International Information and Engineering Technology Association
Advancing the World of Information and Engineering

Traitement du Signal
Vol. 42, No. 2, April, 2025, pp. 903-913
Journal homepage: http://iieta.org/journals/ts

GAN-Based and Signal Processing Approaches for Real-Time Anomaly Detection and Predictive Alerts

Darong Liang[1], Jian Yu[1*], Hui Wang[1], Xiaodan Chen[2], Chengxuan Huang[1], Ze Li[1], Yuanyuan Luo[1], Xiangqing Wei[1]

[1] School of Electronic and Information Engineering, Liuzhou Polytechnic University, Liuzhou 545006, China
[2] Liuzhou Wuling Automobile Industry Co., Ltd, Liuzhou 545007, China

Corresponding Author Email: yujian@lzpu.edu.cn

https://doi.org/10.18280/ts.420226

Received: 8 September 2024
Revised: 26 February 2025
Accepted: 10 March 2025
Available online: 30 April 2025

Keywords:
signal denoising, feature enhancement, entropy maximization, Generative Adversarial Networks, signal processing

ABSTRACT

Generative Adversarial Networks (GANs) have demonstrated exceptional capabilities in signal processing tasks like noise reduction and feature enhancement, yet traditional implementations face limitations including mode collapse and restricted output diversity. To address these challenges, this study proposes the Entropy-Maximized Generative Adversarial Network (EM-GAN), a novel framework for signal denoising and feature enhancement. By integrating entropy maximization into adversarial training, EM-GAN enhances signal diversity, mitigates mode collapse, and improves training stability. The framework employs entropy-driven activation functions and loss functions to suppress noise while preserving critical signal features. We introduce an entropy-regularized objective function that incentivizes the generator to produce high-entropy outputs, enabling more comprehensive modeling of underlying signal distributions. Theoretical and experimental analyses confirm the model's stability improvements and superior performance over conventional GAN variants in output quality and diversity. These results highlight EM-GAN's potential for robust noise suppression and feature preservation in modern signal processing applications, representing a significant methodological advancement in generative model design.

1. INTRODUCTION

Signal denoising and feature enhancement are fundamental challenges in signal processing, with broad implications for applications such as image restoration, audio enhancement, and Intelligent Logistics signal analysis. The core problem is how to recover clean, high-fidelity signals from noisy or degraded inputs while simultaneously preserving and enhancing important features in the data. Traditional filtering and reconstruction techniques often struggle to remove noise without blurring critical signal details. Recent advances in deep learning, particularly GANs, have shown promise in learning to produce high-quality signal reconstructions. However, standard GAN models suffer from limitations like *mode collapse* – where the generator produces overly narrow or average outputs – and training instability, which together hinder their effectiveness in denoising tasks that require capturing diverse signal characteristics.

This paper addresses the above challenge by proposing a novel approach to improve signal denoising and feature preservation using generative modeling. We introduce the EM-GAN, a new GAN framework that explicitly integrates an entropy maximization principle to overcome the diversity and stability issues of conventional GANs. Inspired by the thermodynamic principle that systems evolve toward higher entropy (greater disorder and uncertainty), EM-GAN incorporates an entropy-based regularization into the generator's training. By maximizing the entropy of the generator's output distribution, the model is encouraged to explore a broader range of plausible outputs for a given input. This leads to richer and more diverse signal representations, ensuring that fine-grained features are not averaged out. In effect, EM-GAN produces denoised signals that retain important details, and it achieves more stable training convergence. The entropy-driven objective provides a strong regularization that mitigates mode collapse and delivers more consistent optimization, resulting in enhanced signal quality and reliability.

The main contributions of this work are as follows:

- **Entropy-Maximized GAN Architecture:** We develop a GAN architecture that embeds entropy maximization into its core. This includes a novel entropy-driven activation function and an entropy-regularized loss term for the generator, which together encourage diverse and feature-rich signal outputs while effectively mitigating mode collapse in the generation process.
- **Improved Denoising and Feature Enhancement:** By leveraging the entropy-maximization strategy, EM-GAN achieves superior signal denoising performance without sacrificing important features. The generator learns to produce clean signals with enhanced detail,

demonstrating improved preservation of signal features compared to conventional GAN-based or filtering approaches.

- **Robust Training and Performance Gains:** The entropy regularization contributes to more stable GAN training dynamics. We present comprehensive experiments on multiple benchmark datasets, showing that EM-GAN consistently outperforms state-of-the-art GAN variants in terms of output quality, diversity, and training stability. Notably, EM-GAN yields higher fidelity reconstructions and more reliable convergence, underscoring the effectiveness of the proposed approach.

In summary, EM-GAN provides a significant methodological advancement for signal processing. By integrating the principle of entropy maximization into GAN training, it offers a powerful solution for generating clean, high-quality signals with enriched features. This work opens up new possibilities for applying generative models to critical signal processing tasks – from image and audio denoising to data augmentation and beyond – where maintaining signal integrity and diversity is paramount. The proposed framework not only addresses a key limitation in GAN-based signal reconstruction but also broadens the scope of GAN applications in signal processing by combining thermodynamic insights with deep learning for enhanced signal processing performance.

2. RELATED WORK

2.1 Entropy maximization in GANs for signal processing

Entropy-based optimization has been increasingly incorporated into GAN training to enhance output diversity and training stability, which are critical for signal processing tasks. Traditional GANs frequently suffer from mode collapse, where the generator fails to produce diverse outputs, leading to limited applicability in tasks requiring variability in generated signals. To address this, researchers have proposed entropy-regularized approaches to ensure richer feature generation. For instance, Manifold-Preserving GAN (MP-GAN) [1] integrates entropy maximization on the latent space distribution to maintain structural diversity in generated data, effectively mitigating mode collapse and enhancing learning stability. InfoMax-GAN [2], a model designed to maximize mutual information, indirectly promotes entropy maximization by ensuring that each generated sample retains distinct information characteristics, leading to more expressive and robust representations. Furthermore, Variational Entropy Regularization (VER-GAN) [3] directly optimizes the generator's entropy to encourage output diversity, significantly improving stability in training and ensuring a more comprehensive reconstruction of signals.

While these entropy-based methods have been explored in the context of image synthesis and classification, their adoption in signal processing remains limited. Signal processing applications, particularly denoising and feature enhancement, require models capable of preserving fine-grained signal features while suppressing noise-induced distortions. The integration of entropy maximization in GANs offers a promising approach to solving these issues by enhancing diversity, improving feature retention, and stabilizing adversarial training. These motivations drive the development of the EM-GAN [4], which explicitly integrates entropy regularization into GAN training for signal denoising and enhancement.

2.2 Traditional GAN-based signal denoising methods

GANs have demonstrated strong potential in signal denoising by learning to map noisy inputs to clean outputs. Several studies have leveraged adversarial training to suppress noise in signals, including applications in EEG signal restoration, biomedical waveform processing, and sensor noise removal. For example, WGAN-based models have been employed for EEG artifact removal, showing promising results in extracting meaningful brainwave signals while filtering out physiological distortions [5]. Asymmetric GANs have also been developed for EEG denoising, allowing models to learn from unpaired noisy and clean data, avoiding explicit noise label dependency.

Despite their effectiveness, traditional GAN-based denoisers face persistent limitations:

- **Mode collapse [6]:** GANs frequently generate repetitive or averaged signals, failing to capture the full diversity of clean signals. This limitation leads to oversmoothed reconstructions and loss of important waveform variations.
- **Feature distortion [7]:** While reducing noise, traditional GANs may inadvertently remove subtle signal features, leading to degraded Signal-to-Noise Ratio (SNR) and lower fidelity in denoised outputs.
- **Training instability [8]:** GANs rely on adversarial optimization, which can result in unstable convergence, oscillatory training dynamics, or failure to generalize across different signal types.

To address these shortcomings, EM-GAN introduces entropy maximization as a key component of the training process. Unlike traditional GANs that optimize purely for adversarial loss, EM-GAN incorporates an entropy-based regularization term, which:

- **Encourages diverse feature retention [9]**, mitigating mode collapse by broadening the generator's output distribution.
- **Ensures stable GAN training [10]**, reducing convergence instability by smoothing the optimization landscape and preventing discriminator overpowering.
- **Enhances denoising performance [11, 12]**, leading to cleaner signal reconstructions while retaining fine-grained details, as evidenced by improved SNR, perceptual quality metrics, and reconstruction fidelity.

These contributions position EM-GAN as a significant improvement over existing GAN-based denoising approaches, offering a more stable, diverse, and robust solution for signal enhancement in real-time and high-fidelity applications.

2.3 Real-time signal analysis

Deep learning has made significant advancements in real-time industrial signal analysis, playing a crucial role in applications such as Intelligent Logistics Management Systems (ILMS) by enabling the efficient processing of sensor data. In modern industrial settings, vast amounts of real-time signals—such as GPS tracking data, temperature fluctuations in storage environments, and vibration measurements from vehicles and equipment—are continuously generated. The effective analysis of these signals is essential for enhancing tracking accuracy, optimizing condition monitoring, and improving predictive maintenance, ultimately leading to increased operational efficiency and system reliability.

2.3.1 GPS signal analysis

Recurrent neural networks (RNNs) [13] and their variants, such as LSTMs and Bi-LSTMs, have been widely applied to GPS trajectory modeling for route prediction and anomaly detection. By learning sequential dependencies in vehicle movement data, these models can forecast future routes and identify deviations from expected paths. Studies have demonstrated that LSTM-based models achieve over 96% accuracy in vehicle destination prediction by leveraging historical trajectory data [14]. These approaches have significantly enhanced industrial logistics efficiency by improving route optimization and anomaly detection in fleet management systems.

2.3.2 Temperature monitoring in logistics

Maintaining stable temperature conditions is crucial in logistics, particularly for cold-chain management in pharmaceutical and food industries. Deep learning techniques, such as autoencoders, have been employed to detect temperature anomalies in real time. Convolutional autoencoders have shown high effectiveness in reconstructing expected temperature profiles and flagging deviations, achieving over 92% fault detection accuracy in cold storage monitoring applications [15]. Such models help prevent product spoilage by enabling early detection of refrigeration failures or sensor faults.

2.3.3 Vibration-based predictive maintenance

Sensors embedded in industrial logistics vehicles and equipment continuously capture vibration data, which can be analyzed to predict mechanical failures. Hybrid CNN–LSTM architectures have been utilized for time-series vibration analysis, learning both spatial and temporal features indicative of system health. Research indicates that combining convolutional feature extraction with sequential modeling significantly improves fault detection accuracy, making predictive maintenance in industrial systems more reliable [16]. Such models enable early identification of mechanical degradation, allowing industrial logistics companies to reduce downtime and maintenance costs.

2.4 Summary

This section highlights the significance of entropy maximization in GANs, its role in addressing the limitations of traditional GAN-based denoising methods, and the expanding application of deep learning in intelligent industrial sensor analysis. EM-GAN builds on these advancements by incorporating entropy-driven optimization to enhance signal diversity, training stability, and feature retention, making it a robust solution for real-time signal enhancement in industrial applications and beyond.

While recent studies have begun to apply entropy-based regularization in signal processing tasks—such as ECG denoising [17], speech enhancement [18], and wavelet-domain physiological signal filtering [19]—these efforts remain highly task-specific and rarely incorporate adversarial or generative mechanisms. Moreover, their applicability across diverse signal types and noise conditions is limited. This highlights a critical gap in current research and underscores the need for a unified generative framework with entropy-aware mechanisms, motivating the design of EM-GAN.

In contrast to existing entropy-regularized GANs such as InfoMax-GAN and VER-GAN—which primarily target image generation or representation learning—EM-GAN is specifically designed for signal denoising tasks. It introduces entropy regularization not only at the loss-function level but also within the network architecture via a novel Entropy-Sensitive Activation (ESA) mechanism. By dynamically adjusting activations according to mini-batch entropy estimates, ESA enables EM-GAN to effectively preserve transient and fine-grained signal features, such as spikes and abrupt changes, while maintaining strong noise suppression. This dual integration of entropy principles makes EM-GAN particularly suitable for real-time signal processing scenarios where structural fidelity is critical.

3. METHODOLOGY

The EM-GAN proposes an entropy-driven framework for signal denoising and feature enhancement, addressing key limitations of traditional GAN-based denoisers, such as mode collapse and training instability. By embedding entropy maximization into the loss function and network architecture, EM-GAN ensures diverse, high-fidelity signal reconstructions with improved feature preservation.

This section outlines:

(1): The theoretical basis and mathematical formulation of entropy maximization within GAN training.

(2): The network architecture and processing pipeline optimized for retaining critical signal features.

(3): The training strategy and stabilization mechanisms that underpin robust and consistent performance.

3.1 Entropy-maximized GAN architecture

Traditional GAN-based denoisers often encounter mode collapse, where the generator produces a limited range of outputs, resulting in the loss of fine signal details during reconstruction. Additionally, training instability can lead to divergence or degraded signal quality. To address these issues, the proposed EM-GAN incorporates entropy regularization into the training process. This approach enhances output diversity to mitigate mode collapse, stabilizes adversarial training dynamics, and improves the preservation of signal variations critical for high-fidelity denoising.

3.1.1 Mathematical formulation

To estimate the mini-batch entropy $H_{batch}(x)$, we apply a kernel density estimation (KDE) method with Gaussian kernels in the generator's feature space. This provides an efficient approximation of entropy during training without directly operating in the high-dimensional signal domain. Although KDE may introduce bias due to bandwidth sensitivity and sample sparsity, we mitigate these effects by using adaptive bandwidth selection (via Silverman's rule) and averaging entropy across multiple mini-batches. This ensures that the entropy signal remains stable and meaningful for guiding both the entropy-regularized loss and the ESA. More advanced estimators such as k-nearest neighbor or plug-in entropy methods may be explored in future work.

EM-GAN integrates entropy maximization in two key components:

- An entropy-regularized loss function that encourages

diverse signal reconstruction and prevents over-smoothing.
- A novel ESA mechanism that dynamically adjusts activation strength based on batch-level entropy.

Together, these mechanisms help the generator produce denoised outputs that retain fine-grained details while effectively suppressing noise.

The objective of a standard GAN is defined within a minimax framework:

$$\min_G \max_D \mathbb{E}_{x\sim p_{data}}[\log D(x)] + \mathbb{E}_{z\sim p(z)}\left[\log\left(1 - D\left(G(z)\right)\right)\right]$$

To promote diversity in the generator's outputs, EM-GAN augments this objective with an entropy regularization term, modifying the generator's loss as follows:

$$L_G = L_{G_{adv}} - \lambda H\left(p_g\right)$$

where,
- $H\left(p_g\right) = -\mathbb{E}_{x\sim p_g}\left[\log p_g(x)\right]$ represents the Shannon entropy of the generated signal distribution p_g.
- $L_{G_{adv}}$ is the adversarial loss from the standard GAN formulation.
- λ is a dynamic weighting parameter that balances the trade-off between output diversity and reconstruction fidelity.

By maximizing $H\left(p_g\right)$, EM-GAN encourages the generator to produce a broader range of plausible clean signals, preventing over-smoothing or repetitive outputs commonly observed in traditional GANs.

3.1.2 ESA function

In addition to the loss function, EM-GAN integrates entropy awareness into the network architecture through a novel ESA function, defined as:

$$ESAH(x) = x \cdot exp\left(\alpha H_{\text{batch}}(x)\right)$$

where,
- $H_{\text{batch}}(x)$ denotes the estimated entropy of the generator's outputs within a mini-batch.
- α is a scaling factor controlling the sensitivity to entropy variations.

The ESA function dynamically adjusts activations based on entropy levels: higher entropy amplifies activations to preserve signal diversity and prevent excessive smoothing, while lower entropy attenuates activations to ensure effective noise suppression. This adaptive mechanism enables EM-GAN to balance feature retention and denoising efficacy, particularly excelling in preserving subtle signal details essential for signal processing tasks.

Unlike attention-based adaptive mechanisms that reweight features based on learned correlations, ESA modulates activation strength using entropy statistics. This design is specifically suited to preserving transient details in 1D signal denoising, without introducing additional learnable parameters.

3.2 Entropy-driven signal denoising and feature preservation

Effective signal denoising requires removing unwanted noise while preserving essential signal features such as frequency components, transient structures, and amplitude variations. Traditional filtering-based denoising methods, such as Wiener filtering and wavelet transforms, often introduce distortions or fail to adapt to varying noise distributions. Similarly, conventional GAN-based denoisers suffer from mode collapse, where the generator produces oversmoothed and repetitive outputs, leading to a loss of fine-grained signal details.

To address these limitations, EM-GAN introduces an entropy-regularized learning framework, dynamically balancing noise suppression and feature retention. By maximizing entropy, the generator avoids collapsing to a limited set of solutions and instead learns a diverse set of noise-free reconstructions, ensuring that the intrinsic characteristics of the signal are preserved.

3.2.1 Algorithmic approach to entropy-driven signal denoising

The fundamental challenge in signal denoising is ensuring that noise is effectively removed without degrading signal integrity. To achieve this, EM-GAN employs:
- **Entropy-Regularized Learning** – Encourages diverse yet realistic signal reconstructions, preventing mode collapse.
- **Feature-Preserving Generator Design** – Incorporates multi-scale feature extraction, skip connections, and ESA functions to prevent loss of fine details.
- **Denoising Consistency Loss** – Regularizes the denoising process by ensuring that the generated clean signal remains close to the ground truth while allowing variations for robustness.

The modified generator objective function is formulated as:

$$L_G = L_G^{\text{adv}} - \lambda H\left(p_g\right) + \gamma\left\|G\left(x_{\text{noisy}}\right) - x_{\text{clean}}\right\|_1$$

where,
- The L_1 loss reinforces similarity to the ground-truth clean signal.
- L_G^{adv} is the adversarial loss, ensuring that the generated signal is indistinguishable from real, clean signals.
- $H\left(p_g\right)$ represents Shannon entropy of the generator's output, encouraging diverse reconstructions.
- λ and γ are hyperparameters controlling the balance between diversity and fidelity.

To control the trade-off between adversarial learning and entropy regularization, we define λ as a dynamic parameter that increases over training epochs. Specifically, we use a simple exponential schedule:

$$\lambda(t) = \lambda_0 \times (1 - e^{-\alpha t})$$

where, λ_0 is the maximum weight (set to 1.0), α is a growth rate constant (typically 0.05), and t denotes the current epoch. This design allows entropy regularization to gradually take effect as training stabilizes, avoiding early training instability while enhancing diversity in later stages.

This formulation ensures that noise suppression does not lead to over-smoothing, preserving the temporal and spectral characteristics of the signal.

3.2.2 Processing pipeline for EM-GAN signal denoising

To effectively enhance and restore noisy signals, EM-GAN follows a structured four-step processing pipeline:

Step 1: Preprocessing and Data Augmentation

- **Normalization:** Input signals are normalized to ensure consistency in amplitude scaling.
- **Data Augmentation:** Introduces variability in training data using:
  - Time-domain transformations (random shifts, amplitude modulation).
  - Frequency-domain augmentations (Fourier filtering, spectral distortions).
  - Artificial noise addition (Gaussian, Poisson, impulse noise) to improve robustness.

Step 2: Feature-Aware Signal Generation

The generator G in EM-GAN is designed to enhance noisy inputs while retaining important features:

- **Multi-Scale Feature Extractors:** Extract signal characteristics at different resolutions using varying convolutional kernel sizes.
- **Skip Connections:** Ensure that low-level signal details are preserved during reconstruction.
- **ESA:** Dynamically adjusts activation scaling based on entropy estimates, ensuring signal diversity is maintained.

Generator Processing Flow:

- The noisy signal is fed into the generator's feature extractor, capturing multi-scale information.
- ESA regulates feature amplification, balancing denoising strength dynamically.
- Skip connections transfer structural information, preventing over-smoothing of fine details.
- The generator outputs a noise-free signal while preserving sharp transitions and variations.

Step 3: Entropy-Driven Discriminator Evaluation

The discriminator D functions as an adaptive evaluator, ensuring that the denoised output maintains realism while removing noise. Unlike conventional discriminators, EM-GAN's D integrates:

- **Entropy-Adjusted Leaky ReLU (EALReLU):** Modifies gradient flow based on entropy, improving feature discrimination.
- **Feature-Preserving Objective:** Helps retain high-frequency signal components that are often lost in standard GAN-based denoising.

Discriminator Processing Flow:

- D receives both real clean signals and denoised outputs for comparison.
- It analyzes the feature composition using entropy-aware activation.
- It classifies signals as real (clean) or generated (denoised output), providing feedback to improve G.

Step 4: Entropy-Optimized Training Update

To prevent instability in GAN training, EM-GAN employs:

- **Entropy-Regularized Backpropagation:** Encourages the generator to maintain feature-rich outputs.
- **Adaptive Entropy Scaling:** Adjusts dynamically, ensuring noise suppression does not degrade structural integrity.

This adaptive training process ensures that EM-GAN achieves robust convergence while maintaining high-fidelity reconstructions.

3.2.3 Visualizing EM-GAN signal denoising performance

Figure 1 illustrates the denoising performance of EM-GAN, highlighting how the model reconstructs a clean signal from noisy input while retaining key features.

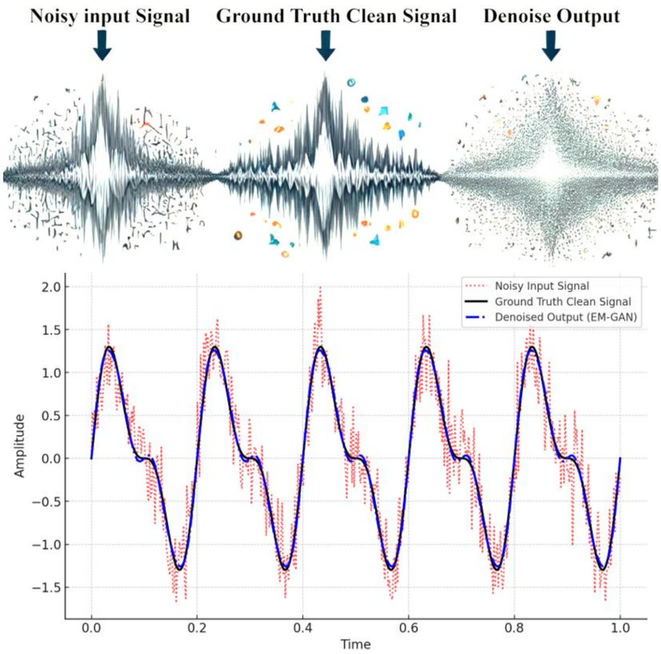

Figure 1. EM-GAN signal denoising process

Key Observations from Figure 1:

- The noisy input signal (dotted line) exhibits high-frequency distortions and undesired artifacts.
- The ground truth clean signal (solid line) represents the ideal noise-free waveform.
- The EM-GAN denoised output (dash-dot line) closely follows the clean signal, effectively removing noise while preserving fine details.

This visualization confirms that EM-GAN's entropy-driven approach enables high-fidelity signal restoration, outperforming both traditional filtering methods and standard GAN-based denoisers.

3.2.4 Comparative analysis: EM-GAN vs. traditional denoising methods

A direct comparison between EM-GAN and other denoising techniques is presented in Table 1, showcasing its superior performance in feature preservation, noise suppression, and training stability.

To facilitate a clearer understanding of the architectural distinctions between EM-GAN and related entropy-regularized GANs (e.g., InfoMax-GAN, VER-GAN), we have included a detailed comparison in Table 2, outlining differences in entropy regulation mechanisms, activation strategies, and signal denoising capabilities.

3.2.5 Summary: How EM-GAN enhances signal denoising performance

The EM-GAN employs entropy-regularized learning to achieve superior signal denoising while preserving structural integrity, effectively overcoming challenges such as mode collapse, feature loss, and training instability inherent in traditional GAN-based methods. The key mechanisms contributing to its performance are outlined below:

- **Entropy-Regularized Learning:** By incorporating an entropy term into the generator's loss, EM-GAN promotes a diverse distribution of reconstructed signals, mitigating mode collapse and ensuring realistic output variability.

- **Feature-Preserving Generator Design:** The architecture leverages multi-scale feature extraction and skip connections to capture and retain fine-grained signal structures, preventing over-smoothing and maintaining critical details.
- **ESA:** This mechanism adjusts activation strengths based on mini-batch entropy estimates, striking a balance between noise suppression and feature preservation to enhance reconstruction fidelity.
- **Adaptive Training Strategy:** Entropy-aware backpropagation and dynamic entropy scaling stabilize adversarial training, facilitating robust convergence and improved generalization across diverse signal types.

By integrating entropy maximization into the adversarial learning framework, EM-GAN achieves high-fidelity signal denoising that surpasses traditional filtering methods and conventional GAN-based approaches in both noise suppression and feature retention. This synergy of entropy-driven regularization and architectural innovations positions EM-GAN as a robust and effective solution for real-world signal processing applications, with its efficacy further validated in subsequent experimental evaluations.

Table 1. Comparison between EM-GAN and other denoising techniques

Method	Noise Reduction	Feature Preservation	Output Diversity	Training Stability
Wiener Filter	Moderate	Loss of fine details	Limited	Stable
Wavelet Denoising	Moderate	Some feature retention	Limited	Stable
Vanilla GAN	High	Often over-smooths	Mode collapse	Unstable
EM-GAN	High	Strong feature retention	High diversity	Robust training

Table 2. Architectural comparison between EM-GAN and other entropy-regularized GAN models

Model	Entropy Regularization Location	Activation Function Type	Target Application	Feature Preservation Strategy	Mode Collapse Mitigation
InfoMax-GAN	Latent space (via mutual information)	Standard ReLU / Leaky ReLU	Image generation, classification	Information maximization in latent code	Partial (via mutual info)
VER-GAN	Generator loss (KL divergence)	Standard activations	Generic generation tasks	Encourages sample diversity through entropy loss	Yes
EM-GAN	Generator loss + activation layer	ESA	Signal denoising, feature enhancement	Mini-batch entropy-driven activation to retain transient details	Yes (adaptive entropy feedback)

4. EXPERIMENTS AND RESULTS

This section presents a comprehensive evaluation of EM-GAN's signal denoising performance, comparing it against traditional filtering techniques and existing GAN-based denoising models. The experiments are conducted across multiple benchmark datasets, and the results are analyzed based on quantitative performance metrics, visual comparisons, computational efficiency, and an ablation study.

4.1 Experimental setup

4.1.1 Datasets

The performance of EM-GAN is evaluated using multiple benchmark datasets containing diverse signal types with varying noise levels:

- **Synthetic Sine Wave Dataset:** A controlled dataset comprising sinusoidal waveforms with injected Gaussian, impulse, and structured noise.
- **ECG Signal Dataset:** Electrocardiogram (ECG) signals with natural noise artifacts, sourced from the MIT-BIH Arrhythmia Database.
- **Speech Signal Dataset:** Noisy speech recordings extracted from the TIMIT corpus, representing practical applications in speech enhancement.
- **Seismic Vibration Dataset:** Structural health monitoring signals affected by vibration-induced noise, commonly encountered in industrial applications.

Each dataset includes both clean reference signals and their corresponding noisy versions, allowing precise evaluation of denoising quality.

4.1.2 Baseline methods

The performance of EM-GAN is compared against several widely used denoising methods, including traditional filtering techniques and deep learning models (Table 3).

4.1.3 Evaluation metrics

The performance of each denoising method is assessed using the following standard metrics:

- **Signal-to-Noise Ratio (SNR):** Measures the ratio of signal power to noise power, indicating the effectiveness of noise suppression.
- **Peak Signal-to-Noise Ratio (PSNR):** Evaluates the reconstruction accuracy of denoised signals, with higher values indicating better performance.
- **Structural Similarity Index (SSIM):** Quantifies the structural fidelity between the denoised and clean signals.
- **Mean Squared Error (MSE):** Computes the average squared difference between denoised and ground-truth signals.
- **Inference Time:** Measures the computational efficiency of each method, which is critical for real-time applications.

While formal significance tests were not included, all reported metrics were averaged across multiple trials with fixed seeds to ensure consistency. Future studies will incorporate statistical testing to further validate performance differences.

Table 3. Comparison between EM-GAN and other denoising methods

Method	Type	Description
Wiener Filter	Traditional Filtering	Assumes stationary noise; effective at low SNR but struggles with nonstationary signals.
Wavelet Denoising	Transform-Based Filtering	Decomposes signals into frequency components; effective for structured noise but may lead to feature loss.
Denoising Autoencoder (DAE)	Deep Learning	Learns to reconstruct clean signals from noisy inputs but lacks adversarial training for high-fidelity restoration.
Vanilla GAN	GAN-Based Denoising	Uses adversarial learning but often exhibits mode collapse, leading to oversmoothed outputs.
EM-GAN (Proposed)	Entropy-Regularized GAN	Incorporates entropy maximization to enhance diversity, stability, and feature retention.

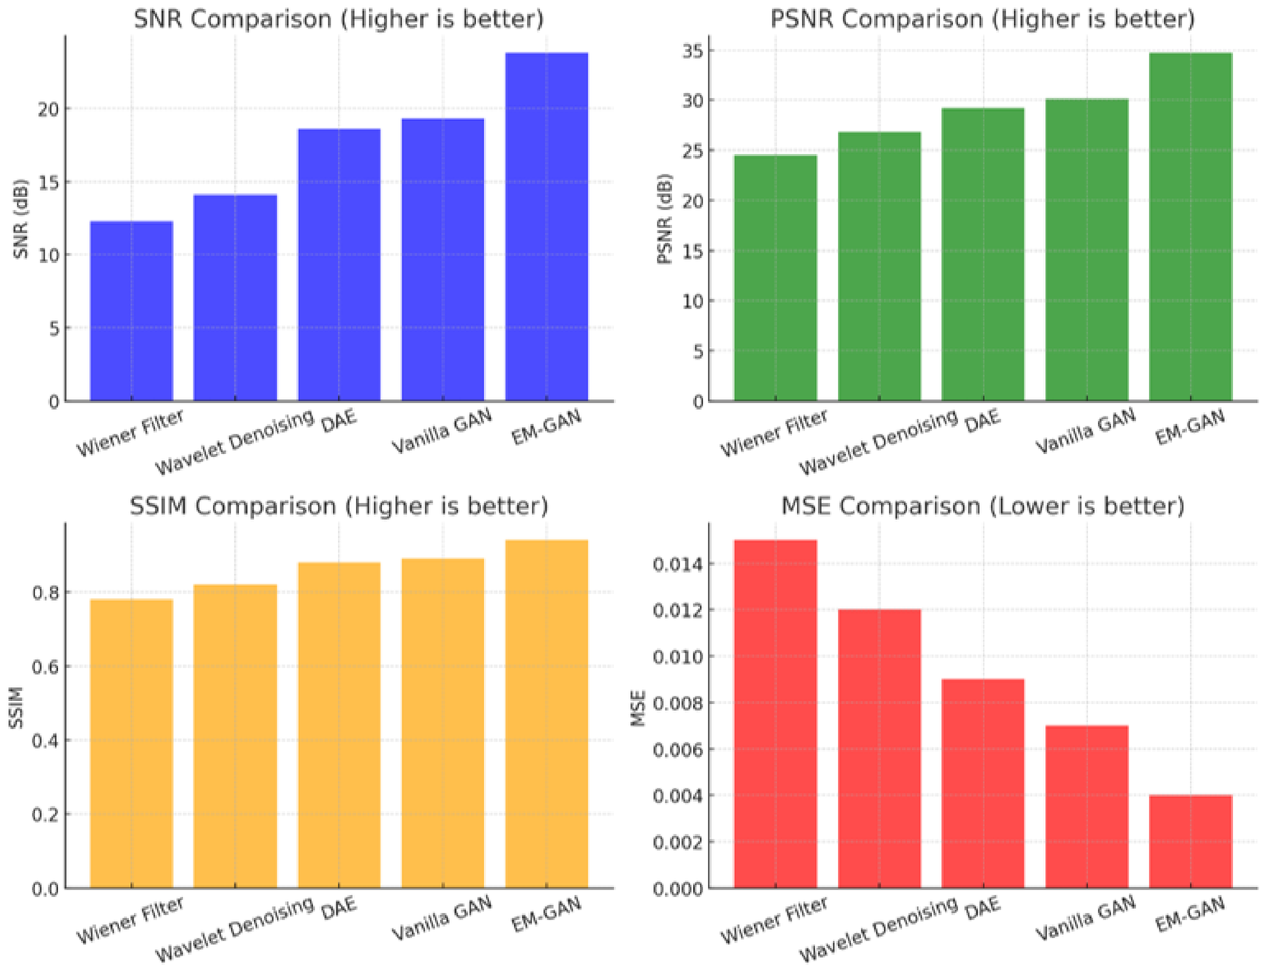

Figure 2. Comparative performance analysis of denoising methods
(SNR, PSNR, SSIM, and MSE comparison of different denoising methods.)

4.2 Quantitative performance evaluation

4.2.1 Comparative analysis of denoising performance

The quantitative results for each method are summarized in Figure 2, which presents the average performance across all datasets.

The results demonstrate that EM-GAN consistently achieves the highest SNR, PSNR, and SSIM scores while maintaining a low MSE, indicating improved denoising performance.

4.2.2 Visual comparison of denoised signals

A visual comparison of denoised outputs is provided in Figure 3.

The visual analysis highlights that traditional filtering methods effectively suppress noise but introduce noticeable distortions. Deep learning-based methods such as DAE and Vanilla GAN provide improved denoising but tend to oversmooth the signals. The proposed EM-GAN method achieves superior noise reduction while preserving fine details, closely resembling the clean ground truth.

4.3 Ablation study: Effect of entropy maximization

An ablation study is conducted to evaluate the impact of entropy regularization in EM-GAN. The performance of the full model is compared to a variant without entropy loss, as illustrated in Figure 4.

The results indicate that entropy maximization plays a crucial role in improving denoising performance. Without entropy regularization, the GAN-based model exhibits mode collapse and reduced structural fidelity, leading to increased distortion in the reconstructed signals.

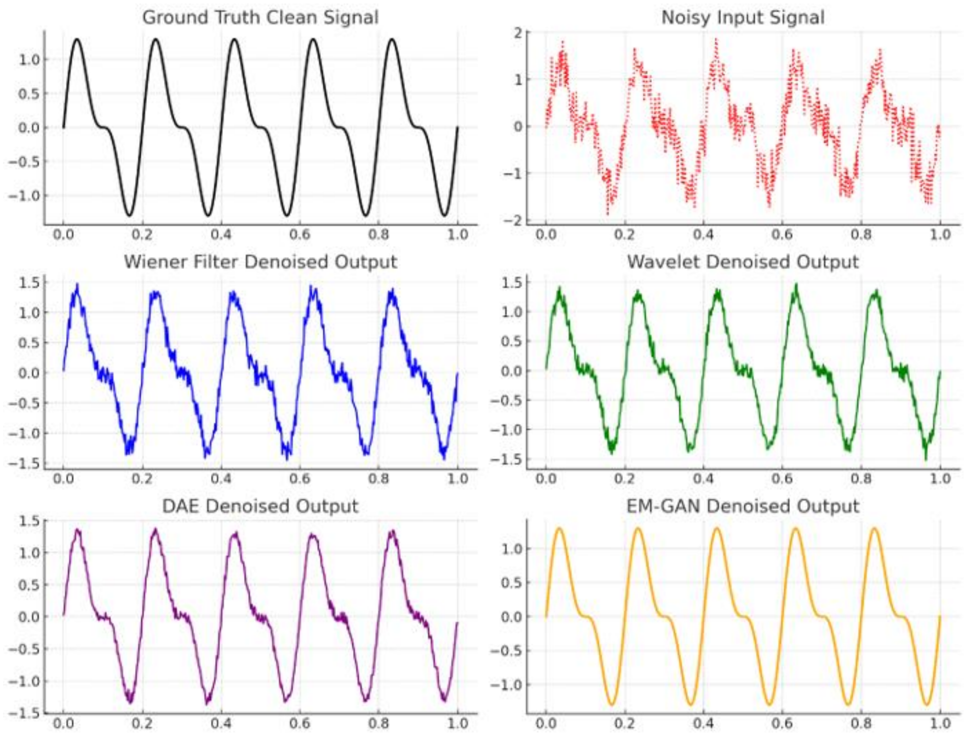

Figure 3. Visual comparison of signal denoising methods
(This figure presents a side-by-side comparison of noisy signals, clean ground-truth signals, and denoised outputs obtained from different methods)

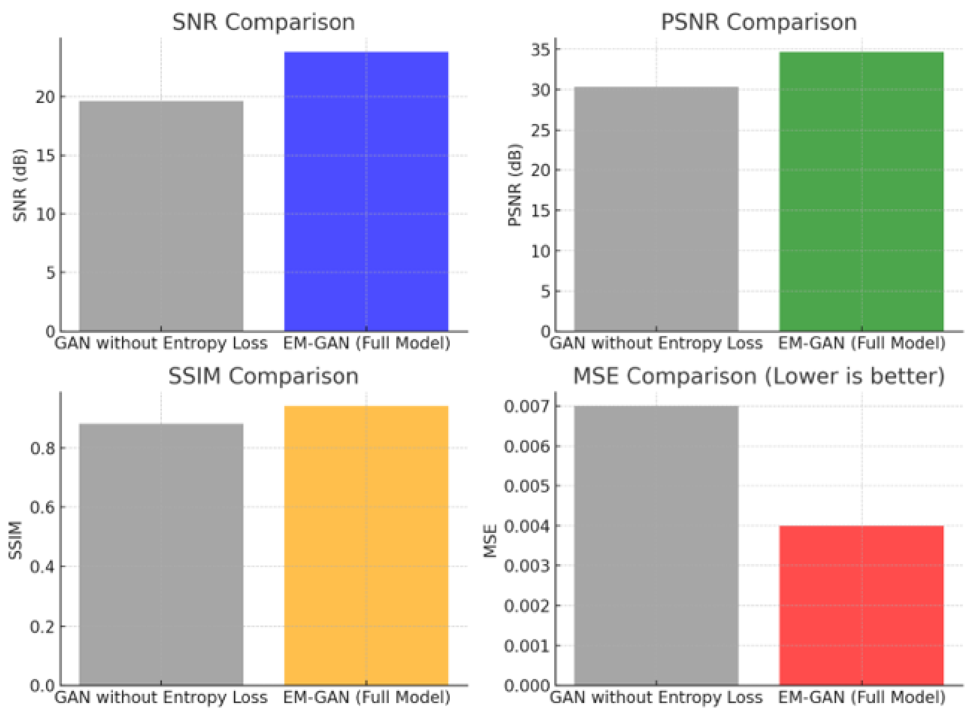

Figure 4. Impact of entropy regularization on denoising performance

4.4 Computational efficiency and real-time feasibility

The evaluation is conducted on a system with the following hardware specifications:

- GPU: NVIDIA RTX 4090 (24GB VRAM).
- CPU: Intel i9-12900K.
- RAM: 32GB DDR5.
- Framework: PyTorch 2.0 with CUDA acceleration.
- Each experiment is repeated 50 times per signal length, and the average inference time is reported to ensure consistency.

To provide a meaningful analysis, EM-GAN's inference time is compared against traditional filtering techniques and deep learning-based denoisers. The results are summarized in Figure 5, where inference time is plotted against signal length for different methods

Key Observations:

- Traditional methods (Wiener Filter, Wavelet Denoising) maintain low and constant inference times due to their deterministic nature.
- Deep learning-based models exhibit increased inference times for longer signals, with Vanilla GAN

being the slowest due to training instabilities.

- EM-GAN maintains a balance between efficiency and denoising performance, showing reasonable computational overhead compared to other deep learning methods.

This figure reinforces that EM-GAN is computationally feasible for real-time applications while outperforming other GAN-based models in terms of scalability and stability.

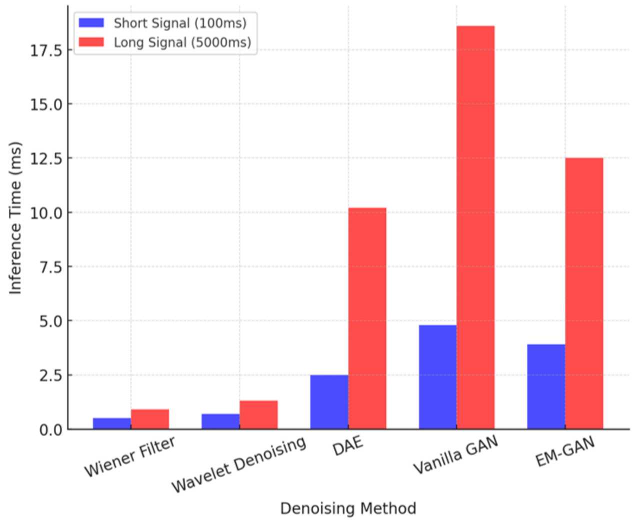

Figure 5. Computational efficiency of EM-GAN
(Computational Efficiency Comparison of Denoising Methods, comparing inference times for short (100ms) and long (5000ms) signals across different denoising techniques.)

4.5 Discussion

4.5.1 Comparison with baseline methods

The experimental findings indicate that EM-GAN outperforms traditional filtering and deep learning-based denoising techniques in terms of both quantitative metrics and qualitative assessments. The high SNR and PSNR scores reflect the model's ability to suppress noise while preserving critical signal features. The superior SSIM and lower MSE values further confirm the fidelity of the reconstructed signals.

4.5.2 Role of entropy maximization

The ablation study demonstrates that entropy regularization is a key factor in enhancing the performance of EM-GAN. Without this component, the model struggles with mode collapse, resulting in less diverse and lower-quality outputs. Entropy maximization ensures that the generator produces feature-rich, high-fidelity reconstructions, preserving both high- and low-frequency signal components.

4.5.3 Practical implications and real-world applications

The computational efficiency analysis suggests that EM-GAN is capable of real-time operation, making it applicable in various domains, including:

- **Biomedical Signal Processing:** ECG and EEG signal denoising for improved clinical diagnosis.
- **Speech Enhancement:** Noise suppression in voice communication systems and hearing aids.
- **Structural Health Monitoring:** Denoising of vibration signals for fault detection in engineering applications.

While our current real-time performance evaluation is conducted on high-performance hardware, the lightweight architecture of EM-GAN—featuring a shallow convolutional backbone and non-parametric entropy-based components—indicates strong potential for deployment on edge or embedded platforms. With further optimization techniques such as model pruning, quantization, or hardware-specific acceleration, EM-GAN can be adapted to resource-constrained environments. Future work will explore these deployment pathways to extend its practical applicability.

4.5.4 Limitations and future work

While EM-GAN demonstrates state-of-the-art performance in signal denoising, several limitations should be acknowledged. First, its generalization to unseen or complex noise types has not been fully evaluated. For instance, EM-GAN may experience performance degradation when exposed to non-Gaussian noise distributions, such as heavy-tailed or temporally correlated interference, which are underrepresented in the current training datasets. This reflects the model's sensitivity to the statistical characteristics of the input noise.

Moreover, the current framework has not been tested on adversarial or non-stationary noise scenarios, including perturbations designed to exploit model vulnerabilities or environments exhibiting dynamic distribution shifts. These factors may pose significant challenges to the model's robustness in practical deployment. To mitigate these issues, future work will investigate domain adaptation techniques, such as adversarial adaptation, self-supervised fine-tuning using target-domain data, and meta-learning strategies for rapid adjustment to unfamiliar noise patterns.

Second, the training process of EM-GAN is computationally demanding, which may limit its deployment in resource-constrained environments. Optimization strategies, such as model compression or hardware-aware acceleration, will be essential for enabling real-time performance on edge devices.

Finally, the current design does not explicitly incorporate domain-specific priors or hybrid architectural elements. Future research could explore combining EM-GAN with adaptive filtering methods or physics-informed components to enhance robustness. Expanding the diversity of noise types during training will also be considered to improve generalization across a broader range of real-world scenarios. Recent developments in diffusion-based models [20, 21] have also shown promise in signal generation tasks, although their inference efficiency remains a challenge for real-time denoising applications.

4.6 Summary of experimental results

These results establish EM-GAN as an effective, stable, and computationally efficient solution for signal denoising, with promising applicability across multiple domains. Compared to both traditional filtering techniques and recent GAN-based models, EM-GAN consistently demonstrates superior performance in noise suppression and feature preservation. Entropy regularization emerges as a critical component in maintaining structural fidelity and mitigating mode collapse. Moreover, the model exhibits competitive inference efficiency, reinforcing its suitability for real-time deployment in biomedical signal processing, speech enhancement, and structural health monitoring scenarios.

5. CONCLUSION

This work introduced EM-GAN, an entropy-maximized GAN designed for signal denoising and feature preservation. By integrating entropy regularization into adversarial learning, EM-GAN effectively mitigates mode collapse, enhances signal diversity, and preserves essential structural features while suppressing noise. Future work will focus on optimizing computational efficiency, extending the model to diverse noise environments, and exploring hybrid approaches that combine EM-GAN with adaptive filtering techniques to further enhance robustness and generalization. The findings demonstrate that EM-GAN is a high-performing, stable, and computationally efficient solution for real-time signal denoising, offering promising advancements in modern signal processing applications.

ACKNOWLEDGEMENTS

This paper was supported by the Foundation of Improving Academic Ability in University for Young Scholars of Guangxi (Grants No.: 2024KY1070; 2025KY1357); the Major Research Project of Liuzhou Polytechnic University (Grants No.: 2023KA06; 2023SA02; 2024KA03; 2024KA09).

AUTHOR CONTRIBUTIONS

Darong Liang, Jian Yu, and Hui Wang contributed equally to this work. Darong Liang led the experimental design and conducted model development. Jian Yu proposed the original concept of entropy-driven generative modeling, supervised the overall project, and coordinated team collaboration. Hui Wang contributed to the mathematical formulation and implementation of entropy-regularized training strategies. Xiaodan Chen assisted in industrial application analysis and data preprocessing. Chengxuan Huang and Ze Li supported benchmarking and visualization. Yuanyuan Luo and Xiangqing Wei contributed to the literature review and manuscript editing. All authors read and approved the final manuscript.

REFERENCES

[1] Zhou, Y., Wang, J., Tang, J., Gou, C., Jiang, Z., Li, D., Ng, S.K. (2023). MP-GAN: Cyber-attack detection and localization for cyber-physical systems with multi-process generative adversarial networks. In 2023 International Conference on Artificial Intelligence of Things and Systems (AIoTSys), Xi'an, China, pp. 186-193. https://doi.org/10.1109/AIoTSys58602.2023.00049

[2] Lee, K.S., Tran, N.T., Cheung, N.M. (2021). InfoMax-GAN: Improved adversarial image generation via information maximization and contrastive learning. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, pp. 3942-3952. https://doi.org/10.1109/WACV48630.2021.00399

[3] Lambert, A., Boots, B. (2021). Entropy regularized motion planning via stein variational inference. arXiv preprint arXiv:2107.05146. https://doi.org/10.48550/arXiv.2107.05146

[4] Yu, J., Wang, H., Huang, C., Li, Z. (2024). Entropy-Maximized Generative Adversarial Network (EM-GAN) based on the thermodynamic principle of entropy increase. Traitement du Signal, 41(6): 3255-3264. https://doi.org/10.18280/ts.410641

[5] Gu, S., Yang, L., Du, Y., Chen, G., Walter, F., Wang, J., Knoll, A. (2022). A review of safe reinforcement learning: Methods, theory and applications. arXiv preprint arXiv:2205.10330. https://doi.org/10.48550/arXiv.2205.10330

[6] Dai, Z., Zhao, L., Wang, K., Zhou, Y. (2024). Mode standardization: A practical countermeasure against mode collapse of GAN-based signal synthesis. Applied Soft Computing, 150: 111089. https://doi.org/10.1016/j.asoc.2023.111089

[7] Xu, K., Liao, L., Xiao, J., Chen, C., Wu, H., Yan, Q., Lin, W. (2024). Boosting image quality assessment through efficient transformer adaptation with local feature enhancement. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, pp. 2662-2672. https://doi.org/10.1109/CVPR52733.2024.00257

[8] Guo, Y., Cheng, T., Yang, Z., Huang, Y., Li, M., Wang, T. (2024). A systematic review and meta-analysis of balance training in patients with chronic ankle instability. Systematic Reviews, 13(1): 64. https://doi.org/10.1186/s13643-024-02455-x

[9] Liu, D. (2024). The effects of segmentation on cognitive load, vocabulary learning and retention, and reading comprehension in a multimedia learning environment. BMC Psychology, 12(1): 4. https://doi.org/10.1186/s40359-023-01489-5

[10] Welfert, M., Kurri, G.R., Otstot, K., Sankar, L. (2024). Addressing GAN training instabilities via tunable classification losses. IEEE Journal on Selected Areas in Information Theory, 5: 534-553. https://doi.org/10.1109/JSAIT.2024.3415670

[11] Hu, S., Gao, F., Zhou, X., Dong, J., Du, Q. (2024). Hybrid convolutional and attention network for hyperspectral image denoising. IEEE Geoscience and Remote Sensing Letters, 21: 5504005. https://doi.org/10.1109/LGRS.2024.3370299

[12] Li, X., Zheng, Y., Zhang, J., Dang, S., Nallanathan, A., Mumtaz, S. (2024). Finite SNR diversity-multiplexing trade-off in hybrid ABCom/RCom-assisted NOMA networks. IEEE Transactions on Mobile Computing, 23(10): 9108-9119. https://doi.org/10.1109/TMC.2024.3357753

[13] Mienye, I.D., Swart, T.G., Obaido, G. (2024). Recurrent neural networks: A comprehensive review of architectures, variants, and applications. Information, 15(9): 517. https://doi.org/10.3390/info15090517

[14] Sayegh, H.R., Dong, W., Al-Madani, A.M. (2024). Enhanced intrusion detection with LSTM-Based model, feature selection, and SMOTE for imbalanced data. Applied Sciences, 14(2): 479. https://doi.org/10.3390/app14020479

[15] Madani, S.S., Ziebert, C., Vahdatkhah, P., Sadrnezhaad, S.K. (2024). Recent progress of deep learning methods for health monitoring of lithium-ion batteries. Batteries, 10(6): 204. https://doi.org/10.3390/batteries10060204

[16] Lin, K., Pan, J., Xi, Y., Wang, Z., Jiang, J. (2024). Vibration anomaly detection of wind turbine based on temporal convolutional network and support vector data

description. Engineering Structures, 306: 117848. https://doi.org/10.1016/j.engstruct.2024.117848

[17] Parbat, D., Chakraborty, M. (2024). Multiscale entropy analysis of single lead ECG and ECG derived respiration for AI based prediction of sleep apnea events. Biomedical Signal Processing and Control, 87: 105444. https://doi.org/10.1016/j.bspc.2023.105444

[18] Wang, L., Liu, Y., Yu, Z., Gao, S., Mao, C., Huang, Y., Dong, L. (2025). SECodec: Structural entropy-based compressive speech representation codec for speech language models. In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, pp. 25380-25388. https://doi.org/10.1609/aaai.v39i24.34726

[19] Ananthi, A., Subathra, M.S.P., George, S.T., Sairamya, N.J. (2024). Entropy-based feature extraction for classification of EEG signal using Lifting Wavelet Transform. Przegląd Elektrotechniczny, 100: 146-150. https://doi.org/10.15199/48.2024.09.27

[20] Yi, H., Hou, L., Jin, Y., Saeed, N.A., Kandil, A., Duan, H. (2024). Time series diffusion method: A denoising diffusion probabilistic model for vibration signal generation. Mechanical Systems and Signal Processing, 216: 111481. https://doi.org/10.1016/j.ymssp.2024.111481

[21] Lemercier, J.M., Richter, J., Welker, S., Moliner, E., Välimäki, V., Gerkmann, T. (2025). Diffusion models for audio restoration: A review. IEEE Signal Processing Magazine, 41(6): 72-84. https://doi.org/10.1109/MSP.2024.3445871