

A physics-constrained spectral generative adversarial network for rolling bearing fault diagnosis

Yitong Liu¹ , Xintian Liu² , Guoqiang Wang^{1,*} , Xia Zhang^{3,*} 

¹ School of Mathematics, Physics and Statistics, Shanghai University of Engineering Science, Shanghai, 201620, P.R. China

² School of Mechanical and Automotive Engineering, Shanghai University of Engineering Science, Shanghai, 201620, P.R. China

³ Office of Scientific Research, Shanghai University of Medicine & Health Sciences, Shanghai, 201318, P.R. China

* Corresponding Author: guoq_wang@hotmail.com, zhangx_10@sumhs.edu.cn

Received: 3 April 2026

Revised: 7 June 2026

Accepted: 25 June 2026

Online: 8 July 2026

This is an open access article
under the [CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Abstract

Accurate bearing fault diagnosis is often hindered by sample scarcity and class imbalance. This paper proposes the Physics-Constrained Spectral Generative Adversarial Network (PCS-GAN), which embeds prior mechanical knowledge into adversarial training to enhance physical interpretability. The PCS-GAN model employs a time-frequency dual-discriminator structure to enforce statistical and spectral realism via physics-informed loss terms. To ensure stability, it integrates WGAN-GP with a progressive loss scheduling strategy. Evaluations on CWRU and MFPT datasets confirm that the PCS-GAN model achieves high spectral integrity while reducing peak GPU memory consumption by 28% compared to Transformer-based architectures. Furthermore, the model exhibits robust noise resilience and elevates macro-F1 scores to over 0.94 in extreme imbalance scenarios. These results demonstrate that the PCS-GAN model provides a computationally efficient and reliable solution for fault diagnosis in data-scarce industrial settings.

Keywords: rolling bearing fault diagnosis, physics-informed generative models, generative adversarial networks, data augmentation, time-frequency spectral consistency

Article citation:

Liu Y, Liu X, Wang G, Zhang X, A physics-constrained spectral generative adversarial network for rolling bearing fault diagnosis, Eksploracja i Niezawodność – Maintenance and Reliability 2027; 29(1) <http://doi.org/10.17531/ein/225075>

Highlights

- A physics-constrained loss integrating prior mechanical knowledge and harmonic alignment is proposed to ensure the physical plausibility of generated signals.
- A time-frequency dual-discriminator based on the Wasserstein Generative Adversarial Network with Gradient Penalty (WGAN-GP) improves training stability and enhances statistical and spectral consistency.
- A progressive loss scheduling strategy for physics regularization accelerates convergence and improves spectral fidelity.

1. Introduction

Rolling element bearings are fundamental components in various types of rotating machinery, and their condition critically affects system reliability and safety. Bearing failures, often caused by fatigue, wear, or lubrication issues, are a leading cause of unscheduled downtime and catastrophic accidents. Industry reports indicate that bearing failures account for a substantial portion of all mechanical failures, with estimates

ranging from 40% to 50%, resulting in significant economic losses and potential safety hazards [1]. Consequently, accurate and early fault diagnosis is of paramount importance for ensuring the operational continuity and safety of industrial processes [2]. The increasing demand for intelligent and automated maintenance has spurred extensive research into advanced fault diagnosis techniques, which have largely transitioned from traditional signal processing to data-driven approaches [3,4].

However, real-world fault data is often sparse due to the infrequent and random nature of bearing failures, while normal operational data is abundant [5]. Bearing failures are random and relatively rare events, particularly in early stages, making fault data collection challenging [6]. This data imbalance (many normal examples, few fault examples) forms a small-sample problem in bearing diagnosis that can severely degrade the performance of data-driven algorithms. Therefore, exploring fault diagnosis methods under small-sample conditions is of increasing research value, especially for improving the

robustness and efficiency of rolling element bearing fault diagnosis in practical settings [7]. This fundamental issue of data scarcity, which is the focus of this work, highlights the need for novel methods that can perform robustly under small-sample conditions. Addressing this data scarcity issue is a major challenge for many data-driven methods, especially advanced artificial intelligence techniques that rely on large datasets.

Artificial intelligence approaches can be further divided into machine learning and deep learning techniques [8,9]. Deep learning methods have shown great success in bearing fault diagnosis by automatically extracting features from vibration signals. However, these models typically require abundant labeled data, which contradicts the scarcity of real fault examples. Recently, diverse generative paradigms have been explored to alleviate this data deficiency. For instance, Diffusion Models [10] and Transformer-based architectures [11] have demonstrated high-fidelity synthesis capabilities for complex sequences. Despite their potential, these paradigms often entail significant computational overhead or risk generating non-physical artifacts due to a lack of explicit mechanical constraints. Generative adversarial networks (GANs), first proposed by Goodfellow et al. [12] in 2014, have emerged as a promising solution by synthesizing realistic fault samples from limited data. By learning the underlying distribution of actual bearing signals, GAN-based augmentation can expand the training dataset without additional manual labeling, and many advanced variants of the basic GAN architecture have been developed. For instance, CycleGAN-based frameworks [13] have been successfully applied to translate fault features across different operating conditions, while Auxiliary Classifier GANs (ACGANs) [14] have improved semi-supervised diagnosis through label-conditioned generation. Conventional GAN approaches, however, have primarily focused on architectural innovations and loss function stabilization [15] in this context. They may generate synthetic signals with spurious or unrealistic characteristics, and the adversarial training process can be unstable. Black-box data augmentation lacks explicit alignment with mechanical vibration principles, undermining diagnostic model trustworthiness [16]. The integration of physical mechanisms with data-driven approaches addresses the limitations of standalone physics-based models or purely data-driven methods

by combining their complementary strengths [17]. Kim et al. [18] proposed the Spectrum-Guided GAN (SGAN) and Density-Direction Sampling (DDS), which leverage both time-domain signals and their frequency spectra during training. Lyu et al. [19] introduced the Gradient Penalty Separated Classifier GAN (GPSC-GAN), a novel data augmentation model that develops independent classifiers to generate multi-mode fault samples under conditions where fault categories are unknown. Despite these advances, existing methods still face two significant limitations: (1) Although some research has attempted to impose spectral constraints during generation, they do not incorporate the physical characteristics of bearing vibrations. Moreover, the approaches for representing physical constraints are relatively simplistic and typically limited to metrics such as Mean Squared Error (MSE). (2) Instability during training due to insufficient data remains an unsolved problem. This issue has been highlighted in prior works, such as the review of GAN in mechanical fault diagnosis by Pan et al. in [15].

To address these limitations, we propose a novel Physics-Constrained Spectral GAN (PCS-GAN) model for rolling bearing fault diagnosis. Building upon the classical Deep Convolutional GAN (DCGAN) framework, this work makes the following key contributions:

- (1) This study proposes a novel physics-constrained loss that integrates spectral consistency and harmonic alignment, ensuring the generated signals are both physically plausible and diagnostically meaningful.
- (2) This study designs a time-frequency dual-discriminator architecture, in which the time-domain discriminator employs the Wasserstein Generative Adversarial Network with Gradient Penalty (WGAN-GP) to enhance training stability, alleviate mode collapse, and support robust cross-domain learning.
- (3) This study develops a progressive loss scheduling strategy for physical regularization, where the generator first learns statistical distributions and then progressively incorporates physical constraints, thereby accelerating convergence and improving spectral fidelity.

In this paper, we systematically evaluate the generated signals using both statistical similarity and physical consistency indicators to reconcile data-driven learning with vibration

theory. Furthermore, the diagnostic utility of the synthetic data is validated through fault dataset augmentation. For both Support Vector Machine (SVM) and Convolutional Neural Network (CNN) classifiers, the inclusion of PCS-GAN-generated samples leads to significant improvements in diagnostic performance.

The remainder of this paper is organized as follows: Section 2 reviews GAN variants and physics-guided machine learning methods; Section 3 details the architecture and training process of the PCS-GAN model; Section 4 presents experimental results and fault diagnosis validation; Section 5 concludes the paper by summarizing the core contributions and quantitative achievements. Finally, Section 6 provides a comprehensive discussion on the diagnostic mechanisms, current limitations, and potential trajectories for future research.

2. Background theory

Rolling bearings are essential components in rotating machinery, and their health is monitored through vibration signals. Faults on the inner race, outer race, or rolling elements generate periodic impulses, which produce distinct features in the frequency spectrum at specific characteristic frequencies [20]. These fault frequencies can be derived analytically based on bearing geometric parameters and are expressed as follows [1]:

$$\begin{cases} f_{IRD} &= f_r \times \frac{Z}{2} \left(1 + \frac{d}{D} \cos(\alpha)\right) \\ f_{ORD} &= f_r \times \frac{Z}{2} \left(1 - \frac{d}{D} \cos(\alpha)\right) \\ f_{RD} &= f_r \times \frac{D}{2d} \left\{1 - \left(\frac{d}{D}\right)^2 \cos^2(\alpha)\right\} \end{cases} \quad (1)$$

where f_{IRD} , f_{ORD} , and f_{RD} denote the fault characteristic frequencies corresponding to the *inner race defect*, *outer race defect*, and *rolling element defect*, respectively. Here, f_r is the shaft rotational frequency, Z is the number of rolling elements, d is the ball diameter, D is the pitch diameter, and α is the contact angle.

These faults induce amplitude modulation in the vibration signal, often masked by noise or system resonance. This modulation creates a high-energy resonance band around the system's resonance frequency, making it difficult to detect fault frequencies in the low-frequency spectrum [1]. Envelope analysis, typically implemented via the Hilbert transform, isolates the amplitude modulation and enhances fault visibility [21,22]. The envelope of the analytic signal $x(t) + j\hat{x}(t)$ is given by:

$$x(t)_{\text{env}} = \sqrt{x(t)^2 + \hat{x}(t)^2} \quad (2)$$

where $\hat{x}(t)$ is the Hilbert transform of $x(t)$. This facilitates extraction of periodic fault-related features from nonstationary signals. Figure 1 shows a typical fault signal and its envelope.




Figure 1. Typical signals and envelope signals from local faults in rolling bearings.




Figure 2. Time-domain waveform and envelope (left column); envelope spectra of four fault types with marked fault frequencies (right column).

To further identify fault signatures, the Fast Fourier Transform (FFT) is employed to convert time-domain data into the frequency domain. FFT is used due to its high computational efficiency, making it well-suited for processing long-duration signals and supporting real-time monitoring in industrial applications [23]. The Fourier transform is defined as:

$$X(f) = \int_{-\infty}^{\infty} x(t) e^{-j2\pi ft} dt \quad (3)$$

Spectral peaks at f_{IRD} , f_{ORD} , and f_{RD} confirm fault presence and type, offering critical insights for condition monitoring [9]. Figure 2 shows the time-domain waveform and its corresponding envelope, while the lower row illustrates the envelope spectrum for four fault types, with fault frequencies marked by red lines.

2.1. Fundamentals of generative adversarial networks

GANs, introduced by Goodfellow et al. [12], consist of a generator G and a discriminator D trained in a two-player minimax game, as illustrated in Figure 3. The generator maps random noise z to synthetic data $G(z)$, while the discriminator

seeks to distinguish real samples $x \sim p_{\text{data}}$ from generated ones.

The objective is:

$$\min_G \max_D \mathcal{L}(G, D) = \mathbb{E}_{x \sim p_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim p_z} [\log(1 - D(G(z)))] \quad (4)$$




Figure 3. Architecture of a standard GAN.

To mitigate the training instability and gradient vanishing problems, WGAN was introduced. It substitutes the Jensen-Shannon divergence with the Wasserstein-1 distance, which provides a more stable gradient signal for the model [24]. A further improved variant, WGAN-GP, enhances convergence by introducing a gradient penalty term to enforce the 1-

Lipschitz constraint. The loss function of WGAN-GP is given by:

$$\mathcal{L}_{\text{WGAN-GP}} = \mathbb{E}_{\hat{x} \sim p_g} [D(\hat{x})] - \mathbb{E}_{x \sim p_{\text{data}}} [D(x)] + \lambda \mathbb{E}_{\hat{x} \sim p_{\hat{x}}} [(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2] \quad (5)$$

where $\hat{x}$ is sampled uniformly along straight lines between real and generated samples, and λ is the gradient penalty coefficient [25].

Conditional GANs (CGANs) extend the basic framework by incorporating additional information (e.g., class labels or attributes) into both G and D , enabling controlled data synthesis [26].

$$\text{MMD}^2(\mathcal{X}, \mathcal{Y}) = \frac{1}{N^2} \sum_{i,j=1}^N k(x_i, x_j) + \frac{1}{M^2} \sum_{i,j=1}^M k(y_i, y_j) - \frac{2}{NM} \sum_{i=1}^N \sum_{j=1}^M k(x_i, y_j) \quad (6)$$

which measures the distributional difference between two sets of samples using kernel-based statistical distance [28]. The lower the MMD value, the better the generated signal distribution approximates the real signal distribution, implying higher fidelity.

Let u and v be a real signal and a generated signal, respectively, and let $\bar{u}$ and $\bar{v}$ be the means of sequences u and v , respectively. Then the definition of PCC is given by [29]:

$$\text{PCC}(u, v) = \frac{\sum_{i=1}^n (u_i - \bar{u})(v_i - \bar{v})}{\sqrt{\sum_{i=1}^n (u_i - \bar{u})^2 \sum_{i=1}^n (v_i - \bar{v})^2}} \quad (7)$$

which quantifies the linear correlation between the real signal u and the generated signal v . A PCC value close to 1 indicates a high degree of similarity in waveform shape.

Let $\Omega = [f_c - \Delta f, f_c + \Delta f]$ be the fault frequency band, and let $A_g(f)$ be the amplitude spectrum of the generated signal. Then the definition of FFER is given by:

$$\text{FFER} = \frac{\sum_{f \in \Omega} A_g(f)}{\sum_{f \in \text{full range}} A_g(f)} \quad (8)$$

which is used to evaluate the preservation of fault features within the generated signal's frequency domain and measures how much spectral energy is concentrated around the fault frequency band Ω . A higher FFER value suggests that the generated signal possesses greater diagnostic relevance, as it more accurately represents the characteristic fault frequencies.

2.3. Evaluation metrics for signal generation quality

Deep learning models have shown remarkable potential in fault signal generation; however, their black-box nature often results in limited interpretability and, more critically, a lack of adherence to domain-specific physical laws. In the context of

2.2. Evaluation metrics for signal generation quality

To assess the quality and realism of generated vibration signals, we employ three key metrics: Maximum Mean Discrepancy (MMD), Pearson Correlation Coefficient (PCC), and Fault Frequency Energy Ratio (FFER). For clarity, we define these metrics as MMD, PCC, and FFER, respectively.

Let $\mathcal{X}$ and $\mathcal{Y}$ be the real and generated sample sets, respectively, and let $k(\cdot, \cdot)$ be a kernel function. Then the definition of MMD is given by [27]:

rotating machinery diagnostics, such physically inconsistent signals may not only degrade data fidelity but also compromise the reliability of downstream fault classification [30].

To address this challenge, physics-guided learning has emerged as a powerful paradigm that incorporates domain knowledge into model training. Existing strategies typically fall into two categories: (i) embedding physical constraints directly into the learning objective, as in Physics-Informed Neural Networks (PINNs) [31], or (ii) guiding generative models through physics-based priors, as in Physics-Guided GANs [32].

Building on this idea, we introduce a physics-constrained spectral regularization into the GAN framework, leveraging prior knowledge of bearing fault mechanics. Specifically, we enforce frequency-domain consistency by encouraging generated signals to exhibit energy concentration at characteristic fault frequencies. This constraint acts as an implicit physical prior, steering the generator toward producing signals that are not only statistically realistic but also diagnostically meaningful.

3. Proposed PCS-GAN model

3.1. Model overview

To address the challenge of data scarcity, we propose a PCS-GAN model. The model incorporates dual discriminators operating in both the time and frequency domains, with additional physical constraints imposed in the frequency domain to enhance the realism and fidelity of the generated data. The complete workflow of the proposed framework is illustrated in Figure 4.

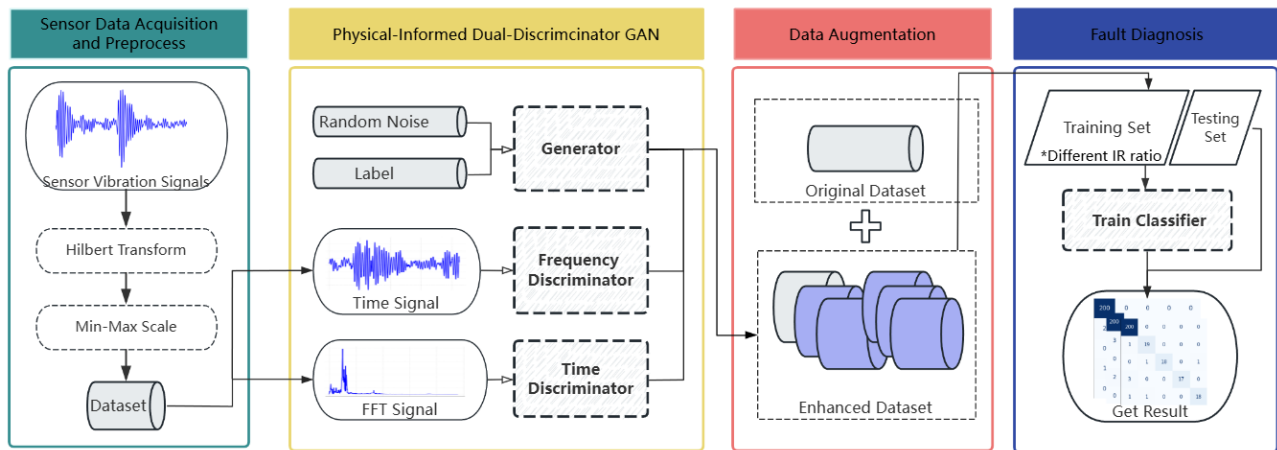


Figure 4. Overall workflow of the proposed generative framework.




Figure 5. Implementation details of the proposed framework.

The overall workflow of the proposed method is summarized as follows:

(i) **Data acquisition and preprocessing:** Raw signals are collected from sensors and subjected to appropriate preprocessing operations.

(ii) **Model construction and training:** A GAN model with dual discriminators is developed, where physical constraints are embedded in the frequency domain during training.

(iii) **Data augmentation:** The trained generator is employed to produce synthetic signals, thereby augmenting rare fault samples and constructing an enhanced dataset.

(iv) **Fault diagnosis:** The enhanced dataset is used to train a classifier, which performs fault diagnosis and yields the final diagnostic results.

Specifically, the implementation details of (ii) model construction and training, including the design of the loss functions, are further depicted in Figure 5.

3.2. Fault characteristic frequency determination

Accurate determination of fault characteristic frequencies is fundamental to embedding physically meaningful constraints. These frequencies, which correspond to specific fault types, are derived from bearing geometry and rotational dynamics. For a given bearing, essential parameters include the number of rolling elements n , ball diameter d , pitch diameter D , contact angle α , and shaft rotational frequency f_r , typically obtained by converting revolutions per minute (RPM) to Hz.

Under the assumption of zero contact angle ($\alpha = 0^\circ$), typical for deep groove ball bearings, the characteristic frequencies—Ball Pass Frequency Outer race (BPFO), Ball Pass Frequency Inner race (BPFI), and Ball Spin Frequency (BSF) are calculated using standard analytical expressions [1]. In this study, we utilize bearings from two widely recognized benchmarks: the Case Western Reserve University (CWRU) dataset [33] and the Machinery Failure Prevention Technology

(MFPT) dataset [34]. For CWRU, we analyze the drive-end (6205) and fan-end (6203) bearings. For MFPT, the NICE bearing parameters are adopted.

Their physical parameters are listed in Table 1.

Table 1. Geometrical parameters of the bearings used in this study.

Bearing Type	n	d (inch)	D (inch)
Drive-end (CWRU 6205)	9	0.3126	1.5370
Fan-end (CWRU 6203)	9	0.2656	1.1220
NICE (MFPT)	8	0.2350	1.2450

The resulting fault frequency orders, expressed relative to the shaft frequency, are shown in Table 2. These serve as domain-informed priors for constructing spectral regularization terms in the proposed model.

Table 2. Characteristic frequency orders (normalized to shaft frequency).

Bearing	BPFO	BPFI	BSF
Drive-end (CWRU 6205)	3.58	5.42	2.36
Fan-end (CWRU 6203)	3.05	4.95	1.99
NICE (MFPT)	3.25	4.76	2.55

It should be noted that in real-world applications, factors such as manufacturing tolerances or wear-induced deviations may cause actual frequencies to shift slightly from these theoretical values. To mitigate the sensitivity to such parameter errors, this study does not enforce point-wise frequency matching; instead, we design the physics-constrained losses to operate over adaptive spectral bands Ω . This integral-based approach, detailed in Section 3.3, provides a strategic tolerance buffer that ensures the model's robustness against minor geometric fluctuations and kinematic stochasticity.

3.3. Physics-guided loss formulation

To enforce physical consistency between generated and real signals, we design a physics-guided loss that integrates prior knowledge of characteristic fault frequencies into the generator's optimization. This loss promotes spectral conformity in regions of interest, ensuring both statistical plausibility and physical realism.

3.3.1. Characteristic frequencies and spectral band definition

The fault characteristic frequencies are calculated based on the bearing's physical parameters, as described in Section 3.2. These frequencies, such as BPFO, BPFI, and BSF, are

determined as follows:

$$\begin{cases} f_{\text{BPFO}} &= \frac{nf_r}{2} \left(1 - \frac{d}{D}\right) \\ f_{\text{BPFI}} &= \frac{nf_r}{2} \left(1 + \frac{d}{D}\right) \\ f_{\text{BSF}} &= \frac{Df_r}{2d} \end{cases} \quad (9)$$

where n is the number of rolling elements, d and D denote the rolling element and pitch diameters, respectively, and f_r is the shaft rotational frequency in Hz (converted from RPM).

To evaluate frequency-domain consistency, we define a fault-related spectral band Ω centered around each fault frequency. The bandwidth Δf is determined through a data-driven calibration mechanism rather than fixed heuristics, accounting for the inherent frequency shifts caused by rolling element slippage. Statistically, this adaptive bandwidth spans 10%–20% of the fault frequency f_{fault} , which provides a robust tolerance that captures the modulated energy and alleviates spectral leakage while excluding irrelevant broadband noise (the detailed calibration procedure is provided in Section 3.3.2):

$$\Omega = [f_{\text{fault}} - \Delta f, f_{\text{fault}} + \Delta f] \quad (10)$$

These fault-centric bands are used in the spectral loss formulation to penalize deviations from expected energy distribution in the frequency domain, thus embedding physical priors into adversarial training.

3.3.2. Loss function design

To enforce spectral consistency between real and generated signals, we define a composite physics-guided loss comprising three terms within the fault frequency band Ω : log-domain spectral distance, peak response alignment, and energy ratio consistency.

These components are strategically designed to constrain the generator's behavior regarding spectral topology, extreme response, and energy distribution, respectively.

Log-domain spectral distance ($\mathcal{L}_{\text{dist}}$) evaluates the point-wise topological similarity of the power distribution within the fault-related frequency band Ω . By operating in the logarithmic scale, it enhances the sensitivity to low-amplitude harmonic components that are often submerged in noise:

$$\mathcal{L}_{\text{dist}} = \mathbb{E} \left[\left\| \log \mathbf{A}_r(f) - \log \mathbf{A}_g(f) \right\|_{2,\Omega} \right] \quad (11)$$

where $\|\cdot\|_{2,\Omega}$ denotes the L_2 -norm restricted to the frequency set Ω , and $\mathbf{A}_r(f), \mathbf{A}_g(f)$ represent the clamped amplitude spectra of real and generated signals, respectively.

Peak response alignment ($\mathcal{L}_{\max}$) ensures that the generated signals maintain the physical characteristic of high-intensity shocks by minimizing the discrepancy between the maximum resonance amplitudes:

$$\mathcal{L}_{\max} = \mathbb{E} \left[\left\| \sup_{f \in \Omega} \mathbf{A}_r(f) - \sup_{f \in \Omega} \mathbf{A}_g(f) \right\| \right] \quad (12)$$

where $\sup$ denotes the supremum (peak amplitude) within the fault-related spectral band.

Energy ratio consistency ($\mathcal{L}_{\text{ratio}}$) constrains the global energy distribution by enforcing the proportion of fault-specific energy to the total spectral energy to match that of real samples:

$$\mathcal{L}_{\text{ratio}} = \mathbb{E} \left[\left\| \frac{\|\mathbf{A}_g(f)\|_{2,\Omega}^2}{\|\mathbf{A}_g(f)\|_2^2} - \frac{\|\mathbf{A}_r(f)\|_{2,\Omega}^2}{\|\mathbf{A}_r(f)\|_2^2} \right\| \right] \quad (13)$$

Composite loss integrates the above components into a unified objective:

$$\mathcal{L}_{\text{physics}} = \lambda_1 \mathcal{L}_{\max} + \lambda_2 \mathcal{L}_{\text{dist}} + \lambda_3 \mathcal{L}_{\text{ratio}} \quad (14)$$

where λ_1, λ_2 , and λ_3 are weighting coefficients designed to balance the multi-dimensional physical constraints. Within $\mathcal{L}_{\text{physics}}$, these coefficients harmonize the complementary roles of each term: $\mathcal{L}_{\max}$ captures extreme transient impulses, $\mathcal{L}_{\text{dist}}$ maintains spectral topology, and $\mathcal{L}_{\text{ratio}}$ supervises energy distribution. To prevent any single term from dominating the optimization due to order-of-magnitude disparities, we determine these weights through a gradient-based calibration mechanism during the initial training phase. Specifically, the gradient norms of $\mathcal{L}_{\max}$, $\mathcal{L}_{\text{dist}}$, and $\mathcal{L}_{\text{ratio}}$ are calculated and scaled to a consistent range ($10^1 \sim 10^2$), effectively aligning the physical constraints with the adversarial objective. This one-time calibration ensures that the generator is guided by rigorous physical boundaries while maintaining robust convergence stability across different datasets.

The bandwidth Δf within Ω is determined through an adaptive calibration process before formal training to account for random variations in bearing kinematics. In industrial applications, actual fault frequencies often deviate from theoretical values due to rolling element slippage. To address these variations, this paper quantifies the frequency deviation δ by analyzing a subset of n real samples, where δ represents the offset between the theoretical peak and the observed spectral maximum. Assuming these deviations follow a normal distribution, the bandwidth is defined as:

$$\Delta f = \text{avg}(\delta) + 2\sigma \quad (15)$$

where σ is the standard deviation. This data-driven approach provides a robust tolerance—typically 10% to 20% of the fault frequency—to capture genuine fault components while filtering out broadband noise. Consequently, this calibration enhances the precision of physics-constrained guidance across various rotational speeds and bearing types.

3.4. Model architecture and training strategy

As illustrated in Figure 6, the PCS-GAN model is developed on the basis of the DCGAN architecture and consists of several essential components, including normalization layers, transposed convolution layers (ConvTranspose1d), convolution layers (Conv1d), dropout, flatten, and linear layers. Each component plays a distinct role in ensuring stable training and effective signal generation. Normalization layers stabilize training and improve convergence. ConvTranspose1d layers in the generator progressively upsample latent vectors into high-resolution signals. Conv1d layers in the discriminators extract temporal features for distinguishing real from generated data. Dropout layers prevent overfitting and avoid excessively strong discriminators. Flatten layers convert feature maps into vectors for dense connections. Linear layers serve as the final classification units for real-fake discrimination.

By integrating adversarial loss in the time domain and spectral-constraint loss in the frequency domain, the PCS-GAN model captures fault-related physical features even with limited data, ensuring balanced training and improved data quality.

The training of the PCS-GAN model is designed to ensure both statistical fidelity to the data distribution and consistency with underlying physical constraints. To this end, we adopt a physics-constrained GAN training strategy based on WGAN-GP, extended with dual discriminators (D_t and D_f) operating in the time and frequency domains. To stabilize the adversarial equilibrium, we introduce a balancing factor γ to moderate the gradient contribution of D_f , preventing it from over-constraining the phase information during time-domain reconstruction. Furthermore, physics-guided losses are introduced following a curriculum-inspired schedule: the weight of the physical constraint is set to zero in early epochs and gradually increased after a warm-up stage.




Figure 6. Overall architecture of the proposed PCS-GAN model.

During this process, we observed that excessive physical guidance might inadvertently lead to mode collapse, where the generator produces highly similar samples that lack stochastic diversity. To mitigate this risk, a diversity-promoting loss $\mathcal{L}_{div}$ is incorporated into the generator's objective. This term incentivizes the generator to explore a broader distribution by penalizing samples with low intra-batch variance:

$$\mathcal{L}_{div} = \frac{1}{\frac{1}{B^2} \sum_{i,j} \|x_i - x_j\|_2 + \epsilon} \quad (16)$$

where B denotes the batch size and ϵ is a small constant for numerical stability. By balancing $\mathcal{L}_{adv}$, $\mathcal{L}_{physics}$, and $\mathcal{L}_{div}$, the model achieves a harmonious trade-off between physical accuracy and sample variety.

The training algorithm is summarized in Table 3.

Table 3. Training algorithm for the PCS-GAN model.

Input: Training dataset X ; batch size B ; gp coeff λ_{gp} ; factors γ, ω_f ; div weight λ_{div} ; initial weights λ_i ; warm-up N_{warm} .
Output: Trained generator G .
1. Initialize G and dual discriminators D_t, D_f . Initialize $\lambda_{phys} = \lambda_{div} = 0$.
2. for epoch = 1 to N_{epochs} do
3. for each mini-batch from X do
4. Discriminator Training (run N_{critic} times):
5. Sample noise z and generate fake signal $x_{fake} = G(z)$.
6. Compute Wasserstein distance with GP for D_t and D_f :
7. $L_{D_t} = E[D_t(x_{fake})] - E[D_t(x_{real})] + \lambda_{gp} \cdot GP_t$
8. $L_{D_f} = E[D_f(x_{fake})] - E[D_f(x_{real})] + \lambda_{gp} \cdot GP_f$
9. Update D_t, D_f by minimizing $L_D = L_{D_t} + \gamma \cdot L_{D_f}$.
10. Generator Training (run once):
11. Generate x_{fake} and compute x_{f_fake} .
12. Compute adversarial loss: $L_{adv} = -E[D_t(x_{fake})] - \gamma \omega_f E[D_f(x_{fake})]$.
13. Compute $L_{physics}$ (via L_{max} , L_{dist} , L_{ratio}) and diversity loss L_{div} .
14. if epoch $\leq N_{warm}$ then $L_G = L_{adv}$;
15. else $L_G = L_{adv} + \alpha L_{physics} + \beta L_{div}$.
16. Update G by minimizing L_G .
17. end for
18. end for

4. Experimental results

4.1. Data description

To rigorously validate the efficacy and generalizability of the

PCS-GAN model, experiments are conducted on two distinct datasets: the CWRU dataset and the MFPT dataset. The integration of these datasets allows for an evaluation across varying sampling rates, load conditions, and bearing geometries.

4.1.1. CWRU dataset configuration

We select five operating conditions from the CWRU drive-end bearing (6205) dataset, sampled at 48 kHz. These conditions cover three categories: Normal, Inner Race fault, and Outer Race fault, with damage diameters of 0.014 and 0.021 inches. As analyzed in Section 3.2, the theoretical fault characteristic frequencies (BPFI and BPFO) serve as the spectral priors. Each signal is segmented into non-overlapping windows of 1024 points. The detailed configuration is summarized in Table 4.

4.1.2. Dataset configuration

The MFPT dataset represents a more complex industrial scenario with a higher sampling rate of 48,828 sps. This study focuses on the NICE bearing under varying load conditions. For each fault category (Inner Race and Outer Race), we aggregate signals across seven different loads (ranging from 0 to 300 lbs) at a constant input shaft rate of 25 Hz. This aggregation ensures that the model learns features robust to load fluctuations. After segmenting the raw signals into 1024-point windows, 600 samples are randomly selected for each fault class to construct the experimental dataset. The fault frequencies are derived from the 25 Hz shaft frequency and the bearing geometry, as detailed in Table 5.

Table 4. Specifications of the processed CWRU dataset for experimental validation.

Fault Location	Fault Width	Sample Length	Speed	Label	Sample Numbers	Fault Frequency (Hz)
Normal	0.000	1024	1772	Normal	472	—
Rolling Element	0.014	1024	1750	B014	475	137.48
Inner Race	0.014	1024	1750	IR014	476	157.94
Outer Race	0.014	1024	1750	OR014	475	104.56
Outer Race	0.021	1024	1730	OR021	474	103.36

Table 5. Specifications of the processed MFPT dataset.

Fault Location	Load (lbs)	Shaft Rate (Hz)	Label	Sample Numbers	Fault Frequency (Hz)
Inner Race	0–300	25	IR	600	119.00
Outer Race	25–300	25	OR	600	81.25

4.2. Baseline models for comparative study

To evaluate the effectiveness and robustness of the PCS-GAN model, we conduct comprehensive comparative experiments against five representative generative models, spanning from classical adversarial networks to modern likelihood-based diffusion architectures. All baselines are trained under consistent experimental protocols to ensure a rigorous and fair performance assessment.

(1) GAN: A standard multilayer perceptron with Tanh activation at the output layer. The model is optimized using Adam with a learning rate of 1×10^{-4} and $\beta_1 = 0.5$.

(2) DCGAN [35]: A convolutional GAN architecture designed to capture local dependencies in the signal. Binary cross-entropy loss is adopted, and optimization is performed with Adam ($\beta_1 = 0.5, \beta_2 = 0.999$).

(3) WGAN-GP [25]: An improved WGAN incorporating gradient penalty to enforce Lipschitz continuity and enhance training stability. Adam optimizers are applied with $\beta_1 = 0.5, \beta_2 = 0.9$; learning rates are set to 1×10^{-4} for the generator and 2×10^{-4} for the discriminator. The number of discriminator updates per generator step is fixed at 5.

(4) DDPM [36]: A denoising diffusion probabilistic model (DDPM) implemented with a 1D-U-Net architecture. As a likelihood-based model, it generates samples by reversing a gradual Gaussian diffusion process over $T = 1000$ steps. The core network utilizes residual blocks and sinusoidal time embeddings to predict the added noise via a MSE objective.

(5) Transformer-Gan [37]: A generative model based on the self-attention mechanism. By employing multi-head attention layers, the model captures long-range correlations in the 1D time series, which are often difficult for CNN-based

architectures to process.

4.3. Experimental setup

To simulate realistic industrial scenarios where fault data are scarce, we construct classification datasets with increasing levels of fault data augmentation.




Figure 7. Dataset partitioning for signal generation and classification tasks.

Data Partitioning: As illustrated in Figure 7, the full dataset is divided into two parts: one for training the generative models and one for evaluating fault diagnosis. For classification, 80% of the data is used for training and 20% for testing. All signals are normalized and segmented into 1024-point non-overlapping samples.

Imbalance Simulation: Each classification dataset is constructed with 1000 normal samples and only 10 real fault samples per class. Synthetic fault data are incrementally added to adjust the imbalance ratio (IR), defined as the ratio of normal to fault samples. IR levels range from 100:1 (no augmentation) to 1:1 (fully balanced). Details are summarized in Table 6.

Table 6. Synthetic fault augmentation under varying IR.

IR Scenario	Normal Samples	Fault Samples (Real / Generated)
100:1	1000	10 / 0
50:1	1000	10 / 10
20:1	1000	10 / 40
10:1	1000	10 / 90
5:1	1000	10 / 190
1:1	1000	10 / 490

Classifiers: To evaluate diagnostic performance under extremely imbalanced regimes, we train four classifiers : SVM, Random Forest (RF), Multilayer Perceptron (MLP), and 1D-CNN on each augmented dataset. Both accuracy and macro-average F1 score are reported across different IRs, enabling a comprehensive comparison of classifier behavior and the contribution of synthetic data.

4.3.1. Quality evaluation of generated signals

In addition to classification performance, we conduct a dedicated evaluation of the fidelity and physical realism of generated fault signals. Both qualitative and quantitative analyses are carried out to assess whether the synthetic data exhibit statistical and spectral properties consistent with real-world signals.

Qualitative Comparison

Figure 8 illustrates a comprehensive qualitative assessment of the signals synthesized by the PCS-GAN model across six

distinct fault categories from the MFPT and CWRU datasets, both operating at high sampling rates (approximately 48 kHz). In the time-domain visualizations, the generated signals successfully replicate the periodic impulsive transients and decay characteristics observed in the real vibration data, demonstrating the model’s ability to capture complex temporal dependencies without falling into mode collapse. Furthermore, the corresponding envelope spectra reveal that the synthetic components exhibit sharp, well-defined peaks that align precisely with the theoretical fault frequency bands. The alignment between generated spectral peaks and theoretical fault signatures, even under significant background noise, demonstrates the efficacy of the integrated physical constraints. By constraining the generator within a physically consistent manifold, the PCS-GAN model ensures that the synthesized data adhere to the mechanical dynamics of the bearing system rather than achieving simple numerical similarity to the training distribution.




Figure 8. Qualitative comparison of real and PCS-GAN generated signals.

Quantitative Metrics

To systematically quantify the fidelity and diagnostic relevance of the synthesized signals, we employ three complementary metrics: MMD, PCC, and FFER. The results across six fault categories from the CWRU and MFPT datasets are summarized in Figures 9–11.

Distributional and Morphological Fidelity: Analysis of the MMD results, presented in Figure 9, indicates that the PCS-

GAN model maintains consistently high performance across all evaluated fault modes. Although likelihood-based diffusion models achieve marginally lower MMD in specific instances, the PCS-GAN model exhibits superior stability and robustly outperforms both vanilla GAN and DCGAN. This consistency underscores that the integration of physical constraints effectively prevents the generator from converging toward non-physical local optima. Regarding the Envelope Spectrum PCC

(Figure 10), which quantifies the alignment of periodic impulsive patterns within diagnostic frequency bands, the PCS-GAN model demonstrates high structural similarity to real signals. While DCGAN occasionally yields high PCC values, a joint analysis considering its elevated MMD suggests a lack of distributional diversity—a characteristic of mode collapse where the generator overfits specific waveform templates. In contrast, the PCS-GAN model achieves a reliable balance between structural fidelity and distributional variety.

Spectral Integrity and Physical Anchoring: The FFER results (Figure 11) further highlight the advantages of our approach. While Transformer-Naive achieves impressive spectral precision due to its inherent long-range dependency capture, the PCS-GAN model closely approaches this performance level, substantially surpassing all other GAN-based variants. This proximity to Transformer-level accuracy suggests that our physical loss terms effectively compensate for the limited receptive field of CNN-based architectures. While we acknowledge the remarkable feature extraction capabilities

of Transformers, the high performance of the PCS-GAN model provides a compelling alternative for resource-constrained environments.

Computational Efficiency and Resource Consumption:

To evaluate the practical applicability of the proposed method, Table 7 summarizes the training costs across different architectures. While the Transformer-Naive model provides impressive generation quality by capturing long-range spectral dependencies, it demands a substantial GPU memory footprint and exhibits the slowest convergence. In contrast, the PCS-GAN model achieves a near-Transformer level of synthetic signal fidelity while reducing peak memory consumption by approximately 28% and maintaining training durations comparable to baseline GANs. This result underscores that incorporating physical priors allows the model to enforce structural realism and spectral alignment without relying on excessively deep or resource-heavy parameters, offering an efficient yet high-quality solution for industrial data augmentation.




Figure 9. Comparison of best MMD scores across different generative models and fault types.




Figure 10. Comparison of best FFT-envelope PCC scores across different generative models and fault types.




Figure 11. Comparison of best FFER scores across different generative models and fault types.

Table 7. Comparison of computational efficiency and resource requirements.

Model	Avg. Conv. Epoch	Avg. Conv. Time (min)	Max GPU Mem (MB)
GAN	233.30	38.84	56.30
WGAN_GP	98.30	14.32	57.90
DCGAN	200.00	24.53	134.20
Diffusion	68.30	4.76	169.40
PCS-GAN (Ours)	170.00	22.02	1084.00
Transformer_Naive	243.30	32.55	1507.70

Table 8. Quantitative results of ablation studies.

Category	Configuration	MMD (↓)	PCC_Env (↑)	FFER (↑)	Time/Epoch (s)
CWRU_B014	Exp0 (Full Model)	0.0135	0.3692	0.0389	11.12
	Exp1 (NoPhysics)	0.0114	0.3426	0.0190	11.51
	Exp2 (NoDiversity)	0.0111	0.3370	0.0319	11.33
	Exp3 (NoDualD)	0.0094	0.3974	0.0339	7.96
	Exp4 (Baseline)	0.0171	0.3355	0.0203	10.66
CWRU_IR014	Exp0 (Full Model)	0.0092	0.4247	0.0285	10.27
	Exp1 (NoPhysics)	0.0070	0.4403	0.0089	10.75
	Exp2 (NoDiversity)	0.0111	0.4222	0.0275	10.74
	Exp3 (NoDualD)	0.0151	0.4000	0.0251	10.48
	Exp4 (Baseline)	0.0098	0.3929	0.0093	10.38
CWRU_OR014	Exp0 (Full Model)	0.0188	0.4346	0.0268	9.16
	Exp1 (NoPhysics)	0.0124	0.4254	0.0103	9.14
	Exp2 (NoDiversity)	0.0126	0.4496	0.0220	9.49
	Exp3 (NoDualD)	0.0194	0.4443	0.0234	8.95
	Exp4 (Baseline)	0.0161	0.3953	0.0109	9.18
CWRU_OR021	Exp0 (Full Model)	0.0180	0.5641	0.0576	9.45
	Exp1 (NoPhysics)	0.0060	0.4080	0.0113	9.45
	Exp2 (NoDiversity)	0.0183	0.4590	0.0423	9.44
	Exp3 (NoDualD)	0.0152	0.5129	0.0507	7.38
	Exp4 (Baseline)	0.0108	0.3865	0.0107	7.42
MFPT_IR	Exp0 (Full Model)	0.0084	0.5997	0.0289	6.36
	Exp1 (NoPhysics)	0.0093	0.4622	0.0051	6.84
	Exp2 (NoDiversity)	0.0151	0.4744	0.0361	6.85
	Exp3 (NoDualD)	0.0083	0.4805	0.0261	8.34
	Exp4 (Baseline)	0.0062	0.3660	0.0050	5.60
MFPT_OR	Exp0 (Full Model)	0.0145	0.3699	0.0295	17.74
	Exp1 (NoPhysics)	0.0036	0.2841	0.0065	17.37
	Exp2 (NoDiversity)	0.0154	0.3464	0.0261	20.97
	Exp3 (NoDualD)	0.0104	0.3804	0.0294	16.27
	Exp4 (Baseline)	0.0129	0.2968	0.0063	20.07

4.3.2. Ablation study and loss balancing analysis

To clarify the contribution of each module and evaluate the generalization of the proposed method, four ablation variants are conducted across multiple fault categories: (1) Exp1 (NoPhysics), (2) Exp2 (NoDiversity), (3) Exp3 (NoDualD), and (4) Exp4 (Baseline). The quantitative results are summarized in Table 8.

An observation of Table 8 reveals a consistent trade-off between statistical proximity and physical integrity across all categories. For instance, in MFPT_OR, Exp1 (NoPhysics) achieves a low MMD, indicating high statistical similarity to the training samples; however, this proximity does not translate into diagnostic reliability, as evidenced by its significantly lower FFER compared to the full model Exp0. This discrepancy confirms that without physical regularization, the generator tends to capture the noise floor rather than the underlying mechanical dynamics.

Beyond these spectral metrics, Figure 12 qualitatively demonstrates the necessity of the diversity-promoting loss using the MFPT-OR category as a representative case. As illustrated in the t-SNE manifolds, the full model Exp0 maintains high structural alignment with the real reference, whereas the absence of diversity constraints in Exp2 leads to severe manifold collapse. Furthermore, while Exp3 achieves higher PCC in certain cases (e.g., CWRU_B) due to its simplified adversarial structure, it fails to match the FFER robustness of Exp0. Consequently, these results reinforce that the integration of dual-domain discriminators and physical constraints is essential for synthesizing signals that are both statistically consistent and diagnostically reliable across diverse industrial conditions.




Figure 12. t-SNE visualization of generated manifolds for ablation configurations.

Mechanistic Synergy and Dynamic Balancing: The synergy between adversarial and physical losses is managed

through a strategic warm-up phase. As illustrated in Figure 13, taking the MFPT-OR category as a representative case, physical constraints are activated at Epoch 20, allowing the generator to first establish a stable global distribution before refining local spectral features. Although the physical constraints comprise multiple terms (L_{max} , L_{dist} , L_{ratio}), the evolution of L_{dist} is presented in Figure 13a as a representative indicator of spectral alignment. Upon its introduction, a sharp decline in spectral error is observed for variants containing physical constraints (Exp0, Exp2, Exp3), effectively steering the generator away from the non-physical manifolds occupied by the baseline models.

Notably, the discriminator loss (d_{loss}) in Figure 13b demonstrates that incorporating physical information helps the model converge to a lower, more stable equilibrium compared to Exp4. This suggests that physical priors provide clearer gradient guidance, which not only enhances generation realism but also optimizes training efficiency. Consequently, the guided search space leads to shorter per-epoch durations in Exp0 and Exp3 compared to the unconstrained and relatively inefficient adversarial process in the baseline DCGAN.




Figure 13. Evolution of training metrics: (a) Physical Distance Loss (l_{dist}), and (b) Discriminator Loss (d_{loss}).

4.3.3. Robustness and data scarcity analysis

Hyper-parameter Sensitivity Analysis

The integration of physical priors necessitates a balancing strategy to align disparate loss terms within the optimization landscape. This study evaluates the model's sensitivity to the physics-constrained weights λ_{max} , λ_{dist} , and λ_{ratio} using three representative configurations on the CWRU dataset:

- (1) Set A (Baseline: $\lambda_{max} = 10, \lambda_{dist} = 0.25, \lambda_{ratio} = 100$)
- (2) Set B (Conservative: 1.5λ)
- (3) Set C (Aggressive: 0.5λ).




Figure 14. Robustness comparison across different physics weight scales.

As illustrated in Figure 14, the performance metrics remain highly consistent across the tested range. Notably, Set B achieves the highest generation fidelity, while the baseline configuration maintains a robust score. These results indicate that the physical constraints function as a flexible guidance mechanism rather than a rigid boundary. This inherent tolerance allows the model to accommodate minor geometric parameter inaccuracies by leveraging the underlying empirical data distribution to refine the generative manifold, thereby ensuring stable performance even under parametric uncertainty.

Minimum data requirements and convergence stability:

To assess the model's resilience against data scarcity, this study quantifies the correspondence between the sample size N and the physical duration of the vibration signals. Utilizing the MFPT dataset as a representative case (48,828 Hz, 3 s total), the data provides a maximum throughput of approximately 143 samples using 1024-point non-overlapping windows. In this context, $N = 50$ and $N = 100$ translate to a reliance on merely 1.05 s and 2.10 s of raw data, respectively.

The relationship between the evolution of total generator

loss ($L_{total,G}$) and sample size is illustrated in Figure 15. While all configurations eventually exhibit a convergent trend, the optimization stability remains sensitive to data volume. Specifically, although the $N = 50$ configuration can proceed with training through adaptive batch-size calibration, it suffers from pronounced stochastic perturbations and a higher loss floor compared to larger datasets. Notably, a distinct performance threshold is observed at $N = 100$ (approximately 2.10 s of signal), beyond which the introduction of physical priors effectively compensates for data scarcity, allowing the generator to establish a convergence level comparable to that of full-scale datasets.




Figure 15. Convergence behavior of total generator loss under limited sample sizes.

This suggests that $N = 100$ provides a sufficiently stable training baseline for the current experimental setup. However, it should be noted that while this threshold ensures technical convergence, the diversity of synthesized patterns remains inherently bounded by the initial training pool. For complex industrial environments with high non-stationarity, expanding the variety of training samples remains a vital direction for future research to further enhance diagnostic reliability.

4.3.4. Comparative evaluation under extremely imbalanced regimes

Effectiveness under Extreme Class Imbalance

To further evaluate the effectiveness of the PCS-GAN under severe data scarcity, we conduct fault classification experiments across varying IR. Table 9 and Table 10 report the classification accuracy and macro-F1 scores achieved by four commonly used classifiers, while Figure 16 directly shows the performance trends. We observe a consistent trend: while the accuracy

remains relatively high even under extreme imbalance (e.g., $IR = 1:100$), the macro-F1 scores are initially very low and increase steadily as the IR improves.

This gap between accuracy and F1 highlights a critical issue in extreme few-shot scenarios—the classifiers tend to overfit the majority class (i.e., normal samples) and fail to detect faults. For example, with $IR = 1:100$, the accuracy is over 95% for all models, but the macro-F1 scores are as low as 0.20 for CNN and MLP. As more synthetic fault samples are introduced, the macro-F1 rises sharply and reaches above 0.94 when the IR improves to 1:10 or 1:5, indicating that the classifiers are better able to recall minority fault classes. The MLP classifier in particular demonstrates the most benefit, reaching 97.5% accuracy and 0.9683 macro-F1 at $IR = 1:5$.

Notably, a performance regression is observed in certain classifiers at $IR = 1:20$, such as the CNN accuracy dropping to 84.72%. This fluctuation suggests that simply increasing the quantity of synthetic data does not always yield linear improvements. Rather, it implies the existence of an optimal augmentation ratio, where excessive synthetic samples may

introduce redundant generative noise that slightly destabilizes the decision boundary, whereas an appropriate ratio (e.g., 1:10 or 1:5) provides the necessary diversity for robust generalization.

To further validate these quantitative trends, the evolution of decision behavior is examined through confusion matrices under various IR settings (see Figure 17). At the most extreme imbalance ($IR = 1:100$), the SVM classifier predominantly assigns samples to the majority (normal) class, resulting in near-zero recall for fault categories despite the high overall accuracy. This visualization corroborates the "accuracy trap" discussed previously. As the representation of the minority classes is enhanced through PCS-GAN augmentation, a progressive shift in the decision boundary is observed, enabling the classifier to resolve distinct fault signatures. These qualitative results align with the macro-F1 improvements reported in Table 10, confirming that the proposed framework effectively restores class balance and enhances diagnostic reliability in data-constrained industrial environments.

Model	1:100	1:50	1:20	1:10	1:5	1:1
CNN	0.9615	0.9259	0.8472	0.9524	0.9417	0.9889
MLP	0.9615	0.9722	0.9500	0.9714	0.9750	0.9900
Random Forest	0.9663	0.9676	0.9500	0.9750	0.9778	0.9917
SVM	0.9856	0.9630	0.9542	0.9714	0.9750	0.9900

Table 10. Macro-F1 score under different IRs.

Model	1:100	1:50	1:20	1:10	1:5	1:1
CNN	0.1961	0.1923	0.2578	0.8629	0.8884	0.9882
MLP	0.1966	0.7461	0.8274	0.9493	0.9683	0.9900
Random Forest	0.2980	0.6324	0.8280	0.9428	0.9675	0.9905
SVM	0.6476	0.5720	0.8302	0.9408	0.9644	0.9895




Figure 16. Performance comparison across classifiers. Left: Accuracy. Right: macro-F1 score.




Figure 17. SVM confusion matrices under different IRs.

Robustness to Noise Interference

To assess noise robustness, we simulate varying SNR levels from ∞ (no noise) to -5dB (extreme) to mimic complex sensor environments. Figure 18 compares the diagnostic performance of models trained with and without PCS-GAN augmentation.

Experimental results indicate that while both models exhibit high diagnostic fidelity under favorable signal-to-noise ratio (SNR) conditions (≥ 20 dB), their performance diverges as noise intensity increases. At an SNR of 10 dB, the baseline model's macro-F1 score decreases to below 0.50, whereas the PCS-GAN-augmented framework maintains a robust performance of 0.98. Under extreme regimes (0 dB and -5 dB), where the noise

power is equivalent to or exceeds the signal power, the baseline model fails to yield statistically significant predictions. In contrast, the proposed model retains a Macro-F1 score of 0.59 at 0 dB. These findings define the operational limits of the diagnostic system under additive interference. Although such low-SNR conditions are typically beyond the conventional requirements for industrial deployment, the results demonstrate that the PCS-GAN extends the reliability boundaries of fault diagnosis. By optimizing on a physics-consistent manifold, the framework achieves enhanced resilience against the stochastic fluctuations and background noise inherent in complex mechanical environments.




Figure 18. Model performance under varying SNR levels.

5. Conclusion

This study presents the PCS-GAN framework for vibration signal synthesis under constrained and imbalanced data regimes.

By integrating data-driven generative capacity with physics-based principles, the architecture addresses the dual challenges of data scarcity and physical implausibility. Experimental evaluations on the CWRU and MFPT datasets demonstrate that

the synergy between the dual-domain discriminator and physics-constrained losses facilitates the synthesis of signals that accurately replicate periodic transients while maintaining spectral integrity. Quantitatively, the model achieves high distributional fidelity comparable to contemporary benchmarks while reducing peak GPU memory consumption by approximately 28%.

Ablation studies confirm that physical constraints stabilize the optimization landscape and define clear operational boundaries for the framework. Notably, PCS-GAN demonstrates superior resilience against noise interference, maintaining diagnostic utility under low-SNR conditions where conventional models typically fail. In extreme class-imbalance scenarios, the proposed augmentation effectively restores diagnostic recall, elevating classification performance to a robust range. Ultimately, these findings verify that embedding domain-specific mechanical dynamics into generative architectures enhances both the reliability and computational efficiency of diagnostic systems in complex industrial environments.

6. Discussion

The performance gains of the PCS-GAN model stem from directing the generative process toward manifolds consistent with mechanical principles. Unlike rigid point-frequency

matching, the integrated physics-constrained losses utilize spectral integration and energy distribution as priors. This approach provides inherent resilience to frequency shifts and geometric inaccuracies typical of industrial components.

Nevertheless, several areas for refinement are identified in this study. While the DCGAN-based backbone maintains computational efficiency, the incorporation of attention mechanisms may better capture long-range dependencies in non-stationary signals. Furthermore, replacing empirical weight calibration with automated hyper-parameter optimization could improve adaptability across diverse mechanical systems. The observed non-linear correlation between augmentation ratios and performance also indicates that identifying an optimal threshold is essential for maximizing generalization.

Future research will focus on the developing hybrid Transformer-based architectures to reconcile modeling capacity with physical interpretability. Specifically, gradient ratio scanning experiments will be conducted to quantitatively determine the exact optimal range of augmentation ratios, addressing the sensitivity of physical consistency to different data scales. Additionally, implementing adaptive calibration mechanisms and refining the model scale remain critical for low-latency deployment on edge computing devices, thereby supporting real-time predictive maintenance requirements.

Acknowledgment

This work was funded by the National Natural Science Foundation of China (Grants No. 11971302).

References

1. Shen S, Lu H, Sadoughi M, Hu C, Nemani V, Thelen A, Webster K, Darr M, Sidon J, Kenny S. A physics-informed deep learning approach for bearing fault detection. *Engineering Applications of Artificial Intelligence* 2021; 103: 104295. <https://doi.org/10.1016/j.engappai.2021.104295>.
2. Meng Z, Zhan X, Li J, Pan Z. An enhancement denoising autoencoder for rolling bearing fault diagnosis. *Measurement* 2018; 130: 448-454. <https://doi.org/10.1016/j.measurement.2018.08.010>.
3. Liu Y, Wen J, Wang G. A comprehensive overview of remaining useful life prediction: From traditional literature review to scientometric analysis. *Machine Learning with Applications* 2025; 21: 100704. <https://doi.org/10.1016/j.mlwa.2025.100704>.
4. Singh V, Gangsar P, Porwal R, Atulkar A. Artificial intelligence application in fault diagnostics of rotating industrial machines: A state-of-the-art review. *Journal of Intelligent Manufacturing* 2023; 34(3): 931-960. <https://doi.org/10.1007/s10845-021-01861-5>.
5. Meng D, Wang H, Yang S, Lv Z, Hu Z, Wang Z. Fault analysis of wind power rolling bearing based on EMD feature extraction. *Computer Modeling in Engineering & Sciences* 2022; 130(1): 543. <https://doi.org/10.32604/cmescs.2022.018123>.
6. Meng D, Nie P, Yang S, Su X, Liao C. Reliability analysis of wind turbine gearboxes: Past, progress and future prospects. *International Journal of Structural Integrity* 2025; 16(1): 4-38. <https://doi.org/10.1108/IJSI-08-2024-0129>.
7. Zhang T, Chen J, Li F, Zhang K, Lv H, He S, Xu E. Intelligent fault diagnosis of machines with small & imbalanced data: A state-of-the-

- art review and possible extensions. *ISA Transactions* 2022; 119: 152-171. <https://doi.org/10.1016/j.isatra.2021.02.042>.
8. Liu R, Yang B, Zio E, Chen X. Artificial intelligence for fault diagnosis of rotating machinery: A review. *Mechanical Systems and Signal Processing* 2018; 108: 33-47. <https://doi.org/10.1016/j.ymssp.2018.02.016>.
 9. Randall RB, Antoni J. Rolling element bearing diagnostics—a tutorial. *Mechanical Systems and Signal Processing* 2011; 25(2): 485-520. <https://doi.org/10.1016/j.ymssp.2010.07.017>.
 10. Zhou P, Wu D, Xu J, Wang Z, Ma D. Fault diagnosis method for rolling bearing based on probabilistic diffusion models under imbalanced data. *IEEE Sensors Journal* 2024; 24(23): 40059-40068. <https://doi.org/10.1109/JSEN.2024.3480135>.
 11. Wang R, Dong E, Cheng Z, Liu Z, Jia X. Transformer-based intelligent fault diagnosis methods of mechanical equipment: A survey. *Open Physics* 2024; 22(1): 20240015. <https://doi.org/10.1515/phys-2024-0015>.
 12. Goodfellow I J, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y. Generative adversarial nets. *Advances in Neural Information Processing Systems* 2014; 27.
 13. Zhao J, Huang W. Transfer learning method for rolling bearing fault diagnosis under different working conditions based on CycleGAN. *Measurement Science and Technology* 2021; 33(2): 025003. <https://doi.org/10.1088/1361-6501/ac3942>.
 14. Fu Z, Liu Z, Ping S, Li W, Liu J. TRA-ACGAN: A motor bearing fault diagnosis model based on an auxiliary classifier generative adversarial network and transformer network. *ISA Transactions* 2024; 149: 381-393. <https://doi.org/10.1016/j.isatra.2024.03.033>.
 15. Pan T, Chen J, Zhang T, Liu S, He S, Lv H. Generative adversarial network in mechanical fault diagnosis under small sample: A systematic review on applications and future perspectives. *ISA Transactions* 2022; 128: 1-10. <https://doi.org/10.1016/j.isatra.2021.11.040>.
 16. Zhang Y, Tiño P, Leonardis A, Tang K. A survey on neural network interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence* 2021; 5(5): 726-742. <https://api.semanticscholar.org/CorpusID:229678413>.
 17. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nature Reviews Physics* 2021; 3(6): 422-440. <https://doi.org/10.1038/s42254-021-00314-5>.
 18. Kim J, Yang Z, Lu Y. A physics-guided deep generative model for predicting melt pool behavior in laser powder bed fusion additive manufacturing. *Journal of Intelligent Manufacturing* 2025; 36(8): 5715-5735. <https://doi.org/10.1007/s10845-024-02504-1>.
 19. Lyu P, Cheng Y, Zhang M, Yu W, Xia L, Liu C. GPSC-GAN: A data enhanced model for intelligent fault diagnosis. *IEEE Transactions on Instrumentation and Measurement* 2024; 73: 1-16. <https://doi.org/10.1109/TIM.2024.3457925>.
 20. Sadoughi M, Hu C. Physics-based convolutional neural network for fault diagnosis of rolling element bearings. *IEEE Sensors Journal* 2019; 19(11): 4181-4192. <https://doi.org/10.1109/JSEN.2019.2898634>.
 21. Cheng W, Gao RX, Wang J, Wang T, Wen W, Li J. Envelope deformation in computed order tracking and error in order analysis. *Mechanical Systems and Signal Processing* 2014; 48(1-2): 92-102. <https://doi.org/10.1016/j.ymssp.2014.03.004>.
 22. McFadden P, Smith J. Model for the vibration produced by a single point defect in a rolling element bearing. *Journal of Sound and Vibration* 1984; 96(1): 69-82. [https://doi.org/10.1016/0022-460X\(84\)90595-9](https://doi.org/10.1016/0022-460X(84)90595-9).
 23. Seryasat O R, Aliyari shoorehdeli M, Honarvar F, Rahmani A. Multi-fault diagnosis of ball bearing using FFT, wavelet energy entropy mean and root mean square (RMS). In: *2010 IEEE International Conference on Systems, Man and Cybernetics*. IEEE; 2010: 4295-4299. <https://doi.org/10.1109/ICSMC.2010.5642389>.
 24. Arjovsky M, Chintala S, Bottou L. Wasserstein generative adversarial networks. In: *International Conference on Machine Learning*. PMLR; 2017: 214-223.
 25. Gulrajani I, Ahmed F, Arjovsky M, Dumoulin V, Courville AC. Improved training of wasserstein gans. *Advances in Neural Information Processing Systems* 2017; 30: 2234-2242. <https://dl.acm.org/doi/10.5555/3295222.3295327>.
 26. Mirza M. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*. Published online 2014. <https://doi.org/10.48550/arXiv.1411.1784>.
 27. Smola A, Gretton A, Song L, Schölkopf B. A hilbert space embedding for distributions. In: Hutter M, Servedio RA, Takimoto E, eds. *Algorithmic Learning Theory*. Springer Berlin Heidelberg; 2007: 13-31. https://doi.org/10.1007/978-3-540-75225-7_5.
 28. Liu S, Jiang H, Wu Z, Li X. Rolling bearing fault diagnosis using variational autoencoding generative adversarial networks with deep regret analysis. *Measurement* 2021; 168: 108371. <https://doi.org/10.1016/j.measurement.2020.108371>.
 29. Zhang T, Chen J, Xie J, Pan T. SASLN: Signals augmented self-taught learning networks for mechanical fault diagnosis under small sample

- condition. *IEEE Transactions on Instrumentation and Measurement* 2020; 70: 1-11. <https://doi.org/10.1109/TIM.2020.3043098>.
30. Jang K, Pilario KES, Lee N, Moon I, Na J. Explainable artificial intelligence for fault diagnosis of industrial processes. *IEEE Transactions on Industrial Informatics* 2023; 21(1): 4-11. <https://doi.org/10.1109/TII.2023.3240601>.
31. Raissi M, Perdikaris P, Karniadakis G E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 2019; 378: 686-707. <https://doi.org/10.1007/s44379-025-00051-x>.
32. Yonekura K. Physics-guided generative adversarial network to learn physical models. *arXiv preprint arXiv:230411488*. Published online 2023. <https://doi.org/10.48550/arXiv.2304.11488>.
33. Smith W A, Randall R B. Rolling element bearing diagnostics using the Case Western Reserve University data: A benchmark study. *Mechanical Systems and Signal Processing* 2015; 64: 100-131. <https://doi.org/10.1016/j.ymssp.2015.04.021>.
34. Bechhoefer E. A Quick Introduction to Bearing Envelope Analysis, MFPT Data. Published online 2012.
35. Radford A, Metz L, Chintala S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:151106434*. Published online 2015. <https://doi.org/10.48550/arXiv.1511.06434>.
36. Ho J, Jain A, Abbeel P. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* 2020; 33: 6840-6851. <https://doi.org/10.48550/arXiv.2006.11239>.
37. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez A N, Kaiser Ł, Polosukhin I. Attention is all you need. *Advances in Neural Information Processing Systems* 2017; 30: 5998-6008. <https://doi.org/10.48550/arXiv.1706.03762>.