

Industrial Anomaly Detecting in Spectral Data with Energy Consistency

Ziqi Wang¹, Hailiang Zhao^{1,*} and Shuiguang Deng²

¹ School of Software Technology, Zhejiang University, Ningbo 315100, China

² College of Computer Science and Technology, Zhejiang University, Hangzhou 310027, China

* Correspondence: hliangzhao@zju.edu.cn

How To Cite: Wang, Z.; Zhao, H.; Deng, S. Industrial Anomaly Detecting in Spectral Data with Energy Consistency. *Journal of Artificial Intelligence for Automation* **2026**, *1*(3), 17. <https://doi.org/10.53941/jaia.2026.100017>

Received: 26 January 2026

Revised: 2 September 2026

Accepted: 22 September 2026

Published: 16 October 2026

Abstract: Anomaly detection is critically important in ensuring the required stability and reliability of industrial equipment. Conventional methods for it typically assume the availability of sufficient samples from both normal and faulty operating conditions. In practice, however, real-world industrial scenarios rarely provide adequate abnormal samples, and fault data are extremely scarce. This scarcity makes it challenging to train robust and generalizable anomaly detection models. Although many studies have attempted to resolve this challenge by generating artificial anomaly samples, such synthetic samples often deviate from genuine fault patterns. As a result, models trained on them may overfit to unrealistic features and lower their generalization capability in practical applications. To address such issue, we propose Spectral-Regularized Denoising AutoEncoder (SR-DAE), that is designed for anomaly detection in high-dimensional spectral signals under anomaly-scarce conditions. Specifically, it introduces a band-structured spectral regularization strategy that guides the reconstruction process to preserve the coarse-grained energy distribution across different frequency sub-bands. SR-DAE improves sensitivity to subtle abnormal patterns while enhancing overall detection performance. We evaluate it on two turbocharger vibration datasets and compare it with several representative baseline methods. Experimental results demonstrate that SR-DAE outperforms its peers in terms of several commonly-used evaluation metrics. It can well balance between reducing false alarms and accurately identifying anomalous samples.

Keywords: industrial fault detection; denoising autoencoder; industrial monitoring; predictive maintenance; spectral analysis

1. Introduction

Anomaly detection is essential to ensure the reliable operation of industrial systems [1]. It can improve maintenance planning strategies [2] and reduce operational costs [3]. By enabling early fault identification, anomaly detection can help prevent severe equipment failures and support more stable industrial processes. This capability is particularly critical for key industrial devices such as turbochargers, pumps, and motors, which often exhibit subtle abnormal behaviors before major failures occur. Detecting these early warning signs can effectively reduce unexpected downtime, reduce maintenance expenses, and improve overall operational efficiency.

Several challenges hinder its practical implementation in real-world environments. A major challenge lies in the scarcity of anomalous data [4]. In industrial systems, abnormal events occur infrequently, making it difficult to collect sufficient fault samples for model training. Many conventional anomaly detection methods rely on normal and abnormal data during training [5]. However, practical industrial scenarios typically provide only a limited number of abnormal samples. In some cases, it is not even possible to obtain failure data. To mitigate this issue, researchers have commonly employed artificially generated anomalies. However, such synthetic anomalies often fail to fully capture the complexity and diversity of real fault patterns. Consequently, models trained with these samples



learn unrealistic abnormal characteristics, which weakens their ability to generalize to unseen fault conditions. Therefore, developing robust anomaly detection models capable of identifying unknown anomalies under limited fault-data conditions remains a challenging problem.

Another critical challenge lies in the effective analysis of high-dimensional spectral data [6]. Industrial equipment such as pumps and motors continuously generates vibration signals that reflect the underlying mechanical condition of the system. Since many fault-related patterns are more distinguishable in the frequency domain, these signals are commonly transformed into spectral representations for condition monitoring and fault diagnosis. In such representations, different frequency components are associated with different physical behaviors, while abnormal conditions may appear as variations in specific frequency bands. In practical industrial scenarios, spectral data contains hundreds or even thousands of frequency bins, forming a high-dimensional feature space. In addition, fault-related variations are usually subtle, localized, and easily obscured by irrelevant frequency components. Therefore, capturing fault-relevant spectral characteristics while suppressing irrelevant frequency variations remains a key challenge.

Under high-dimensional spectral conditions, traditional methods often struggle to identify features that are truly informative for anomaly detection [7]. Many existing approaches treat frequency bins in a nearly uniform manner, which may lead to over-smoothing effects during representation learning [8]. Consequently, important spectral details can be weakened or even obscured. For example, narrow spectral peaks associated with fault behavior are smoothed out during reconstruction, reducing the model's sensitivity to subtle but meaningful spectral changes. Such variations are often indicative of early-stage mechanical defects. In many industrial applications, mechanical faults are reflected not only by individual frequency components but also by changes in the energy distribution within specific frequency regions. Therefore, preserving the overall spectral energy profile is crucial for reliable fault diagnosis. For example, an abnormal increase or decrease in vibration energy within certain frequency bands indicates early equipment degradation. However, many commonly used reconstruction-based loss functions do not explicitly consider the energy structure embedded in spectral representations [9]. When the physical significance of different frequency bands is ignored, the reconstruction process fails to preserve fault-relevant energy variations, overlooking critical abnormal patterns.

To address these challenges, we propose Spectral-Regularized Denoising AutoEncoder (SR-DAE), a reconstruction-based anomaly detection framework for high-dimensional spectral signals under anomaly-scarce conditions. The proposed framework incorporates a Differentiable Filterbank layer, which adaptively assigns different levels of importance to frequency bands according to their physical relevance. It enables the model to focus on spectral components that are more closely associated with anomalous behaviors. Furthermore, we design a Spectral Band Loss to preserve energy consistency across different frequency bands. It is particularly important for detecting mechanical faults, which often manifest as energy variations within specific spectral regions. By constraining the reconstruction process at the band level, it helps to retain meaningful spectral energy patterns and improves sensitivity to fault-related changes. During inference, SR-DAE adopts a Conservative Dual-Score Inference Strategy, which jointly considers the reconstruction error and spectral energy error. By integrating these two complementary criteria, SR-DAE reduces false positive detections and improves the distinction between true anomalies and normal signal variations. The main contributions of this work are summarized as follows:

- We develop a novel Differentiable Filterbank layer that adaptively assigns importance weights to different frequency bands according to their physical relevance. By explicitly emphasizing informative spectral regions while suppressing less relevant components, the model can capture fault-related spectral patterns more effectively.
- We design a Spectral Band Loss to preserve band-level energy consistency during reconstruction. Unlike conventional reconstruction objectives, which mainly focus on point-wise errors, it constrains the reconstructed spectrum to maintain the original energy distribution across frequency bands. This design helps to retain physically significant spectral structures and improves the detection of mechanical faults characterized by localized energy variations.
- We introduce a Conservative Dual-Score Inference Strategy that jointly considers fine-grained reconstruction error and band-level spectral energy error. By integrating these complementary anomaly indicators, SR-DAE reduces the influence of normal signal fluctuations, suppresses false positive detections, and yields more reliable anomaly decisions under practical industrial conditions.

The rest of this paper is organized as follows. Section 2 reviews existing studies on industrial anomaly detection. Section 3 presents SR-DAE in detail. Section 4 introduces the experimental settings and reports the corresponding results. Finally, Section 5 concludes the paper.

2. Related Work

Various deep learning-based methods are developed to identify abnormal patterns across different application domains. Representative approaches include autoencoders, convolutional neural networks (CNNs), and generative adversarial networks (GANs). These methods show potential in modeling complex data distributions and detecting anomalous behaviors.

2.1. Autoencoder for Anomaly Detection

Autoencoders are widely adopted for anomaly detection due to their ability to learn compact representations of normal patterns and detect deviations through reconstruction errors. Ma et al. [10] propose a multiscale memory autoencoder with spatial filtering for hyperspectral anomaly detection. Its memory modules reconstruct background and anomalous pixels at multiple scales. However, the reliance on spatial filtering may limit its robustness in industrial scenarios with subtle anomalies and complex backgrounds. Tu et al. [11] develop a self-supervised Masked Graph AutoEncoder (MGAE) to reduce the dependence on anomaly annotations in hyperspectral images. By constructing a topological graph and using graph attention networks, MGAE captures long-range dependencies and non-Euclidean relationships during reconstruction, while a re-masking strategy further improves representation learning. However, its capability to model dynamic temporal characteristics remains limited. For industrial Internet of Things (IIoT) anomaly detection, Zeng et al. [12] introduce an adversarial autoencoder that integrates a variational autoencoder with particle swarm optimization for hyperparameter tuning. Although it achieves improved performance on several industrial datasets, the additional optimization procedure increases computational complexity and limits its applicability to real-time industrial scenarios.

For time-series anomaly detection, Lin et al. [13] propose a denoising autoencoder model based on the broad learning system. By constructing an ensemble of diverse denoising autoencoders, the method improves detection performance while maintaining low computational cost. Wang et al. [14] develop a multi-classification decision fusion model based on the Stacked Sparse Shrink AutoEncoder, which extracts multiple feature sets and integrates them through decision-level fusion to enhance anomaly classification. In hyperspectral anomaly detection, Wang et al. [15] introduce a Gaussian-inspired attention mechanism into an autoencoder framework. The attention layer focuses on the local neighborhood of each pixel, enabling more accurate feature extraction and improved detection performance. However, this locality-driven design limits its adaptability to dynamic industrial environments, where anomaly patterns and background conditions often change over time.

2.2. Convolutional Neural Networks for Anomaly Detection

CNNs are adopted in anomaly detection because of their strong ability to extract spatial features from complex data. Fu et al. [16] propose DeCNN-AD, a plug-and-play framework that integrates low-rank representation with a CNN denoiser to distinguish anomalous and background pixels. It clusters background pixels to construct a refined dictionary and exploits spatial correlations in representation coefficients. However, its reliance on clustering may reduce adaptability to dynamic datasets with evolving anomaly distributions. For few-shot anomaly detection, Zhu et al. [17] propose ADFormer, a dual CNN-Transformer architecture for limited-data scenarios. It combines CNN-based local feature extraction with Transformer-based global modeling, enabling more generalizable anomaly representations when only a few abnormal samples are available.

Beyond spatial feature extraction, CNN-based models are extended to capture more complex spatial-spectral structures and to support secure industrial deployment. Zhang et al. [18] propose a tensor-based transferred CNN for hyperspectral anomaly detection, where tensor structures and convolution kernels are used to jointly capture spatial-spectral features and measure discrepancies between test and learned feature tensors. However, the method is unsuitable for non-hyperspectral scenarios or noisy industrial environments. Cheng et al. [19] develop SecureAD, a secure CNN-based anomaly detection framework for edge computing environments. It employs additive secret sharing for secure CNN training and a Bloom-filter-based mechanism for fine-grained access control, protecting raw industrial data while reducing response latency and maintaining reliable detection performance in industrial edge-cloud systems.

2.3. Generative Adversarial Networks for Anomaly Detection

GANs are explored for anomaly detection, especially in high-dimensional or unlabeled data scenarios. By learning complex data distributions through adversarial training, GAN-based models can identify abnormal patterns as deviations from the learned normal behavior. Shim et al. [20] propose a dual-GAN framework for induction motor anomaly detection, where mel-spectrograms with different time and frequency resolutions are processed in parallel.

An ensemble strategy is then used to compute anomaly scores based on reconstruction loss and comprehensive feature distances. However, it depends heavily on mel-spectrogram representations, which restricts its applicability to other industrial sensor data. Kong et al. [21] develop AMBi-GAN for IIoT anomaly detection by integrating GANs, bidirectional LSTM networks, and an attention mechanism. It captures temporal dependencies in large-scale multidimensional time-series data and improves anomaly detection in complex industrial systems. However, its increased model complexity and computational cost limit its use in real-time industrial scenarios with constrained computational resources.

Beyond time-series and motor-monitoring applications, GAN-based models are explored for industrial security and visual defect detection. In industrial wireless sensor networks, Yang et al. [22] propose a GAN-based trust management system for malicious node detection without labeled data. It combines type-2 fuzzy logic for sensor-node reputation evaluation with a GAN structure for malicious node detection. Although it achieves a detection rate of 96% with a false positive rate below 8%, its dependence on the initial reputation evaluation reduces robustness in dynamic industrial environments. Long et al. [23] propose a multiscale GAN framework for defect detection in fabric labels. It generates templates from qualified FPL samples and adopts adaptive region-based strategies, including a statistical region subtractor and a region-adaptive thresholding mechanism, to improve defect localization. This framework enhances the identification of defective regions under complex industrial conditions.

2.4. Multimodal and Adaptive Fusion Methods for Industrial Anomaly Detection

Recent studies have increasingly explored multimodal and adaptive fusion methods to exploit complementary information from heterogeneous sensing channels. Representative approaches include cross-modal feature mapping between RGB and 3D signals [24], cross-modal distillation under incomplete modalities [25], fusion-restoration frameworks for reconstructing cross-domain embeddings [26], and geometry-guided score fusion for industrial anomaly detection [27]. In predictive maintenance, multi-sensor fusion improves fault diagnosis by integrating signals from different sensing locations or modalities, providing more comprehensive fault representations. These methods enhance anomaly representation through feature alignment, feature fusion, or knowledge distillation across multiple modalities.

SR-DAE focuses on single-modal spectral anomaly detection. In vibration-based monitoring, mechanical faults appear as subtle and localized energy variations within specific frequency regions. Therefore, an effective model should preserve both fine-grained spectral details and the overall energy distribution of spectral data. To this end, SR-DAE integrates spectral fidelity and physically informed spectral consistency into a unified reconstruction framework. Specifically, a Differentiable Filterbank layer is introduced to preserve band-wise spectral energy consistency across frequency sub-bands, enabling the model to retain fault-relevant energy shifts during reconstruction. Moreover, a conservative dual-score inference mechanism jointly considers fine-grained reconstruction error and coarse-grained spectral energy error, reducing false alarms and improving sensitivity to abnormal spectral patterns in high-dimensional industrial spectral data.

3. Methodological Framework

SR-DAE integrates representation learning with physically motivated spectral regularization to preserve the intrinsic structure of spectral data. As shown in Figure 1, it consists of two key components: A Denoising Autoencoder (DAE) backbone and a Differentiable Filterbank. The DAE backbone is responsible for learning robust latent representations by reconstructing clean spectra from corrupted inputs, which enables the model to capture fine-grained spectral characteristics while enhancing noise robustness. In parallel, the Differentiable Filterbank enforces energy consistency across physically meaningful frequency sub-bands, guiding the reconstruction process to preserve the global spectral energy distribution.




Figure 1. Architecture of SR-DAE. It consists of a DAE backbone that performs fine-grained spectral reconstruction and a Differentiable Filterbank branch that imposes coarse-grained spectral energy consistency during training.

3.1. Denoising Autoencoder Backbone

Let $\mathcal{X} = \{\mathbf{x}_i\}$ denote a set of vibration spectral snapshots collected from industrial equipment monitoring systems. Each snapshot $\mathbf{x}_i \in \mathbb{R}^D$ is a normalized power spectral density (PSD) vector at a specific time instant. It is obtained by transforming vibration signals from industrial machinery into the frequency domain using the Fourier transform [28]. D denotes the number of frequency bins in the PSD representation, with each bin corresponding to a specific frequency range. Thus, each $\mathbf{x}_i$ provides a high-dimensional spectral characterization of the vibration state of the monitored system. To enhance robustness and promote informative latent representation learning, Gaussian noise [29] is injected into the input spectra during training. For each spectral snapshot $\mathbf{x}_i$, a perturbation $\mathbf{n} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ is sampled to obtain the corrupted input $\tilde{\mathbf{x}}_i = \mathbf{x}_i + \mathbf{n}$, where σ^2 is the noise variance and $\mathbf{I}$ denotes the identity matrix, implying independent noise injection across frequency bins. The model is then trained to reconstruct the clean spectrum $\mathbf{x}_i$ from $\tilde{\mathbf{x}}_i$. This denoising design improves the robustness of learned spectral representations against noise interference in industrial vibration signals.

The backbone adopts a DAE architecture composed of an encoder and a decoder. The encoder E_ϕ projects the spectral input $\tilde{\mathbf{x}}_i$ into a compact latent representation $\mathbf{z}_i \in \mathbb{R}^d$:

$$\mathbf{z}_i = \sigma_E(\mathbf{W}_e \tilde{\mathbf{x}}_i + \mathbf{b}_e), \quad (1)$$

where $\mathbf{W}_e$ and $\mathbf{b}_e$ denote the encoder parameters, and σ_E is the ReLU activation function. The decoder D_θ then maps the latent representation back to the reconstructed spectrum:

$$\hat{\mathbf{x}}_i = \tanh(\mathbf{W}_d \mathbf{z}_i + \mathbf{b}_d), \quad (2)$$

where $\mathbf{W}_d$ and $\mathbf{b}_d$ are the decoder parameters. The output-layer $\tanh$ activation constrains $\hat{\mathbf{x}}_i$ to the interval $[-1, 1]$, matching the normalization range of the input spectra. This constraint helps ensure scale consistency between the input and reconstructed spectra.

3.2. Physically-Informed Spectral Regularization

A key limitation of the mean squared error (MSE) in reconstruction-based anomaly detection is that it assigns equal importance to all frequency bins. This uniform treatment leads to over-smoothing, weakening sharp spectral peaks that contain critical fault-related information in mechanical fault diagnosis, such as resonance responses or characteristic frequency components associated with specific defects. To overcome this issue, a Differentiable Filterbank layer is introduced to guide the model to emphasize informative frequency sub-bands while preserving fine-grained spectral details. This band-aware design helps to retain fault-relevant spectral structures that are suppressed by point-wise reconstruction objectives.

We define a fixed projection matrix $\mathbf{M} \in \mathbb{R}^{K \times D}$, where each row represents a triangular filter over the frequency bins. The element $M_{k,i}$ denotes the contribution weight of the i -th frequency bin to the k -th filter and is defined as

$$M_{k,i} = \max \left(0, 1 - \left| \frac{i - c_k}{w_k} \right| \right), \quad (3)$$

where c_k and w_k are the center and bandwidth of the k -th filter, respectively. Accordingly, bins closer to c_k receive larger weights, while the weights decrease linearly to zero as the distance from the center increases. The triangular filterbank is adopted as the default structured partition because rotating machinery spectra exhibit localized and smoothly varying energy redistribution, including resonance-related peaks, narrowband frequency shifts, and energy transitions across neighboring frequency regions. Overlapping triangular filters provide smooth inter-band transitions and allow the model to capture gradual energy transfer across adjacent spectral regions.

Then, the coarse-grained sub-band energy vector $\mathbf{e}$ is obtained by projecting the spectral vector $\mathbf{x}$ through the filterbank matrix $\mathbf{M}$, i.e., $\mathbf{e} = \mathbf{M}\mathbf{x}$. It summarizes the energy distribution across frequency sub-bands and provides a compact description of global spectral patterns relevant to fault detection. Since $\mathbf{M}$ is deterministic and differentiable, the resulting energy-related errors can be back-propagated through the network during training. To preserve the sub-band energy distribution during reconstruction, we define the *Spectral Band Loss* as

$$\mathcal{L}_{band} = \|\mathbf{M}\mathbf{x} - \mathbf{M}\hat{\mathbf{x}}\|_2^2, \quad (4)$$

where $\hat{\mathbf{x}}$ denotes the reconstructed spectrum. By penalizing band-level energy discrepancies, it enforces spectral consistency between the original and reconstructed spectra. This regularization is well suited to mechanical fault diagnosis, where incipient faults often appear as localized energy changes in specific spectral regions, including resonance-related components and wear-induced spectral signatures.

3.3. Training Loss

The overall training objective is designed to jointly preserve fine-grained spectral fidelity and coarse-grained energy consistency. The reconstruction loss encourages accurate point-wise recovery of the original spectrum, while the Spectral Band Loss regularizes the reconstruction at the frequency-band level to maintain the global spectral energy distribution. The complete objective is defined as

$$\mathcal{L}_{total} = \underbrace{\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2}_{\mathcal{L}_{rec}} + \lambda_{train} \underbrace{\|\mathbf{M}(\mathbf{x} - \hat{\mathbf{x}})\|_2^2}_{\mathcal{L}_{band}}, \quad (5)$$

where λ_{train} is a trade-off hyperparameter that balances point-wise reconstruction accuracy and coarse-grained spectral energy preservation.

3.4. Conservative Dual-Score Inference

A conservative decision strategy is adopted during inference to reduce false alarms. Given an input spectrum $\mathbf{x}$, the anomaly score $A(\mathbf{x})$ is computed by combining fine-grained reconstruction error and coarse-grained spectral energy discrepancy:

$$A(\mathbf{x}) = \mathcal{L}_{rec}(\mathbf{x}, \hat{\mathbf{x}}) + \lambda_{infer} \cdot \mathcal{L}_{band}(\mathbf{x}, \hat{\mathbf{x}}), \quad (6)$$

where λ_{infer} controls the contribution of the spectral band term during inference. The reconstruction error captures point-wise spectral deviations, whereas the spectral band loss measures inconsistencies in the sub-band energy distribution. This dual-score design reduces the influence of minor normal fluctuations that do not substantially affect the overall energy profile, suppressing false alarms and improving the reliability of anomaly detection.

To determine whether an input sample is normal or anomalous, the anomaly score is compared with a decision threshold τ . Specifically, a sample is classified as anomalous if its anomaly score exceeds the threshold:

$$y(\mathbf{x}) = \begin{cases} 1 \text{ (Anomaly)}, & \text{if } A(\mathbf{x}) > \tau, \\ 0 \text{ (Normal)}, & \text{otherwise.} \end{cases} \quad (7)$$

The threshold τ is estimated from the training set as the 95th percentile of the anomaly-score distribution. Since the training data are assumed to primarily represent normal operating conditions, this percentile-based threshold provides a conservative decision boundary that tolerates normal fluctuations while detecting samples with substantial deviations. The overall inference process is illustrated in Figure 2.




Figure 2. Inference-stage decision process of SR-DAE. The anomaly score is computed by jointly integrating the reconstruction error and the band-level spectral energy discrepancy. The final anomaly decision is obtained by comparing the resulting dual-score with a predefined threshold.

4. Experimental Evaluation

4.1. Experimental Setup

4.1.1. Dataset

We evaluate SR-DAE on real industrial vibration data collected from key rotating equipment in a nitrogen production plant. The datasets cover two machines under practical operating conditions: a first-stage turbo-compressor shaft with 4667 samples collected from 27 September to 27 October, and a low-pressure nitrogen compressor with 7001 samples collected from 1 to 30 October. Each vibration signal is transformed into the frequency domain using the Fast Fourier Transform (FFT), producing a 501-dimensional spectral feature vector. 80% of the samples from Device 1 and 70% from Device 2 are used for training, and the remaining samples are used for testing. Because labeled fault samples are unavailable, we introduce synthetic anomalies into the test set by adding Gaussian perturbations to normal samples. In the main performance evaluation, the anomaly ratios are approximately 33.3% for Device 1 and 66.7% for Device 2. In the mechanism-attribution study, an anomaly ratio of approximately 66.7% is used for both devices.

4.1.2. Parameter Setting

Table 1 summarizes the hyperparameter settings of SR-DAE. Latent dimension d is set to 16 to obtain compact representations of normal operating conditions while retaining sensitivity to abnormal spectral patterns. The filterbank consists of $K = 64$ triangular filters, which provides a suitable trade-off between spectral resolution and robustness. During training, Gaussian noise with $\sigma = 0.05$ is added to the input spectra to mitigate memorization and enhance noise robustness. The training coefficient is set to $\lambda_{train} = 7.0$, assigning greater weight to the preservation of the spectral energy envelope. During inference, $\lambda_{infer} = 2.0$ is adopted to increase sensitivity to band-level spectral discrepancies.

Table 1. Parameter settings for SR-DAE.

Parameter	Symbol	Value
Input dimension (frequency bins)	D	501
Latent dimension	d	16
Number of filterbank bands	K	64
Noise standard deviation	σ	0.05
Band loss weight (training)	λ_{train}	7.0
Band loss weight (inference)	λ_{infer}	2.0

4.1.3. Baseline Models

SR-DAE is compared with five anomaly detection baselines:

- MemAE [30,31]: A memory-augmented autoencoder that retrieves representative normal patterns from an external memory to suppress anomaly reconstruction.
- Isolation Forest [32]: A tree-based method that detects anomalies by recursively isolating samples through random feature partitioning.
- OC-SVM [33]: A kernel-based one-class classifier that learns a decision boundary enclosing normal samples.
- DCdetector [34]: A contrastive attention-based framework for time-series anomaly detection.
- MCM [35]: A multi-scale memory-based method that identifies anomalies according to deviations from stored normal representations.

These baselines include classical tree-based and kernel-based methods as well as recent deep anomaly detection models, enabling a broad evaluation of SR-DAE. All methods are tested under the same spectral snapshot anomaly detection protocol, with identical training/testing splits, spectral input features, and evaluation metrics. Methods originally designed for temporal or modality-specific settings are adapted to the snapshot-based spectral input.

4.2. Performance Comparison

Tables 2 and 3 present the quantitative performance comparisons of different methods. Model performance is evaluated by using (1) the area under the ROC curve (ROC-AUC); (2) precision-recall AUC (PR-AUC); (3) F1-score; (4) precision; (5) recall; and (6) accuracy. We use AUC-based metrics to assess the overall discriminative ability of a method in separating normal and anomalous samples. We use precision, recall, and F1-score to show the trade-off between anomaly detection and false alarm control.

On Device 1, SR-DAE achieves the best overall performance among all compared methods, obtaining the highest ROC-AUC, PR-AUC, and F1-score. It has a strong ability to distinguish anomalous samples from normal operating data. It also achieves a favorable balance between precision and recall, suggesting that it can effectively identify anomalies while maintaining a relatively low false alarm rate. Among the baselines, MCM shows the closest performance to SR-DAE. In contrast, Isolation Forest and DCdetector exhibit a more evident trade-off between precision and recall. Although DCdetector achieves the highest recall, its precision decreases to 0.6016, resulting in an F1-score of 0.6956. MemAE shows the opposite tendency, achieving relatively high precision (0.8643) but lower recall (0.6274). OC-SVM performs the weakest, with a ROC-AUC of 0.8014 and an F1-score of 0.6352, indicating limited effectiveness in this task.

On Device 2, SR-DAE achieves the best overall performance. It achieves a ROC-AUC of 0.8501 and a PR-AUC of 0.9361, demonstrating robust discriminative capability under different equipment conditions. SR-DAE also obtains a strong F1-score of 0.8291 and a precision of 0.8415, further confirming its effectiveness in identifying anomalous samples. Its relatively lower recall is consistent with the conservative inference design, which emphasizes false positive reduction, *i.e.*, avoiding the incorrect classification of normal samples as anomalies, while still maintaining an acceptable level of anomaly detection capability. Although MCM and MemAE remain competitive in F1-score and precision, their overall accuracy is lower than that of SR-DAE. OC-SVM again shows weaker performance, suggesting that conventional one-class boundary estimation is insufficient for high-dimensional industrial spectral data.

Table 2. Comparison of Metrics on Device 1. The best results are highlighted in bold.

Model	ROC-AUC	PR-AUC	F1-Score	Precision	Recall	Accuracy
SR-DAE	0.9447	0.9317	0.8638	0.9558	0.7880	0.9172
DCdetector	0.8688	0.8724	0.6956	0.6016	0.8244	0.7595
MCM	0.8785	0.8657	0.7495	0.7495	0.7495	0.8330
MemAE	0.8589	0.8324	0.7270	0.8643	0.6274	0.8430
Isolation Forest	0.8749	0.7863	0.7280	0.6703	0.7966	0.8016
OC-SVM	0.8014	0.7614	0.6352	0.5769	0.7066	0.7295

Table 3. Comparison of Metrics on Device 2. The best results are highlighted in bold.

Model	ROC-AUC	PR-AUC	F1-Score	Precision	Recall	Accuracy
SR-DAE	0.8501	0.9361	0.8291	0.8415	0.8170	0.7754
DCdetector	0.7526	0.8858	0.7680	0.7252	0.8161	0.6712
MCM	0.8298	0.9255	0.8101	0.8101	0.8101	0.7468
MemAE	0.8022	0.9111	0.8015	0.6698	0.9977	0.6706
Isolation Forest	0.7372	0.8563	0.8033	0.6875	0.9660	0.6846
OC-SVM	0.7559	0.8837	0.8002	0.6735	0.9854	0.6718

4.2.1. Operational Safety and False Positive Analysis

ROC-AUC provides an overall assessment of anomaly detection performance, but practical industrial applications require particular attention to the false positive rate (FPR). Excessive false alarms can impose substantial operational burdens and increase unnecessary maintenance costs. To further evaluate the practical reliability of SR-DAE, we compare the trade-off between detection sensitivity and false alarm control across SR-DAE and the baseline methods. The corresponding results are shown in Table 4. SR-DAE achieves a favorable balance between suppressing false alarms and maintaining strong anomaly detection capability.

Table 4. Comparison of FPR, recall, precision, and ROC-AUC on Device 1. The results evaluate the trade-off between false alarm control and anomaly detection sensitivity. The best results are highlighted in bold.

Model	FPR ↓	Recall ↑	Precision	ROC-AUC
OC-SVM	0.258	0.707	0.577	0.801
DCdetector	0.272	0.824	0.602	0.869
Isolation Forest	0.195	0.797	0.670	0.875
MCM	0.125	0.750	0.750	0.879
MemAE	0.049	0.627	0.864	0.859
SR-DAE (Ours)	0.018 (1.8%)	0.788	0.956	0.945

SR-DAE obtains an exceptionally low FPR of 1.8%, achieving the best result among all compared methods. In contrast, DCdetector produces an FPR of 27.2%, which is approximately 15 times higher than that of SR-DAE. Isolation Forest exhibits a similar limitation, with an FPR about 10.7 times greater than that of SR-DAE. These results suggest that some baseline methods improve recall by adopting a more aggressive anomaly detection strategy. However, it inevitably introduces a large number of unnecessary alarms into the industrial monitoring process, reducing its practical applicability. For the detection sensitivity, SR-DAE also compares favorably with more conservative methods. For example, MemAE maintains a relatively low FPR of 4.9%, but its recall decreases substantially to 62.7%. In contrast, SR-DAE achieves a much lower FPR while preserving a strong recall of 78.8%. These results demonstrate that the spectral regularization mechanism produces more reliable anomaly alarms by reducing the influence of normal operational fluctuations, while still retaining sensitivity to critical fault-related events.

4.3. Parameter Sensitivity

A sensitivity study is conducted on four key parameters, as shown in Figure 3. Specifically, we vary the training regularization coefficient $\lambda_{\text{train}} \in \{0.1, 1.0, 3.0, 5.0, 7.0, 10.0, 15.0\}$, the inference coefficient $\lambda_{\text{infer}} \in \{0.1, 0.3, 0.5, 0.7, 1.0, 1.5, 2.0\}$, the latent dimension $d \in \{16, 32, 48, 64, 96, 128, 192\}$, and the number of spectral bands $K \in \{8, 16, 24, 32, 48, 64, 96\}$. The sensitivity analysis is reported separately from the main performance comparison.




Figure 3. Parameter sensitivity analysis of SR-DAE on Device 1. Each stacked bar shows three performance metrics (ROC-AUC, PR-AUC, and F1-score) evaluated under different parameter settings: (a) λ_{train} controls the band energy reconstruction loss weight during training; (b) λ_{infer} balances reconstruction and band-level anomaly scores during inference; (c) d denotes the latent representation dimension; (d) K represents the number of spectral bands in the filterbank.

As shown in Figure 3a, SR-DAE achieves the best performance when $\lambda_{\text{train}} = 7.0$. It indicates that strong band-energy regularization during training improves anomaly discrimination by preserving coarse-grained spectral energy structures. For the inference coefficient, Figure 3b indicates that $\lambda_{\text{infer}} = 2.0$ achieves the best performance. This suggests that incorporating band-level spectral consistency into anomaly scoring improves the robustness of inference-stage decision-making. Regarding the latent dimension, Figure 3c shows that the compact setting $d = 16$ achieves the best overall performance. This suggests that SR-DAE does not require a high-dimensional latent space to achieve effective anomaly detection, as the proposed spectral regularization helps preserve fault-relevant information even under compact representation learning. Finally, as shown in Figure 3d, the best performance is obtained when $K = 64$, where the model reaches a ROC-AUC of 0.960, a PR-AUC of 0.947, and an F1-score of 0.881. 64 spectral bands provide an effective balance between frequency resolution and robustness, allowing SR-DAE to capture fault-related spectral deviations across multiple sub-band scales.

4.4. Ablation Study

We conduct an ablation study on Device 1 (Figure 4). The complete SR-DAE framework is compared with three degraded variants: (1) *w/o Band Loss*: the band-energy reconstruction term is removed from the training objective by setting $\lambda_{\text{train}} = 0$; (2) *w/o Band Score*: the band-level energy score is excluded during inference; and (3) *w/o Noise*: Gaussian noise injection is disabled and the model is trained directly on clean samples.

Figure 4 shows that all components make positive contributions to performance. SR-DAE achieves a ROC-AUC of 0.909, an F1-score of 0.786, and a recall of 0.762. Removing the band-energy loss leads to a noticeable

performance degradation, with ROC-AUC and F1-score decreasing by 5.7% and 9.4%, respectively. It confirms that the spectral regularization mechanism helps the model to learn latent representations that preserve frequency-domain structural information. Excluding the band-level score during inference causes a 6.0% reduction in ROC-AUC, indicating that band-level energy discrepancy provides complementary discriminative information beyond point-wise reconstruction error. Among all degraded variants, disabling noise injection results in the largest recall decline, from 0.762 to 0.478. This demonstrates the importance of the denoising strategy for learning robust representations that generalize better to unseen anomalies.




Figure 4. Ablation study results on Device 1. We evaluate the contribution of each key component by comparing the full SR-DAE model against three variants: (a) removing band energy loss from training ($\lambda_{\text{train}} = 0$); (b) removing band score from inference; and (c) training without input noise.

We compare several alternative grouping strategies under the same backbone architecture to further examine whether the performance gains depend on the specific triangular partition. As reported in Table 5, structured grouping strategies consistently outperform the random projection variant, suggesting that the improvement mainly comes from meaningful band-level spectral organization rather than arbitrary projection or generic regularization. Meanwhile, triangular, equal-width, log-triangular, and learnable grouped projection strategies all achieve competitive performance. Therefore, we do not claim that the triangular filterbank is the only effective implementation. Instead, it is adopted as the default design because it offers a simple, stable, and interpretable realization of the proposed band-structured spectral constraint.

Table 5. Mechanism-attribution study of spectral grouping strategies. All variants are evaluated under matched settings with the same DAE backbone, training budget, and training-only threshold selection rule. The anomaly ratio in this study is approximately 66.7%.

Device	Variant	F1	Precision	Recall	FPR
Device 1	SR-DAE (triangular)	0.7373	0.9490	0.6028	0.0648
	Learnable grouped projection	0.7383	0.9487	0.6043	0.0654
	Equal-width grouping	0.7316	0.9465	0.5963	0.0675
	Log-triangular grouping	0.7311	0.9474	0.5953	0.0662
	Harmonic grouping	0.7234	0.9453	0.5858	0.0678
	Plain DAE	0.7184	0.9474	0.5786	0.0643
	Generic latent L_2	0.7070	0.9426	0.5657	0.0689
	Random projection grouping	0.6849	0.9519	0.5349	0.0541
Device 2	SR-DAE (triangular)	0.7777	0.9592	0.6540	0.0557
	Learnable grouped projection	0.7733	0.9583	0.6482	0.0564
	Equal-width grouping	0.7778	0.9575	0.6550	0.0582
	Log-triangular grouping	0.7774	0.9577	0.6542	0.0578
	Harmonic grouping	0.7670	0.9592	0.6390	0.0544
	Plain DAE	0.7730	0.9560	0.6487	0.0597
	Generic latent L_2	0.7741	0.9583	0.6494	0.0565
	Random projection grouping	0.7483	0.9576	0.6142	0.0544

4.5. Case Study

To demonstrate the anomaly localization capability of SR-DAE under practical industrial conditions, we conduct a case study using an injected anomalous sample that is successfully detected by the model. As shown in Figure 5, the detection process is visualized by decomposing the anomaly score into interpretable components.

Figure 5a compares the normal and anomalous spectra, where the injected abnormal region is highlighted in green as the reference region. Figure 5b presents the frequency-wise reconstruction error $\mathcal{L}_{rec}$ together with the 95th percentile threshold. It shows that the reconstruction error captures fine-grained spectral deviations but does not fully reflect broader energy changes across frequency bands. Figure 5c further shows the band-level energy deviation $\mathcal{L}_{band}$ computed using the triangular filterbank, with anomalous spectral bands highlighted in red. The results indicate that $\mathcal{L}_{band}$ can identify abnormal frequency regions that are not completely captured by point-wise reconstruction error. This observation demonstrates that the reconstruction error and band-level energy discrepancy provide complementary evidence, enabling more comprehensive and interpretable fault localization.




Figure 5. Case study of SR-DAE anomaly detection on Device 1. (a) Normal v.s. anomalous spectrum with ground-truth injected region highlighted in green. (b) Per-frequency reconstruction error $\mathcal{L}_{rec}$. (c) Band-wise energy deviation $\mathcal{L}_{band}$ through triangular filterbank with anomalous bands highlighted.

The adopted Gaussian perturbations provide a controlled benchmark for evaluating anomaly detection under anomaly-scarce conditions. However, real mechanical faults manifest as more structured spectral deviations. To examine whether SR-DAE is overly dependent on Gaussian-type perturbations, we evaluate its robustness under three representative perturbation families: Gaussian noise, resonance bumps, and sidebands. As reported in Table 6, SR-DAE achieves the highest F1-score across all perturbation settings, indicating that the effectiveness of SR-DAE is not limited to Gaussian perturbations.

Table 6. F1-score comparison under different perturbation settings. The best results are highlighted in bold.

Perturbation	DAE	Isolation Forest	PCA	SR-DAE
Gaussian noise	0.9853	0.7436	0.9856	0.9863
Resonance bump	0.9747	0.5759	0.9503	0.9857
Sidebands	0.3097	0.1658	0.3921	0.4126

4.6. Computational Complexity Analysis

We evaluate SR-DAE's computational efficiency in terms of FLOPs, memory consumption, and detection accuracy. It achieves an F1-score of 0.864 with only 0.012 GB of memory consumption and 11M FLOPs per iteration, indicating both strong detection performance and low computational overhead. Unlike conventional methods such as Isolation Forest and OC-SVM, which may require storing or accessing large amounts of training data during inference, SR-DAE relies on a compact encoder-decoder architecture with spectral regularization. It supports low-latency inference and reduces hardware requirements, making it well suited for continuous industrial monitoring on resource-constrained edge devices.

5. Conclusions

We propose Spectral-Regularized Denoising AutoEncoder (SR-DAE), an anomaly detection framework that integrates deep representation learning with physical spectral consistency. It is designed to address two major challenges in industrial anomaly detection: the scarcity of anomalous samples and the difficulty of modeling high-dimensional spectral data. By incorporating a Differentiable Filterbank and a Spectral Band Loss, SR-DAE preserves band-wise energy consistency across frequency sub-bands and retains fault-relevant spectral structures during reconstruction. It enables the model to capture subtle spectral variations associated with incipient mechanical defects. Extensive experiments on real industrial vibration datasets demonstrate that SR-DAE consistently outperforms its peers across several evaluation metrics. It further shows that the spectral regularization mechanism improves detection reliability by reducing false alarms while maintaining strong sensitivity to abnormal spectral patterns.

In future work, we will extend SR-DAE to larger-scale industrial datasets and investigate its applicability to broader sensing modalities and operating conditions. Moreover, integrating it with predictive maintenance systems can enable more proactive fault management of industrial equipment.

Author Contributions

Z.W.: conceptualization, methodology, software; H.Z.: data curation, original draft preparation; S.D.: visualization, investigation. All authors have read and agreed to the published version of the manuscript.

Funding

This research was funded by National Key Research and Development Program of China under Grant 2024YFB3309400.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The data will be made available upon reasonable request.

Conflicts of Interest

The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies

During the preparation of this work, the authors used GPT to polish the language. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

1. Wang, G.; Zou, Y.; He, S.; et al. Anomaly Detection and Localization via Reverse Distillation with Latent Anomaly Suppression. *IEEE Trans. Circuits Syst. Video Technol.* **2025**, *35*, 9592–9607.
2. Maggipinto, M.; Beghi, A.; Susto, G.A. A Deep Convolutional Autoencoder-Based Approach for Anomaly Detection with Industrial, Non-Images, 2-Dimensional Data: A Semiconductor Manufacturing Case Study. *IEEE Trans. Autom. Sci. Eng.* **2022**, *19*, 1477–1490.

3. Yuan, S.; Song, M.; Song, J.; et al. Robust Industrial Anomaly Detection via Style Shift Estimation and Cascade Distillation. *IEEE Trans. Instrum. Meas.* **2025**, *74*, 5034014.
4. Liu, B.; Guo, T.; Luo, B.; et al. Cross-Attention Regression Flow for Defect Detection. *IEEE Trans. Image Process.* **2024**, *33*, 5183–5193.
5. Cao, G.; Feng, Z.; Wang, T. A Semi-Supervisory Anomaly Detection Method for Industrial Networks Security. In Proceedings of the 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC), Toronto, ON, Canada, 11–14 October 2020; pp. 2125–2129.
6. Cheng, D.; Jiang, X.; Xia, S.; et al. Pseudo-Label-Based Unsupervised Granular-Ball Division and Fast Spectral Clustering for High-Dimensional Data. In Proceedings of the 2025 IEEE 41st International Conference on Data Engineering (ICDE), Hong Kong, 19–23 May 2025; pp. 4106–4119.
7. Kong, W.; Shen, D.; Ma, X.; et al. Hyperspectral Anomaly Detection Based on High-Order Coupled Fully Connected Tensor Network Decomposition. In Proceedings of the IGARSS 2025—2025 IEEE International Geoscience and Remote Sensing Symposium, Brisbane, Australia, 3–8 August 2025; pp. 8672–8676.
8. Jin, W.; Dang, F.; Zhu, L. Feature Enhancement with Reverse Distillation for Hyperspectral Anomaly Detection. *IEEE Geosci. Remote Sens. Lett.* **2024**, *21*, 5509705.
9. Ma, X.; Zhang, X.; Huyan, N.; et al. Hyper-Laplacian Regularized Low-Rank Tensor Decomposition for Hyperspectral Anomaly Detection. In Proceedings of the IGARSS 2018—2018 IEEE International Geoscience and Remote Sensing Symposium, Valencia, Spain, 22–27 July 2018; pp. 6380–6383.
10. Ma, Z.; Zhang, Y.; Lian, Y.; et al. Multiscale Memory Autoencoder and Spatial Filtering for Hyperspectral Anomaly Detection. *IEEE Geosci. Remote Sens. Lett.* **2025**, *22*, 5501705.
11. Tu, B.; He, B.; He, Y.; et al. Self-Supervised Masked Graph Autoencoder for Hyperspectral Anomaly Detection. *IEEE Trans. Image Process.* **2025**, *34*, 6714–6729.
12. Zeng, G.Q.; Yang, Y.W.; Lu, K.D.; et al. Evolutionary Adversarial Autoencoder for Unsupervised Anomaly Detection of Industrial Internet of Things. *IEEE Trans. Reliab.* **2025**, *74*, 3454–3468.
13. Lin, Y.; Yu, Z.; Yang, K.; et al. Ensemble Denoising Autoencoders Based on Broad Learning System for Time-Series Anomaly Detection. *IEEE Trans. Neural Netw. Learn. Syst.* **2025**, *36*, 13913–13926.
14. Wang, Z.; Guan, Z.; Wu, X.; et al. Multi-Classification Decision Fusion Based on Stacked Sparse Shrink AutoEncoder and GS-Tabnet for Network Intrusion Detection. In Proceedings of the 2024 10th International Conference on Control, Decision and Information Technologies (CoDIT), Vallette, Malta, 1–4 July 2024; pp. 2560–2565.
15. Wang, R.; Hu, J. Gaussian-Inspired Attention Mechanism for Hyperspectral Anomaly Detection. *IEEE Geosci. Remote Sens. Lett.* **2025**, *22*, 5500705.
16. Fu, X.; Jia, S.; Zhuang, L.; et al. Hyperspectral Anomaly Detection via Deep Plug-and-Play Denoising CNN Regularization. *IEEE Trans. Geosci. Remote Sens.* **2021**, *59*, 9553–9568.
17. Zhu, B.; Gu, Z.; Zhu, G.; et al. ADFormer: Generalizable Few-Shot Anomaly Detection with Dual CNN-Transformer Architecture. *IEEE Trans. Instrum. Meas.* **2025**, *74*, 1–16.
18. Zhang, L.; Cheng, B. Transferred CNN Based on Tensor for Hyperspectral Anomaly Detection. *IEEE Geosci. Remote Sens. Lett.* **2020**, *17*, 2115–2119.
19. Cheng, H.; Liu, X.; Wang, H.; et al. SecureAD: A Secure Video Anomaly Detection Framework on Convolutional Neural Network in Edge Computing Environment. *IEEE Trans. Cloud Comput.* **2022**, *10*, 1413–1427.
20. Shim, J.; Seo, J.; Lee, S.; et al. Ensemble-Based Anomaly Detection with Comprehensive Feature Distances for Induction Motors. *IEEE Trans. Ind. Appl.* **2024**, *60*, 6033–6044.
21. Kong, F.; Li, J.; Jiang, B.; et al. Integrated Generative Model for Industrial Anomaly Detection via Bidirectional LSTM and Attention Mechanism. *IEEE Trans. Ind. Inform.* **2023**, *19*, 541–550.
22. Yang, L.; Yang, S.X.; Li, Y.; et al. Generative Adversarial Learning for Trusted and Secure Clustering in Industrial Wireless Sensor Networks. *IEEE Trans. Ind. Electron.* **2023**, *70*, 8377–8387.
23. Long, J.; Lan, X.; Yan, S.; et al. Multiscale GAN with Region Adaptive Schemes for Online Inspection of Fabric Flexographic Printing Labels. *IEEE Trans. Instrum. Meas.* **2025**, *74*, 2510118.
24. Brenner, M.; Reyes, N.H.; Susnjak, T.; et al. RGB-D and Thermal Sensor Fusion: A Systematic Literature Review. *IEEE Access* **2023**, *11*, 82410–82442.
25. Wei, S.; Luo, Y.; Ma, X.; et al. MSH-Net: Modality-Shared Hallucination with Joint Adaptation Distillation for Remote Sensing Image Classification Using Missing Modalities. *IEEE Trans. Geosci. Remote Sens.* **2023**, *61*, 4402615.
26. Karam, L.J.; Sadaka, N.G.; Ferzli, R.; et al. An Efficient Selective Perceptual-Based Super-Resolution Estimator. *IEEE Trans. Image Process.* **2011**, *20*, 3470–3482.
27. Tan, Q.; Wei, R.; Xi, Z.; et al. Consistent3D: Diffusion-Driven Sparse View Completion and 3D Reconstruction with Geometric Priors. In Proceedings of the 2025 International Conference on Digital Image Computing: Techniques and Applications (DICTA), Adelaide, Australia, 3–5 December 2025; pp. 1–8.
28. Pei, S.C.; Hsue, W.L.; Ding, J.J. Discrete Fractional Fourier Transform Based on New Nearly Tridiagonal Commuting Matrices. *IEEE Trans. Signal Process.* **2006**, *54*, 3815–3828.

29. Kim, Y.H. Feedback Capacity of the First-Order Moving Average Gaussian Channel. *IEEE Trans. Inf. Theory* **2006**, 52, 3063–3079.
30. Gong, D.; Liu, L.; Le, V.; et al. Memorizing Normality to Detect Anomaly: Memory-Augmented Deep Autoencoder for Unsupervised Anomaly Detection. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, South Korea, 27 October–2 November 2019; pp. 1705–1714.
31. Wang, F.; Jiang, Y.; Zhang, R.; et al. A Survey of Deep Anomaly Detection in Multivariate Time Series: Taxonomy, Applications, and Directions. *Sensors* **2025**, 25, 190.
32. Liu, F.T.; Ting, K.M.; Zhou, Z.H. Isolation Forest. In Proceedings of the 2008 Eighth IEEE International Conference on Data Mining, Pisa, Italy, 15–19 December 2008; pp. 413–422.
33. Schölkopf, B.; Platt, J.C.; Shawe-Taylor, J.; et al. Estimating the Support of a High-Dimensional Distribution. *Neural Comput.* **2001**, 13, 1443–1471.
34. Yang, Y.; Zhang, C.; Zhou, T.; et al. DCDetector: Dual Attention Contrastive Representation Learning for Time Series Anomaly Detection. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Long Beach, CA, USA, 6–10 August 2023 ; pp. 3033–3045.
35. Jezequel, L.; Vu, N.S.; Beaudet, J.; et al. Unified Anomaly Detection via Multi-Scale Contrasted Memory. *IEEE Trans. Image Process.* **2026**, 35, 2802–2815.