ENABLING UNSUPERVISED TRAINING OF DEEP EEG DENOISERS WITH INTELLIGENT PARTITIONING

Anonymous authors
Paper under double-blind review

ABSTRACT

Denoising electroencephalogram (EEG) is an inherently challenging task, since neural activity is not only subtle but also inseparable from spectrally overlapping noise artifacts. Today, effective EEG denoising is more important than ever, given the rapid adoption of wearables across various applications. Deep learning methods have shown promising results in decomposition-free denoising that handles the time-varying pervasive EEG artifacts. However, training highly expressive neural networks requires artifact-free EEG, which is inherently unobtainable. To address this, we propose Intelligent Partitioning for Self-supervised Denoising (iPSD). Our method eliminates the need for clean references by learning to partition an input EEG segment into independent noisy realizations with the same underlying signal. This enables self-supervision of deep learning denoisers, even in zero-shot settings where only a single EEG segment to be denoised is available. We validate iPSD through extensive experiments, including validations on wearable EEG from in-ear sensors. The results show that iPSD achieves state-of-the-art performance, most notably under extremely low signal-to-noise ratios (down to -10 dB) and challenging artifacts (e.g., EMG), with spectral fidelity orders of magnitude higher than competitive baselines.

1 INTRODUCTION

Wearable technologies are transforming continuous health monitoring, with significant advances in electrocardiogram (ECG) (Tian et al., 2025) and photoplethysmogram (PPG) (Zhang et al., 2023) tracking. However, progress in wearable electroencephalogram (EEG) monitoring has been comparatively slow: on top of the poor electrode-skin contact, motion artifacts, and physiological interference inherent to wearable recordings, the brain signals themselves are also subtle, broadband, and non-stationary (Chu et al., 2021; Huhn et al., 2022; Rao et al., 2025).

A range of signal processing methods has been proposed to denoise EEG through decomposition and selective reconstruction. One line of work decomposes the signal using wavelet bases and reconstructs a denoised version by thresholding wavelet coefficients (Krishnaveni et al., 2006; Alyasseri et al., 2019). However, this approach struggles when the noise spectrum overlaps with that of the EEG, as is the case for electromyographic (EMG) artifacts from muscle contractions, a prevalent noise source in wearable recordings. Another approach partially alleviates this by decomposing the signal into data-driven mode functions (Huang et al., 1998; Dragomiretskiy & Zosso, 2013). However, they are sensitive to hyperparameter choices and abrupt perturbations common in wearable EEG.

Deep learning (DL) offers a compelling alternative to decomposition-based denoising, as highly expressive neural networks can directly map a noisy EEG to its clean counterpart without any fixed or heuristic basis (Yang et al., 2016; Sun et al., 2020; Jurczak et al., 2022; Chuang et al., 2022). However, supervised training of DL denoisers requires clean reference recordings, which are unattainable since neural signals are inseparable from physiological and environmental noise at the measurement point. Self-supervision, using only independent noisy realizations, has proven effective for denoising images with the hugely successful Noise2Noise (N2N) (Lehtinen et al., 2018) method. N2N learns to restore the clean image by training a neural network to map one noisy realization of the image to another. Mansour & Heckel (2023) later extended N2N to settings without multiple measurements of the same image via interleaved downsampling, which assigns alternating pixels to two sub-images to form a pair of noisy realizations. This heuristic, however, does not straightforwardly transfer to EEG: unlike

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

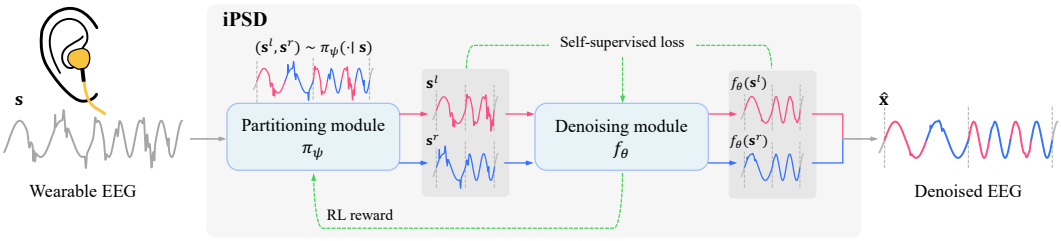

Figure 1: Schematic overview of iPSD. The partitioning module π_ψ splits the input signal into two sub-signals, $\mathbf{s}^l$ and $\mathbf{s}^r$, which are denoised by the denoising module f_θ. Training of f_θ uses the self-supervised N2N loss, $\left\|f_\theta(\mathbf{s}^l) - \mathbf{s}^r\right\|_2^2 + \left\|f_\theta(\mathbf{s}^r) - \mathbf{s}^l\right\|_2^2$. The negative of the converged f_θ loss is then used as a reward to update π_ψ. We alternate the updates of π_ψ and f_θ until convergence.

on locally smooth images, EEG's rapid oscillatory structure causes interleaved downsampling to yield realizations with distinct underlying clean signals, thus violating N2N's core assumption.

To restore clean EEG signals without any reference, we propose Intelligent Partitioning for Self-supervised Denoising (iPSD). The core idea is to learn to assign samples of a noisy input into two sub-signals with the same underlying clean signal. (Rather than assignment via a predetermined rule, such as in interleaved downsampling, the intelligent partitioning enables learning signal-specific assignment.) The two sub-signals then act as a pair of noisy realizations for N2N-style self-supervised denoising. The intuition is that, for two sub-signals sharing the same clean signal but carrying independent noise, the noise in one is unpredictable from the other; so a denoiser network trained to map one to the other is driven to converge to output their shared component, i.e., the clean signal. In iPSD, the partitioning is optimized via reinforcement learning (RL), jointly with the denoiser, using the negative of the converged denoising loss as the reward. We show in Section 2.2 that, under reasonable assumptions, this drives the partitioning module toward sub-signals that are maximally informative of one another's underlying signal. We further develop a zero-shot variant, named iPSD-Zero, that recovers the clean signal from a single noisy segment without any prior training.

We evaluate iPSD on EEG signals corrupted by White Gaussian Noise (WGN) and real-world EMG artifacts—across noise levels up to ten times the signal power—as well as real noisy wearable EEG data from custom-built in-ear sensors. Across all conditions, iPSD consistently outperforms state-of-the-art methods with spectral fidelity orders of magnitude higher than established baselines. On a downstream sleep-stage classification task, wearable EEG denoised by iPSD achieves accuracy comparable to clinical-grade scalp EEG, underscoring the method's practical value for accessible health monitoring. To the best of our knowledge, iPSD is the first self-supervised EEG denoiser to demonstrate such effectiveness across noise types, SNR levels, and real-world wearable recordings.

2 Methodology

Let $\mathbf{x} \in \mathbb{R}^L$ denote a clean EEG signal of length L, and let $\mathbf{s} = \mathbf{x} + \mathbf{n}$ be a noise-corrupted measurement of $\mathbf{x}$, where $\mathbf{n} \in \mathbb{R}^L$ denotes additive noise that is independent of $\mathbf{x}$, i.e., $\mathbb{E}\left[\mathbf{n}^\top \mathbf{x}\right] = 0$. Denoising then refers to recovering the underlying clean signal $\mathbf{x}$ from its noisy measurement $\mathbf{s}$. We focus on deep-learning denoisers, which employ expressive neural networks to directly map the noisy $\mathbf{s}$ to its clean estimate $\hat{\mathbf{x}}$, bypassing the fixed or heuristic bases of classical signal processing.

2.1 iPSD formulation

In most practical scenarios, the clean reference $\mathbf{x}$ is unavailable, making supervised training of the denoiser infeasible. Unlike images or audio, where clean reference signals can be synthetically generated, EEG recordings are inherently contaminated. This is because the true neural signal $\mathbf{x}$ is inseparable from physiological and environmental noise at the point of acquisition, no matter how precise the measurement device. The design of iPSD is optimized for this fundamentally constrained setting, where the true target is physically inaccessible rather than merely scarce.

The schematic in Fig. 1 depicts the architecture of iPSD, which consists of two learnable modules: the *partitioning module* takes a noisy input $\mathbf{s}$ and uses a neural network π_ψ parametrized by ψ to generate a pair of noisy sub-signals $(\mathbf{s}^l, \mathbf{s}^r)$; and the *denoising module* uses a neural network f_θ parametrized by θ to map a noisy signal to its clean estimate. The parameters of the two neural networks are jointly optimized via the following self-supervised objective

$$\theta^\star, \psi^\star = \arg\min_{\theta,\psi} \mathbb{E}_{(\mathbf{s}^l,\mathbf{s}^r)\sim\pi_\psi(\cdot|\mathbf{s})} \left[\left\| f_\theta(\mathbf{s}^l) - \mathbf{s}^r \right\|_2^2 + \left\| f_\theta(\mathbf{s}^r) - \mathbf{s}^l \right\|_2^2 \right]. \tag{1}$$

Self-supervised denoising The objective in Eq. (1) trains the denoising model f_θ using only the noisy sub-signals $(\mathbf{s}^l, \mathbf{s}^r)$ with no clean reference required. This self-supervision is valid when the two sub-signals share the same underlying clean signal: the independent noise components in $\mathbf{s}^l$ and $\mathbf{s}^r$ are each unpredictable from the other, so f_θ trained to map one sub-signal to the other is driven to output their shared component, i.e., the clean signal. The theorem below, adapted from Mansour & Heckel (2023), formalizes this:

Theorem 1. *Suppose* $(\mathbf{s}^l, \mathbf{s}^r) \sim \pi_\psi(\cdot \mid \mathbf{s})$ *are independent noisy realizations of the same underlying signal* $\mathbf{x}$*. Then the* $\theta^\star$ *that optimizes the self-supervised loss in Eq.* (1) *also minimizes the expected* $L2$ *distance between the network outputs and the ground-truth* $\mathbf{x}$*, i.e.,*

$$\theta^\star = \arg\min_{\theta} \mathbb{E}\Big[\left\| f_\theta(\mathbf{s}^l) - \mathbf{x} \right\|_2^2 + \left\| f_\theta(\mathbf{s}^r) - \mathbf{x} \right\|_2^2 \Big].$$

The proof of Theorem 1 is provided in Appendix A. This theorem shows that, given the partition $(\mathbf{s}^l, \mathbf{s}^r)$, f_θ can be optimized via gradient descent on the loss

$$\mathcal{L}(\theta) = \left\| f_\theta(\mathbf{s}^l) - \mathbf{s}^r \right\|_2^2 + \left\| f_\theta(\mathbf{s}^r) - \mathbf{s}^l \right\|_2^2. \tag{2}$$

Crucially, the validity of this self-supervised training scheme hinges on the two sub-signals $(\mathbf{s}^l, \mathbf{s}^r)$ sharing the same underlying clean signal. Achieving this is particularly challenging for EEG, whose non-stationarity and fast oscillatory structure demand a flexible, targeted partitioning strategy. This motivates the learnable partitioning module at the core of iPSD.

Intelligent partitioning Here, we describe the partitioning module that obtains the pair $(\mathbf{s}^l, \mathbf{s}^r)$ from the input noisy signal $\mathbf{s}$ (see Fig. 2). Two ideas are combined here: (i) the input is segmented into small windows where the underlying signal is approximately stationary, and (ii) the model π_ψ learns how to best partition each window to optimize the denoising performance by stochastically exploring all candidate partitions.

Formally, the input $\mathbf{s} \in \mathbb{R}^L$, with index set $I = \{1, \cdots, L\}$, is segmented into non-overlapping local windows of an appropriate even length W (that divides L). The index set of the k-th window is $I^{(k)} = \{(k-1)W+1, \cdots, kW\}$, with $k = 1, \cdots, L/W$. The partitioning model π_ψ learns to extract from each $I^{(k)}$ (corresponding to window k), a subset $I^{(k),l}$, constrained to cardinality $W/2$ (and thus its complement $I^{(k),r} = I^{(k)} \setminus I^{(k),l}$). We use $(\mathbf{s}^l, \mathbf{s}^r) \sim \pi_\psi(\cdot \mid \mathbf{s})$ to denote the extraction of the signal pair $(\mathbf{s}^l, \mathbf{s}^r)$ according to the stochastic policy π_ψ: $\mathbb{R}^L \to \mathcal{P}(I^l)$ by

$$\mathbf{s}^l = [\mathbf{s}_i]_{i\in I^l} \in \mathbb{R}^{L/2}, \qquad \mathbf{s}^r = [\mathbf{s}_i]_{i\in I^r} \in \mathbb{R}^{L/2}, \tag{3}$$

where $I^l = \bigcup_{k=1}^{L/W} I^{(k),l}$ and $I^r = \bigcup_{k=1}^{L/W} I^{(k),r}$.

Proposition 1. *For any signal window* $I^{(k)}$*, the partitioning strategy used in iPSD can always achieve a partition that is at least as good as interleaved downsampling (Mansour & Heckel, 2023), in terms of the mismatch between the underlying clean components of the sub-signals* $(\mathbf{s}_{I^{(k),l}}, \mathbf{s}_{I^{(k),r}})$.

We provide the proof of Proposition 1 in Appendix B.

2.2 Training iPSD

The combinatorial nature of the partitioning policy π_ψ renders the objective in Eq. (1) non-differentiable with respect to the parameters ψ. To address this, iPSD uses an alternating optimization scheme, where f_θ is optimized via gradient descent and π_ψ via RL.

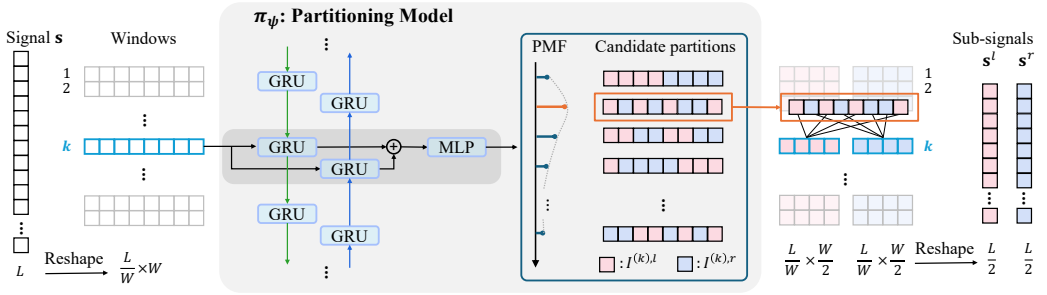

Figure 2: Workflow of the partitioning module. The input signal $\mathbf{s}$ of length L is reshaped into L/W non-overlapping windows of length W, each admitting $\binom{W}{W/2}/2$ candidate partitions into two equal-length subsets. The policy π_ψ selects a partition for each window, and the corresponding subsets are concatenated across windows to form the sub-signal pair. The denoising module f_θ then processes the sub-signals and provides feedback used to optimize π_ψ (see Fig. 1).

Algorithm 1. Training iPSD

Initialization: f_θ, π_ψ, $\mathcal{S}$; learning rates η^θ, η^ψ; $N = 0$.
while not converged **do**
  Sample $\mathbf{s} \in \mathcal{S}$ and partition $\mathbf{s}$ into $(\mathbf{s}^l, \mathbf{s}^r) \sim \pi_\psi(\cdot \mid \mathbf{s})$
  while θ not converged **do**
    $J_n^d = \left\| f_\theta(\mathbf{s}^l) - \mathbf{s}^r \right\|_2^2 + \left\| f_\theta(\mathbf{s}^r) - \mathbf{s}^l \right\|_2^2$
    $\theta \leftarrow \theta - \eta^\theta \nabla_\theta J_n^d$; $N = N + 1$
  end while
  $\psi \leftarrow \psi + \eta^\psi \nabla_\psi \log \pi_\psi(\mathbf{s}^l, \mathbf{s}^r \mid \mathbf{s})\, R(\mathbf{s}^l, \mathbf{s}^r)$, $\quad R(\mathbf{s}^l, \mathbf{s}^r) = -\frac{1}{10} \sum_{n=N-9}^{N} J_n^d$
end while

Let $\mathcal{S}$ denote a set of noisy training signals. In each iPSD training iteration, we sample $\mathbf{s} \in \mathcal{S}$ and use π_ψ to partition it into $(\mathbf{s}^l, \mathbf{s}^r)$. We then update f_θ via gradient descent on the loss in Eq. (2) until convergence, defined as the variance of loss values in the last 10 steps falling below 10^{-6}. The negative of the converged denoising loss serves as the reward signal for the RL-based optimization of π_ψ. To reduce reward variance, we apply an averaging scheme. Suppose θ converges at the Nth step, and let J_n^d denote the value of $\mathcal{L}(\theta)$ at step n. The reward $R(\mathbf{s}^l, \mathbf{s}^r)$ is then defined as:

$$R(\mathbf{s}^l, \mathbf{s}^r) \;=\; -\frac{1}{10} \sum_{n=N-9}^{N} J_n^d \,. \tag{4}$$

In Appendix C, we show that under reasonable assumptions, this RL reward formulation drives π_ψ toward $\mathbf{s}^l$ and $\mathbf{s}^r$ that are maximally informative of each other's underlying signal.

Using the scalar reward $R(\mathbf{s}^l, \mathbf{s}^r)$ enables gradient-based optimization of the partitioning policy π_ψ. To find π_ψ that maximizes the reward, we apply the policy gradient theorem (see Appendix D), which shows that the gradient of the expected reward with respect to ψ can be written as

$$\nabla_\psi \, \mathbb{E}_{(\mathbf{s}^l, \mathbf{s}^r)} \left[R(\mathbf{s}^l, \mathbf{s}^r) \right] = \mathbb{E}_{(\mathbf{s}^l, \mathbf{s}^r)} \left[\nabla_\psi \log \pi_\psi(\mathbf{s}^l, \mathbf{s}^r \mid \mathbf{s})\, R(\mathbf{s}^l, \mathbf{s}^r) \right] . \tag{5}$$

We therefore update π_ψ via gradient ascent on the right-hand side of Eq. (5). Algorithm 1 summarizes the full self-supervised denoising procedure. In practice, we employ Proximal Policy Optimization (PPO) (Schulman et al., 2017) and sample a batch of partitions in parallel to further reduce the variance in the observed $R(\mathbf{s}^l, \mathbf{s}^r)$, ensuring stable and efficient policy updates.

Zero-shot setting As an extension, we develop iPSD-Zero, a variant of iPSD designed for the challenging *zero-shot* setting, where a clean signal is recovered directly from a single noisy recording without any prior training. This capability is particularly valuable for wearable EEG, where instant deployment is desirable, yet large variabilities in individual physiology, noise characteristics, and hardware make it impractical to form a representative training set before deployment. One significant

challenge when only the test signal $\mathbf{s}$ is available is to rapidly identify the optimal partition from a large search space. We address this by enforcing a shared partitioning strategy across all localized windows, reducing the search space to $\binom{W}{W/2}/2$ candidates. Selecting the optimal partition can be naturally formulated as a multi-armed bandit problem, where each candidate partition corresponds to an *arm*. Pulling each arm yields a reward, defined as in Eq. (4), which is stochastic due to the inherent randomness of training f_θ. The objective of iPSD-Zero is to identify the best arm, measured by the expected reward, with high probability using as few trials as possible. This is a canonical pure-exploration objective, for which we employ lil'UCB (Jamieson et al., 2014), whose sample complexity matches the theoretical lower bound up to constant factors. By iteratively pulling different arms, iPSD-Zero gradually builds empirical estimates of each arm's expected reward and eventually identifies the arm whose corresponding partition of $\mathbf{s}$ best allows f_θ to recover the clean signal. Further details are provided in Appendix E.

3 Experiments

We begin with some notes on the implementation of iPSD. For the partitioning module, we use a window length of 8 samples, chosen by ablation (see details in Appendix F) to achieve sufficient partitioning flexibility while keeping the search space tractable. The policy network π_ψ uses a bidirectional GRU architecture, which outperforms alternatives including MLP, U-Net, and RNN (see Appendix F for detailed comparisons). We train π_ψ via PPO using the Adam optimizer with a learning rate of 10^{-4}, a batch size of 64, and a total of $20{,}000$ update steps; other hyperparameters follow Huang et al. (2022). We observe stable convergence of iPSD across independent runs with different random seeds (see results in Appendix I.2).

For the denoising module f_θ, we use the loss function $\mathcal{L}(\theta)$ in Eq. (2) together with two regularization terms proposed by Mansour & Heckel (2023). Given a partition $(\mathbf{s}^l, \mathbf{s}^r) \sim \pi_\psi(\cdot \mid \mathbf{s})$, we define functions $\pi^l, \pi^r : \mathbb{R}^L \to \mathbb{R}^{L/2}$ such that $\mathbf{s}^l = \pi^l(\mathbf{s})$ and $\mathbf{s}^r = \pi^r(\mathbf{s})$. The loss function is

$$\mathcal{L}_{\text{full}}(\theta) = \mathcal{L}(\theta) + \|f_\theta(\pi^l(\mathbf{s})) - \pi^l(f_\theta(\mathbf{s}))\|_2^2 + \|f_\theta(\pi^r(\mathbf{s})) - \pi^r(f_\theta(\mathbf{s}))\|_2^2, \tag{6}$$

where the two regularization terms encourage iPSD's output to obey a consistency property that the underlying clean signal should satisfy: partitioning and then denoising should yield the same result as denoising and then partitioning. We implement f_θ as a lightweight fully convolutional network with three layers and LeakyReLU activations, totaling approximately 7k parameters. The channel dimensions progress as $1 \to 48 \to 48 \to 1$, with all layers using a kernel size of 3 and padding of 1. The compact architecture of f_θ leads to fast convergence (in around 200 epochs) and permits parallel training of multiple instances, allowing the partitioning policy π_ψ to quickly explore diverse strategies and identify high-quality partitions.

Experiments on synthetic data We benchmark iPSD against state-of-the-art baselines including wavelet-based (Alyasseri et al., 2019; Houamed et al., 2020), mode decomposition-based (Dora & Biswal, 2020; Colominas et al., 2014; Ranjan et al., 2022), hybrid (Kerechanin & Bobrov, 2022), and self-supervised DL (Chen et al., 2025) approaches (see a summary in Appendix G). For quantitative evaluation, we use three complementary metrics standard in signal denoising: signal-to-noise ratio (SNR), peak signal-to-noise ratio (PSNR), and the mean squared error of the power spectrum (Spectral MSE). SNR provides a global measure of time-domain fidelity, PSNR emphasizes robustness to large transient spikes, and Spectral MSE assesses spectral fidelity, which is critical because EEG analysis depends heavily on spectral features such as relative band power, peak frequency, and spectral entropy. Together, these metrics provide a balanced evaluation covering both the temporal and spectral domains. The definitions of all metrics are provided in Appendix H. Note that the clean signal is solely used for metric computation and is not accessible to any denoising method.

We use EEG signals from the CHB-MIT Scalp EEG Database (Goldberger et al., 2010; Shoeb, 2009), which was collected at the Children's Hospital Boston and contains over 900 h long-term scalp EEG recordings from 22 pediatric patients (5 males aged 3–22 years; 17 females aged 1.5–19) with intractable seizures. All signals are sampled at 256 Hz with 16-bit resolution, using the international 10–20 electrode placement system with a bipolar montage. To generate synthetic data, we first split all recordings into 10-second segments, which are long enough for iPSD and the baselines to capture meaningful neural rhythms (Rosiani et al., 2021). To create contaminated EEG signals,

Table 1: Comparison of denoising performance.

	WGN						EMG					
	−5 dB Input SNR			0 dB Input SNR			−5 dB Input SNR			0 dB Input SNR		
Method	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE
mVMD	−5.05	9.05	$1.46_{\times 10^4}$	2.18	16.28	$1.13_{\times 10^4}$	−6.05	8.05	$5.40_{\times 10^4}$	0.93	15.03	$3.06_{\times 10^4}$
CEEMDAN	−3.85	10.25	$0.98_{\times 10^4}$	4.13	18.23	$0.78_{\times 10^4}$	0.25	14.35	$0.36_{\times 10^3}$	3.26	17.36	$0.22_{\times 10^3}$
FrWT	−3.54	10.56	$0.96_{\times 10^4}$	3.43	17.53	$0.83_{\times 10^4}$	−0.06	14.04	$0.55_{\times 10^3}$	3.53	17.63	$0.62_{\times 10^3}$
EMD-LoG	2.28	16.38	$0.80_{\times 10^4}$	4.40	18.50	$0.80_{\times 10^4}$	0.32	14.42	$0.38_{\times 10^3}$	3.55	17.65	$0.22_{\times 10^3}$
WPT-ICA	−4.89	9.21	$1.63_{\times 10^4}$	1.74	15.84	$1.23_{\times 10^4}$	−0.09	14.01	$0.66_{\times 10^3}$	2.36	16.46	$0.31_{\times 10^3}$
Chen et al.	2.01	16.11	$0.28_{\times 10^4}$	3.25	17.35	$0.21_{\times 10^4}$	−0.05	14.05	$0.42_{\times 10^3}$	3.01	17.12	$0.19_{\times 10^3}$
Optimal-WT	3.34	17.44	$0.36_{\times 10^4}$	5.07	19.18	$0.27_{\times 10^4}$	0.33	14.42	$0.69_{\times 10^3}$	3.89	18.10	$0.43_{\times 10^3}$
iPSD	3.97	18.07	55.14	5.78	19.90	41.13	3.63	17.73	53.04	5.95	20.05	39.21
iPSD-Zero	3.89	18.00	58.34	5.76	19.86	48.67	3.27	17.37	56.62	5.31	19.41	44.12

* Pink (blue) highlights the best (second-best) performance; S-MSE stands for spectral MSE.

we introduce WGN and real-world EMG artifacts. The WGN is generated by sampling from a Gaussian distribution, while the EMG artifacts are from Rantanen et al. (2016), recorded on 15 healthy participants (8 female, 7 male; aged 26–57) performing facial movements such as chewing, smiling, lip puckering, and frowning. Since the EMG artifacts were recorded at a different sampling rate, we downsample them to 256 Hz before adding them to the EEG. The synthetic dataset was divided into training and test sets with an 80/20 split.

Table 1 compares the denoising performance of iPSD and its zero-shot variant, iPSD-Zero, against the baselines at input SNR levels of -5 dB and 0 dB. (The zero-shot iPSD is optimized directly on each test signal.) The proposed iPSD consistently and significantly outperforms all baselines across all evaluation metrics, under both WGN and EMG contamination. Compared to the strongest baseline, Optimal-WT, iPSD consistently achieves better denoising performance, with up to 3.3 dB higher output SNR, corresponding to more than 53% lower residual noise power after denoising. We also observe that iPSD is more robust under low-SNR conditions (at -5 dB, signal power is only approximately 32% of noise power), where the baselines show limited denoising capability, with several outputting negative SNR values. Most notably, while the spectral MSEs of the baselines are on the order of thousands under WGN and hundreds under EMG noise, the spectral MSE of iPSD remains consistently below 60. To provide an intuitive view of the strong spectral fidelity of iPSD, we additionally present visual comparisons of iPSD against Optimal-WT in Appendix J.

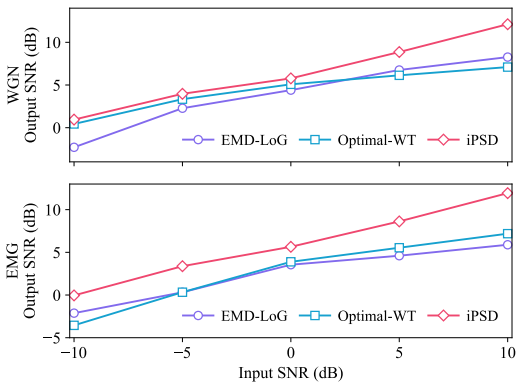

Figure 3: Performance over different noise levels.

Figure 3 compares the output SNR of iPSD against two representative baselines across input SNR levels ranging from the critically low -10 dB to 10 dB. The baselines are Optimal-WT from the wavelet-based family and EMD-LoG from the mode decomposition-based family. Across all noise levels, iPSD outperforms both baselines. We also observe that the baselines' output SNR gradually plateau as input SNR increased, whereas iPSD maintains a steady increase throughout the range.

These results highlight the benefit of the learning-based iPSD over fixed or heuristic partitioning strategies. Performance of the baselines varies across noise types: mode decomposition-based methods (CEEMDAN and EMD-LoG) are less effective for WGN, due to the highly unstructured temporal form, whereas wavelet-based methods (Optimal-WT and FrWT) are less effective for EMG, due to the substantial spectral overlap with EEG activity. The iPSD method, in contrast, maintains robust performance under these challenging conditions. We also note that iPSD-Zero

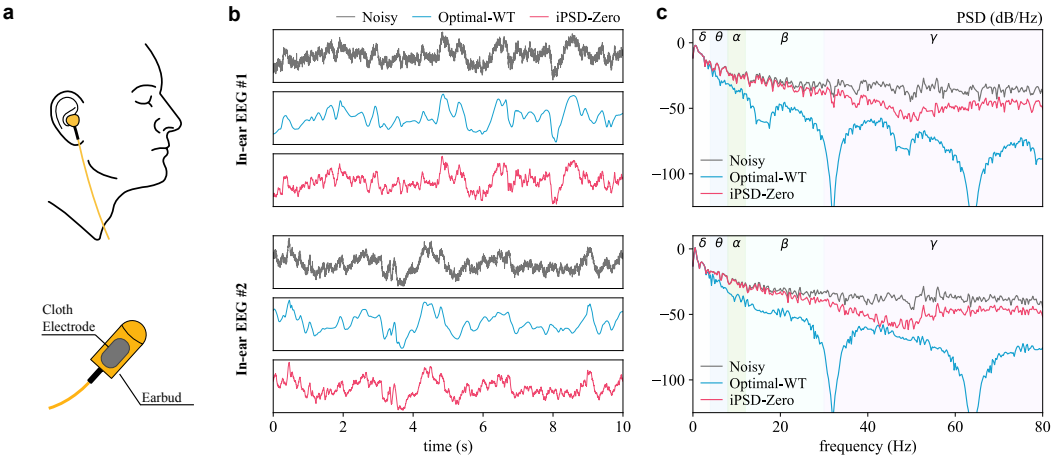

Figure 4: Real-world in-ear EEG denoising with iPSD. (**a**) Wearable EEG acquisition setup using the in-ear sensor, which integrates a viscoelastic earbud with a stretchable soft-cloth electrode. (**b**) Noisy in-ear EEG (gray) and the denoised output from Optimal-WT (blue) and iPSD-Zero (pink). (**c**) PSDs of the noisy in-ear EEG (gray), the Optimal-WT output (blue), and the iPSD-Zero output (pink). The key frequency bands for EEG analysis: δ (0.5–4 Hz), θ (4–8 Hz), α (8–12 Hz), β (12–30 Hz), and γ (30–80 Hz) are color-coded in the background.

achieves competitive performance, outperforming all baseline methods and ranking only second to the original iPSD. This slight performance degradation comes with a huge saving in computational cost. Starting from scratch, iPSD-Zero denoises a 10-second EEG signal in 5 seconds, making it suitable for real-time and resource-constrained scenarios where only the incoming signal is accessible. Additionally, while the independence assumption allows for theoretical proof of iPSD's convergence, yet in practice, iPSD can recover the clean signal well without the assumption holding strictly. In fact, iPSD's gain over baselines is most significant on the temporally correlated, real-world EMG noise, which violates the assumption. In Appendix I.6, we further quantify iPSD's robustness under different levels of assumption violation using autoregressive (AR) noise with different correlation levels, where iPSD retains competitive performance under markedly correlated noise.

Experiments on real-world wearable data The proposed iPSD method is also evaluated on real-world EEG recordings collected using our custom-built wearable sensors. The dataset comprised 684 hours of recordings from 37 participants (aged 65–83 years; 17 females, 20 males), sampled at 256 Hz with 16-bit resolution and segmented into 10-second windows. Figure 4a illustrates the in-ear sensor used for data collection. To maintain stable skin contact and minimize motion artifacts, the in-ear sensor uses a viscoelastic foam earbud with a stretchable, low-impedance cloth electrode on its outer surface. We also applied conductive gel to the electrode before inserting the sensor into the ear canal to further improve the skin-electrode contact.

Figures 4b and 4c show noisy in-ear EEG recordings, together with the denoising outputs of Optimal-WT and iPSD-Zero. While both iPSD and iPSD-Zero are applicable here, we only show iPSD-Zero in the figure: it operates under the more constrained zero-shot setting, yet produces outputs that are visually indistinguishable from those of iPSD. Figure 4b shows time-domain waveforms. The Optimal-WT outputs, although visually smooth, substantially corrupt the original EEG signal. For example, the spindle-like waveform between 7–8.5 s in the lower panel of Fig. 4b (In-ear EEG #2), characteristic of non-REM sleep, is corrupted by Optimal-WT but preserved by iPSD-Zero. The corresponding power spectral densities (PSDs) in Fig. 4c further support that Optimal-WT is oversmoothing: it exhibits spurious dips at multiple frequencies, indicating that genuine spectral content is being corrupted. In contrast, iPSD-Zero effectively suppresses the spiky, fluctuating noise artifacts while preserving the underlying EEG. Since the in-ear EEG was recorded during sleep, its spectral power is expected to concentrate below 30 Hz, with the elevated high-frequency power reflecting noise contamination. As shown, iPSD-Zero suppresses the high-frequency noise while preserving the low-frequency EEG components. This visual comparison is consistent with the quantitative results in Table 1, where iPSD substantially outperforms the baselines in spectral fidelity.

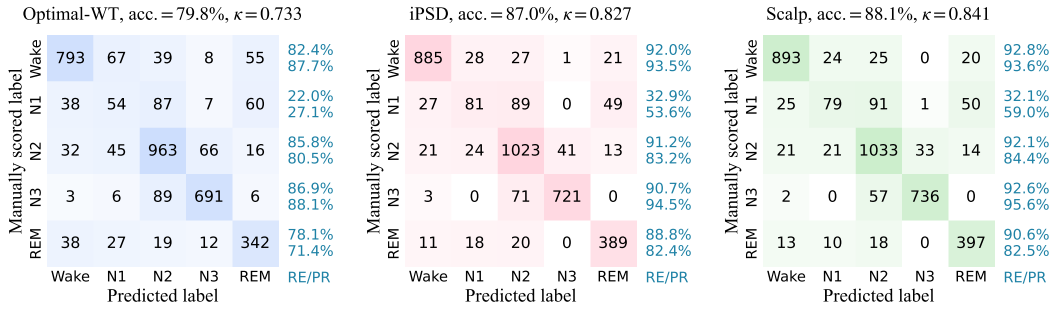

Figure 5: Confusion matrices for sleep-stage classification using the output of the baseline Optimal-WT and the proposed iPSD. Classification results using concurrently recorded high-quality scalp EEG are also included as a reference performance target. The recall and precision (RE/PR) for each sleep stage are shown on the right of each row.

Table 2: Ablating iPSD partitioning by replacing π_ψ with ID and DR.

	WGN						EMG					
	−5 dB Input SNR			0 dB Input SNR			−5 dB Input SNR			0 dB Input SNR		
Method	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE
ID	2.58	16.68	62.33	5.13	19.24	67.18	1.62	15.73	67.68	4.31	18.41	60.52
DR	2.74	16.84	59.43	5.30	19.41	53.62	2.20	16.31	63.88	4.73	18.83	52.27
iPSD	3.97	18.07	55.14	5.78	19.90	41.13	3.63	17.73	53.04	5.95	20.05	39.21
iPSD-Zero	3.89	18.00	58.34	5.76	19.86	48.67	3.27	17.37	56.62	5.31	19.41	44.12

* Pink (blue) highlights the best (second-best) performance; S-MSE stands for spectral MSE.

Because the real-world in-ear recordings lack clean references, we perform quantitative evaluation by performing sleep stage classification, a cornerstone task in automatic sleep monitoring. The classification is performed over 3,563 in-ear EEG segments, each annotated by experts into one of five sleep stages: Wake, N1, N2, N3, and REM. We denoise the segments using iPSD and the baselines in Appendix G. For each method, an XGBoost classifier is trained using a standard set of features (Vallat & Walker, 2021) extracted from the denoising output. Figure 5 shows the confusion matrices of (i) Optimal-WT, the best-performing baseline, (ii) iPSD, and (iii) classification based on features extracted from concurrently recorded scalp EEG (C3-M2). The scalp EEG, recorded with the clinical-grade SomnoHD system, is of high quality and serves as a performance reference that the denoising methods aim to match. We observe that classification based on iPSD substantially outperforms that based on Optimal-WT (accuracy 87.0%, Cohen's $\kappa = 0.827$ vs. 79.8%, $\kappa = 0.733$), with the largest gains in sleep stages characterized by higher-frequency activity (Wake, N1, and REM). This pattern aligns with Fig. 4c, where iPSD shows superior spectral fidelity above 4 Hz. N1 is consistently the most difficult stage to classify, but iPSD's recall and precision substantially exceed those of Optimal-WT and match the scalp EEG, indicating that the difficulty is inherent to N1 classification (Lee et al., 2022) with EEG alone rather than a limitation of denoising. Overall, the per-class metrics of iPSD-denoised in-ear EEG fall within a few points of scalp EEG across all stages, demonstrating that iPSD recovers clinically meaningful information from the noisy wearable recordings at a level comparable to clinical-grade recordings.

Ablation study on intelligent partition We ablate the effect of the proposed intelligent partitioning module by replacing π_ψ with interleaved downsampling (ID) or differentiable relaxation (DR). ID assigns alternating samples to $(\mathbf{s}^l, \mathbf{s}^r)$ and was proven highly effective in self-supervised image denoising (Mansour & Heckel, 2023). DR uses the standard straight-through Gumbel Softmax, which we detail in Appendix I.3. Experiments are conducted on the same synthetic EEG signals with WGN and EMG contamination. Table 2 shows the denoising performance, where the intelligent partitioning module (in both iPSD and iPSD-Zero) outperforms ID and DR. The fixed ID partitioning performs the worst, as its local-smoothness assumption does not transfer well to rapidly oscillating

EEG signals, a limitation that is particularly pronounced for wearable EEG recordings with lower sampling rates (see Appendix I.4). DR is only about 3.5% faster than the RL-based partitioning in iPSD, but at a significant cost to performance (up to 1.33 dB lower SNR and 33% higher S-MSE).

A major motivation for the proposed intelligent partitioning module is its flexibility to produce sub-signals with high underlying signal consistency and noise independence, assumptions that underlie most, if not all, reference-free denoising methods (Lehtinen et al., 2018; Huang et al., 2021; Mansour & Heckel, 2023). We explicitly verify whether iPSD mitigates violations of these assumptions by measuring both properties in the sub-signals generated by ID, DR, iPSD, and iPSD-Zero (see Table 10 in Appendix I.3). Compared to ID and DR, iPSD partitioning consistently produces sub-signals that better satisfy the core assumptions of Theorem 1, reducing underlying signal mismatch by 42–54%, noise correlation under WGN by 34–43%, and noise correlation under real-world EMG artifacts by 78–82%. These results indicate that iPSD is more robust to the assumption violations that arise in real-world EEG recordings, and that the proposed intelligent partitioning is not merely an improvement to N2N but a key enabler of self-supervised denoising for EEG.

4 Related work

Conventional scalp EEG processing widely uses band-pass and notch filters, which have limited use for low-SNR wearable EEG due to its substantial spectral overlap with prominent artifacts (EMG and EOG) (Urigüen & Garcia-Zapirain, 2015). To handle the localized and nonstationary nature of these artifacts, wavelet-based methods decompose the EEG signal to attenuate or discard components that correlate weakly with a predefined set of wavelet bases before reconstructing the denoised signal (Krishnaveni et al., 2006). However, these methods are sensitive to the choice of wavelet basis, thresholding strategy, and decomposition depth (Grobbelaar et al., 2022). Mode decomposition methods overcome the restriction of fixed wavelet bases by decomposing signals into data-driven mode functions (Huang et al., 1998; Dragomiretskiy & Zosso, 2013), but they rely either on perturbation-sensitive local extrema and signal envelopes, or on hyperparameter-sensitive non-convex optimization. Despite the improvements by hybrid methods such as VMD-wavelet and WPT-ICA (Kaur et al., 2021; Kerechanin & Bobrov, 2022), these approaches remain heavily reliant on handcrafted bases and heuristic decomposition strategies that do not generalize well.

Recent advances in DL have opened a promising new avenue for EEG denoising. Highly expressive neural architectures allow direct mapping of noisy EEG signals to their clean counterparts, eliminating the need for handcrafted decomposition strategies (Zhang et al., 2021; Gao et al., 2022; Pu et al., 2022; Chen et al., 2023). However, these approaches require a large collection of paired clean-noisy EEG data for training, which is unobtainable in practice. Noise2Noise (N2N) (Lehtinen et al., 2018; Kashyap et al., 2021) offers a compelling self-supervised image denoising approach based on learning to map independent noisy realizations of a signal to one another. Yet N2N requires paired noisy measurements of the same signal, which is virtually impossible to obtain for EEG signals, as they are unique, unrepeatable observations. Subsequent works (Krull et al., 2019; Huang et al., 2021; Mansour & Heckel, 2023; Chen et al., 2025) seek to learn from a single noisy observation by relying on rule-based partitioning or masking strategies, which are not suited to non-stationary, fast-oscillating EEG signals. In contrast, iPSD learns to generate paired noisy realizations sharing the same underlying clean signal via flexible partitioning of a single measurement.

5 Conclusion

The proposed iPSD (and its zero-shot variant, iPSD-Zero) enables training of highly expressive deep learning EEG denoisers without clean ground-truth signals. Across extensive experiments on synthetic and real EEG data, iPSD and iPSD-Zero consistently outperformed state-of-the-art baselines with markedly higher spectral fidelity. The value of denoising is in empowering tasks that depend on high-quality signals. We demonstrated this in a sleep-stage classification task, where the application of iPSD enhanced classification accuracy up to the standards of scalp EEG. These results suggest that the methods we developed here can effectively narrow the gap between accessible wearable measurements and clinical-grade recordings. Moreover, iPSD-Zero's ability to denoise a single signal without prior training makes it deployable in real-time and on-device, bringing continuous, accessible neural monitoring and diagnosis closer to reality.

AI USE STATEMENT

AI tools were only used to improve the wording and check the grammar of individual sentences. All final text in the manuscript was reviewed and finalized by the authors, who take full responsibility for the content of this work.

ETHICS STATEMENT

This work analyzes data collected from human participants. The study was conducted in accordance with the Declaration of Helsinki and Good Clinical Practice guidelines and was approved by the relevant research ethics committee. All participants received a detailed explanation of the study procedures and provided written informed consent prior to enrollment. The study site and ethics approval reference are omitted for double-blind review and will be clearly stated upon acceptance.

REPRODUCIBILITY STATEMENT

Detailed descriptions of the proposed algorithms are given in Section 2 and 3, together with Appendices E and F. The complete implementation is submitted as a supplementary zip file, which is sufficient for training and validating the proposed method. The CHB-MIT and EMG datasets are publicly available (Goldberger et al., 2010; Shoeb, 2009; Rantanen et al., 2016); the in-ear EEG dataset cannot be released due to ethical restrictions on patient confidentiality.

REFERENCES

Zaid Abdi Alkareem Alyasseri, Ahamad Tajudin Khader, Mohammed Azmi Al-Betar, Ammar Kamal Abasi, and Sharif Naser Makhadmeh. EEG signals denoising using optimal wavelet transform hybridized with efficient metaheuristic methods. *IEEE Access*, 8:10584–10605, 2019.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

Junfu Chen, Dechang Pi, Xiaoyi Jiang, Yue Xu, Yang Chen, and Xixuan Wang. Denosieformer: A transformer-based approach for single-channel EEG artifact removal. *IEEE Transactions on Instrumentation and Measurement*, 73:1–16, 2023.

Wensheng Chen, Cong Yu, Zhenhua Zhao, Nan Zheng, Han Li, and Yurong Li. Self-supervised EEG denoising via dual-branch consistency learning with masked reconstruction. *Knowledge-Based Systems*, pp. 114703, 2025.

Lei Chu, Ling Pei, and Robert Qiu. Ahed: A heterogeneous-domain deep learning model for IoT-enabled smart health with few-labeled EEG data. *IEEE Internet of Things Journal*, 8(23): 16787–16800, 2021.

Chun-Hsiang Chuang, Kong-Yi Chang, Chih-Sheng Huang, and Tzyy-Ping Jung. IC-U-Net: a U-Net-based denoising autoencoder using mixtures of independent components for automatic EEG artifact removal. *NeuroImage*, 263:119586, 2022.

Marcelo A Colominas, Gastón Schlotthauer, and María E Torres. Improved complete ensemble EMD: A suitable tool for biomedical signal processing. *Biomedical Signal Processing and Control*, 14: 19–29, 2014.

Chinmayee Dora and Pradyut Kumar Biswal. Correlation-based ECG artifact correction from single channel EEG using modified variational mode decomposition. *Computer Methods and Programs in Biomedicine*, 183:105092, 2020.

Konstantin Dragomiretskiy and Dominique Zosso. Variational mode decomposition. *IEEE Transactions on Signal Processing*, 62(3):531–544, 2013.

Tengfei Gao, Dan Chen, Yunbo Tang, Zhekai Ming, and Xiaoli Li. EEG reconstruction with a dual-scale CNN-LSTM model for deep artifact removal. *IEEE Journal of Biomedical and Health Informatics*, 27(3):1283–1294, 2022.

Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. CHB-MIT Scalp EEG Database. *PhysioNet*, June 2010. doi: 10.13026/C2K01R. URL https://doi.org/10.13026/C2K01R. Version 1.0.0.

Maximilian Grobbelaar, Souvik Phadikar, Ebrahim Ghaderpour, Aaron F Struck, Nidul Sinha, Rajdeep Ghosh, and Md Zaved Iqubal Ahmed. A survey on denoising techniques of electroencephalogram signals using wavelet transform. *Signals*, 3(3):577–586, 2022.

Ibtissem Houamed, Lamir Saidi, and Fawzi Srairi. ECG signal denoising by fractional wavelet transform thresholding. *Research on Biomedical Engineering*, 36(3):349–360, 2020.

Norden E Huang, Zheng Shen, Steven R Long, Manli C Wu, Hsing H Shih, Quanan Zheng, Nai-Chyuan Yen, Chi Chao Tung, and Henry H Liu. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings: Mathematical, Physical and Engineering Sciences*, pp. 903–995, 1998.

Shengyi Huang, Rousslan Fernand Julien Dossa, Chang Ye, Jeff Braga, Dipam Chakraborty, Kinal Mehta, and João G.M. Araújo. CleanRL: High-quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18, 2022. URL http://jmlr.org/papers/v23/21-1342.html.

Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. Neighbor2Neighbor: Self-supervised denoising from single noisy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14781–14790, 2021.

Sophie Huhn, Miriam Axt, Hanns-Christian Gunga, Martina Anna Maggioni, Stephen Munga, David Obor, Ali Sié, Valentin Boudo, Aditi Bunker, Rainer Sauerborn, et al. The impact of wearable technologies in health research: Scoping review. *JMIR mHealth and uHealth*, 10(1):e34384, 2022.

Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil' UCB: An optimal exploration algorithm for multi-armed bandits. In *Proceedings of the 27th Annual Conference on Learning Theory*, volume 35, pp. 423–439, 2014.

Marcin Jurczak, Marcin Kołodziej, and Andrzej Majkowski. Implementation of a convolutional neural network for eye blink artifacts removal from the electroencephalography signal. *Frontiers in Neuroscience*, 16:782367, 2022.

Madhav Mahesh Kashyap, Anuj Tambwekar, Krishnamoorthy Manohara, and S. Natarajan. Speech denoising without clean training data: A Noise2Noise approach. In *Interspeech 2021*, pp. 2716–2720, 2021. doi: 10.21437/Interspeech.2021-1130.

Chamandeep Kaur, Amandeep Bisht, Preeti Singh, and Garima Joshi. EEG signal denoising using hybrid approach of variational mode decomposition and wavelets for depression. *Biomedical Signal Processing and Control*, 65:102337, 2021.

Yaroslav Kerechanin and Pavel Bobrov. EEG denoising using wavelet packet decomposition and independent component analysis. In *Proceedings of the IEEE Fourth International Conference Neurotechnologies and Neurointerfaces*, pp. 68–70, 2022.

V Krishnaveni, S Jayaraman, S Aravind, V Hariharasudhan, and K Ramadoss. Automatic identification and removal of ocular artifacts from EEG using wavelet transform. *Measurement Science Review*, 6(4):45–57, 2006.

Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2Void-learning denoising from single noisy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2129–2137, 2019.

Yun Ji Lee, Jae Yong Lee, Jae Hoon Cho, and Ji Ho Choi. Interrater reliability of sleep stage scoring: a meta-analysis. *Journal of Clinical Sleep Medicine*, 18(1):193–202, 2022.

Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning image restoration without clean data. In *Proceedings of the International Conference on Machine Learning*, pp. 2965–2974, 2018.

Youssef Mansour and Reinhard Heckel. Zero-shot Noise2Noise: Efficient image denoising without any data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14018–14027, 2023.

Xiaorong Pu, Peng Yi, Kecheng Chen, Zhaoqi Ma, Di Zhao, and Yazhou Ren. EEGDnet: Fusing non-local and local self-similarity for EEG signal denoising with transformer. *Computers in Biology and Medicine*, 151:106248, 2022.

Rakesh Ranjan, Bikash Chandra Sahana, and Ashish Kumar Bhandari. Motion artifacts suppression from EEG signals using an adaptive signal denoising method. *IEEE Transactions on Instrumentation and Measurement*, 71:1–10, 2022.

Ville Rantanen, Mirja Ilves, Antti Vehkaoja, Anton Kontunen, Jani Lylykangas, Eeva Mäkelä, Markus Rautiainen, Veikko Surakka, and Jukka Lekkala. A survey on the feasibility of surface EMG in facial pacing. In *Proceedings of the 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1688–1691, 2016.

Qiyu Rao, Zdenka Babic, Scott C Douglas, and Danilo P Mandic. Panorama: An enabling technology for Hearables. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1–5, 2025.

Ulla Delfana Rosiani, Pramana Yoga Saputra, and Muhammad Afif Hendrawan. The impact of segment length on EEG based biometric system. In *Proceedings of the IEEE 7th Information Technology International Seminar*, pp. 1–5, 2021.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Ali Hossam Shoeb. *Application of machine learning to epileptic seizure onset detection and treatment*. PhD thesis, Massachusetts Institute of Technology, 2009.

Weitong Sun, Yuping Su, Xia Wu, and Xiaojun Wu. A novel end-to-end 1D-ResCNN model to remove artifact from EEG signals. *Neurocomputing*, 404:108–121, 2020.

Haozhe Tian, Qiyu Rao, Nina Moutonnet, Pietro Ferraro, and Danilo Mandic. Machine intelligence on the edge: Interpretable cardiac pattern localisation using reinforcement learning. *arXiv preprint arXiv:2508.21652*, 2025.

Jose Antonio Urigüen and Begoña Garcia-Zapirain. EEG artifact removal—state-of-the-art and guidelines. *Journal of Neural Engineering*, 12(3):031001, 2015.

Raphael Vallat and Matthew P Walker. An open-source, high-performance tool for automated sleep staging. *eLife*, 10:e70092, 2021.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992.

Banghua Yang, Kaiwen Duan, and Tao Zhang. Removal of EOG artifacts from EEG using a cascade of sparse autoencoder and recursive least squares adaptive filter. *Neurocomputing*, 214:1053–1060, 2016.

Haoming Zhang, Chen Wei, Mingqi Zhao, Quanying Liu, and Haiyan Wu. A novel convolutional neural network model to remove muscle artifacts from EEG. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1265–1269, 2021.

Liping Zhang, Anzi Li, Shukai Chen, Wei Ren, and Kim-Kwang Raymond Choo. A secure, flexible, and PPG-based biometric scheme for healthy IoT using homomorphic random forest. *IEEE Internet of Things Journal*, 11(1):612–622, 2023.

APPENDIX

A PROOF OF THEOREM 1

Theorem 1. *Suppose* $(\mathbf{s}^l, \mathbf{s}^r) \sim \pi_\psi(\cdot \mid \mathbf{s})$ *are independent noisy realizations of the same underlying signal* $\mathbf{x}$*. Then the* $\theta^\star$ *that optimizes the self-supervised loss in Eq.* (1) *also minimizes the expected* $L2$ *distance between the network outputs and the ground-truth* $\mathbf{x}$*, i.e.,*

$$\theta^\star = \arg\min_\theta \mathbb{E}\Big[\big\|f_\theta(\mathbf{s}^l) - \mathbf{x}\big\|_2^2 + \big\|f_\theta(\mathbf{s}^r) - \mathbf{x}\big\|_2^2\Big].$$

Proof. The independent noisy realization pairs $(\mathbf{s}^l, \mathbf{s}^r)$ can be written as $\mathbf{s}^l = \mathbf{x} + \mathbf{n}^l$ and $\mathbf{s}^r = \mathbf{x} + \mathbf{n}^r$, where $\mathbf{n}^l$ and $\mathbf{n}^r$ are zero-mean and mutually independent, and each is independent of $\mathbf{x}$. Expanding the first term of the self-supervised loss in Eq. (1) gives

$$\begin{aligned}\mathbb{E}\big[\|f_\theta(\mathbf{s}^l) - \mathbf{s}^r\|_2^2\big] &= \mathbb{E}\big[\|f_\theta(\mathbf{s}^l) - \mathbf{x} - \mathbf{n}^r\|_2^2\big] \\ &= \mathbb{E}\big[\|f_\theta(\mathbf{s}^l) - \mathbf{x}\|_2^2\big] - 2\,\mathbb{E}\Big[{\mathbf{n}^r}^\top\big(f_\theta(\mathbf{s}^l) - \mathbf{x}\big)\Big] + \mathbb{E}\big[\|\mathbf{n}^r\|_2^2\big] \\ &= \mathbb{E}\big[\|f_\theta(\mathbf{s}^l) - \mathbf{x}\|_2^2\big] + \mathbb{E}\big[\|\mathbf{n}^r\|_2^2\big]\,,\end{aligned}$$

where the last equality follows from $\mathbb{E}\big[{\mathbf{n}^r}^\top f_\theta(\mathbf{s}^l)\big] = 0$ and $\mathbb{E}\big[{\mathbf{n}^r}^\top \mathbf{x}\big] = 0$, since $\mathbf{n}^r$ is zero-mean and independent of both $\mathbf{x}$ and $\mathbf{s}^l$. By the same argument with the roles of $\mathbf{s}^l$ and $\mathbf{s}^r$ swapped, the second term of Eq. (1) satisfies

$$\mathbb{E}\big[\|f_\theta(\mathbf{s}^r) - \mathbf{s}^l\|_2^2\big] = \mathbb{E}\big[\|f_\theta(\mathbf{s}^r) - \mathbf{x}\|_2^2\big] + \mathbb{E}\big[\|\mathbf{n}^l\|_2^2\big]\,.$$

Summing the two terms and noting that $\mathbb{E}\big[\|\mathbf{n}^l\|_2^2\big] + \mathbb{E}\big[\|\mathbf{n}^r\|_2^2\big]$ does not depend on θ, we obtain

$$\arg\min_\theta \mathbb{E}\big[\|f_\theta(\mathbf{s}^l) - \mathbf{s}^r\|_2^2 + \|f_\theta(\mathbf{s}^r) - \mathbf{s}^l\|_2^2\big] = \arg\min_\theta \mathbb{E}\big[\|f_\theta(\mathbf{s}^l) - \mathbf{x}\|_2^2 + \|f_\theta(\mathbf{s}^r) - \mathbf{x}\|_2^2\big]\,,$$

which proves the theorem. □

B PROOF OF PROPOSITION 1

Proposition 1. *For any signal window* $I^{(k)}$*, the partitioning strategy used in iPSD can always achieve a partition that is at least as good as interleaved downsampling (Mansour & Heckel, 2023), in terms of the mismatch between the underlying clean components of the sub-signals* $(\mathbf{s}_{I^{(k),l}}, \mathbf{s}_{I^{(k),r}})$.

Proof. The fixed interleaved partition used in Mansour & Heckel (2023) is one particular case in the iPSD search space, which contains all $\binom{W}{W/2}/2$ candidate within-window partitions of $I^{(k)}$. It follows immediately that iPSD can achieve a partition at least as good as the fixed interleaved partition, since the latter is always available as a candidate.

Moreover, the iPSD partition can be strictly better than interleaved downsampling. Consider a clean signal window

$$\mathbf{x}_{I^{(k)}} = (1, -1, 1, -1).$$

Under interleaved partitioning, the resulting clean sub-signals are

$$\mathbf{x}_{I^{(k),l}} = (1, 1), \qquad \mathbf{x}_{I^{(k),r}} = (-1, -1),$$

or equivalently in reversed order. These two sub-signals clearly do not share the same underlying clean component. However, iPSD may choose the feasible within-window partition

$$I^{(k),l} = \{1, 2\}, \qquad I^{(k),r} = \{3, 4\},$$

which yields sub-signals with the same underlying signal:

$$\mathbf{x}_{I^{(k),l}} = (1, -1), \qquad \mathbf{x}_{I^{(k),r}} = (1, -1).$$

This and many other examples, such as those illustrated in Fig. 1 and the real-life recordings in Table 10 of Appendix I.3, show that for signals with periodic structure, the flexibility of iPSD's search space yields strictly better partitions than fixed interleaved downsampling. □

C JUSTIFICATION FOR THE RL REWARD SIGNAL

In this section, we show that under reasonable assumptions, using the converged denoiser loss as RL reward drives π_ψ toward $\mathbf{s}^l$ and $\mathbf{s}^r$ that are maximally informative of each other's underlying clean signal. For conciseness, we present the $l \to r$ direction; the reverse direction follows symmetrically.

We follow the standard N2N assumption that the noise $\mathbf{n}$ is zero-mean and independent of the clean signal $\mathbf{x}$. In the iPSD setting, for the partitioned pair $(\mathbf{s}^l, \mathbf{s}^r)$ induced by π_ψ, we assume that the right-side noise $\mathbf{n}^r$ remains conditionally zero-mean given the left-side observation $\mathbf{s}^l$ and the right-side clean signal $\mathbf{x}^r$, i.e., $\mathbb{E}[\mathbf{n}^r \mid \mathbf{s}^l, \mathbf{x}^r] = \mathbf{0}$. We also assume that the noise power $\mathbb{E}[\|\mathbf{n}^r\|_2^2]$ does not depend on the partition policy π_ψ, which is reasonable since the two partitions have fixed cardinality and the noise statistics are approximately homogeneous across samples. For $(\mathbf{s}^l, \mathbf{s}^r) \sim \pi_\psi(\cdot \mid \mathbf{s})$, the N2N denoising objective can be rewritten as:

$$\begin{aligned}
\arg\min_\theta \mathbb{E}\left[\|f_\theta(\mathbf{s}^l) - \mathbf{s}^r\|_2^2\right] &= \arg\min_\theta \mathbb{E}\left[\left\|\left(f_\theta(\mathbf{s}^l) - \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]\right) - \left(\mathbf{s}^r - \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]\right)\right\|_2^2\right] \\
&= \arg\min_\theta \mathbb{E}\left[\left\|f_\theta(\mathbf{s}^l) - \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]\right\|_2^2\right] \\
&\quad - 2\,\mathbb{E}\left[\left(f_\theta(\mathbf{s}^l) - \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]\right)^\top \left(\mathbf{s}^r - \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]\right)\right] \\
&\quad + \mathbb{E}\left[\left\|\mathbf{s}^r - \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]\right\|_2^2\right] \\
&\overset{(i)}{=} \arg\min_\theta \mathbb{E}\left[\left\|f_\theta(\mathbf{s}^l) - \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]\right\|_2^2\right],
\end{aligned} \tag{7}$$

where (i) follows because the third term does not depend on θ and can be dropped from the $\arg\min$, and the cross term vanishes by applying the tower property $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X \mid \mathbf{s}^l]]$:

$$\begin{aligned}
&\mathbb{E}\left[\left(f_\theta(\mathbf{s}^l) - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l]\right)^\top \left(\mathbf{s}^r - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l]\right)\right] \\
&\quad = \mathbb{E}\left[\mathbb{E}\left[\left(f_\theta(\mathbf{s}^l) - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l]\right)^\top \left(\mathbf{s}^r - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l]\right) \;\middle|\; \mathbf{s}^l\right]\right] \\
&\quad \overset{(ii)}{=} \mathbb{E}\left[\left(f_\theta(\mathbf{s}^l) - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l]\right)^\top \mathbb{E}\left[\mathbf{s}^r - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l] \;\middle|\; \mathbf{s}^l\right]\right] \\
&\quad = \mathbb{E}\left[\left(f_\theta(\mathbf{s}^l) - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l]\right)^\top \cdot \mathbf{0}\right] = 0,
\end{aligned}$$

where (ii) holds because $f_\theta(\mathbf{s}^l) - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l]$ is a deterministic function of $\mathbf{s}^l$ that can be pulled out of the conditional expectation given $\mathbf{s}^l$. The remaining inner factor then vanishes:

$$\mathbb{E}\left[\mathbf{s}^r - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l] \;\middle|\; \mathbf{s}^l\right] = \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l] - \mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l] = \mathbf{0},$$

where we used linearity of conditional expectation together with the fact that $\mathbb{E}[\mathbf{s}^r \mid \mathbf{s}^l]$ is itself a deterministic function of $\mathbf{s}^l$ and therefore equals its own conditional expectation given $\mathbf{s}^l$. Let $f_{\theta^\star}$ denote the optimal denoiser.

Assuming sufficient expressiveness of the denoising module f_θ, Eq. (7) shows that

$$f_{\theta^\star}(\mathbf{s}^l) = \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]. \tag{8}$$

Let $\mathcal{L}^\star(\pi_\psi)$ denote the converged N2N loss in the $l \to r$ direction for $(\mathbf{s}^l, \mathbf{s}^r) \sim \pi_\psi(\cdot \mid \mathbf{s})$. Following Eq. (8),

$$\begin{aligned}
\mathcal{L}^\star(\pi_\psi) &= \mathbb{E}\left[\|f_{\theta^\star}(\mathbf{s}^l) - \mathbf{s}^r\|_2^2\right] \\
&= \mathbb{E}\left[\left\|\mathbf{s}^r - \mathbb{E}\left[\mathbf{s}^r \mid \mathbf{s}^l\right]\right\|_2^2\right] \\
&= \mathbb{E}\left[\left\|\mathbf{x}^r + \mathbf{n}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right] - \mathbb{E}\left[\mathbf{n}^r \mid \mathbf{s}^l\right]\right\|_2^2\right] \\
&\overset{(iii)}{=} \mathbb{E}\left[\left\|\left(\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]\right) + \mathbf{n}^r\right\|_2^2\right] \\
&\overset{(iv)}{=} \mathbb{E}\left[\left\|\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]\right\|_2^2\right] + \mathbb{E}\left[\|\mathbf{n}^r\|_2^2\right],
\end{aligned} \tag{9}$$

where (iii) uses $\mathbb{E}[\mathbf{n}^r \mid \mathbf{s}^l] = \mathbf{0}$, which follows from the tower property: $\mathbb{E}[\mathbf{n}^r \mid \mathbf{s}^l] = \mathbb{E}\big[\mathbb{E}[\mathbf{n}^r \mid \mathbf{s}^l, \mathbf{x}^r] \mid \mathbf{s}^l\big] = \mathbf{0}$ by our assumption. Step (iv) follows because the cross term vanishes. To see this, apply the tower property with conditioning on $(\mathbf{s}^l, \mathbf{x}^r)$:

$$\begin{aligned}\mathbb{E}\left[\left(\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]\right)^\top \mathbf{n}^r\right] &= \mathbb{E}\left[\mathbb{E}\left[\left(\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]\right)^\top \mathbf{n}^r \;\middle|\; \mathbf{s}^l, \mathbf{x}^r\right]\right] \\ &= \mathbb{E}\left[\left(\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]\right)^\top \mathbb{E}\left[\mathbf{n}^r \mid \mathbf{s}^l, \mathbf{x}^r\right]\right] \\ &= \mathbb{E}\left[\left(\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]\right)^\top \cdot \mathbf{0}\right] = 0,\end{aligned}$$

where the second equality follows because $\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]$ is determined by $\mathbf{s}^l$ and $\mathbf{x}^r$, and therefore behaves as a constant under the conditional expectation given $(\mathbf{s}^l, \mathbf{x}^r)$; the third uses the assumption that $\mathbb{E}[\mathbf{n}^r \mid \mathbf{s}^l, \mathbf{x}^r] = \mathbf{0}$.

Upon combining the two results, we obtain the following regarding the converged loss of the denoiser:

$$\mathcal{L}^\star(\pi_\psi) = \underbrace{\mathbb{E}\left[\left\|\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]\right\|_2^2\right]}_{\text{depends on } \pi_\psi} + \underbrace{\mathbb{E}\left[\|\mathbf{n}^r\|_2^2\right]}_{\text{constant in } \pi_\psi}. \tag{10}$$

Note that the first term depends on π_ψ through the joint distribution of $(\mathbf{s}^l, \mathbf{x}^r)$ induced by the partition policy. By our assumption that $\mathbb{E}[\|\mathbf{n}^r\|_2^2]$ does not depend on π_ψ, the second term is constant with respect to ψ. Therefore, maximizing the RL reward $-\mathcal{L}^\star(\pi_\psi)$ is equivalent to minimizing the first term of Eq. (10):

$$\arg\max_\psi \left(-\mathcal{L}^\star(\pi_\psi)\right) = \arg\min_\psi \mathbb{E}\left[\left\|\mathbf{x}^r - \mathbb{E}\left[\mathbf{x}^r \mid \mathbf{s}^l\right]\right\|_2^2\right], \tag{11}$$

which is the minimum mean squared error (MMSE) of predicting the right-side clean signal $\mathbf{x}^r$ from the left-side observation $\mathbf{s}^l$. Hence, using the negative converged N2N loss as the RL reward drives π_ψ toward partitions under which $\mathbf{s}^l$ is maximally predictive of $\mathbf{x}^r$ in the MMSE sense; by symmetry, $\mathbf{s}^r$ is maximally predictive of $\mathbf{x}^l$.

D Derivation of the policy gradient

We now derive the gradient of the expected denoising reward $\mathbb{E}_{(\mathbf{s}^l,\mathbf{s}^r)\sim\pi(\cdot|\mathbf{s})}\left[R(\mathbf{s}^l, \mathbf{s}^r)\right]$ with respect to the parameters ψ. Applying the log-derivative identity $\nabla_x f(x) = f(x)\,\nabla_x \log f(x)$ yields

$$\begin{aligned}\nabla_\psi \mathbb{E}_{(\mathbf{s}^l,\mathbf{s}^r)\sim\pi(\cdot|\mathbf{s})}\left[R(\mathbf{s}^l, \mathbf{s}^r)\right] &= \nabla_\psi \sum_{(\mathbf{s}^l,\mathbf{s}^r)} \pi_\psi(\mathbf{s}^l, \mathbf{s}^r \mid \mathbf{s}) R(\mathbf{s}^l, \mathbf{s}^r) \\ &= \sum_{(\mathbf{s}^l,\mathbf{s}^r)} \pi_\psi(\mathbf{s}^l, \mathbf{s}^r \mid \mathbf{s}) \left[\nabla_\psi \log \pi_\psi(\mathbf{s}^l, \mathbf{s}^r \mid \mathbf{s}) R(\mathbf{s}^l, \mathbf{s}^r)\right] \\ &= \mathbb{E}_{(\mathbf{s}^l,\mathbf{s}^r)}\left[\nabla_\psi \log \pi_\psi(\mathbf{s}^l, \mathbf{s}^r \mid \mathbf{s}) R(\mathbf{s}^l, \mathbf{s}^r)\right].\end{aligned} \tag{12}$$

Equation (12) is the standard REINFORCE estimator (Williams, 1992): the intractable sum over all partitions is replaced by a Monte Carlo expectation, making the gradient estimable from sampled partitions $(\mathbf{s}^l, \mathbf{s}^r) \sim \pi_\psi(\cdot \mid \mathbf{s})$.

E Detailed implementation of iPSD-Zero

In this section, we provide details of iPSD-Zero. Let arm a denote the partition that corresponds to $(\mathbf{s}_a^l, \mathbf{s}_a^r)$. Pulling a leads to a certain reward $R(\mathbf{s}_a^l, \mathbf{s}_a^r)$, defined as in Eq. (4) as the negated denoiser f_θ loss at convergence. Due to the stochastic nature of training neural networks, $R(\mathbf{s}_a^l, \mathbf{s}_a^r)$ is also stochastic for each arm. To balance the exploration of less frequently sampled partitions with the exploitation of those that have already demonstrated strong denoising performance, the lil' UCB index for arm a at round t is defined as

$$l_a(t) = \hat{\mu}_a(t) + (1+\beta)(1+\sqrt{\epsilon})\sqrt{\frac{2\sigma^2(1+\epsilon)\log\left(\frac{\log((1+\epsilon)T_a(t))}{\delta}\right)}{T_a(t)}}, \tag{13}$$

Algorithm 2. iPSD-Zero

Require: Noisy signal $\mathbf{s}$; hyperparameters $(\alpha, \beta, \sigma, \epsilon, \delta)$; learning rate η^θ

Initialization: pull each arm once to obtain $\hat{\mu}_a(0)$ for all a; set $T_a(0) \leftarrow 1$ for all a; let $t \leftarrow 0$
while $T_a(t) < 1 + \alpha \sum_{i \neq a} T_i(t)$ for all arms a **do**
  Compute $l_a(t)$ for all a with Eq. (13)
  Select arm $a_t = \arg\max_a l_a(t)$ and partition $\mathbf{s}$ into $(\mathbf{s}^l_{a_t}, \mathbf{s}^r_{a_t})$
  Initialize f_θ; set $N = 0$
  while θ not converged **do**
    $J^d_n = \left\| f_\theta(\mathbf{s}^l_{a_t}) - \mathbf{s}^r_{a_t} \right\|_2^2 + \left\| f_\theta(\mathbf{s}^r_{a_t}) - \mathbf{s}^l_{a_t} \right\|_2^2$
    $\theta \leftarrow \theta - \eta^\theta \nabla_\theta J^d_n; \quad N = N + 1$
  end while
  Reward $r_t = R(\mathbf{s}^l_{a_t}, \mathbf{s}^r_{a_t}) = -\frac{1}{10} \sum_{n=N-9}^{N} J^d_n$
  $\hat{\mu}_a(t+1) \leftarrow \begin{cases} \dfrac{T_a(t)\,\hat{\mu}_a(t) + r_t}{T_a(t)+1}, & a = a_t \\ \hat{\mu}_a(t), & a \neq a_t \end{cases}$
  $T_a(t+1) \leftarrow \begin{cases} T_a(t) + 1, & a = a_t \\ T_a(t), & a \neq a_t \end{cases}$
  $t \leftarrow t + 1$
end while
$a^\star \leftarrow \arg\max_a T_a(t)$

where $\hat{\mu}_a(t)$ denotes the empirical mean reward of arm a up to round t, $T_a(t)$ is the number of times arm a has been pulled up to round t, and σ^2 is the sub-Gaussian variance proxy of the reward distribution. The confidence parameter $\delta \in (0, 1)$ and the exploration parameters $\beta, \epsilon > 0$ are user-specified hyperparameters that jointly control the tightness of the confidence interval. The first term in Eq. (13) promotes the exploitation of arms with higher observed rewards, while the second term promotes the exploration of other arms with decreasing intensity as $T_a(t)$ increases. Algorithm 2 summarizes our adaptation of lil' UCB to the denoising process in the zero-shot setting.

In iPSD-Zero, hyperparameters are set to $\beta = 1$, $\alpha = [(2 + \beta)/\beta]^2$, $\sigma = 0.008$, $\epsilon = 0.01$, and $\delta = 0.55 \times 10^{-4}$. The iPSD-Zero method is initialized by pulling each arm once. We then iteratively apply lil' UCB to select the next arm. This results in increasingly confident estimates of partition quality for each arm, reflected by the average reward. At convergence, the algorithm identifies the optimal partition, $(\mathbf{s}^l_{a^\star}, \mathbf{s}^r_{a^\star})$, which is then used to train the final denoiser $f_{\theta^\star}$. The design of iPSD-Zero transforms the zero-shot denoising problem into a structured exploration-exploitation task, where, by leveraging lil' UCB, we achieve principled and sample-efficient identification of the partition that yields the most reliable self-supervision.

F Notes on iPSD implementation

Hardware The iPSD is implemented in Python 3.12.5 and tested on an Ubuntu 22.04 machine with a 13th Gen Intel Core i7-13850HX CPU, an Nvidia RTX 3500 Ada GPU (12 GB VRAM), and 32 GB of RAM. Compared with the state-of-the-art unsupervised DL baseline (Chen et al., 2025), iPSD trains in 5.2 h versus >12 h, uses only 4% of the parameters, and denoises a 10 s EEG segment in 90 ms. Even on a CPU-only machine, iPSD denoises comfortably within real-time constraints due to its lightweight architecture.

Window length We evaluate different lengths of the localized window (Table 3). Since the window length must divide the signal length $L = 2560$ and be even (so that $\mathbf{s}^l$ and $\mathbf{s}^r$ have equal length), we ablate over window lengths of 4, 8, and 10 samples. Longer windows (e.g., the next feasible length, 16 samples) produce too many candidate partitions, inflating the parameter count of π_ψ to an intractable size. Results show that a window length of 8 samples is optimal, providing sufficient flexibility while keeping the search space tractable for efficient policy training ($W = 8$ already admits 35 per window and 11200 per signal distinct partitions).

Table 3: Ablation study on localized window lengths.

	WGN						EMG					
	−5 dB Input SNR			0 dB Input SNR			−5 dB Input SNR			0 dB Input SNR		
Window length	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE
4 samples	2.61	16.71	60.03	5.38	19.48	55.27	1.78	15.88	60.07	4.77	18.87	45.13
8 samples	3.97	18.07	55.14	5.78	19.90	41.13	3.63	17.73	53.04	5.95	20.05	39.21
10 samples	3.96	18.06	55.47	5.52	19.62	46.92	3.37	17.47	56.62	5.66	19.76	42.12

S-MSE denotes the spectral MSE.

Table 4: Partitioning module architectures

Model	**Layer**	**Channel Progression**	**Hidden dim**
iPSD-MLP	FC-1	$8 \to 64$	-
	FC-2	$64 \to 128$	-
	FC-3	$128 \to 35$	-
iPSD-U-Net	Conv-Enc-1	$8 \to 64 \to 64$	-
	Conv-Enc-2	$64 \to 128 \to 128$	-
	Conv-Enc-3	$128 \to 256 \to 256$	-
	Conv-Dec-3	$512 \to 128 \to 128$	-
	Conv-Dec-2	$256 \to 64 \to 64$	-
	Conv-Dec-1	$128 \to 64 \to 64$	-
	Conv-Out	$64 \to 35$	-
iPSD-RNN	RNN-1	$8 \to 128$	64
	RNN-2	$128 \to 128$	64
	FC-1	$128 \to 256$	-
	FC-2	$256 \to 256$	-
	FC-3	$256 \to 35$	-
iPSD-GRU	GRU-1	$8 \to 128$	64
	GRU-2	$128 \to 128$	64
	FC-1	$128 \to 256$	-
	FC-2	$256 \to 256$	-
	FC-3	$256 \to 35$	-

FC = Fully Connected; Conv-Enc/Dec = 1D Convolutional encoder/decoder block (2 conv layers each).
FC and Conv layers use the ReLU activation; Batch normalization is applied in all Conv layers.
iPSD-RNN and iPSD-GRU are bidirectional.

Partitioning module Before settling on the architecture of the partitioning model π_ψ, we compared four candidates: (i) iPSD-MLP, a window-independent baseline without temporal context; (ii) iPSD-U-Net, a multi-scale fully convolutional architecture; (iii) iPSD-RNN, a recurrent neural network capturing sequential context across windows; and (iv) iPSD-GRU, a gated RNN variant with selective retention of long-range context. Table 4 summarizes the detailed architectures of the four networks. All four variants are trained using the same PPO hyperparameter setting.

Table 5 compares the denoising performance of all the iPSD variants. Under WGN contamination, the iPSD-GRU achieved the best overall performance, benefiting from its gating mechanisms, which effectively capture sequential context while mitigating vanishing/exploding gradients. Notably, the other sequential model, iPSD-RNN, performed worse than iPSD-MLP, which does not use any sequential context. This indicates that while sequential context is indeed important for partitioning $\mathbf{s}$, long sequences must be handled appropriately, which GRU achieves successfully. The iPSD-U-Net model does not perform well compared to the other variants, which may be attributed to its relatively limited receptive field when applied to long EEG signals. Based on these results, we chose GRU as the architecture for the partitioning model.

Table 5: Denoising performance of different partitioning module architectures.

Method	-5 dB Input SNR			0 dB Input SNR		
	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE
WGN artifacts						
iPSD-MLP	3.88	17.97	63.66	5.73	19.83	54.11
iPSD-U-Net	3.76	17.87	69.07	5.69	19.79	52.46
iPSD-RNN	3.83	17.94	65.89	5.71	19.81	51.24
iPSD-GRU	3.97	18.07	55.14	5.78	19.90	41.13
EMG artifacts						
iPSD-MLP	3.19	17.29	58.46	5.23	19.33	49.48
iPSD-U-Net	3.03	17.13	59.34	5.18	19.29	50.01
iPSD-RNN	2.94	17.04	61.30	5.05	19.15	51.55
iPSD-GRU	3.63	17.73	53.04	5.95	20.05	39.21

Table 6: Baselines used in quantitative comparisons

Method	**Description**
mVMD (Dora & Biswal, 2020)	modified Variational Mode Decomposition for correlated noise
CEEMDAN (Colominas et al., 2014)	ensemble Empirical Mode Decomposition with noise injection
FrWT (Houamed et al., 2020)	thresholding of fractional Wavelet Transform coefficients
EMD-LoG (Ranjan et al., 2022)	Laplacian-of-Gaussian filtering of Empirical Mode Decomposition components
WPT-ICA (Kerechanin & Bobrov, 2022)	Wavelet Packet Transform and Independent Component Analysis
Optimal-WT (Alyasseri et al., 2019)	optimizing the thresholding of the Discrete Wavelet Transform
Chen et al. (Chen et al., 2025)	self-supervised consistency learning with masked reconstruction

G BASELINES

Table 6 provides the details of the baselines used in our quantitative evaluation. These baselines cover the dominant paradigms in EEG artifact removal, spanning wavelet-based methods such as FrWT by Houamed et al. (2020) and Optimal-WT by Alyasseri et al. (2019); mode decomposition methods such as mVMD by Dora & Biswal (2020), CEEMDAN by Colominas et al. (2014), and EMD-LoG by Ranjan et al. (2022); hybrid methods such as WPT-ICA by Kerechanin & Bobrov (2022); and self-supervised DL methods by Chen et al. (2025). All baselines are implemented using the authors' released code where available, otherwise re-implemented from the published descriptions. Hyperparameters follow each method's original specification. For CEEMDAN, we use the PyEMD implementation with default parameters and discard the first 2 intrinsic mode functions. For Chen et al., we retrain it on the same training set that iPSD uses with the authors' default hyperparameters.

H EVALUATION METRICS

Signal-to-Noise Ratio (SNR) SNR is defined as the ratio of the clean signal power to the residual noise power after denoising. A higher SNR indicates better noise suppression while retaining the power of the original signal. The SNR (in dB) is computed as

$$\mathrm{SNR} = 10 \log_{10} \left(\frac{\|\mathbf{x}\|_2^2}{\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2} \right), \tag{14}$$

where $\mathbf{x}$ denotes the clean reference signal and $\hat{\mathbf{x}}$ the denoising output. This metric directly measures reconstruction accuracy in the time domain.

Table 7: Standard deviation of denoising performance. Values are reported as mean ± std.

	Method	-5 dB Input SNR			0 dB Input SNR		
		SNR	PSNR	S-MSE	SNR	PSNR	S-MSE
WGN	mVMD	$-5.05_{\pm 1.93}$	$9.05_{\pm 3.67}$	$(1.46 \pm 0.31)_{\times 10^4}$	$2.18_{\pm 2.85}$	$16.28_{\pm 3.62}$	$(1.13 \pm 0.36)_{\times 10^4}$
	CEEMDAN	$-3.85_{\pm 1.61}$	$10.25_{\pm 3.31}$	$(0.98 \pm 0.24)_{\times 10^4}$	$4.13_{\pm 1.84}$	$18.23_{\pm 2.67}$	$(0.78 \pm 0.25)_{\times 10^4}$
	FrWT	$-3.54_{\pm 1.41}$	$10.56_{\pm 2.91}$	$(0.96 \pm 0.21)_{\times 10^4}$	$3.43_{\pm 1.92}$	$17.53_{\pm 2.73}$	$(0.83 \pm 0.27)_{\times 10^4}$
	EMD-LoG	$2.28_{\pm 1.79}$	$16.38_{\pm 3.58}$	$(0.80 \pm 0.17)_{\times 10^4}$	$4.40_{\pm 1.87}$	$18.50_{\pm 2.65}$	$(0.80 \pm 0.25)_{\times 10^4}$
	WPT-ICA	$-4.89_{\pm 1.87}$	$9.21_{\pm 2.90}$	$(1.63 \pm 0.28)_{\times 10^4}$	$1.74_{\pm 1.55}$	$15.84_{\pm 2.46}$	$(1.23 \pm 0.25)_{\times 10^4}$
	Chen et al.	$2.01_{\pm 1.28}$	$16.11_{\pm 2.52}$	$(0.28 \pm 0.07)_{\times 10^4}$	$3.25_{\pm 1.67}$	$17.35_{\pm 2.69}$	$(0.21 \pm 0.04)_{\times 10^4}$
	Optimal-WT	$3.34_{\pm 1.50}$	$17.44_{\pm 3.18}$	$(0.36 \pm 0.11)_{\times 10^4}$	$5.07_{\pm 1.72}$	$19.18_{\pm 2.73}$	$(0.27 \pm 0.11)_{\times 10^4}$
	iPSD	$3.97_{\pm 0.64}$	$18.07_{\pm 2.85}$	$55.14_{\pm 27.62}$	$5.78_{\pm 1.25}$	$19.90_{\pm 2.38}$	$41.13_{\pm 18.33}$
	iPSD-Zero	$3.89_{\pm 0.79}$	$18.00_{\pm 2.52}$	$58.34_{\pm 25.13}$	$5.76_{\pm 1.32}$	$19.86_{\pm 2.60}$	$48.67_{\pm 18.65}$
EMG	mVMD	$-6.05_{\pm 2.31}$	$8.05_{\pm 3.45}$	$(5.40 \pm 1.14)_{\times 10^4}$	$0.93_{\pm 2.24}$	$15.03_{\pm 3.25}$	$(3.06 \pm 0.97)_{\times 10^4}$
	CEEMDAN	$0.25_{\pm 1.31}$	$14.35_{\pm 2.75}$	$(0.36 \pm 0.14)_{\times 10^3}$	$3.26_{\pm 2.01}$	$17.36_{\pm 3.02}$	$(0.22 \pm 0.08)_{\times 10^3}$
	FrWT	$-0.06_{\pm 1.42}$	$14.04_{\pm 3.21}$	$(0.55 \pm 0.17)_{\times 10^3}$	$3.53_{\pm 1.43}$	$17.63_{\pm 3.12}$	$(0.62 \pm 0.21)_{\times 10^3}$
	EMD-LoG	$0.32_{\pm 1.63}$	$14.42_{\pm 3.20}$	$(0.38 \pm 0.13)_{\times 10^3}$	$3.55_{\pm 1.58}$	$17.65_{\pm 3.31}$	$(0.22 \pm 0.11)_{\times 10^3}$
	WPT-ICA	$-0.09_{\pm 1.22}$	$14.01_{\pm 2.47}$	$(0.66 \pm 0.13)_{\times 10^3}$	$2.36_{\pm 2.07}$	$16.46_{\pm 3.59}$	$(0.31 \pm 0.08)_{\times 10^3}$
	Chen et al.	$-0.05_{\pm 0.18}$	$14.05_{\pm 2.56}$	$(0.42 \pm 0.14)_{\times 10^3}$	$3.01_{\pm 1.77}$	$17.12_{\pm 3.23}$	$(0.19 \pm 0.09)_{\times 10^3}$
	Optimal-WT	$0.33_{\pm 2.40}$	$14.42_{\pm 3.69}$	$(0.69 \pm 0.14)_{\times 10^3}$	$3.89_{\pm 1.95}$	$18.10_{\pm 3.46}$	$(0.43 \pm 0.10)_{\times 10^3}$
	iPSD	$3.63_{\pm 1.16}$	$17.73_{\pm 2.44}$	$53.04_{\pm 29.91}$	$5.95_{\pm 1.31}$	$20.05_{\pm 2.55}$	$39.21_{\pm 17.17}$
	iPSD-Zero	$3.27_{\pm 1.14}$	$17.37_{\pm 2.32}$	$56.62_{\pm 29.75}$	$5.31_{\pm 1.22}$	$19.41_{\pm 2.65}$	$44.12_{\pm 17.68}$

* Pink (blue) highlights the best (second-best) performance; S-MSE stands for spectral MSE.

Peak Signal-to-Noise Ratio (PSNR) PSNR is defined as the ratio of the squared peak amplitude of the clean signal to the average power of the residual noise after denoising. A higher PSNR indicates better noise suppression while retaining extreme values of the original signal, which are often clinically relevant information for EEG signals (e.g., spikes or sharp waves). The PSNR (in dB) is defined as

$$\mathrm{PSNR} = 10 \log_{10} \left(\frac{x_{\max}^2}{\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 / L} \right), \tag{15}$$

where L is the signal length and $x_{\max}$ is the maximum amplitude of the clean signal. Compared to SNR, which focuses on average power, PSNR is more sensitive to extreme values in the signal.

Mean Squared Error of the Spectrum (Spectral MSE) Since the clinical interpretation of EEG signals often relies on frequency-domain information (e.g., power in δ, θ, α, β, and γ bands) (Vallat & Walker, 2021), we also measure reconstruction accuracy in the frequency domain. Spectral MSE is defined as the mean squared error between PSDs of the clean and denoised signals. The PSD of a signal $\mathbf{x}$ (in dB) is defined as

$$\mathrm{spec}_{\mathbf{x}}(f) = 10 \log_{10} \left(\sum_{m=-\infty}^{\infty} R_{\mathbf{xx}}[m] \mathrm{e}^{-\mathrm{j}2\pi \mathrm{fm}} \right), \tag{16}$$

where $R_{\mathbf{xx}}[m] = \mathbb{E}\left[x_i \cdot x_{i+m}\right]$ denotes the autocorrelation of $\mathbf{x}$, and f denotes the frequency. In practice, we estimate the PSD of finite-length signals using Welch's method, which returns discrete frequency bins $\{f_k\}_{k=1}^{N_f}$. The Spectral MSE is then computed as

$$\text{Spectral MSE} = \frac{1}{N_f} \sum_{k=1}^{N_f} \left(\mathrm{spec}_{\mathbf{x}}(f_k) - \mathrm{spec}_{\hat{\mathbf{x}}}(f_k) \right)^2. \tag{17}$$

This metric assesses how well the denoising algorithm removes noise while preserving the spectral distribution of the EEG signal, which is critical for downstream EEG analysis.

Table 8: Stability of iPSD convergence. Values are reported as mean ± std over independent training runs across 5 distinct random seeds.

	Input SNR	SNR	PSNR	S-MSE	Partition entropy
WGN	−5 dB	$3.96_{\pm 0.05}$	$18.06_{\pm 0.05}$	$55.26_{\pm 0.36}$	$0.58_{\pm 0.03}$
	0 dB	$5.75_{\pm 0.06}$	$19.86_{\pm 0.06}$	$41.30_{\pm 0.51}$	$0.53_{\pm 0.04}$
EMG	−5 dB	$3.64_{\pm 0.06}$	$17.74_{\pm 0.06}$	$53.01_{\pm 0.62}$	$0.74_{\pm 0.02}$
	0 dB	$5.94_{\pm 0.03}$	$20.04_{\pm 0.03}$	$39.42_{\pm 0.23}$	$0.69_{\pm 0.02}$

Table 9: Ablation on the iPSD partitioning strategy. Values are reported as mean ± std.

		-5 dB Input SNR			0 dB Input SNR		
	Method	SNR	PSNR	S-MSE	SNR	PSNR	S-MSE
WGN	ID	$2.58_{\pm 0.95}$	$16.68_{\pm 2.92}$	$62.33_{\pm 37.51}$	$5.13_{\pm 1.79}$	$19.24_{\pm 3.27}$	$67.18_{\pm 37.21}$
	DR	$2.74_{\pm 0.94}$	$16.84_{\pm 2.89}$	$59.43_{\pm 33.21}$	$5.30_{\pm 1.19}$	$19.41_{\pm 3.05}$	$53.62_{\pm 24.73}$
	iPSD	$3.97_{\pm 0.64}$	$18.07_{\pm 2.85}$	$55.14_{\pm 27.62}$	$5.78_{\pm 1.25}$	$19.90_{\pm 2.38}$	$41.13_{\pm 18.33}$
	iPSD-Zero	$3.89_{\pm 0.79}$	$18.00_{\pm 2.52}$	$58.34_{\pm 25.13}$	$5.76_{\pm 1.32}$	$19.86_{\pm 2.60}$	$48.67_{\pm 18.65}$
EMG	ID	$1.62_{\pm 1.00}$	$15.73_{\pm 2.78}$	$67.68_{\pm 25.53}$	$4.31_{\pm 1.57}$	$18.41_{\pm 3.07}$	$60.52_{\pm 35.34}$
	DR	$2.20_{\pm 1.01}$	$16.31_{\pm 2.88}$	$63.88_{\pm 32.73}$	$4.73_{\pm 1.28}$	$18.83_{\pm 2.96}$	$52.27_{\pm 22.11}$
	iPSD	$3.63_{\pm 1.16}$	$17.73_{\pm 2.44}$	$53.04_{\pm 29.91}$	$5.95_{\pm 1.31}$	$20.05_{\pm 2.55}$	$39.21_{\pm 17.17}$
	iPSD-Zero	$3.27_{\pm 1.14}$	$17.37_{\pm 2.32}$	$56.62_{\pm 29.75}$	$5.31_{\pm 1.22}$	$19.41_{\pm 2.65}$	$44.12_{\pm 17.68}$

* Pink (blue) highlights the best (second-best) performance; S-MSE stands for spectral MSE.

I Additional results

I.1 Comparison of denoising performance

In Table 7, we report the standard deviation across all EEG segments for the comparison in Table 1 of Section 3. We observe that both iPSD and iPSD-Zero significantly outperform baselines on average, while remaining highly consistent across segments, demonstrating the robustness of the proposed framework.

I.2 Independent runs with different random seeds

For iPSD, we conduct independent training runs across 5 distinct random seeds and report the mean ± std under different noise conditions in Table 8. Across the runs, SNR and PSNR have std ≤ 0.06 dB and Spectral MSE has std $< 1.5\%$, indicating that iPSD consistently converges to high performance, with a clear gain over the interleaved downsampling baseline (by 1.41 dB higher SNR and 26.6% lower S-MSE on average). We also observe no degenerate partitions, as the entropies of selected partitions across all runs lie in the $[0.5, 0.75]$ range, far from uniform (entropy$= 1$) or deterministic (entropy$= 0$).

I.3 Ablation on intelligent partition

In Table 2 of Section 3, we ablate the partitioning module by switching the π_ψ of iPSD to simpler heuristics. In ID, we assign alternating samples to the two sub-signals, which was proven effective for locally smooth images (Mansour & Heckel, 2023). In DR, we maintain W continuous, differentiable Gumbel-Softmax logits for the samples in each window, denoted by p_{soft}. We assign the top $W/2$ samples to one sub-signal and the remainder to the other, yielding a binary assignment $p_{\text{hard}} \in \{0,1\}^W$. We then use the straight-through estimator in Bengio et al. (2013):

$$p = p_{\text{hard}} + p_{\text{soft}} - \text{stopgrad}(p_{\text{soft}}). \tag{18}$$

This means the forward pass uses discrete assignments p_{hard}, while gradients propagate through p_{soft}.

In Table 9, we report the standard deviation across all EEG segments for the ablation in Table 2, which shows that the intelligent partitioning is consistently effective across EEG segments. Notably,

Table 10: Comparing the quality of sub-signals $(\mathbf{s}^l, \mathbf{s}^r)$ generated by different partition strategies. Values are reported as mean $\pm$ std, scaled by 10^2.

		-5 dB Input SNR		0 dB Input SNR	
	Method	Signal mismatch	Noise correlation	Signal mismatch	Noise correlation
WGN	ID	$32.52_{\pm 24.29}$	$2.22_{\pm 1.55}$	$32.52_{\pm 24.29}$	$2.22_{\pm 1.55}$
	DR	$22.32_{\pm 16.62}$	$2.28_{\pm 1.61}$	$33.27_{\pm 20.53}$	$2.79_{\pm 2.38}$
	iPSD	$13.04_{\pm 11.81}$	$1.47_{\pm 0.70}$	$14.95_{\pm 12.30}$	$1.26_{\pm 0.60}$
	iPSD-Zero	$14.32_{\pm 13.18}$	$1.80_{\pm 1.72}$	$15.81_{\pm 13.28}$	$1.31_{\pm 0.91}$
EMG	ID	$32.52_{\pm 24.29}$	$12.35_{\pm 7.87}$	$32.52_{\pm 24.29}$	$12.35_{\pm 7.87}$
	DR	$34.43_{\pm 15.01}$	$14.16_{\pm 8.69}$	$28.81_{\pm 16.18}$	$11.23_{\pm 11.08}$
	iPSD	$16.23_{\pm 12.81}$	$2.25_{\pm 0.87}$	$16.40_{\pm 12.94}$	$2.44_{\pm 0.82}$
	iPSD-Zero	$18.72_{\pm 15.32}$	$3.16_{\pm 0.91}$	$18.81_{\pm 13.90}$	$2.70_{\pm 0.93}$

* Pink (blue) highlights the best (second-best) performance.

Table 11: Output SNR of iPSD and Interleaved Downsampling (ID) on EEG signals with lower sampling rate.

		WGN		EMG	
Method	Sampling rate	-5 dB	0 dB	-5 dB	0 dB
iPSD	128 Hz	$3.99_{\pm 0.70}$	$5.76_{\pm 1.23}$	$3.56_{\pm 1.12}$	$5.89_{\pm 1.27}$
	256 Hz	$3.97_{\pm 0.64}$	$5.78_{\pm 1.25}$	$3.63_{\pm 1.16}$	$5.95_{\pm 1.31}$
ID	128 Hz	$2.32_{\pm 0.90}$	$4.85_{\pm 1.23}$	$0.21_{\pm 1.63}$	$3.99_{\pm 1.44}$
	256 Hz	$2.58_{\pm 0.95}$	$5.13_{\pm 1.79}$	$1.62_{\pm 1.00}$	$4.31_{\pm 1.57}$

DR gains only 3.5% speed increase compared to RL, but at significant cost to performance. This is because DR optimizes a biased approximation of an otherwise non-smooth objective, while RL optimizes a smooth, differentiable objective by construction.

Additionally, we show that the intelligent iPSD partitioning generates sub-signals that better support self-supervised denoising by computing signal mismatch and noise correlation, which directly correspond to the central assumptions of Theorem 1 that determine denoising performance. On synthetic signals, the underlying signal $(\mathbf{x}^l, \mathbf{x}^r)$ and noise $(\mathbf{n}^l, \mathbf{n}^r)$ are recovered by applying the partition indices to $\mathbf{x}$ and $\mathbf{n}$. Signal mismatch is measured by the normalized mean squared error (NMSE) between $\mathbf{x}^l$ and $\mathbf{x}^r$, where a lower value indicates more similar underlying signals in the two sub-signals. Noise correlation is measured by the absolute Pearson correlation between $\mathbf{n}^l$ and $\mathbf{n}^r$, where a lower value indicates more independent noise in the two sub-signals. As shown in Table 10, across all noise conditions, iPSD and iPSD-Zero generate sub-signals with more similar underlying signals and less correlated noise, leading to their strong denoising performance in Table 9.

I.4 Applicability to lower sampling frequency

Many real-world EEG signals are recorded at a lower sampling frequency with reduced local smoothness. To evaluate the robustness of iPSD under such settings, we use the same synthetic dataset in the main text, but downsample it from 256 Hz to 128 Hz. Table 11 compares iPSD against the ID baseline in terms of output SNR. Because interleaved downsampling hinges on local smoothness, we observe a large performance degradation under low sampling frequency, whereas iPSD remains robust, falling by less than 0.07 dB across all noise conditions. On average, iPSD's gains over ID are 1.96 dB at 128 Hz, compared to 1.42 dB at 256 Hz. The results corroborate that flexible partitioning overcomes the strong local smoothness assumption of simpler heuristics, such as ID, and allows splits with higher underlying signal similarity.

I.5 Supervised learning baselines

Since there is no other ground-truth-free DL denoiser that we are aware of, we compare iPSD with supervised DL methods that, in a sense, indicate the upper limit to self-supervised denoising

Table 12: Comparing iPSD (trained without any ground-truth) and DL EEG denoisers (trained with full supervision using ground-truth signals).

			WGN		EMG	
Method	GT-free	Params	-5 dB	0 dB	-5 dB	0 dB
DuoCL	✗	8.67 m	$4.12_{\pm 1.31}$	$5.90_{\pm 1.55}$	$3.71_{\pm 1.21}$	$6.02_{\pm 1.58}$
DenoiseFormer	✗	6.90 m	$4.24_{\pm 1.25}$	$6.36_{\pm 1.43}$	$3.94_{\pm 1.00}$	$6.43_{\pm 1.34}$
iPSD	✓	0.33 m	$3.97_{\pm 0.64}$	$5.78_{\pm 1.25}$	$3.63_{\pm 1.16}$	$5.95_{\pm 1.31}$

Table 13: Output SNR (dB) under AR(1) noise at an input SNR of -5 dB for different autoregressive coefficients ρ. Δ stands for the improvement of iPSD over ID.

ρ	-0.6	-0.4	-0.2	0	0.2	0.4	0.6
ID	$0.95_{\pm 0.93}$	$1.89_{\pm 0.92}$	$2.44_{\pm 0.88}$	$2.58_{\pm 0.95}$	$2.42_{\pm 0.83}$	$1.70_{\pm 0.70}$	$0.73_{\pm 0.60}$
iPSD	$3.23_{\pm 0.93}$	$3.69_{\pm 0.87}$	$3.94_{\pm 0.71}$	$3.97_{\pm 0.64}$	$3.94_{\pm 0.69}$	$3.87_{\pm 0.83}$	$3.74_{\pm 0.95}$
Δ	$+2.28$	$+1.80$	$+1.50$	$+1.39$	$+1.52$	$+2.17$	$+3.01$

performance. The baselines are DuoCL (Gao et al., 2022) and DenoiseFormer (Chen et al., 2023), which rely on highly expressive DL architectures. The results in Table 12 show that, without the clean-reference supervision that these supervised baselines hinge on, iPSD comes within 0.6 dB of their performance, all while using far fewer parameters. We further train the denoiser f_θ of iPSD with clean signals, which acts as an architecture-matched upper bound. On WGN, supervised training yields 4.17 dB and 6.03 dB at -5 dB and 0 dB. On EMG, supervised training yields 3.81 dB and 6.11 dB at -5 dB and 0 dB. iPSD comes within 0.25 dB of this upper bound despite training with no access to the ground-truths.

I.6 Robustness against AR noise

Although Theorem 1 assumes that the two partitioned views are independent noisy realizations of the same underlying signal, iPSD can effectively recover the clean signal even when this assumption is not strictly satisfied. For example, iPSD achieves robust denoising performance in our main experiments despite the temporal correlation in the EMG noise. To systematically assess its sensitivity to temporal correlation and investigate potential failure modes, we conduct a controlled experiment using AR(1) noise at an input SNR of -5 dB across a range of autoregressive coefficients ρ. From Table 13, we observe the performance (indicated by SNR value) change with increasing absolute temporal correlation; the proposed iPSD retains competitive performance even in the presence of markedly correlated noise ($\rho = \pm 0.6$). Compared with the interleaved downsampling (ID) baseline, iPSD achieves an even larger performance advantage under more correlated noise. This is due to the flexible partitioning of iPSD, making it more robust to violations of the independence assumption compared to the ablated ID counterpart without the learnable partitioning component.

J Visual comparisons between iPSD and Optimal-WT

Figure 6 uses example synthetic EEG segments to illustrate the denoising performance of both iPSD and iPSD-Zero. For reference, we also include the strongest baseline, Optimal-WT. The input signals are corrupted with synthetically added WGN or EMG artifacts, with input SNR levels set to either -5 dB or 0 dB. For each input signal, we show the noisy input EEG (orange), the clean ground-truth EEG (green), the Optimal-WT output (gray), the iPSD output (pink), and the iPSD-Zero output (blue). The lower right panel for each signal displays the PSDs of the original and denoised signals, allowing visual comparison of all the methods in the frequency domain.

Under WGN contamination, although the Optimal-WT reconstruction appears visually smooth in the time domain, it incurs substantial spectral distortion, sharply attenuating components above 10 Hz and introducing spurious dips at a few frequencies, resulting in the loss of physiologically meaningful neural activity. In contrast, both iPSD and iPSD-Zero faithfully recover the underlying clean EEG, preserving the temporal morphology (e.g., the ridges and spikes) and the spectral characteristics. This

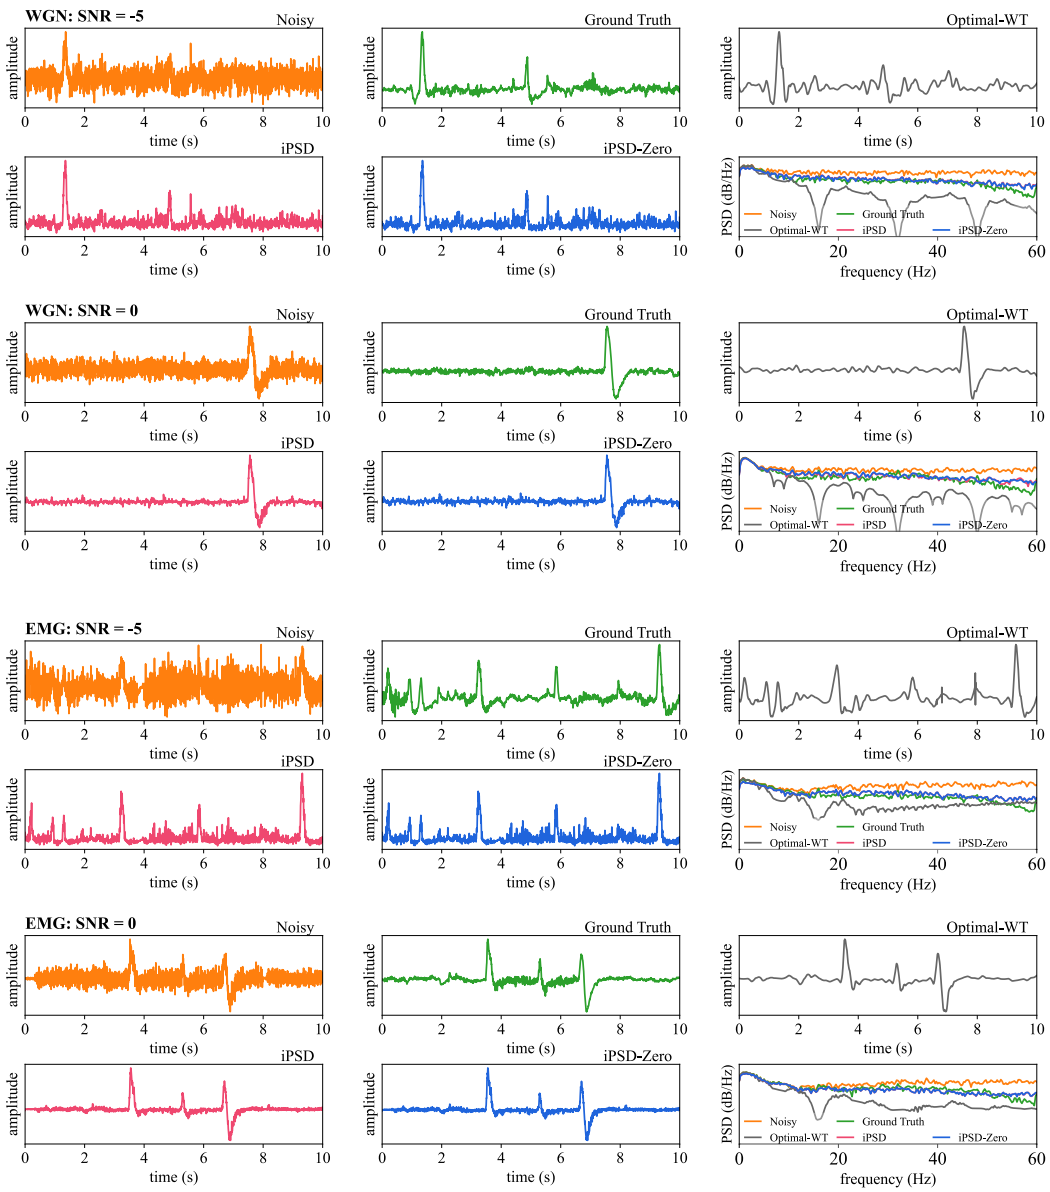

Figure 6: Visualizing the denoising of example EEG segments contaminated by WGN and EMG artifacts at input SNRs of -5 dB and 0 dB. For each example, the 6 panels show: the noisy input EEG (orange), the clean ground-truth EEG (green), the Optimal-WT output (gray), the iPSD output (pink), the iPSD-Zero output (blue), and the corresponding PSDs (lower-right panels).

highlights the advantage of iPSD's expressive DL denoiser, which captures the underlying signal structure for effective denoising, over Optimal-WT's fixed decomposition strategy, which corrupts the original EEG content.

Under EMG contamination, the advantage of iPSD becomes even more pronounced. At both input SNR levels, the iPSD and iPSD-Zero outputs closely track the ground-truth in the time domain, with EMG-induced artifacts substantially suppressed. The corresponding PSDs further confirm that iPSD selectively attenuates broadband EMG components while recovering the underlying EEG spectrum. In contrast, Optimal-WT fails to recover the clean signal, introducing spurious spikes in the time domain that are absent from the ground-truth and causing noticeable distortion in the frequency domain. Overall, iPSD-Zero matches the performance of iPSD, making it well-suited for settings that demand immediate deployment without prior training.

K CROSS-SUBJECT GENERALIZABILITY

To investigate whether the learned partitioning strategy generalizes to unseen subjects, we conduct a leave-one-subject-out evaluation on five subjects from the CHB-MIT dataset. Each subject's recordings are corrupted with distinct EMG noise at an input SNR of -5 dB. In each fold, one subject is held out for evaluation, while the partitioning policy is learned using the remaining four subjects. The held-out subject and its associated noise are therefore unseen during training. As shown in Table 14, iPSD achieves consistent denoising performance across all five held-out subjects, with output SNRs ranging from 3.49 to 3.55 dB and an average of 3.52 dB. This indicates iPSD learns a general, transferable partitioning principle rather than overfitting to subjects or conditions.

Table 14: Leave-one-subject-out evaluation on five CHB-MIT subjects with EMG noise at an input SNR of -5 dB.

Held-out subject	#1	#2	#3	#4	#5
Output SNR (dB)	3.55	3.51	3.49	3.52	3.53