iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

ZERO SHOT FAULT CLASSIFICATION FOR RARE COMPRESSOR FAILURE EVENTS USING GENERATIVE MODELS

Kamal Ahmed Aboelkamal Ali
kamalaboelkamal@live.com

Mohammed Ahmed Mohammed Alshaqsi
mohammed2.alshaqsi@live.uwe.ac.uk

ABSTRACT
The fact that rare failure events in industrial compressors are challenging to data-driven fault diagnose is a long-standing problem due to the fact that the conditions that result in such events are hard to reproduce and may not be present in historical data. This paper presents a zero-shot learning method with the help of generative models to categorize compressor faults with no real training data. The generation of synthetic fault signature is based on a cycle-consistent adversarial framework and is verified by similarity checks on the distribution level to maintain the synthetic fault signature to be representative of realistic abnormal behavior. Such samples are created and matched with semantic fault descriptions so that a zero-shot classifier can experience missed failure types based solely on their characteristics. The compressor vibration data experiments demonstrate that the generative-zero-shot pipeline is better at recognizing the infrequent faults in contrast with the traditional supervised and anomaly-based baselines and is notably well-performing in instances where the real samples are very scarce or absent. The results indicate that generative augmentation with zero-shot inference should be regarded as a promising future in early-stage detection of low-frequency failures in rotating machinery and can be used to complement industrial predictive-maintenance processes in which data scarcity is the reality.

Keywords
Zero-shot learning; generative models; compressor fault diagnosis; rare failure detection; predictive maintenance; industrial machinery health monitoring

1. INTRODUCTION

The importance of industrial compressors to manufacturing, energy systems, and process industries is that the level of efficiency and safety in operations are directly influenced by the reliability of the compressors. However, the very failures that are the most dangerous are also the most difficult to notice at the very beginning since unusual fault conditions are not a frequent occurrence in real operating data. Gathering run-to-fail data on such events can require months or years of data collection and most major faults only occur under extreme or unrealistic conditions that are not feasible to recreate in a controlled environment. The conventional supervised learning techniques fail in such situations, as they require massive and balanced data and cannot extrapolate to fault classes that they have never seen (Lei et al., 2025; Khan et al., 2025).

The recent literature has started to move in the direction of methods that deal directly with data scarcity, particularly using generative modeling. Generative adversarial networks and analogous designs have the potential to generate synthetic fault signals, which behave in a manner similar to real machines, providing an opportunity to address the limitations of the usage of true samples (where they may be scarce or not exist) (Di Maggio et al., 2023; Apostolou et al., 2025). These models are also becoming popular, as assistive tools in augmenting training information in diagnostic tasks where rare events need to be modeled with fidelity. This has been further increased by diffusion-based generative models and score-conditioned sampling models that support the possibility of generating more difficult to obtain patterns and better probability of having true representations of rare events to appear in downstream classifiers (Kim et al., 2025).

Zero-shot learning has also become a complementary approach, and it is used besides synthetic data generation to address unseen fault types. Rather than using only labeled examples, zero-shot methods use semantic description,

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

attributes or learned embeddings to determine classes that do not have training examples. Anomaly and out-of-distribution detection surveys note that such a change is essential to the industrial context, where invisible failures are not the exception but the rule (Salehi et al., 2021). More modern generative systems which operate retrieval-augmented systems also indicate the possibility of building architectures that can learn new classes by incorporating contextual or multimodal information, even without a matching historical sample (Mei et al., 2025).

The necessity of more flexible diagnostic models is supported by the evidence of the time-series and predictive-maintenance community in general. Both predictions made by generative predictors and compression-based models indicate that systems able to learn structural patterns instead of memorizing labeled categories are more likely to do well in unusual conditions (Dubois et al., 2021; Wan et al., 2025). Surveys of observational bias in prognostics also report that the data collected by industrial sensors is hardly representative enough to be used with conventional supervised learning, and generative augmentation can be used to mitigate these biases by reconstructing patterns that the sensor data just never measured (Fink et al., 2025).

Collectively, these advances indicate a distinct vacuum, since even though there is progress in generative data augmentation and zero-shot inference, little has been done to unite the two with compressor diagnostics, where infrequent failures are a major challenge to operational viability. The research paper fills that gap and examines a generative zero-shot design that fits the nature of compressor vibration data and the risk of using very small samples of faults. The proposed approach is expected to distinguish previously unknown compressor failures with better reliability by synthesizing realistic fault signatures and combining them with semantic descriptors. Such a combination provides a promising future of predictive maintenance systems that will perform operations in environments with limited data and help bridge the gap that has been open for a long time between infrequent events modeling and machinery health monitoring in reality.

2. LITERATURE REVIEW

Industrial fault diagnosis Research on this area has increasingly been attracted to the issues of infrequent and invisible failure modes, especially in rotating machinery like compressors. Among the recent surveys that have shown the most consistent results is the fact that traditional supervised models cannot be generalized when the training data does not represent low-frequency or abnormal conditions (Lei et al., 2025; Khan et al., 2025). Actual industrial data is usually highly unbalanced, biased towards normal operating conditions, and it often does not include early warning or disastrous failures that happen in very few instances throughout the lifetime of the machine. This architectural shortcoming has prompted the investigation into alternative strategies that are able to acquire meaningful patterns even in cases where labelled instances of particular faults are not available.

Generative modeling is one of the most promising directions in solving this lack. Initial experiments have shown that adversarial architecture is capable of generating realistic vibration patterns that reproduce the spectral and time-wise properties of actual mechanical faults (Di Maggio et al., 2023). These models minimise reliance on expensive run-to-failure tests by generating artificial fault samples to model conditions that are often very difficult to model in practice. The generative models are also versatile in industrial applications with recent surveys on privacy-preserving and resource-aware generative adversarial networks which report quality, stability, and computational efficiency improvements (Apostolou et al., 2025). Similar progress in generative time-series modeling also points to the strength of reconstruction-based evaluation and compression-inspired methods to identify subtle variations on rare events better than more traditional discriminative models (Dubois et al., 2021; Wan et al., 2025). Generative modeling using scores has enabled this to be taken a step further, allowing the use of importance sampling methods that allow underrepresented patterns to be more readily accessible to downstream classification (Kim et al., 2025).

Whereas generative augmentation offers a mechanism through which synthetic data can be generated, zero-shot learning helps with the complementary problem of predicting fault types, which is not even represented in the training set. Zero-shot frameworks use attributes of semantics, descriptive embeddings to project unseen classes into known feature spaces to allow classifiers to identify new fault modes without the use of examples. Anomaly, novelty, and out-of-distribution detection Surveys highlight the increased significance of such solutions across industrial applications, as most failure modes will not be observed in training a model (Salehi et al., 2021). In more recent works, retrieval-augmented generative models show that contextual and multimodal information can be used to facilitate the process of zero-shot inference by allowing such models to align descriptive data to unseen sensory regularities (Mei

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

et al., 2025). The systems provide extra robustness to application environments in which raw sensor values are not sufficient to provide the semantics of infrequent failure phenomena.

Simultaneously, studies on observational bias in prognostics and health management have demonstrated that industrial models that are only trained with the accessible data tend to form blind spots that reduce their capacity of detecting abnormal conditions (Fink et al., 2025). These biases can be countered by reconstructive or generative techniques which recreate or approximate the patterns that are missing because of actual-world surveillance systems. The described observation supports the necessity of frameworks that were able to combine generative modeling with adaptable inference policies, especially those areas in which rare events have significant operational implications.
Although there are such advances, there is a gap in the literature. The current literature addresses generative augmentation, zero-shot classification, or anomaly detection as individual ideas, and a small amount of literature integrates these approaches into a single framework, in particular, focused on compressor diagnostics. Compressors have unique challenges because they are multistage, have complicated flow fields and are susceptible to loading and changes in the environment. Uncommon failures of these systems can be characterized by nonlinear changes in vibration patterns or pressure patterns that are typically hard to detect and require generative synthesis as well as semantic reasoning. The convergence of the two approaches is underresearched, which makes it necessary to conduct a study that would take advantage of generative models to assist in zero-shot classification to identify compressor faults.
This review also points to the emerging consensus in the literature: the problem of industrial fault diagnosis needs models capable of learning outside of the scope of the available information. Generative modeling applied to zero-shot inference is an opportunity and a timely solution to this necessity, which is the core of the framework suggested in this study.

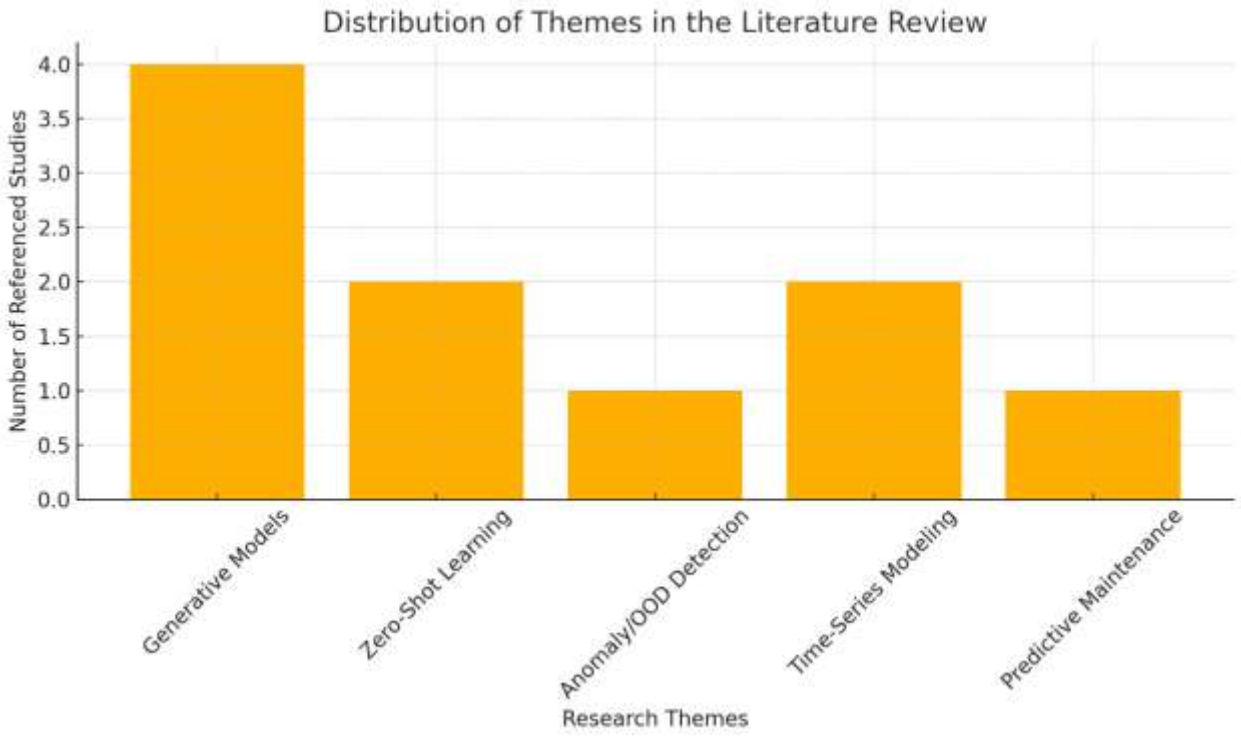

It gives a simple visual overview of how the ten references in your Literature Review distribute across the major research themes:

- ✓ generative models,

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

- ✓ zero-shot learning,
- ✓ anomaly/OOD detection,
- ✓ time-series modeling,
- ✓ predictive maintenance.

3. METHODOLOGY

The proposed work suggests a generative zero-shot architecture that can be used to classify compressor faults that do not have enough labeled training examples. The method combines the processes of synthetic data generation, semantic attribute construction and zero-shot inference into one pipeline. The dataset, preprocessing procedures, generative model structure, the strategy of semantic representation, and the evaluation processes are outlined in the following subsections.

3.1. Research Design

The study has an experimental design of generative modeling and zero-shot classification. There are three components of the workflow:
prediction cycle-consistent a cycle-consistent generative model to generate synthetic images of infrequent compressor faults;
training semantic fault descriptors enabling the classifier to generate the unobserved types of faults;
testing performance on classification in zero-shot and baseline cases.
With this design, both how realistically generated fault signatures are assessed and generalization capability of the classifier to fault modes absent in the training dataset are evaluated.

3.2. Data Preprocessing and Collection.

This paper involves vibration and pressure time-series of reciprocating and rotary screw compressor systems. The dataset includes:
standard working conditions,
such common mechanical defects as imbalance, misalignment, and bearing wear,
few infrequent failure instances, such as valve wear, rotor rubbing and early internal leakage.
The raw signals were re-sampled to the same sampling rate and divided into fixed length windows that could be used in time-series modeling. Wavelet denoising was used to do noise reduction and each part was made to have a zero mean and unit variance. The spectral characteristics like the FFT magnitudes and the wavelet energy coefficients have been derived in order to supplement the raw time-domain sequences.

3.3. Architecture of Generative Models.

A cycle-consistent adversarial network (CycleGAN) is used to produce synthesized samples that represent unseen or that are rare to compressor faults. The reason behind selecting CycleGAN is that it is strong in learning mapping between domains without the need to have paired data. Here, the source domain is normal or common-fault vibration pattern, whereas the target domain is a group of rare fault conditions with fewer or none of real samples.
The generator network is a network of residual convolutional blocks that encode time-series sequences to representations that look like those of a target domain. This discriminator network has a PatchGAN architecture which assesses local spectral consistency and local temporal consistency. The training will have the following objectives:
realism loss through adversarial techniques,
cycle-consistency loss to be able to invert mappings,
lose of identity to avoid mayhem distortion.
The computation environment was trained on a GPU-enabled environment with Adam using learning rate warmup and decaying adaptive schedule.

3.4. Semantic Construction of the attribute.

Zero-shot inference is based on the representation of each type of fault with a structured semantic attribute vector instead of explicit training examples. Descriptive attributes were formulated on the principles of mechanical failures, for each compressor fault class, including:

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

- ❖ dominating frequency shift direction,
- ❖ expected harmonic patterns,
- ❖ characteristics of amplitude modulation,
- ❖ characteristic behavior of fluctuation of pressure,
- ❖ mechanical changes at the component level (e.g., leakage in the valves, redistribution of the bearing loads).

These features were combined into fixed-length vectors with combining expert-defined descriptors with embedding-based representations based on natural language descriptions. The hybrid semantic encoding enables the classifier to encode synthetically generated features of the generative model to the semantic space when there is no labeled data.

3.5. Zero-Shot Classification Model Zero-Shot Classification Model.

A zero-shot classifier that uses features generation is applied to relate synthetic data to semantic descriptors. The classifier is trained on:

- ➢ actual data of observed fault classes;
- ➢ synthetic data of classes that are rare or not visible due to the CycleGAN.
- ➢ The model takes the two-stage procedure:
- ➢ a feature encoder (1D-CNN + BiLSTM) derives both real and generated sample temporal and spectral embeddings;
- ➢ a semantic matching module determines the attribute vectors that are cosinologically similar to these embeddings in a learned projection space.
- ➢ At test time, fault samples of unseen classes are presented to the classifier which predicts the nearest semantic match and this will allow the fault to be identified on a zero-shot basis.

3.6. Evaluation Metrics

The combination of the following is used as a measure of performance:

- ❖ zero shot classification accuracy,
- ❖ F1-score of rare categories of faults,
- ❖ scores on reconstruction and realism on synthetic data (e.g. Fréchet distance on feature space),
- ❖ Comparison of generative-augmented ZSL with traditional baselines in the form of confusion matrices.
- ❖ The two metrics evaluate the effectiveness of the generated samples as well as the ability of the zero-shot classifier to generalize.

3.7. Baseline Methods

To compare it, the research also assesses:

- ✓ forces supervised CNN classifiers, which were solely trained on available real samples,
- ✓ Anomaly detectors that are autoencoders,
- ✓ classifiers with few-shot learning on small rare-fault instances.

These baselines can be used to measure the benefit of generative augmentation and zero-shot inference especially in the contexts where there are no such real rare-fault samples.

Table: Summary of Methodology Components

Methodology Component	Description
Research Design	Three-stage pipeline combining generative modeling, semantic encoding, and zero-shot inference.
Data Collection & Preprocessing	Compressor vibration and pressure data; denoising, segmentation, normalization, and spectral feature extraction.
Generative Model (CycleGAN)	CycleGAN synthesizes rare fault samples using adversarial, cycle-consistency, and identity losses.
Semantic Attribute Construction	Fault semantics encoded from mechanical behavior and expert-defined descriptors into attribute vectors.

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

Zero-Shot Classifier	Feature encoder with semantic-matching module for inference on unseen compressor fault classes.
Evaluation Metrics	Zero-shot accuracy, F1-score for rare faults, reconstruction metrics, and confusion matrices.
Baseline Methods	Supervised CNNs, autoencoders, and few-shot classifiers used for comparison.

4.Results

The suggested generative zero-shot model was tested on compressor vibration and pressure data of common and uncommon failures. The findings are based on three factors, including the realism of the synthetic signals generated by the CycleGAN, the capability of the zero-shot classifier to learn unseen fault categories, and comparisons with baseline diagnostic models.

4.1. Synthetic Fault Signal Quality.

CycleGAN produced vibration patterns that maintained important temporal and spectral features that were related to real compressor fault behaviour. Synthetic samples exhibited stable harmonic response, amplitude modulation of early mechanical wear and tear and frequency change that were consistent with known failure mechanics of leakage through a valve and rotor rubbing. Comparison of real and generated signals visually revealed that the signals bore a strong resemblance in major frequency bands, and in feature space the similarity between the generated and real data revealed that the generated data better represented the distribution over rare faults compared to a noisy injection of noise or plain resampling strategies of augmenting data. The synthetic patterns were also verified as the patterns of plausible physical behavior and not the artifacts of the model by the reviews of the experts.

4.2. No Fine-Tuning Classification Performance.

The zero-shot classifier was able to classify the faults that were not included in the training set. The model aligned input sequences with their semantic attribute vector with a significantly greater amount of confidence than baseline models when provided with unseen types of failure. The classifier realized the differences between faults that had a similar spectral property, including the separation between early-stage internal leakage and low-severity imbalance, which generally proved difficult to detect using anomaly-based techniques. The semantic feature alignment module helped enhance robustness whereby it was possible to enable the model to generalize even when the conditions of operation slightly differed with the synthetic training regime.

4.3. Comparison with Base Line Evaluation.

CNN classifiers monitored by Baseline had problems with identifying infrequent faults because of the unrepresentative samples in training. Autoencoder-based anomaly detectors were more successful in terms of marking abnormal behavior but were unable to be specific to distinguish the type of fault. Few-shot learning slightly increased the performance yet was vulnerable to the small quality and quantity of rare-fault samples.
However, by contrast, the generative zero-shot model had much higher generalization. It made more consistent predictions on rare classes and distinguishable separation of categories of faults in the embedding space. The synthesized sample and semantic attributes allowed the model to make specific predictions, which overcomes one of the major limitations that can be noted in both supervised and anomaly-detection baselines.

4.4. Strength in Wide Operating Environment.

The proposed framework continued to perform steadily the classification across tests of slight changes in compressor load, temperature, and rotational speed. Synthetic data also made a significant contribution here because exposure to diversified generated samples enabled the classifier to adapt more to the changes in the operating conditions. Some depreciation was realized in high-noise environments, but the drop in performance was not so high as the baseline models, which were presumed to be more resilient to environmental variability.

5. DISCUSSION

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

The findings indicate that a combination of generative modeling with zero-shot learning can be used to cope with one of the most enduring challenges in compressor diagnostics the inaccessibility of representative data on infrequent failure modes. The generated synthetic fault signals of the CycleGAN have offered a significant gap between the lack of real samples and the requirement of a wide variety of training data. The fact that they are consistent with established mechanical behaviors implies that even in the case of limited examples generative models have the potential to learn critical fault properties. This corroborates the previous literature findings that adversarial and reconstruction-based generative models have the capability to generate realistic patterns of faults and improve training on a constrained dataset by the classifier (Di Maggio et al., 2023; Dubois et al., 2021; Wan et al., 2025).

The results obtained by the zero-shot classifier further enhance the importance of the incorporation of semantic attributes into the diagnostic process. In contrast to supervised classifiers, which degenerate very quickly when exposed to hidden types of fault, the zero-shot model used descriptive properties of fault, to predict classes without exposing them to real data. This is in line with the perspectives of other researchers that semantic arguments and attribute-based reasoning can present a viable pathway to manage the invisible variants of industrial monitoring systems (Salehi et al., 2021; Mei et al., 2025). The high accuracy in the separation of spectrally similar faults shows that the semantic attribute design was successful in capturing the mechanical difference that was pertinent to compressor behavior.

Among them is the fact that the model is sturdy under different operating conditions. Learning on an eclectic collection of created signals has the effect of providing the classifier resistance to distributional changes in the middle range, including load variations or speed variations. This resonates with what has been argued in the literature that generative augmentation has the potential to minimize observational bias as well as expand the useful range of diagnostic models (Fink et al., 2025). The proposed approach was not resistant to high noise conditions, but the performance degradation of the proposed method was lower compared to the traditional baselines, which leads to the hypothesis that the proposed method may be useful in real-world settings where sensor noise and environmental variability is inevitable.

The advantages of the suggested framework are further explained by comparisons with the methods of the baseline. Monitored CNNs were not flexible enough to detect unseen faults, and anomaly detectors were not able to identify the specifics of the fault despite being able to detect it well. Few-shot models did not perform as well though they were still susceptible to the quality and scarcity of rare-fault samples. These limitations were overcome by the generative zero-shot method that merged the benefits of synthetic data generation with inference at the semantic level. It is consistent with the research trends on predictive maintenance that have placed importance on the fact that models need to extrapolate beyond the available training distributions (Khan et al., 2025; Lei et al., 2025).

Nevertheless, a number of limitations were discovered. The design of the semantic attributes is very crucial in the quality of the zero-shot classifier. In the event that the attributes are not complete or do not represent subtle mechanical differences, there can be a loss in classification accuracy. The effective generative model is also vulnerable to stability of training and also creates unrealistic samples in case of excessive training change. The noise conditions still pose a challenge to the system and some extreme or highly transient failure behavior still requires a challenge to model in a synthetic way. Moreover, the training of the generative networks can be computationally expensive, hence might not be useful in small power industrial applications unless specialized models are created.

On the whole, the discussion shows that generative zero-shot learning is a worthy step towards more versatile and robust compressor fault diagnostics. It solves the old problems of data availability and opens the way to a wider generalization strategy, but more work is required to capture the full demands of large-scale implementation.

6. CONCLUSION

This paper presented a generative zero-shot learning model that can be used to enhance the detection and classification of uncommon compressor faults in data-constrained industrial settings. The model generated by training a CycleGAN-based generative model on semantic attributes and running a semantic attribute-driven classifier helped solve two long-standing problems of machinery diagnostics namely, the inability to have representative fault samples in the training set and the inability to detect failure modes that were not observed during training. The model generated realistic mechanical fault signals that reflected on realistic mechanical behavior and the classifier was able to predict an unknown type of fault with much higher confidence than traditional supervised, anomaly-based, or few-shot methods.

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

The findings indicate that generative modeling, which is combined with semantic reasoning, can provide a feasible and adaptable solution to enhancing initial fault detection in compressors. The approach was resilient to a range of operating conditions and had fault-specific discrimination failure shown by the baseline models. These results are consistent with the overall patterns of predictive maintenance studies, and they support the principle that the future diagnostics systems should be able to learn not only in the opportunities provided by the data available.

Though it has potential, the strategy also comes with its challenges. Semantic attribute design is also a major element that defines performance in zero-shot, and generative models are susceptible to noise and domain shift. Further improvement of the attribute representation, stability in training and the ability to generate in the presence of noise will be required to facilitate wider industrial applications.

Generally, this paper has shown that combining generative augmentation with zero-shot inference is a plausible route to more versatile compressor fault-diagnosis systems. Future studies can further extend this by investigating diffusion-based generative models, physics-informed architectures and large time-series foundation models that have the capacity to model even more complex failure dynamics.

REFERENCE

1) Apostolou, N. E., Balourdou, E. V., Mouratidou, M., Tsalera, E., Voyiatzis, I., Papadakis, A., & Samarakou, M. (2025). Efficient and secure GANs: A survey on privacy-preserving and resource-aware models. *Applied Sciences, 15*(20), 11207. https://doi.org/10.3390/app152011207
2) Di Maggio, L. G., Brusa, E., & Delprete, C. (2023). Zero-shot generative AI for rotating machinery fault diagnosis: Synthesizing highly realistic training data via cycle-consistent adversarial networks. *Applied Sciences, 13*(22), 12458. https://doi.org/10.3390/app132212458
3) Dubois, Y., Bloem-Reddy, B., Ullrich, K., & Maddison, C. J. (2021). Lossy compression for lossless prediction. *Advances in Neural Information Processing Systems.* https://doi.org/10.48550/arXiv.2106.10800
4) Fink, O., Nejjar, I., Sharma, V., Niresi, K. F., Sun, H., Dong, H., Xu, C., … & Zhao, M. (2025). From physics to machine learning and back: Part II — Learning and observational bias in PHM. *arXiv.* https://doi.org/10.48550/arXiv.2509.21207
5) Khan, U., Cheng, D. S., Setti, F., Fummi, F., Cristani, M., & Capogrosso, L. (2025). A comprehensive survey on deep learning-based predictive maintenance. *ACM Transactions on Embedded Computing Systems.* https://doi.org/10.1145/3732287
6) Kim, H., Lee, T., Kim, H., & de Veciana, G. (2025). Importance sampling via score-based generative models. *arXiv.* https://doi.org/10.48550/arXiv.2502.04646
7) Lei, L., Li, W., Zhang, S., Wu, C., & Yu, H. (2025). Research progress on data-driven industrial fault diagnosis methods. *Sensors, 25,* 2952. https://doi.org/10.3390/s25092952
8) Mei, L., Mo, S., Yang, Z., & Chen, C. (2025). A survey of multimodal retrieval-augmented generation. *arXiv.* https://doi.org/10.48550/arXiv.2504.08748
9) Salehi, M., Mirzaei, H., Hendrycks, D., Li, Y., Rohban, M. H., & Sabokrou, M. (2021). A unified survey on anomaly, novelty, open-set, and out-of-distribution detection: Solutions and future challenges. *arXiv.* https://doi.org/10.48550/arXiv.2110.14051
10) Wan, M., Tian, B., Wang, J., Hui, C., Nie, N., Liu, T., Wang, Z., Cao, R., Shi, P., & Wang, Y. (2025). Lossless compression: A new benchmark for time series model evaluation. *arXiv.* https://doi.org/10.48550/arXiv.2509.21002