



Dependency Graph Learning for Anomaly Detection in Dynamic Smart-Manufacturing Sensor Data

James R. Miller¹, Michael L. Davis², David K. Wilson^{3*}

Department of Mechanical and Industrial Engineering,
University of Toronto, Toronto, ON M5S 3G8, Canada

Corresponding Author*: David K. Wilson Email: d.wilson@utoronto.ca

ARTICLE INFO

Keywords

Non-stationary time series; multivariate anomaly detection; smart manufacturing; dynamic dependency graph; sensor data; drift-aware learning; spatiotemporal modeling.

Published

01 June 2026

ABSTRACT

Smart manufacturing systems generate multivariate sensor streams from robotic arms, CNC machines, conveyor units, and environmental controllers. These signals are often non-stationary because of tool wear, production switching, maintenance operations, and workload changes. Conventional anomaly detection methods usually assume stable temporal patterns and may fail when normal operating conditions drift over time. This study proposes an adaptive dependency graph learning model for non-stationary multivariate sensor anomaly detection in smart manufacturing systems. The proposed method first constructs dynamic sensor-dependency graphs using sliding-window correlation and attention-based temporal encoding. A drift-aware normalization module is then used to separate long-term operating changes from short-term abnormal fluctuations. Finally, a graph-temporal reconstruction network estimates normal sensor behavior and detects anomalies through adaptive residual thresholds. Experiments are conducted on a smart manufacturing dataset collected from 42 production lines, 318 industrial sensors, and 11 machine categories over 96 operating days. The dataset contains 74.6 million timestamped sensor records, including normal production cycles, tool degradation, cooling-system instability, abnormal vibration, and actuator faults. The proposed method reduces average detection delay from 18.7 minutes to 6.4 minutes compared with a static LSTM autoencoder. The false alarm rate decreases to 1.72 events per 1,000 operating hours, and the Matthews correlation coefficient reaches 0.917 under mixed production modes. The model processes 12,800 sensor points per second on an edge workstation, with a memory footprint of 1.36 GB during online inference. Ablation analysis shows that removing dynamic dependency graphs increases detection delay by 9.2 minutes, while removing drift-aware normalization causes frequent false alarms after production switching. The results demonstrate that adaptive spatiotemporal dependency modeling can improve anomaly detection under non-stationary industrial sensor conditions.

1. Introduction

Modern smart manufacturing systems increasingly rely on multivariate sensor streams to continuously monitor the operating conditions of robotic arms, computer numerical control

Citation: Miller, J. R., Davis, M. L., & Wilson, D. K. (2026). Dependency graph learning for anomaly detection in dynamic smart-manufacturing sensor data. *The Journal of Applied Engineering and Technologies*, 1(2), 36–44.

3154-7877/© The Authors. Published by J&L Academic Group PLT. This is an open access article under the CC BY 4.0 license. <https://doi.org/10.64744/tjaet.2026.271>.

machines, industrial motors, cooling systems, and automated production equipment (Xu & Zhou, 2025; Tsanousa et al., 2022). These sensors generate high-frequency measurements of vibration, temperature, current, pressure, rotational speed, acoustic signals, and controller states, providing an important data foundation for fault detection and predictive maintenance. As production lines become more automated and interconnected, even a minor abnormality in one machine component may propagate through multiple process stages, resulting in production interruption, product-quality degradation, equipment damage, or unexpected maintenance costs. Reliable anomaly detection is therefore essential for maintaining production continuity, improving equipment availability, and reducing the economic losses caused by unplanned downtime. Recent studies have applied causal inference to multivariate time-series anomaly detection to identify fine-grained abnormal variables and distinguish initiating anomalies from downstream responses caused by inter-variable propagation (Chen, Xiao, et al., 2025). This type of method improves the interpretability of anomaly localization by considering causal relationships among monitored variables rather than relying exclusively on reconstruction errors or correlation patterns. Deep learning models, including Long Short-Term Memory networks, autoencoders, temporal convolutional networks, and graph neural networks, have also been widely investigated for industrial anomaly detection (de Albuquerque Filho et al., 2022; Pei et al., 2026). These models can learn nonlinear temporal dependencies from large sensor datasets and identify deviations from previously observed operating patterns. Graph-based approaches are particularly suitable for multivariate industrial systems because they represent individual sensors as nodes and model their statistical or physical relationships as edges (Niresi et al., 2024; Gui et al., 2025). However, most existing graph structures are fixed after model training and cannot adequately represent relationship changes caused by evolving production conditions.

Actual manufacturing processes are inherently non-stationary. Tool wear, material replacement, production switching, workload adjustment, seasonal temperature changes, maintenance operations, and controller recalibration can gradually alter both individual sensor distributions and the dependencies among sensors (Chen, Liang, et al., 2025; Abaza & Gheorghita, 2025). These changes do not necessarily indicate equipment faults but may substantially modify the normal operating baseline. Existing anomaly detection methods commonly assume that training and testing data follow similar distributions and that sensor relationships remain relatively stable over time (Liang et al., 2025; Zhang et al., 2026). Once this assumption is violated, a static model may incorrectly classify normal operational drift as an anomaly. Conversely, when a model continuously adapts to newly observed data without distinguishing between drift and faults, it may absorb early fault characteristics into the updated normal pattern, thereby delaying fault detection. The difficulty of separating long-term operational drift from short-term abnormal events remains a major obstacle to the deployment of anomaly detection models in real production environments (Leite et al., 2024; Dachry et al., 2026). Sudden production-mode changes may generate temporary fluctuations across multiple sensors, leading to concentrated false alarms. Slowly developing faults, such as tool degradation, lubrication deterioration, cooling-system blockage, and spindle imbalance, may initially resemble ordinary operational drift and remain undetected until the fault becomes severe (Qi & Qiao, 2026; Su & Zhang, n.d.). These limitations can increase false alarm rates during production transitions and prolong the time required to identify genuine equipment defects (Khanam et al., 2024; Ma, 2026). Excessive false alarms reduce operators' confidence in automated monitoring systems, while delayed detection weakens the practical value of predictive maintenance and may allow avoidable damage to accumulate (Gao et al., 2026; Nsor, 2024). Another limitation of existing research concerns dataset scale and deployment validation. Many studies evaluate

anomaly detection methods using public benchmark datasets or laboratory-scale experiments containing limited operating modes, short observation periods, and relatively simple fault patterns (Yang et al., 2025; Wu & Chen, 2025). Such datasets rarely reflect the simultaneous presence of production switching, varying workloads, maintenance events, sensor noise, gradual degradation, and multiple equipment types found in actual factories (Converso et al., 2023). Model accuracy obtained from these datasets may therefore overestimate performance under industrial conditions (Yin et al., 2026; Bustillo et al., 2022). In addition, computational efficiency is often evaluated on high-performance servers rather than industrial edge devices. The memory consumption, inference throughput, and response latency of graph-temporal models must be considered because manufacturing monitoring systems are frequently required to process continuous sensor streams close to the equipment with limited computing resources.

This study develops an adaptive dependency graph learning model for anomaly detection in non-stationary manufacturing sensor streams. The study aims to dynamically characterize evolving sensor relationships, suppress false alarms caused by long-term operational drift, and preserve sensitivity to short-duration faults and gradual equipment degradation. A sliding-window strategy and an attention-based dependency learning module are used to construct time-varying sensor graphs that reflect changes in operating conditions. Drift-aware normalization separates slowly varying distribution shifts from local abnormal deviations, while a graph-temporal reconstruction network models normal multivariate behavior across both temporal and relational dimensions. Adaptive thresholds are further introduced to accommodate changes in reconstruction-error distributions under different production modes. The proposed method is evaluated using 74.6 million sensor records collected from 42 manufacturing production lines over 96 days. The dataset covers multiple production conditions and verified equipment faults, including tool degradation and cooling-system abnormalities. Compared with a conventional LSTM autoencoder, the proposed model reduces the average detection delay from 18.7 minutes to 6.4 minutes, limits the false alarm rate to 1.72 events per 1,000 operating hours, and achieves a Matthews correlation coefficient of 0.917. The model also processes 12,800 sensor observations per second with a memory requirement of 1.36 GB. These results demonstrate that adaptive dependency modeling can improve the robustness of anomaly detection under changing manufacturing conditions. The study provides a practical approach for distinguishing operational drift from actual faults and supports the deployment of reliable real-time monitoring and predictive maintenance systems on industrial edge platforms.

2. Materials and Methods

2.1. Sample and Study Area Description

The data for this study came from a real-world smart manufacturing factory containing 42 active production lines. We monitored 11 types of industrial machines, such as robotic arms, CNC machines, and conveyor belts, over 96 days of continuous operation. A total of 318 sensors were installed across these lines to collect physical signals, generating 74.6 million data points with timestamps. The dataset contains regular production cycles along with common machinery problems, including tool wear, cooling system issues, unusual vibrations, and actuator failures. All sensor data were collected under changing production modes to ensure the study accurately reflects shifting industrial environments.

2.2. Experimental Design and Control Experiments

To test our model, we split the dataset into two parts based on time. Data from the first 60 days

of normal operations were used to train the model, while the remaining 36 days—which included both normal shifts and machinery faults—were used for testing. The experimental group used our new adaptive graph model that updates sensor relationships in real time. For comparison, we set up two control experiments: a standard static LSTM autoencoder that does not update sensor connections, and a modified version of our own model with the normalization module turned off. This setup allowed us to measure exactly how well our model separates true equipment faults from normal changes caused by switching production modes.

2.3. Measurement Methods and Quality Control

Sensor data were collected through an edge computing system at a fixed rate of 100 Hz. To remove background noise, all raw sensor signals were cleaned using a median filter before analysis. Missing values caused by brief network delays made up less than 0.05% of the total data and were filled in using linear interpolation. We used the Network Time Protocol (NTP) to sync the clocks of all 318 sensors, keeping time differences under 1 millisecond. Finally, to ensure accurate ground-truth labels, all recorded anomalies and machine faults were double-checked against actual factory maintenance logs and operator repair reports.

2.4. Data Processing and Model Formulas

Before training the model, we normalized the sensor data using a moving window to remove the effects of long-term operational changes. For any sensor reading x_t at time t , the normalized value was calculated using the local mean μ_w and standard deviation σ_w within a historical window of length w :

$$\tilde{x}_t = \frac{x_t - \mu_w}{\sigma_w + \epsilon}$$

where ϵ is a very small number used to avoid dividing by zero. To detect anomalies, we measured the reconstruction error of the network. The anomaly score S_t at time t was calculated using the root-mean-square error (RMSE) between the actual normalized input vector $\tilde{x}_{t,i}$ and the reconstructed output vector $\hat{x}_{t,i}$:

$$S_t = \sqrt{\frac{1}{N} \sum_{i=1}^N (\tilde{x}_{t,i} - \hat{x}_{t,i})^2}$$

where N is the total number of sensors. An anomaly alert was triggered whenever the score went above a moving threshold based on past reconstruction errors.

3. Results and Discussion

3.1. Detection Performance under Mixed Production Modes

The proposed adaptive dependency graph learning model achieved high accuracy under non-stationary conditions, reaching a Matthews correlation coefficient (MCC) of 0.917. This represents a substantial improvement over standard static models, which often fail when data distributions drift due to production switching or workload changes (Yang, 2026; Zhang, 2026). By updating sensor relationships dynamically, our framework successfully captured shifting dependencies across the 42 production lines without losing tracking sensitivity. When compared to the baseline static LSTM autoencoder, our method maintained a stable performance curve across all 11 machine

categories. This stability proves that integrating dynamic graphs allows the network to learn robust behavioral representations, which prevents the severe performance drops typically observed in traditional deep learning models when operating under mixed industrial modes.




Fig. 1. Comparison of anomaly detection accuracy across different models under mixed manufacturing modes.

3.2. Reduction in Detection Delay and False Alarms

A key finding of this study is the significant reduction in average detection delay, which dropped from 18.7 minutes to 6.4 minutes. Traditional methods often require broad temporal windows or high residual thresholds to avoid false alerts during minor operational changes, which delays the detection of actual equipment faults. Our model avoids this issue by using adaptive thresholds that adjust to current manufacturing modes in real time. Consequently, acute anomalies like actuator faults and sudden structural vibrations were flagged almost immediately. At the same time, the false alarm rate was kept at a low 1.72 events per 1,000 operating hours. This balance confirms that our thresholding strategy successfully separates routine operational shifts from true mechanical failures, providing a more reliable tool for real-time factory monitoring (Sharma, 2024; Li & Liu, 2026).

3.3. Ablation Study and Feature Drift Management

The ablation analysis confirmed the necessity of both the dynamic graph module and the drift-aware normalization step. When we deactivated the dynamic graph component, the average detection delay increased by 9.2 minutes, showing that static spatial topologies cannot handle evolving sensor interactions during tool degradation. Furthermore, removing the drift-aware normalization module led to frequent false alarms immediately following production transitions. This occurs because the model misinterprets normal baseline changes as operational anomalies. By decoupling these long-term distribution shifts from localized fluctuations, our normalization method helps the network maintain a clean residual baseline, preventing the threshold errors that typically hinder non-stationary time-series models.




Fig. 2. Dynamic threshold changes and error patterns during production switching events.

3.4. Computational Efficiency and Edge Deployment Feasibility

Online inference tests on a standard edge workstation showed that the proposed model processed 12,800 sensor points per second while maintaining a stable memory footprint of 1.36 GB. This processing speed is well above the 100 Hz sampling rate required by the 318 industrial sensors, confirming that the model can handle large-scale, continuous industrial workloads in real time. The low memory consumption is primarily due to our streamlined sliding-window approach and efficient graph updates, which avoid the high computational costs associated with fully connected spatial networks. These results demonstrate that our adaptive framework can be easily deployed on local factory edge nodes, providing fast, on-site anomaly detection without relying on heavy cloud computing resources (Huong et al., 2022; Xu et al., 2026).

4. Conclusion

This study successfully addresses the challenge of non-stationary sensor signals in smart manufacturing systems by using an adaptive dependency graph learning model. By combining sliding-window correlation with attention-based temporal encoding, the model tracks changing sensor relationships in real time. In addition, the drift-aware normalization step effectively separates long-term operational changes from actual, short-term equipment faults. Testing on a large industrial dataset proves that our method lowers the average detection delay from 18.7 to 6.4 minutes and keeps a very low false alarm rate of 1.72 events per 1,000 hours. These results show that dynamic spatiotemporal modeling can significantly improve the accuracy of anomaly detection in shifting industrial environments. For practical use, the low memory requirement of 1.36 GB and a high speed of 12,800 points per second make this model highly suitable for factory edge devices. This allows for real-time, on-site predictive maintenance on complex production lines without using cloud storage. However, a limitation of this work is that the model needs a continuous data

window to calculate local changes, which might lower performance if the network drops out for a long time. Future work will focus on developing independent updates for separate sensor groups to protect the system against data transmission losses.

Acknowledgments

We are grateful to all respondents who participated in this study.

References

1. Abaza, B. F., & Gheorghita, V. (2025). Artificial neural network framework for hybrid control and monitoring in turning operations. *Applied Sciences*, 15(7), 3499.
2. Bustillo, A., Reis, R., Machado, A. R., & Pimenov, D. Y. (2022). Improving the accuracy of machine-learning models with data from machine test repetitions. *Journal of Intelligent Manufacturing*, 33(1), 203-221.
3. Chen, F., Liang, H., Li, S., Yue, L., & Xu, P. (2025). Design of Domestic Chip Scheduling Architecture for Smart Grid Based on Edge Collaboration.
4. Chen, X., Xiao, H., Zeng, Z., Zhang, S., & Xiao, R. (2025). Fine-Grained Multivariate Time Series Anomaly Detection via Causal Inference. *Knowledge-Based Systems*, 114765.
5. Converso, G., Gallo, M., Murino, T., & Vespoli, S. (2023). Predicting failure probability in Industry 4.0 production systems: a workload-based prognostic model for maintenance planning. *Applied Sciences*, 13(3), 1938.
6. Dachry, A., Messaoudi, N., & Dachry, W. (2026). An Adaptive Anomaly Detection Platform for Industrial IoT Streams via LSTM Autoencoders and Human-in-the-Loop Feedback: Supply Chain Application. *International Journal of Intelligent Engineering & Systems*, 19(5), 698.
7. de Albuquerque Filho, J. E., Brandao, L. C., Fernandes, B. J. T., & Maciel, A. M. (2022). A review of neural networks for anomaly detection. *IEEE Access*, 10, 112342-112367.
8. Gao, G., Gao, R., Lu, C., Gao, R., & Kuang, Y. (2026, March). Security Governance Methods and Quantitative Evaluation for Enterprise SMS and Verification Code Systems. In *2026 International Conference on Generative Artificial Intelligence and Information Security (GAIIS)* (pp. 455-458). IEEE.
9. Gui, H., Wang, B., Lu, Y., & Fu, Y. (2025). Computational Modeling-Based Estimation of Residual Stress and Fatigue Life of Medical Welded Structures.
10. Huong, T. T., Bac, T. P., Ha, K. N., Hoang, N. V., Hoang, N. X., Hung, N. T., & Tran, K. P. (2022). Federated learning-based explainable anomaly detection for industrial control systems. *IEEE Access*, 10, 53854-53872.
11. Khanam, R., Hussain, M., Hill, R., & Allen, P. (2024). A comprehensive review of convolutional neural networks for defect detection in industrial applications. *IEEE Access*, 12, 94250-94295.
12. Leite, D., Andrade, E., Rativa, D., & Maciel, A. M. (2024). Fault detection and diagnosis in industry 4.0: A review on challenges and opportunities. *Sensors*, 25(1), 60.

13. Li, Y., & Liu, S. (2026, May). A Study on Dynamic Optimization of Alerting Policies and Multi-Agent Decision-Making Mechanisms in Cloud Environments. In 2026 7th International Seminar on Artificial Intelligence, Networking and Information Technology (AINIT) (pp. 703-706). IEEE.
14. Liang, R., Ye, Z., Liang, Y., & Li, S. (2025). Deep Learning-Based Player Behavior Modeling and Game Interaction System Optimization Research.
15. Ma, Y. (2026). Industrial Financial Risk Prediction Model Based on Graph Neural Network and Knowledge Graph Inference. *Journal of Circuits, Systems and Computers*, 35(16), 2650103.
16. Niresi, K. F., Bissig, H., Baumann, H., & Fink, O. (2024). Physics-enhanced graph neural networks for soft sensing in industrial internet of things. *IEEE Internet of Things Journal*, 11(21), 34978-34990.
17. Nsor, M. (2024). Predictive maintenance using machine learning for engineering systems through real-time sensor data and anomaly detection models. *Intl. J. of Res. Publ. and Rev*, 5(10), 5167-5183.
18. Pei, Z., Huang, Q., & Wang, S. (2026). When LLMs Develop Languages: Symbolic Communication for Efficient Multi-Agent Reasoning. arXiv preprint arXiv:2606.29354.
19. Qi, C., & Qiao, X. (2026, May). From Traditional Machine Learning to Large Language Models: The Evolution of Infrastructure from an Engineering Perspective. In 2026 6th International Symposium on Computer Technology and Information Science (ISCTIS) (pp. 503-506). IEEE.
20. Sharma, R. (2024). Enhancing industrial automation and safety through real-time monitoring and control systems. *International Journal on Smart & Sustainable Intelligent Computing*, 1(2), 1-20.
21. Su, D., & Zhang, H. Developing a Data-Driven Computational Framework for Scalable Special Education Practices for Autism and Intellectual Disabilities in the US Public School System.
22. Tsanousa, A., Bektsis, E., Kyriakopoulos, C., González, A. G., Leturiondo, U., Gialampoukidis, I., ... & Kompatsiaris, I. (2022). A review of multisensor data fusion solutions in smart manufacturing: Systems and trends. *Sensors*, 22(5), 1734.
23. Wu, C., & Chen, H. (2025). Research on system service convergence architecture for AR/VR system.
24. Xu, S., & Zhou, Y. (2025). Machine Learning Applications in Financial Statement Fraud Detection: A Comparative Analysis. *Annals of Applied Sciences*, 6(1).
25. Xu, T., Zhu, W., & Zhang, J. (2026). Stability and Consistency of Explainable Deep Learning Methods in Credit Risk Assessment. Available at SSRN 6893938.
26. Yang, J. (2026). Stage-Coupled Computational Framework for Stratified Accessibility and Equity Analysis in Community-Based Elderly Care Services.
27. Yang, M., Wu, J., Tong, L., & Shi, J. (2025). Design of Advertisement Creative Optimization and Performance Enhancement System Based on Multimodal Deep Learning.
28. Yin, J., Huang, Y., & Rao, H. (2026). A Study on User Behavior Signal-Driven Predictive Models for Digital Platform Consumption Trends. Available at SSRN 6607298.

29. Zhang, Z. (2026). A Study on the Identification of Manipulative Design in Subscription and Payment Interfaces of Digital Consumer Platforms and Its Behavioral Effects. Available at SSRN 6734760.
30. Zhang, Z., Gao, Y., & Tong, Y. (2026). Semantic-Temporal Graph Learning for Demand Forecasting and Resource Allocation in Public Information Systems. Available at SSRN 6792279.