

NEURAL BRAIN FIELDS: A NERF-INSPIRED APPROACH FOR GENERATING NONEXISTENT EEG ELECTRODES

Anonymous authors

Paper under double-blind review

ABSTRACT

Electroencephalography (EEG) data present unique modeling challenges because recordings vary in length, exhibit very low signal to noise ratios, differ significantly across participants, drift over time within sessions, and are rarely available in large and clean datasets. Consequently, developing deep learning methods that can effectively process EEG signals remains an open and important research problem. To tackle this problem, this work presents a new method inspired by Neural Radiance Fields (NeRF Mildenhall et al. (2021)). In computer vision, NeRF techniques train a neural network to memorize the appearance of a 3D scene and then uses its learned parameters to render and edit the scene from any viewpoint. We draw an analogy between the discrete images captured from different viewpoints used to learn a continuous 3D scene in NeRF, and EEG electrodes positioned at different locations on the scalp, which are used to infer the underlying representation of continuous neural activity. Building on this connection, we show that a neural network can be trained on a single EEG sample in a NeRF style manner to produce a fixed size and informative weight vector that encodes the entire signal. Moreover, via this representation we can render the EEG signal at previously unseen time steps and spatial electrode positions. We demonstrate that this approach enables continuous visualization of brain activity at any desired resolution, including ultra high resolution, and reconstruction of raw EEG signals. Finally, our empirical analysis shows that this method can effectively simulate nonexistent electrodes data in EEG recordings, allowing the reconstructed signal to be fed into standard EEG processing networks to improve performance.

1 INTRODUCTION

Electroencephalography (EEG) records millisecond scale voltage fluctuations generated by groups of brain cells and remains the most widely used non-invasive technique for tracking real-time brain activity. Its scientific and societal value rests on three pillars: (i) it enables real-time brain technology interaction through brain-computer interfaces and neurofeedback systems (Yun, 2024), (ii) it serves as a clinical gold standard for diagnosing and monitoring neurological conditions such as epilepsy, sleep disorders, and coma, and (iii) it provides a high temporal resolution discovery tool for cognitive and systems neuroscience, illuminating processes such as perception, attention, and memory through event related potentials and oscillatory analyses.

EEG signal processing is inherently difficult due to its low signal to noise ratio, which is susceptible to various artifacts (Niedermeyer & da Silva, 2005), and its limited spatial resolution caused by volume conduction (Sanei & Chambers, 2013). This spatial limitation arises from sparse sensor placement and the smearing of signals across the scalp (Nunez & Srinivasan, 2006), constituting a core constraint of EEG despite its widespread use and excellent temporal resolution. As a result, the signals captured at the scalp are coarse, spatially blurred reflections of the underlying cortical dynamics, with the limited electrode count often insufficient to resolve the true complexity of neural activity patterns.

Early EEG pipelines relied on band-pass filtering, Independent Component Analysis, and hand crafted features such as power spectra or common spatial patterns, which were then fed to shallow classifiers. Over the past decade, deep learning has transformed EEG research. Convolutional

neural networks extract spectrotemporal patterns directly from raw signals (Schirrneister et al., 2017; Lawhern et al., 2018b; Svantesson et al., 2021a), recurrent models capture temporal dynamics (Ruffini et al., 2016; Roy et al., 2019), graph neural networks leverage the spatial topology of the electrode array (Klepl et al., 2024; Demir et al., 2021), and transformer architectures combine attention mechanisms with self supervised pre-training to improve cross subject generalization (Wang et al., 2024). Together, these approaches now define the SoTA in many EEG tasks such as classification, super resolution, denoising, and others.

Yet, despite steady gains, EEG research has not achieved breakthroughs. Two major bottlenecks continue to hinder comparable advances: First, EEG datasets are limited in both size and spatial coverage, leaving much intracranial brain activity unrecorded. Second, recordings vary across participants, sessions, hardware setups, and environments, and this heterogeneity hinders cross sample and cross dataset generalization. To overcome these bottlenecks, inspired by NeRF (Mildenhall et al. (2021)), we introduce Neural Brain Fields (NBF), a deep learning framework that implicitly learns the underlying dynamics of each recording by mapping a continuous four dimensional space and time field to voltage values. Once trained, the model can estimate voltage at any intracranial location and time point, including regions where no physical electrodes are present, and directly provides tools for reconstruction and super resolution. Our method draws inspiration from techniques in the domain of 3-D scene representation, which use a set of 2-D images to encode the scene into neural weights by learning a continuous function over space and viewing direction, thus enabling scene rendering from arbitrary viewpoints.

The core idea is to overcome the challenges of cross sample generalization caused by non uniform EEG representations by learning the underlying representation of each sample separately, rather than learning shared representations across samples. Moreover, we hypothesize that the remarkable sample efficiency of NeRF based techniques can help mitigate the data constraints commonly found in the EEG domain by producing less noisy signals and supporting high resolution reconstructions, including at locations where no physical electrodes are present and direct access is not possible.

Our main contributions encompass the following three aspects: (i) We propose Neural Brain Fields (NBF), a NeRF like approach designed to learn a latent representation of individual EEG recordings, which can be used after training to render voltage signals at arbitrary spatial and temporal locations. (ii) We empirically demonstrate that, with appropriate architectural choices, such as suitable positional encoding and effective normalization, NBF can accurately predict voltages at non existent electrode positions. This capability amplifies the EEG signal and improves the performance of downstream neural networks on tasks such as decoding speech perception from brain activity, emotion classification, and a broad range of other cognitive and affective state decoding tasks. (iii) Our technique is broadly applicable and provides a set of tools for a wide range of EEG processing tasks, including super resolution, signal reconstruction, learning fixed size and continuous representations through model weights, and producing high quality, smooth visualizations.

2 BACKGROUND AND RELATED WORKS

Neural Radiance Fields NeRF (Mildenhall et al., 2021)) is a widely used technique in computer vision for learning 3D scene representations, enabling both photorealistic rendering and scene manipulation. Given a set of posed images of a single scene denoted by s , NeRF constructs a Ray sample dataset by casting rays through image pixels and sampling 3D points along each ray. This dataset can be formally defined as:

$$\mathcal{D}_s = \left\{ \left(\mathbf{x}_\ell^{(s)}, \theta_s, \phi_s, y_s \right) \mid \ell = 1, \dots, L \right\}, \quad \mathbf{x}_\ell^{(s)} \in \mathbb{R}^3, \theta_s \in [0, 2\pi), \phi_s \in [0, \pi], y_s \in \mathbb{R}^3 \quad (1)$$

where x, y, z are physical 3D coordinates in the visual space and Θ and Φ denote the azimuthal and polar angles, respectively. Using this dataset, NeRF learns a continuous volumetric scene representation via a neural function f_w , typically implemented as a simple MLP. The function maps each 3D location $\mathbf{x} \in \mathbb{R}^3$ and viewing direction (θ, Φ) to a color and volume density:

$$f_w(\mathbf{x}, \mathbf{d}) \mapsto (\mathbf{c}, \sigma), \quad \mathbf{c} \in \mathbb{R}^3, \sigma \in \mathbb{R}_{\geq 0} \quad (2)$$

The predicted outputs are aggregated along rays using differentiable volumetric rendering, resulting in a predicted color for each data point. The rendered color $\hat{\mathbf{C}}(\mathbf{r})$ is then compared to the ground truth color y_s , and the network is trained to reconstruct the observed pixel color. Through this process, the network implicitly learns to encode both the geometry and appearance of the scene within its

108 weights. Once the network is trained, it can render the scene from novel viewpoints by querying the
109 learned function. This rendering process enables photorealistic novel view synthesis and has been
110 applied to a wide range of tasks, including 3D reconstruction, relighting, scene editing (Yuan et al.,
111 2022), virtual reality, and content creation for films and games. Inspired by this line of work, our
112 approach aims to learn the underlying dynamics of individual EEG signals to obtain compact, fixed
113 size neural representations encoded directly in the network’s weights. Additionally, we hypothesize
114 that advanced NeRF based rendering techniques such as super resolution, nonexistent electrode syn-
115 thesis, and EEG signal reconstruction can play a crucial role in amplifying the underlying sources of
116 EEG activity. These techniques may help mitigate data inefficiency and noise, two major challenges
117 in EEG modeling.

118 **DL-based EEG Processing** In recent years, many DL methods have been proposed to improve
119 EEG analysis, classification, and generation. This line of work spans a wide spectrum from
120 simple, shallow networks to large scale architectures based on pre-trained Transformers, such as
121 EEGPT (Wang et al., 2024). Notable examples include: (i) CNN-based methods (Schirrmester
122 et al., 2017; Lawhern et al., 2018b; Svantesson et al., 2021a), (ii) RNN-based methods (Tan et al.,
123 2020), (iii) GNN-based methods (Klepl et al., 2024; Demir et al., 2021), and (iv) Transformer mod-
124 els (Wan et al., 2023; Wang et al., 2024), and (iv) modern architectures such as Mamba-based mod-
125 els (Gui et al., 2024; Yang & Jia, 2024) and diffusion models (Lopez et al., 2025). Moreover,
126 adaptation of DL tools such as design dedicated data augmentation and positional encoding tech-
127 niques have also been explored for EEG (Rommel et al., 2022; Torma & Szegletes, 2025; Jia et al.,
128 2024). However, rather than investigating which architectures can improve classification accuracy
129 or cross sample generalization, we focus on a less explored research direction: Investigate whether
130 it is possible to learn the underlying source of a single EEG sample, and use this insight to build
131 robust EEG representations that enable signal amplification via rendering techniques.

132 **Relevant Advances in NeRF and Neural EEG Modeling** Although our work focuses on adapt-
133 ing NeRF style representations to EEG data, it builds on several recent developments in neural
134 rendering. Mip-NeRF (Barron et al., 2021) and Mip-NeRF 360 (Barron et al., 2022) introduced
135 multiscale sampling and anti-aliasing for unbounded scenes. Instant-NGP (Müller et al., 2022),
136 KiloNeRF (Reiser et al., 2021), FastNeRF (Garbin et al., 2021), and TinyNeRF (Zhao et al., 2023)
137 proposed runtime and architectural optimizations for efficient rendering. Editable NeRF (Yuan et al.,
138 2022) enabled relighting and geometry editing, while Mega-NeRF (Turki et al., 2022) demonstrated
139 large scale scene reconstruction. These advances highlight the scalability and flexibility of neural
140 fields and motivate their adaptation to EEG tasks such as spatial super resolution, signal reconstruc-
141 tion, and virtual electrode generation.

141 Parallel progress in neural EEG modeling has introduced generative and representation learning
142 methods, including diffusion based approaches have been proposed for EEG denoising (Torma &
143 Szegletes, EasyChair, 2023), and deep networks have been used for EEG to image reconstruction
144 and visual representation learning (Singh et al., 2023; 2024). Several studies have also explored
145 conditioning NeRF style or 3D generative models on EEG signals for object and scene reconstruc-
146 tion (Xiang et al., 2024; Ge et al., 2025). Finally, foundation models such as LUNA leverage NeRF
147 inspired positional encodings for robust cross subject decoding (Döner et al.).

148 Unlike prior approaches that rely on generative models or representation learning for tasks such as
149 denoising and classification, our method directly models the continuous spatial structure of EEG
150 voltage across the scalp using neural fields.

151 **EEG Signal Super Resolution** In the context of EEG, super resolution refers to the task of recon-
152 structing or synthesizing EEG signals at higher spatial resolution than what was originally recorded,
153 typically by estimating signals at unmeasured electrode locations. Several recent studies have ex-
154 plored learnable approaches for enhancing EEG spatial resolution, including the generation of sig-
155 nals at unmeasured or missing electrode positions. Notably, Svantesson et al. (2021b) introduced
156 *Virtual EEG-electrodes: Convolutional neural networks as a method for upsampling or restoring*
157 *channels*, which proposed CNN based mappings to restore high density EEG from reduced mon-
158 tages (e.g., 4→21, 14→21, or 20→21 electrodes). While their work is an important contribution
159 toward EEG upsampling, it relies on fixed mappings defined for a specific dataset and does not
160 generalize beyond the predefined input output configurations. Similarly, the work *Deep EEG super-*
161 *resolution: Up sampling EEG spatial resolution with Generative Adversarial Networks* Corley &
Huang (2018) applied GANs to EEG super resolution on a limited dataset with constrained input

162 output mappings. Although promising, the authors note that training GANs for EEG data was com-
 163 putationally intensive and challenging to optimize, reflecting the inherent complexity of learning
 164 realistic signal distributions across different subjects and electrode locations.

165 In contrast, our method is designed for per subject and per recording modeling, allowing the genera-
 166 tion of EEG signals at arbitrary spatial locations without being restricted to fixed channel to channel
 167 mappings. This flexibility enables the synthesis of electrodes across the entire scalp topology, adapt-
 168 ing to the substantial variability in individual head shapes and signal characteristics. By leveraging
 169 neural implicit representations, our approach learns a continuous spatial field of EEG activity that
 170 generalizes to unseen positions during inference, offering a subject and record specific alternative
 171 for EEG signal super resolution.

172 3 METHOD

173 We begin by defining the learning problem and its motivation in Sec. 3.1. We then present the
 174 architecture designed to solve this problem in Sec. 3.2. Finally, we describe how the learned per-
 175 sample network can be applied to multiple rendering tasks in Sec. 3.3.

176 3.1 LEARNING PROBLEM

177 To learn the underlying latent representation of EEG, we introduce a NeRF-inspired DL architecture
 178 specifically tailored to model spatiotemporal EEG dynamics. The objective is to approximate a
 179 continuous function $f : \mathbb{R}^4 \rightarrow \mathbb{R}$, which maps a four dimensional vector comprising three spatial
 180 coordinates (x, y, z) and one temporal coordinate t to a scalar voltage value v : $(x, y, z, t) \mapsto v$.

181 Motivated by NeRF and the inherent difficulty of cross-sample generalization in EEG, where signal
 182 characteristics vary widely across participants, sessions, and recording conditions, we fit a separate
 183 network for each sample. This approach circumvents the variability introduced by such non-uniform
 184 representations and instead focuses on faithfully capturing the latent structure of individual record-
 185 ings, leading to more robust and accurate reconstructions.

186 **Input Normalization** To ensure numerical stability and consistent scaling across spatial and tem-
 187 poral dimensions, we apply min-max normalization to the coordinate vector (x, y, z, t) . The spatial
 188 coordinates are normalized jointly using their minimum and maximum values, while the temporal
 189 dimension is normalized independently:

$$190 (x', y', z', t') = \left(2 \frac{x - s_{\min}}{s_{\max} - s_{\min}} - 1, 2 \frac{y - s_{\min}}{s_{\max} - s_{\min}} - 1, 2 \frac{z - s_{\min}}{s_{\max} - s_{\min}} - 1, \frac{t - t_{\min}}{t_{\max} - t_{\min}} \right), \quad (3)$$

191 where $s_{\min} = \min\{x, y, z\}$, and $s_{\max} = \max\{x, y, z\}$. The normalized input is denoted as $\mathbf{v}' =$
 192 $(x', y', z', t') \in \mathbb{R}^4$.

193 **Voltage Normalization and Denormalization** Target voltage values are standardized using z-
 194 score normalization inspired by (Apicella et al., 2023):

$$195 v_{\text{norm}} = \frac{v - \mu}{\sigma}, \quad (4)$$

196 where μ and σ denote the mean and standard deviation of training voltages. The predicted voltage
 197 is denormalized during inference by $\hat{v} = \hat{v}_{\text{norm}} \cdot \sigma + \mu$, ensuring that outputs remain consistent with
 198 real world EEG magnitudes.

199 **Loss Function** To train the model, we employ the Huber loss function Huber (1992), which pro-
 200 vides a robust alternative to the Mean Squared Error (MSE). The Huber loss behaves quadratically
 201 near the origin and transitions to linear growth for large residuals:

$$202 \mathcal{L}_{\delta}(x, y) = \begin{cases} \frac{1}{2}(x - y)^2 & \text{if } |x - y| < \delta \\ \delta \cdot (|x - y| - \frac{1}{2}\delta) & \text{otherwise,} \end{cases} \quad (5)$$

203 where δ is a tunable threshold. This formulation is particularly suitable for EEG signals, which
 204 may contain occasional transient spikes or artifacts. Unlike the MSE loss, which is highly sensitive
 205 to outliers, the Huber loss moderates their influence. This robustness allows the model to focus
 206 on learning the underlying neural patterns without being disproportionately affected by noise or
 207 corrupted samples that are common in EEG data.

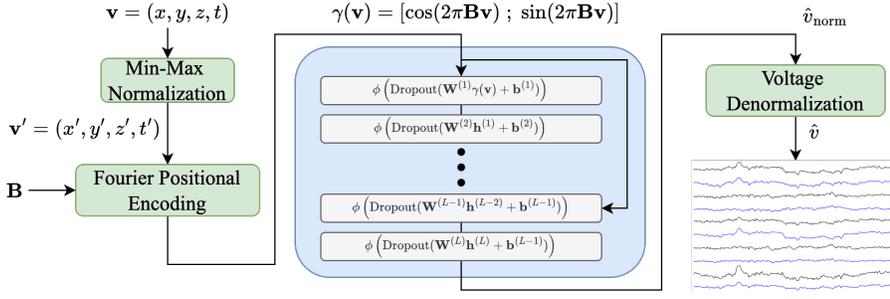


Figure 1: Diagram of the proposed NBF architecture.

3.2 ARCHITECTURE

We propose an architecture tailored to capture the continuous spatiotemporal dynamics of individual EEG recordings, with its main components outlined below.

Fourier Positional Encoding (PE) To overcome the spectral bias of standard MLPs, which are biased toward learning low frequency functions Rahaman et al. (2019), we employ Fourier-based PE Tancik et al. (2020). This approach maps each normalized input $\mathbf{v} \in \mathbb{R}^4$ to a high-dimensional embedding using sinusoidal functions across multiple frequencies:

$$\gamma(\mathbf{v}) = [\cos(2\pi\mathbf{B}\mathbf{v}) ; \sin(2\pi\mathbf{B}\mathbf{v})] \in \mathbb{R}^{2m}, \quad (6)$$

where $\mathbf{B} \in \mathbb{R}^{m \times 4}$ is a matrix with entries sampled i.i.d from a Gaussian distribution, $B_{ij} \sim \mathcal{N}(0, \sigma^2)$. The hyperparameter σ controls the frequency scale and affects the model’s ability to learn high frequency signal components.

Multi Layer Perceptron (MLP) The encoded vector $\gamma(\mathbf{v})$ is passed through a fully connected MLP composed of L layers of width W , each using ReLU activation, dropout, and periodic skip connections. The network is defined as:

$$\mathbf{h}^{(0)} = \gamma(\mathbf{v}), \quad \text{for } l = 1, \dots, L - 1: \quad \mathbf{h}^{(l)} = \phi \left(\text{Dropout}(\mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}) \right), \quad (7)$$

with the output computed as:

$$\hat{v}_{\text{norm}} = \mathbf{W}^{(L)}\mathbf{h}^{(L-1)} + \mathbf{b}^{(L)}. \quad (8)$$

For skip-connected layers, the input $\gamma(\mathbf{v})$ is concatenated with intermediate features:

$$\mathbf{h}^{(l)} = \phi \left(\text{Dropout}(\mathbf{W}^{(l)}[\mathbf{h}^{(l-1)}; \gamma(\mathbf{v})] + \mathbf{b}^{(l)}) \right). \quad (9)$$

Model Initialization To accelerate the overall training process, we utilize a progressive fine tuning approach. The model is first trained from scratch on the initial window for a full number of epochs and then saved. For each subsequent window, we initialize training from the previously saved model and fine tune it for a reduced number of epochs, subsequently overwriting the saved checkpoint. This iterative strategy reduces the total training time by approximately 30%, while simultaneously improving reconstruction accuracy. The improvement is attributed to the temporal coherence that exists across neighboring windows, which the model exploits when initialized with weights from a temporally adjacent segment.

3.3 RENDERING APPLICATIONS

Inference Once trained, the model is capable of predicting voltage values at arbitrary spatiotemporal coordinates (x, y, z, t) , including positions that were not part of the original training dataset. These target locations can be selected based on task specific requirements. For instance, in speech decoding tasks, one may focus on regions proximal to the auditory cortex, such as those corresponding to electrodes T7 and T8. In cognitive neuroscience experiments that involve executive function or attention, predictions may be concentrated in frontal regions represented by electrodes such as Fz. The model’s continuous formulation allows for smooth spatial and temporal interpolation, enabling flexible voltage estimation across a range of configurations.

Data Driven Electrode Augmentation Strategy To enhance EEG spatial resolution beyond electrode placement limits, we use a data driven augmentation strategy that synthesizes virtual signals at additional scalp sites. This augments the sensor grid with plausible voltages, improving spatial coverage and enabling comprehensive downstream analyses. The model is trained on temporal windows of length T , where the number of windows equals the recording duration divided by T . In each window, training uses all available electrodes, after which the model synthesizes voltages at new scalp locations not coinciding with the originals. Repeating this across windows produces high resolution synthetic data that captures both spatial proximity and localized temporal dynamics.

4 EXPERIMENTS

Decoding Speech from EEG Défossez et al. (2023) developed a model that decodes speech perception from noninvasive brain recordings (MEG and EEG) using contrastive learning, employing a dedicated brain module and a pretrained speech module based on wav2vec 2.0. Top 1 accuracy quantifies the proportion of instances in which the model’s highest ranked prediction matches the actual speech segment, while top 10 accuracy reflects the frequency with which the correct segment appears among the ten most probable predictions.

Experimental Protocol and Training Strategy To capture subject specific neural dynamics efficiently, we employed progressive fine tuning of one Neural Brain Field (NBF) model per subject using temporally segmented EEG/MEG data. Recordings were split into non overlapping 3s windows. For each subject, training began from scratch on the first window with all electrodes. After training on the first window, the model then predicted voltages at novel auditory cortex electrodes for that window, generating synthetic 3s signals integrated into the data stream. Model parameters were checked, reloaded for each subsequent window, and fine tuned for fewer epochs. After each step, predictions at the same novel electrodes sites were again generated and inserted per time window. This iterative scheme progressively adapted weights while preserving subject specific features, capturing localized spatiotemporal dynamics and exploiting temporal coherence. Avoiding reinitialization cut training time by $\sim 30\%$, while per subject training maintained physiological fidelity, yielding realistic and stable scalp wide electrode synthesis.

Evaluating Functional Benefits via Emotion Recognition To further assess the practical utility of our NBF model in real world EEG applications, we conducted an experiment on emotion recognition. The goal was to determine whether augmenting the spatial coverage of EEG recordings with virtual electrodes synthesized by our model could improve performance in a representative downstream task. We selected the DREAMER dataset (Katsigiannis & Ramzan, 2018), a widely used benchmark for emotion recognition based on EEG, and adopted a standard classification framework using the EEGNet architecture (Lawhern et al., 2018a). Virtual electrodes were synthesized at central scalp locations commonly associated with affective processing, and the augmented data were used to train the classifier following established protocols (Kukhilava et al., 2025). This experiment allows us to evaluate whether the spatially enriched representations generated by our model offer measurable benefits in functional decoding tasks beyond interpolation accuracy alone.

Hyperparameters We optimized the NBF model via an extensive grid search over dropout rate, batch size, depth, width, PE levels, and learning rate, using a 90/10 train-validation split. Validation sets were selected for maximal cortical coverage to ensure robust generalization. The search was conducted independently on all three datasets using subsets of participants, and the best configuration was retained. Final settings for most experiments were dropout 0.1, width 1450, depth 8, and batch size 32, trained with Adam and gradient clipping for stability. For the Gwilliams dataset, a larger batch size of 250 was adopted to match dataset scale; this choice was validated on larger participant subsets, confirming robustness across datasets and configurations.

5 RESULTS

5.1 ENHANCED DECODING PERFORMANCE VIA ELECTRODE AUGMENTATION STRATEGY

As shown in Table 1, the additional electrode data generated by our NBF model significantly enhanced the performance of the Brainmagick (Défossez et al., 2023) and outperformed Transformer-VAE (Chen et al., 2025) across all datasets and metrics in the task of decoding speech perception from noninvasive brain recordings. On the Brennan dataset (Brennan & Hale, 2019), our model achieves a Top-1 accuracy of 9.64 and a Top-10 accuracy of 36.94, compared to 5.2 and 25.7 for

Table 1: Top-1 and Top-10 accuracy comparison of three methods (Original, VAE, and Ours) across three evaluation benchmarks: Brennan, Gwilliams, and Borderick. Results show that our method consistently improves performance across multiple datasets and evaluation metrics.

Method / Dataset	Brennan		Gwilliams		Borderick	
	Top-1 (\uparrow)	Top-10 (\uparrow)	Top-1 (\uparrow)	Top-10 (\uparrow)	Top-1 (\uparrow)	Top-10 (\uparrow)
Brainmagick (Défossez et al., 2023)	5.2 \pm 0.8	25.7 \pm 2.9	41.3 \pm 0.1	70.7 \pm 0.1	5.0 \pm 0.4	17.0 \pm 0.6
Transformer-VAE (Chen et al., 2025)	4.1	26.82	–	–	–	–
NBF (Our)	9.64 \pm 0.33	36.94 \pm 0.1	42.1 \pm 0.1	71.33 \pm 0.1	7.44 \pm 0.13	22.68 \pm 0.08

Brainmagick. On the Gwilliams dataset (Gwilliams et al., 2022), our model reaches 42.1 and 71.33, slightly surpassing Brainmagick’s 41.3 and 70.7. On the Broderick dataset, we observe similar improvements. These gains are attributed to our method of augmenting the input space with synthetic electrodes near auditory regions. By enriching the spatial representation, the model gains better access to neural dynamics relevant for speech perception, enabling more accurate decoding.

5.2 IMPACT OF ELECTRODE AUGMENTATION ON EMOTION RECOGNITION PERFORMANCE

To assess the utility of our NBF model in augmenting spatial EEG information for emotion recognition, we conducted experiments on the DREAMER dataset (Katsigiannis & Ramzan, 2018). We employed our model to synthesize EEG signals at five central electrode sites: Cz, Fz, Pz, C3, and C4. These locations were chosen based on prior studies that emphasize their importance in affective processing and emotion classification (Singh & Singh, 2017; Koelstra et al., 2012). By introducing virtual signals at these strategically selected sites, our model effectively extends spatial coverage in brain regions known to carry emotion discriminative features.

To quantify the contribution of the augmented signals to downstream tasks, we trained the EEGNet architecture (Lawhern et al., 2018a) on the DREAMER dataset (Katsigiannis & Ramzan, 2018), supplemented with our synthesized signals from five additional central electrodes (Cz, Fz, Pz, C3, and C4). The training followed the protocol outlined by Kukhilava et al. (Kukhilava et al., 2025), which aligns with contemporary benchmarking standards in EEG based emotion recognition. As shown in Table 2, the inclusion of virtual electrode signals led to consistent improvements in classification accuracy for both arousal and valence dimensions. These findings highlight the effectiveness of the NBF model in enhancing EEG spatial resolution and suggest its potential to improve performance across diverse downstream EEG decoding tasks, including but not limited to affective computing.

5.3 EVALUATION OF EEG SIGNAL PREDICTION AT UNSEEN ELECTRODE POSITIONS

To assess the model’s ability to predict EEG at unseen electrodes, it was trained on 55 of 60 electrodes and tested on the 5 held out positions, repeating three times with different disjoint validation sets across all Brennan dataset subjects. Evaluation used one third of each subject’s recordings to reduce computational cost. Performance was compared with Spline Surface Interpolation (SSI) Perrin et al. (1989) and Radial Basis Function (RBF) interpolation (Jäger et al., 2016), using identical train test splits and preprocessing. Subjects with extreme baseline errors were excluded; the remaining 30 satisfied SSI’s spherical assumptions via MNE-Python’s `_check_pos_sphere` (Gramfort et al., 2013). Negative R^2 values were clamped to zero, and extreme SNR outliers discarded. Evaluation was conducted over 80 non overlapping 3 s windows per subject, with averaged results. Table 2 shows the NBF model consistently outperformed baselines, demonstrating robust reconstruction and interpolation at spatially unseen electrodes.

Figure 2: Performance comparison of EEGNet on the DREAMER dataset with and without the addition of virtual electrodes generated by our NBF model. Values are reported as mean \pm standard deviation across subjects.

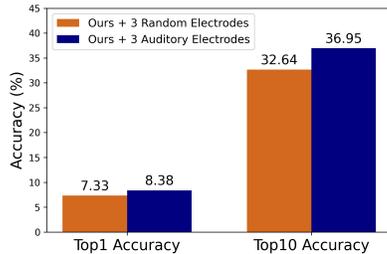
Model	Valence Accuracy	Arousal Accuracy
EEGNet	0.60 \pm 0.08	0.46 \pm 0.09
EEGNet W. Ours	0.63 \pm 0.03	0.49 \pm 0.01

5.4 MODEL EVALUATION ON AUDITORY-CENTRIC ELECTRODE AUGMENTATION

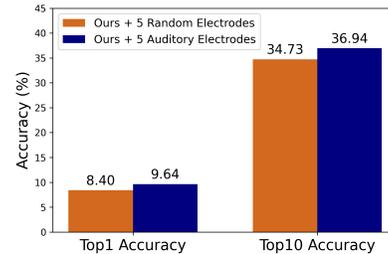
To evaluate the effectiveness of our model in generating meaningful synthetic EEG electrode data, we conducted a comparative analysis using the Brainmagick model on the Brennan dataset. We examined the impact of augmenting the original electrode configuration with either three or five additional electrodes. These electrodes were either selected at random or strategically placed near the auditory cortex. As illustrated in the graphs below (Figures 3a and 3b), the inclusion of audi-

Table 2: Performance comparison between our model and baselines across various metrics calculated using the Scikit-learn library (Pedregosa et al., 2011).

	MSE (\downarrow)	MAE (\downarrow)	R^2 (\uparrow)	PCC (\uparrow)	SNR (\uparrow)	NMSE (\downarrow)
RBF (Jäger et al., 2016)	2.80e−8	1.78e−5	0.456	0.746	3.93	0.544
SSI (Perrin et al., 1989)	4.19e−8	2.06e−5	0.399	0.720	3.33	0.601
NBF (Ours)	1.32e−8	1.07e−5	0.483	0.755	4.29	0.517



(a) Accuracy with 3 additional electrodes



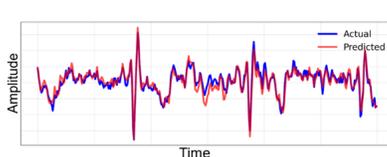
(b) Accuracy with 5 additional electrodes

Figure 3: Reconstruction accuracy improvement with synthetic auditory electrodes.

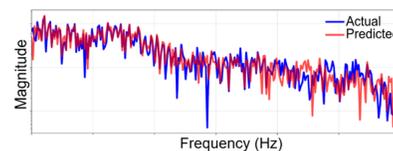
tory region electrodes consistently outperformed the addition of randomly placed electrodes in both augmentation settings (three or five electrodes). These findings are consistent with the nature of the task, decoding auditory stimuli, as the auditory cortex is directly involved in speech perception. This suggests that our model has the potential to support task specific electrode augmentation.

6 MODEL ANALYSIS

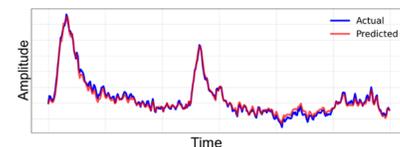
Reconstructed EEG Signals To illustrate the model’s ability to reconstruct neural activity at spatially unseen sites, we trained it from scratch on a single 3 s window from one subject, using a 90%/10% spatial split of electrodes (training/validation). Training and prediction were restricted to this window and subject, with no leakage from held out electrodes. Figure 4 show representative results: Signal 1 is an MEG channel from the Gwilliams dataset (Gwilliams et al., 2022), and Electrode 2 an EEG channel from the Broderick dataset (Broderick et al., 2018). Time domain plots (4a, 4c) reveal strong alignment between predictions and ground truth, capturing waveform morphology and neural dynamics, while frequency domain plots (4b, 4d) demonstrate spectral coherence across low and high frequencies. These results highlight the model’s capacity to synthesize meaningful neural signals at previously unseen spatial locations from limited temporal and spatial data.



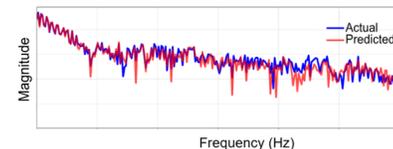
(a) Electrode 1 – Time Domain



(b) Electrode 1 – Frequency Domain



(c) Electrode 2 – Time Domain



(d) Electrode 2 – Frequency Domain

Figure 4: Electrode signals are presented in each row: in (a,c), the time domain; and in (b,d), the frequency domain.

Qualitative Assessment of Synthesized Electrode Signals Figure 5 shows synthesized electrode signals for a representative Brennan subject. Each original electrode (black) is paired with its nearest synthesized auditory-related electrode (blue) based on Euclidean distance in the standard montage:

Table 3: Ablation study results for different normalization, encoding, and architectural configurations. Metrics were calculated using Pedregosa et al. (2011).

Ablation	MSE (\downarrow)	MAE (\downarrow)	R2 (\uparrow)	PCC (\uparrow)	SNR (\uparrow)	NMSE (\downarrow)
Ours	8.18e-11	5.20e-06	0.680	0.8330	5.40	0.32
Ours – PE (16 levels)	8.26e-11	5.39e-06	0.670	0.8311	5.31	0.33
Ours – PE (8 levels)	8.45e-11	5.50e-06	0.661	0.8300	5.32	0.34
Ours – w/o Positional Encoding	1.41e-10	7.88e-06	0.380	0.6360	2.64	0.62
Ours – MSE loss	8.32e-11	5.59e-06	0.657	0.8290	5.35	0.34
Ours – w/o Dropout	8.38e-11	5.54e-06	0.657	0.8290	5.30	0.34
Ours – w/o Model Initialization	8.31e-11	5.64e-06	0.653	0.6526	5.20	0.35
Ours – w/o Coord. Normalization	9.01e-11	5.65e-06	0.630	0.8140	5.03	0.37
Ours – w/o Z-Score Normalization	2.62e-10	1.09e-05	-0.049	0.0300	0.04	1.05
Ours – w/o Skip Connection	8.32e-11	5.54e-06	0.658	0.8290	5.32	0.34

59–62, 53–63, 58–64, 54–65, and 60–66. This side by side comparison enables evaluation of temporal structure and morphology. The synthetic signals preserve key waveform patterns and dynamics of the originals, demonstrating physiologically plausible data at novel sites and supporting the model’s spatial interpolation validity.

Ablation Study As shown in Table 3, we conducted an ablation study on a high quality subset of subjects from the Brennan dataset, using 55 electrodes for training and 5 unseen electrodes for validation. Replacing the Huber loss with MSE slightly reduced performance, confirming Huber’s robustness to transient artifacts. Disabling progressive initialization degraded accuracy and increased training time, whereas initializing from adjacent windows reduced training time by $\sim 30\%$ by leveraging temporal continuity. Removing Z score or coordinate normalization impaired performance, with Z score removal causing catastrophic failure. Similarly, reducing PE depth (e.g., 8 or 16 levels) impaired generalization, while omitting PEs or skip connections substantially reduced accuracy. Overall, each architectural and normalization component proved critical for efficient and accurate EEG reconstruction.

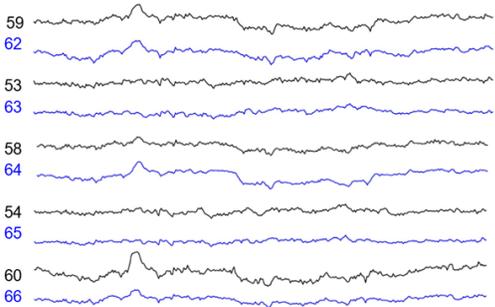


Figure 5: Qualitative comparison of original (black) and synthesized (blue) EEG signals for spatially nearest electrode pairs (59–62, 53–63, 58–64, 54–65, and 60–66) in the Brennan dataset, illustrating preserved waveform morphology and temporal dynamics at novel electrode locations.

7 CONCLUSION

We introduce Neural Brain Fields (NBF), a NeRF-inspired framework for continuous spatiotemporal EEG representations. Trained per subject and recording, NBF generates signals at arbitrary scalp locations, enabling spatial interpolation, synthetic electrode placement, and adaptation to individual head variability. NBF outperforms SSI and RBF interpolation, improves speech decoding in Brainmagick (Défossez et al., 2023) with auditory cortex electrodes, and enhances EEGNet (Lawhern et al., 2018a) emotion classification on DREAMER (Katsigiannis & Ramzan, 2018) via central virtual electrodes, demonstrating generalizable spatial resolution gains. Formally, NBF predicts voltages v from spatiotemporal inputs (x, y, z, t) , establishing a flexible framework for other voltage-mapped domains (e.g., geophysics, environmental monitoring) requiring dense spatial interpolation. Current computational cost limits real-time use; however, runtime acceleration is feasible via NeRF advances: multiresolution hash encoding and CUDA optimization (Instant-NGP (Müller et al., 2022)), architectural simplification (tinyNeRF (Zhao et al., 2023)), parallel MLP inference (KiloNeRF (Reiser et al., 2021)), modular computation separation (FastNeRF (Garbin et al., 2021)), and multiscale sampling (Mip-NeRF (Barron et al., 2021)). Adapting these may enable both offline and low latency deployment. Future work will pursue such optimizations and extend NBF to other neuroimaging modalities and downstream applications, translating continuous neural fields into practical neuroscience tools.

8 ETHICS STATEMENT

This research utilizing publicly available and requested access EEG datasets commits to maintaining scientific integrity by clearly distinguishing between original and synthetically generated electrode data in all publications, ensuring proper attribution and data licensing compliance, validating synthetic signal quality through rigorous metrics, and acknowledging limitations of the NBF approach. We recognize that synthetic neural signals, despite being derived from legitimate datasets, require transparent reporting to prevent misrepresentation in downstream applications, particularly in clinical or brain computer interface contexts where safety and efficacy depend on authentic neural recordings. The research prioritizes responsible sharing of findings that prevent misuse of synthetic EEG generation technology and adherence to existing institutional ethics frameworks while establishing appropriate boundaries between research applications and potential clinical deployment.

9 REPRODUCIBILITY STATEMENT

The reproducibility package provides all necessary instructions and scripts to replicate the experiments reported in our paper. It includes environment setup files, preprocessing pipelines for the Brennan2019 dataset, and end to end scripts for training and evaluating the proposed NBF model. Step by step usage is documented, covering data preparation, model execution, electrode augmentation, and integration with downstream baselines such as Brainmagick. By following the provided guidelines, researchers can reproduce the reported results and further explore the effectiveness of the proposed approach.

REFERENCES

- Andrea Apicella, Francesco Isgrò, Andrea Pollastro, and Roberto Prevete. On the effects of data normalization for domain adaptation on eeg data. *Engineering Applications of Artificial Intelligence*, 123:106205, 2023. ISSN 0952-1976. doi: <https://doi.org/10.1016/j.engappai.2023.106205>. URL <https://www.sciencedirect.com/science/article/pii/S0952197623003895>.
- Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5855–5864, 2021.
- Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5470–5479, 2022.
- Jonathan R. Brennan and John T. Hale. Hierarchical structure guides rapid linguistic predictions during naturalistic listening. *PLoS ONE*, 14(1):e0207741, 2019. doi: 10.1371/journal.pone.0207741. URL <https://doi.org/10.1371/journal.pone.0207741>.
- Michael P. Broderick, Adam J. Anderson, Giovanni M. Di Liberto, Michael J. Crosse, and Edmund C. Lalor. Electrophysiological correlates of semantic dissimilarity reflect the comprehension of natural, narrative speech. *Current Biology*, 28(5):803–809.e3, 2018. doi: 10.1016/j.cub.2018.01.080.
- Terrance Yu-Hao Chen, Yulin Chen, Pontus Soederhaell, Sadrishya Agrawal, and Kateryna Shapovalenko. Decoding eeg speech perception with transformers and vae-based data augmentation, 2025. URL <https://arxiv.org/abs/2501.04359>.
- Isaac A Corley and Yufei Huang. Deep eeg super-resolution: Upsampling eeg spatial resolution with generative adversarial networks. In *2018 IEEE EMBS international conference on biomedical & health informatics (BHI)*, pp. 100–103. IEEE, 2018.
- Alexandre Défossez, Charlotte Caucheteux, Jérémy Rapin, Ori Kabeli, and Jean-Rémi King. Decoding speech perception from non-invasive brain recordings. *Nature Machine Intelligence*, 5(10):1097–1107, 2023.

- 540 Andac Demir, Toshiaki Koike-Akino, Ye Wang, Masaki Haruna, and Deniz Erdogmus. Eeg-gnn:
541 Graph neural networks for classification of electroencephalogram (eeg) signals. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*,
542 pp. 1061–1067. IEEE, 2021.
- 543
544 Berkay Döner, Thorir Mar Ingolfsson, Luca Benini, and Yawei Li. Luna: Efficient and topology-
545 agnostic foundation model for eeg signal analysis. In *1st ICML Workshop on Foundation Models*
546 *for Structured Data*.
- 547
548 Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fast-
549 nerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF international*
550 *conference on computer vision*, pp. 14346–14355, 2021.
- 551
552 Yuxiang Ge, Jionghao Cheng, Ruiquan Ge, Zhaojie Fang, Gangyong Jia, Xiang Wan, Nannan Li,
553 Ahmed Elazab, and Changmiao Wang. 3d-telepathy: Reconstructing 3d objects from eeg signals.
554 *arXiv preprint arXiv:2506.21843*, 2025.
- 555
556 Alexandre Gramfort, Martin Luessi, Eric Larson, Denis A. Engemann, Daniel Strohmeier, Christian
557 Brodbeck, Roman Goj, Mainak Jas, Teon Brooks, Lauri Parkkonen, and Matti S. Hämäläinen.
558 MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7(267):1–13, 2013.
doi: 10.3389/fnins.2013.00267.
- 559
560 Yiyu Gui, Mingzhi Chen, Yuqi Su, Guibo Luo, and Yuchao Yang. Eegmamba: Bidirectional
561 state space model with mixture of experts for eeg multi-task classification. *arXiv preprint*
562 *arXiv:2407.20254*, 2024.
- 563
564 Laura Gwilliams, Jean-Rémi King, Alec Marantz, and David Poeppel. Neural dynamics of phoneme
565 sequences reveal position-invariant code for content and order. *Nature Communications*, 13
566 (1):6606, 2022. doi: 10.1038/s41467-022-34326-1. URL <https://doi.org/10.1038/s41467-022-34326-1>.
- 567
568 Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodol-*
569 *ogy and distribution*, pp. 492–518. Springer, 1992.
- 570
571 Janin Jäger, Alexander Klein, Martin Buhmann, and Wolfgang Skrandies. Reconstruction of elec-
572 troencephalographic data using radial basis functions. *Clinical Neurophysiology*, 127(4):1978–
573 1983, 2016.
- 574
575 Hao Jia, Zihao Huang, Cesar F Caiafa, Feng Duan, Yu Zhang, Zhe Sun, and Jordi Solé-Casals.
576 Assessing the potential of data augmentation in eeg functional connectivity for early detection of
577 alzheimer’s disease. *Cognitive Computation*, 16(1):229–242, 2024.
- 578
579 Stamos Katsigiannis and Naeem Ramzan. Dreamer: A database for emotion recognition through
580 eeg and ecg signals from wireless low-cost off-the-shelf devices. *IEEE Journal of Biomedical*
581 *and Health Informatics*, 22(1):98–107, 2018. doi: 10.1109/JBHI.2017.2688239.
- 582
583 Dominik Klepl, Min Wu, and Fei He. Graph neural network-based eeg classification: A survey.
584 *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 32:493–503, 2024.
- 585
586 Sander Koelstra, Christian Muhl, Mohammad Soleymani, Jong-Seok Lee, Ashkan Yazdani, Touradj
587 Ebrahimi, Thierry Pun, Anton Nijholt, and Ioannis Patras. Deap: A database for emotion analysis
588 ;using physiological signals. *IEEE Transactions on Affective Computing*, 3(1):18–31, 2012. doi:
589 10.1109/T-AFFC.2011.15.
- 590
591 Natia Kukhilava, Tatia Tsmindashvili, Rapael Kalandadze, Anchit Gupta, Sofio Katamadze,
592 François Brémond, Laura M. Ferrari, Philipp Müller, and Benedikt Emanuel Wirth. Evaluation
593 in eeg emotion recognition: State-of-the-art review and unified framework, 2025. URL
<https://arxiv.org/abs/2505.18175>.
- 594
595 Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and
596 Brent J Lance. Eegnet: a compact convolutional neural network for eeg-based brain–computer
597 interfaces. *Journal of Neural Engineering*, 15(5):056013, July 2018a. ISSN 1741-2552. doi: 10.
598 1088/1741-2552/aace8c. URL <http://dx.doi.org/10.1088/1741-2552/aace8c>.

- 594 Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and
595 Brent J Lance. Eegnet: a compact convolutional neural network for eeg-based brain–computer
596 interfaces. *Journal of neural engineering*, 15(5):056013, 2018b.
597
- 598 Eleonora Lopez, Luigi Sigillo, Federica Colonnese, Massimo Panella, and Danilo Comminiello.
599 Guess what i think: Streamlined eeg-to-image generation with latent diffusion models. In
600 *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing*
601 (*ICASSP*), pp. 1–5. IEEE, 2025.
- 602 Christoph M Michel, Micah M Murray, Göran Lantz, Sara Gonzalez, Laurent Spinelli, and
603 Rolando Grave De Peralta. Eeg source imaging. *Clinical neurophysiology*, 115(10):2195–2222,
604 2004.
605
- 606 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
607 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*
608 *of the ACM*, 65(1):99–106, 2021.
609
- 610 Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics prim-
611 itives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15,
612 2022.
- 613 Ernst Niedermeyer and FH Lopes da Silva. *Electroencephalography: basic principles, clinical*
614 *applications, and related fields*. Lippincott Williams & Wilkins, 2005.
615
- 616 Paul L Nunez and Ramesh Srinivasan. *Electric fields of the brain: the neurophysics of EEG*. Oxford
617 university press, 2006.
618
- 619 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Pretten-
620 hofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and
621 E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*,
622 12:2825–2830, 2011.
- 623 François Perrin, Jacques Pernier, Olivier Bertrand, and Jean Francois Echallier. Spherical splines
624 for scalp potential and current density mapping. *Electroencephalography and clinical neurophysi-*
625 *ology*, 72(2):184–187, 1989.
626
- 627 Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred A. Hamprecht,
628 Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks, 2019. URL
629 <https://arxiv.org/abs/1806.08734>.
- 630 Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radi-
631 ance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF international conference*
632 *on computer vision*, pp. 14335–14345, 2021.
633
- 634 Cédric Rommel, Joseph Paillard, Thomas Moreau, and Alexandre Gramfort. Data augmentation for
635 learning predictive models on eeg: a systematic comparison. *Journal of Neural Engineering*, 19
636 (6):066020, 2022.
637
- 638 Subhrajit Roy, Isabell Kiral-Kornek, and Stefan Harrer. Chrononet: A deep recurrent neural net-
639 work for abnormal eeg identification. In *Artificial Intelligence in Medicine: 17th Conference on*
640 *Artificial Intelligence in Medicine, AIME 2019, Poznan, Poland, June 26–29, 2019, Proceedings*
641 *17*, pp. 47–56. Springer, 2019.
- 642 Giulio Ruffini, David Ibanez, Marta Castellano, Stephen Dunne, and Aureli Soria-Frisch. Eeg-driven
643 rnn classification for prognosis of neurodegeneration in at-risk patients. In *Artificial Neural Net-*
644 *works and Machine Learning–ICANN 2016: 25th International Conference on Artificial Neural*
645 *Networks, Barcelona, Spain, September 6-9, 2016, Proceedings, Part I 25*, pp. 306–313. Springer,
646 2016.
647
- Saeid Sanei and Jonathon A Chambers. *EEG signal processing*. John Wiley & Sons, 2013.

- 648 Robin Tibor Schirrmeister, Jost Tobias Springenberg, Lukas Dominique Josef Fiederer, Martin
649 Glasstetter, Katharina Eggenberger, Michael Tangermann, Frank Hutter, Wolfram Burgard, and
650 Tonio Ball. Deep learning with convolutional neural networks for eeg decoding and visualization.
651 *Human brain mapping*, 38(11):5391–5420, 2017.
- 652
653 Moon Inder Singh and Mandeep Singh. Development of a real time emotion classifier based on
654 evoked eeg. *Biocybernetics and Biomedical Engineering*, 37(3):498–509, 2017. ISSN 0208-5216.
655 doi: <https://doi.org/10.1016/j.bbe.2017.05.004>. URL <https://www.sciencedirect.com/science/article/pii/S0208521616303035>.
- 656
657 Prajwal Singh, Pankaj Pandey, Krishna Miyapuram, and Shanmuganathan Raman. Eeg2image: im-
658 age reconstruction from eeg brain signals. In *ICASSP 2023-2023 IEEE International Conference*
659 *on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2023.
- 660
661 Prajwal Singh, Dwip Dalal, Gautam Vashishtha, Krishna Miyapuram, and Shanmuganathan Raman.
662 Learning robust deep visual representations from eeg brain recordings. In *Proceedings of the*
663 *IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 7553–7562, 2024.
- 664
665 Mats Svantesson, Håkan Olausson, Anders Eklund, and Magnus Thordstein. Virtual eeg-electrodes:
666 Convolutional neural networks as a method for upsampling or restoring channels. *Journal of*
667 *Neuroscience Methods*, 355:109126, 2021a.
- 668
669 Mats Svantesson, Håkan Olausson, Anders Eklund, and Magnus Thordstein. Virtual eeg-electrodes:
670 Convolutional neural networks as a method for upsampling or restoring channels. *Journal*
671 *of Neuroscience Methods*, 355:109126, 2021b. ISSN 0165-0270. doi: <https://doi.org/10.1016/j.jneumeth.2021.109126>. URL <https://www.sciencedirect.com/science/article/pii/S0165027021000613>.
- 672
673 Ke Tan, Buye Xu, Anurag Kumar, Eliya Nachmani, and Yossi Adi. Sagramn: Self-attentive gated rnn
674 for binaural speaker separation with interaural cue preservation. *IEEE Signal Processing Letters*,
675 28:26–30, 2020.
- 676
677 Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh
678 Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn
679 high frequency functions in low dimensional domains. *Advances in neural information processing*
680 *systems*, 33:7537–7547, 2020.
- 681
682 Szabolcs Torma and Luca Szegletes. Generative modeling and augmentation of eeg signals using
683 improved diffusion probabilistic models. *Journal of Neural Engineering*, 22(1):016001, 2025.
- 684
685 Szabolcs Torma and Luca Szegletes. Eegwave: a denoising diffusion probabilistic approach for eeg
686 signal generation. EasyChair Preprint 10275, EasyChair, 2023.
- 687
688 Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction
689 of large-scale nerfs for virtual fly-throughs. In *Proceedings of the IEEE/CVF conference on*
computer vision and pattern recognition, pp. 12922–12931, 2022.
- 690
691 Zhijiang Wan, Manyu Li, Shichang Liu, Jiajin Huang, Hai Tan, and Wenfeng Duan. Eegformer: A
692 transformer-based brain activity classification method using eeg signal. *Frontiers in neuroscience*,
693 17:1148855, 2023.
- 694
695 Guangyu Wang, Wenchao Liu, Yuhong He, Cong Xu, Lin Ma, and Haifeng Li. Eegpt: Pretrained
696 transformer for universal and reliable representation of eeg signals. *Advances in Neural Informa-*
tion Processing Systems, 37:39249–39280, 2024.
- 697
698 Xin Xiang, Wenhui Zhou, and Guojun Dai. Eeg-driven 3d object reconstruction with style consis-
699 tency and diffusion prior. *arXiv preprint arXiv:2410.20981*, 2024.
- 700
701 Xiaoxiao Yang and Ziyu Jia. Spatial-temporal mamba network for eeg-based motor imagery classi-
fication. In *International Conference on Advanced Data Mining and Applications*, pp. 418–432.
Springer, 2024.

702 Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing:
703 geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Com-*
704 *puter Vision and Pattern Recognition*, pp. 18353–18364, 2022.

706 Seokho Yun. Advances, challenges, and prospects of electroencephalography-based biomarkers for
707 psychiatric disorders: a narrative review. *Journal of Yeungnam Medical Science*, 41(4):261–268,
708 2024.

709 Tianli Zhao, Jiayuan Chen, Cong Leng, and Jian Cheng. Tinynerf: Towards 100 x compression
710 of voxel radiance fields. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol-
711 ume 37, pp. 3588–3596, 2023.

714 A APPENDIX

715
716 You may include other additional sections here.

718 **Datasets & Preprocessing** We utilized three publicly available neural datasets. Broderick (Bro-
719 derick et al., 2018) includes EEG recordings from 19 subjects using 128 sensors (19.2 hours total),
720 collected during continuous narrative speech listening. Brennan (Brennan & Hale, 2019) consists
721 of EEG data from 33 participants (60 sensors, 6.7 hours), recorded during passive listening to a
722 12.4 minute audiobook, aimed at analyzing hierarchical linguistic prediction. Gwilliams (Gwilliams
723 et al., 2022) provides MEG recordings from 27 subjects (208 sensors, 56.2 hours), focused on
724 phoneme sequence encoding during natural speech perception. All datasets were downsampled
725 to 120 Hz and preprocessed with a 0.5 to 59.5 Hz bandpass filter.

726 The DREAMER dataset was preprocessed with a bandpass filter in the range 0.3–45 Hz to retain
727 relevant neural frequencies, and a 50 Hz notch filter was applied to remove power line interference.

729 A.1 COMPLEXITY ANALYSIS

730
731 To assess the computational demands of our method, we measured the total processing time for each
732 dataset. The pipeline involved training on sequential non overlapping 3 second windows, generating
733 synthetic electrode data per window, and iterating through the dataset. Processing the Brennan
734 dataset required approximately 62 hours, Broderick around 172 hours, and Gwilliams the largest
735 and most complex approximately 525 hours. These durations reflect the cumulative runtime across
736 all 3 second windows. All experiments were performed on a parallelized setup with 8 NVIDIA RTX
737 6000 GPUs, allowing the effective per GPU runtime to be estimated by dividing the total time by
738 8. This parallelization substantially reduced wall clock time and demonstrates the scalability of our
739 approach on multi GPU systems.

740 **Biophysical Background** The feasibility of super resolving EEG signals is grounded in well es-
741 tablished biophysical properties of volume conduction, which induce spatial smoothness and strong
742 correlations across neighboring electrode sites (Nunez & Srinivasan, 2006; Michel et al., 2004).
743 This inductive prior is precisely what enables super resolution. Theoretical work has long estab-
744 lished that scalp EEG is dominated by smooth, spatially diffuse fields resulting from the physics of
745 current propagation through the brain and skull (Nunez & Srinivasan, 2006). Our results empiri-
746 cally validate that these correlations are present and can be effectively leveraged using a spatially
747 continuous representation learned per subject and per recording. Furthermore, with appropriate fine
748 tuning, our model can adapt to new montages and patient specific geometries, offering a flexible and
749 extensible framework for enhancing spatial resolution beyond traditional interpolation methods or
750 fixed supervised learning paradigms.

751 **Brennan and Gwilliams Datasets (Brennan & Hale, 2019; Gwilliams et al., 2022)** Figures 6
752 and 7 show Top-1 and Top-10 accuracy as functions of added electrodes. On the Brennan dataset,
753 accuracy rises with more electrodes, peaking at 5 before declining, indicating that NBF effectively
754 exploits limited spatial information and generalizes well by interpolating neural activity, validating
755 its utility under sensor constraints. On the Gwilliams dataset, gains are smaller but consistent up
to 11 electrodes; despite high baseline performance, NBF improves Top-1 by $\sim 0.8\%$ and Top-10



Figure 6: Top-1 and Top-10 accuracy on Brennan dataset as a function of added electrodes.



Figure 7: Top-1 and Top-10 accuracy on the Gwilliams dataset as a function of added electrodes.

by $\sim 0.6\%$, i.e., $8\times$ and $6\times$ the reported 0.1% STD (Gwilliams et al., 2022), strongly surpassing baseline variability and confirming substantial improvement over prior methods.