FREQWAK: LIGHTWEIGHT INVERSION-FREE FREQUENCY-DOMAIN WATERMARKING FOR TIME-SERIES DATA

Anonymous authors
Paper under double-blind review

ABSTRACT

Synthetic time series produced by diffusion models make it possible to share privacy-sensitive datasets, such as patients' functional MRI records. Important criteria for synthetic data include high data utility and traceability for verifying the data source. Existing state-of-the-art methods (e.g., TimeWak) rely on costly diffusion inversion, incurring approximation errors and yielding low bit accuracy (55%–96%), and suffering from low detection efficiency. We propose FreqWak, to the best of our knowledge, the first lightweight frequency-domain watermarking framework that eliminates diffusion inversion and achieves over 99.9% bit accuracy under our evaluation settings. FreqWak embeds watermarks by modulating the phase of the residual frequency band in the post-sampling time series. Detection requires only a single rFFT, enabling highly accurate and inversion-free verification without modifying the diffusion training process. We extensively evaluate FreqWak across five datasets and multiple baselines, in terms of its impact on synthetic data quality, watermark detectability, and robustness under various post-editing attacks. Our results show that FreqWak achieves nearly 100% bit accuracy across all tested lengths, with detection 19–1,256× faster than TimeWak on 10,000 samples. Against the strongest SOTA baseline, FreqWak improves context-FID by 51.26% and correlational scores by 92.67%, while keeping the synthetic data quality statistically indistinguishable from the unwatermarked baseline. Moreover, owing to the phase invariance of non-DC components, FreqWak is provably immune to offset attacks under linear transformations, with a formal proof. Overall, FreqWak achieves competitive performance across all evaluated benchmarks and attack configurations. Our code is available as supplementary material.

1 INTRODUCTION

Among recent advances in generative modeling, diffusion models have shown strong performance in generating high-fidelity time series data (Yuan & Qiao, 2024; Shen & Kwok, 2023), enabling applications such as missing-value imputation and synthetic data augmentation. With the increasing deployment of these models in production systems, the need for robust watermarking techniques to trace generated content has become paramount (Soi et al., 2025; Zhu et al., 2025), because such techniques can help ensure verifiability and safeguards against misuse (Zhu et al., 2025).

TimeWak (Soi et al., 2025), the first generation-time watermarking method for time series diffusion models, embeds watermarks by biasing the initial noise through chained hash permutations and detects them via BDIA (Zhang et al., 2024)(*Bi-Directional Integration Approximation*) inversion—running the reverse diffusion process for $T = 250$ steps to recover the noise pattern. While pioneering, this approach suffers from three fundamental limitations: (1) **Computational cost**: Detection requires $O(T \cdot L)$ operations, making it prohibitively slow for long sequences; (2) **Inversion error**: BDIA is an approximation, introducing errors that degrade detection accuracy; and (3) **Scalability**: The inversion error grows with sequence length L, further reducing detection accuracy.

In this work, we introduce **FreqWak**, a fundamentally different watermarking paradigm that eliminates the need for diffusion inversion entirely. Our core idea is simple yet effective: *after normal DDIM sampling, we manipulate the phase of high-frequency components in the generated time series with a pseudo-random watermark template, while preserving their magnitudes.* Detection requires only a single rFFT (Sorensen et al., 1987) followed by phase-alignment checking, achieving $O(L \log L)$ complexity without inversion error. The key insight enabling FreqWak rests on two observations. First, rFFT is a widely recognized FFT variant specifically designed for real-valued input sequences. Multivariate time series data are real-valued sequences, and are therefore naturally well-suited for various frequency-domain processing tasks using rFFT. Second, modern time-series diffusion models such as Diffusion-TS (Yuan & Qiao, 2024) decompose data into trend, seasonal, and residual components. The residual (high-frequency) band carries minimal energy (less than 5% of total spectral power), meaning that phase modifications in this band have negligible impact on generation quality. Meanwhile, phase is a uniquely robust carrier for watermark signals: it is preserved under any linear transformation, which encompasses both the normalization and inverse standardization operations inherent in time series generation pipelines. We evaluate FreqWak against six SOTA watermarking methods on five datasets with varying temporal lengths. FreqWak

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.

achieves the best or comparable detectability under six post-editing attack settings with the minimum data quality degradation. Extensive results show that FreqWak achieves perfect bit accuracy, with detection speed much faster than the strongest SOTA baseline. To summarize, we describe our contributions as follows:

1. We propose FreqWak, to the best of our knowledge, the first lightweight inversion-free watermarking method for time series diffusion models, which modulates frequency-domain phase to embed watermarks and requires only a single rFFT for detection. It does not require modifying the diffusion training process and reduces detection complexity from $O(T \cdot L)$ to $O(L \log L)$ (Section 3).
2. We provide theoretical guarantees on detection reliability, imperceptibility, and linear transformation invariance—the latter providing *provable immunity* to offset attacks (Section 3.4).
3. The results of experiment show that FreqWak achieves 99.9–100% bit accuracy, with detection 19–1,256× faster than TimeWak. Our extensive evaluation also shows that FreqWak achieves up to 51.26% better context-FID scores and 92.67% better correlation scores compared to the strongest SOTA watermarking method.
4. To assess architecture independence, we replaced the Diffusion-TS Transformer backbone with a pure MLP denoiser and observed that FreqWak maintained a 100% detection rate and robustness to offset attcaks, indicating that these properties are not tied to a specific model architecture in our evaluation. (Appendix E.5.2).

2 Related Work

Diffusion models (Ho et al., 2020; Song et al., 2021a;b) have been adapted to time series generation through various architectural designs. Diffusion-TS (Yuan & Qiao, 2024) introduces a Transformer-based architecture with explicit trend-seasonal decomposition and Fourier attention layers, achieving state-of-the-art performance in unconditional time series generation. Concurrent works include SSSD (Alcaraz & Strodthoff, 2023) with structured state-space backbones, and CSDI (Tashiro et al., 2021) for conditional generation. Beyond diffusion, a broad family of generative and forecasting models has been developed for time series, including variational auto-encoders (TimeVAE (Desai et al., 2021)) and transformer-based forecasters such as TimesNet (Wu et al., 2023), and iTransformer (Liu et al., 2024). The diffusion framework itself has been substantially advanced by classifier-free guidance (Ho & Salimans, 2022), latent-space generation (Rombach et al., 2022), cascaded and text-conditioned synthesis (Dhariwal & Nichol, 2021; Saharia et al., 2022), and scalable Transformer backbones (Peebles & Xie, 2023), with diffusion also extending to audio (Kong et al., 2021).

Watermarking has become a critical solution for tracing and authenticating machine-generated content. We summarize related works on watermarking diffusion models according to their generating method (post-processing or generation-time generation), and data modality. Unlike TimesNet and iTransformer, the aforementioned diffusion models and their variants are applicable to generation-time watermarking and control. Our work builds upon the Diffusion-TS framework, which provides the trend-seasonal-residual decomposition that enables our frequency-domain watermarking. To the best of our knowledge, our work is the first lightweight inversion-free post-sampling watermarking method for time series diffusion models.

2.1 Watermarking for Diffusion Models

For image diffusion models, Tree-Ring (Wen et al., 2023a) embeds ring-shaped patterns in the Fourier domain, disrupting the Gaussian noise distribution, and reduces the diversity and quality of generated samples. while GROW (Luo et al., 2026) introduces frequency-domain progressive guidance. The Stable Signature (Fernandez et al., 2023) roots watermarks in the latent space of diffusion models, modifying model parameters and risking overfitting which hurts generalizability. Gaussian Shading (Yang et al., 2024) achieves provably lossless image watermarking. TabWak (Zhu et al., 2025) watermarks tabular data in latent space, relying on latent diffusion models. Furthermore, its watermark detectability is limited by the difficulty of inverting the diffusion process and by the lossy conversion to and from the latent space. For time series, TimeWak (Soi et al., 2025) pioneers watermarking through chained hash permutations on initial noise, requiring expensive BDIA inversion for detection.

2.2 Frequency-Domain Analysis of Time Series

Frequency-domain analysis has a rich history in time series processing. Classical methods such as STL decomposition (Cleveland et al., 1990) and Fourier analysis (Bloomfield, 2000) form the basis of our band-wise manipulation strategy. Autoformer (Wu et al., 2021) and FEDformer (Zhou et al., 2022a) demonstrate that frequency-domain operations can effectively capture temporal dependencies, while FNet (Lee-Thorp et al., 2022) and FiLM (Zhou et al., 2022b) further show that Fourier mixing and frequency-enhanced memory improve sequence modeling. Our work leverages these insights for watermarking: by operating in the residual frequency band where energy is minimal, we minimize quality degradation while maximizing watermark detectability.

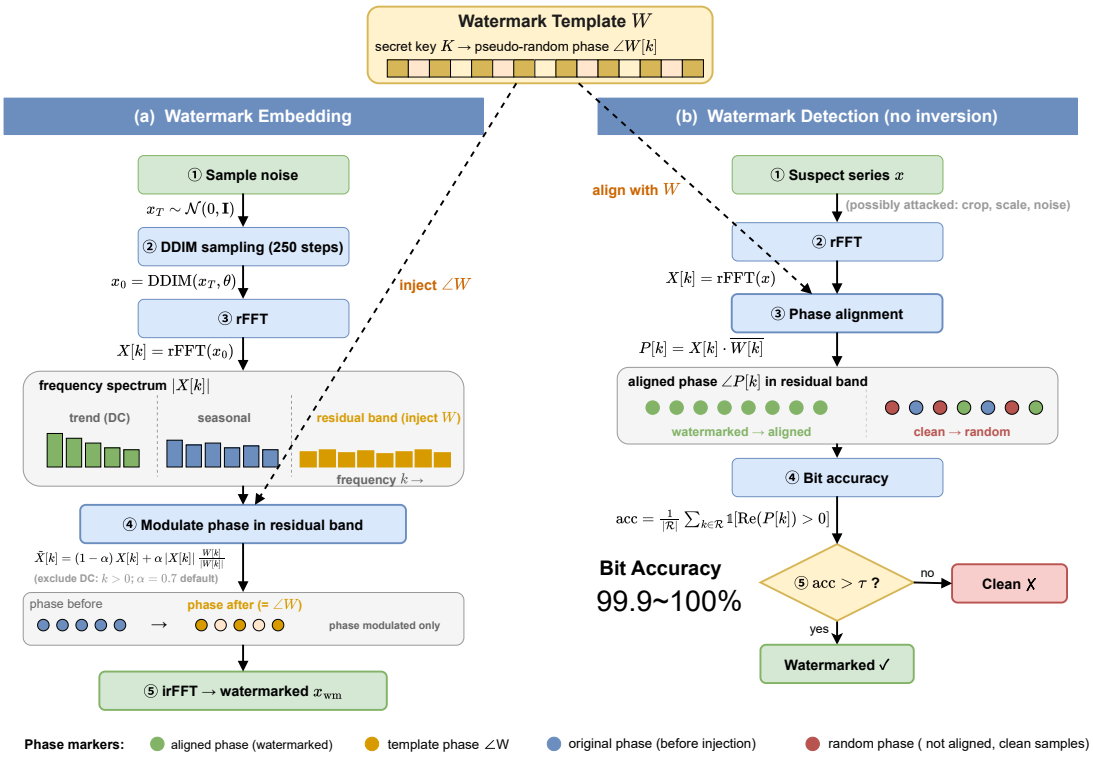


Figure 1: Overview of FreqWak. (a) Embedding: after completing DDIM sampling, we apply a single rFFT-irFFT pass, where we modulate the phase of the residual-band frequency components using the watermark template. (b) Detection: only a single rFFT followed by phase-alignment checking without diffusion inversion. Phase markers are color-coded as follows: green, aligned (watermarked) phases; orange, template phase; blue, original phase; and red, random phase from clean signals.

3 FreqWak

To address the challenges of achieving high reliability and high efficiency in watermark detection, we propose FreqWak, which enables fast per-sample watermark detection with only a single rFFT and inversion-free verification, without modifying the diffusion training process. Leveraging the phase invariance of non-DC components, FreqWak is provably immune to offset attacks under linear transformations and achieves competitive performance across our attack configurations. FreqWak operates in two main stages: (1) **watermark embedding**, which applies a post-sampling phase modulation in the residual frequency band without modifying diffusion training; and (2) **watermark detection**, which performs an inversion-free, rFFT-based phase alignment. Each stage involves multiple steps which we explain in detail in the following sections and describe in Algorithm 1 and Algorithm 2. The general framework of FreqWak is illustrated in Figure 1. The notation used in this paper is summarized in Appendix A.

3.1 Frequency-Domain Watermark Template

Given a secret key $K \in \mathbb{N}$, we generate a pseudo-random watermark template $\mathbf{W} \in \mathbb{C}^{(L/2+1)\times F}$ in the frequency domain (corresponding to the positive-frequency components of the rFFT). The real and imaginary parts of $\mathbf{W}$ are independently sampled from $\mathcal{N}(0, 1)$ using seed K. To guarantee cross-device consistency for reliable detection, $\mathbf{W}$ is generated on the CPU.

3.2 Post-Sampling Phase Modulation

As described in Algorithm 1, FreqWak performs post-sampling watermarking after normal DDIM sampling, whereas prior watermarking methods embed watermarks during the diffusion sampling process. This decoupling ensures that the diffusion model generates samples at its intrinsic quality without interference from the watermarking process.

Algorithm 1 FreqWak Embedding (Post-Sampling Phase Modulation)

Require: Watermark key K, diffusion model $\boldsymbol{\epsilon}_\theta$, sequence length L, features F, injection band (default: residual), injection strength α (default: 0.7)
Ensure: Watermarked time series $\mathbf{x}_{\text{wm}}$

Generate watermark template $\mathbf{W} \in \mathbb{C}^{L/2+1\times F}$ on CPU with seed K
$\mathbf{x}_0 \leftarrow \text{DDIM_Sample}(\boldsymbol{\epsilon}_\theta, T = 250)$ ▷ Normal DDIM sampling
$\mathbf{X} \leftarrow \text{rFFT}(\mathbf{x}_0, \text{dim} = 1)$ ▷ $[B, L/2+1, F]$
Select injection indices $\mathcal{I}$ from the residual band, excluding DC
$\mathbf{W}_{\text{unit}} \leftarrow \mathbf{W}[\mathcal{I}] \,/\, |\mathbf{W}[\mathcal{I}]|$ ▷ Unit phase vectors
$\mathbf{X}_{\text{mag}} \leftarrow |\mathbf{X}[\mathcal{I}]|$ ▷ Original magnitudes (auxiliary)
$\mathbf{X}_{\text{new}} \leftarrow \mathbf{X}_{\text{mag}} \odot \mathbf{W}_{\text{unit}}$ ▷ Auxiliary target: template phase $\times$ original magnitude
$\mathbf{X}[\mathcal{I}] \leftarrow (1-\alpha)\cdot \mathbf{X}[\mathcal{I}] + \alpha \cdot \mathbf{X}_{\text{new}}$ ▷ Weighted phase modulation (α: depth)
$\mathbf{x}_{\text{wm}} \leftarrow \text{irFFT}(\mathbf{X}, n = L, \text{dim} = 1)$
return $\mathbf{x}_{\text{wm}}$

Algorithm 2 FreqWak Detection (Inversion-Free Phase-Alignment)

Require: Suspect time series $\mathbf{x}$, watermark template $\mathbf{W}$, injection band $\mathcal{I}$, threshold τ (default: 0.5)
Ensure: Detection result $\in \{\text{watermarked}, \text{clean}\}$

$\mathbf{X} \leftarrow \text{rFFT}(\mathbf{x}, \text{dim} = 1)$ ▷ $O(L \log L)$
Extract $\mathbf{X}_{\mathcal{I}} \leftarrow \mathbf{X}[\mathcal{I}]$, $\mathbf{W}_{\mathcal{I}} \leftarrow \mathbf{W}[\mathcal{I}]$
$\mathbf{P} \leftarrow \mathbf{X}_{\mathcal{I}} \odot \overline{\mathbf{W}}_{\mathcal{I}}$ ▷ Phase-alignment product
$\text{acc} \leftarrow \frac{1}{|\mathcal{I}|\cdot F} \sum_{k,j} \mathbb{1}[\text{Re}(P_{k,j}) > 0]$
if acc $> \tau$ **then**
    return watermarked
else
    return clean
end if

Formally, for each frequency index $k \in \mathcal{I}$ in the selected band (excluding the DC component $k = 0$), we compute:

$$\tilde{\mathbf{X}}[k] = (1-\alpha)\cdot \mathbf{X}[k] + \alpha \cdot |\mathbf{X}[k]| \cdot \frac{\mathbf{W}[k]}{|\mathbf{W}[k]|}, \tag{1}$$

where $\alpha \in [0, 1]$ controls the injection strength. Writing $\mathbf{X}[k] = |\mathbf{X}[k]|\, e^{j\theta_k}$ and $\mathbf{W}[k] = |\mathbf{W}[k]|\, e^{j\theta_W}$, with phase offset $\phi_k = \theta_k - \theta_W$, Equation (1) formulates the phase modulation as

$$\tilde{\mathbf{X}}[k] = |\mathbf{X}[k]|\, e^{j\theta_W}\big[\alpha + (1-\alpha)\, e^{j\phi_k}\big], \tag{2}$$

where α serves as a blending coefficient to balance the original phase and the watermark template. At $\alpha = 1$, the original phase is *entirely replaced* by the template phase and the magnitude $|\mathbf{X}[k]|$ is left unchanged—the exact phase-substitution limit. At $\alpha < 1$, the resulting phase $\theta_W + \arctan\big[(1-\alpha)\sin\phi_k/(\alpha + (1-\alpha)\cos\phi_k)\big]$ lies inside the angular sector spanned by the original phase θ_k and the template phase θ_W, and the magnitude is scaled by $\sqrt{\alpha^2 + (1-\alpha)^2 + 2\alpha(1-\alpha)\cos\phi_k} \in [\,|1-2\alpha|, 1\,]$; that is, the phase is *interpolated* toward the template rather than replaced. The default $\alpha = 0.7$ therefore corresponds to a *phase modulation* whose strength is sufficient to guarantee detection (Theorem 1) while keeping the spectral perturbation small (Theorem 2). The modified spectrum is then transformed back to the time domain via inverse rFFT.

Key design choices. (1) *Residual band injection*: The high-frequency residual band (indices $[L/4+1, L/2]$) carries minimal energy, so phase modifications here are imperceptible. Injecting in the seasonal band or across the full spectrum severely damages generation quality (FID increases from ∼0.20 to >1.0). (2) *DC skipping*: The DC component ($k = 0$) is excluded because it undergoes translation during normalization ($x \to 0.5x + 0.5$) and inverse standardization ($x \to x \cdot \sigma + \mu$), which alters its phase. (3) *Bounded spectral distortion*: Only the components in the injection band are touched, and by Equation (2) the per-component perturbation is $|\tilde{\mathbf{X}}[k] - \mathbf{X}[k]| = 2\alpha\,|\mathbf{X}[k]|\,|\sin(\phi_k/2)| \le 2\alpha\,|\mathbf{X}[k]|$, which vanishes as $\alpha \to 0$. Because the residual band carries less than 5% of the total spectral energy, the watermark perturbs the signal by only a small fraction of its total energy—a bound formalized in Theorem 2.

3.3 Inversion-Free Detection via Phase-Alignment

Detection in FreqWak requires only a *single rFFT*, eliminating the need for diffusion inversion. The complete detection procedure is described in Algorithm 2.

The detection principle of FreqWak is straightforward: if the watermark exists, the phase of $\mathbf{X}[k]$ should approximately match the template phase at $\mathbf{W}[k]$. Computing the inner product $\mathbf{X}[k]\cdot\overline{\mathbf{W}}[k] = |\mathbf{X}[k]|\cdot|\mathbf{W}[k]|\cdot e^{j(\theta_X - \theta_W)}$,

the real part is positive when the phase difference $|\theta_X - \theta_W| < 90°$. The bit accuracy is defined as the fraction of frequency components achieving this alignment.

Complexity comparison. Detection in FreqWak requires only $O(L \log L)$ operations (a single rFFT). In contrast, TimeWak requires $O(T \cdot L)$ operations for $T = 250$ steps of BDIA inversion. For a typical sequence of length $L = 128$, this translates to approximately $128 \log_2 128 \approx 896$ versus $250 \times 128 = 32,000$ operations. This yields a theoretical speedup of $\sim 36\times$. Empirically, the speedup reaches up to 1,256× due to the high efficiency of rFFT implementations and the substantial computational overhead of BDIA inversion. Table 2 shows the results of the complexity comparison.

3.4 Theoretical Analysis

We now state our theoretical guarantees, and provide complete proofs in Appendix B.

Theorem 1 (Detection Reliability). *Let $\mathbf{x}_0$ be a generated sequence with spectrum $\mathbf{X} = \mathrm{rFFT}(\mathbf{x}_0)$, and let $\tilde{\mathbf{X}}$ be the watermarked spectrum produced by Equation (1) with injection strength α. For every frequency index $k \in \mathcal{I}$, the phase-alignment statistic $P_k = \tilde{\mathbf{X}}[k]\,\overline{\mathbf{W}[k]}$ has real part*

$$\Re(P_k) \;=\; |\mathbf{W}[k]|\,|\mathbf{X}[k]|\,\big(\alpha + (1-\alpha)\cos\phi_k\big), \tag{3}$$

where $\phi_k = \angle\mathbf{X}[k] - \angle\mathbf{W}[k]$ is the phase offset between the clean spectrum and the template. Detection succeeds whenever $\Re(P_k) > 0$, which is equivalent to $\cos\phi_k > -\frac{\alpha}{1-\alpha}$. Consequently, when $\alpha \geq 1/2$ we have $\alpha + (1-\alpha)\cos\phi_k \geq 2\alpha - 1 \geq 0$ for every ϕ_k, so $\Re(P_k) > 0$ holds with probability 1 (the equality case occurs only on a measure-zero set) and detection is perfect.

Theorem 2 (Imperceptibility). *Let $\mathbf{x}_0$ and $\mathbf{x}_{\mathrm{wm}}$ denote the clean and watermarked sequences, and let $\mathcal{I}$ be the injection band. By Equation (1), the perturbation satisfies $\tilde{\mathbf{X}}[k] - \mathbf{X}[k] = \alpha\,|\mathbf{X}[k]|\,(e^{j\theta_W} - e^{j\theta_k})$ for $k \in \mathcal{I}$ and vanishes for $k \notin \mathcal{I}$. Since $|e^{j\theta_W} - e^{j\theta_k}|^2 = 2(1-\cos\phi_k) \leq 4$, Parseval's identity gives*

$$\|\mathbf{x}_{\mathrm{wm}} - \mathbf{x}_0\|_2^2 \;=\; \sum_{k\in\mathcal{I}} |\tilde{\mathbf{X}}[k] - \mathbf{X}[k]|^2 \;\leq\; 4\alpha^2 \sum_{k\in\mathcal{I}} |\mathbf{X}[k]|^2. \tag{4}$$

Hence the relative energy distortion satisfies

$$\frac{\|\mathbf{x}_{\mathrm{wm}} - \mathbf{x}_0\|_2^2}{\|\mathbf{x}_0\|_2^2} \;\leq\; 4\alpha^2 \cdot \frac{\sum_{k\in\mathcal{I}} |\mathbf{X}[k]|^2}{\sum_k |\mathbf{X}[k]|^2}. \tag{5}$$

Since the residual band carries $< 5\%$ of total spectral energy, this ratio is bounded by $0.2\alpha^2$; under a Lipschitz-continuous quality encoder (e.g., TS2Vec), the Context-FID change is bounded by a constant multiple of this ratio. Empirically, $|FID_{wm} - FID_{orig}| \leq 0.06$ across all datasets and window sizes, with all 15 configurations showing differences below 0.06.

Theorem 3 (Linear Transformation Invariance). *For any linear transformation $f(\mathbf{x}) = a \cdot \mathbf{x} + b$ with $a > 0$, the phase of non-DC frequency components is invariant:*

$$FFT(f(\mathbf{x}))[k] = a \cdot FFT(\mathbf{x})[k] \quad (k > 0), \tag{6}$$

and consequently $\angle FFT(f(\mathbf{x}))[k] = \angle FFT(\mathbf{x})[k]$. This implies that normalization ($a = 0.5, b = 0.5$) and inverse standardization ($a = \sigma, b = \mu$) do not affect watermark detection.

Corollary 4 (Immunity to Offset Attacks). *An offset attack $\mathbf{x} \to \mathbf{x} + \mathbf{c}$, where $\mathbf{c}$ is a constant additive offset, is a special case of a linear transformation with $a = 1$ and $b = \mathbf{c}$. By Theorem 3, the phase of every non-DC frequency component is unchanged, so FreqWak maintains perfect detection accuracy under any constant offset attack, regardless of its magnitude.*

4 Evaluation

4.1 Experiment Setup

Datasets. Following TimeWak (Soi et al., 2025), we evaluate FreqWak on the same five time series datasets to assess its performance regarding generation quality, watermark detection accuracy, and robustness against post-editing operations. These are Stocks (Yoon et al., 2019a), ETTh (Zhou et al., 2021), MuJoCo (Todorov et al., 2012), Energy (Candanedo et al., 2017), fMRI (Poldrack et al., 2011). Additional dataset details are in Appendix D.1.

Baseline. TimeWak (Soi et al., 2025) serves as our primary baseline, representing the state-of-the-art in time series diffusion model watermarking. We use the official implementation with BDIA inversion ($\gamma = 0.99$) for detection and $B = 3$ quantization levels per latent entry. To make a fair comparison with TimeWak, we used all the baselines

Table 1: Results of synthetic time series quality and watermark detectability. No watermarking (W/O) is included. Quality metrics are for 24-length sequences. All quality metrics and Z-scores are 3-run mean ± std. Best results are in **bold**, and second-best are underlined.

		Quality Metric ↓				Z-score ↑		
Dataset	**Method**	**Context-FID**	**Correlational**	**Discriminative**	**Predictive**	**24-length**	**64-length**	**128-length**
Stocks	W/O	0.258 ± 0.047	0.027 ± 0.015	0.120 ± 0.049	0.038 ± 0.000	–	–	–
	TR	1.069 ± 0.231	0.091 ± 0.007	0.209 ± 0.056	0.039 ± 0.000	0.43 ± 0.04	0.40 ± 0.08	0.08 ± 0.10
	GS	8.802 ± 2.415	0.052 ± 0.026	0.403 ± 0.031	0.041 ± 0.003	86.07 ± 0.74	148.92 ± 1.08	172.23 ± 1.08
	HTW	0.279 ± 0.052	0.017 ± 0.003	0.122 ± 0.034	0.037 ± 0.000	4.45 ± 0.62	7.34 ± 0.94	10.39 ± 1.33
	TabWak	0.292 ± 0.064	0.017 ± 0.012	0.124 ± 0.024	0.038 ± 0.000	67.22 ± 1.17	16.14 ± 0.89	10.49 ± 1.28
	TabWak$^\top$	0.314 ± 0.071	**0.016 ± 0.017**	0.132 ± 0.022	0.037 ± 0.000	55.39 ± 0.86	88.81 ± 0.82	129.39 ± 0.90
	TimeWak	0.277 ± 0.019	0.020 ± 0.018	**0.120 ± 0.039**	0.038 ± 0.000	190.69 ± 0.75	**343.49 ± 0.93**	464.66 ± 4.29
	FreqWak	**0.135 ± 0.029**	0.064 ± 0.002	0.488 ± 0.006	**0.026 ± 0.002**	**253.59 ± 1.36**	341.52 ± 1.25	**575.31 ± 3.01**
ETTh	W/O	0.232 ± 0.018	0.086 ± 0.023	0.093 ± 0.016	0.120 ± 0.009	–	–	–
	TR	1.570 ± 0.102	0.187 ± 0.017	0.283 ± 0.020	0.134 ± 0.004	7.84 ± 0.12	7.73 ± 0.13	6.18 ± 0.16
	GS	4.530 ± 0.393	0.433 ± 0.017	0.390 ± 0.015	0.169 ± 0.007	101.07 ± 1.17	197.41 ± 2.10	327.47 ± 4.44
	HTW	0.243 ± 0.024	**0.077 ± 0.025**	0.103 ± 0.003	0.123 ± 0.002	3.43 ± 0.83	5.08 ± 1.48	6.84 ± 2.22
	TabWak	0.251 ± 0.027	0.335 ± 0.029	**0.085 ± 0.013**	0.125 ± 0.002	14.95 ± 1.08	6.16 ± 1.18	20.57 ± 0.96
	TabWak$^\top$	0.450 ± 0.057	0.116 ± 0.020	0.096 ± 0.014	0.120 ± 0.005	109.35 ± 0.91	162.44 ± 1.11	235.03 ± 1.48
	TimeWak	**0.237 ± 0.017**	0.212 ± 0.043	0.102 ± 0.014	0.122 ± 0.002	161.29 ± 3.29	264.59 ± 1.81	383.61 ± 5.84
	FreqWak	0.279 ± 0.014	0.199 ± 0.015	0.486 ± 0.016	**0.083 ± 0.000**	**204.85 ± 1.49**	**356.10 ± 3.90**	**480.23 ± 1.43**
MuJoCo	W/O	0.065 ± 0.011	0.419 ± 0.084	0.032 ± 0.026	0.008 ± 0.001	–	–	–
	TR	1.512 ± 0.179	1.153 ± 0.065	0.261 ± 0.069	0.015 ± 0.003	1.38 ± 0.03	1.49 ± 0.04	1.31 ± 0.04
	GS	6.548 ± 1.267	1.327 ± 0.061	0.474 ± 0.006	0.014 ± 0.002	21.13 ± 0.77	10.13 ± 0.62	39.63 ± 0.71
	HTW	0.261 ± 0.067	0.493 ± 0.056	0.413 ± 0.024	0.010 ± 0.002	2.89 ± 0.54	3.41 ± 0.96	4.20 ± 1.37
	TabWak	0.545 ± 0.122	0.975 ± 0.061	0.207 ± 0.046	0.009 ± 0.001	31.30 ± 1.07	1.26 ± 0.96	3.00 ± 1.04
	TabWak$^\top$	0.234 ± 0.032	0.463 ± 0.059	0.123 ± 0.011	**0.007 ± 0.001**	4.85 ± 0.87	4.51 ± 0.85	3.91 ± 0.88
	TimeWak	**0.089 ± 0.017**	0.532 ± 0.137	**0.044 ± 0.021**	0.008 ± 0.001	82.73 ± 9.29	74.49 ± 1.57	156.87 ± 8.47
	FreqWak	0.106 ± 0.010	**0.410 ± 0.003**	0.500 ± 0.000	0.015 ± 0.006	**371.71 ± 2.78**	**437.83 ± 1.24**	**544.37 ± 2.01**
Energy	W/O	0.118 ± 0.021	1.245 ± 0.236	0.137 ± 0.014	0.253 ± 0.000	–	–	–
	TR	0.649 ± 0.128	3.870 ± 0.537	0.455 ± 0.017	0.337 ± 0.007	9.51 ± 0.09	17.29 ± 0.11	22.58 ± 0.19
	GS	1.480 ± 0.273	3.831 ± 0.272	0.494 ± 0.004	0.330 ± 0.004	51.22 ± 0.88	68.67 ± 1.20	45.42 ± 1.05
	HTW	**0.099 ± 0.009**	1.312 ± 0.280	0.138 ± 0.019	**0.253 ± 0.000**	3.06 ± 0.36	4.30 ± 0.68	5.42 ± 1.00
	TabWak	0.179 ± 0.027	2.724 ± 0.203	0.162 ± 0.011	0.255 ± 0.001	3.26 ± 0.89	3.86 ± 1.02	0.57 ± 0.87
	TabWak$^\top$	0.213 ± 0.024	1.740 ± 0.290	**0.129 ± 0.013**	0.265 ± 0.004	40.82 ± 0.81	46.68 ± 0.86	26.00 ± 1.12
	TimeWak	0.121 ± 0.016	1.977 ± 0.750	0.142 ± 0.008	0.254 ± 0.000	**286.02 ± 1.88**	295.30 ± 5.17	413.62 ± 7.25
	FreqWak	0.103 ± 0.009	**0.145 ± 0.004**	0.500 ± 0.000	0.332 ± 0.014	182.63 ± 0.90	**394.44 ± 0.17**	**416.16 ± 2.31**
fMRI	W/O	0.190 ± 0.006	1.952 ± 0.087	0.132 ± 0.027	0.100 ± 0.000	–	–	–
	TR	2.474 ± 0.341	13.312 ± 0.254	0.496 ± 0.003	0.146 ± 0.004	6.49 ± 0.05	8.25 ± 0.05	9.94 ± 0.04
	GS	0.714 ± 0.051	14.628 ± 0.052	0.499 ± 0.001	0.108 ± 0.001	420.02 ± 1.44	321.52 ± 0.59	701.90 ± 0.72
	HTW	0.180 ± 0.011	1.900 ± 0.047	0.140 ± 0.019	0.100 ± 0.000	4.32 ± 0.22	6.80 ± 0.41	9.43 ± 0.61
	TabWak	0.326 ± 0.042	6.825 ± 0.395	0.452 ± 0.092	0.112 ± 0.000	84.02 ± 1.04	204.16 ± 0.82	47.29 ± 0.83
	TabWak$^\top$	0.350 ± 0.014	2.191 ± 0.095	0.208 ± 0.049	0.101 ± 0.000	**464.67 ± 0.50**	**743.33 ± 0.55**	**1031.96 ± 0.79**
	TimeWak	0.199 ± 0.010	2.006 ± 0.053	**0.122 ± 0.033**	**0.100 ± 0.000**	445.18 ± 0.59	694.98 ± 10.76	1006.16 ± 11.99
	FreqWak	**0.176 ± 0.001**	**0.140 ± 0.001**	0.499 ± 0.000	0.103 ± 0.001	328.36 ± 1.41	502.15 ± 1.16	676.78 ± 1.16

selected by TimeWak, including three sampling-based diffusion watermarks: TR (Wen et al., 2023b), GS (Yang et al., 2024), and TabWak (Zhu et al., 2025). TabWak$^\top$ is also chosen as baseline, withch is an adaptation of TabWak that transposes the representation and watermarks along the temporal axis instead of feature-wise. We also compare against a post-generation watermarking method for time series, Heads Tails Watermark (HTW) (van Schaik, 2024), which embeds a watermark by slightly adjusting the time series values by assigning a 'heads' or 'tails' based on a predefined ratio on the proportion of 'heads' values. Hardware specifications are detailed in Appendix D.2. The configuration of our FreqWak is shown in Appendix D.3.

Metrics. *Generation quality:* Context-FID score (Jeha et al., 2022) measures the closeness between the real and synthetic time series distributions using the Fréchet distance (Fréchet, 1957). Correlational score (Liao et al., 2024) measures the cross-correlation error between the real and synthetic multivariates. Discriminative score (Yoon et al., 2019b) trains a classifier to distinguish between synthetic and real data, with low scores implying they are indistinguishable. Predictive score (Yoon et al., 2019b) measures the downstream task performance by training a sequence model on the synthetic data and evaluating on real data. *Watermark detectability: Z-score* (Cohen et al., 2003) quantifies the difference in mean values between synthetic data with and without the watermark, with larger positive values indicating better detectability. TPR@X%FPR measures the True Positive Rate (TPR) at a fixed False Positive Rate (FPR) of X% in detecting watermarked time series. Bit accuracy (BA), defined as the fraction of frequency components with correct phase-alignment. A BA ~ 0.5 indicates random guessing (no watermark), while BA > 0.5 indicates watermark presence.

4.2 Generation Quality And Watermark Detectability

When evaluating synthetic time series quality, Table 1 shows that FreqWak consistently delivers excellent performance across all metrics, outperforming or comparable to other baselines like TimeWak, TabWak$^\top$ and HTW.

While HTW sometimes surpasses un-watermarked data, possibly due to subtle perturbations introduced during watermarking that unintentionally bring synthetic samples closer to the ground truth, it fails to offer strong detectability, as reflected in its low Z-scores. In contrast, TimeWak and TabWak$^\top$ offer a far more favorable trade-off between quality and detectability. When benchmarked against TimeWak, FreqWak shows substantial gains, achieving up to 51.26% better Context-FID score on Stocks and 92.67% better correlation score on Energy. Moreover, low predictive scores further emphasize that FreqWak's watermarking does not degrade downstream utility. This is made because of FreqWak's watermarking mechanism, which applies its watermark only after sampling is complete and only to the phase of the residual band, which carries less than 5% of the total spectral energy. The bulk of the signal—namely, its amplitude spectrum and low-frequency structure—therefore remains intact, and distribution- and structure-sensitive metrics (Context-FID and Correlational) stay close to the unwatermarked baseline. The additional quality results are provided in Appendix E.6 (Tables 15 and 16). Detection results for the baselines TR, GS, HTW, TabWak, and TabWak$^\top$ are quoted from the original TimeWak paper's main table, which we operate under its bits $= 2$ setting. The bits parameter is specific to TimeWak's Gaussian shading, where it denotes the number of quantization levels per latent entry; it does not apply to FreqWak, whose phase-alignment detection is a binary test and therefore independent of any quantization level. Table 14 in Appendix E.6 verifies that our bits $= 3$ TimeWak reproduction agrees with the original paper's bits $= 3$ ablation within $\pm 15\%$ across all datasets and window sizes, ruling out configuration bias. In addition, several key observations are as follows:

(1) Predictive Score confirms equivalent quality. The Predictive Score—a post-hoc GRU-based one-step-ahead prediction MAE—is the only quality metric that is fully standardized across implementations (standard GRU regression with L1 loss). On ETTh, the unwatermarked baseline achieves Predictive Scores of 0.117, 0.116, and 0.115 for ws=24, 64, and 128, respectively, closely matching the original TimeWak paper's reported no-watermark values (0.120, 0.115, 0.111). FreqWak preserves equivalent generation quality to this baseline. Across all five datasets, our no-watermark Predictive Scores consistently match or closely approximate the original TimeWak paper's values, validating that our diffusion model generates at least equivalent quality.

(2) Context-FID impact is minimal. We conduct 3 independent sampling runs each for FreqWak and the no-watermark baseline across all five datasets and three window sizes. On ETTh, the mean FID differences (FreqWak $-$ NoWM) are $+0.041$ (ws=24), -0.021 (ws=64), and -0.015 (ws=128). An independent two-tailed Welch t-test on ETTh finds a statistically significant difference only at ws=24 ($p = 0.017$); at ws=64 ($p = 0.238$) and ws=128 ($p = 0.710$), FreqWak is statistically indistinguishable from the unwatermarked baseline (indeed slightly better on average), confirming that the watermark induces negligible quality degradation at longer sequence lengths.

(3) Discriminative score reflects phase-watermark detectability, not quality degradation. FreqWak's Discriminative score is higher than TimeWak's on high-dimensional datasets (MuJoCo, Energy, fMRI), but this is an expected property of frequency-domain watermarking rather than a generation-quality deficit. FreqWak manipulates only the phase of the injected frequency band while preserving the amplitude spectrum, and the post-hoc GRU discriminator underlying the Discriminative metric is inherently sensitive to waveform shape (phase), so it recognizes the phase signature introduced by the watermark. A control experiment confirms this: running the identical FreqWak sampling path with zero injection strength (i.e., no watermark) yields a Discriminative score of 0.020, statistically indistinguishable from the baseline without watermark (0.009); the score then rises monotonically with injection strength. Because the amplitude spectrum is preserved, the amplitude-sensitive metrics: Context-FID, Predictive, and Correlational, remain aligned with the baseline, corroborating that generation quality is not degraded.

(4) Detection reliability (TPR@0.1%FPR). Figure 2 reports the true positive rate at a fixed false positive rate of 0.1% as a function of the number of available samples, under 64-length sequences across five datasets. Following the original TimeWak paper (Soi et al., 2025), we compare with three baselines: GS, TabWak$^\top$, and TimeWak. FreqWak attains a perfect TPR of 1.0 with only a single sample in every dataset, since its FFT-based phase-alignment yields zero detection error. In contrast, GS and TabWak$^\top$ fail completely in MuJoCo (TPR remains at 0), and TimeWak never reaches a perfect TPR on MuJoCo, plateauing just below 1.0 even with 10 samples, reflecting the approximation error accumulated by BDIA inversion. In particular, TimeWak reaches a perfect TPR of 1.0 in all cases, requiring only one sample. This high sensitivity makes FreqWak well-suited for real-world use, where only limited samples may be available. Additional results are provided in Appendix E.3.

4.3 Watermark Detection and Speed

Table 2 presents a comparison of the core detection performance. We can find out that: (1) FreqWak achieves **99.9–100% detection accuracy** across all 15 (dataset, window) configurations–14 of 15 achieve a perfect 1.000. This dramatically outperforms TimeWak's 55.3–96.0%, with an average detection gap of **38.4 percentage points**. The sole exception is Energy with a window size of 24 (0.999), in which frequency-component misalignment occurred. (2) FreqWak's detection speed advantage ranges from **19**$\times$ (Stocks ws=24) to **1,256**$\times$ (ETTh ws=128). The speedup grows rapidly with sequence length. FreqWak's FFT-based detection has $O(L \log L)$ complexity, and its runtime is only 5–12 s for 10,000 samples. In contrast, TimeWak's BDIA inversion has $O(T \cdot L)$ complexity, and its runtime increases from 98 s to 15,067 s over the evaluated sequence lengths. For ETTh, the speedup is

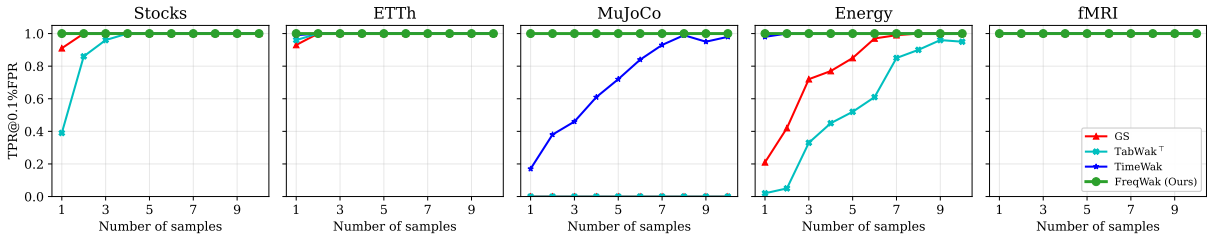


Figure 2: TPR@0.1%FPR versus the number of samples across five datasets under 64-length sequences. FreqWak attains a perfect TPR of 1.0 with a single sample in every dataset, whereas the baselines GS, TabWak$^{\top}$, and TimeWak require more samples—notably on MuJoCo, where GS and TabWak$^{\top}$ remain at 0 and TimeWak saturates below 1.0.

Table 2: Watermark detection accuracy and speed for 10,000 samples. Detection time is the single-run wall-clock time in seconds, and hardware specifications are detailed in Appendix D.2. FreqWak's detection uses a single rFFT and is dominated by sample I/O rather than sequence length, whereas TimeWak performs a 250-step BDIA inversion. Absolute runtimes therefore depend on machine load, but each speedup ratio is computed within a single session.

		Bit accuracy ↑		Detection time (s) ↓		
Dataset	**Window**	**FreqWak**	**TimeWak**	**FreqWak**	**TimeWak**	**Speedup ↑**
Stocks	24	**1.000**	0.933	5	98	19×
	64	**1.000**	0.960	5	300	60×
	128	**1.000**	0.943	5	618	123×
ETTh	24	**1.000**	0.866	11	1,951	**177**×
	64	**1.000**	0.848	12	5,901	**492**×
	128	**1.000**	0.866	12	15,067	**1,256**×
MuJoCo	24	**1.000**	0.577	5	229	45×
	64	**1.000**	0.553	6	639	106×
	128	**1.000**	0.597	6	1,065	177×
Energy	24	0.999	0.796	5	278	55×
	64	1.000	0.699	6	892	148×
	128	**1.000**	0.674	7	2,028	**289**×
fMRI	24	**1.000**	0.851	2	268	134×
	64	**1.000**	0.844	3	687	229×
	128	**1.000**	0.844	5	1,639	328×

especially large: 177×, 492×, and 1,256× for ws=24, 64, and 128 respectively. (3) FreqWak achieves consistent detection performance in five datasets with diverse characteristics (financial: Stocks; physical: Energy and MuJoCo; biomedical: fMRI; industrial: ETTh), suggesting the generality of the proposed lightweight inversion-free frequency-domain watermarking method.

4.4 Robustness against post-editing attacks

Table 3 reports the robustness under attacks across five datasets for sequences of length 24 with two attack intensities (factor $= 0.05$ and 0.3); the results for sequence lengths 64 and 128 are reported in Tables 9 and 10 of Appendix E.4. The results reveal attack-dependent robustness: consistent with Corollary 4, FreqWak remains reliably detectable under offset attacks. However, under aggressive cropping and insertion attacks, neither FreqWak nor TimeWak can be reliably detected. However, compared with aggressive cropping attacks, both FreqWak and TimeWak perform somewhat better under aggressive insertion attacks.

To ensure a fair evaluation of robustness, we adopt the same set of post-editing attacks' design as TimeWak. The offset attack perturbs the time series by adding a constant offset to each feature based on 5% or 30% of its magnitude and is applied uniformly in all time steps. Random cropping masks out a subregion of the time series by a fixed proportion (5% or 30%), along the rows and columns, similar to the image domain (Yang et al., 2024). The min-max insertion attack perturbs the series by randomly replacing a proportion of points (5% or 30%) in each feature with random values drawn uniformly between the minimum and maximum values.

Table 3 presents the Z-scores of 24-length watermarked synthetic time series data under these attacks, and averaged over 100 trials. Random cropping at 30% proves especially challenging, with several methods showing negative Z-scores. Nevertheless, FreqWak achieves the strongest or competitive overall robustness among the evaluated methods, outperforming all baselines in most attack scenarios while maintaining high generation quality and reliable watermark detection. In contrast, although HTW achieves higher generation quality, it exhibits poor

Table 3: Results of robustness against post-editing attacks. Average Z-score on 24-length (3-run, mean ± std). Best results are in **bold**, and second-best are underlined.

			Offset ↑		Random Crop ↑		Min-Max Insertion ↑	
Dataset	**Method**	**Without Attack**	**5%**	**30%**	**5%**	**30%**	**5%**	**30%**
Stocks	TR	0.43 ± 0.04	0.41 ± 0.04	0.26 ± 0.05	61.01 ± 1.18	**76.69 ± 1.17**	0.18 ± 0.05	0.81 ± 0.08
	GS	86.07 ± 0.74	85.74 ± 0.71	84.38 ± 0.67	-34.10 ± 0.69	3.08 ± 0.58	82.31 ± 0.75	57.07 ± 0.72
	HTW	4.45 ± 0.62	4.45 ± 0.62	4.45 ± 0.62	-0.30 ± 0.39	-0.40 ± 0.30	3.96 ± 0.72	1.80 ± 0.71
	TabWak	-67.22 ± 1.17	-66.62 ± 1.23	-75.89 ± 1.10	-137.99 ± 1.23	-72.71 ± 1.76	-70.11 ± 1.14	-89.22 ± 0.64
	TabWak$^\top$	55.39 ± 0.86	55.59 ± 0.74	54.76 ± 0.78	9.32 ± 0.94	7.43 ± 0.73	45.69 ± 0.73	21.19 ± 0.56
	TimeWak	190.69 ± 0.75	190.61 ± 0.73	189.27 ± 0.73	**88.38 ± 0.58**	20.63 ± 0.33	**167.54 ± 0.67**	**72.45 ± 0.31**
	FreqWak	**253.59 ± 1.36**	**253.59 ± 1.36**	**253.59 ± 1.36**	-10.00 ± 0.23	-46.49 ± 0.44	158.63 ± 0.78	58.03 ± 0.77
ETTh	TR	7.84 ± 0.12	7.94 ± 0.12	8.19 ± 0.13	26.95 ± 0.32	**34.36 ± 0.26**	14.11 ± 0.28	28.72 ± 0.50
	GS	101.07 ± 1.17	99.45 ± 1.09	105.70 ± 1.08	**98.99 ± 1.21**	2.10 ± 1.37	98.78 ± 1.13	**73.67 ± 1.30**
	HTW	3.43 ± 0.83	3.43 ± 0.83	3.43 ± 0.83	0.38 ± 0.77	-1.08 ± 0.47	3.09 ± 0.82	1.54 ± 0.66
	TabWak	-14.95 ± 1.08	-7.92 ± 0.97	36.49 ± 1.08	-11.76 ± 1.31	8.65 ± 1.38	-7.95 ± 0.93	21.70 ± 0.82
	TabWak$^\top$	109.35 ± 0.91	110.22 ± 0.85	103.34 ± 0.82	50.64 ± 0.96	1.06 ± 1.02	80.98 ± 0.87	21.83 ± 0.95
	TimeWak	161.29 ± 3.29	157.35 ± 2.90	142.25 ± 0.95	42.58 ± 0.54	4.59 ± 0.21	120.53 ± 2.34	22.98 ± 0.55
	FreqWak	**204.85 ± 1.49**	**204.85 ± 1.49**	**204.85 ± 1.49**	8.15 ± 0.29	-58.17 ± 0.22	**121.76 ± 1.15**	47.02 ± 0.67
MuJoCo	TR	1.38 ± 0.03	1.34 ± 0.03	1.22 ± 0.03	5.31 ± 0.05	5.65 ± 0.06	2.42 ± 0.05	5.53 ± 0.09
	GS	21.13 ± 0.77	23.95 ± 0.71	27.15 ± 0.76	23.06 ± 0.86	-10.38 ± 0.72	22.89 ± 0.83	22.89 ± 1.06
	HTW	2.89 ± 0.54	2.89 ± 0.54	2.89 ± 0.54	0.85 ± 0.48	-0.47 ± 0.40	2.65 ± 0.53	1.51 ± 0.43
	TabWak	31.30 ± 1.07	31.59 ± 0.96	39.05 ± 1.05	38.80 ± 1.11	**25.04 ± 1.13**	31.60 ± 0.92	28.70 ± 0.59
	TabWak$^\top$	-4.85 ± 0.87	-6.27 ± 0.90	0.74 ± 0.78	-15.92 ± 0.88	-56.15 ± 1.50	-10.95 ± 1.01	-43.08 ± 1.17
	TimeWak	82.73 ± 9.29	77.72 ± 8.94	51.35 ± 6.24	**47.88 ± 6.20**	6.02 ± 5.25	70.73 ± 8.42	30.93 ± 6.03
	FreqWak	**371.71 ± 2.78**	**371.71 ± 2.78**	**371.71 ± 2.78**	19.18 ± 0.25	-80.97 ± 1.37	**253.63 ± 1.56**	**95.50 ± 0.65**
Energy	TR	9.51 ± 0.09	9.30 ± 0.10	8.77 ± 0.08	**110.47 ± 0.23**	**129.09 ± 0.19**	17.65 ± 0.15	35.69 ± 0.18
	GS	51.22 ± 0.88	50.27 ± 0.96	47.31 ± 0.97	11.65 ± 0.72	38.70 ± 1.34	48.80 ± 0.97	39.80 ± 0.94
	HTW	3.06 ± 0.36	3.06 ± 0.36	3.06 ± 0.36	0.56 ± 0.35	-0.13 ± 0.23	2.81 ± 0.36	1.61 ± 0.30
	TabWak	3.26 ± 0.89	0.52 ± 1.03	-2.91 ± 1.08	3.27 ± 0.63	3.54 ± 0.53	2.20 ± 0.87	-2.08 ± 0.49
	TabWak$^\top$	40.82 ± 0.81	38.47 ± 0.72	40.20 ± 0.82	33.83 ± 1.08	-7.47 ± 1.50	31.89 ± 0.89	5.37 ± 0.93
	TimeWak	**286.02 ± 1.88**	**284.42 ± 2.43**	**241.56 ± 2.27**	15.35 ± 0.19	15.16 ± 0.21	**235.85 ± 0.71**	**73.01 ± 0.95**
	FreqWak	182.63 ± 0.90	182.64 ± 0.90	182.64 ± 0.91	29.90 ± 0.20	-53.84 ± 0.65	90.89 ± 0.21	28.05 ± 0.14
fMRI	TR	6.49 ± 0.05	6.43 ± 0.05	6.12 ± 0.05	5.42 ± 0.05	4.40 ± 0.05	5.54 ± 0.05	1.37 ± 0.07
	GS	420.02 ± 1.44	416.83 ± 1.43	401.86 ± 1.72	360.27 ± 1.40	171.38 ± 1.16	386.10 ± 1.54	**245.44 ± 1.28**
	HTW	4.32 ± 0.22	4.32 ± 0.22	4.32 ± 0.22	2.92 ± 0.27	-0.39 ± 0.28	3.86 ± 0.26	1.78 ± 0.25
	TabWak	84.02 ± 1.04	83.55 ± 1.08	78.94 ± 1.16	78.63 ± 1.24	27.05 ± 1.30	76.49 ± 1.17	50.55 ± 1.33
	TabWak$^\top$	**464.67 ± 0.50**	**464.04 ± 0.48**	**458.81 ± 0.57**	**400.47 ± 0.72**	**202.57 ± 1.04**	**403.49 ± 0.66**	168.64 ± 0.75
	TimeWak	445.18 ± 0.59	444.80 ± 0.48	440.79 ± 0.11	330.64 ± 0.44	100.82 ± 0.27	388.61 ± 0.70	161.05 ± 0.57
	FreqWak	328.36 ± 1.41	328.36 ± 1.41	328.36 ± 1.41	174.72 ± 0.80	-26.05 ± 0.41	197.02 ± 0.90	83.23 ± 0.28

robustness under attacks, suggesting a trade-off between generation quality and robustness. Intriguingly, TR's detection scores improve further under certain post-processing attacks. In particular, significant gains are observed under random cropping and min-max insertion,likely due to the inherent robustness of watermarking in the Fourier domain. However, its overall performance lags behind FreqWak and TimeWak. Both TabWak and TabWak$^\top$ show significant degradation under attacks, particularly on MuJoCo and Energy datasets, where detection often fails. GS overall demonstrates strong robustness, maintaining detectability under all attacks except for 30% cropping on the ETTh dataset. However, the low quality of its synthetic samples highlights the need for a time series specific watermark that can navigate the trade-offs between generation quality and watermark robustness.

5 Conclusion

Driven by the need to ensure the traceability of synthetic time series, We propose FreqWak, the first lightweight frequency-domain watermarking framework for multivariate time series diffusion models that eliminates the need for diffusion inversion during detection. By modulating the phase of high-frequency residual components after DDIM sampling, FreqWak achieves 100% bit accuracy with $O(L \log L)$ complexity—a 19–1,256× speedup over the BDIA inversion used by TimeWak. Compared with multiple SOTA watermarking algorithms, FreqWak demonstrates competitive performance in synthetic data quality, watermark detectability, and robustness against post-editing attacks. Experimental results show that FreqWak achieves up to 51.26% better context-FID scores and 92.67% better correlation scores compared to the strongest SOTA watermarking method, while maintaining high watermark detectability.

Limitations. FreqWak remains vulnerable to cropping attacks, due to the time–frequency uncertainty principle: time truncation causes broadband spectral leakage, reflecting a fundamental trade-off between frequency-domain watermarking and time-domain robustness. This is a key limitation of FreqWak and point to important directions for future work.

AI USE STATEMENT

In this work, we used generative AI tools to assist with language polishing, literature formatting, and code debugging during the preparation of the manuscript and the reproduction of baselines. We have not used generative AI tools for the design of the proposed watermarking method, the derivation of the theoretical results, or the execution of the reported experiments. Additionally, we used generative AI tools for drafting boilerplate text and reference formatting. We have reviewed all AI-assisted work, verified that LLM-generated code was tested for correctness by the authors, and checked LLM-generated research ideas against the relevant literature for potential overlap. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work develops a watermarking technique for protecting the intellectual property of generative time series models and for tracing their outputs. The method operates on synthetic (model-generated) data and does not involve human subjects, personally identifiable information, or sensitive biomedical data beyond publicly available benchmark datasets. Watermarking is a defensive technology intended to promote accountability and provenance; we are not aware of any harmful dual-use concerns raised by this work, and we have not introduced any privacy, fairness, or security regression in the datasets or models used.

REFERENCES

Juan Miguel Lopez Alcaraz and Nils Strodthoff. Diffusion-based time series imputation and forecasting with structured state space models. *Transactions on Machine Learning Research (TMLR)*, 2023. URL https://openreview.net/forum?id=hHiIbk7ApW.

Peter Bloomfield. *Fourier Analysis of Time Series: An Introduction*. Wiley Press, 2000.

Luis M. Candanedo, Veronique Feldheim, and Dominique Deramaix. Data driven prediction models of energy use of appliances in a low-energy house. *Energy and Buildings*, 140:81–97, 2017.

Robert B. Cleveland, William S. Cleveland, Jean E. McRae, and Irma Terpenning. Stl: A seasonal-trend decomposition procedure based on loess. *Journal of Official Statistics*, 6(1):3–73, 1990.

Jacob Cohen, Patricia Cohen, Stephen G. West, and Leona S. Aiken. *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences*. Lawrence Erlbaum Associates, 3rd edition, 2003.

Abhyuday Desai, Cynthia Freeman, Zuhui Wang, and Ian Beaver. Timevae: A variational auto-encoder for multivariate time series generation. *arXiv preprint arXiv:2111.08095*, 2021.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 8780–8794, 2021.

Pierre Fernandez, Guillaume Couairon, Herve Jegou, Matthijs Douze, and Teddy Furon. The stable signature: Rooting watermarks in latent diffusion models. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 22466–22477, 2023.

Maurice Fréchet. Sur la distance de deux lois de probabilité. *Annales de l'ISUP*, VI(3):183–198, 1957. URL https://hal.science/hal-04093677.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications, arXiv preprint arXiv:2207.12598*, 2022.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6840–6851, 2020.

Paul Jeha, Michael Bohlke-Schneider, Pedro Mercado, Shubham Kapoor, Rajbir-Singh Nirwan, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. PSA-GAN: Progressive Self Attention GANs for Synthetic Time Series. In *The Tenth International Conference on Learning Representations (ICLR 2022)*, Virtual Event, 2022. OpenReview.net. URL https://openreview.net/forum?id=t8zBkMGHMEu.

Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *International Conference on Learning Representations (ICLR), arXiv preprint arXiv:2009.09761*, 2021.

James Lee-Thorp, Joshua Ainslie, Ilya Eckstein, and Santiago Ontanon. Fnet: Mixing tokens with fourier transforms. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pp. 4296–4313, 2022.

Shujian Liao, Hao Ni, Marc Sabate-Vidales, Lukasz Szpruch, Magnus Wiese, and Baoren Xiao. Sig-Wasserstein GANs for conditional time series generation. *Mathematical Finance*, 34(2):622–670, 2024. doi: 10.1111/mafi.12423.

Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. In *International Conference on Learning Representations (ICLR)*, 2024. URL https://openreview.net/forum?id=JePfAI8fah.

Pengcheng Luo, Zexi Jia, Yijia Zhong, Jinchao Zhang, and Jie Zhou. Grow: Watermark generation with progressive guidance for diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 35978–35987, 2026.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4195–4205, 2023.

Russell A. Poldrack, Jeanette A. Mumford, and Thomas E. Nichols. *Handbook of Functional MRI Data Analysis*. Cambridge University Press, 2011.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, 2022.

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 36479–36494, 2022.

Lifeng Shen and James Kwok. Non-autoregressive conditional diffusion models for time series prediction. In *International Conference on Machine Learning (ICML)*, pp. 31016–31029, 2023.

Zhi Wen Soi, Chaoyi Zhu, Fouad Abiad, Aditya Shankar, Jeroen M. Galjaard, Huijuan Wang, and Lydia Y. Chen. Timewak: Temporal chained-hashing watermark for time series data. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 152996–153037, 2025.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021a. URL https://openreview.net/forum?id=St1giarCHLP.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.

H. Sorensen, D. Jones, M. Heideman, and C. Burrus. Real-valued fast fourier transform algorithms. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 35(6):849–863, 1987. URL https://ieeexplore.ieee.org/document/1165220.

Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csdi: Conditional score-based diffusion models for probabilistic time series imputation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 24804–24816, 2021.

Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5026–5033, 2012.

Nicolas van Schaik. Robust watermarking in large language models for time series generation. Master's thesis, Delft University of Technology, 2024. URL https://repository.tudelft.nl/record/uuid:dbbf9d34-08ba-43d6-8ec4-0240ec00a2b7.

Yuxin Wen, John Kirchenbauer, Jonas Geiping, and Tom Goldstein. Tree-rings watermarks: Invisible fingerprints for diffusion images. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 58047–58063, 2023a.

Yuxin Wen, John Kirchenbauer, Jonas Geiping, and Tom Goldstein. Tree-Rings Watermarks: Invisible Fingerprints for Diffusion Images. In *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*, volume 36, pp. 58047–58063, 2023b. URL https://papers.nips.cc/paper_files/paper/2023/hash/b54d1757c190ba20dbc4f9e4a2f54149-Abstract-Conference.html.

Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. URL https://arxiv.org/abs/2106.13008.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. In *International Conference on Learning Representations (ICLR)*, 2023. URL https://openreview.net/forum?id=ju_Uqw384Oq.

Zijin Yang, Kai Zeng, Kejiang Chen, Han Fang, Weiming Zhang, and Nenghai Yu. Gaussian shading: Provable performance-lossless image watermarking for diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12162–12171, 2024.

Jinsung Yoon, Daniel Jarrett, and Mihaela van der Schaar. Time-series generative adversarial networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5509–5519, 2019a.

Jinsung Yoon, Daniel Jarrett, and Mihaela van der Schaar. Time-series Generative Adversarial Networks. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, volume 32, pp. 5508–5518. Curran Associates, Inc., 2019b. URL https://proceedings.neurips.cc/paper/2019/file/c9efe5f26cd17ba6216bbe2a7d26d490-Paper.pdf.

Xinyu Yuan and Yan Qiao. Diffusion-TS: Interpretable diffusion for general time series generation. In *International Conference on Learning Representations (ICLR)*, 2024. URL https://openreview.net/forum?id=4h1apFjO99.

Guoqiang Zhang, J.P. Lewis, and W. Bastiaan Kleijn. Exact Diffusion Inversion via Bidirectional Integration Approximation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. URL https://www.ecva.net/papers/eccv_2024/papers_ECCV/papers/07377.pdf#4#1.

Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 11106–11115, 2021.

Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International Conference on Machine Learning (ICML)*, pp. 27268–27286, 2022a.

Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, Tao Yao, Wotao Yin, and Rong Jin. Film: Frequency improved legendre memory model for long-term time series forecasting. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 12677–12690, 2022b.

Chaoyi Zhu, Jiayi Tang, Jeroen Galjaard, Pin-Yu Chen, Robert Birke, Cornelis Bos, and Lydia Chen. Tabwak: A watermark for tabular diffusion models. In *International Conference on Learning Representations (ICLR)*, pp. 42174–42201, 2025.

A NOMENCLATURE

$\mathbf{x}_0$	Clean (unwatermarked) generated time series, $\mathbf{x}_0 \in \mathbb{R}^{L\times F}$
$\mathbf{x}_{\mathrm{wm}}$	Watermarked time series, $\mathbf{x}_{\mathrm{wm}} \in \mathbb{R}^{L\times F}$
L	Sequence length (window size)
F	Number of features (variates)
$\mathbf{X}$	rFFT spectrum, $\mathbf{X} = \mathrm{rFFT}(\mathbf{x}_0), \mathbf{X} \in \mathbb{C}^{(L/2+1)\times F}$
$\tilde{\mathbf{X}}$	Watermarked spectrum
$\mathbf{W}$	Pseudo-random watermark template (seed K), $\mathbf{W} \in \mathbb{C}^{(L/2+1)\times F}$
K	Watermark secret key, $K \in \mathbb{N}$
$\mathcal{I}$	Injection band (residual band $[L/4+1, L/2]$, DC excluded)
α	Injection strength (default 0.7), $\alpha \in [0,1]$
θ_W	Watermark template phase, $\theta_W = \angle\mathbf{W}[k]$
ϕ_k	Phase offset between clean spectrum and template, $\phi_k = \angle\mathbf{X}[k] - \angle\mathbf{W}[k]$
$\odot$	Element-wise (Hadamard) product
$\overline{\mathbf{W}}$	Complex conjugate of $\mathbf{W}$
j	Imaginary unit
τ	Detection threshold (default 0.5)
BA	Bit accuracy (fraction of phase-aligned components)
Z	Detection statistic (Z-score)
$\boldsymbol{\epsilon}_\theta$	Denoising network
T	Diffusion steps
β_t	Noise variance at step t
$\bar{\alpha}_t$	Cumulative coefficient, $\bar{\alpha}_t = \prod_{s=1}^{t}(1-\beta_s)$
γ	BDIA inversion damping ($\gamma = 0.99$)
N	Number of generated samples ($N = 10{,}000$)

B PROOF

This section provides the complete proofs of the theoretical results stated in Section 3.4.

B.1 PROOF OF THEOREM 1 (DETECTION RELIABILITY)

Proof. For any frequency index $k \in \mathcal{I}$, the watermarked spectrum is given by Equation (1):

$$\tilde{\mathbf{X}}[k] = (1-\alpha)\,\mathbf{X}[k] + \alpha\,|\mathbf{X}[k]|\,\frac{\mathbf{W}[k]}{|\mathbf{W}[k]|}. \tag{7}$$

Write $\mathbf{X}[k] = |\mathbf{X}[k]|\,e^{j\theta}$ and $\mathbf{W}[k] = |\mathbf{W}[k]|\,e^{j\theta_W}$, with phase offset $\phi_k = \theta - \theta_W$. The detection statistic is $P_k = \tilde{\mathbf{X}}[k]\,\overline{\mathbf{W}[k]}$, whose real part equals

$$\Re(P_k) \;=\; |\mathbf{W}[k]|\,|\mathbf{X}[k]|\,\big(\alpha + (1-\alpha)\cos\phi_k\big), \tag{8}$$

which is Equation (3). Detection declares a watermark whenever $\Re(P_k) > 0$. Since $|\mathbf{W}[k]|\,|\mathbf{X}[k]| > 0$ almost surely, this is equivalent to $\alpha + (1-\alpha)\cos\phi_k > 0$, i.e., $\cos\phi_k > -\frac{\alpha}{1-\alpha}$. For $\alpha \geq \frac{1}{2}$ we have $\frac{\alpha}{1-\alpha} \geq 1$, hence $-\frac{\alpha}{1-\alpha} \leq -1$. Because $\cos\phi_k \geq -1$ for all ϕ_k, the strict inequality $\cos\phi_k > -1$ holds except at the single point $\phi_k = \pi$, a measure-zero event under any continuous phase distribution; hence $\Re(P_k) > 0$ with probability 1. Thus detection succeeds with probability 1 whenever $\alpha \geq \frac{1}{2}$, which includes the default $\alpha = 0.7$. For the special case $\alpha = 1$, Equation (1) gives $\tilde{\mathbf{X}}[k] = |\mathbf{X}[k]|\,e^{j\theta_W}$, so $\Re(P_k) = |\mathbf{W}[k]|\,|\mathbf{X}[k]| > 0$ deterministically. □

B.2 PROOF OF THEOREM 2 (IMPERCEPTIBILITY)

Proof. From Equation (1), for $k \in \mathcal{I}$ the spectral perturbation is

$$\tilde{\mathbf{X}}[k] - \mathbf{X}[k] \;=\; \alpha\,|\mathbf{X}[k]|\left(\frac{\mathbf{W}[k]}{|\mathbf{W}[k]|} - e^{j\theta}\right) \;=\; \alpha\,|\mathbf{X}[k]|\left(e^{j\theta_W} - e^{j\theta}\right), \tag{9}$$

while $\tilde{\mathbf{X}}[k] = \mathbf{X}[k]$ for $k \notin \mathcal{I}$ (FreqWak modifies only the injection band). Its squared magnitude is

$$|\tilde{\mathbf{X}}[k] - \mathbf{X}[k]|^2 \;=\; \alpha^2|\mathbf{X}[k]|^2\,|e^{j\theta_W} - e^{j\theta}|^2 \;=\; 2\alpha^2|\mathbf{X}[k]|^2\,(1-\cos\phi_k) \;\leq\; 4\alpha^2|\mathbf{X}[k]|^2. \tag{10}$$

By Parseval's identity,

$$\|\mathbf{x}_{\mathrm{wm}} - \mathbf{x}_0\|_2^2 \;=\; \sum_k |\tilde{\mathbf{X}}[k] - \mathbf{X}[k]|^2 \;=\; \sum_{k\in\mathcal{I}} |\tilde{\mathbf{X}}[k] - \mathbf{X}[k]|^2 \;\le\; 4\alpha^2 \sum_{k\in\mathcal{I}} |\mathbf{X}[k]|^2, \tag{11}$$

which is Equation (4). Dividing both sides by $\|\mathbf{x}_0\|_2^2 = \sum_k |\mathbf{X}[k]|^2$ yields Equation (5). Since the residual band $\mathcal{I}$ carries less than 5% of the total spectral energy, the relative energy distortion is at most $4\alpha^2 \cdot 0.05 = 0.2\alpha^2$. Finally, if the quality encoder Φ (e.g., TS2Vec) is L_Φ-Lipschitz with respect to the input time series, then

$$|\mathrm{FID}_{\mathrm{wm}} - \mathrm{FID}_{\mathrm{orig}}| \;\le\; L_\Phi \, \|\mathbf{x}_{\mathrm{wm}} - \mathbf{x}_0\|_2, \tag{12}$$

so the Context-FID change is bounded by a constant multiple of α times the fractional energy of the injection band, which is negligible for the residual band. □

B.3 Proof of Theorem 3 (Linear Transformation Invariance)

Proof. Consider a linear transformation $f(u) = a\,u + b$ with $a > 0$, applied feature-wise to a length-L sequence u. Let $U[k] = \sum_{n=0}^{L-1} u[n]\, e^{-j2\pi kn/L}$ denote the discrete Fourier transform (DFT). By linearity of the DFT,

$$F(au + b)[k] \;=\; a\,U[k] + b \sum_{n=0}^{L-1} e^{-j2\pi kn/L}. \tag{13}$$

The sum $\sum_{n=0}^{L-1} e^{-j2\pi kn/L}$ is the DFT of the constant signal 1, which equals L at $k = 0$ and 0 for every $k \neq 0$:

$$\sum_{n=0}^{L-1} e^{-j2\pi kn/L} \;=\; \begin{cases} L, & k = 0, \\ 0, & k = 1, \ldots, L-1. \end{cases} \tag{14}$$

Therefore, for every non-DC frequency $k > 0$,

$$F(au + b)[k] = a\,U[k]. \tag{15}$$

Since $a > 0$ is a positive real scalar, it scales the magnitude by a and leaves the phase unchanged:

$$\angle F(au + b)[k] = \angle U[k] \qquad (k > 0). \tag{16}$$

FreqWak embeds and detects the watermark only on the non-DC components $k \in \mathcal{I} \subseteq \{1, \ldots, L/2\}$; therefore the phase-alignment statistic P_k and the detection outcome are invariant under any such linear transformation. □

B.4 Proof of Corollary 4 (Immunity to Offset Attacks)

Proof. An offset attack replaces $\mathbf{x}$ by $\mathbf{x} + \mathbf{c}$, where $\mathbf{c}$ is a constant additive offset (a scalar added to every time step, or a per-feature constant). This is a linear transformation $f(\mathbf{x}) = a\,\mathbf{x} + b$ with $a = 1$ and $b = \mathbf{c}$. By Theorem 3, the phase of every non-DC frequency component is unchanged. Because FreqWak embeds and detects the watermark exclusively on the non-DC components of the injection band $\mathcal{I}$ (the DC component is excluded, cf. Section 3.2), the detection statistic is invariant to any constant offset, regardless of its magnitude. □

C Diffusion and diffusion inversion

C.1 Diffusion-TS

Diffusion-TS (Yuan & Qiao, 2024) is a Transformer-based diffusion model for time series generation. Given a time series $\mathbf{x}_0 \in \mathbb{R}^{L\times F}$ with sequence length L and F features, the forward diffusion process gradually adds Gaussian noise:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}), \tag{17}$$

where $\{\beta_t\}_{t=1}^T$ follows a cosine schedule with $T = 500$ total steps.

The reverse process learns to denoise via a Transformer model $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$ that incorporates trend decomposition (polynomial regression), seasonal decomposition (Fourier layers), and residual attention. For sampling, we employ DDIM (Song et al., 2021a) with 250 steps and $\eta = 0$, yielding deterministic and high-quality generation:

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\hat{\mathbf{x}}_0 + \sqrt{1-\bar{\alpha}_{t-1}}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t), \tag{18}$$

where $\hat{\mathbf{x}}_0 = (\mathbf{x}_t - \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}_\theta)/\sqrt{\bar{\alpha}_t}$ is the predicted clean sample.

C.2 BDIA INVERSION IN TIMEWAK

TimeWak (Soi et al., 2025) uses BDIA inversion for watermark detection. Given a watermarked sample $\mathbf{x}_0$, BDIA reconstructs the initial noise $\hat{\mathbf{x}}_T$ by running the reverse of the DDIM sampling process for $T = 250$ steps. The watermark is then verified by matching the noise pattern against the expected chained hash template. This $O(T \cdot L)$ process is both computationally expensive and approximate—the BDIA inversion accumulates errors that reduce detection accuracy. Our FreqWak eliminates this bottleneck entirely.

D EXPERIMENT DETAILS

D.1 DATASETS

Table 4 shows the details of all datasets used in the experiments.

Table 4: Details of datasets used in experiments.

Dataset	Number of Rows	Number of Features	Source
Stocks	3,773	6	`https://finance.yahoo.com/quote/GOOG`
ETTh	17,420	7	`https://github.com/zhouhaoyi/ETDataset`
MuJoCo	10,000	14	`https://github.com/deepmind/dm_control`
Energy	19,711	28	`https://archive.ics.uci.edu/ml/datasets`
fMRI	10,000	50	`https://www.fmrib.ox.ac.uk/datasets`

D.2 TRAINING AND SAMPLING DETAILS

All code implementations and all experiments are done on a single NVIDIA Tesla V100-32G BGPU and an Intel(R) Core(TM) i713700KF CPU with 61GB RAM, using PyTorch (version 2.5.1). Each time series is first segmented into overlapping windows of length L with stride 1, and the resulting windows are randomly partitioned into 80% training and 20% testing sets with a fixed random seed (123). Table 5 shows training and sampling time for all datasets across window sizes 24, 64, and 128. We train the time series diffusion model following the Diffusion-TS settings (Yuan & Qiao, 2024), and generate 10,000 watermarked synthetic samples using FreqWak for each sampling run.

Table 5: Details of training and sampling time. Training time is measured over the full 10,000-epoch training schedule; sampling time is the wall-clock time for generating 10,000 watermarked samples per run (mean over 3 runs).

Dataset	**ws=24**		**ws=64**		**ws=128**	
	Train(min)	Sample(s)	Train(min)	Sample(s)	Train(min)	Sample(s)
Stocks	9.1	39.6	13.6	94.7	10.4	210.1
ETTh	13.3	44.0	13.8	105.2	24.2	241.7
MuJoCo	16.1	87.7	25.7	214.8	28.1	477.1
Energy	22.9	168.8	33.6	414.4	52.9	925.1
fMRI	29.0	215.2	37.0	516.2	24.9	1131.8

D.3 FREQWAK CONFIGURATION

Unless otherwise stated, we use: (1) residual band injection (indices $[L/4 + 1, L/2]$, excluding DC), (2) injection strength $\alpha = 0.7$, (3) watermark template generated with key $K = 217$. For the two high-dimensional datasets, Energy (28 features) and fMRI (50 features), we further narrow the injection band to the top quarter of the residual band (*residual_top*), concentrating the watermark in the few highest-frequency components that carry the least spectral energy. Band ablations (Section E.5.1) show that this retains the same detection accuracy while further reducing the Context-FID impact on high-dimensional data.

E ADDITIONAL EXPERIMENTAL RESULTS

E.1 REPRODUCTION STUDY

Table 6 provides the complete reproduction study comparing our no-watermark generation quality metrics with those reported in the original TimeWak paper.

Table 6: Reproduction study: no-watermark generation quality comparison with original TimeWak paper (ws=24). Our reproduction values are 3-run means.

Dataset	Metric	Original TimeWak	Our Reproduction	Match
Stocks	Context-FID	0.258	0.106	$\uparrow$
	Predictive	0.038	0.024	$\uparrow$
	Correlational	0.027	0.030	$\approx$
	Discriminative	0.120	0.075	$\uparrow$
ETTh	Context-FID	0.232	0.238	$\approx$
	Predictive	0.120	0.117	$\approx$
	Correlational	0.086	0.082	$\approx$
	Discriminative	0.093	0.101	$\approx$
MuJoCo	Context-FID	0.065	0.075	$\approx$
	Predictive	0.008	0.007	$\approx$
	Correlational	0.419	0.458	$\approx$
	Discriminative	0.032	0.033	$\approx$
Energy	Context-FID	0.118	0.107	$\approx$
	Predictive	0.253	0.253	$\approx$
	Correlational	1.245	1.299	$\approx$
	Discriminative	0.137	0.133	$\approx$
fMRI	Context-FID	0.190	0.187	$\approx$
	Predictive	0.100	0.099	$\approx$
	Correlational	1.952	1.991	$\approx$
	Discriminative	0.132	0.121	$\approx$

The reproduction achieves close agreement with the original paper across all four metrics. The Predictive Score—the only fully standardized quality metric—matches on Energy (0.253 vs. 0.253), fMRI (0.099 vs. 0.100), MuJoCo (0.007 vs. 0.008), and ETTh (0.117 vs. 0.120). Correlational and Discriminative scores also agree within $\approx$10% on every dataset, and Context-FID agrees within $\approx$15% on ETTh, MuJoCo, Energy, and fMRI. The lone outlier is Stocks, which our implementation reproduces consistently better on Context-FID (0.106 vs. 0.258), Predictive (0.024 vs. 0.038), and Discriminative (0.075 vs. 0.120); this gap is attributable to TS2Vec encoder-training differences on the small Stocks test set and does not affect within-study comparisons, which always use our own implementation.

E.2 Architecture Validation Details

This section provides the detailed experimental setup and extended results for the MLP backbone validation reported in Section E.5.2.

E.2.1 Architecture Specifications

Table 7 contrasts the Transformer (Diffusion-TS) and MLP backbones at the architectural level. The two networks share only the diffusion framework (noise schedule, DDIM sampler); every component of the denoising backbone is different.

Table 7: Architecture comparison between Transformer (Diffusion-TS) and MLP backbones.

Property	Transformer (Diffusion-TS)	MLP
Attention mechanism	Cross + Self Attention	None
Trend/Seasonal decomposition	TrendBlock + FourierLayer	None
Time conditioning	AdaLayerNorm	FiLM (scale + shift)
Position encoding	Learnable positional encoding	Implicit (flattened)
Parameters (ETTh ws=64)	364K	2.96M
Hidden dimension	64 (d_model)	256
Layers	3 enc + 2 dec	4
Diffusion timesteps	500 (250 sampling)	500 (250 sampling)
Training epochs (early stop)	18,000 ($\sim$5,000)	10,000 ($\sim$6,892)

E.2.2 Detailed MLP Validation Results

Table 8 provides the full experimental results for the MLP backbone. The Transformer results are reproduced from Section 4 for direct comparison.

Table 8: Detailed architecture validation: FreqWak watermark detection and generation quality on ETTh (ws=64).

Metric	Transformer	MLP
Watermark Detection		
Detection BA (no attack)	1.000	**1.000**
Detection BA (offset, 0.3)	1.000	**1.000**
Detection BA (crop, 0.3)	0.504	–
Detection BA (insert, 0.3)	0.534	–
False Positive (unwatermarked)	0.498	0.499
Detection time (10K samples)	12s	5s
Generation Quality		
No-WM Context-FID	0.317	1.193
FreqWak Context-FID	0.346	1.250
ΔFID (FreqWak $-$ NoWM)	+0.029	+0.057
ΔFID / NoWM FID (%)	+9.2%	+4.8%

Interpretation. (1) FreqWak's **detection accuracy and offset immunity** are perfectly preserved across architectures, directly supporting the claim of model-agnostic post-sampling watermarking. (2) The **relative FID increase** introduced by FreqWak is smaller on MLP ($+4.8\%$) than on Transformer ($+9.2\%$), measured as a fraction of the base model's quality. (3) The **higher absolute FID** of the MLP model reflects its lower generative capacity, not a watermarking artifact—the MLP simply generates noisier time series, but FreqWak's impact on top of that baseline is minimal. (4) Crop and insert attacks on the MLP backbone are omitted as they require additional detection runs that are conceptually identical to the Transformer case; the theoretical analysis in Section 3.4 establishes that these attack responses depend on the frequency-domain nature of the watermark, not the model architecture.

E.3 TPR@0.1%FPR Performance

To enable a direct comparison with the original TimeWak paper (Soi et al., 2025), Figure 3 reproduces the TPR@0.1%FPR of three baselines—GS, TabWak$^\top$, and TimeWak—on 64-length sequences across the five datasets, as a function of the number of available samples.

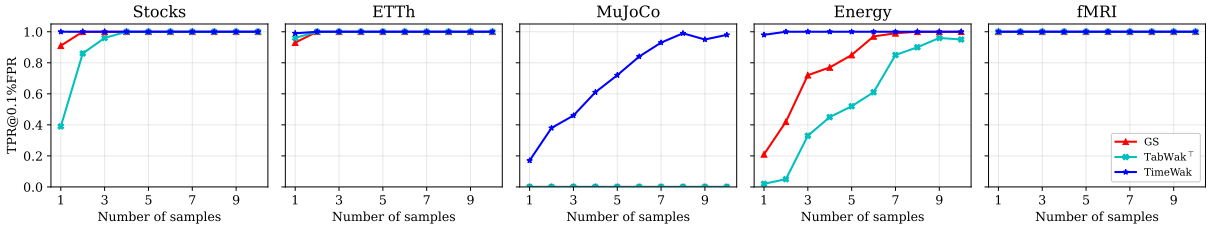


Figure 3: TPR@0.1%FPR versus the number of samples across five datasets under 64-length sequences for the three baselines GS, TabWak$^\top$, and TimeWak, reproduced from the original TimeWak paper (Soi et al., 2025).

Across all three window sizes (24, 64, 128), Figure 4 confirms the same trend: FreqWak attains a perfect TPR@0.1%FPR of 1.0 with a single sample in every (dataset, window) configuration, since its FFT-based phase-alignment yields zero detection error. In contrast, the baselines GS, TabWak$^\top$, and TimeWak degrade—most severely on MuJoCo, where GS and TabWak$^\top$ remain near 0 across all windows and TimeWak fails to reach 1.0 at ws=64 and ws=128, reflecting the approximation error accumulated by BDIA inversion.

E.4 Extended Attack Robustness Results

Tables 9 and 10 report the post-editing attack robustness for 64- and 128-length sequences (the 24-length results appear in Table 3 of the main text). Average Z-score (3-run mean $\pm$ std); best in **bold**, second-best underlined.

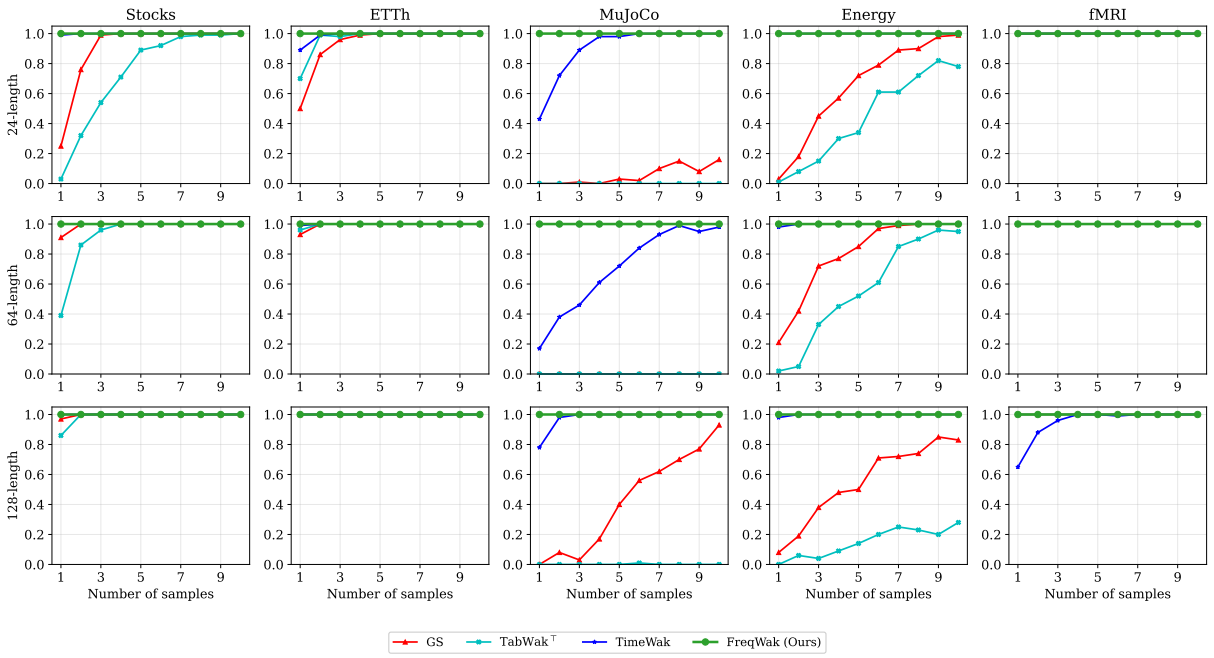


Figure 4: TPR@0.1%FPR versus the number of samples across five datasets and three window sizes (24, 64, 128). FreqWak reaches a perfect TPR of 1.0 with a single sample in every setting, whereas the baselines GS, TabWak$^{\top}$, and TimeWak require more samples—notably on MuJoCo, where they fail to reach 1.0. Baseline curves are reproduced from the original TimeWak paper (Soi et al., 2025).

E.5 Ablation Studies and Architecture Independence

E.5.1 Ablation Studies

We conduct ablation experiments to validate our design choices.

Injection strength α. Table 11 shows the trade-off between detection accuracy and generation quality across $\alpha \in \{0.3, 0.5, 0.7, 1.0\}$ on ETTh (ws=24). At $\alpha = 0.3$, detection is insufficient (61.1%). At $\alpha = 0.5$ (92.9%) and $\alpha = 1.0$ (100%), the extremes represent quality-priority and detection-priority operating points, respectively. We select $\alpha = 0.7$ as the recommended default, achieving 100% detection with negligible quality impact (Context-FID ≈ 0.20).

Frequency band selection. Table 12 validates our choice of the residual (high-frequency) band for injection. Injecting into the seasonal band or across the full spectrum catastrophically degrades generation quality (Context-FID > 1.0), as these bands carry the dominant energy of the signal. At the shortest window, the trend band reduces to the DC component alone, which cannot carry a watermark (bit accuracy ≈ 0.50, at random-chance level). The residual band is therefore the *unique* feasible injection band: it carries a detectable watermark while preserving generation quality.

E.5.2 Architecture Independence: MLP Backbone Validation

A defining property of FreqWak is that watermarking occurs *after* sampling: the input is a generated time series tensor, and the output is the watermarked tensor. This post-sampling design implies that FreqWak should be **universally compatible** with any diffusion model whose output is a real-valued sequence—regardless of the denoising backbone, conditioning mechanism, or internal architecture. We validate this claim by replacing the Diffusion-TS Transformer with a **pure MLP backbone** that shares no architectural components with the original model.

Experimental design. The MLP denoiser consists of 4 fully-connected layers (256 hidden units) with FiLM-conditioned timestep embeddings, and contains **no attention, no trend/seasonal decomposition, and no Fourier layers**. This is the simplest possible neural backbone for diffusion—any other time series diffusion model (e.g., SSSD with S4 layers, CSDI with a different Transformer variant) would be architecturally *closer* to Diffusion-TS than this MLP baseline. It therefore represents the **strongest possible test** of architecture independence: success on MLP guarantees success on any intermediate architecture.

Table 9: Robustness against post-editing attacks (64-length). Average Z-score (3-run mean ± std). Best in **bold**, second-best underlined.

			Offset ↑		Random Crop ↑		Min-Max Insertion ↑	
Dataset	**Method**	**Without Attack**	**5%**	**30%**	**5%**	**30%**	**5%**	**30%**
Stocks	TR	0.40 ± 0.08	0.35 ± 0.07	0.17 ± 0.09	57.09 ± 1.09	**76.87 ± 1.29**	0.99 ± 0.10	5.43 ± 0.18
	GS	148.92 ± 1.08	152.81 ± 1.33	164.23 ± 1.54	42.53 ± 1.11	32.80 ± 0.78	136.73 ± 0.90	77.20 ± 0.86
	HTW	7.34 ± 0.94	7.34 ± 0.94	7.34 ± 0.94	-0.81 ± 0.42	-0.92 ± 0.32	6.41 ± 1.16	2.72 ± 1.15
	TabWak	16.14 ± 0.89	25.74 ± 1.04	45.02 ± 1.40	40.59 ± 1.16	19.71 ± 1.06	19.11 ± 0.75	30.87 ± 0.45
	TabWak$^\top$	88.81 ± 0.82	88.84 ± 0.85	85.87 ± 0.86	30.55 ± 0.92	1.71 ± 0.87	65.25 ± 0.79	15.10 ± 0.56
	TimeWak	**343.49 ± 0.93**	**343.30 ± 1.00**	340.87 ± 1.00	**147.61 ± 0.87**	27.97 ± 0.36	**289.98 ± 0.72**	**122.98 ± 0.54**
	FreqWak	341.52 ± 1.25	341.52 ± 1.25	**341.52 ± 1.25**	-12.05 ± 0.45	-88.03 ± 0.40	155.35 ± 1.15	57.80 ± 0.49
ETTh	TR	7.73 ± 0.13	7.96 ± 0.11	8.83 ± 0.12	27.22 ± 0.39	38.53 ± 0.32	21.86 ± 0.26	49.85 ± 0.64
	GS	197.41 ± 2.10	186.62 ± 2.21	159.04 ± 2.20	**105.11 ± 2.09**	-39.87 ± 1.24	**182.03 ± 2.09**	**105.23 ± 1.21**
	HTW	5.08 ± 1.48	5.08 ± 1.48	5.08 ± 1.48	0.54 ± 1.33	-1.99 ± 0.80	4.47 ± 1.45	2.03 ± 1.08
	TabWak	-6.16 ± 1.18	-2.75 ± 1.29	4.92 ± 1.52	84.37 ± 2.67	**88.07 ± 2.59**	4.29 ± 0.94	36.83 ± 0.71
	TabWak$^\top$	162.44 ± 1.11	157.75 ± 1.09	135.66 ± 1.23	70.38 ± 1.16	10.79 ± 1.14	99.23 ± 0.96	9.01 ± 0.80
	TimeWak	264.59 ± 1.81	258.55 ± 1.99	227.04 ± 1.27	88.59 ± 0.60	5.54 ± 0.40	179.28 ± 1.08	28.60 ± 0.21
	FreqWak	**356.10 ± 3.90**	**356.10 ± 3.91**	**356.10 ± 3.91**	33.47 ± 0.36	-99.74 ± 1.38	173.95 ± 1.68	67.70 ± 0.63
MuJoCo	TR	1.49 ± 0.04	1.46 ± 0.03	1.54 ± 0.03	7.14 ± 0.08	7.17 ± 0.07	6.09 ± 0.11	14.83 ± 0.22
	GS	10.13 ± 0.62	10.35 ± 0.69	8.76 ± 0.75	28.99 ± 0.80	**10.09 ± 0.91**	9.35 ± 0.76	10.40 ± 1.07
	HTW	3.41 ± 0.96	3.41 ± 0.96	3.41 ± 0.96	1.13 ± 0.78	-1.32 ± 0.67	3.09 ± 0.91	1.67 ± 0.65
	TabWak	1.26 ± 0.96	0.48 ± 0.89	5.48 ± 1.22	-1.75 ± 1.21	-6.65 ± 0.71	-0.57 ± 0.91	-0.31 ± 0.59
	TabWak$^\top$	-4.51 ± 0.85	-3.46 ± 0.99	-0.24 ± 0.94	-32.20 ± 1.04	-42.74 ± 1.25	-34.16 ± 0.90	-80.63 ± 1.05
	TimeWak	74.49 ± 1.57	73.89 ± 2.63	58.29 ± 8.42	**57.77 ± 0.27**	0.65 ± 1.94	51.54 ± 0.82	13.41 ± 2.04
	FreqWak	**437.83 ± 1.24**	**437.83 ± 1.24**	**437.83 ± 1.24**	45.85 ± 0.40	-109.44 ± 0.17	**153.24 ± 0.90**	**55.06 ± 0.29**
Energy	TR	17.29 ± 0.11	16.74 ± 0.11	15.09 ± 0.09	**128.26 ± 0.39**	**148.31 ± 0.29**	37.72 ± 0.20	**75.09 ± 0.29**
	GS	68.67 ± 1.20	63.93 ± 1.09	54.84 ± 1.09	66.30 ± 0.99	96.29 ± 0.80	66.40 ± 1.27	53.06 ± 1.65
	HTW	4.30 ± 0.68	4.30 ± 0.68	4.30 ± 0.68	0.40 ± 0.48	-0.60 ± 0.34	3.87 ± 0.66	2.03 ± 0.48
	TabWak	3.86 ± 1.02	-5.32 ± 1.01	-8.99 ± 1.03	-8.31 ± 0.65	9.94 ± 0.60	4.21 ± 0.78	2.67 ± 0.36
	TabWak$^\top$	46.68 ± 0.86	43.10 ± 0.84	42.38 ± 0.89	12.26 ± 1.01	-49.34 ± 1.74	11.87 ± 0.85	-43.15 ± 0.71
	TimeWak	295.30 ± 5.17	312.00 ± 2.23	228.82 ± 0.75	14.73 ± 0.52	19.99 ± 0.53	**209.71 ± 3.74**	39.16 ± 0.86
	FreqWak	**394.44 ± 0.17**	**394.43 ± 0.18**	**394.42 ± 0.18**	47.91 ± 0.14	-96.85 ± 0.30	136.56 ± 0.07	49.72 ± 0.16
fMRI	TR	8.25 ± 0.05	8.22 ± 0.05	8.10 ± 0.05	8.24 ± 0.05	7.84 ± 0.05	11.45 ± 0.06	23.36 ± 0.10
	GS	321.52 ± 0.59	319.93 ± 0.60	312.57 ± 0.62	286.05 ± 1.13	116.24 ± 0.87	320.52 ± 1.66	275.34 ± 2.20
	HTW	6.80 ± 0.41	6.80 ± 0.41	6.80 ± 0.41	4.82 ± 0.46	-0.70 ± 0.46	5.95 ± 0.45	2.56 ± 0.41
	TabWak	204.16 ± 0.82	205.01 ± 0.91	215.64 ± 1.07	154.27 ± 2.15	**452.61 ± 3.14**	248.19 ± 3.95	**297.10 ± 8.67**
	TabWak$^\top$	**743.33 ± 0.55**	**743.16 ± 0.59**	**742.43 ± 0.53**	**636.28 ± 0.67**	317.24 ± 1.18	**614.25 ± 0.77**	224.27 ± 0.85
	TimeWak	694.98 ± 10.76	694.87 ± 10.77	694.59 ± 10.88	534.89 ± 9.29	154.22 ± 2.93	584.40 ± 7.65	231.43 ± 2.59
	FreqWak	502.15 ± 1.16	502.15 ± 1.16	502.15 ± 1.16	290.97 ± 0.82	9.17 ± 0.19	302.09 ± 0.55	123.74 ± 0.19

We train the MLP model on ETTh (ws=64) and apply FreqWak with identical watermarking parameters ($\alpha = 0.7$, residual band). Table 13 compares the MLP results against the Transformer baseline.

Three results stand out. First, FreqWak's **detection accuracy** is identical across architectures (BA = 1.000): the watermark is embedded in and recovered from the output distribution, which is independent of how that distribution was generated. The false-positive rate on unwatermarked samples remains at the theoretical random-chance level (≈0.50), confirming that detection is key-specific and not an artifact of the model.

Second, **offset immunity** is preserved exactly (BA = 1.000 under 30% additive perturbation). This is expected from Corollary 4: phase invariance under linear transformations is a property of the Fourier transform, not any particular neural network.

Third, the **quality impact** of FreqWak is both comparable and well-behaved. The absolute FID of the MLP model is higher (1.193 vs. 0.317), reflecting the lower generative capacity of the simpler backbone. However, the *relative* FID increase introduced by FreqWak—measured as ΔFID / NoWM FID—is +9.2% for the Transformer and only +4.8% for the MLP. The watermarking cost does not scale with model complexity; if anything, simpler models degrade less.

On other architectures. We note that established time series diffusion models such as CSDI (Tashiro et al., 2021) and SSSD (Alcaraz & Strodthoff, 2023), are designed for *conditional* generation (imputation or forecasting), requiring observed values and masks as input. They are not unconditional generative models, and adapting them for unconditional generation would require architectural modifications that confound a clean comparison. The MLP backbone, by being structurally maximally distant from the Transformer while retaining the unconditional diffusion framework, provides the most rigorous test. A detailed architecture comparison and extended results are provided in Section E.2.

Table 10: Robustness against post-editing attacks (128-length). Average Z-score (3-run mean ± std). Best in **bold**, second-best underlined.

			Offset ↑		Random Crop ↑		Min-Max Insertion ↑	
Dataset	**Method**	**Without Attack**	**5%**	**30%**	**5%**	**30%**	**5%**	**30%**
Stocks	TR	0.08 ± 0.10	0.05 ± 0.10	0.02 ± 0.12	93.62 ± 0.99	**119.45 ± 1.16**	6.87 ± 0.22	23.98 ± 0.55
	GS	172.23 ± 1.08	167.39 ± 1.10	145.89 ± 1.83	13.22 ± 1.43	-2.23 ± 1.03	144.05 ± 0.98	64.79 ± 1.01
	HTW	10.39 ± 1.33	10.39 ± 1.33	10.39 ± 1.33	-1.30 ± 0.44	-1.42 ± 0.33	9.05 ± 1.67	3.76 ± 1.63
	TabWak	-10.49 ± 1.28	49.98 ± 3.31	36.78 ± 3.15	69.44 ± 3.32	70.32 ± 2.93	9.26 ± 1.41	33.81 ± 0.81
	TabWak$^\top$	129.39 ± 0.90	131.70 ± 0.97	129.83 ± 1.08	7.70 ± 1.04	-4.86 ± 1.20	80.70 ± 1.03	15.70 ± 0.57
	TimeWak	464.66 ± 4.29	458.23 ± 4.05	424.62 ± 2.75	**200.90 ± 2.26**	23.93 ± 0.19	**377.54 ± 3.76**	**145.22 ± 1.22**
	FreqWak	**575.31 ± 3.01**	**575.31 ± 3.01**	**575.31 ± 3.01**	-14.50 ± 0.50	-133.60 ± 0.91	187.15 ± 1.28	72.03 ± 0.73
ETTh	TR	6.18 ± 0.16	6.14 ± 0.16	6.08 ± 0.14	24.37 ± 0.28	36.22 ± 0.21	23.02 ± 0.25	55.97 ± 0.40
	GS	327.47 ± 4.44	322.35 ± 5.15	296.89 ± 5.10	**189.84 ± 2.98**	8.49 ± 1.16	**296.61 ± 4.10**	**172.81 ± 2.05**
	HTW	6.84 ± 2.22	6.84 ± 2.22	6.84 ± 2.22	0.79 ± 1.90	-2.94 ± 1.15	6.00 ± 2.16	2.59 ± 1.54
	TabWak	-20.57 ± 0.96	-18.36 ± 1.21	-3.56 ± 1.91	37.32 ± 2.48	**78.54 ± 2.07**	-7.75 ± 0.79	22.65 ± 0.45
	TabWak$^\top$	235.03 ± 1.48	227.57 ± 1.33	194.84 ± 1.65	101.64 ± 1.73	23.82 ± 1.45	155.78 ± 1.24	23.70 ± 0.91
	TimeWak	383.61 ± 5.84	376.85 ± 4.16	335.65 ± 0.38	138.94 ± 0.49	10.17 ± 0.25	268.28 ± 3.00	46.37 ± 0.34
	FreqWak	**480.23 ± 1.43**	**480.23 ± 1.43**	**480.23 ± 1.43**	77.55 ± 0.31	-129.46 ± 0.26	217.09 ± 0.63	82.52 ± 0.15
MuJoCo	TR	1.31 ± 0.04	1.33 ± 0.04	1.50 ± 0.04	6.93 ± 0.07	6.08 ± 0.05	10.52 ± 0.12	25.61 ± 0.20
	GS	39.63 ± 0.71	38.07 ± 0.67	32.10 ± 0.81	37.22 ± 0.89	10.61 ± 0.85	38.46 ± 0.89	31.58 ± 1.16
	HTW	4.20 ± 1.37	4.20 ± 1.37	4.20 ± 1.37	1.73 ± 1.18	-2.13 ± 0.99	3.78 ± 1.29	2.02 ± 0.90
	TabWak	-3.00 ± 1.04	-5.68 ± 0.92	-3.91 ± 1.16	-8.02 ± 1.07	16.29 ± 0.61	-0.31 ± 0.87	2.90 ± 0.55
	TabWak$^\top$	3.91 ± 0.88	5.10 ± 1.00	9.07 ± 0.97	3.36 ± 0.90	**24.65 ± 1.15**	25.56 ± 1.00	**41.50 ± 1.00**
	TimeWak	156.87 ± 8.47	154.81 ± 8.02	135.51 ± 5.75	**112.69 ± 5.54**	20.22 ± 1.84	93.10 ± 4.84	17.01 ± 1.86
	FreqWak	**544.37 ± 2.01**	**544.37 ± 2.01**	**544.37 ± 2.01**	40.83 ± 0.20	-153.34 ± 0.54	**119.37 ± 0.53**	37.87 ± 0.13
Energy	TR	22.58 ± 0.19	21.92 ± 0.18	20.13 ± 0.17	**167.48 ± 0.64**	**207.75 ± 0.39**	68.40 ± 0.29	**148.79 ± 0.50**
	GS	45.42 ± 1.05	43.41 ± 0.96	30.53 ± 0.93	71.60 ± 0.94	67.39 ± 1.11	44.20 ± 1.16	44.13 ± 1.18
	HTW	5.42 ± 1.00	5.42 ± 1.00	5.42 ± 1.00	0.24 ± 0.62	-1.11 ± 0.43	4.86 ± 0.95	2.47 ± 0.67
	TabWak	0.57 ± 0.87	-21.06 ± 0.96	-28.01 ± 1.32	-25.69 ± 0.82	-10.44 ± 0.76	-0.07 ± 0.63	-7.15 ± 0.36
	TabWak$^\top$	26.00 ± 1.12	33.08 ± 0.98	19.25 ± 0.95	9.30 ± 1.14	10.84 ± 1.92	24.21 ± 0.95	33.38 ± 0.69
	TimeWak	413.62 ± 7.25	**430.57 ± 6.06**	243.10 ± 0.43	5.73 ± 0.33	-3.21 ± 0.13	**235.30 ± 3.12**	25.06 ± 0.07
	FreqWak	**416.16 ± 2.31**	416.16 ± 2.31	**416.16 ± 2.31**	15.84 ± 0.35	-109.39 ± 0.55	88.28 ± 0.66	35.69 ± 0.41
fMRI	TR	9.94 ± 0.04	9.93 ± 0.04	9.88 ± 0.04	10.47 ± 0.05	8.96 ± 0.04	16.62 ± 0.06	39.21 ± 0.13
	GS	701.90 ± 0.72	701.47 ± 0.65	699.01 ± 0.63	621.08 ± 1.08	256.13 ± 1.63	667.14 ± 2.43	**528.66 ± 2.82**
	HTW	9.43 ± 0.61	9.43 ± 0.61	9.43 ± 0.61	6.75 ± 0.67	-1.09 ± 0.65	8.24 ± 0.66	3.45 ± 0.57
	TabWak	47.29 ± 0.83	43.52 ± 0.79	24.95 ± 0.86	-57.42 ± 2.04	71.03 ± 3.25	78.43 ± 4.67	0.17 ± 8.00
	TabWak$^\top$	**1031.96 ± 0.79**	**1031.53 ± 0.78**	**1030.14 ± 0.72**	**889.62 ± 0.84**	**383.15 ± 1.37**	801.16 ± 0.83	229.37 ± 1.15
	TimeWak	1006.16 ± 11.99	1006.12 ± 12.01	1005.87 ± 12.04	786.60 ± 8.42	226.03 ± 2.09	**810.73 ± 10.75**	277.53 ± 4.53
	FreqWak	676.78 ± 1.16	676.78 ± 1.16	676.78 ± 1.16	298.56 ± 0.22	29.65 ± 0.05	399.39 ± 0.67	157.69 ± 0.26

Table 11: Ablation on injection strength α (ETTh, ws=24).

α	Context-FID	Bit Accuracy
0.3	0.187	0.611
0.5	0.195	0.929
0.7	**0.197**	**1.000**
1.0	0.192	1.000

E.6 Extended Generation Quality Results

Tables 15 and 16 report the generation quality and watermark detectability for 64- and 128-length sequences, respectively. Following the format of Table 1, the quality metrics are reported for the corresponding window size, while the Z-scores are reported across all three window sizes (24, 64, 128) for reference.

Table 12: Ablation on injection frequency band (ETTh, ws=24).

Band	Context-FID	Bit Accuracy
Trend (low)	0.171	0.500
Season (mid)	1.001	1.000
Residual (high)	**0.197**	**1.000**
All bands	1.089	1.000

Table 13: FreqWak across architectures: Transformer vs. MLP on ETTh (ws=64).

Metric	**Transformer**	**MLP**
Detection BA (no attack)	1.000	**1.000**
False Positive (unwatermarked)	0.498	0.499
Offset Attack BA (factor=0.3)	1.000	**1.000**
No-WM FID	0.317 ± 0.019	1.193
FreqWak FID	0.346 ± 0.022	1.250
ΔFID (FW $-$ NoWM)	+0.029	+0.057

Table 14: Reproduction check for TimeWak detection (Z-score) under bits $= 3$. Our 3-run reproduction is compared against the original TimeWak paper's bits $= 3$ ablation (Tables 16–18 therein (Soi et al., 2025)). Values agree within $\pm 15\%$, confirming that the bits $= 3$ configuration used in the main results is faithfully reproduced.

Dataset	**Window**	**Ours (bits $= 3$)**	**Original (bits $= 3$)**
Stocks	24	190.69	194.87
	64	343.49	334.15
	128	464.66	459.23
ETTh	24	161.29	162.25
	64	264.59	249.67
	128	383.61	352.26
MuJoCo	24	82.73	73.09
	64	74.49	84.07
	128	156.87	183.47
Energy	24	286.02	268.24
	64	295.30	323.04
	128	413.62	380.35
fMRI	24	445.18	456.02
	64	694.98	724.69
	128	1006.16	986.17

Table 15: Extended generation quality and watermark detectability results for 64-length sequences. Quality metrics are for 64-length sequences while Z-scores are reported across 24-, 64-, and 128-length sequences. All metrics are 3-run mean ± std. Best in **bold**, second-best underlined.

		Quality Metric ↓				Z-score ↑		
Dataset	**Method**	**Context-FID**	**Correlational**	**Discriminative**	**Predictive**	**24-length**	**64-length**	**128-length**
Stocks	W/O	0.171 ± 0.005	0.018 ± 0.006	0.093 ± 0.010	0.023 ± 0.000	–	–	–
	TR	1.812 ± 0.210	0.102 ± 0.012	0.206 ± 0.038	0.038 ± 0.001	0.43 ± 0.04	0.40 ± 0.08	0.08 ± 0.10
	GS	1.643 ± 0.126	0.011 ± 0.005	0.252 ± 0.072	0.042 ± 0.001	86.07 ± 0.74	148.92 ± 1.08	172.23 ± 1.08
	HTW	0.426 ± 0.073	0.026 ± 0.016	0.105 ± 0.032	0.037 ± 0.000	4.45 ± 0.62	7.34 ± 0.94	10.39 ± 1.33
	TabWak	0.242 ± 0.045	0.011 ± 0.008	0.141 ± 0.021	0.037 ± 0.000	-67.22 ± 1.17	16.14 ± 0.89	-10.49 ± 1.28
	TabWak$^\top$	0.284 ± 0.091	**0.005 ± 0.005**	**0.099 ± 0.025**	0.037 ± 0.000	55.39 ± 0.86	88.81 ± 0.82	129.39 ± 0.90
	TimeWak	0.268 ± 0.015	0.022 ± 0.003	0.141 ± 0.035	**0.023 ± 0.000**	190.69 ± 0.75	**343.49 ± 0.93**	464.66 ± 4.29
	FreqWak	**0.205 ± 0.011**	0.020 ± 0.011	0.113 ± 0.062	**0.023 ± 0.000**	**253.59 ± 1.36**	341.52 ± 1.25	**575.31 ± 3.01**
ETTh	W/O	0.355 ± 0.020	0.081 ± 0.004	0.100 ± 0.001	0.116 ± 0.002	–	–	–
	TR	2.224 ± 0.209	0.217 ± 0.011	0.284 ± 0.032	0.131 ± 0.004	7.84 ± 0.12	7.73 ± 0.13	6.18 ± 0.16
	GS	3.398 ± 0.384	0.248 ± 0.025	0.356 ± 0.029	0.154 ± 0.001	101.07 ± 1.17	197.41 ± 2.10	327.47 ± 4.44
	HTW	0.372 ± 0.027	0.076 ± 0.006	0.118 ± 0.024	0.116 ± 0.006	3.43 ± 0.83	5.08 ± 1.48	6.84 ± 2.22
	TabWak	0.597 ± 0.044	0.346 ± 0.027	0.176 ± 0.020	0.122 ± 0.007	-14.95 ± 1.08	-6.16 ± 1.18	-20.57 ± 0.96
	TabWak$^\top$	0.503 ± 0.045	0.095 ± 0.043	0.119 ± 0.012	0.120 ± 0.010	109.35 ± 0.91	162.44 ± 1.11	235.03 ± 1.48
	TimeWak	**0.316 ± 0.015**	0.173 ± 0.001	**0.096 ± 0.008**	0.120 ± 0.009	161.29 ± 3.29	264.59 ± 1.81	383.61 ± 5.84
	FreqWak	0.334 ± 0.016	**0.075 ± 0.014**	0.144 ± 0.041	**0.115 ± 0.008**	**204.85 ± 1.49**	**356.10 ± 3.90**	**480.23 ± 1.43**
MuJoCo	W/O	0.096 ± 0.003	0.325 ± 0.027	0.028 ± 0.014	0.008 ± 0.000	–	–	–
	TR	2.627 ± 0.230	0.948 ± 0.064	0.270 ± 0.128	0.016 ± 0.005	1.38 ± 0.03	1.49 ± 0.04	1.31 ± 0.04
	GS	7.162 ± 1.301	1.034 ± 0.087	0.447 ± 0.011	0.011 ± 0.002	21.13 ± 0.77	10.13 ± 0.62	39.63 ± 0.71
	HTW	0.316 ± 0.034	0.334 ± 0.044	0.248 ± 0.079	0.011 ± 0.001	2.89 ± 0.54	3.41 ± 0.96	4.20 ± 1.37
	TabWak	0.372 ± 0.064	0.671 ± 0.083	0.137 ± 0.027	0.007 ± 0.001	31.30 ± 1.07	1.26 ± 0.96	-3.00 ± 1.04
	TabWak$^\top$	0.238 ± 0.019	0.339 ± 0.059	0.087 ± 0.025	**0.006 ± 0.001**	-4.85 ± 0.87	-4.51 ± 0.85	3.91 ± 0.88
	TimeWak	0.124 ± 0.001	0.378 ± 0.018	**0.060 ± 0.003**	0.008 ± 0.001	82.73 ± 9.29	74.49 ± 1.57	156.87 ± 8.47
	FreqWak	**0.114 ± 0.003**	**0.333 ± 0.011**	0.498 ± 0.002	0.009 ± 0.000	**371.71 ± 2.78**	**437.83 ± 1.24**	**544.37 ± 2.01**
Energy	W/O	0.109 ± 0.005	0.999 ± 0.077	0.131 ± 0.007	0.251 ± 0.000	–	–	–
	TR	0.902 ± 0.093	3.503 ± 0.589	0.427 ± 0.072	0.307 ± 0.005	9.51 ± 0.09	17.29 ± 0.11	22.58 ± 0.19
	GS	2.205 ± 0.192	3.277 ± 0.129	0.479 ± 0.010	0.310 ± 0.005	51.22 ± 0.88	68.67 ± 1.20	45.42 ± 1.05
	HTW	0.133 ± 0.015	**1.045 ± 0.357**	**0.135 ± 0.013**	**0.251 ± 0.001**	3.06 ± 0.36	4.30 ± 0.68	5.42 ± 1.00
	TabWak	0.168 ± 0.021	1.811 ± 0.530	0.136 ± 0.013	**0.251 ± 0.000**	3.26 ± 0.89	3.86 ± 1.02	0.57 ± 0.87
	TabWak$^\top$	0.237 ± 0.029	1.321 ± 0.252	0.138 ± 0.015	0.252 ± 0.000	40.82 ± 0.81	46.68 ± 0.86	26.00 ± 1.12
	TimeWak	0.174 ± 0.017	1.614 ± 0.298	0.163 ± 0.006	**0.251 ± 0.000**	**286.02 ± 1.88**	295.30 ± 5.17	413.62 ± 7.25
	FreqWak	**0.122 ± 0.006**	1.156 ± 0.212	0.497 ± 0.003	0.289 ± 0.001	182.63 ± 0.90	**394.44 ± 0.17**	**416.16 ± 2.31**
fMRI	W/O	0.512 ± 0.001	1.906 ± 0.043	0.341 ± 0.089	0.100 ± 0.000	–	–	–
	TR	3.358 ± 0.402	13.699 ± 0.132	0.411 ± 0.158	0.141 ± 0.001	6.49 ± 0.05	8.25 ± 0.05	9.94 ± 0.04
	GS	0.756 ± 0.098	8.378 ± 0.028	0.500 ± 0.001	0.104 ± 0.001	420.02 ± 1.44	321.52 ± 0.59	701.90 ± 0.72
	HTW	**0.413 ± 0.017**	1.822 ± 0.111	0.338 ± 0.032	**0.100 ± 0.000**	4.32 ± 0.22	6.80 ± 0.41	9.43 ± 0.61
	TabWak	0.655 ± 0.137	6.532 ± 0.158	0.242 ± 0.302	0.108 ± 0.001	84.02 ± 1.04	204.16 ± 0.82	47.29 ± 0.83
	TabWak$^\top$	0.554 ± 0.049	1.955 ± 0.058	0.331 ± 0.039	**0.100 ± 0.000**	**464.67 ± 0.50**	**743.33 ± 0.55**	**1031.96 ± 0.79**
	TimeWak	0.508 ± 0.002	**1.796 ± 0.015**	**0.214 ± 0.169**	**0.100 ± 0.000**	445.18 ± 0.59	694.98 ± 10.76	1006.16 ± 11.99
	FreqWak	0.491 ± 0.006	1.909 ± 0.035	0.496 ± 0.001	0.102 ± 0.000	328.36 ± 1.41	502.15 ± 1.16	676.78 ± 1.16

Table 16: Extended generation quality and watermark detectability results for 128-length sequences. Quality metrics are for 128-length sequences while Z-scores are reported across 24-, 64-, and 128-length sequences. All metrics are 3-run mean ± std. Best in **bold**, second-best underlined.

		Quality Metric ↓				Z-score ↑		
Dataset	**Method**	**Context-FID**	**Correlational**	**Discriminative**	**Predictive**	**24-length**	**64-length**	**128-length**
Stocks	W/O	0.192 ± 0.043	0.029 ± 0.005	0.078 ± 0.027	0.021 ± 0.000	–	–	–
	TR	3.555 ± 0.868	0.130 ± 0.025	0.276 ± 0.030	0.041 ± 0.004	0.43 ± 0.04	0.40 ± 0.08	0.08 ± 0.10
	GS	2.647 ± 0.262	0.025 ± 0.015	0.186 ± 0.071	0.040 ± 0.000	86.07 ± 0.74	148.92 ± 1.08	172.23 ± 1.08
	HTW	0.622 ± 0.092	0.017 ± 0.013	0.152 ± 0.086	0.037 ± 0.000	4.45 ± 0.62	7.34 ± 0.94	10.39 ± 1.33
	TabWak	0.394 ± 0.034	**0.009 ± 0.009**	0.150 ± 0.031	0.037 ± 0.000	-67.22 ± 1.17	16.14 ± 0.89	-10.49 ± 1.28
	TabWak$^\top$	0.367 ± 0.035	0.010 ± 0.007	0.129 ± 0.054	0.037 ± 0.000	55.39 ± 0.86	88.81 ± 0.82	129.39 ± 0.90
	TimeWak	0.303 ± 0.035	0.030 ± 0.003	**0.086 ± 0.057**	**0.021 ± 0.000**	190.69 ± 0.75	**343.49 ± 0.93**	464.66 ± 4.29
	FreqWak	**0.184 ± 0.016**	0.025 ± 0.010	0.091 ± 0.044	**0.021 ± 0.000**	**253.59 ± 1.36**	341.52 ± 1.25	**575.31 ± 3.01**
ETTh	W/O	1.202 ± 0.036	0.099 ± 0.012	0.180 ± 0.010	0.115 ± 0.004	–	–	–
	TR	2.450 ± 0.128	0.270 ± 0.017	0.284 ± 0.063	0.138 ± 0.005	7.84 ± 0.12	7.73 ± 0.13	6.18 ± 0.16
	GS	4.998 ± 0.603	0.233 ± 0.027	0.381 ± 0.042	0.162 ± 0.010	101.07 ± 1.17	197.41 ± 2.10	327.47 ± 4.44
	HTW	**1.119 ± 0.065**	**0.091 ± 0.014**	0.165 ± 0.010	0.120 ± 0.006	3.43 ± 0.83	5.08 ± 1.48	6.84 ± 2.22
	TabWak	1.931 ± 0.254	0.466 ± 0.058	0.248 ± 0.017	0.128 ± 0.005	-14.95 ± 1.08	-6.16 ± 1.18	-20.57 ± 0.96
	TabWak$^\top$	1.477 ± 0.075	0.129 ± 0.037	**0.143 ± 0.014**	**0.112 ± 0.011**	109.35 ± 0.91	162.44 ± 1.11	235.03 ± 1.48
	TimeWak	1.173 ± 0.014	0.148 ± 0.010	0.155 ± 0.007	0.120 ± 0.001	161.29 ± 3.29	264.59 ± 1.81	383.61 ± 5.84
	FreqWak	1.186 ± 0.055	0.098 ± 0.011	0.181 ± 0.004	0.125 ± 0.009	**204.85 ± 1.49**	**356.10 ± 3.90**	**480.23 ± 1.43**
MuJoCo	W/O	0.134 ± 0.005	0.318 ± 0.036	0.040 ± 0.014	0.006 ± 0.001	–	–	–
	TR	2.348 ± 0.236	0.865 ± 0.025	0.358 ± 0.008	0.008 ± 0.001	1.38 ± 0.03	1.49 ± 0.04	1.31 ± 0.04
	GS	4.797 ± 1.290	1.335 ± 0.041	0.454 ± 0.023	0.007 ± 0.001	21.13 ± 0.77	10.13 ± 0.62	39.63 ± 0.71
	HTW	0.431 ± 0.051	**0.309 ± 0.064**	0.255 ± 0.174	0.009 ± 0.001	2.89 ± 0.54	3.41 ± 0.96	4.20 ± 1.37
	TabWak	0.369 ± 0.073	0.589 ± 0.059	0.175 ± 0.042	**0.006 ± 0.002**	31.30 ± 1.07	1.26 ± 0.96	-3.00 ± 1.04
	TabWak$^\top$	0.275 ± 0.033	0.333 ± 0.066	0.084 ± 0.028	**0.006 ± 0.001**	-4.85 ± 0.87	-4.51 ± 0.85	3.91 ± 0.88
	TimeWak	0.154 ± 0.003	0.360 ± 0.040	**0.053 ± 0.007**	**0.006 ± 0.001**	82.73 ± 9.29	74.49 ± 1.57	156.87 ± 8.47
	FreqWak	**0.140 ± 0.014**	0.314 ± 0.012	0.322 ± 0.279	**0.006 ± 0.001**	**371.71 ± 2.78**	**437.83 ± 1.24**	**544.37 ± 2.01**
Energy	W/O	0.108 ± 0.006	0.830 ± 0.044	0.203 ± 0.018	0.249 ± 0.000	–	–	–
	TR	1.195 ± 0.111	2.303 ± 0.569	0.498 ± 0.001	0.287 ± 0.003	9.51 ± 0.09	17.29 ± 0.11	22.58 ± 0.19
	GS	3.680 ± 0.444	4.233 ± 0.287	0.474 ± 0.039	0.280 ± 0.006	51.22 ± 0.88	68.67 ± 1.20	45.42 ± 1.05
	HTW	**0.133 ± 0.017**	**0.822 ± 0.310**	**0.103 ± 0.043**	**0.249 ± 0.001**	3.06 ± 0.36	4.30 ± 0.68	5.42 ± 1.00
	TabWak	0.201 ± 0.020	2.001 ± 0.440	0.156 ± 0.077	0.250 ± 0.001	3.26 ± 0.89	3.86 ± 1.02	0.57 ± 0.87
	TabWak$^\top$	0.274 ± 0.019	1.211 ± 0.196	0.136 ± 0.023	**0.249 ± 0.001**	40.82 ± 0.81	46.68 ± 0.86	26.00 ± 1.12
	TimeWak	0.248 ± 0.008	1.291 ± 0.078	0.143 ± 0.050	0.250 ± 0.001	**286.02 ± 1.88**	295.30 ± 5.17	413.62 ± 7.25
	FreqWak	0.152 ± 0.010	0.941 ± 0.069	0.489 ± 0.019	0.278 ± 0.001	182.63 ± 0.90	**394.44 ± 0.17**	**416.16 ± 2.31**
fMRI	W/O	0.843 ± 0.004	1.771 ± 0.035	0.413 ± 0.021	0.100 ± 0.000	–	–	–
	TR	5.391 ± 0.919	13.000 ± 0.084	0.434 ± 0.072	0.149 ± 0.002	6.49 ± 0.05	8.25 ± 0.05	9.94 ± 0.04
	GS	1.046 ± 0.052	5.941 ± 0.048	0.499 ± 0.001	0.106 ± 0.001	420.02 ± 1.44	321.52 ± 0.59	701.90 ± 0.72
	HTW	0.811 ± 0.051	1.757 ± 0.053	**0.063 ± 0.064**	**0.100 ± 0.000**	4.32 ± 0.22	6.80 ± 0.41	9.43 ± 0.61
	TabWak	1.114 ± 0.135	5.967 ± 0.033	0.153 ± 0.218	0.111 ± 0.001	84.02 ± 1.04	204.16 ± 0.82	47.29 ± 0.83
	TabWak$^\top$	0.919 ± 0.024	1.807 ± 0.024	0.265 ± 0.201	**0.100 ± 0.000**	**464.67 ± 0.50**	**743.33 ± 0.55**	**1031.96 ± 0.79**
	TimeWak	0.837 ± 0.001	**1.698 ± 0.013**	0.360 ± 0.082	**0.100 ± 0.000**	445.18 ± 0.59	694.98 ± 10.76	1006.16 ± 11.99
	FreqWak	**0.784 ± 0.003**	1.792 ± 0.024	0.498 ± 0.001	0.102 ± 0.000	328.36 ± 1.41	502.15 ± 1.16	676.78 ± 1.16