
‘EEGReXferNet’ – A Lightweight Gen-AI Framework for EEG Subspace Reconstruction via Cross-Subject Transfer Learning and Channel-Aware Embedding

Shantanu Sarkar^{1*}, Piotr Nabrzyski^{1,2}, Saurabh Prasad¹, Jose Luis Contreras-Vidal¹

¹ IUCRC BRAIN Center, Cullen College of Engineering, University of Houston, Houston, TX, USA

² Department of Computer Engineering, Purdue University, West Lafayette, IN, USA

*Corresponding author: Shantanu Sarkar (shantanu75@gmail.com)

Abstract

Electroencephalography (EEG) is a widely used non-invasive technique for monitoring brain activity, but low signal-to-noise ratios (SNR) due to various artifacts often compromise its utility. Conventional artifact removal methods require manual intervention or risk suppressing critical neural features during filtering/reconstruction. Recent advances in generative models, including Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs), have shown promise for EEG reconstruction; however, these approaches often lack integrated temporal-spectral-spatial sensitivity and are computationally intensive, limiting their suitability for real-time applications like brain-computer interfaces (BCIs). To overcome these challenges, we introduce *EEGReXferNet*, a lightweight Gen-AI framework for EEG subspace reconstruction via cross-subject transfer learning - developed using Keras TensorFlow (v2.15.1). *EEGReXferNet* employs a modular architecture that leverages volume conduction across neighboring channels, band-specific convolution encoding, and dynamic latent feature extraction through sliding windows. By integrating reference-based scaling, the framework ensures continuity across successive windows and generalizes effectively across subjects. This design improves spatial-temporal-spectral resolution (mean PSD correlation ≥ 0.95 ; mean spectrogram RV-Coefficient ≥ 0.85), reduces total weights by $\sim 45\%$ to mitigate overfitting, and maintains computational efficiency for robust, real-time EEG pre-processing in BCI applications. The model and its components are openly shared to support reproducible research and cross-domain applications.

1 Introduction

Electroencephalography (EEG) is a non-invasive method that records brain activity via scalp electrodes by amplifying spontaneous potentials, providing insight into spatial, temporal, and spectral patterns associated with sensory and cognitive processes [1; 2]. EEG is widely used in clinical and neuroscience, cognitive and psychiatric studies due to its non-invasive, safe, and relatively inexpensive nature, its ability to directly measure neuronal electrical activity, and its real-time capability, which has also made it increasingly popular in brain-computer interface (BCI) applications[3; 4; 5].

EEG are highly susceptible to artifacts from physiological (e.g., ocular, muscular) and non-physiological (e.g., powerline, motion) sources, resulting in low signal-to-noise ratios (SNR) with reduced interpretability[4; 6; 7; 8; 9]. While denoising EEG signals by removing components from various unwanted sources, precise filtering is essential to filter artifacts that share frequency content similar to the underlying EEG features [10]. For closed-loop implementation of the BCI system, the

event-locked nature of many types of artifacts complicate real-time neural decoding [10]. Substantial efforts have been devoted to developing algorithms for EEG artifact removal, each offering distinct advantages and limitations [4; 7; 8; 9]. Independent Component Analysis (ICA) utilizing blind source separation (BSS) techniques is the most widely used methods; however, it requires manual intervention to exclude artifacts [8; 9]. To automate artifact removal, early approaches employed channel-wise statistical thresholding to eliminate abnormal activity patterns [11] - one such method is wavelet-based filtering with sparsity constraints [12]; however, these techniques often compromise the full feature representation of the affected channel or suppress spectral features in sparsity-enforced bands. Adaptive filtering techniques, such as H-Infinity [13; 14], have demonstrated promising results in real-time artifact denoising, but the need for a reference noise signal limits their application. The BSS methods like ICA and Principal Component Analysis (PCA) estimate the sources of artifacts considering a mixing matrix for original and observed signals by decomposing the EEG into transformed spaces[7; 8; 9]. The Artifact Subspace Reconstruction (ASR) decomposes EEG into a principal component (PC) space derived from clean data covariance and removes artifacts by suppressing components with abnormally high variance[11]. Transformed spaces like PC space are linear combinations of all channels; suppressing contaminated components in transformed space may alter the data structure, potentially losing essential features [15]. With recent developments in Generative AI, the Variational Autoencoder (VAE)- and Generative Adversarial Network (GAN)-based models are gaining popularity as an EEG pre-processing method [16; 17]. Generative reconstruction methods often overlook spatial channel relationships, use heavy encoder-decoder models with weak temporal-spectral coupling, and lack consistent mapping across successive sliding windows. To address these limitations, we propose *EEGReXferNet* - a novel, lower memory footprint architecture for EEG channel subspace reconstruction that leverages volume conduction across neighboring channels[1; 18], integrates spatial structure and dynamic temporal-spectral encoding, and applies reference scaling (μ_{Ref} , σ_{Ref}) for temporal continuity. *EEGReXferNet* is trained via Cross-Subject Transfer Learning to enhance generalizability across individuals.

2 Model Architecture

EEGReXferNet adopts a modular design with task-specific layers and a custom loss function, as shown in Figure 1. Implemented in Keras TensorFlow (v2.15.1), it supports scalable, interpretable, GPU-efficient modeling. Each layer and loss function are detailed below.

Neighborhood-Driven Input Selection: The model takes multichannel EEG windows (B, C, W) as input, from which it selects neighboring channels via a predefined dictionary. This dictionary maps each EEG channel to the indices of its nearest neighbors (L2 distance < 0.05) in the 10–20 system [19]. During training, conditional channel dropout is applied using `SpatialDropout1D`, which drops one or two neighbors based on channel counts (≤ 3). Finally, depth-wise convolution aggregates spatial dependencies into (B, 1, W). Here, B is the batch size, C is the total number of channels, and W is the size of the sliding window.

Sub-Window Convolution Encoding Block: The output of the “Neighborhood-Driven Input Selection” is passed through a series of custom `SubWindowConv1D` layers. Hartmann et al. showed that stacked convolutions extract fine-grained spectral features from EEG signals, with each layer specializing in distinct spectral characteristics [20]. Building on this insight, we implemented band-specific stacked convolutions using a custom-designed `SubWindowConv1D` layer, parameterized by kernel size, stride, filters, sub-window size, and tanh activation. To enforce spectral selectivity, we tailored the kernel sizes, strides, and sub-window configurations for each iteration. Parameters and targeted frequency bands are shown in Figure 1.

Sliding Stats Layer: This layer applies a sliding window mechanism to segment the input into overlapping temporal frames, where two lightweight dense layers estimate latent statistics. It enhances temporal resolution while reducing parameters by 45% compared to a dense layer for a 32-D latent space, thereby lowering memory use and mitigating overfitting—making it well suited for real-time, low-SNR applications. We used 160-ms sliding windows with a 40-ms stride to capture fine-grained, microstate-level temporal dynamics [21].

Sampling Layer and Latent Regularization: The `SamplingLyr` performs reparametrized sampling following the classical VAE framework by injecting Gaussian noise, enabling differentiable transformations of encoder outputs [22]. For latent regularization, we replace KL-divergence (KLD) with the geometry-aware, sample-based Sliced Wasserstein Distance (SWD)[23; 24] using 50 projections, improving gradient stability and reducing min–max conflicts in high-dimensional latent spaces.

Base Decoding Using Transposed Convolution: This block uses transposed convolution (deconvolution) to progressively up-sample the latent vector z into a structured feature map. Unlike dense layers, it provides spatially structured, parameter-efficient up-sampling[25].

Sub-Window Convolution Decoding Block: The output from the “Base Decoding” is passed through a series of custom SubWindowConv1D layers. Like encoding, each layer is parameterized by tailored kernel size, stride, number of filters, and sub-window size to implement band-specific stacked convolutions. The selected parameters for each iteration are shown in Figure 1 for the respective targeted frequency bands.

Linear Activation Layer: The output from the “Sub-Window Convolution Decoding Block” is flattened and passed to a linear-dense layer to produce the continuous-valued reconstruction.

Filtering and Scaling Layer: After decoding, the reconstructed outputs are passed sequentially through the RemoveOutlier and ScaleOutput Layer layers to produce a sample-wise adaptive reconstruction. The RemoveOutlier layer clips values beyond a z-score threshold, mitigating any spurious deviations in the decoded signal. Subsequently, the ScaleOutput Layer normalizes by sample-wise mean and SD and then rescales it using reference statistics (μ_{Ref} , σ_{Ref}). During training, reference statistics are computed from clean EEG segments, whereas for reconstruction of noisy subspaces, the layer utilizes statistics from previously clean segments. Random scaling factors ($\pm 10\%$) were applied to achieve stochastic regularization.

Loss Function: We used a latent regularization loss (\mathcal{L}_{Latent}) via SWD [23; 24], combined with a custom loss aggregating domain-specific Mean Square Error (MSE) and Hjorth mobility[26]. Temporal-MSE ($\mathcal{L}_{mse}^{\omega}$) and Magnitude-Spectrum-MSE ($\mathcal{L}_{mag}^{\omega}$) are weighted by a trainable uncertainty parameter[27] and multiplicatively coupled[28] with Phase-Spectrum-MSE (\mathcal{L}_{phase}) and Hjorth mobility loss ($\mathcal{L}_{mobility}$). The total loss function is defined by Equ. (1).

$$\mathcal{L}_{Total} = (\mathcal{L}_{mse}^{\omega} + \mathcal{L}_{mag}^{\omega}) \cdot (\mathcal{L}_{mobility} + 1) \cdot (\mathcal{L}_{phase} + 1) + \mathcal{L}_{latent} \quad (1)$$

3 Materials and Methods

3.1 Dataset

Our goal is to reconstruct contaminated EEG subspaces in real time and improve motor imagery (MI) decoding, focusing on BCI Competition IV (Dataset 1) [29]. The dataset includes 59-channel EEG from 7 subjects (a, b, g, f: human; others synthetic) performing binary MI tasks. Data were band-pass

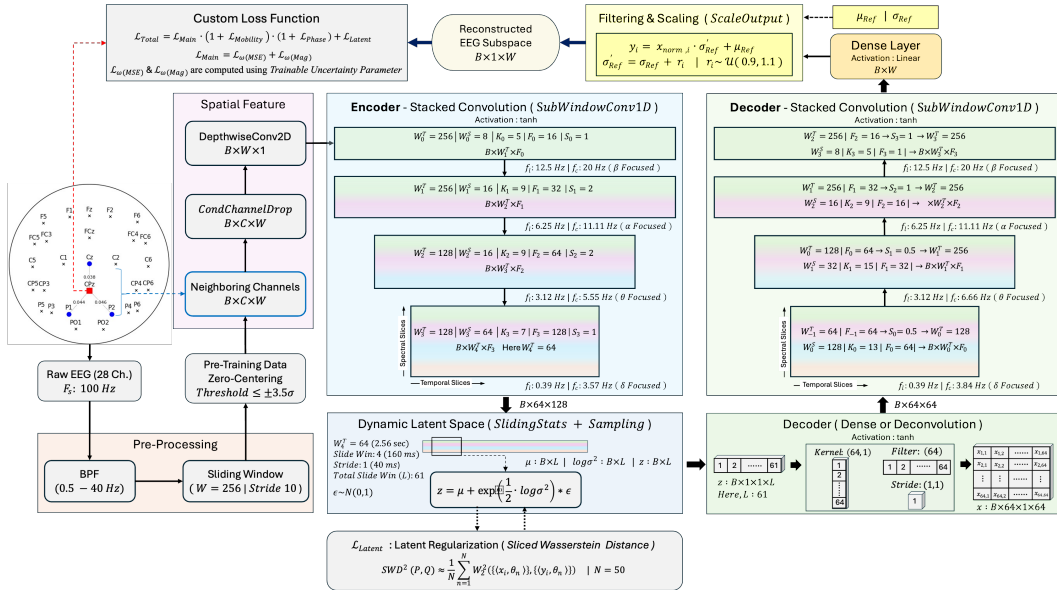


Figure 1: Overview of *EEGReXferNet* architecture illustrating key processing blocks and workflow.

filtered (BPF, 0.05–200 Hz) and digitized at 1000/100 Hz. This study used 100 Hz calibration data from human subjects, selecting 28 channels (Figure 1) matching our paradigm.

3.2 Pre-processing

To further preprocess the EEG, we applied a 6th-order zero-phase Butterworth BPF (0.5 - 40 Hz). Following BPF, to create the EEG subspace, we used a sliding window of 2.56 seconds to ensure the minimum frequency $\geq 0.4\text{Hz}$, with a stride of 100 ms. Within each sliding window, the channel-wise signal was re-centered to align with the global channel mean and added with a small random perturbation ($\pm 10\%$). We stratified the sliding windows into ‘Clean’ and ‘Noisy’ categories based on amplitude thresholding. Under ‘Clean’, we considered only those EEG windows, where the channel-wise amplitude range fell within specified limits ($> \pm 3.5\sigma$) and were utilized to train the model. The remaining windows were categorized as ‘Noisy’ and used for evaluation.

3.3 Methodology

To evaluate the key components of *EEGReXferNet*, we performed an ablation study across four configurations: (i) KLD (Model A) vs. SWD (Model B), (ii) fixed (Models A & B) vs. dynamic latent space (Models C & D), and (iii) dense (Model C) vs. deconvolution decoding (Model D). Models were trained on clean EEG data from three subjects, leaving one out for evaluation. Clean segment reconstruction used window/channel-wise scaling based on their stats (μ, σ). For reconstruction error metrics, we followed FDA work on EEG feature consistency [30; 31] and used Symmetric Mean Absolute Percentage Error (sMAPE) across channels/windows for relative $\delta, \theta, \alpha, \beta$ power (Figure 2), temporal/spectral entropy, and mobility, along with JS-Divergence. EEG channel-wise probability densities were estimated using the Dual Polynomial Regression method from the *estimatePDF* Python package[32]. Furthermore, we evaluated the MSE across the time, frequency (phase and magnitude), and time–frequency (spectrogram) domains. Noisy-window reconstruction utilized reference statistics ($\mu_{\text{Ref}}, \sigma_{\text{Ref}}$) from preceding clean windows, and was evaluated using PSD correlation (Pearson) and spectrogram similarity (RV coefficient). Model performance was compared using metrics across all channels with Friedman tests and post-hoc Nemenyi and Wilcoxon tests ($\alpha = 0.01$). For downstream classification, we trained *EEGNet-8-2*[33] with clean EEG windows and tested noisy ones as the baseline; misclassified windows were then reconstructed with Model D and re-evaluated. All models were trained with a batch size of 64 using Adam optimizer (Learning rate=0.001), early stopping (patience=25), up to 250 epochs, with 20% validation. For reproducibility, the backend was reset with a fixed seed prior to training for each channel, and a seed was similarly used for downstream evaluation. To train *EEGReXferNet* used Sabine Cluster (1 node, 28 cores, 1 GPU [16 GB]) and the rest of the task was done with a MacBook Pro (Apple M1 Max, 10 cores, 1 GPU).

4 Results

Friedman tests with post-hoc Nemenyi indicated that metrics differed significantly across models in most cases, which was further supported by Wilcoxon pairwise tests. Figure 3 presents the Wilcoxon rank-based heatmaps with overall mean scores. Models C and D, with dynamic latent space

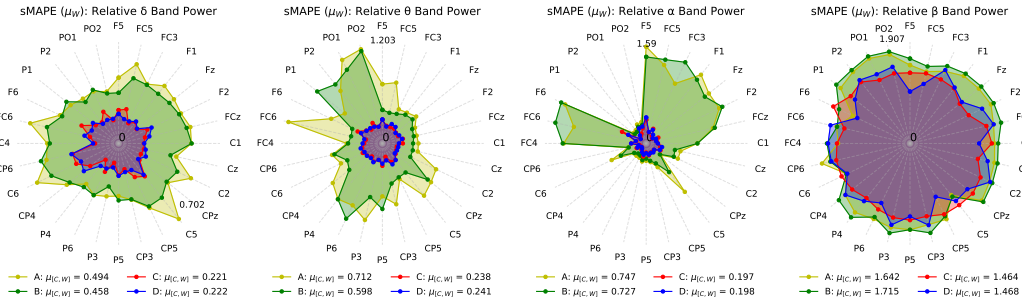


Figure 2: Subject ‘a’: sMAPE across EEG channels/windows for relative δ, θ, α , and β band power.

	(a) Subject-'a'				(b) Subject-'b'				(c) Subject-'f'				(d) Subject-'g'			
Temporal MSE	113.6	93.29	46.30 *	46.79	36.23	28.52	13.14 *	13.28	80.68	64.30	28.61 *	28.99	135.8	109.3	62.56 *	62.92
Magnitude Spectrum MSE	300.9	259.6	152.0 *	155.3	523.5	375.7	207.2 *	208.9	398.9	313.3	155.8 *	157.1	386.4	300.9	207.5 *	208.2
Phase Spectrum MSE	2.48	2.39	2.29	2.25 *	2.28	2.18	2.06	2.05 *	2.36	2.25	2.18 *	2.20	2.38	2.29	2.20 *	2.21
Spectrogram MSE	324.9	276.3	98.74 *	105.0	67.08	46.66	17.95 *	18.40	154.2	116.6	32.25 *	33.01	588.0	459.6	210.5 *	212.5
Relative δ Band Power	0.494	0.458	0.221 *	0.222	0.184	0.174	0.104	0.100 *	0.342	0.301	0.126 *	0.130	0.334	0.289	0.187	0.183 *
Relative θ Band Power	0.712	0.598	0.238 *	0.241	0.960	0.944	0.213 *	0.219	1.71	1.50	0.165 *	0.168	1.21	1.18	0.259 *	0.264
Relative α Band Power	0.747	0.727	0.197 *	0.198	0.738	0.530	0.342	0.328 *	0.481	0.411	0.214	0.212 *	0.531	0.373	0.235 *	0.246
Relative β Band Power	1.64	1.71	1.46 *	1.47	1.26	1.42	1.05 *	1.09	0.861	0.924	0.751 *	0.759	1.22	1.42	1.22	1.18 *
Entropy	0.011	0.010	0.009 *	0.009	0.012	0.011	0.009	0.009 *	0.012	0.011	0.009 *	0.009	0.013	0.012	0.010 *	0.010
JS Divergence	0.082	0.067	0.035 *	0.036	0.063	0.051	0.026 *	0.026	0.083	0.067	0.033 *	0.033	0.053	0.043	0.026 *	0.026
Spectral Entropy	0.071	0.091	0.061 *	0.062	0.075	0.076	0.062 *	0.065	0.067	0.067	0.044 *	0.046	0.086	0.102	0.083	0.083 *
Hjorth Mobility	0.044 *	0.069	0.054	0.054	0.062 *	0.087	0.119	0.119	0.045 *	0.055	0.061	0.062	0.084 *	0.101	0.106	0.103
Correlation (p) - PSD	0.920	0.942	0.964 #	0.963	0.952	0.962	0.961	0.966 #	0.908	0.933	0.970 #	0.970	0.881	0.917	0.956 #	0.954
RV-Coefficient - Spectrogram	0.810	0.855	0.907 #	0.901	0.834	0.864	0.911	0.915 #	0.789	0.846	0.927 #	0.926	0.667	0.718	0.858 #	0.855
	Model-A	Model-B	Model-C	Model-D	Model-A	Model-B	Model-C	Model-D	Model-A	Model-B	Model-C	Model-D	Model-A	Model-B	Model-C	Model-D

Figure 3: Model comparison across subjects using EEG metrics. Heatmaps show Wilcoxon ranks for (top) clean and (bottom) noisy data. In (top), * marks the best (min), in (bottom), # marks the best (max). Cells show mean scores. Color scale: green (min) → cyan → yellow → gray (max).

(45% fewer weights), consistently outperformed Models A and B across metrics, except for Hjorth mobility, where differences were minimal (≤ 0.05). Between the latent regularization strategies, SWD outperformed KLD, yielding more consistent improvements across subjects. Subject- and metric-specific differences were observed between Model C (dense) and Model D (deconvolution); however, their overall performance remained statistically consistent. Correlation analyses on noisy windows mirrored the clean-window findings: Models C and D exhibited higher correlations than Models A and B for both PSD and spectrogram-based measures. Downstream classification showed a marked improvement in accuracy metrics across all subjects when previously misclassified noisy EEG windows were reconstructed using Model-C and Model-D, and subsequently re-evaluated via EEGNet-8-2, as illustrated in Figure 4. Model-C yielded superior performance for Subjects-‘a’ and ‘b’, while Model-D outperformed for Subjects-‘f’ and ‘g’. Model-D exhibited the lowest mean training time across EEG channels (~ 11 min), in contrast to Model-A, which showed the highest (~ 16 min). Mean inference times per sliding window across EEG channels showed minimal variation: Model-A (0.75 ms), Model-B (0.77 ms), Model-C (0.78 ms), and Model-D (0.78 ms).




Figure 4: Comparison of accuracy metrics (Downstream classification using EEGNet-8-2) across subjects - Baseline vs. Reconstructed misclassified EEG windows via Model-C and D.

5 Conclusion

This work introduced *EEGReXferNet*, a novel lightweight Gen-AI framework for EEG subspace reconstruction that integrates temporal, spectral, and spatial sensitivity while maintaining computational efficiency. By leveraging cross-subject transfer learning and reference-based scaling, the model demonstrated robust generalization and continuity across successive windows. Models C and D, leveraging dynamic latent spaces with SWD regularization, consistently outperformed baselines across all metrics, except Hjorth mobility (minimal difference). For Model-C and D, although performance varied across subjects and metrics, both models exhibited consistent statistical behavior across the full evaluation spectrum. Furthermore, *EEGReXferNet* maintained consistently low inference latency (< 1 ms), underscoring its suitability for real-time neurophysiological applications. Future work should explore larger and more diverse datasets beyond MI, integrate adaptive artifact detection for real-time deployment, and analyze internal representations to better understand learned features. The custom layers and loss functions developed in Keras TensorFlow are GPU-optimized, modular, and broadly applicable—offering a scalable foundation for building advanced deep learning models across domains.

Code Availability

The code is available on GitHub (*ShanSarkar75/EEGReXferNet*), and the reconstructed EEG window data (.npz) are shared via FigShare (DOI: 10.6084/m9.figshare.30343642).

Acknowledgments and Disclosure of Funding

Supported by NSF IUCRC BRAIN and the NSF REU Site on Regulatory Science (Award #2349657). We thank Ramin Bighamian (FDA) and Xiaoqian Jiang (UTHealth) for their insightful feedback

References

- [1] Michael X. Cohen. Where Does EEG Come From and What Does It Mean? *Trends in Neurosciences*, 40: 208–218, 4 2017. ISSN 1878108X. URL <https://doi.org/10.1016/j.tins.2017.02.004>.
- [2] S. Siuly, Y. Li, and Y. Zhang. *Electroencephalogram (EEG) and Its Background*, pages 3–21. Springer International Publishing, 2016. ISBN 978-3-319-47653-7. URL https://doi.org/10.1007/978-3-319-47653-7_1.
- [3] Alexander Craik, Yongtian He, and Jose L. Contreras-Vidal. Deep Learning for Electroencephalogram (EEG) Classification Tasks: A Review. *Journal of Neural Engineering*, 16:31001, 2019. URL <https://doi.org/10.1088/1741-2552/ab0ab5>.
- [4] Alexander Craik, Juan José González-España, Ayman Alamir, David Edquilang, Sarah Wong, Lianne Sánchez Rodríguez, Jeff Feng, Gerard E. Francisco, and Jose L. Contreras-Vidal. Design and Validation of a Low-Cost Mobile EEG-Based Brain–Computer Interface. *Sensors*, 23, 2023. ISSN 14248220. URL <https://doi.org/10.3390/s23135930>.
- [5] Richard James Sugden, Viet-Linh Luke Pham-Kim-Nghiem-Phu, Ingrid Campbell, Alberto Leon, and Phedias Diamandis. Remote collection of electrophysiological data with brain wearables: opportunities and challenges. *Bioelectronic Medicine*, 9:12, 2023. ISSN 2332-8886. URL <https://doi.org/10.1186/s42234-023-00114-5>.
- [6] Avinash Tandle, Nandini Jog, Pancham D’cunha, and Monil Chheta. Classification of Artefacts in EEG Signal Recordings and EOG Artefact Removal using EOG Subtraction. *Communications on Applied Electronics*, 4:12–19, 1 2016. ISSN 2394-4714. URL <https://doi.org/10.5120/cae2016651997>.
- [7] Jose A. Uriguen and Begonya Garcia-Zapirain. EEG artifact removal—state-of-the-art and guidelines. *Journal of Neural Engineering*, 12:31001, 2015. URL <https://doi.org/10.1088/1741-2560/12/3/031001>.
- [8] Xianyun Jiang, Gang-Bing Bian, and Zhi Tian. Removal of Artifacts from EEG Signals: A Review. *Sensors (Basel, Switzerland)*, 19:987, 2019. URL <https://doi.org/10.3390/s19050987>.
- [9] C. R. Rashmi and C. P. Shantala. EEG artifacts detection and removal techniques for brain computer interface applications: a systematic review. *International Journal of Advanced Technology and Engineering Exploration*, 9:354–383, 3 2022. ISSN 23947454. URL <https://doi.org/10.19101/IJATEE.2021.874883>.
- [10] Atilla Kilicarslan and Jose L. Contreras-Vidal. *Neuro-Robotics: Rehabilitation and Restoration of Walking Using Exoskeletons via Non-invasive Brain–Machine Interfaces*, pages 143–166. Springer International Publishing, 2021. ISBN 978-3-030-68545-4. URL https://doi.org/10.1007/978-3-030-68545-4_6.
- [11] Chi-Yuan Chang, Sheng-Hsiou Hsu, Luca Pion-Tonachini, and Tzyy-Ping Jung. Evaluation of Artifact Subspace Reconstruction for Automatic Artifact Components Removal in Multi-Channel EEG Recordings. *IEEE Transactions on Biomedical Engineering*, 67:1114–1121, 2020. URL <https://doi.org/10.1109/TBME.2019.2930186>.
- [12] G. Geetha and S. N. Geethalakshmi. EEG De-noising using SURE Thresholding based on Wavelet Transforms. *International Journal of Computer Applications*, 24:29–33, 6 2011. ISSN 0975-8887. URL <https://doi.org/10.5120/2948-3935>.
- [13] Atilla Kilicarslan, Robert G. Grossman, and Jose L. Contreras-Vidal. A Robust Adaptive Denoising Framework for Real-Time Artifact Removal in Scalp EEG Measurements. *Journal of Neural Engineering*, 13:26013, 2 2016. URL <https://doi.org/10.1088/1741-2560/13/2/026013>.

- [14] Atilla Kilicarslan and Jose L. Contreras-Vidal. Characterization and Real-Time Removal of Motion Artifacts from EEG Signals. *Journal of Neural Engineering*, 16:56027, 9 2019. URL <https://doi.org/10.1088/1741-2552/ab2b61>.
- [15] Abdulrahman Oladipupo Ibraheem. Correlation of Data Reconstruction Error and Shrinkages in Pair-wise Distances under Principal Component Analysis (PCA). *arXiv*, 1412.6752, 2014. URL <https://doi.org/10.48550/arXiv.1412.6752>.
- [16] Jamal F. Hwaidi and Thomas M. Chen. A Noise Removal Approach from EEG Recordings Based on Variational Autoencoders. In *2021 13th International Conference on Computer and Automation Engineering (ICCAE)*, pages 19–23, 2021. URL <https://doi.org/10.1109/ICCAE51876.2021.9426150>.
- [17] Congzhong Sun and Chaozhou Mou. Survey on the research direction of EEG-based signal processing. *Frontiers in Neuroscience*, 17, 7 2023. ISSN 1662-453X. URL <https://doi.org/10.3389/fnins.2023.1203059>.
- [18] Terrence D. Lagerlund, Devon I. Rubin, and Jasper R. Daube. *Volume Conduction*, pages 929–946. Oxford University Press, 4 edition, 6 2016. URL <https://doi.org/10.1093/med/9780190259631.003.0054>.
- [19] Koen B. E. Böcker, Jurgen A. G. van Avermaete, and Margaretha M. C. van den Berg-Lenssen. The international 10–20 system revisited: Cartesian and spherical co-ordinates. *Brain Topography*, 6:231–235, 1994. ISSN 1573-6792. URL <https://doi.org/10.1007/BF01187714>.
- [20] Kay Gregor Hartmann, Robin Tibor Schirrmester, and Tonio Ball. Hierarchical internal representation of spectral features in deep convolutional networks trained for EEG decoding. In *2018 6th International Conference on Brain-Computer Interface, BCI 2018*, volume 2018-January, 2018. URL <https://doi.org/10.1109/IWW-BCI.2018.8311493>.
- [21] Christoph M. Michel and Thomas Koenig. EEG microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: A review. *NeuroImage*, 180:577–593, 2018. ISSN 10959572. URL <https://doi.org/10.1016/j.neuroimage.2017.11.062>.
- [22] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv*, 1312.6114, 2013. URL <https://doi.org/10.48550/arXiv.1312.6114>.
- [23] Soheil Kolouri, Yang Zou, and Gustavo K. Rohde. Sliced Wasserstein kernels for probability distributions. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5258–5267, 2016. URL <https://doi.org/10.1109/CVPR.2016.568>.
- [24] Soheil Kolouri, Phillip E. Pope, Charles E. Martin, and Gustavo K. Rohde. Sliced-Wasserstein Autoencoder: An Embarrassingly Simple Generative Model. *arXiv*, 1804.01947, 2018. URL <https://doi.org/10.48550/arXiv.1804.01947>.
- [25] Matthew D. Zeiler, Dilip Krishnan, Graham W. Taylor, and Rob Fergus. Deconvolutional networks. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010. URL <https://doi.org/10.1109/CVPR.2010.5539957>.
- [26] Raja Majid Mehmood, Ruoyu Du, and Hyo Jong Lee. Optimal feature selection and deep learning ensembles method for emotion recognition from human brain EEG sensors. *IEEE Access*, 5:14797–14806, 2017. ISSN 21693536. URL <https://doi.org/10.1109/ACCESS.2017.2724555>.
- [27] Roberto Cipolla, Yarin Gal, and Alex Kendall. Multi-task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 7482–7491, 2018. URL <https://doi.org/10.1109/CVPR.2018.00781>.
- [28] Charles D. Coleman. Total Loss Functions for Measuring the Accuracy of Nonnegative Cross-Sectional Predictions. *arXiv*, 2507.15136, 2025. URL <https://doi.org/10.48550/arXiv.2507.15136>.
- [29] Benjamin Blankertz, Guido Dornhege, Matthias Krauledat, Klaus-Robert Müller, and Gabriel Curio. The non-invasive Berlin Brain–Computer Interface: Fast acquisition of effective performance in untrained subjects. *NeuroImage*, 37:539–550, 8 2007. ISSN 1053-8119. URL <https://doi.org/10.1016/j.neuroimage.2007.01.051>.
- [30] David O. Nahmias, Kimberly L. Kontson, David A. Soltysik, and Eugene F. Civillico. Consistency of quantitative electroencephalography features in a large clinical data set. *Journal of Neural Engineering*, 16:066044, 2019. URL <https://doi.org/10.1088/1741-2552/ab4af3>.

- [31] David O. Nahmias and Kimberly L. Kontson. Quantifying Signal Quality From Unimodal and Multimodal Sources: Application to EEG With Ocular and Motion Artifacts. *Frontiers in Neuroscience*, 15, 2021. ISSN 1662-453X. URL <https://doi.org/10.3389/fnins.2021.566004>.
- [32] Shantanu Sarkar. estimatePDF: Probability Density Estimation using Dual Polynomial Regression (DPR). <https://pypi.org/project/estimatePDF/>, 2025. Python package.
- [33] 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*, 15:56013, 7 2018. URL <https://doi.org/10.1088/1741-2552/aace8c>.

A Technical Appendices and Supplementary Material

A.1 Additional Tables

Table 1: Comparison of model configurations used for ablation study.

Model	Latent & Decoding Approach	Regularization	Weights	Reduction(%)
A	Standard 32 Latent Space Base Decode: Dense	KL Divergence	896,198	0.00%
B	Standard 32 Latent Space Base Decode: Dense	Sliced Wasserstein	896,198	0.00%
C	Dynamic Latent from Sliding Window (160 ms, Stride=40 ms) Base Decode: Dense	Sliced Wasserstein	491,656	45.13% ↓
D	Dynamic Latent from Sliding Window (160 ms, Stride=40 ms) Base Decode: Deconvolution	Sliced Wasserstein	487,624	45.58% ↓

Note: Typical case using a set of five neighboring channels.

Table 2: Parameters (w.r.t. target band) used in Sub-Window Convolution Encoding Block

Iter.	Target Band (Hz)	Filter	Inp Win W_i^T	F_S (Hz)	Kernel K	f_c (F_S/K)	Sub-Win W_S	f_L (F_S/W_S)	Stride	Out Win W_{i+1}^T
1	Beta (12–30)	16	256	100	5	20.00 Hz	8	12.50 Hz	1	256
2	Alpha (8–12)	32	256	100	9	11.11 Hz	16	6.25 Hz	2	128
3	Theta (4–8)	64	128	50	9	5.55 Hz	16	3.12 Hz	2	64
4	Delta (0.5–4)	128	64	25	7	3.57 Hz	64	0.39 Hz	1	64

Table 3: Parameters (w.r.t. target band) used in Sub-Window Convolution Decoding Block

Iter.	Target Band (Hz)	Filter	Inp Win W_{i-1}^T	Stride	Out Win W_i^T	F_S (Hz)	Kernel K	f_c (F_S/K)	Sub-Win W_S	f_L (F_S/W_S)
1	Delta (0.5–4.0)	64	64	0.5	128	50	13	3.84 Hz	128	0.39 Hz
2	Theta (4.0–8.0)	32	128	0.5	256	100	15	6.66 Hz	32	3.12 Hz
3	Alpha (8.0–12.0)	16	256	1	256	100	9	11.11 Hz	16	6.25 Hz
4	Beta (12–30)	1	256	1	256	100	5	20.00 Hz	8	12.50 Hz

A.2 Additional Figures




Figure 5: Channel-wise mean and standard deviation of EEG signals for subject ‘a’ before and after band-pass filtering (BPF).




Figure 6: Comparison of reconstruction error across models while reconstructing clean EEG segments ($\leq \pm 3.5\sigma$) of subject ‘a’ using four MSE-based metrics. Radar plots illustrate the performance of Models A–D across EEG channels for (a) Temporal MSE, (b) Magnitude Spectrum MSE, (c) Phase Spectrum MSE, and (d) Spectrogram MSE.

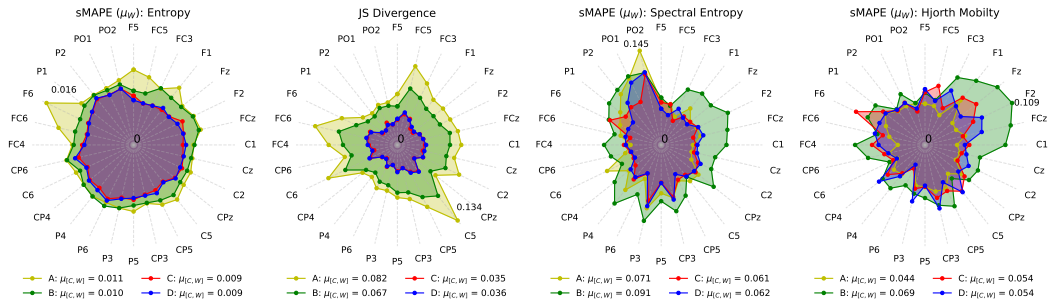


Figure 7: Comparison of model performance across EEG channels using entropy-based and distributional metrics. Radar plots show the reconstruction accuracy of Models A–D for (a) Entropy, (b) JS Divergence, (c) Spectral Entropy, and (d) Hjorth Mobility, evaluated on clean EEG segments ($\leq \pm 3.5\sigma$) of Subject ‘a’. Metrics are reported as symmetric Mean Absolute Percentage Error (sMAPE) or divergence values, with lower values indicating better performance.

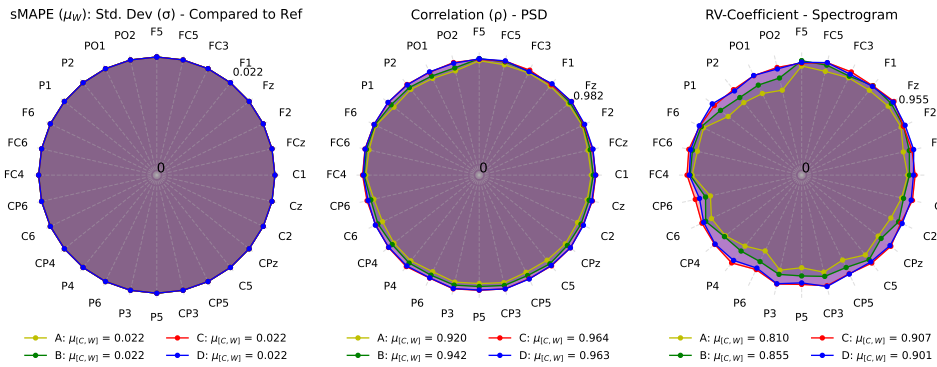


Figure 8: Comparison of reconstruction quality across models for noisy EEG signals ($> \pm 3.5\sigma$) from Subject ‘a’ data. Signals were scaled using reference statistics (μ_{Ref}, σ_{Ref}) derived from preceding clean segments. Radar plots show model-wise performance across EEG channels for: (a) sMAPE of standard deviation relative to reference (σ_{Ref}) (b) Pearson correlation (ρ) of power spectral density (PSD), and (c) RV-Coefficient of the spectrogram. All models show consistent scaling, confirming the effectiveness of the ScaleOutput layer.

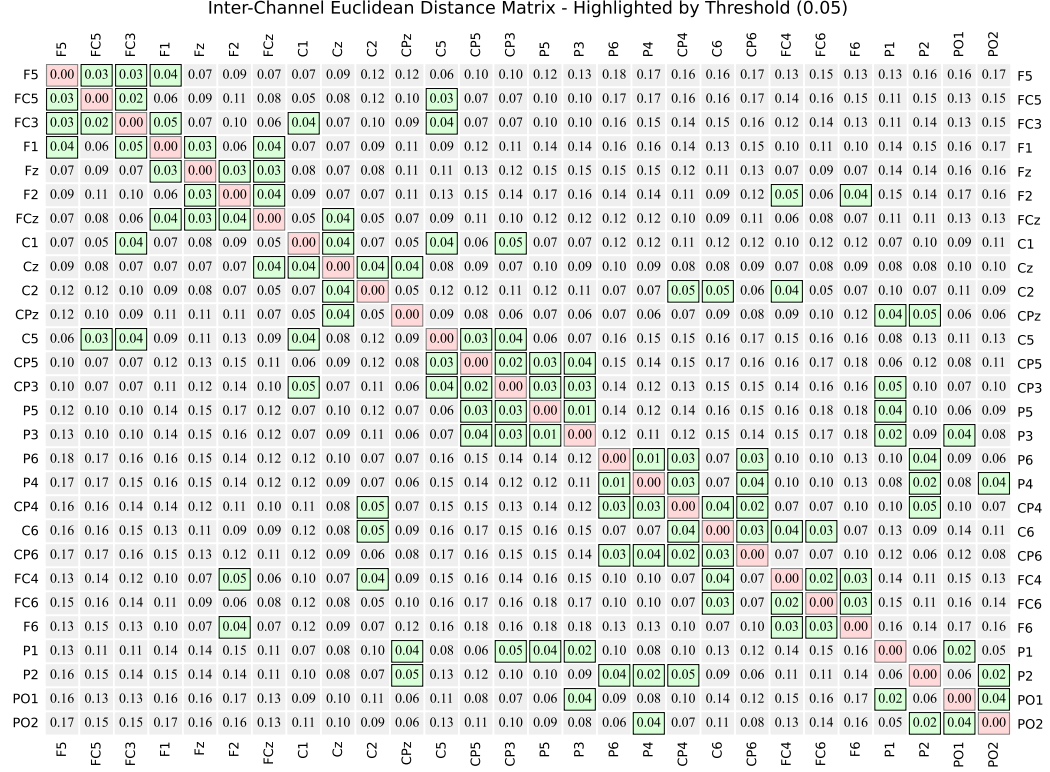


Figure 9: Inter-channel Euclidean distance matrix based on the 10–20 EEG electrode placement standard. Each row represents a reference EEG channel (red cell), and green cells indicate neighboring channels whose Euclidean distance from the reference is below or equal to the threshold of 0.05. This spatial neighborhood mapping was used to define channel-specific subspaces.




Figure 10: Training and validation loss curves for four models (A–D), trained using clean EEG windows from subjects ‘b’, ‘c’, and ‘d’ to reconstruct the EEG subspace of Channel Cz for subject ‘a’. The blue line represents training loss, while the green line denotes validation loss. Key metrics—including the Elbow Point, Best Validation Loss, Maximum Absolute Difference (Δ), and Final Epoch Δ —are annotated for each model. Model A exhibits prolonged training with poor generalization and pronounced overfitting. Model B achieves improved validation loss but still shows signs of overfitting. Models C and D demonstrate superior generalization, characterized by lower validation loss and minimal overfitting.




Figure 11: Scatter plots showing the correlation between Total Validation Loss (y-axis) and individual validation loss components (x-axis) across Models A–D at the best epoch, evaluated over 28 channels. Each color represents a different loss component: Temporal MSE (blue), Magnitude Spectrum MSE (orange), Phase Spectrum MSE (red), Hjorth Mobility (green), and Latent Regularization (purple; KLD in Model A, SWD in Models B–D). Trend lines indicate the strength and direction of correlation.

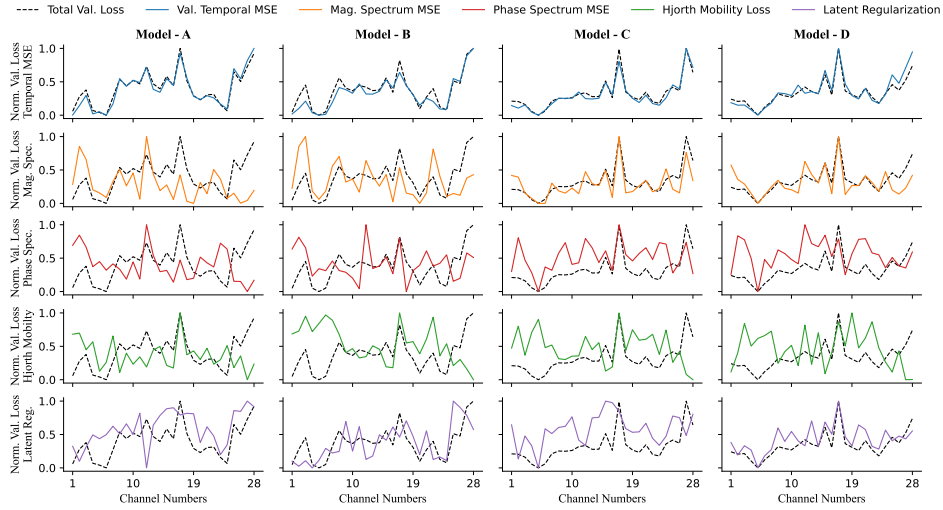


Figure 12: Subplots showing channel-wise comparison of individual validation loss components (Temporal MSE, Magnitude Spectrum MSE, Phase Spectrum MSE, Hjorth Mobility, Latent Regularization-KLD/SWD) against Total Validation Loss (dashed black line) across Models A–D at the best epoch. Each row corresponds to a specific loss component, and each column represents a model. Models were trained independently per channel using clean data from subjects 'b', 'c', and 'd' (tested on subject 'a'). The layout highlights how each component tracks with total loss across channels.

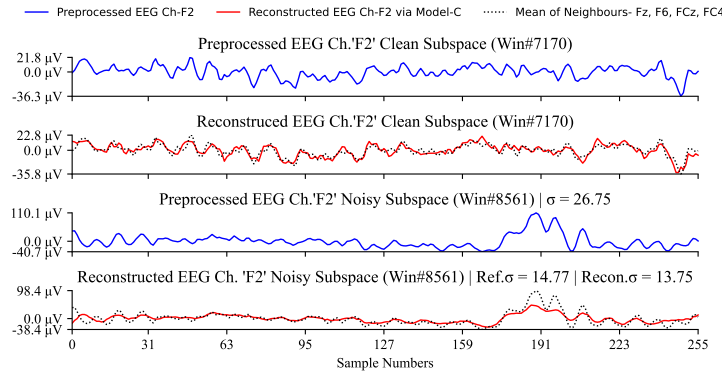


Figure 13: Comparison of clean and noisy EEG signals from channel F2 (subject 'a') with Model-C reconstructions, highlighting alignment with neighboring channels and signal variability.




Figure 14: Normalized validation metrics across training epochs for Models A–D, trained on clean EEG data from subjects ‘b’, ‘c’, and ‘d’ (tested on subject ‘a’) for Cz-channel reconstruction, reveal distinct convergence behaviors. Each model’s Total Validation Loss reflects contributions from Temporal MSE, Magnitude/Phase Spectrum MSE, Mobility, and Latent Regularization (SWD), alongside auxiliary metrics—Kurtosis, Skewness, and Variance—normalized to [0, 1] for comparability. Model-A demonstrates early high performance but exhibits pronounced fluctuations beyond the initial epochs, suggesting potential overfitting and reduced generalization. In contrast, Model-B and Model-D converge more rapidly and maintain stable validation trajectories, indicating better optimization dynamics and robustness. The ‘Best’ epoch was identified as the point of minimum total validation loss, while training termination was governed by a patience setting of 25 epochs.




Figure 15: Training time per model across EEG channels and subjects. Mean training durations are highlighted for each model: Model-A (16 min), Model-B (12.5 min), Model-C (12.8 min), and Model-D (11.1 min), illustrating comparative computational efficiency.




Figure 16: Channel-wise average inference time per model across subjects. Solid lines indicate mean inference times: Model-A (0.75 ms), Model-B (0.77 ms), Model-C (0.78 ms), and Model-D (0.78 ms). All models demonstrate sub-millisecond latency with minimal channel-wise variation, supporting suitability for real-time EEG applications.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract and introduction clearly outline the limitations of existing EEG artifact removal and EEG subspace reconstruction methods and generative approaches, and consistently present EEGReXferNet as a lightweight, modular framework for EEG subspace reconstruction. We describe the key features, such as volume conduction across neighboring channels, band-specific convolution encoding, dynamic latent feature extraction, and reference-based scaling, as well as cross-subject transfer learning for generalization. The contributions and scope, including improved temporal, spectral, and spatial resolution and suitability for real-time BCI applications, are clearly stated and align with the methods and expected outcomes presented in the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Yes, in the paper we addressed limitations and outlined directions for future work in the conclusion. Specifically, it notes that the model should be evaluated on larger and more diverse datasets (beyond MI), integrated with adaptive artifact detection for real-time applicability, and further analyze the features from inner layers to understand its learned representations. The model EEGReXferNet is built on assumptions of volume conduction between neighboring EEG channels, which are explicitly stated and supported by prior literature. EEGReXferNet is designed with a modular architecture with task-specific layers with scalability in mind. And we also shared the computational efficiency metrics (training and inference times) across models with detailed hardware specifications and parameters. While limitations are not presented in a standalone section, they are meaningfully embedded throughout the paper, particularly in the methodology, results, and conclusion.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be

used reliably to provide closed captions for online lectures because it fails to handle technical jargon.

- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with data set size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

The main body of the paper focuses on empirical evaluation, model development, and EEG reconstruction experiments, while theoretical assumptions such as volume conduction or gradient stability in Sliced Wasserstein Distance (SWD) are supported by prior work.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper fully discloses the information needed to reproduce the main experimental results that support its key claims and conclusions. It provides a detailed description of the model architecture, including all custom layers, layer parameters (kernel sizes, strides, filters, sub-window sizes, activations), and loss functions. The preprocessing steps, EEG channel selection, sliding-window procedure, and criteria for "Clean" versus "Noisy" windows are clearly explained. Training procedures—including dataset splits, batch size, learning rate, early stopping, number of epochs, and hardware used—are specified.

To ensure reproducibility, the backend was reset with a fixed seed prior to training for each channel, and a seed was similarly used for downstream evaluation. This allows other researchers to reproduce both the model training and the evaluation pipeline with minimal variability.

Evaluation metrics and statistical analyses are provided in sufficient detail, including sMAPE, MSE, PSD correlation, spectrogram RV-Coefficient, and relevant post-hoc tests.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.

- If the contribution is a data set and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same data set, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source data set or instructions for how to construct the data set).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: The study utilizes the publicly available BCI Competition IV Dataset 1 and provides full access to the code and reproduction instructions via GitHub (<https://github.com/ShanSarkar75/EEGReXferNet/>). Reconstructed EEG window data (.npz) are shared through FigShare (DOI: 10.6084/m9.figshare.30343642), enabling faithful replication of all main experimental results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: The paper specifies all key training and test details necessary to understand the results. This includes data splits (leave-one-subject-out), hyperparameters (batch size 64, learning rate 0.001), type of optimizer (Adam), early stopping (patience=25, max 250 epochs), and validation fraction (20%). Hyperparameters were chosen based on standard practice for stable convergence and prior EEGNet training experience. Preprocessing steps, such as window/channel-wise scaling using μ and σ , are detailed. The methodology also describes ablation study configurations, downstream classification evaluation, computational resources, and reproducibility measures (backend resets and fixed seeds), enabling faithful replication of the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: The paper reports statistical significance using well-defined error measures and appropriate tests. Channel- and window-wise reconstruction errors were quantified using Symmetric Mean Absolute Percentage Error (sMAPE) for relative δ , θ , α , and β power, temporal/spectral entropy, and Hirsch mobility, along with JS-Divergence and MSE computed on temporal, spectral (magnitude and phase), and spectrogram to capture reconstruction error.

For noisy-window reconstructions, we verified the fidelity of the reconstructed signals by assessing their correlation with the original signals, despite amplitude differences caused by scaling with varying reference statistics from preceding clean windows. Specifically, Pearson correlation was used to evaluate similarity in power spectral density (PSD), and the RV-Coefficient quantified the similarity between reconstructed and original spectrograms, capturing the preservation of the time–frequency structure.

Statistical comparisons were performed using Friedman tests with post-hoc Nemenyi and Wilcoxon tests ($\alpha = 0.01$) on flattened arrays including all channels.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The paper provides detailed information on the computational resources used for each experiment. EEGReXferNet training was performed on the Sabine Cluster using 1 node with 28 CPU cores and 1 GPU [16 GB], while all other training and evaluation tasks—including EEGNet-8-2 training and downstream analyses — were conducted on a MacBook Pro (Apple M1 Max, 10 CPU cores, integrated GPU). Channel-wise training and inference time per model across subjects are illustrated via Figure 15 and Figure 16 [Technical Appendices and Supplementary Material].

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification:

- Research involving human subjects or participants: Not applicable in current scope (Publicly available EEG datasets used)
- Data-related concerns:
 - Privacy: Datasets contain no personally identifiable information
 - Consent: Not applicable (Only publicly released datasets were used)
 - Deprecated datasets: Publicly available datasets (BCI Competition IV Dataset 1) were used
 - Copyright and Fair Use: Cited - BCI Competition IV Dataset 1
 - Representative evaluation practice: The dataset represents motor imagery EEG from healthy subjects; claims are limited to this population
- Societal Impact and Potential Harmful Consequences:
 - Safety & Security: Focus on brain-computer interfaces for motor imagery; no direct safety risks
 - Discrimination: No discrimination; analyses are neutral regarding gender, race, or protected characteristics
 - Surveillance: Not applicable (No bulk surveillance data used)
 - Deception & Harassment: Study cannot be used to deceive, impersonate, or harm individuals

- Environment: Experiments conducted on standard computational hardware; minimal impact
- Human Rights: No violations
- Bias and fairness: Limitations acknowledged; evaluation restricted to dataset population
- Impact Mitigation Measures:
 - Data and model documentation: Fully documented in paper and supplemental material
 - Data and model licenses: Code and models will be released under an open-access license for research
 - Secure and privacy-preserving storage & distribution: Only public datasets used; no sensitive data collected
 - Responsible release and publication strategy: Code and trained models to be released via GitHub with instructions for responsible use
 - Allowing access to research artifacts: Code, models, hyperparameters, and computational requirements provided for reproducibility
 - Disclose essential elements for reproducibility: Model architecture, training procedures, data splits, hyperparameters, and computational resources fully reported
 - Ensure legal compliance: Data use complies with the original dataset terms

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: The paper highlights potential positive impacts, such as advancing brain-computer interface methods for motor imagery decoding, which could support future assistive technologies for individuals with motor impairments. Potential negative impacts are minimal as the work uses publicly available EEG datasets and focuses on assistive applications; no misuse is anticipated.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.

- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The study uses publicly available, anonymized EEG datasets. No pretrained language models or other sensitive datasets are used, so there is no risk of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All assets used in this paper, including datasets and algorithms, are publicly available and have been carefully reviewed for license and terms of use. The original creators are properly credited in the manuscript. Specifically, this study uses the BCI Competition IV Dataset 1, which is publicly distributed for academic research purposes. The dataset does not specify an explicit license, but its use is restricted to non-commercial research. The dataset is cited as: Blankertz, B., Dornhege, G., Krauledat, M., Müller, K.-R., & Curio, G. (2007). The non-invasive Berlin Brain-Computer Interface: Fast acquisition of effective performance in untrained subjects. *NeuroImage*, 37(2), 539–550.

Dataset Link: https://www.bbci.de/competition/iv/desc_1.html

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The study provides full access to the source code and detailed reproduction instructions via GitHub (<https://github.com/ShanSarkar75/EEGReXferNet/>). Additionally, the reconstructed EEG window data (.npz) are publicly available on FigShare (DOI: 10.6084/m9.figshare.30343642), enabling faithful replication of all main experimental results. The documentation clearly describes the training details, version, and known limitations to ensure transparency and reproducibility. No human subject data is directly released, so consent considerations are not applicable.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The study does not involve crowdsourcing experiments or direct interactions with human participants. Only publicly available, anonymized EEG datasets.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This study exclusively utilizes publicly available, anonymized EEG datasets that were previously collected under appropriate institutional approvals. No new data were gathered from human participants for this work, and therefore, no direct risks to participants were incurred or disclosed. For future real-time testing of the model, an approved IRB exists, but this is outside the scope of the current study.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The research does not use any large language models (LLMs) as part of the core methodology. No LLMs were involved in the experiments, model development, or analysis.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.