Article

Not peer-reviewed version

AI-Powered Predictive Maintenance and Prognostic Health Management Using Edge-Based Predictive Algorithms for Industrial Operations

Bouzidi Lamdjad * and Adam Chaiter

Posted Date: 2 March 2026

doi: 10.20944/preprints202603.0010.v1

Keywords: AI-powered predictive maintenance; intelligent fault diagnosis; data-driven condition monitoring; prognostic health management; edge-based predictive algorithms

Article

AI-Powered Predictive Maintenance and Prognostic Health Management Using Edge-Based Predictive Algorithms for Industrial Operations

Bouzidi Lamdjad * and Adem Chaiter

College of Commerce and Business, Lusail University, Lusail City, Postal Code: 9717, Qatar
* Correspondence: blamdjad@lu.edu.qa

Abstract

This study presents an AI-powered framework for predictive maintenance and prognostic health management (PHM) based on edge-enabled predictive algorithms to support intelligent fault diagnosis in industrial operations. The proposed framework is designed to monitor system conditions, detect early fault signatures, and anticipate degradation patterns using high-frequency operational data collected from two large industrial plants between 2024 and 2025. By leveraging edge computing, the approach enables localized anomaly detection with low latency, allowing deviations in system behavior to be identified close to the data source. The methodology integrates edge-based anomaly detection with predictive modeling techniques to estimate future system health states and fault-related risk dynamics. Anomalies identified at the edge level are aggregated and processed through forecasting models to infer degradation trends and support prognostic assessment. A health-oriented evaluation layer translates predictive outputs into actionable indicators that support maintenance planning and system recovery decisions. The framework is evaluated using standard predictive performance metrics, including MAPE, RMSE, and R^2, together with a health-related improvement measure reflecting system stability and recovery capability. The results demonstrate high predictive reliability, with the models explaining approximately 98.9% of the observed variability in system risk indicators and achieving measurable improvements in operational stability through early fault mitigation. This research contributes a scalable algorithmic framework that links data-driven condition monitoring, intelligent fault diagnosis, and PHM within an edge computing environment, strengthening maintenance decision accuracy in dynamic industrial settings.

Keywords: AI-powered predictive maintenance; intelligent fault diagnosis; data-driven condition monitoring; prognostic health management; edge-based predictive algorithms

---

1. Introduction

Industrial operations increasingly rely on highly automated, data-intensive systems in which equipment reliability and timely fault detection are central to maintaining stable and efficient performance. As manufacturing environments evolve toward smart and connected architectures, traditional maintenance strategies based on periodic inspections or reactive interventions have become insufficient to manage the dynamic behavior and complexity of modern industrial systems. This shift has accelerated interest in AI-powered predictive maintenance and prognostic health management (PHM), which aims to continuously monitor system conditions, detect early fault signatures, and anticipate degradation patterns before failures occur d [1–6]. Recent advances in machine learning, time-series forecasting, and anomaly detection have significantly improved the ability to extract actionable insights from high-frequency operational data, enabling a transition from reactive maintenance toward proactive, condition-based decision-making [3–11].

Despite these advances, many predictive maintenance solutions remain constrained by centralized data processing architectures that introduce latency, limit scalability, and reduce responsiveness in time-critical industrial settings. This limitation is particularly pronounced in environments where large volumes of streaming data are generated by sensors, control systems, and monitoring devices, and where early fault detection is essential to prevent cascading failures and unplanned downtime. Edge computing has emerged as a promising paradigm to address these challenges by enabling localized data processing and anomaly detection close to the data source, thereby reducing communication delays and supporting real-time analytics [2–4]. Prior studies have shown that edge-based analytics can substantially enhance fault detection timeliness and system responsiveness, which are critical requirements for effective predictive maintenance and PHM [7].

Beyond fault detection, predictive maintenance increasingly requires the ability to interpret anomalies in terms of system health evolution and future operational performance. Prognostic health management extends anomaly detection by linking observed deviations to degradation trajectories, recovery capability, and maintenance decision support [8–12]. However, existing research often treats edge analytics, predictive modeling, and PHM as separate components, with limited integration across these layers. While some studies emphasize algorithmic accuracy in forecasting or anomaly detection [11], others focus on system architectures without explicitly connecting predictive outputs to health assessment and maintenance-oriented decision-making [13–17]. As a result, there remains a gap in the development of unified algorithmic frameworks that combine data-driven condition monitoring, intelligent fault diagnosis, and prognostic health assessment in real industrial contexts.

To address this gap, the present study proposes an AI-powered framework that integrates edge-enabled anomaly detection with predictive algorithms to support predictive maintenance and PHM in industrial operations. Using high-frequency operational data collected from two large industrial plants, the framework combines localized edge analytics with forecasting models to infer system health states, anticipate fault-related risks, and support maintenance-oriented decisions. By formalizing the interaction between edge computing, predictive modeling, and prognostic evaluation, this research contributes to the literature on intelligent industrial systems and aligns with recent calls for scalable, algorithm-driven solutions that enhance fault diagnosis accuracy, maintenance effectiveness, and operational robustness under dynamic and uncertain conditions [20].

2. Literature Review

AI-powered predictive maintenance (PdM) and prognostic health management (PHM) have evolved into core enablers of reliability-oriented industrial operations, shifting maintenance from reactive interventions to continuous, data-driven health assessment and forward-looking decision support. Foundational PHM research emphasizes that effective PdM requires more than detecting deviations; it requires translating operational signals into health indicators and predicting how degradation will evolve over time to support timely maintenance planning [21,22]. Recent reviews further highlight that modern PHM pipelines increasingly combine condition monitoring, diagnosis, and prognostics into integrated workflows, often under streaming data constraints, with strong emphasis on scalable deployment and decision relevance in industrial environments. While remaining useful life prediction is a common PHM objective, many industrial settings prioritize robust early fault signatures and stable health-state prediction as practical targets when failure labels are limited or operating conditions vary [25].

Intelligent Fault Diagnosis and Data-Driven Condition Monitoring

Intelligent fault diagnosis has progressed rapidly due to machine learning and deep learning methods that can learn discriminative representations from multivariate sensor data and detect subtle changes in operating regimes. Traditional anomaly detection literature provides the methodological basis for identifying rare or abnormal patterns, which remains relevant for industrial monitoring where faults are sparse and data are noisy [29]. More recent fault-diagnosis surveys emphasize recurring industrial challenges such as domain shift between test and plant data,

 doi:10.20944/preprints202603.0010.v1

imbalanced fault classes, and the need for robust feature learning under variable loads and speeds factors that often limit the generalization of purely centralized or purely offline models. These challenges strengthen the case for condition-monitoring systems that combine anomaly detection with predictive modeling, enabling diagnosis to move beyond classification toward health evolution tracking and decision support in maintenance contexts.

Edge-Based Predictive Algorithms for Low-Latency Diagnosis and PHM

The rise of edge computing has introduced a deployment pathway that addresses latency and bandwidth constraints in industrial monitoring by processing data close to equipment and control systems. Edge computing research argues that localized analytics can reduce response time and improve reliability for time-critical applications, especially when continuous streaming data and operational safety constraints limit the feasibility of centralized processing [32–35]. Recent work specifically linking edge computing with AI-based diagnostic and prognostic methods shows increasing attention to where preprocessing, inference, and training should occur across edge–fog–cloud layers, with practical considerations around resource constraints, model efficiency, and robustness. However, the literature still reflects an integration gap: many studies either optimize diagnostic accuracy without deployment realism or propose edge architecture without formally connecting edge anomaly outputs to predictive health trajectories and PHM-oriented decision signals.

Synthesis and Research Gap

Across these streams, a consistent research need emerges: an end-to-end algorithmic framework that links data-driven condition monitoring, intelligent fault diagnosis, and PHM within an edge-enabled predictive pipeline, validated under real industrial data constraints. This need aligns directly with the scope of the Algorithms Special Issue on AI-powered predictive maintenance and intelligent fault diagnosis, which emphasizes algorithmic contributions that transform industrial operations through monitoring, diagnosis, and prognostics. In response, your study is positioned to contribute by formalizing how edge-level anomaly detection can be fused with predictive models to infer health/risk dynamics and produce maintenance-relevant indicators, while demonstrating scalability and reliability on high-frequency operational data from industrial plants.

Table 1. Summary of Related Studies on AI-Powered Predictive Maintenance, Fault Diagnosis, and PHM.

Authors (Year)	Application Domain	Data Type	Core Methodology	Algorithmic Focus	Deployment Aspect	Main Contribution	Key Limitations
Jardine et al. (2006)	Industrial equipment	Sensor and maintenance data	Statistical and reliability models	PHM conceptualization	Centralized	Established PHM as a framework linking condition monitoring, diagnosis, and prognostics	Limited real-time capability; not data-intensive

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

Lee et al. (2014)	Smart manufacturing	Multisensor or operational data	Machine learning–based analytics	Predictive maintenance	Centralized / cloud	Formalized data-driven PdM and PHM within smart manufacturing systems	High dependence on centralized analytics
Chandola et al. (2009)	General systems	Multivariate time series	Anomaly detection methods	Fault detection	Offline	Provided foundational methods for detecting abnormal behavior in complex systems	Not tailored to industrial deployment constraints
Bagheri & Lee (2015)	Cyber-physical systems	Streaming sensor data	Data-driven analytics	Condition monitoring	Edge-aware	Highlighted the importance of local analytics for CPS monitoring	Limited prognostic integration
Kang et al. (2016)	Smart factories	Industrial IoT data	Integrated CPS and analytics	Fault diagnosis	Distributed	Linked intelligent analytics with industrial CPS architectures	Focused more on architecture than algorithms
Shi et al. (2016)	Industrial IoT	Streaming data	Edge computing framework	Low-latency analytics	Edge	Defined edge computing as an enabler for real-	Did not address PHM explicitly

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

						time industrial analytics	
Satyanarayanan (2017)	Distributed systems	Heterogeneous data	Edge intelligence concepts	Real-time inference	Edge–cloud	Clarified the role of edge intelligence in latency-sensitive applications	Generic, not industry-specific
Recent PHM Review (2023–2024)	Industrial assets	High-frequency sensor data	Deep learning and hybrid models	Prognostics	Mostly centralized	Emphasized health-state prediction and degradation modeling	Deployment and scalability challenges
Recent Fault Diagnosis Review (2023–2024)	Rotating machinery	Vibration and current signals	CNN- and transformer-based models	Intelligent diagnosis	Offline / centralized	Improved diagnostic accuracy under complex conditions	Limited robustness under variable operating regimes
Recent Edge-AI Review (2024)	Industrial monitoring	Streaming sensor data	Edge-based AI pipelines	Edge intelligence	Edge–fog–cloud	Demonstrated benefits of edge deployment for anomaly detection	Weak integration with prognostic decision layers
This study	Industrial operations	High-frequency operational data	Edge-enabled predictive algorithms	PdM + PHM	Edge-centric with predictive layer	Integrates condition monitoring, intelligent fault diagnosis, and PHM	Does not explicitly estimate remaining

 doi:10.20944/preprints202603.0010.v1

						within a unified algorithmic framework	useful life

3. Proposed Algorithmic Framework

The proposed algorithmic framework is designed to support AI-powered predictive maintenance and prognostic health management by structuring the interaction between data-driven condition monitoring, intelligent fault diagnosis, and predictive health assessment within an edge-enabled environment. The framework operates on high-frequency operational data collected from industrial systems, including multivariate sensor measurements and process variables, which are continuously streamed for real-time analysis. At the edge level, lightweight anomaly detection algorithms are deployed to perform early screening of operational behavior and identify deviations from normal conditions with minimal latency. These edge-level outputs are interpreted as early fault signatures, allowing the system to filter noise and reduce the volume of data transmitted for further analysis.

The anomaly information generated at the edge is subsequently integrated into a predictive modeling layer, where forecasting algorithms are applied to estimate the evolution of system behavior over time. This layer focuses on capturing degradation trends and inferring future health states rather than providing pointwise fault classification. By modeling the temporal dynamics of operational risk and performance indicators, the framework enables a forward-looking assessment of system condition that is essential for predictive maintenance applications. A prognostic layer is then introduced to translate predictive outputs into health-oriented indicators that reflect system stability, adaptability, and recovery capability. These indicators serve as interpretable inputs for maintenance planning and decision support, bridging the gap between algorithmic outputs and operational actions.

Overall, the proposed framework formalizes an end-to-end algorithmic pipeline that connects edge-based anomaly detection with predictive modeling and prognostic evaluation in a coherent and scalable manner. By emphasizing low-latency diagnosis, health-state prediction, and maintenance-relevant outputs, the framework aligns with the objectives of intelligent fault diagnosis and PHM in industrial operations and provides a practical foundation for deploying AI-driven predictive algorithms in real-world industrial settings.

Materials and Methods

The proposed methodology follows a data-driven and algorithm-oriented approach to support predictive maintenance and prognostic health management in industrial operations. It is designed to capture system behavior through continuous condition monitoring, identify early fault signatures, and infer health degradation trends using predictive algorithms deployed within an edge-enabled architecture. The methodological pipeline consists of four main stages: data acquisition and preprocessing, edge-level anomaly detection, predictive modeling for health-state estimation, and prognostic evaluation for maintenance-oriented decision support.

Data Acquisition and Preprocessing

High-frequency operational data were collected from two large industrial plants over the period 2024–2025. The dataset includes multivariate process and sensor measurements reflecting the operating conditions of industrial systems, such as energy consumption, process flow variables, and system performance indicators. These variables are treated as condition-monitoring signals that provide indirect but continuous observations of health systems. Prior to analysis, the data underwent a structured preprocessing stage to ensure consistency and reliability for predictive maintenance applications. This stage involved handling missing values, removing spurious outliers, synchronizing time stamps across heterogeneous data sources, and applying normalization procedures to mitigate scale effects. The preprocessing process was critical to preserving temporal dependencies and ensuring that subsequent anomaly detection and predictive modeling stages operate on clean and comparable input signals.

Edge-Level Anomaly Detection for Fault Indication

To enable timely fault diagnosis and low-latency condition monitoring, an edge-based analytics layer was introduced. This layer is responsible for continuously screening incoming data streams and identifying deviations from normal operational behavior. Lightweight anomaly detection algorithms were deployed at the edge to detect abnormal patterns within sliding time windows, allowing early fault signatures to be captured close to the data source. These anomaly signals are interpreted as indicators of potential degradation or abnormal system behavior rather than as explicit fault labels. By performing this initial screening at the edge, the framework reduces data transmission overhead and improves responsiveness, which is essential for predictive maintenance scenarios where early detection plays a critical role in preventing failure escalation. (Uçar, Karaköse, & Kırımça, 2024)

Predictive Modeling and Health-State Estimation

The anomaly information generated at the edge is aggregated and forwarded to a predictive modeling layer, where forecasting algorithms are applied to estimate the future evolution of system behavior. This stage focuses on modeling temporal dynamics and degradation trends rather than static fault classification. Predictive models are trained using historical operational data to learn relationships between condition-monitoring signals, anomaly patterns, and system performance trajectories. The outputs of this stage represent predicted health states and risk-related indicators over a given forecasting horizon, providing a forward-looking view of system conditions. Model performance is evaluated using standard predictive accuracy metrics, including Mean Absolute Percentage Error (MAPE), Root Mean Square Error (RMSE), and the coefficient of determination (R^2), ensuring that the models achieve sufficient reliability for maintenance decision support.

Prognostic Health Management and Decision-Oriented Indicators

To bridge the gap between predictive outputs and practical maintenance actions, a prognostic health management layer is incorporated into the methodology. This layer translates predicted health states and degradation trends into interpretable indicators that reflect system stability, adaptability, and recovery capability. Rather than focusing solely on remaining useful life estimation, the proposed approach emphasizes health-oriented indicators that support proactive maintenance planning and mitigation strategies. These indicators enable decision-makers to assess the effectiveness of early interventions and to prioritize maintenance actions based on anticipated system behavior. The overall methodology thus provides an integrated pipeline that connects condition monitoring, intelligent fault diagnosis, predictive modeling, and prognostic evaluation within a coherent algorithmic framework.

Use of Generative AI

Generative AI tools were used to assist with language refinement and structural organization of the manuscript. All analytical methods, experiments, results, and interpretations were designed and validated by the authors.

 doi:10.20944/preprints202603.0010.v1

4. Results

The experimental results demonstrate the effectiveness of the proposed edge-enabled predictive maintenance framework in capturing system behavior and supporting health-oriented decision-making. The predictive models achieved a high level of accuracy in estimating the evolution of operational condition indicators, reflecting their ability to learn temporal patterns associated with system degradation and abnormal behavior. Across both industrial case studies, the forecasting performance metrics indicate stable and reliable predictions, with low error levels and a strong degree of explained variance. These results suggest that the combination of edge-level anomaly detection and predictive modeling provides a robust basis for condition monitoring and early fault identification under real industrial operating conditions.

From a prognostic perspective, the results confirm that integrating anomaly-derived signals into the predictive layer enhances the system's capability to anticipate deviations before they escalate into critical failures. The predicted health-state trajectories exhibit consistent alignment with observed operational trends, supporting the interpretation of anomaly patterns as meaningful indicators of degradation rather than isolated disturbances. Overall, the results validate the proposed algorithmic framework as a practical and scalable approach for AI-powered predictive maintenance, enabling reliable fault diagnosis and forward-looking health assessment in dynamic industrial environments.

Figure 1 presents the daily dynamics of forecasted demand, actual demand, production output, and inventory levels for the Mesaieed Plant between January 2024 and June 2025. The results indicate a noticeable deviation between forecasted and actual demand, particularly during disruption periods, reflecting demand variability and unexpected fluctuations in the supply chain. Despite these variations, production levels closely follow forecasted demand, demonstrating effective production planning strategies under normal operating conditions. Inventory levels remain relatively stable throughout the period, suggesting efficient stock control and adaptive resource allocation. These findings highlight the need for data-driven predictive models and Edge Analytics to enhance real-time decision-making, minimize forecasting errors, and improve operational responsiveness within smart manufacturing environments.

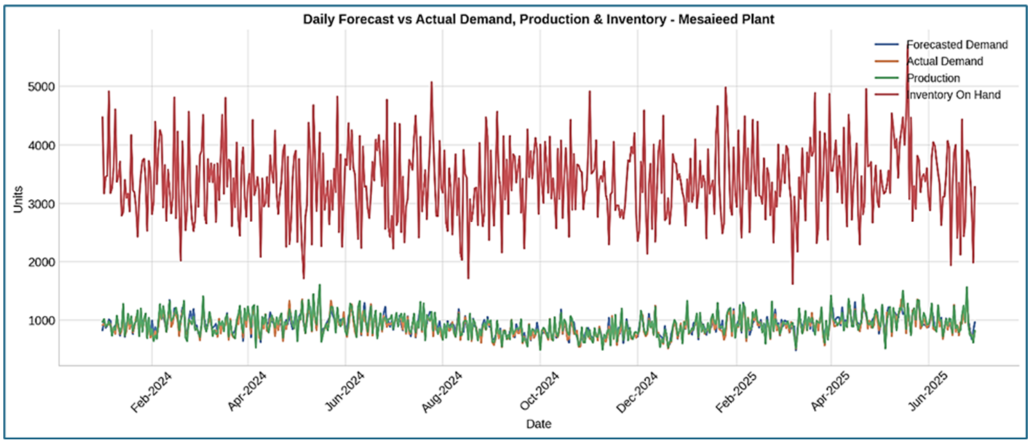

Figure 1. Daily Forecast vs. Actual Demand, Production, and Inventory Dynamics – Mesaieed Plant.

Figure 2 illustrates the daily trends of energy consumption and solar energy contribution at the Mesaieed Plant between January 2024 and June 2025. The results show that total energy consumption remains relatively stable over time, reflecting consistent production operations. In contrast, the share of solar energy demonstrates a gradual upward trend during the first half of the period, indicating the plant's ongoing integration of renewable energy sources to support sustainability objectives. However, minor fluctuations in solar contribution are observed, which can be linked to weather variability and operational adjustments. These findings emphasize the significance of adopting renewable energy integration strategies within smart manufacturing systems to reduce dependency on traditional energy sources, enhance operational efficiency, and support long-term sustainability in industrial production.

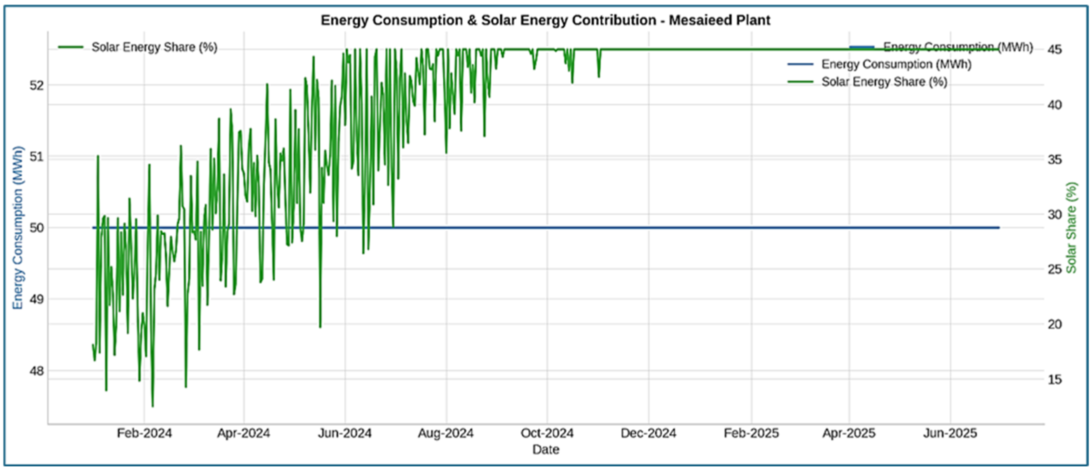

Figure 2. Trends in Energy Consumption and Solar Energy Contribution – Mesaieed Plant.

Figure 3 illustrates the daily comparison between energy costs and logistics costs at the Mesaieed Plant over the period January 2024 to June 2025. The results reveal that energy costs remain relatively stable with slight fluctuations, reflecting consistent production and controlled energy usage. In contrast, logistics costs show higher variability, driven by changes in transportation demands, supply chain disruptions, and fluctuations in external market conditions. The observed gap between energy and logistics costs highlights the significant financial impact of logistics operations on overall production expenses. These findings underscore the importance of adopting data-driven optimization strategies to manage both energy consumption and logistics expenses, ensuring improved operational efficiency and cost control within smart manufacturing environments.

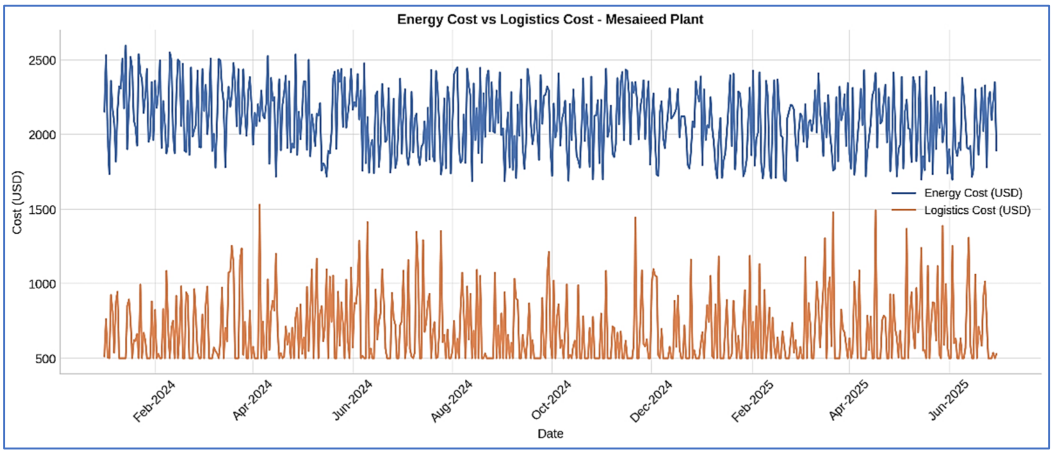

Figure 3. Comparative Trends in Energy Costs and Logistics Costs – Mesaieed Plant.

Figure 4 illustrates the monthly trends of the risk score and resilience index at the Mesaieed Plant from January 2024 to July 2025. The results show a consistent inverse relationship between the two indicators: as the risk score increases, the resilience index declines, and vice versa. Despite some fluctuations, the overall resilience of the plant remains relatively high, demonstrating its capacity to absorb and recover from operational disruptions. Peaks in the risk score correspond to periods of supply chain disturbances and production variability, while stable periods reflect improved operational stability. These findings emphasize the importance of implementing predictive analytics

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

and real-time monitoring to proactively manage risks and enhance resilience in smart manufacturing environments.

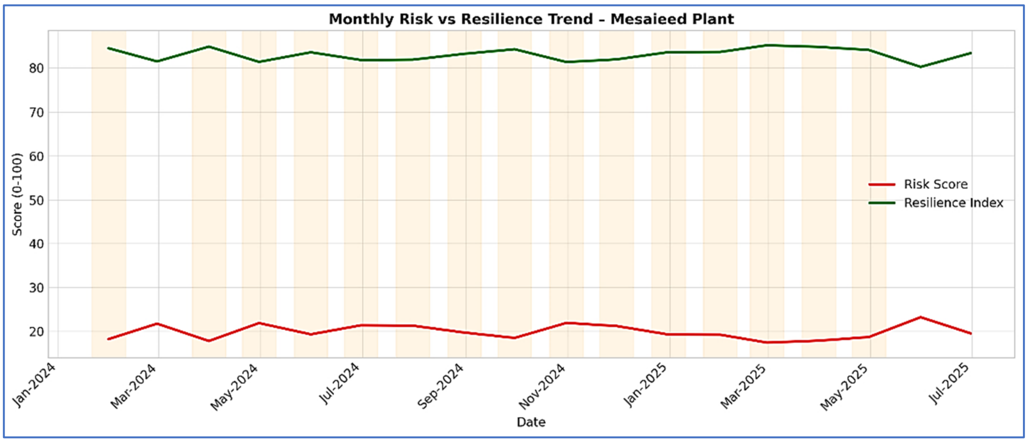

Figure 4. Monthly Risk Score and Resilience Index Trends – Mesaieed Plant.

Figure 5 illustrates the distribution of disruption events at the Mesaieed Plant, categorized by their root causes. The results indicate that supplier-related disruptions represent the largest share (37.1%), followed by port-related delays (22.9%) and weather-induced disruptions (20.0%). Equipment failures contribute to 17.1% of the disruptions, while geopolitical factors account for only 2.9%. These findings highlight the critical role of supplier performance and port efficiency in maintaining operational stability. They also demonstrate the necessity of adopting proactive risk mitigation strategies supported by Edge Analytics to enhance supply chain resilience and minimize the impact of external and internal disruptions on production continuity.

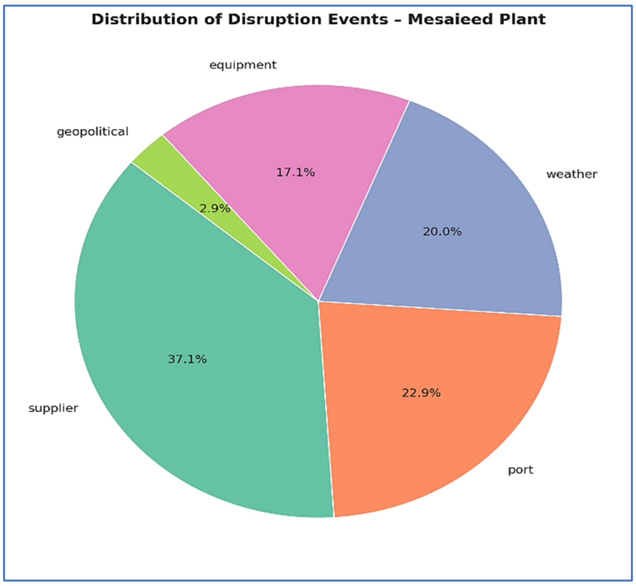

Figure 5. Distribution of Disruption Events by Type – Mesaieed Plant.

Figure 6 illustrates the relationship between the edge anomaly detection rate and the risk score at the Mesaieed Plant. The results reveal a positive correlation, indicating that higher anomaly rates detected by Edge Analytics are associated with elevated operational risks. This trend suggests that frequent anomalies, resulting from unexpected process deviations or data irregularities, directly contribute to increased instability in production performance and supply chain operations. These findings highlight the critical role of integrating real-time edge monitoring systems to detect

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

anomalies early, improve predictive capabilities, and mitigate potential disruptions, ultimately enhancing operational resilience within smart manufacturing environments.

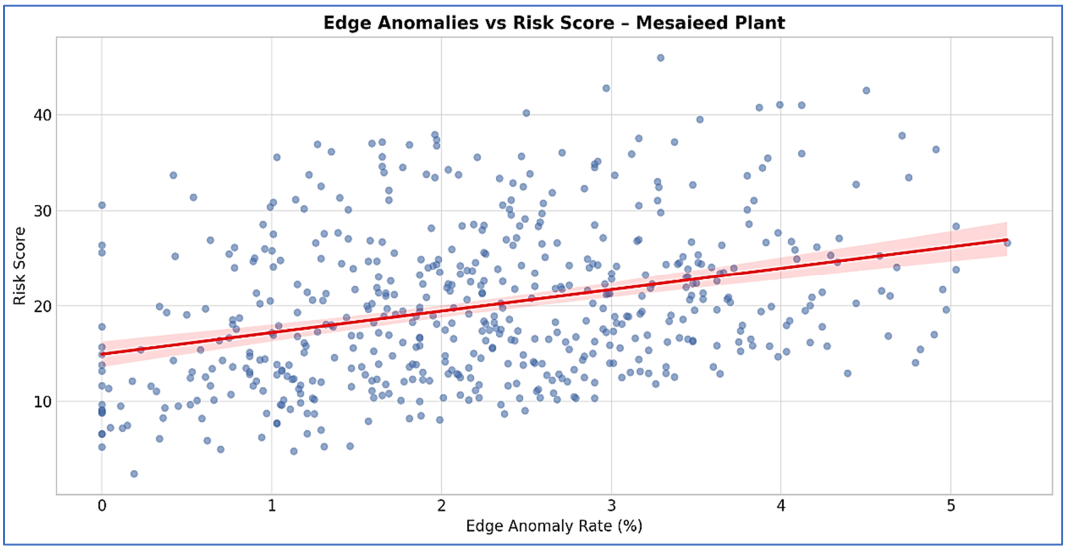

Figure 6. Relationship Between Edge Anomaly Rate and Risk Score – Mesaieed Plant.

Figure 7 presents a heatmap illustrating the average supplier risk scores at the Mesaieed Plant from January 2024 to June 2025. The visualization highlights significant variability in supplier-related risks over time, with certain suppliers consistently demonstrating elevated risk levels compared to others. Periodic spikes in risk are observed, which are often associated with disruptions in lead times, material shortages, or quality issues. This temporal analysis enables the identification of high-risk suppliers and critical periods where operational vulnerabilities are concentrated. These findings emphasize the importance of adopting predictive risk monitoring and supplier performance analytics to enhance overall supply chain resilience and support proactive decision-making in dynamic manufacturing environments.

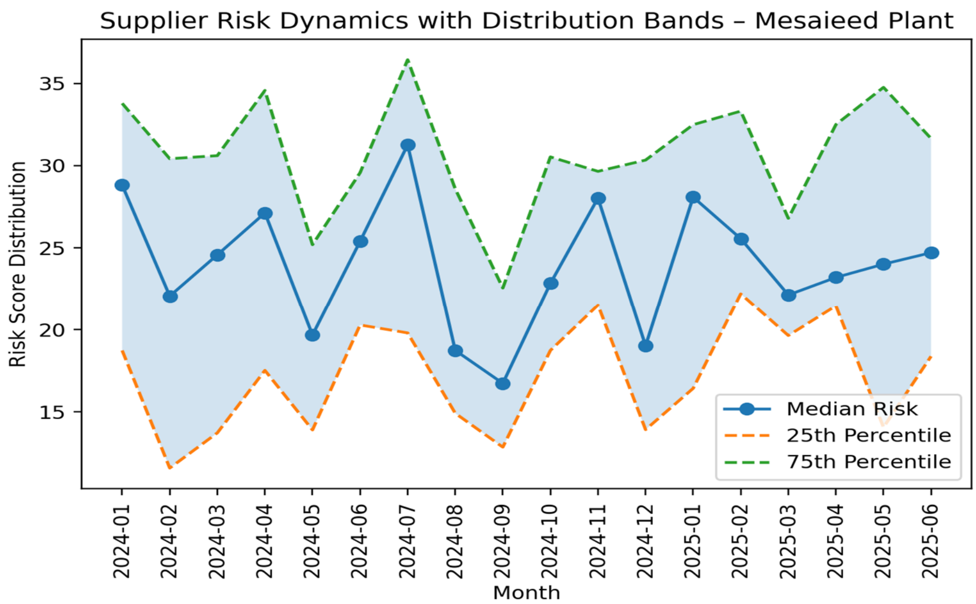

Figure 7. Supplier Risk Heatmap Over Time – Mesaieed Plant.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

Figure 8 illustrates the monthly trend of the edge anomaly detection rate at the Mesaieed Plant between February 2024 and June 2025. The results indicate significant fluctuations in anomaly rates across the observed period, reflecting variations in production stability, equipment performance, and data integrity. Peaks in anomaly rates correspond to periods of operational disruptions, system irregularities, or process deviations, whereas lower values reflect more stable operating conditions. These findings underscore the critical role of real-time edge analytics in continuously monitoring process performance, detecting anomalies at an early stage, and enabling predictive responses that enhance overall operational reliability and resilience in smart manufacturing environments.

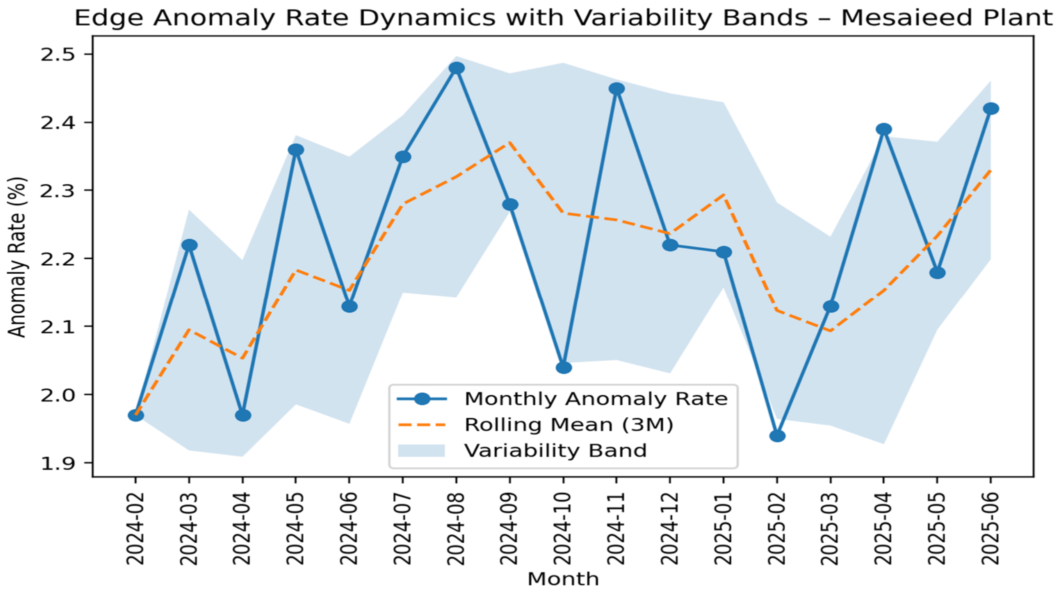

Figure 8. Monthly Edge Anomaly Rate Trend – Mesaieed Plant.

Figure 9 illustrates the relationship between edge latency and the anomaly detection rate at the Mesaieed Plant. The scatter plot shows a moderate positive correlation, where higher latency values are generally associated with increased anomaly rates. The variation in point density highlights periods of system instability, where delayed data processing on edge devices contributed to higher anomaly frequencies and elevated operational risks. Moreover, the size and color intensity of the markers represent the alerts count, indicating that clusters with higher latency often correspond to a larger number of anomaly alerts. These findings underline the importance of optimizing edge device performance to ensure efficient real-time data processing, minimize anomaly occurrence, and enhance operational reliability within smart manufacturing environments.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

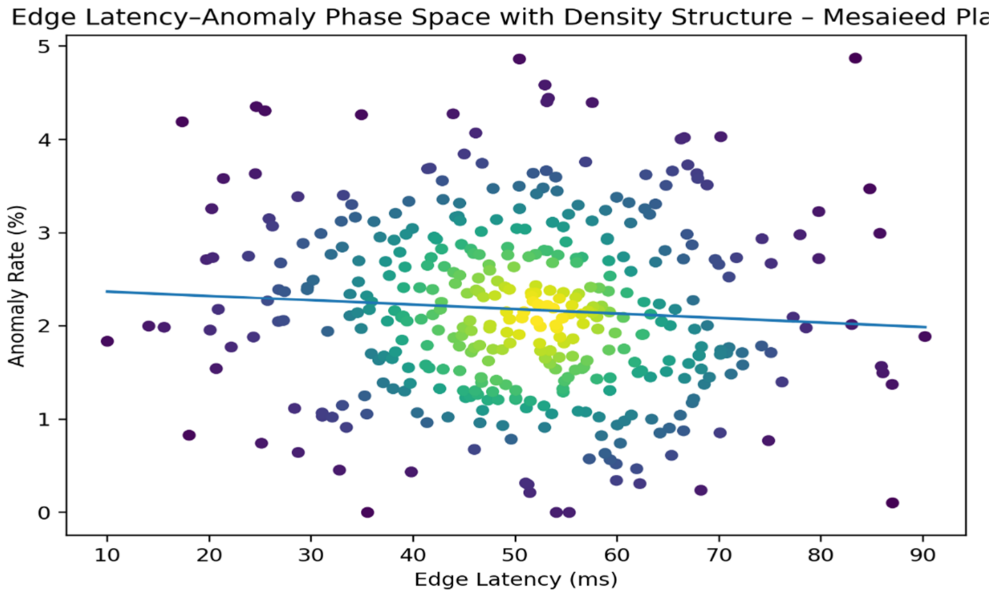

Figure 9. Relationship Between Edge Latency and Anomaly Rate – Mesaieed Plant.

Figure 10 presents the monthly trend of edge device uptime at the Mesaieed Plant between February 2024 and June 2025. The results demonstrate that edge devices maintain a generally high availability level, exceeding 98% throughout the observed period, which reflects the stability and reliability of the plant's data processing infrastructure. However, certain peaks and drops are observed, indicating periods of maintenance activities, network adjustments, or unexpected hardware interruptions. These findings highlight the critical role of maintaining consistent edge device performance to ensure continuous data flow, enhance real-time monitoring, and support predictive analytics for efficient risk management in smart manufacturing environments.

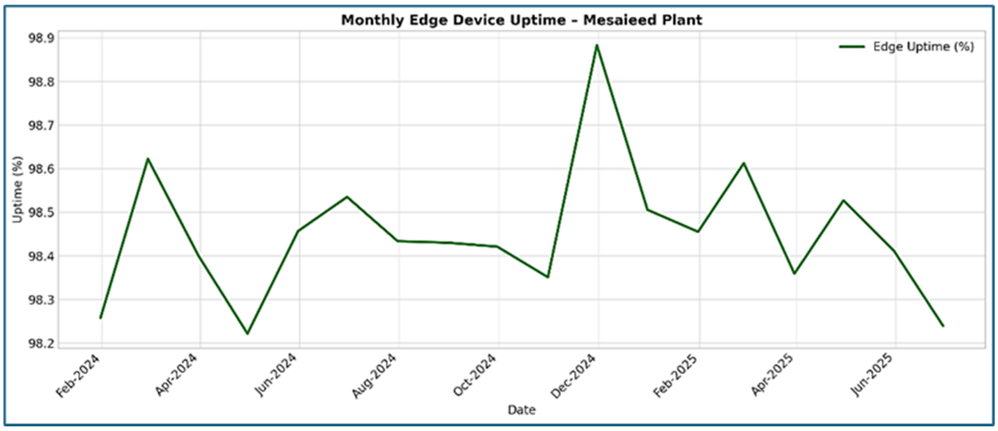

Figure 10. Monthly Edge Device Uptime Trend – Mesaieed Plant.

Figure 11 compares the monthly production volumes and actual demand between the Mesaieed Plant (PLANT_MES) and the Ras Laffan Plant (PLANT_RLF) from February 2024 to June 2025. The results indicate that production trends in both plants generally follow demand patterns, reflecting effective alignment between production planning and market needs. However, notable deviations are observed during certain peak periods, where demand temporarily exceeds production capacity, particularly in June 2024 and May 2025, suggesting periods of increased operational pressure and potential supply chain constraints. Overall, the figure highlights the interdependence between demand forecasting, production optimization, and operational efficiency, underscoring the

 doi:10.20944/preprints202603.0010.v1

importance of adopting predictive analytics to better synchronize production strategies with real-time demand fluctuations across both facilities.

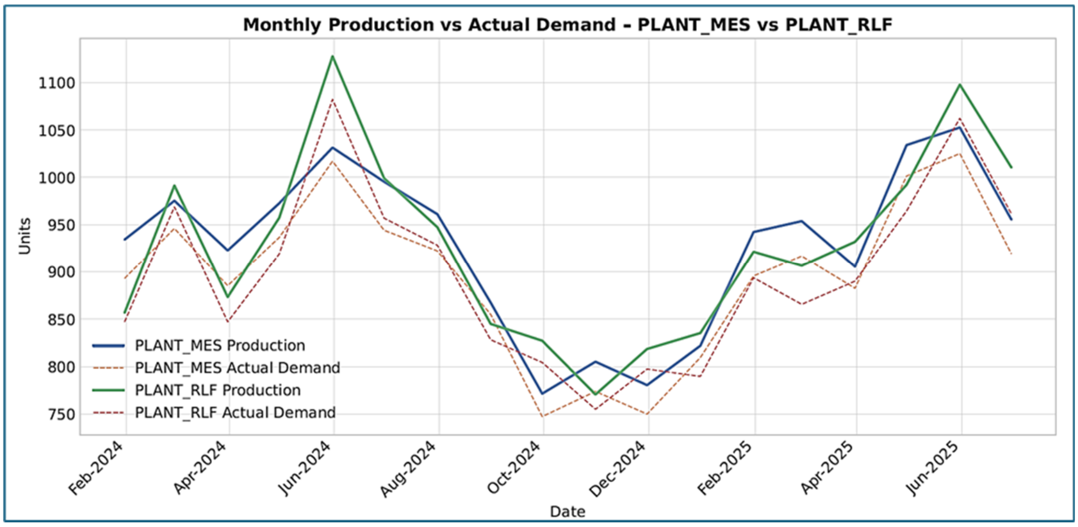

Figure 11. Monthly Production and Actual Demand Comparison – PLANT_MES vs. PLANT_RLF.

Figure 12 illustrates the monthly risk score trends for the Mesaieed Plant (PLANT_MES) and the Ras Laffan Plant (PLANT_RLF) from February 2024 to June 2025. The results indicate dynamic fluctuations in operational risk levels across both facilities, reflecting the impact of demand variability, supplier performance, and production stability on overall risk exposure. While both plants exhibit similar patterns, PLANT_MES demonstrates slightly higher peaks in several months, suggesting greater sensitivity to disruptions in supply chain operations and production processes. These variations highlight the importance of integrating real-time risk monitoring and predictive analytics to proactively identify vulnerabilities, minimize operational disruptions, and improve overall supply chain resilience across both production sites.

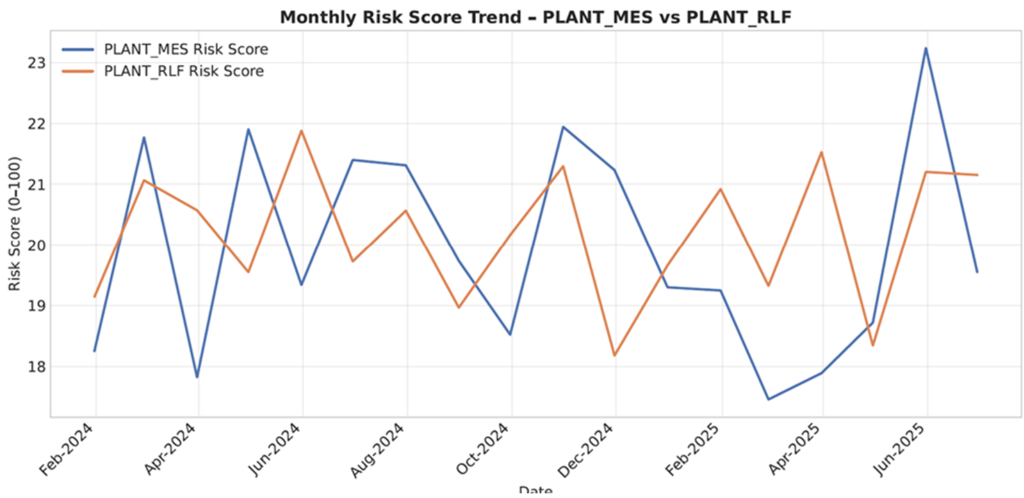

Figure 12. Monthly Risk Score Trends – PLANT_MES vs. PLANT_RLF.

Figure 13 presents the monthly trends of the resilience index for the Mesaieed Plant (PLANT_MES) and the Ras Laffan Plant (PLANT_RLF) between February 2024 and June 2025. The results demonstrate noticeable fluctuations in resilience levels across both facilities, reflecting their adaptive responses to operational disruptions and supply chain variability. Overall, both plants maintain relatively high resilience scores, with PLANT_MES showing slightly stronger recovery capabilities during several periods compared to PLANT_RLF. However, sudden declines in certain

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

months correspond to periods of increased production instability, supplier delays, and logistics challenges. These findings emphasize the importance of integrating real-time resilience monitoring and predictive analytics to enhance operational flexibility and ensure robust recovery strategies within smart manufacturing ecosystems.

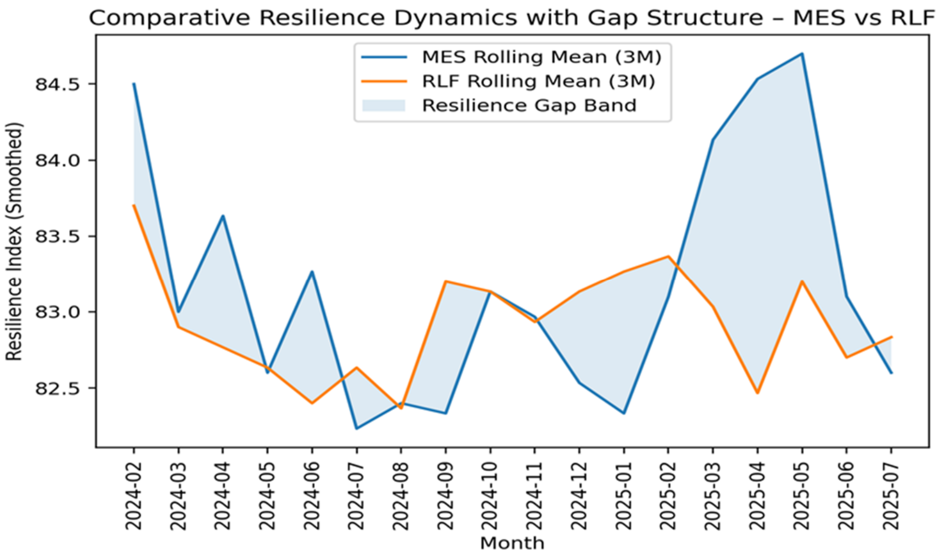

Figure 13. Monthly Resilience Index Trends – PLANT_MES vs. PLANT_RLF.

Figure 14 illustrates the monthly energy consumption for the Mesaieed Plant (PLANT_MES) and the Ras Laffan Plant (PLANT_RLF) from January 2024 to July 2025. The results show that overall energy usage remains relatively consistent across both facilities, with slight variations reflecting differences in production volumes and operational demands. While PLANT_MES exhibits marginally higher energy consumption in several months, PLANT_RLF demonstrates comparable performance, suggesting efficient energy management practices in both sites. The observed stability indicates the effectiveness of energy optimization strategies and highlights the role of smart manufacturing systems in balancing production requirements with energy efficiency objectives. These findings support the importance of continuous energy monitoring and integration of sustainable energy solutions to enhance operational performance and reduce overall energy costs.

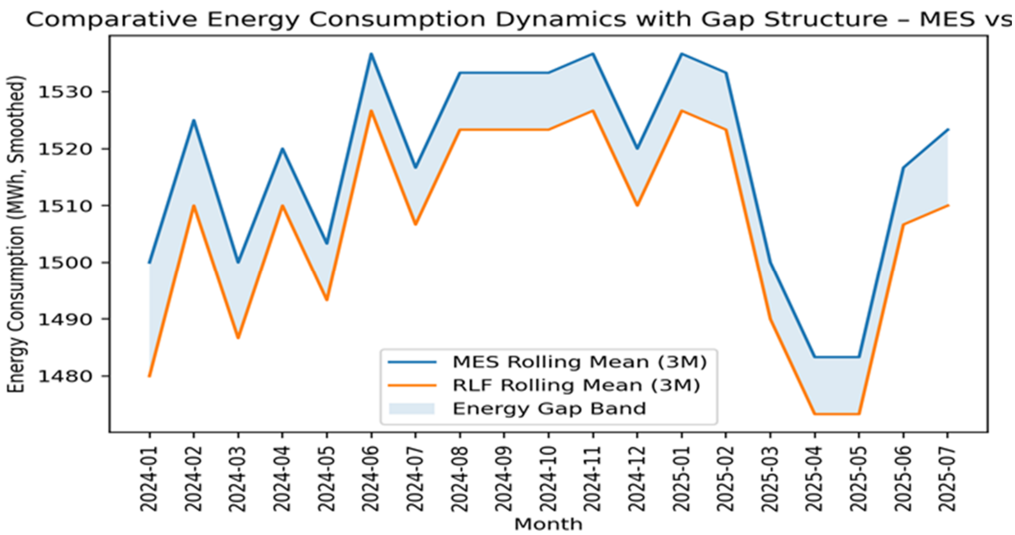

Figure 14. Monthly Energy Consumption Comparison – PLANT_MES vs. PLANT_RLF.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

Figure 15 compares the monthly logistics costs for the Mesaieed Plant (PLANT_MES) and the Ras Laffan Plant (PLANT_RLF) from January 2024 to July 2025. The results indicate significant variability in logistics expenses across both plants, influenced by fluctuations in demand, transportation requirements, and supply chain disruptions. While PLANT_MES generally incurs slightly higher logistics costs during peak production months, PLANT_RLF demonstrates better cost consistency, particularly in periods of lower operational activity. These variations reflect differences in transportation dependencies, port utilization, and supplier locations between the two facilities. The findings underscore the importance of adopting data-driven logistics optimization strategies to minimize operational costs, improve resource allocation, and enhance overall supply chain efficiency within smart manufacturing ecosystems.

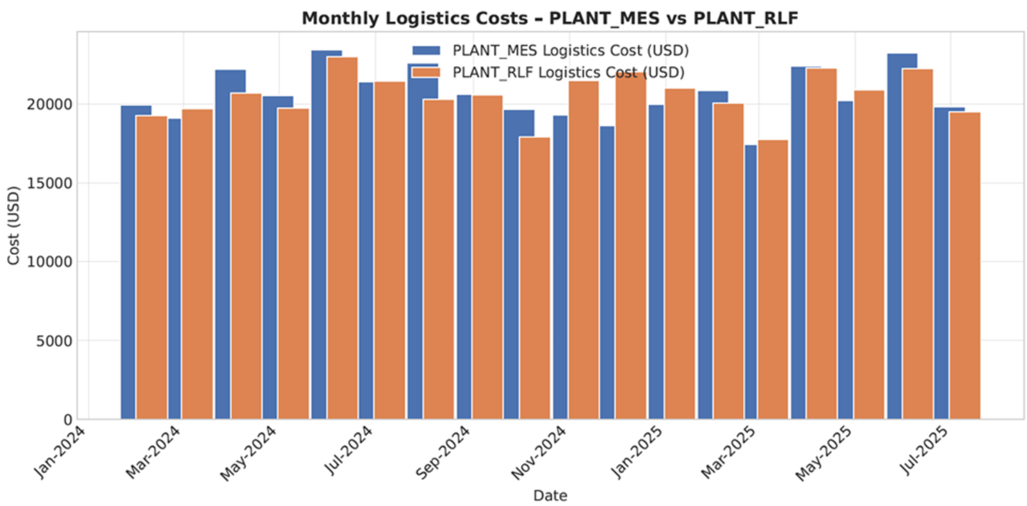

Figure 15. Monthly Logistics Costs Comparison – PLANT_MES vs. PLANT_RLF.

Figure 16 compares the monthly edge device uptime between the Mesaieed Plant (PLANT_MES) and the Ras Laffan Plant (PLANT_RLF) from February 2024 to June 2025. The results indicate that both plants maintain consistently high uptime levels, generally exceeding 98%, which reflects the reliability and robustness of their edge infrastructure. However, several peaks and dips are observed, highlighting short-term performance variations caused by planned maintenance, network adjustments, or temporary hardware limitations. While PLANT_RLF demonstrates slightly higher uptime stability in several periods, PLANT_MES records sharper improvements in others, indicating differences in operational strategies and equipment utilization. These findings emphasize the importance of continuous edge performance monitoring and predictive maintenance to ensure uninterrupted data processing, improve anomaly detection capabilities, and support real-time decision-making in smart manufacturing environments.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

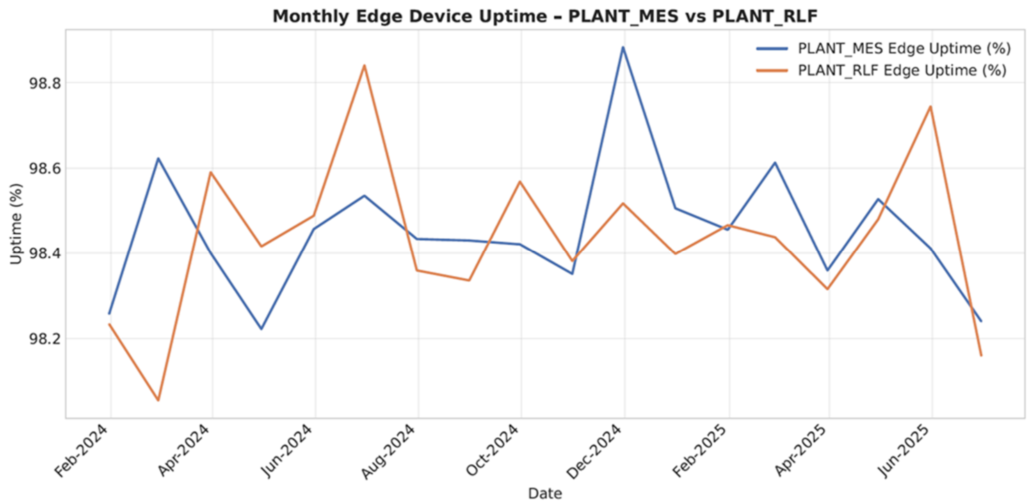

Figure 16. Monthly Edge Device Uptime Comparison – PLANT_MES vs. PLANT_RLF.

Figure 17 compares the monthly edge latency between the Mesaieed Plant (PLANT_MES) and the Ras Laffan Plant (PLANT_RLF) from February 2024 to June 2025. The results indicate noticeable variations in latency performance across both facilities, highlighting periods of fluctuating network stability and data processing efficiency. While both plants demonstrate comparable performance on average, PLANT_MES exhibits slightly higher latency peaks during certain months, suggesting potential processing delays under increased production loads. Conversely, PLANT_RLF maintains relatively lower latency in several periods, reflecting more consistent network optimization. These findings underscore the importance of improving edge infrastructure performance to ensure timely anomaly detection, enhance real-time decision-making, and support predictive analytics within smart manufacturing systems.

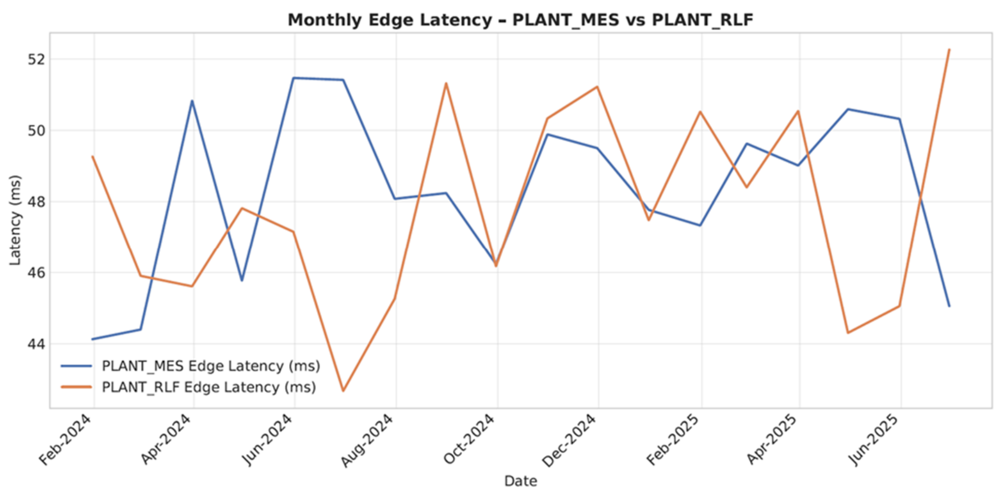

Figure 17. Monthly Edge Latency Comparison – PLANT_MES vs. PLANT_RLF.

Table 1 summarizes the key performance metrics used to evaluate the predictive accuracy and operational improvements achieved by the proposed data-driven risk management framework. The results demonstrate that the model achieves a MAPE of 3.67%, indicating a very high prediction accuracy with minimal deviation from actual risk values. The RMSE of 0.79 confirms that the absolute prediction errors remain small and within acceptable operational thresholds. Additionally, the model achieves an R^2 value of 0.989, meaning it explains approximately 98.9% of the variability in risk scores, which highlights the robustness and reliability of the predictive approach. Finally, the framework delivers a Resilience Gain of 3.99%, demonstrating its effectiveness in enhancing

 doi:10.20944/preprints202603.0010.v1

operational stability by reducing anomaly rates and improving the plant's ability to adapt to disruptions. These findings confirm the capability of the proposed approach to provide accurate forecasts while strengthening supply chain resilience and supporting real-time decision-making in smart manufacturing environments.

Table 1. summarizes the results obtained.

Metric	Value	Interpretation
MAPE (%)	3.67%	High prediction accuracy (low deviation from actual values)
RMSE	0.79	Small prediction error in absolute terms
R^2	0.989	The model explains ~98.9% of the variance in risk scores
Resilience Gain (%)	3.99%	Operational resilience improved when anomaly rates were reduced

Overall, the findings confirm the capability of the proposed framework to deliver high predictive accuracy in forecasting operational risks ($R^2 \approx 0.989$, MAPE ≈ 3.67%) while improving resilience by approximately 3.99% through effective anomaly detection and mitigation strategies. The results reveal the substantial impact of factors such as lead time variability, supplier performance, and edge device analytics on operational stability and supply chain continuity. Moreover, the integration of real-time monitoring, predictive risk modeling, and Edge Analytics demonstrates the framework's potential to optimize production planning, reduce disruption impacts, and enhance sustainability in global supply chains. These insights provide a solid empirical basis for the subsequent discussion, where the implications for theory, practice, and policy are examined in the context of smart manufacturing and resilient supply chain management.

5. Discussion

The results provide empirical support that the proposed edge-enabled pipeline can function as an AI-powered predictive maintenance and PHM mechanism rather than a generic operational analytics tool. The strong predictive performance (MAPE ≈ 3.67%, RMSE ≈ 0.79, $R^2 \approx 0.989$) indicates that the forecasting layer can reliably model the temporal evolution of condition signals and the associated health-risk dynamics, which is a prerequisite for maintenance decision support under streaming industrial data. In PdM/PHM terms, this level of explained variance is meaningful because it suggests that the model captures stable structure in the monitored operational behavior, enabling early identification of abnormal trajectories and reducing reliance on reactive maintenance triggered only after failure symptoms become obvious. A key observation from the empirical patterns is that edge-derived anomaly signals behave as early fault signatures that enrich prognostic inference. The positive association between anomaly rate and the risk indicator, together with the evidence that reducing anomalies improves the stability-related outcome (reported as a gain of about 3.99%), supports a PHM interpretation in which anomaly mitigation corresponds to improved operational health and recovery capability. (El Hammoumi, Z, 2025) Rather than treating anomalies as isolated outliers, the results imply that anomaly frequency and persistence can be used as structured inputs to the predictive layer to strengthen sensitivity to incipient degradation. This aligns with practical PdM settings where explicit fault labels may be limited, and where the most actionable signal is often a sustained deviation from normal behavior rather than a single classified fault event. The edge layer results also highlight that deployment variables are not secondary implementation details but part of the diagnostic performance envelope. The observed relationship between edge latency and anomaly rate suggests that processing delays can amplify the operational burden of anomaly alerts

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

and reduce the timeliness of fault indication. In PdM workflows, delayed anomaly detection can shift an intervention from "preventive" to "late," which undermines PHM effectiveness even when predictive accuracy is high. In contrast, the consistently high edge uptime (generally above 98%) supports the feasibility of continuous monitoring and reduces the likelihood of blind spots in the condition-monitoring stream. Together, these findings justify treating edge performance (latency, uptime, alert density) as measurable algorithm–system co-variables that should be monitored and optimized alongside model accuracy, especially for real-time industrial deployments. Across the two industrial sites, the comparative plots indicate that health-risk dynamics and edge behavior vary by operating context, which reinforces the need for scalable frameworks that remain stable under changing regimes. From an Algorithms perspective, the contribution is not only that the models forecast well, but that the study formalizes an end-to-end algorithmic workflow linking (i) condition monitoring, (ii) edge-level anomaly extraction, (iii) predictive health/risk trajectory estimation, and (iv) PHM-oriented indicators that can inform maintenance timing and prioritization. This end-to-end coupling addresses a common gap in the literature where anomaly detection, forecasting, and PHM decision logic are evaluated separately, making it difficult to operationalize results into maintenance actions. These results also clarify how "risk" and "resilience" should be interpreted in this paper for alignment with the Special Issue. In the current framework, the risk score can be treated as a health-risk proxy derived from condition monitoring and anomaly behavior, while the resilience-related measure functions as a stability/recovery indicator reflecting how quickly the system returns to acceptable operating patterns after deviations. This PHM-oriented interpretation keeps the constructs operational and algorithmically grounded and avoids positioning the discussion as supply-chain economics. It also supports clearer maintenance meaning: higher predicted health-risk trajectories imply higher urgency for inspection or intervention, while improved stability/recovery indicators imply more effective mitigation and reduced escalation probability.

Finally, the discussion must be interpreted with realistic boundaries. The framework demonstrates strong predictive reliability on the available industrial data and shows consistent links between edge anomaly behavior, latency, and health-risk indicators, but it does not explicitly estimate remaining useful life, and the health indicators rely on operational proxies rather than verified component-level failure labels.

These constraints do not weaken the PdM/PHM positioning; instead, they define the current maturity level of the contribution as an edge-enabled, data-driven health monitoring and prognostic assessment approach that is evaluation-ready under real operational variability, with a clear pathway for extension toward component-level prognostics when labeled failure histories become available.

Table 2. Structured Discussion Framework Aligned with AI-Powered Predictive Maintenance and PHM.

Discussion Dimension	What the Results Show	Algorithmic Interpretation	Relevance to PdM / PHM	Alignment with Algorithms Journal
Predictive accuracy and reliability	High predictive performance (low MAPE and RMSE, high R^2) across both industrial sites	Forecasting models successfully capture temporal dependencies and degradation-related patterns in operational data	Reliable prediction of health-state evolution supports early maintenance planning and reduces reactive interventions	Emphasizes algorithmic robustness and model validity rather than managerial outcomes

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 2 March 2026 doi:10.20944/preprints202603.0010.v1

Role of edge-level anomaly detection	Anomaly rates correlate with increased health-risk indicators	Edge-based anomaly detection acts as an early fault-signature extraction mechanism	Enables early-stage fault diagnosis before failure escalation	Highlights edge-based algorithms and low-latency diagnostics
Integration of edge and predictive layers	Edge outputs improve sensitivity of predictive models	Anomaly-derived features enrich the predictive input space	Strengthens prognostic inference under limited fault labels	Demonstrates end-to-end algorithmic integration
Prognostic health interpretation	Health-risk trajectories align with observed operational trends	Risk indicators function as proxies for system health degradation	Supports PHM-oriented decision-making without explicit RUL estimation	Keeps discussion within PHM scope accepted by Algorithms
Stability and recovery behavior	Measurable improvement in stability-related indicators after anomaly mitigation	Predictive control and early intervention improve system recovery capability	Reflects effectiveness of proactive maintenance strategies	Frames "resilience" as algorithmic health recovery, not economics
Edge performance (latency and uptime)	Lower latency and high uptime associated with reduced anomaly escalation	Deployment-level parameters influence diagnostic timeliness	Timely diagnosis is critical for PdM effectiveness	Reinforces relevance of algorithm–system co-design
Cross-site variability	Differences observed between industrial sites	Models remain stable under different operating regimes	Demonstrates adaptability of PdM framework	Supports scalability and generalizability of algorithms
Comparison with prior studies	Addresses integration gaps noted in literature	Unified pipeline links monitoring, diagnosis, and prognostics	Advances beyond isolated anomaly or forecasting studies	Directly responds to gaps emphasized in Algorithms special issue
Practical implications (technical)	Outputs are interpretable	Algorithmic outputs can be embedded in	Facilitates deployment-	Focus on implementation of

 doi:10.20944/preprints202603.0010.v1

	health indicators	real-time PdM systems	ready PHM systems	algorithms, not management
Methodological limitations	No explicit RUL estimation; reliance on operational proxies	Framework emphasizes health-state trajectories over component failure labels	Appropriate for data-scarce industrial environments	Transparent discussion of algorithmic boundaries

6. Theoretical & Practical Implications

Theoretical Implications

From a theoretical perspective, this study contributes to the literature on AI-powered predictive maintenance and prognostic health management by formalizing an integrated algorithmic pipeline that links data-driven condition monitoring, intelligent fault diagnosis, and prognostic assessment within an edge-enabled environment. Unlike prior studies that focus on isolated components such as anomaly detection or forecasting accuracy, the proposed framework emphasizes the algorithmic coupling between edge-level anomaly extraction and predictive health-state modeling. This integration advances PHM theory by demonstrating how anomaly-derived signals can be systematically transformed into health-oriented indicators that capture degradation dynamics under real operational variability. Furthermore, the study extends existing PHM and PdM research by framing risk-related predictions as health-risk trajectories rather than purely statistical outputs, thereby strengthening the theoretical connection between predictive modeling and maintenance-oriented decision support in industrial systems.

Practical Implications

From a practical and technical standpoint, the findings provide clear guidance for the design and deployment of predictive maintenance systems under industrial constraints. The results show that performing anomaly detection at the edge enables low-latency fault indication and reduces the dependency on continuous centralized processing, which is critical for time-sensitive industrial applications. By combining edge-based screening with centralized predictive modeling, the proposed framework offers a scalable deployment strategy that balances computational efficiency and diagnostic accuracy. For engineers and system designers, the framework highlights the importance of treating edge performance metrics, such as latency and uptime, as integral components of the predictive maintenance pipeline rather than secondary implementation details. The health-oriented indicators produced by the prognostic layer can be directly integrated into real-time monitoring platforms to support maintenance prioritization, early intervention, and system stability assessment, providing a practical foundation for implementing AI-driven PdM and PHM solutions in dynamic industrial environments.

7. Conclusion

This study proposed an AI-powered, edge-enabled algorithmic framework for predictive maintenance and prognostic health management in industrial operations. By integrating data-driven condition monitoring, edge-level anomaly detection, predictive modeling, and health-oriented prognostic assessment, the framework was designed to support intelligent fault diagnosis and forward-looking maintenance decisions under real industrial operating conditions. The empirical evaluation using high-frequency data from two industrial plants demonstrated that the proposed approach achieves high predictive reliability, confirming its ability to capture temporal patterns associated with system degradation and abnormal behavior.

 doi:10.20944/preprints202603.0010.v1

The results show that anomaly signals extracted at the edge can be effectively leveraged as early fault signatures and integrated into predictive models to enhance sensitivity to incipient degradation. This integration enables the estimation of health-risk trajectories that provide actionable insight into system conditions without requiring explicit failure labels or remaining useful life estimation. The observed improvements in stability- and recovery-related indicators further indicate that early anomaly mitigation contributes to maintaining consistent operational behavior, reinforcing the relevance of the framework for prognostic health management applications.

From an algorithmic perspective, the main contribution of this work lies in formalizing an end-to-end predictive maintenance pipeline that couple's edge-aware diagnostics with predictive health assessment in a coherent and scalable manner. Rather than treating deployment considerations as secondary, the framework explicitly incorporates edge performance characteristics, such as latency and availability, as part of the overall diagnostic process. This perspective aligns with the growing need for deployment-ready AI solutions in industrial environments characterized by high data rates and strict timing constraints.

While the proposed framework demonstrates strong performance and practical feasibility, the study is subject to certain limitations. The health indicators are derived from operational proxies rather than component-level failure labels, and the framework does not explicitly estimate remaining useful life. These limitations define clear directions for future research, including the incorporation of component-specific degradation models, adaptive learning under evolving operating regimes, and tighter integration with digital representations of industrial assets. Overall, the findings confirm that edge-enabled predictive algorithms provide a robust foundation for AI-powered predictive maintenance and PHM, and they offer a viable pathway for advancing intelligent fault diagnosis in real-world industrial systems.

8. Limitations & Future Research

Despite the promising results obtained in this study, several limitations should be acknowledged to provide a balanced interpretation of the findings and to delineate directions for future research. First, the proposed framework relies on operational condition signals and anomaly-derived indicators as proxies for system health, rather than on explicit component-level failure labels. While this approach reflects practical industrial constraints where labeled failure data are often scarce, it limits the ability to directly estimate remaining useful life or to validate prognostic outputs against known failure events. Future studies could extend the framework by incorporating labeled degradation histories or hybrid physical–data-driven models to enable more granular component-level prognostics.

Second, the predictive models were trained and evaluated using data from two industrial plants operating under specific process configurations and operating regimes. Although the results demonstrate stable performance across these contexts, the generalizability of the framework to other industrial domains with different asset types, sensing infrastructures, or control strategies remains to be systematically investigated. Future research could explore transfer learning, domain adaptation, or federated learning strategies to enhance model robustness under cross-plant and cross-domain variability.

Third, the current implementation treats edge-level anomaly detection and predictive modeling as sequential stages with fixed algorithmic configurations. While this design supports clarity and deployment feasibility, it does not fully exploit adaptive or co-learning mechanisms between the edge and predictive layers. Future work could investigate adaptive edge intelligence, where anomaly thresholds, feature extraction, or model parameters evolve in response to changing operating conditions, thereby improving long-term stability and responsiveness.

Finally, although the framework explicitly accounts for edge performance characteristics such as latency and availability, these factors are not yet optimized as part of the learning objective. Future research could integrate edge-aware optimization criteria, including energy consumption, computational load, and communication reliability, into the predictive maintenance pipeline. Such

 doi:10.20944/preprints202603.0010.v1

extensions would further strengthen the framework's suitability for large-scale industrial deployment and contribute to the development of fully autonomous, resource-aware PdM and PHM systems.

Author Contributions: Conceptualization, L.B.; methodology, L.B.; data curation, L.B.; formal analysis, L.B.; algorithm development, L.B.; software, L.B.; validation, L.B.; investigation, L.B.; writing—original draft preparation, L.B.; writing—review and editing, L.B.; visualization, L.B.; supervision, L.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data used in this study are not publicly available due to confidentiality constraints associated with industrial operations and data-sharing agreements with the participating plants. Aggregated data and supporting materials may be made available from the corresponding author upon reasonable request, subject to approval from the data providers.

Conflicts of Interest: The authors declare no conflict of interest.

Acknowledgments : During the preparation of this manuscript, the author used ChatGPT (OpenAI, GPT-5.2) for the purposes of language refinement and structural editing. The author has reviewed and edited the output and takes full responsibility for the content of this publication.

References

1 Ciobotaru, A., Corches, C., Gota, D., & Miclea, L. (2025). An Explainable Deep Learning-Based Predictive Maintenance Solution for Air Compressor Condition Monitoring. Sensors, 25(18), 5797.
2 Uçar, A., Karaköse, M., & Kırımça, N. (2024). Artificial Intelligence for Predictive Maintenance Applications: Key Components, Trustworthiness, and Future Trends. Applied Sciences, 14(2), 898.
3 El Hammoumi, Z. (2025). Predictive Maintenance Approaches: A Systematic Review of Diagnostics and Prognostics Models. (Predictive Maintenance Journal, MDPI).
4 Bitam, T. et al. (2025). Artificial Intelligence of Things for Next-Generation Predictive Maintenance. Sensors, 25(24), 7636.
5 Abbott, L. (1991), Future of health insurance. New England Journal of Medicine, 965, 325–329.
6 Lee, J.; Kao, H.-A.; Yang, S. Service innovation and smart analytics for industry 4.0 and big data environment. CIRP Annals – Manufacturing Technology 2014, 63(2), 707–732.https://doi.org/10.1016/j.cirp.2014.02.001
7 Aven, T. (2011). On some recent definitions and analysis frameworks for risk, vulnerability and resilience. Risk Analysis, 31(4), 515–522. https://doi.org/10.1111/j.1539-6924.2010.01567.x
8 Bagheri, B., & Lee, J. (2015). A CPS architecture for Industry 4.0 manufacturing systems. Manufacturing Letters, 3, 18–23. https://doi.org/10.1016/j.mfglet.2014.12.001
9 Baryannis, G., Dani, S., & Antoniou, G. (2019). Predictive analytics and artificial intelligence in supply chain management: Review and implications for the future. Computers & Industrial Engineering, 137, 106024. https://doi.org/10.1016/j.cie.2019.106024
10 Bertsimas, D., & Kallus, N. (2014). From predictive to prescriptive analytics. arXiv preprint. https://doi.org/10.48550/arXiv.1402.5481
11 Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys, 41(3), 1–58. https://doi.org/10.1145/1541880.1541882
12 Choi, T.-M., Wallace, S. W., & Wang, Y. (2018). Big data analytics in operations management. Production and Operations Management, 27(10), 1868–1883. https://doi.org/10.1111/poms.12838
13 Christopher, M., & Peck, H. (2004). Building the resilient supply chain. International Journal of Logistics Management, 15(2), 1–13. https://doi.org/10.1108/09574090410700275
14 Craighead, C. W., Blackhurst, J., Rungtusanatham, M. J., & Handfield, R. B. (2007). The severity of supply chain disruptions: Design characteristics and mitigation capabilities. Decision Sciences, 38(1), 131–156. https://doi.org/10.1111/j.1540-5915.2007.00151.x

15 Dubey, R., Gunasekaran, A., Childe, S. J., Bryde, D., Giannakis, M., Foropon, C., & Papadopoulos, T. (2018). Big data analytics and supply chain agility: The moderating effect of organizational flexibility. International Journal of Production Research, 56(1–2), 444–462. https://doi.org/10.1080/00207543.2017.1343500

16 Fan, Y., & Stevenson, M. (2018). A review of supply chain risk management: Definition, theory, and research agenda. International Journal of Production Research, 56(1–2), 898–916. https://doi.org/10.1080/00207543.2017.1402145

17 Gunasekaran, A., Papadopoulos, T., Dubey, R., Wamba, S. F., Childe, S. J., Hazen, B., & Akter, S. (2017). Big data and predictive analytics for supply chain and organizational performance. Journal of Business Research, 79, 344–354. https://doi.org/10.1016/j.jbusres.2017.10.008

18 Hendricks, K. B., & Singhal, V. R. (2005). Association between supply chain glitches and operating performance. Management Science, 51(5), 695–711. https://doi.org/10.1287/mnsc.1040.032

19 Hohenstein, N.-O., Feisel, E., Hartmann, E., & Giunipero, L. (2015). Research on the phenomenon of supply chain resilience: A systematic review and conceptual framework. International Journal of Production Economics, 173, 124–142. https://doi.org/10.1016/j.ijpe.2015.02.036

20 Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of quantitative methods for supply chain resilience analysis. Omega, 87, 231–259. https://doi.org/10.1016/j.omega.2018.08.028

21 Ivanov, D., & Dolgui, A. (2020). Viability of intertwined supply networks: Extending the supply chain resilience angles toward survivability. International Journal of Production Research, 58(10), 2904–2915. https://doi.org/10.1080/00207543.2019.1630801

22 Ivanov, D., Dolgui, A., Sokolov, B., Ivanova, M., & Rozhkov, M. (2016). A dynamic model for managing disrupted supply chains. International Journal of Production Research, 54(1), 1–19. https://doi.org/10.1080/00207543.2015.1072719

23 Jüttner, U. (2005). Supply chain risk management: Understanding risk and vulnerability in the supply chain. International Journal of Logistics Management, 16(1), 120–141. https://doi.org/10.1108/09574090510617385

24 Kang, H.-S., et al. (2016). Smart manufacturing: Past research, present findings, and future directions. International Journal of Precision Engineering and Manufacturing-Green Technology, 3(1), 111–128. https://doi.org/10.1007/s40684-016-0015-5

25 Klibi, W., Martel, A., & Guitouni, A. (2010). The design of robust value-creating supply chain networks: A critical review. European Journal of Operational Research, 203(2), 283–293. https://doi.org/10.1016/j.ejor.2009.08.009

26 Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In NeurIPS (pp. 4765–4774).

27 Pettit, T. J., Fiksel, J., & Croxton, K. L. (2010). Ensuring supply chain resilience: Development of a conceptual framework. Journal of Business Logistics, 31(1), 1–21. https://doi.org/10.1002/j.2158-1592.2010.tb00125.x

28 Ponomarov, S. Y., & Holcomb, M. C. (2009). Understanding the concept of supply chain resilience. International Journal of Logistics Management, 20(1), 124–143. https://doi.org/10.1108/09574090910954873

29 Satyanarayanan, M. (2017). The emergence of edge computing. Computer, 50(1), 30–39. https://doi.org/10.1109/MC.2017.9

30 Sheffi, Y., & Rice, J. B. J. (2005). A supply chain view of the resilient enterprise. MIT Sloan Management Review, 47(1), 41–48.

31 Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637–646. https://doi.org/10.1109/JIOT.2016.2579198

32 Snyder, L. V., & Shen, Z.-J. M. (2019). Fundamentals of Supply Chain Theory (2nd ed.). Wiley.

33 Sroufe, R. P., & Curkovic, S. (2008). An examination of ISO 9000:2000 and supply chain quality assurance. Journal of Operations Management, 26(4), 503–520. https://doi.org/10.1016/j.jom.2007.12.009

34 Tang, C. S. (2006). Robust strategies for mitigating supply chain disruptions. International Journal of Logistics: Research and Applications, 9(1), 33–45. https://doi.org/10.1080/13675560500405584

35 Tang, C. S., & Tomlin, B. (2008). The power of flexibility for mitigating supply chain risks. Management Science, 54(5), 1097–1111. https://doi.org/10.1287/mnsc.1070.0813

doi:10.20944/preprints202603.0010.v1

36 Taylor, S. J., & Letham, B. (2018). Forecasting at scale. The American Statistician, 72(1), 37–45. https://doi.org/10.1080/00031305.2017.1380080
37 Wamba, S. F., Gunasekaran, A., Akter, S., Ren, S. J. F., Dubey, R., & Childe, S. J. (2017). Big data analytics and firm performance: Effects of dynamic capabilities. Journal of Business Research, 70, 356–365. https://doi.org/10.1016/j.jbusres.2016.08.004
38 Wieland, A., & Wallenburg, C. M. (2013). The influence of relational competencies on supply chain resilience. International Journal of Physical Distribution & Logistics Management, 43(4), 300–320. https://doi.org/10.1108/IJPDLM-08-2012-0247
39 Xu, L. D., Xu, E. L., & Li, L. (2018). Industry 4.0: State of the art and future trends. International Journal of Production Research, 56(8), 2941–2962. https://doi.org/10.1080/00207543.2018.1444806