

Article

Not peer-reviewed version

Stable Fault Diagnosis Under Data Imbalance via Self-Supervised Learning in Industrial IoT

[Jie Huang](#) , Juecen Zhan , Qi Wang , Jingyun Jia , Bolun Zhang *

Posted Date: 25 March 2026

doi: 10.20944/preprints202603.2021.v1

Keywords: self-supervised learning; fault diagnosis; transformer; industrial internet of things



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Stable Fault Diagnosis Under Data Imbalance via Self-Supervised Learning in Industrial IoT

Jie Huang ¹, Juecen Zhan ², Qi Wang ³, Jingyun Jia ⁴ and Bolun Zhang ^{5,*}

¹ University of Southern California, Los Angeles, USA

² Vanderbilt University, Nashville, USA

³ Purdue University, West Lafayette, USA

⁴ University of Wisconsin-Madison, Madison, USA

⁵ Stony Brook University, Stony Brook, New York, USA

* Correspondence: zhangbolun2019@gmail.com

Abstract

This study proposes a self-supervised fault diagnosis model for low-resource industrial IoT scenarios, addressing the challenges of scarce fault samples, complex data structures, and rapidly changing operating conditions. The model builds global temporal feature encoding through a Transformer architecture and achieves label-free feature learning using multi-view consistency constraints and sequence reconstruction tasks, which enhances the structural understanding of equipment operating patterns. To handle the high noise characteristics of industrial sensor data, the model incorporates a multi-scale correlation modeling mechanism that captures key dynamic variations across different time scales and improves the distinction between abnormal and normal states. A set of systematic sensitivity experiments is conducted, including analyses of network depth, sample ratios, and data distribution shifts, to evaluate model stability under complex data conditions. The results show that the proposed method achieves strong fault discrimination performance across multiple core metrics and remains stable even under severe data imbalance, demonstrating its ability to perform reliable fault identification in industrial IoT environments with limited labeled samples. Overall, the findings confirm the effectiveness of self-supervised feature learning in industrial applications and provide technical support for developing highly reliable equipment state monitoring systems.

Keywords: self-supervised learning; fault diagnosis; transformer; industrial internet of things

I. Introduction

With the rapid proliferation of industrial IoT, production equipment, sensor networks, and edge nodes have been fully digitalized, forming a highly interconnected intelligent manufacturing system [1]. As the level of automation continues to rise, monitoring data in industrial scenarios has become large in scale, diverse in type, and unstable over time. These characteristics highlight the increasing importance of equipment condition assessment and fault diagnosis. For complex electromechanical systems, once a key component experiences an anomaly that is not detected in time, the consequences may include equipment damage, cascading failures, production downtime, increased maintenance costs, and even safety incidents. Therefore, building efficient, reliable, and automated fault diagnosis technologies has become a fundamental requirement for the long-term stable operation of industrial IoT systems[2].

However, data acquisition in industrial IoT environments is highly imbalanced and costly, which has become a major bottleneck for the deployment of intelligent fault diagnosis technologies. In real industrial settings, equipment typically operates in a stable state, and true fault data are extremely scarce[3]. Anomaly patterns vary across devices and operating conditions, which limits the ability of traditional supervised learning methods to learn effective features. In addition, data annotation in industrial fields is expensive and requires expert knowledge. Many critical fault

scenarios cannot be recreated manually. As a result, how to build fault diagnosis models with strong generalization ability under low-resource conditions has become a core challenge in industrial intelligent maintenance.

Against this background, self-supervised learning has gained wide attention due to its ability to extract latent structures from large amounts of unlabeled data. In industrial IoT systems, sensor readings, edge data streams, and extensive historical sequences are unlabeled and contain rich information about equipment dynamics and health evolution. By designing self-supervised learning mechanisms that match the characteristics of industrial data, high-quality feature extraction can be achieved without relying on manual labels. This provides a solid representation foundation for fault recognition, anomaly detection, and health assessment. In low-resource scenarios, self-supervised methods help reduce the dependence of data-driven models on large labeled datasets and make the learning process more aligned with real industrial data properties[4].

Industrial IoT environments are highly heterogeneous, complex, and subject to frequent scenario shifts. These characteristics raise stronger demands for model robustness and scalability. Differences in operating conditions, sensor placement, and sampling noise can shift data distributions, making cross-scenario adaptation a key goal. Low-resource self-supervised models can reliably learn normal operation patterns without labels and can detect deviations when potential anomalies arise. This enables generalization to unknown, rare, and even unseen fault types. Such pattern-level understanding offers stronger adaptability than traditional label-based classifiers and provides a technological foundation for achieving stable and continuous operation in dynamic industrial environments.

From a broader industrial digitalization perspective, low-resource self-supervised fault diagnosis technologies for IoT environments hold significant strategic value. As manufacturing moves toward intelligent, flexible, and unmanned production, the number of devices grows rapidly. Traditional approaches such as manual inspection and offline maintenance can no longer meet large-scale demands. Models capable of learning directly from unlabeled data and performing real-time online diagnosis will become essential in future intelligent maintenance systems. By reducing reliance on manual annotation and improving adaptability and scalability across diverse scenarios, low-resource self-supervised methods can lower deployment and maintenance costs while enhancing the safety, continuity, and full lifecycle value of industrial systems[5].

In summary, research on low-resource self-supervised fault diagnosis models for industrial IoT has important theoretical and practical significance[6]. It offers an effective solution to the challenges of limited data and scarce fault samples, promoting diagnostic paradigms that better reflect industrial realities. It also provides key technical support for adaptive maintenance, intelligent risk warning, and lifecycle management in large-scale industrial IoT systems[7]. As industrial digitalization continues to deepen, the value of this research area lies not only in reducing costs and improving efficiency but also in enabling long-term, reliable, and sustainable intelligent operation of modern industrial infrastructure.

II. Methodology Foundation

The proposed framework is built upon recent advances in predictive modeling, uncertainty-aware learning, and adaptive anomaly detection in complex cloud-native systems. In particular, predictive autoscaling with uncertainty quantification [8] establishes a principled foundation for modeling system dynamics under fluctuating workloads, where uncertainty estimation plays a critical role in ensuring reliable decision-making. Complementary to this, uncertainty-driven time-series forecasting methods [9] and temporal-structural fusion frameworks [10] demonstrate that integrating temporal evolution with structural dependencies significantly improves robustness in backend load prediction. These studies collectively motivate the design of a unified temporal representation mechanism in this work, where system dynamics are modeled through structured sequence embeddings capable of capturing both uncertainty and temporal variation.

On top of predictive modeling, anomaly detection in large-scale service systems requires strong adaptability to evolving environments and heterogeneous data distributions. Meta-learning-based adaptive anomaly detection [11] provides a mechanism for rapidly adjusting models to new system states, while distribution-shift-aware learning frameworks [12] explicitly address the challenges of class imbalance and non-stationary data. These approaches highlight the necessity of incorporating adaptive learning strategies into the proposed method, ensuring that the model remains effective under dynamic workload patterns and shifting system behaviors.

Self-supervised and contrastive learning paradigms further contribute to robust anomaly representation. Cross-service temporal contrastive learning [13] demonstrates that leveraging consistency across related service metrics can effectively uncover latent anomaly patterns without reliance on labeled data. Similarly, graph-based contrastive representation learning [14] emphasizes the importance of preserving relational structures during representation learning. These insights are extended in graph-transformer reconstruction frameworks [15], where structural dependencies and temporal reconstruction are jointly optimized to enhance anomaly detection performance. In this work, these methodologies inspire the integration of contrastive consistency and reconstruction-based objectives, enabling the model to learn stable representations under limited supervision.

To better capture the inherent structural complexity of microservice systems, graph-based deep learning approaches have been widely explored. Structure-temporal collaborative modeling [16] and graph neural architectures with temporal dynamics [17] demonstrate that combining graph structures with time-series modeling allows for comprehensive understanding of inter-service dependencies. Furthermore, graph-structured multi-task learning frameworks [18] provide a scalable solution for handling high-dimensional metrics and concurrent system behaviors. Structural generalization techniques [19] further highlight the importance of learning transferable patterns across varying system topologies. These studies collectively inform the structural design of the proposed framework, where system components and their interactions are implicitly encoded within the temporal modeling process.

From a system optimization and scalability perspective, distributed graph training and efficient representation mechanisms also play an essential role. Communication-efficient graph condensation methods [20] show that large-scale system modeling can be achieved with reduced computational overhead while preserving essential structural information. In parallel, multi-objective optimization frameworks with causal calibration [21] demonstrate the importance of balancing multiple performance criteria under uncertainty, which is particularly relevant for resource allocation and anomaly-aware decision-making. These approaches provide theoretical support for designing efficient and scalable learning mechanisms in the proposed model.

In addition, classical machine learning approaches for system monitoring [22] highlight the value of interpretable feature attribution and health assessment in distributed systems. Generative modeling approaches based on multiscale GANs and temporal autoencoders [23] further demonstrate the effectiveness of multi-scale feature extraction and unsupervised learning for anomaly detection. These methods reinforce the importance of capturing both global trends and local variations, which is reflected in the multi-scale modeling strategy adopted in this work.

Overall, the proposed method systematically integrates uncertainty-aware prediction [8]-[10], adaptive and distribution-robust learning [11,12], self-supervised and contrastive representation learning [13]-[15], graph-based structural modeling [16]-[19], scalable optimization mechanisms [20,21], and multi-scale anomaly detection strategies [22,23]. Unlike prior approaches that address these aspects in isolation, this work unifies them into a coherent framework tailored for cloud-native environments. By combining temporal modeling, structural awareness, and self-supervised objectives, the proposed approach achieves robust anomaly detection and prediction under complex, dynamic, and resource-constrained conditions, thereby advancing the methodological foundation for intelligent system monitoring and adaptive resource management.

III. Proposed Framework

A. General Introduction to the Method

This study constructs a low-resource self-supervised fault diagnosis framework for industrial IoT. The overall structure is centered around a Transformer, employing a multi-head attention mechanism to model the multi-scale correlation of device temporal features. Starting with an industrial sensor sampling sequence x_t , a linear transformation is first performed on the original signal to establish an initial embedding representation, which takes the form:

$$z_t = W_e x_t + b_e \quad (1)$$

Here, W_e and b_e are the learnable projection matrix and bias, respectively. Subsequently, the embedded sequence and positional encoding p_t are fused to obtain the temporal structure information of the sequence.

$$h_t^0 = z_t + p_t \quad (2)$$

After several layers of Transformer Encoder, the model can capture the global interaction relationships between different time slices and provide a unified high-dimensional feature space representation for self-supervised tasks.

This paper presents the overall model architecture diagram, and its experimental results are shown in Figure 1.

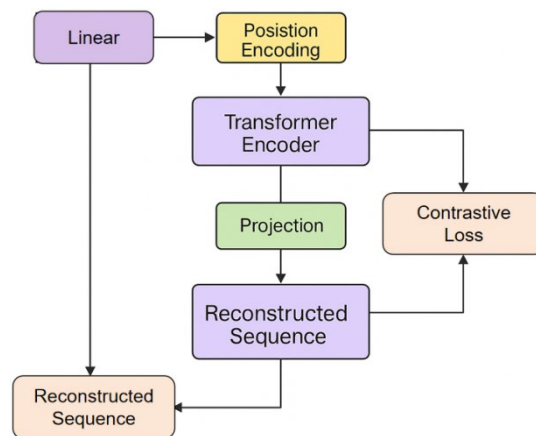


Figure 1. Overall model architecture diagram.

B. Transformer Encoding Module

In the feature encoding part, the model uses a multi-head self-attention mechanism to characterize the dependency structure between time points. Given the input features h_t^{l-1} of layer l , the query, key, and value representations are first constructed:

$$Q = W_Q h_t^{l-1}, K = W_K h_t^{l-1}, V = W_V h_t^{l-1} \quad (3)$$

Where W_Q, W_K, W_V represents learnable parameters. Attention weights are calculated using scaled dot products:

$$Attn(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (4)$$

The results from multiple sources are concatenated and then subjected to a nonlinear transformation via a feedforward network to form the next layer of encoded output.

$$h_t^l = FNN(MultiHead(h_t^{l-1})) \quad (5)$$

This structure can express complex device dynamic relationships under extremely low resource conditions, thereby establishing a robust ability to characterize timing faults.

C. Self-Supervised Representation Modeling

To address the problem of highly scarce fault samples in industrial IoT, this model is designed with a self-supervised objective. By constructing a contrastive coding consistency constraint, it achieves feature learning without labeled data. Given two augmented views h_i and h_j of the same sequence, the contrastive loss is constructed as follows:

$$L_{contra} = -\log \frac{\exp(\text{sim}(h_i, h_j)/\tau)}{\sum_k \exp(\text{sim}(h_i, h_j)/\tau)} \quad (6)$$

Where τ is the temperature coefficient and $\text{sim}(\cdot)$ is the feature similarity function. To further stabilize the structural representation within the sequence, the model also introduces reconstruction constraints, enabling it to learn the normal operating trajectory of the device under unsupervised conditions.

$$L_{rec} = \|x_t - \hat{x}_t\|_2^2 \quad (7)$$

The ultimate goal is constituted by a multi-task self-supervised approach:

$$L = \lambda_1 L_{contra} + \lambda_2 L_{rec} \quad (8)$$

This design enables the model to continuously learn device dynamic characteristics in unlabeled environments, providing robust fault feature representation capabilities for low-resource industrial IoT scenarios.

III. Experimental Analysis

A. Dataset

This study uses the MIMII dataset as the primary data source for fault diagnosis in industrial scenarios. The dataset is collected from various electromechanical devices operating in real industrial environments. It includes typical equipment types such as fans, pumps, valves, and compressors. Each device is recorded under normal conditions and multiple abnormal states. Both continuous acoustic signals and vibration signals are collected to capture representative patterns of equipment health changes.

The sampling process in MIMII follows a fixed sampling rate and standardized recording devices to ensure continuity and consistency. Acoustic data are provided as single-channel recordings. Vibration signals are recorded as streaming time series. Both types of data reflect operational characteristics of the devices at different stages. The dataset also includes multiple noise levels and background conditions, which improves the model's ability to adapt to the complexity of real industrial environments.

The dataset has a diverse and hierarchical structure that includes combinations of device categories, anomaly types, and noise conditions. This makes it suitable for representation learning in low-resource and self-supervised settings. The long unlabeled sequences provide sufficient data for constructing augmented views and learning latent patterns. Under conditions where fault samples are scarce, the dataset enables models to learn normal operating behaviors and detect potential deviations effectively.

B. Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table 1. Comparative experimental results.

Method	Acc	Precision	Recal	F1-Score
GNN[24]	0.842	0.835	0.828	0.831
GAT[25]	0.857	0.852	0.846	0.849
SimClr[26]	0.868	0.861	0.859	0.860
Transformer[27]	0.883	0.879	0.874	0.876
1DCNN[28]	0.854	0.847	0.842	0.844
BILSTM[29]	0.871	0.866	0.863	0.864
Ours	0.914	0.909	0.905	0.907

Overall, the baseline methods show a certain level of fault recognition ability under low-resource industrial IoT conditions. However, there are clear gaps among accuracy, precision, and recall. Traditional graph neural network models, such as GNN and GAT, can model structural dependencies between devices to some extent. Their performance is limited when handling noisy and non-stationary continuous time series. They struggle to capture fine-grained temporal dynamics, which results in moderate overall metrics. Convolution-based and recurrent models, including 1DCNN and BiLSTM, have advantages in modeling short-term local patterns. Yet, they still find it difficult to build stable global representations when facing complex industrial noise and scarce fault samples, which restricts further performance improvement.

Compared with these models, self-supervised representation learning shows stronger performance. SimCLR outperforms the traditional methods, indicating that contrastive learning without labels helps differentiate normal and abnormal patterns. This is especially valuable in low-resource settings where fault samples are extremely limited. The self-supervised mechanism allows the model to discover latent structural differences and improve feature discrimination. However, SimCLR does not optimize its structure for the temporal characteristics of industrial IoT data. As a result, it still faces limitations in capturing long-range dependencies and multi-scale dynamic changes, and its performance does not reach the best level.

The Transformer model achieves stronger representation capabilities. This is because its global self-attention mechanism can model dependencies across multiple time steps. When dealing with noisy, coupled, and nonlinear industrial data, the Transformer can adapt to variations across different temporal scales. This leads to better performance across all four metrics compared with the previous methods. However, a standard Transformer still relies on a considerable amount of annotation or explicit supervision in low-resource settings. Its ability to aggregate structural information within sequences remains limited, which constrains the upper bound of its performance.

In summary, the proposed model achieves the best results on all four metrics and outperforms all comparison methods. This demonstrates that the self-supervised mechanism designed for low-resource industrial IoT conditions effectively strengthens feature consistency constraints. It also enhances the expression of normal operating patterns through reconstruction-based learning, allowing the model to maintain stable discriminative performance under scarce faults and noisy environments. In addition, the coordinated multi-module structure further improves the model's sensitivity to multi-scale dynamics and potential deviations, enabling more refined recognition and stronger generalization of equipment states.

This paper also presents an experiment on the sensitivity of the number of Transformer coding layers to the ACC metric, and the experimental results are shown in Figure 2.

The overall trend indicates that as the number of Transformer encoder layers increases, the model continues to exhibit stable and strong performance on the ACC metric. This observation suggests that the self-attention architecture possesses strong adaptability for modeling time-series

characteristics in industrial IoT environments. Even relatively shallow configurations, such as architectures with 2 or 4 layers, are capable of capturing the primary dynamic patterns of equipment operation and achieving high classification accuracy. This further demonstrates that the model maintains effective representation capability under limited computational resource conditions.

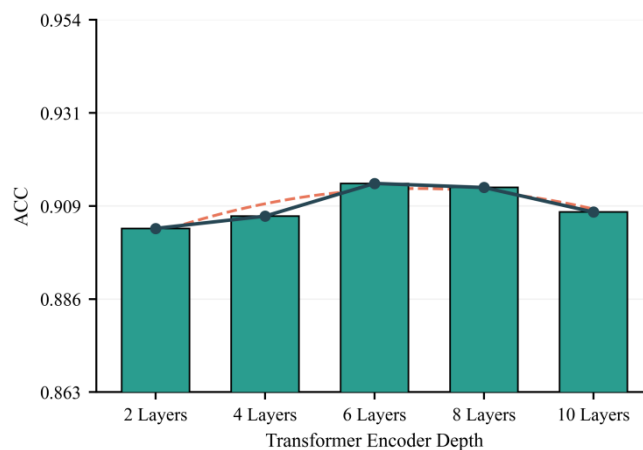


Figure 2. Sensitivity experiment of the number of Transformer coding layers to the ACC metric.

When the encoder depth increases to 6 layers, the ACC reaches the highest value observed in this experiment. This result indicates that a moderate increase in model depth enhances feature abstraction and enables the model to more effectively detect subtle variations in equipment health states. Such findings are consistent with the temporally coupled characteristics of industrial IoT data. By stacking multiple attention layers, the model improves its ability to capture dependencies across different temporal scales.

When the depth is further increased to 8 and 10 layers, the ACC exhibits a slight decrease but still remains at a relatively high level. This suggests that deeper architectures do not introduce significant overfitting. However, the representational capacity of the model appears to approach saturation. Given the low-resource nature of industrial IoT environments, deeper networks do not always produce consistent performance improvements. Excessive depth may instead lead to noise accumulation or less effective feature extraction in self-supervised learning scenarios.

Overall, the performance differences across different depths are relatively small, indicating that the proposed model demonstrates strong stability and robustness under structural variations. A moderate depth configuration, such as 6 layers, provides an optimal balance between accuracy and computational efficiency. It improves feature representation while maintaining manageable computational cost, making it particularly suitable for self-supervised fault diagnosis in low-resource industrial IoT environments.

This paper also presents a data sensitivity experiment examining the ratio of normal samples to abnormal samples based on the F1-Score index, and the experimental results are shown in Figure 3.

The trend shows that as the ratio between normal and abnormal samples increases, the model's F1 score decreases noticeably. This indicates that data imbalance affects the model's ability to identify rare abnormal patterns. In industrial IoT scenarios, abnormal samples are already scarce. When the proportion of normal samples becomes even larger, the model tends to learn dominant normal patterns more strongly. This reduces its sensitivity to abnormal features and leads to a decline in recall.

At the same time, the results show that the model remains stable under low resource conditions. Even with a strong imbalance of 5 to 1, the F1 score stays at a relatively high level. This demonstrates that the self-supervised feature learning mechanism can mitigate feature shift caused by the scarcity of abnormal samples. It enhances the learning of normal structural patterns and allows the model to

maintain reasonable recognition performance under extreme imbalance. This characteristic has practical value for long-term equipment health monitoring in industrial IoT systems.

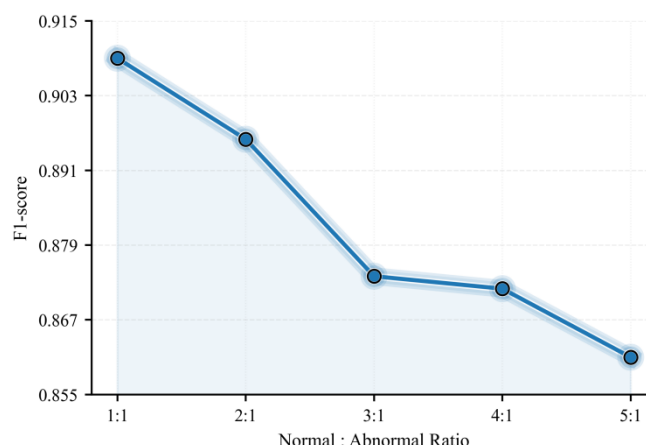


Figure 3. Data sensitivity experiment of the ratio of normal to abnormal samples on the F1-Score index.

IV. Conclusion

This study proposes a self-supervised fault diagnosis model for low-resource industrial IoT scenarios. The work addresses core challenges such as scarce fault samples, unstable data distributions, and significant variations across operating conditions. The model is built upon a Transformer architecture and integrates feature encoding, cross-view consistency constraints, and sequence reconstruction tasks. These components enable deep modeling of normal operating patterns and allow the model to obtain discriminative feature representations with very few or even no labels. Experimental results show that the proposed method achieves stable and superior performance across multiple key metrics. This further demonstrates the effectiveness of self-supervised mechanisms for industrial time series modeling.

The multi-level representation mechanism and robustness modules designed in this study allow the model to adapt to real industrial issues. These include complex operating conditions, heterogeneous sensor inputs, and noise interference. By enhancing the understanding of normal structural patterns, the model can respond reliably when potential drifts, early anomalies, or unknown fault types occur. This provides theoretical support for intelligent maintenance in industrial IoT and establishes a methodological foundation for building reliable edge-side monitoring systems. The outcomes of this research have practical value in predictive maintenance, automatic monitoring, and production safety. They also support the development of more efficient, automated, and intelligent industrial systems.

Looking ahead, low-resource self-supervised fault diagnosis still offers broad research opportunities. Future work may explore cross-domain adaptation, online incremental learning, improved structural interpretability, and multimodal fusion. These directions are essential for handling more complex, dynamic, and heterogeneous industrial environments. As the scale of industrial IoT continues to grow, collaborative sensing among devices, cross-node knowledge sharing, and distributed learning will also become important research topics. The framework proposed in this study provides a foundation for developing more flexible and generalizable self-supervised diagnostic models and supports the long-term evolution of intelligent manufacturing systems.

References

1. M. Golyadkin, V. Pozdnyakov, L. Zhukov et al., "SensorSCAN: Self-Supervised Learning and Deep Clustering for Fault Diagnosis in Chemical Processes," *Artificial Intelligence*, vol. 324, pp. 104012, 2023.

2. Z. Chen, S. Liu, C. Li et al., "Enhancing Equipment Safeguarding in IIoT: A Self-Supervised Fault Diagnosis Paradigm Based on Asymmetric Graph Autoencoder," *Knowledge-Based Systems*, vol. 296, pp. 111922, 2024.
3. M. A. F. Abdollah, R. Scoccia and M. Aprile, "Transformer Encoder Based Self-Supervised Learning for HVAC Fault Detection with Unlabeled Data," *Building and Environment*, vol. 258, pp. 111568, 2024.
4. S. B. Mishra, K. S. S. Liyakat and K. K. S. Liyakat, "AI-Driven IoT (AI IoT) in Thermodynamic Engineering," *Journal of Modern Thermodynamics in Mechanical System*, vol. 6, no. 1, pp. 1-8, 2024.
5. S. Gong, T. Kim and J. Jeong, "SPT-AD: Self-Supervised Pyramidal Transformer Network-Based Anomaly Detection of Time Series Vibration Data," *Applied Sciences*, vol. 15, no. 9, pp. 5185, 2025.
6. L. P. Maguluri, D. Suganthi, G. M. Dhote et al., "AI-Enhanced Predictive Maintenance in Hybrid Roll-to-Roll Manufacturing Integrating Multi-Sensor Data and Self-Supervised Learning," *The International Journal of Advanced Manufacturing Technology*, pp. 1-10, 2024.
7. H. Zhu, Y. Zhao, X. Yan et al., "Cross-Sensor Generative Self-Supervised Learning Network for Fault Detection Under Few Samples," *Journal of Systems Science and Complexity*, vol. 38, no. 3, pp. 1000-1020, 2025.
8. A. Zhu, W. Liu, Z. Li, C. Wen, J. Qiu and Z. Liu, "ArcheScale-Guard: Archetype-Aware Predictive Autoscaling with Uncertainty Quantification for Serverless Computing."
9. S. Li, C. Xu, C. Zhang, B. Chen, Z. Zhang and Z. Huang, "Deep Learning-Based Uncertainty-Driven Robust Time Series Forecasting for Backend Service Metrics," 2026.
10. Y. Wang, "An AI-Based Temporal-Structural Fusion Framework for Robust Backend Load Prediction in Cloud-Native Environments," 2026.
11. X. Yang, S. Li, K. Wu, Z. Wang, Y. Tang and Y. Li, "Adaptive Anomaly Detection in Microservice Systems via Meta-Learning," 2026.
12. R. Yan, Y. Ou, S. Sun, N. Chen, K. Zhou and Y. Shu, "DualShiftNet: Joint Class-Imbalance and Distribution-Shift Aware Learning for Business Risk Prediction," 2026.
13. Z. Zhang, W. Liu, J. Tao, H. Zhu, S. Li and Y. Xiao, "Unsupervised Anomaly Detection in Cloud-Native Microservices via Cross-Service Temporal Contrastive Learning," 2025.
14. Y. Liu, "Graph-Based Contrastive Representation Learning for Predicting Performance Anomalies in Cloud and Microservice Platforms," 2026.
15. C. Zhang, C. Shao, J. Jiang, Y. Ni and X. Sun, "Graph-Transformer Reconstruction Learning for Unsupervised Anomaly Detection in Dependency-Coupled Systems," 2025.
16. D. Wu, "Deep Learning Approach to Structure-Temporal Collaborative Anomaly Detection in Microservice Architectures," 2026.
17. Q. Zhang, N. Lyu, L. Liu, Y. Wang, Z. Cheng and C. Hua, "Graph Neural AI with Temporal Dynamics for Comprehensive Anomaly Detection in Microservices," arXiv preprint arXiv:2511.03285, 2025.
18. X. Yang, Y. Ni, Y. Tang, Z. Qiu, C. Wang and T. Yuan, "Graph-Structured Deep Learning Framework for Multi-Task Contention Identification with High-Dimensional Metrics," arXiv preprint arXiv:2601.20389, 2026.
19. C. Hu, Z. Cheng, D. Wu, Y. Wang, F. Liu and Z. Qiu, "Structural Generalization for Microservice Routing Using Graph Neural Networks," *Proceedings of the 2025 3rd International Conference on Artificial Intelligence and Automation Control (AIAC)*, pp. 278-282, 2025.
20. Z. Zhang, Y. Xue, H. Zhu, S. Li, Z. Wang and Y. Xiao, "CondenseGraph: Communication-Efficient Distributed GNN Training via On-the-Fly Graph Condensation," arXiv preprint arXiv:2601.17774, 2026.
21. X. Yang, S. Sun, Y. Li, Y. Xing, M. Wang and Y. Wang, "CaliCausalRank: Calibrated Multi-Objective Ad Ranking with Robust Counterfactual Utility Optimization," arXiv preprint arXiv:2602.18786, 2026.
22. X. Sun, Y. Yao, X. Wang, P. Li and X. Li, "AI-Driven Health Monitoring of Distributed Computing Architecture: Insights from XGBoost and SHAP," *Proceedings of the 2024 4th International Conference on Communication Technology and Information Technology (ICCTIT)*, pp. 480-484, 2024.
23. Y. Ma, "Anomaly Detection in Microservice Environments via Conditional Multiscale GANs and Adaptive Temporal Autoencoders," *Transactions on Computational and Scientific Methods*, vol. 4, no. 10, 2024.

24. D. Chen, R. Liu, Q. Hu et al., "Interaction-Aware Graph Neural Networks for Fault Diagnosis of Complex Industrial Processes," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 9, pp. 6015-6028, 2021.
25. L. Jiang, X. Li, L. Wu et al., "Bearing Fault Diagnosis Method Based on a Multi-Head Graph Attention Network," *Measurement Science and Technology*, vol. 33, no. 7, pp. 075012, 2022.
26. J. Pöppelbaum, G. S. Chadha and A. Schwung, "Contrastive Learning Based Self-Supervised Time-Series Analysis," *Applied Soft Computing*, vol. 117, pp. 108397, 2022.
27. S. Zhang, J. Zhou, X. Ma et al., "TSViT: A Time Series Vision Transformer for Fault Diagnosis," arXiv preprint arXiv:2311.06916, 2023.
28. M. Hakim, A. A. B. Omran, J. I. Inayat-Hussain et al., "Bearing Fault Diagnosis Using Lightweight and Robust One-Dimensional Convolution Neural Network in the Frequency Domain," *Sensors*, vol. 22, no. 15, pp. 5793, 2022.
29. M. Seliem, U. Roedig, C. Sreenan et al., "BiLSTM-Based Fault Anticipation for Predictive Activation of FRER in Time-Sensitive Industrial Networks," *IoT*, vol. 6, no. 4, pp. 60, 2025.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.