Conditional Diffusion-Based EEG Channel Reconstruction for Motor Imagery Decoding

Anonymous Authors[1]

Abstract

This paper proposes a diffusion-based framework using Denoising Diffusion Probabilistic Models (DDPM) for EEG signal reconstruction in Motor Imagery (MI)-based Brain–Computer Interface (BCI) applications. The framework is integrated with CSP+LDA, FBCSP+LDA, and EEGNet decoding baselines to improve robustness while preserving decoding performance. The proposed method reconstructs predefined EEG channels from spatially adjacent signals to maintain inter-channel dependencies, and was evaluated using training sessions and an independent held-out test session. The work highlights areas for further improvement in EEG-based Motor Imagery decoding models in Latin America.

1. Introduction

Motor imagery–based brain–computer interfaces(MI-BCIs) enable users to control external systems by decoding sensorimotor patterns during imagined movements, providing a non-invasive and communication paradigm (1). Nonetheless, the practical application of MI-BCIs still faces critical challenges (2), particularly in achieving stable and high-quality EEG signals due to their susceptibility to noise (2), sensor dropout, and electrode failure (4). These factors directly affect the motor imagery (MI) signal quality in terms of spatially distributed patterns and, consequently, decoding performance. Traditional methods, although widely used, often struggle to capture the complex, high-dimensional probability distributions that characterize real MI-related neural activity, leading to inconsistent task execution and degraded EEG signal quality (5). Therefore, addressing these limitations is essential for improving MI-BCI performance and advancing EEG-based brain state recognition.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

In MI-BCI systems, generative models have emerged as a promising research direction due to their ability to enhance the reconstruction quality of EEG signals used for decoding (6; 7). Recently, Denoising Diffusion Probabilistic Models (DDPMs) have been explored with the goal of expanding training datasets while capturing the intrinsic complexity of real EEG signals (8; 9; 10). In its standard formulation, DDPM learns the marginal distribution without additional contextual information (11). Although this formulation offers modeling simplicity, it may compromise reconstruction fidelity in EEG signals. In MI tasks, preserving class-specific spatial patterns, such as event-related desynchronization (ERD), is essential for maintaining discriminative information (12; 13). To overcome this limitation, conditional DDPMs incorporate condition inputs reconstruction (14). For example, in MI tasks, unconditional generation may fail to maintain class-specific spatial patterns, such as the lateralized event-related desynchronization (ERD) over contralateral sensorimotor regions observed during left or right hand imagery (15). In contrast, the conditional DDPM framework performs reconstruction by conditioning on observed channel data or task labels, thereby facilitating improved preservation of task-specific spatial and temporal dynamics (16; 17). The key contributions are summarized as follows:

(1) Proposed a conditional diffusion-based framework for EEG channel reconstruction in MI decoding.

(2) Developed a controlled evaluation strategy using the following classifiers: CSP (18) +LDA, FBCSP (19) +LDA, and EEGNet (20) under a rigorous experimental protocol.

(3) Experimental results demonstrating that framework preserves EEG signal reconstruction fidelity and task-relevant spatio-temporal patterns.

2. Method

We propose a novel DDPM-based framework for robust EEG target-channel reconstruction integrated into conventional MI decoding pipelines. The method conditionally reconstructs predefined target (also referred to as ”hybrid”) channels from spatially adjacent reference signals while preserving the original channels, maintaining inter-channel dependencies and stable decoding performance. The workflow of the entire framework is shown in Fig. 1.

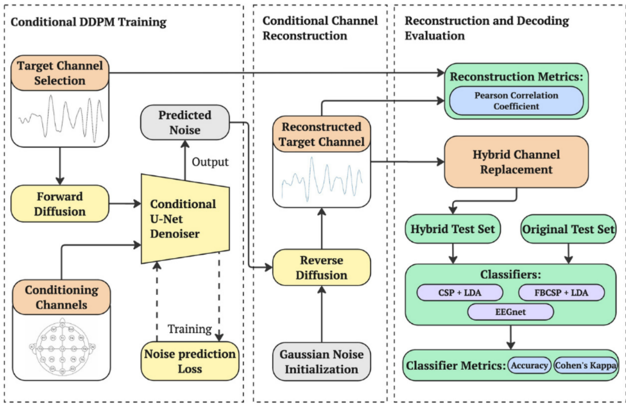

Figure 1. Workflow of the proposed conditional diffusion-based EEG channel reconstruction framework.

2.1. Problem Formulation

Let $\mathbf{X} \in \mathrm{R}^{Q \times T}$ represent an original MI EEG signal, where Q denotes the number of channels and T the temporal length. Given a predefined subset of target channels $\mathcal{T} \subset \{1, \ldots, Q\}$, the objective is to reconstruct each target channel signal $x_0^{(j)} \in \mathrm{R}^T$ for $j \in \mathcal{T}$ from a set of spatially adjacent conditioning channels $\mathbf{C}^{(j)} \in \mathrm{R}^{K \times T}$, where K denotes the number of conditioning channels, while all remaining channels remain unchanged. The original channel signal is defined as $\mathbf{X}_j^{\text{original}}$, which corresponds to the real EEG recording of channel j.

After reconstruction, the hybrid EEG representation is defined as

$$X_j^{\text{hybrid}} = m_j \hat{x}_0^{(j)} + (1 - m_j) X_j^{\text{original}} \tag{1}$$

where $m_j \in \{0, 1\}$ denotes a binary mask indicating reconstructed channels.

2.2. Conditional DDPM Training

Target Channel Selection. To enable controlled reconstruction, a predefined subset of electrodes is selected. The target channels are defined as $\mathcal{T} = \{\mathrm{Fp1}, \mathrm{Fp2}, \mathrm{AF3}, \mathrm{AF4}, \mathrm{F7}, \mathrm{F8}, \mathrm{T7}, \mathrm{T8}\}$, where each target index $j \in \mathcal{T}$ corresponds to electrodes selected based on spatial proximity under the international 10–20 montage.

Conditioning Channels. Using a nearest-neighbor spatial interpolation technique, missing central sensorimotor channels (C3, Cz, C4) can be approximated based on signals from adjacent electrodes, partially preserving activity in the μ (8–12 Hz) and β (13–30 Hz) frequency bands associated with MI tasks.

Forward Diffusion Process. For each target channel $j \in \mathcal{T}$, reconstruction is formulated as a conditional DDPM (9), where the forward process is applied to $x_0^{(j)}$. The forward diffusion process corrupts the channel defined as:

$$x_t^{(j)} = \sqrt{\bar{\alpha}_t}\, x_0^{(j)} + \sqrt{1 - \bar{\alpha}_t}\, \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \tag{2}$$

where $\bar{\alpha}_t = \prod_{s=1}^{t}(1 - \beta_s)$ and $\{\beta_s\}_{s=1}^{T_d}$ represents a predefined linear variance schedule (14).

Conditional U-Net Denoiser. To parameterize the reverse diffusion process, a conditional denoising network is trained to predict the injected Gaussian noise:

$$\epsilon_\theta(x_t^{(j)}, t, \mathbf{C}^{(j)}) \rightarrow \hat{\epsilon}_t^{(j)} \tag{3}$$

where the inputs correspond to the noisy target signal $x_t^{(j)}$, the diffusion timestep t, and the conditioning channels $\mathbf{C}^{(j)}$. The network parameters are optimized using a mean squared error loss defined as:

$$\mathcal{L}(\theta) = \mathbb{E}_{x_0^{(j)}, \epsilon, t}\left[\|\epsilon - \epsilon_\theta(x_t^{(j)}, t, \mathbf{C}^{(j)})\|^2\right]. \tag{4}$$

The learned denoiser outputs the predicted noise $\hat{\epsilon}_t^{(j)}$, which is later employed during the reverse process to reconstruct the target EEG signal. The denoising function is implemented as a one-dimensional conditional U-Net (21) for temporal EEG signals, composed of Conv1D residual blocks with batch normalization and ReLU activations. Sinusoidal time-step embeddings are incorporated into each residual block, while conditioning signals $\mathbf{C}^{(j)}$ are integrated through channel-wise concatenation at the input layer. The diffusion process is defined over $T_d = 1000$ steps with a linear noise schedule β_t ranging from 10^{-4} to 2×10^{-2}.

2.3. Conditional Channel Reconstruction

Gaussian Noise. During inference, reconstruction starts from a Gaussian noise sample $z \sim \mathcal{N}(0, \mathbf{I})$, which serves as the initial state of the reverse diffusion process.

Reverse Diffusion. The reverse process is defined as a conditional Markov chain that progressively removes noise from $x_t^{(j)}$:

$$p_\theta(x_{t-1}^{(j)} \mid x_t^{(j)}, \mathbf{C}^{(j)}). \tag{5}$$

Given the predicted noise $\hat{\epsilon}_t^{(j)}$ obtained in (3), the reverse update follows the standard DDPM formulation:

$$x_{t-1}^{(j)} = \frac{1}{\sqrt{\alpha_t}} \left(x_t^{(j)} - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}} \hat{\epsilon}_t^{(j)} \right) + \sigma_t z, \quad z \sim \mathcal{N}(0, \mathbf{I}). \tag{6}$$

The reverse diffusion chain iteratively reconstructed the target channel $\hat{x}_0^{(j)}$. This conditional reverse diffusion process enables controlled reconstruction of target EEG channels while preserving spatial dependencies with central sensorimotor regions.

Reconstructed Target Channel. Iterative denoising produces the reconstructed target signal $\hat{x}_0^{(j)}$, enabling recovery of peripheral EEG channels while preserving spatial-temporal dependencies.

2.4. Reconstruction and Decoding Evaluation

Hybrid Channel Replacement. As a result, the reconstruction procedure yields two aligned datasets: an original dataset and a hybrid dataset with reconstructed target channels. Both representations are used for MI decoding evaluation.

Evaluation. To evaluate the impact of diffusion-based channel reconstruction, this study tests three classical MI classifiers.

CSP + LDA. Common Spatial Pattern (CSP) is employed to extract discriminative spatial filters that maximize variance differences between MI classes. The resulting features are classified using Linear Discriminant Analysis (LDA).

FBCSP + LDA. Filter Bank Common Spatial Pattern (FBCSP) extends CSP by applying multiple band-pass filters within the MI frequency range before spatial filtering. The combined features are subsequently classified using LDA.

EEGNet. Implemented as a convolutional neural network designed for EEG decoding, its architecture comprises temporal convolutions that learn frequency-selective filters, spatial convolutions that capture channel-specific patterns, and separable convolutions that integrate temporal and spatial features.

3. Experiments and Results

3.1. Dataset and Data Processing

Experiments are performed on the BCI Competition III Dataset V, a continuous EEG benchmark collected from three participants using a 32-channel BioSemi system with a sampling frequency of 512 Hz. While the dataset includes three distinct mental tasks, this work focuses on left-hand and right-hand MI.

Following the evaluation protocol, the continuous recordings are segmented into overlapping 1 s windows with a stride of 0.5 s. After, the resulting segments are preprocessed via a fold-safe Independent Component Analysis (ICA) pipeline designed to avoid information leakage between train and test partitions. The ICA decomposition is computed solely on the training set after applying a temporary 1 Hz high-pass filter.

Table 1. Decoding performance on the independent held-out session using the DDPM framework.

Method	Dataset	Accuracy ($\pm$ std)	κ	F1-L	F1-R
CSP+LDA	Original	0.774 ± 0.096	0.551	0.754	0.787
FBCSP+LDA	Original	0.786 ± 0.086	0.568	0.754	0.808
EEGNet	Original	0.733 ± 0.095	0.470	0.687	0.762
CSP+LDA	Hybrid	0.768 ± 0.105	0.537	0.792	0.732
FBCSP+LDA	Hybrid	0.772 ± 0.087	0.549	0.784	0.759
EEGNet	Hybrid	0.721 ± 0.109	0.442	0.649	0.764

After ICA, signals are band-pass filtered between 8–30 Hz to highlight the sensorimotor μ and β bands. Lastly, z-score normalization is performed using statistics derived only from the training data to mitigate amplitude variability across recording sessions.

3.2. Experimental Setup and Results

Experimental Details. The first three training sessions are used for model development, while the fourth session is reserved as an independent held-out test set. The diffusion model is trained using the Adam optimizer with learning rate 1×10^{-4} and batch size of 64. Training is conducted for 200 epochs, and the model with lowest validation loss is selected. CSP employs 4 spatial components, with log-variance features classified using LDA with shrinkage. FBCSP uses the same number of spatial components across fixed MI bands (8–30 Hz), followed by LDA with shrinkage. EEGNet is trained for 200 epochs using Adam (learning rate 10^{-3}, weight decay 10^{-4}), batch size 128, dropout 0.5, and patience 80.

Experimental results. Reconstruction quality is quantified using the Pearson correlation coefficient (PCC) between reconstructed and reference EEG signals, averaged across trials. Decoding performance is evaluated using classification Accuracy and Cohen's kappa coefficient (κ) on the held-out session. We assess the impact of structured channel reconstruction on MI decoding performance using the mentioned MI classifiers.

As shown in Table 1, decoding performance remains stable under hybrid channel replacement across all classifiers. Particularly, FBCSP+LDA achieves the highest performance in both original and hybrid conditions, indicating that spectro-spatial representations are largely preserved after reconstruction. In addition, CSP+LDA exhibits comparable robustness with a moderate performance decrease. Finally, EEGNet shows slightly higher sensitivity to channel replacement, although differences remain small. Overall, the results indicate that diffusion-based structured channel reconstruction maintains discriminative information necessary for within-session MI decoding.

Overall, accuracy reductions are limited to approximately $\leq 1.5\%$, and remain small relative to the reported standard deviations, indicating minimal performance degradation under within-subject evaluation. Similar behavior is observed in κ, suggesting only marginal changes in class agreement. Analysis of F1-scores shows no systematic class-specific drop using hybrid reconstruction. Although minor redistributions in left and right-hand performances are observed, specifically, a slight increase in left-hand F1 score for CSP-based approaches, these variations remain balanced across classifiers.

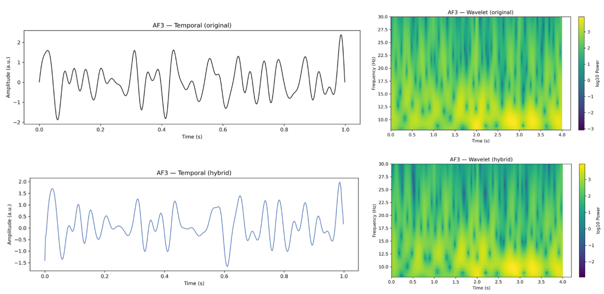

Figure 2. Temporal and time–frequency comparison between the original and reconstructed hybrid EEG signal of the AF3 channel (PCC = 0.935).

Qualitative comparisons in Fig. 2 further illustrate the visual similarity between the original and reconstructed hybrid signals for one target channel. From this example, it could be seen that the temporal shape of the reconstructed signals exhibit comparable amplitude variations and similar time–frequency characteristics to the original EEG signals.

3.3. Ablation Analysis

We conduct ablation studies to evaluate the impact of spatial conditioning design and diffusion formulation on the fidelity of EEG channel reconstruction. Reconstruction quality is measured using the PCC between reconstructed and original target channels.

Spatial Topology. As shown in Fig. 3(a), reconstruction performance strongly depends on spatial adjacency. The NEAR topology, defined as conditioning on neighboring adjacent channels according to the 10–20 montage, achieves the highest PCC (0.732). In contrast, FAR (0.321) and

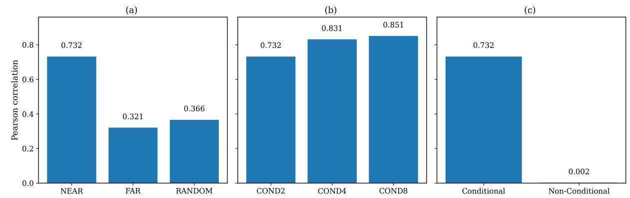

Figure 3. Ablation of Spatial Topology, Conditioning Size, and Diffusion Mode for Conditional EEG Channel Reconstruction.

RANDOM (0.366) configurations might exhibit substantial degradation. These results suggest that preserving physiologically meaningful local context is essential for maintaining channel-specific characteristics during reconstruction.

Conditioning Size. Fig. 3(b) examines the effect of conditioning-set size on reconstruction performance. PCC increases from 0.732 (COND2) to 0.831 (COND4) and 0.851 (COND8), indicating that richer spatial context improves structural recovery. However, the incremental gains decrease as more conditioning channels are added.

Diffusion Mode. Fig. 3(c) demonstrates the importance of conditional modeling. While the conditional framework achieves a PCC of 0.732, the non-conditional variant yields a PCC of 0.002, failing to preserve signal correspondence. This result reflects the limitation of modeling the marginal distribution without structured conditioning, highlighting the necessity of spatial context for accurate EEG reconstruction.

Overall, these ablations validate that spatially structured conditioning is fundamental to accurate EEG channel reconstruction and forms the basis for robust MI decoding.

4. Conclusion

This work proposes a conditional diffusion framework for channel reconstruction in MI decoding. The proposed approach reconstructs predefined target channels by conditioning the reverse diffusion process on spatially adjacent electrode observations, enabling targeted channel imputation while preserving inter-channel dependencies.

Experimental evaluation under an independent held-out protocol demonstrates that decoding performance remains stable across CSP+LDA, FBCSP+LDA, and EEGNet classifiers and κ exhibits consistent trends, indicating that hybrid reconstruction does not introduce systematic class degradation. Results suggest that the proposed conditional diffusion strategy preserves discriminative spatio-temporal structure essential for MI decoding, supporting its suitability for structured missing-channel scenarios without performance compromise.

References

[1] Singh, A., Hussain, AA., Lal, S., Guesgen, HW.: A comprehensive review on critical issues and possible solutions of motor imagery based electroencephalography brain-computer interface. *Sensors*. **21**(6), 2173 (2021). 10.3390/s21062173

[2] Zhu, T., Tang, H., Jiang, L., Li, Y., Li, S., Wu, Z.: A study of motor imagery EEG classification based on feature fusion and attentional mechanisms. *Frontiers in Human Neuroscience* **19**, 1611229 (2025). 10.3389/fnhum.2025.1611229

[3] Yang, B., Rong, F., Xie, Y., Li, D., Zhang, J., Li, F., Shi, G., Gao, X.: A multi-day and high-quality EEG dataset for motor imagery brain-computer interface. *Scientific Data* **12**(1), 488 (2025). 10.1038/s41597-025-04826-y

[4] Fahimi, F., Dosen, S., Ang, K.K., Mrachacz-Kersting, N., Guan, C.: Generative adversarial networks-based data augmentation for brain-computer interface. *IEEE Transactions on Neural Networks and Learning Systems* **32**(9), 4039-4051 (2021).

[5] Wang, X., Zhang, X., Zhao, K., Zhang, S.: GCDE: Graph-embedded conditional diffusion for EEG data augmentation. In: *2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 4636-4643 (2025).

[6] Shimizu, H., Srinivasan, R.: Improving classification and reconstruction of imagined images from EEG signals. In: *PLOS ONE* **17**(9), e0274847 (2022). 10.1371/journal.pone.0274847

[7] Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems* **27** (2014).

[8] Penava, S., Radivojević, A., Ščepanović, D., Milić, D., Milovanović, B.: A novel approach for EEG data augmentation using generative adversarial networks. *Neural Networks* **157**, 263-274 (2023).

[9] Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. In: *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS 2021)*, 1-15 (2021).

[10] Song, Y., Ermon, S.: Score-based generative modeling through stochastic differential equations. In: *International Conference on Learning Representations* (2020).

[11] Ho, J., Jain, A., Abbeel, P.: Denoising Diffusion Probabilistic Models. In: *Advances in Neural Information Processing Systems* (2020).

[12] Gómez-Morales, Ó.W., Collazos-Huertas, D.F., Álvarez-Meza, A.M., Castellanos-Dominguez, C.G.: EEG signal prediction for motor imagery classification in brain–computer interfaces. *Sensors* **25**(7), 2259 (2025). 10.3390/s25072259

[13] Habashi, A.G., Azab, A.M., Eldawlatly, S., Aly, G.M.: Generative adversarial networks in EEG analysis: an overview. In: *Journal of NeuroEngineering and Rehabilitation* **20**(1), 40 (2023).

[14] De Lima Alexandre, H., De Moraes Lima, C.A.: Synthetic EEG generation using diffusion model for motor imagery task. In: *Lecture Notes in Computer Science*, 516-530 (2026).

[15] Angulo-Sherman, I.N., León-Domínguez, U., Martinez-Torteya, A., Fragoso-González, G.A., Martínez-Pérez, M.V.: Proficiency in motor imagery is linked to the lateralization of focused ERD patterns and beta PDC. In: *Journal of NeuroEngineering and Rehabilitation* **22**(1), 30 (2025).

[16] Li, K., El-Fiqi H., Wang, M.: Gate control mechanisms of autoencoders for EEG signal reconstruction. *Sensors* **25**(11), 3389 (2025). 10.3390/s25113389

[17] Luo, T., Fan, Y., Chen, L., Guo, G., Zhou, C.: EEG signal reconstruction using a generative adversarial network with Wasserstein distance and temporal-spatial-frequency loss. *Front Neuroinform* **14**(15) (2020). 10.3389/fninf.2020.00015

[18] Blankertz, B., Tomioka, R., Lemm, S., Kawanabe, M., Müller, K.-R.: Optimizing spatial filters for robust EEG single-trial analysis. In: *IEEE Signal Processing Magazine* **25**(1), 41-56 (2008).

[19] Ang, K.K., Chin, Z.Y., Wang, C., Guan, C., Zhang, H.: Filter bank common spatial pattern algorithm on BCI Competition IV datasets 2a and 2b. In: *Frontiers in Neuroscience* **6**(1), 39 (2012).

[20] Lawhern, V. J., Solon, A. J., Waytowich, N. R., Gordon, S. M., Hung, C. P., Lance, B. J.: EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces. In: *Journal of Neural Engineering* **15**(5), 056013 (2018). 10.1088/1741-2552/aace8c

[21] Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted

Intervention – MICCAI 2015. *Lecture Notes in Computer Science* **9351**, 234–241 (2015). 10.1007/978-3-319-24574-4_28