed over time for each query waveform and then averaged over $\mathcal{I}$.

Each metric is first computed over all test-set query waveforms for one support mask. For random missing-channel reconstruction, the reported mean and standard deviation are then taken over the 50 test-time support masks at each support ratio. For strict unseen-electrode reconstruction, these results are additionally aggregated over the predefined held-out-electrode configurations. The same masks and target electrodes are shared across applicable methods to enable paired comparison.

In addition to pointwise reconstruction metrics, we evaluate feature-level distribution fidelity using EEG-FID and task-relevant signal preservation through downstream decoding accuracy. EEG-FID compares reconstructed and ground-truth high-density EEG in a frozen EEG representation space, as described in Section E. For downstream evaluation, observed support electrodes retain their measured signals, while reconstructed values are used for the missing electrodes. Computational efficiency is characterized by trainable parameter count, GFLOPs, and per-sample forward latency.

C.5 Training and Evaluation Setup

Architecture and optimization. Unless otherwise specified, ScalpINR uses a 96-dimensional support-token representation and a 96-dimensional coordinate-decoder hidden space. The support

encoder contains two set-attention layers with four attention heads, and the coordinate decoder contains three support-conditioned blocks. The waveform encoder uses 32 intermediate channels, while the temporal support encoder produces a 96-dimensional time-resolved context. We use 24 learnable spatial anchors, four Fourier-frequency bands for the three-dimensional electrode coordinates, and eight Fourier-frequency bands for the temporal coordinate. The functional temporal branch contains four depthwise temporal operators with kernel sizes 3, 7, 15, and 31, and the spectral branch uses eight Gaussian basis functions. The hidden dimension of the functional parameter generator is 32. The resulting ScalpINR contains 0.438M trainable parameters.

ScalpINR is trained with AdamW using an initial learning rate of 5×10^{-4} and weight decay of 10^{-4}. We use cosine learning-rate annealing over 10 epochs, with a minimum learning rate of 2.5×10^{-5}. ScalpINR is trained with a batch size of 12 and random seed 2026. Model selection is performed using validation NMSE, and the checkpoint with the lowest validation NMSE is used for final evaluation.

For the compared baselines, we follow the optimization protocols and hyperparameter settings specified in their official implementations or original papers. All methods use the same dataset partitions, support masks, query targets, and evaluation pipeline wherever their model assumptions permit.

Training losses. Following Eq. equation **??**, we use the mean absolute error (L1 distance) for waveform reconstruction, observed-channel fidelity, and the cross-support consistency distance in $\mathcal{L}_{\text{topo}}$.

For the random missing-channel experiments, we use $\lambda_{\text{fid}} = 0.25$ and $\lambda_{\text{topo}} = 0.15$. The topology-consistency weight is linearly warmed up from zero to its final value during the first six epochs.

For the strict unseen-electrode experiments, we use $\lambda_{\text{fid}} = 0.10$ and $\lambda_{\text{topo}} = 0.08$, with the topology-consistency weight linearly warmed up during the first five epochs. In all settings, permanently held-out electrodes are excluded from every training loss and from model selection.

D Random Missing-Channel Reconstruction Across Support Ratios

The main paper summarizes reconstruction performance across the evaluated support ratios. Here, we provide the complete results for each dataset and support ratio to examine how reconstruction quality changes as progressively fewer electrodes remain visible. We report NMSE, PCC, and frequency-domain reconstruction error under the same random missing-channel protocol described in Section C.2.

Table 5 reports the complete results on SEED. ScalpINR consistently achieves the lowest NMSE and highest PCC across all three support ratios, while also obtaining the lowest frequency error. More importantly, its reconstruction quality degrades only moderately as the support becomes increasingly sparse: when r decreases from 0.5 to 0.125, NMSE increases from 0.48 to 0.54 and PCC decreases from 0.72 to 0.67. At $r = 0.125$, ScalpINR outperforms the strongest competing baseline, MCMA-EEG, reducing NMSE from 0.64 to 0.54, increasing PCC from 0.60 to 0.67, and reducing frequency error from 0.11 to 0.09. These results show that ScalpINR maintains relatively stable reconstruction quality as the number of observed electrodes decreases.

Table 6 reports the complete results on ESAA, including the more severe $r = 0.0625$ setting. At $r = 0.5$, ESTformer obtains the lowest NMSE (0.34) and highest PCC (0.81), while ScalpINR achieves the lowest frequency error (0.04). As the support becomes sparser, however, ScalpINR shows substantially more stable reconstruction performance. From $r = 0.25$ onward, it achieves the lowest NMSE and highest PCC among all compared methods, while also maintaining the lowest frequency error.

The difference becomes most pronounced at $r = 0.0625$, where ScalpINR obtains an NMSE of 0.41, PCC of 0.76, and frequency error of 0.06. In comparison, MCMA-EEG, the strongest competing method at this support ratio, obtains 0.53, 0.68, and 0.10, respectively. Thus, the relative advantage of ScalpINR becomes more apparent under severe spatial undersampling, supporting its ability to reconstruct EEG from highly sparse and variable electrode supports.

Table 5: Random missing-channel reconstruction on SEED under different support ratios r. Results are reported as mean $\pm$ standard deviation over the random support masks. Lower NMSE and frequency error and higher PCC indicate better performance.

Method	$r = 0.5$			$r = 0.25$			$r = 0.125$		
	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓
Spherical Spline	0.70 ± 0.03	0.64 ± 0.03	0.36 ± 0.05	1.14 ± 0.07	0.53 ± 0.04	0.46 ± 0.06	2.06 ± 0.07	0.40 ± 0.05	0.60 ± 0.06
ZUNA	0.97 ± 0.05	0.10 ± 0.02	0.66 ± 0.06	0.98 ± 0.05	0.07 ± 0.02	0.66 ± 0.06	0.99 ± 0.05	0.02 ± 0.02	0.67 ± 0.06
Neural Brain Fields	0.72 ± 0.05	0.60 ± 0.04	0.36 ± 0.07	0.96 ± 0.05	0.55 ± 0.04	0.42 ± 0.06	1.34 ± 0.07	0.51 ± 0.05	0.50 ± 0.06
CNP	1.02 ± 0.08	0.49 ± 0.05	0.27 ± 0.05	1.41 ± 0.06	0.44 ± 0.03	0.32 ± 0.04	1.95 ± 0.08	0.42 ± 0.05	0.37 ± 0.05
ANP	1.12 ± 0.08	0.41 ± 0.05	0.28 ± 0.05	1.63 ± 0.08	0.37 ± 0.06	0.33 ± 0.06	2.51 ± 0.08	0.34 ± 0.05	0.39 ± 0.05
ImputeINR	0.61 ± 0.05	0.66 ± 0.04	0.13 ± 0.04	0.86 ± 0.05	0.57 ± 0.05	0.16 ± 0.04	1.54 ± 0.08	0.47 ± 0.05	0.21 ± 0.04
CSDI	0.75 ± 0.05	0.59 ± 0.05	0.24 ± 0.06	1.26 ± 0.08	0.48 ± 0.05	0.31 ± 0.05	2.38 ± 0.08	0.35 ± 0.05	0.42 ± 0.06
ESTformer	0.53 ± 0.04	0.68 ± 0.03	0.10 ± 0.04	0.60 ± 0.06	0.63 ± 0.04	0.11 ± 0.04	0.69 ± 0.05	0.57 ± 0.05	0.12 ± 0.04
MCMA-EEG	0.52 ± 0.03	0.69 ± 0.04	0.09 ± 0.02	0.57 ± 0.05	0.65 ± 0.04	0.10 ± 0.02	0.64 ± 0.05	0.60 ± 0.04	0.11 ± 0.02
SRGDiff	0.67 ± 0.06	0.65 ± 0.04	0.20 ± 0.04	1.08 ± 0.08	0.54 ± 0.05	0.25 ± 0.05	1.87 ± 0.08	0.42 ± 0.05	0.33 ± 0.05
ScalpINR	0.48 ± 0.04	0.72 ± 0.07	0.07 ± 0.01	0.50 ± 0.05	0.70 ± 0.07	0.08 ± 0.01	0.54 ± 0.04	0.67 ± 0.08	0.09 ± 0.02

Table 6: Random missing-channel reconstruction on ESAA under different support ratios r. Results are reported as mean $\pm$ standard deviation over the random support masks. Lower NMSE and frequency error and higher PCC indicate better performance.

Method	$r = 0.5$			$r = 0.25$			$r = 0.125$			$r = 0.0625$		
	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓
Spherical Spline	0.39 ± 0.03	0.78 ± 0.04	0.43 ± 0.06	0.47 ± 0.04	0.75 ± 0.05	0.47 ± 0.07	0.55 ± 0.03	0.70 ± 0.04	0.51 ± 0.06	0.69 ± 0.03	0.53 ± 0.05	0.57 ± 0.06
ZUNA	0.75 ± 0.05	0.50 ± 0.04	0.56 ± 0.06	0.78 ± 0.06	0.45 ± 0.05	0.56 ± 0.08	0.75 ± 0.05	0.50 ± 0.05	0.57 ± 0.06	0.80 ± 0.06	0.43 ± 0.05	0.59 ± 0.08
Neural Brain Fields	0.38 ± 0.05	0.79 ± 0.04	0.41 ± 0.06	0.43 ± 0.06	0.77 ± 0.04	0.44 ± 0.06	0.48 ± 0.05	0.76 ± 0.04	0.46 ± 0.06	0.54 ± 0.05	0.74 ± 0.05	0.49 ± 0.05
CNP	0.40 ± 0.04	0.70 ± 0.03	0.22 ± 0.04	0.47 ± 0.05	0.68 ± 0.02	0.23 ± 0.03	0.54 ± 0.05	0.66 ± 0.04	0.25 ± 0.05	0.62 ± 0.06	0.63 ± 0.04	0.27 ± 0.06
ANP	0.58 ± 0.05	0.51 ± 0.05	0.28 ± 0.07	0.65 ± 0.06	0.48 ± 0.04	0.29 ± 0.05	0.72 ± 0.05	0.46 ± 0.05	0.31 ± 0.05	0.79 ± 0.07	0.42 ± 0.05	0.32 ± 0.05
ImputeINR	0.37 ± 0.05	0.80 ± 0.04	0.10 ± 0.06	0.47 ± 0.05	0.75 ± 0.04	0.12 ± 0.04	0.57 ± 0.05	0.67 ± 0.04	0.13 ± 0.04	0.67 ± 0.06	0.58 ± 0.05	0.14 ± 0.04
CSDI	0.41 ± 0.05	0.75 ± 0.05	0.22 ± 0.05	0.57 ± 0.07	0.62 ± 0.04	0.26 ± 0.06	0.68 ± 0.05	0.57 ± 0.05	0.28 ± 0.05	0.77 ± 0.05	0.46 ± 0.06	0.30 ± 0.05
ESTformer	0.34 ± 0.06	0.81 ± 0.04	0.09 ± 0.04	0.46 ± 0.05	0.76 ± 0.04	0.11 ± 0.04	0.58 ± 0.05	0.66 ± 0.04	0.12 ± 0.04	0.68 ± 0.05	0.57 ± 0.05	0.13 ± 0.05
MCMA-EEG	0.38 ± 0.05	0.80 ± 0.03	0.07 ± 0.03	0.40 ± 0.06	0.78 ± 0.06	0.08 ± 0.04	0.45 ± 0.05	0.75 ± 0.04	0.09 ± 0.05	0.53 ± 0.04	0.68 ± 0.04	0.10 ± 0.06
SRGDiff	0.37 ± 0.04	0.79 ± 0.04	0.05 ± 0.05	0.49 ± 0.05	0.74 ± 0.05	0.06 ± 0.02	0.59 ± 0.05	0.68 ± 0.04	0.07 ± 0.02	0.72 ± 0.06	0.50 ± 0.05	0.08 ± 0.03
ScalpINR	0.36 ± 0.05	0.80 ± 0.06	0.04 ± 0.01	0.37 ± 0.04	0.79 ± 0.06	0.05 ± 0.02	0.39 ± 0.04	0.78 ± 0.07	0.05 ± 0.01	0.41 ± 0.04	0.76 ± 0.06	0.06 ± 0.01

Table 7 presents the complete results on BCI2000. ScalpINR achieves the lowest NMSE at every support ratio and remains best or tied for best in PCC and frequency error. At $r = 0.5$, ScalpINR and MCMA-EEG both achieve a PCC of 0.89, while ScalpINR obtains a lower NMSE (0.20 versus 0.21) and frequency error (0.04 versus 0.05). As the support ratio decreases, the difference between the two methods becomes larger. ScalpINR's NMSE increases only from 0.20 at $r = 0.5$ to 0.25 at $r = 0.125$, whereas MCMA-EEG increases from 0.21 to 0.31. At $r = 0.125$, ScalpINR further improves PCC from 0.83 to 0.86 and reduces frequency error from 0.08 to 0.06 relative to MCMA-EEG. Overall, the BCI2000 results further demonstrate that ScalpINR maintains strong reconstruction fidelity as the observed montage becomes increasingly sparse.

E Feature-Level Distribution Fidelity

Pointwise reconstruction metrics quantify sample-wise waveform agreement, but do not directly measure whether reconstructed EEG follows the representation distribution of real high-density recordings. We therefore complement the signal-level evaluation with EEG-FID, following the feature-level evaluation protocol of SRGDiff (Liu et al., 2026). EEG-FID measures the distributional discrepancy between reconstructed and ground-truth EEG in a learned EEG representation space, providing a complementary assessment of whether the reconstructed signals preserve the statistical structure of real EEG.

Following SRGDiff, we train an EEGNet feature extractor separately for each dataset using only its training split. The feature extractor is then frozen throughout evaluation, so that neither reconstructed test signals nor ground-truth test signals are used to update its parameters. Let $f(\cdot)$ denote the frozen EEGNet encoder. Given ground-truth high-density EEG samples $\{X_i\}_{i=1}^{N}$ and their corresponding reconstructions $\{\hat{X}_i\}_{i=1}^{N}$, we extract feature representations

$$\mathbf{h}_i = f(X_i), \qquad \hat{\mathbf{h}}_i = f(\hat{X}_i). \tag{18}$$

We then estimate the empirical means and covariance matrices of the two feature distributions, denoted by $(\boldsymbol{\mu}_{\mathrm{gt}}, \boldsymbol{\Sigma}_{\mathrm{gt}})$ and $(\boldsymbol{\mu}_{\mathrm{rec}}, \boldsymbol{\Sigma}_{\mathrm{rec}})$, respectively. EEG-FID is computed as the Fréchet distance

$$\begin{aligned} \text{EEG-FID} = \; & \|\boldsymbol{\mu}_{\mathrm{gt}} - \boldsymbol{\mu}_{\mathrm{rec}}\|_2^2 \\ & + \operatorname{Tr}\left(\boldsymbol{\Sigma}_{\mathrm{gt}} + \boldsymbol{\Sigma}_{\mathrm{rec}} - 2\left(\boldsymbol{\Sigma}_{\mathrm{gt}}\boldsymbol{\Sigma}_{\mathrm{rec}}\right)^{1/2}\right). \end{aligned} \tag{19}$$

Table 7: Random missing-channel reconstruction on BCI2000 under different support ratios r. Results are reported as mean $\pm$ standard deviation over the random support masks. Lower NMSE and frequency error and higher PCC indicate better performance.

Method	$r = 0.5$			$r = 0.25$			$r = 0.125$		
	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓	NMSE ↓	PCC ↑	Freq. Err. ↓
Spherical Spline	0.27 ± 0.04	0.86 ± 0.04	0.33 ± 0.06	0.41 ± 0.02	0.80 ± 0.03	0.40 ± 0.07	0.68 ± 0.03	0.71 ± 0.04	0.51 ± 0.06
ZUNA	0.96 ± 0.05	0.20 ± 0.01	0.37 ± 0.05	0.99 ± 0.06	0.03 ± 0.02	0.38 ± 0.06	0.99 ± 0.05	0.01 ± 0.02	0.39 ± 0.05
Neural Brain Fields	0.90 ± 0.05	0.33 ± 0.05	0.29 ± 0.05	0.90 ± 0.05	0.32 ± 0.05	0.30 ± 0.06	0.90 ± 0.05	0.32 ± 0.05	0.32 ± 0.05
CNP	0.99 ± 0.05	0.11 ± 0.02	0.25 ± 0.05	0.99 ± 0.07	0.11 ± 0.02	0.26 ± 0.05	0.99 ± 0.05	0.10 ± 0.02	0.28 ± 0.05
ANP	1.00 ± 0.04	0.06 ± 0.02	0.40 ± 0.06	1.00 ± 0.04	0.06 ± 0.02	0.42 ± 0.07	1.00 ± 0.05	0.06 ± 0.02	0.44 ± 0.06
ImputeINR	0.23 ± 0.05	0.88 ± 0.04	0.08 ± 0.04	0.35 ± 0.05	0.85 ± 0.05	0.10 ± 0.04	0.68 ± 0.05	0.71 ± 0.03	0.13 ± 0.04
CSDI	0.33 ± 0.06	0.82 ± 0.04	0.09 ± 0.04	0.51 ± 0.05	0.77 ± 0.06	0.11 ± 0.04	0.82 ± 0.05	0.67 ± 0.04	0.14 ± 0.05
ESTformer	0.24 ± 0.07	0.87 ± 0.04	0.06 ± 0.04	0.48 ± 0.05	0.81 ± 0.04	0.08 ± 0.05	0.70 ± 0.05	0.73 ± 0.04	0.11 ± 0.04
MCMA-EEG	0.21 ± 0.04	0.89 ± 0.06	0.05 ± 0.02	0.25 ± 0.06	0.87 ± 0.04	0.06 ± 0.02	0.31 ± 0.05	0.83 ± 0.04	0.08 ± 0.02
SRGDiff	0.25 ± 0.05	0.87 ± 0.05	0.05 ± 0.02	0.37 ± 0.04	0.83 ± 0.05	0.06 ± 0.03	0.71 ± 0.05	0.69 ± 0.04	0.08 ± 0.02
ScalpINR	0.20 ± 0.05	0.89 ± 0.07	0.04 ± 0.01	0.23 ± 0.04	0.88 ± 0.07	0.04 ± 0.01	0.25 ± 0.04	0.86 ± 0.06	0.06 ± 0.01

Table 8: EEG-FID under the random missing-channel reconstruction setting corresponding to Table 1. The same test splits, support masks, and reconstructed signals are used for evaluation. Lower EEG-FID indicates better agreement with the feature distribution of real high-density EEG. Best results are in bold and second-best results are underlined.

Method	EEG-FID ↓		
	SEED	ESAA	BCI2000
Spherical Spline	291	146	96
ZUNA	260	181	133
Neural Brain Fields	228	87	118
CNP	68	53	100
ANP	139	104	129
ImputeINR	46	28	26
CSDI	91	51	30
ESTFormer	21	23	21
SRGDiff	60	15	15
MCMA-EEG	23	16	14
ScalpINR	20	11	10

A lower EEG-FID indicates that reconstructed EEG is closer to real high-density EEG in the learned feature distribution.

We evaluate EEG-FID on exactly the same reconstructed test signals used in Tables 1 and 2, without changing the data splits, support masks, or reconstruction outputs. Thus, the following results provide a feature-level counterpart to the signal-level metrics reported above.

Table 8 shows that the signal-level reconstruction gains of ScalpINR are accompanied by improved feature-level distribution fidelity. ScalpINR achieves the lowest EEG-FID on all three datasets, with scores of 20, 11, and 10 on SEED, ESAA, and BCI2000, respectively. In particular, its advantage over methods with competitive pointwise reconstruction suggests that the recovered signals are not merely close to individual targets in waveform space, but also better preserve the representation distribution of real high-density EEG.

The distribution-level advantage becomes particularly informative under the strict unseen-electrode protocol. As shown in Table 9, ScalpINR obtains the lowest EEG-FID across all datasets and support ratios, including when target electrode coordinates are entirely absent from training. For example, at $r = 1$, ScalpINR reduces EEG-FID to 17, 25, and 14 on SEED, ESAA, and BCI2000, respectively. The advantage is retained as the support becomes progressively sparser. These results complement the NMSE and PCC results in Table 2 and indicate that coordinate-queryable reconstruction at unseen locations preserves not only pointwise signal fidelity but also the feature distribution characteristic of real EEG.

Table 9: EEG-FID under strict unseen-electrode reconstruction corresponding to Table 2. Held-out electrodes are never observed during training and are evaluated only as query targets. The same held-out splits, support masks, and reconstructed signals as in Table 2 are used. Lower is better. Best results are in bold and second-best results are underlined.

r	Method	EEG-FID ↓		
		SEED	ESAA	BCI2000
1	Spherical Spline	32	27	30
	CNP	145	120	137
	ANP	141	129	105
	ZUNA	88	99	115
	Neural Brain Fields	126	115	128
	ScalpINR	17	25	14
0.5	Spherical Spline	70	77	49
	CNP	124	126	118
	ANP	155	133	132
	ZUNA	85	100	115
	Neural Brain Fields	136	99	103
	ScalpINR	26	17	13
0.25	Spherical Spline	62	61	79
	CNP	140	129	122
	ANP	149	122	137
	ZUNA	156	72	152
	Neural Brain Fields	101	93	73
	ScalpINR	25	14	45
0.125	Spherical Spline	153	111	72
	CNP	146	136	126
	ANP	158	142	138
	ZUNA	61	91	123
	Neural Brain Fields	91	70	92
	ScalpINR	48	43	61

F Downstream Decoding Across Support Ratios

To complement the downstream evaluation at $r = 0.25$ in the main paper, we further evaluate task-specific decoding under both denser ($r = 0.5$) and sparser ($r = 0.125$) electrode supports. For each dataset, the task-specific classifier is trained and selected using only real high-density EEG from the training and validation partitions and is then frozen for all reconstruction methods. At test time, observed support electrodes retain their measured signals, while only the missing electrodes are replaced by their reconstructed values. Accuracy is reported as mean (%) $\pm$ standard deviation. The GT-HD reference is evaluated directly on the original high-density EEG and is therefore independent of the support ratio.

Across both additional support ratios, ScalpINR achieves the highest downstream accuracy among all reconstruction methods on all three datasets. On SEED, it obtains 41.33% and 39.74% accuracy at $r = 0.5$ and $r = 0.125$, respectively, compared with 38.24% and 36.62% for ESTformer, the strongest competing reconstruction method. On ESAA, ScalpINR achieves 54.16% and 52.61%, compared with 52.13% and 50.64% for MCMA-EEG. The same pattern is observed on BCI2000, where ScalpINR reaches 70.43% and 68.86%, compared with 68.94% and 67.38% for MCMA-EEG.

Together with the $r = 0.25$ results reported in the main paper, these results show that the downstream performance of ScalpINR remains stable across a broad range of support densities. Importantly, the reconstructed signals remain close to the GT-HD reference even under severe spatial undersampling. At $r = 0.125$, the absolute accuracy gaps between ScalpINR and GT-HD are only 0.63, 0.51, and 1.98 percentage points on SEED, ESAA, and BCI2000, respectively. This consistency between reconstruction quality and downstream decoding performance indicates that the signals recovered by ScalpINR preserve task-relevant EEG structure even when only a small fraction of the original montage is observed.

Table 10: Downstream decoding accuracy (%) under additional support ratios on SEED, ESAA, and BCI2000. The corresponding results at $r = 0.25$ are reported in the main paper. GT-HD denotes evaluation on the original high-density EEG, while Sparse input denotes evaluation without reconstructing the missing electrodes.

	SEED		ESAA		BCI2000	
Input / Reconstruction	$r = 0.5$	$r = 0.125$	$r = 0.5$	$r = 0.125$	$r = 0.5$	$r = 0.125$
GT-HD (reference)	40.37 ± 5.24	40.37 ± 5.24	53.12 ± 3.06	53.12 ± 3.06	70.84 ± 3.42	70.84 ± 3.42
Sparse input	36.05 ± 4.81	34.91 ± 5.17	53.02 ± 3.64	50.88 ± 3.24	52.95 ± 4.34	51.37 ± 4.06
Spherical Spline	34.21 ± 5.11	32.86 ± 4.96	51.08 ± 3.18	49.62 ± 3.01	54.63 ± 3.95	53.14 ± 3.76
ImputeINR	34.87 ± 5.22	33.28 ± 5.02	51.67 ± 3.26	50.05 ± 3.08	67.08 ± 3.63	65.42 ± 3.44
CSDI	34.08 ± 5.18	33.31 ± 4.99	51.02 ± 3.37	49.71 ± 3.19	51.47 ± 4.20	50.03 ± 4.01
ESTformer	38.24 ± 4.91	36.62 ± 4.72	52.01 ± 2.95	50.58 ± 2.77	68.35 ± 3.69	66.71 ± 3.50
SRGDiff	33.61 ± 5.38	32.79 ± 5.21	51.74 ± 3.29	50.16 ± 3.11	51.26 ± 4.31	49.74 ± 4.12
MCMA-EEG	35.42 ± 4.99	33.91 ± 4.81	52.13 ± 3.02	50.64 ± 2.84	68.94 ± 3.56	67.38 ± 3.37
ZUNA	34.55 ± 4.94	33.11 ± 4.76	51.24 ± 3.10	49.88 ± 2.92	51.71 ± 4.06	50.28 ± 3.87
Neural Brain Fields	31.93 ± 5.49	30.58 ± 5.31	50.46 ± 3.46	49.03 ± 3.28	46.18 ± 4.55	44.69 ± 4.36
ANP	32.47 ± 5.11	31.06 ± 4.93	50.92 ± 4.36	49.51 ± 4.18	44.72 ± 3.94	43.21 ± 3.75
CNP	32.19 ± 5.27	30.82 ± 5.08	50.71 ± 3.64	49.28 ± 3.46	46.57 ± 4.23	45.06 ± 4.04
ScalpINR	41.33 ± 5.36	39.74 ± 5.18	54.16 ± 3.20	52.61 ± 3.02	70.43 ± 3.44	68.86 ± 3.25

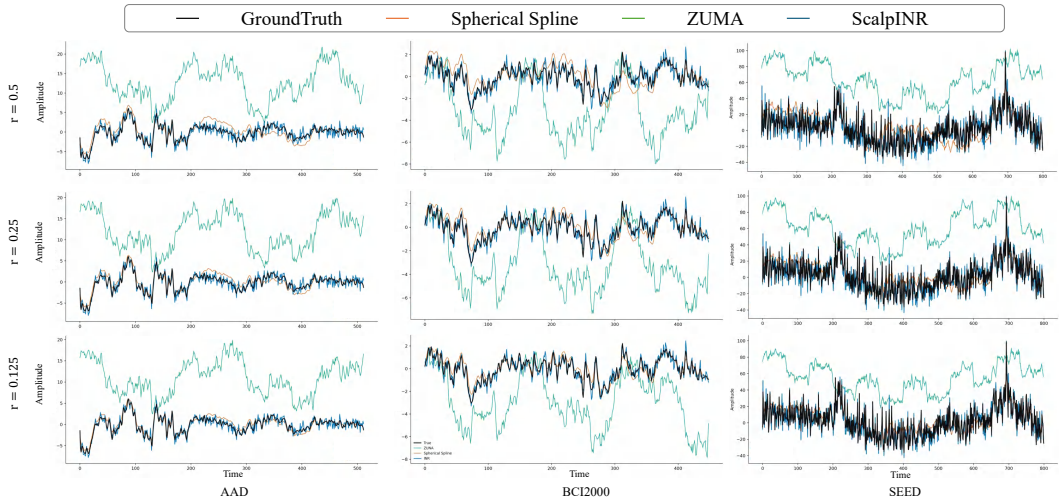


Figure 2: Qualitative waveform comparison under strict unseen-electrode generation. The target electrodes are permanently held out from the support set during training and are queried only during evaluation. Rows correspond to different support ratios r, and columns show representative examples from ESAA, BCI2000, and SEED. Compared with spherical spline interpolation and ZUNA, ScalpINR produces waveforms that are more consistent with the ground truth, demonstrating its ability to synthesize EEG signals at unseen electrode coordinates.

F.1 Qualitative Visualization

Figure 2 visualizes strict unseen-electrode generation. Under this more challenging setting, spherical spline interpolation can capture coarse trends but often over-smooths local waveform details, while ZUNA shows clear deviations from the ground truth in several cases. In contrast, ScalpINR remains closely aligned with the real signals across SEED, ESAA, and BCI2000, even under low support ratios. These examples are consistent with the quantitative held-out results and show that ScalpINR preserves waveform morphology when queried at electrode locations excluded from training.

Table 11: Computational cost and per-sample forward latency for a SEED window with 62 electrodes and 800 time samples at support ratio $r = 0.125$. All newly added coordinate-queryable and set-conditioned baselines are profiled under the same setting on an RTX 4090. Runtime denotes the mean forward latency per sample. Spherical spline is non-parametric and therefore has no trainable parameter count or GFLOP measurement.

Method	#Params (M)	GFLOPs	Runtime (ms)
Spherical Spline	N/A	N/A	1.842
ZUNA	380+	46.218	118.734
Neural Brain Fields	0.504	49.800	6.732
CNP	0.190	0.276	0.903
ANP	0.431	0.310	1.551
ImputeINR	1.247	2.684	7.913
CSDI	4.414	38.572	286.441
ESTformer	2.111	4.302	5.039
SRGDiff	8.421	5.439	92.541
MCMA-EEG	4.836	4.917	6.284
ScalpINR	0.438	0.270	1.804

G Efficiency Analysis

We compare model size, computational cost, and inference latency under the same SEED input shape and support ratio. For coordinate-queryable methods, the reported runtime includes both support encoding and querying all missing electrodes. All methods are evaluated using the same input configuration to provide a consistent comparison of computational efficiency.

As shown in Table 11, ScalpINR contains only 0.438M trainable parameters and requires 0.270 GFLOPs per sample, making it one of the most compact and computationally efficient learned reconstruction models considered here. Its parameter count is substantially smaller than those of the fixed-layout reconstruction baselines and ZUNA, while remaining comparable to the lightweight set-conditioned controls.

ScalpINR achieves a per-sample forward latency of 1.804 ms. Although the lightweight CNP and ANP controls are faster at 0.903 ms and 1.551 ms, respectively, ScalpINR remains considerably faster than the other learned reconstruction baselines. In particular, compared with SRGDiff, its latency decreases from 92.541 ms to 1.804 ms, corresponding to a reduction of approximately 98.1%. It is also substantially faster than ZUNA (118.734 ms), CSDI (286.441 ms), ImputeINR (7.913 ms), Neural Brain Fields (6.732 ms), MCMA-EEG (6.284 ms), and ESTformer (5.039 ms).

Notably, ScalpINR achieves this efficiency while retaining support-dependent field instantiation and coordinate-queryable reconstruction. Once the support set has been encoded and the recording-specific field parameters have been generated, missing electrodes can be reconstructed through deterministic coordinate queries without iterative test-time optimization. Together with the reconstruction results reported in the main text, these results indicate that ScalpINR provides a favorable accuracy–efficiency trade-off rather than achieving its reconstruction gains through increased model size or computational cost.