



Article

Optimizing Predictive and Prescriptive Maintenance Using Unified Namespace (UNS) for Industrial Equipments

Renjithkumar Surendran Pillai ^{1,*} , Patrick Denny ² , Eoin O'Connell ³ , Adam Dooley ³ and Mihai Penica ³

¹ Faculty of Science and Engineering, University of Limerick, V94 T9PX Limerick, Ireland

² Department of Computer Science and Information Systems (CSIS), D2ICE Research Centre, Faculty of Science and Engineering, University of Limerick, V94 T9PX Limerick, Ireland; patrick.denny@ul.ie

³ Department of Electronics and Computer Engineering (E&CE), Faculty of Science and Engineering, University of Limerick, V94 T9PX Limerick, Ireland; eoin.oconnell@ul.ie (E.O.); adam.dooley@ul.ie (A.D.); mihai.penica@ul.ie (M.P.)

* Correspondence: surendranpillai.renjithkumar@ul.ie

Abstract

This paper proposes a new Unified Namespace (UNS)-based architecture to improve predictive and prescriptive maintenance of industrial equipment and overcome challenges such as incomplete data, poor interoperability, and disconnected IT/OT environments. The framework combines interoperable data formats in real-time sensor data, predictive modeling, prescriptive analytics, and simulations of digital twins, using UNS as a centralized, protocol-agnostic data layer that is scalable and complies with Industry 4.0 and Pharma 4.0 standards. The suggested methodology increases data accessibility, reduces integration complexity, and allows low-latency analytics and automated decision-making. Machine learning predictive models achieved more than 94% accuracy in predicting equipment failures. Prescriptive analytics provides maintenance recommendations to reduce downtime and risks. The feedback loops of digital twins can enhance the accuracy of predictions and allow decision optimization through what-if analysis. A test-bench deployment showed a higher performance compared to traditional point-to-point integration, with lower latency (approximately 18 ms vs. approximately 31 ms), decreasing packet loss (0.40% vs. 3.11%), and higher model accuracy (94.20% vs. 87.51%). The structure avoided more than 4000 simulated breakdowns in the test-bench environment, indicating dependability. The study connects the theoretical applications of the UNS with the actual maintenance processes and provides a sound approach to the industrial analytics and optimization of the equipment.

Keywords: Unified Namespace; predictive maintenance; prescriptive maintenance; digital twin



Academic Editors: Marco Rossi and Shuhui Li

Received: 12 February 2026

Revised: 11 March 2026

Accepted: 13 March 2026

Published: 19 March 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Industrial maintenance strategies have evolved from reactive and preventive approaches toward data-driven predictive and prescriptive methodologies enabled by Industry 4.0 technologies [1]. Predictive maintenance (PdM) identifies early degradation using real-time sensing and machine learning, while prescriptive maintenance (RxM) recommends optimal interventions to mitigate failure risks [2]. Despite these advances, practical deployment remains constrained by fragmented data sources, heterogeneous protocols, and persistent IT/OT separation. These limitations hinder closed-loop decision-making, digital-twin validation, and scalable analytics [3].

The Unified Namespace (UNS) has emerged as a promising architectural solution, offering a protocol-agnostic, hierarchical, real-time data layer [4]. However, the existing literature primarily discusses UNS conceptually, without demonstrating how it can operationalize PdM/RxM workflows or support digital-twin-based decision validation in real environments.

This study addresses these gaps by presenting the first experimentally validated UNS-based predictive–prescriptive maintenance architecture. The key contributions are as follows:

- A fully implemented UNS architecture integrating MQTT, OPC-UA, Modbus, predictive models, and a digital twin into a unified workflow.
- Quantitative comparison of UNS vs. point-to-point (P2P) communication, demonstrating lower latency (18 ms vs. 31 ms), reduced packet loss (0.40% vs. 3.11%), and improved predictive accuracy (94.20% vs. 87.51%).
- A digital-twin feedback loop that validates predictions, monitors drift and evaluates prescriptive actions before deployment.
- A reproducible methodology linking raw sensor data, UNS ingestion, predictive modeling, prescriptive optimization, and simulation-based validation.

While UNS has been discussed in prior work as a conceptual integration pattern for Industry 4.0 [5,6], existing studies do not demonstrate how UNS can operationally support predictive modeling, prescriptive analytics, or digital-twin-based decision validation.

The novelty of this work lies in

- (i) implementing a fully functional UNS-based PdM/RxM workflow,
- (ii) experimentally validating it on a physical test bench, and
- (iii) integrating a digital-twin feedback loop directly into the UNS for real-time prescriptive action verification. To our knowledge, this is the first empirical demonstration of a UNS-driven predictive–prescriptive maintenance architecture evaluated within a controlled test-bench environment, providing an initial but important step toward validating UNS-enabled PdM/RxM workflows in real industrial settings.

2. Methodology

This study adopts a system-based experimental methodology using a controlled electromechanical test bench designed to emulate the behavior of a small industrial pump or motor. The objective of the setup is to generate realistic operational data under both normal and fault-induced conditions and to evaluate how a Unified Namespace (UNS) can support predictive and prescriptive maintenance workflows [7]. As noted earlier in the manuscript, the test bench continuously streams vibration, temperature, pressure, and power measurements, and “the UNS receives this data in real time.” The methodology integrates data acquisition, preprocessing, predictive modeling, digital-twin simulation, and prescriptive decision optimization into a unified and reproducible workflow.

2.1. Test Bench and Data

The test bench consists of a motor-driven rotating shaft instrumented with a three-axis accelerometer, a temperature probe positioned near the bearing assembly, a pressure transmitter installed in a closed pneumatic loop, and an electrical power meter measuring voltage, current, and real power. These sensors were selected to reflect common industrial monitoring points in rotating machinery [8]. All sensors operate at a sampling rate of 100 Hz, resulting in approximately 33,000 samples per sensor for each experimental run. To evaluate the robustness of the proposed maintenance architecture, controlled disturbances were introduced into the system. Mechanical imbalance was created by adding mass to the rotating shaft, thermal overload was induced by restricting airflow around the

motor housing, and pressure restriction was simulated by partially closing a valve in the pneumatic line. These interventions produced realistic degradation patterns while maintaining safe operating conditions [9].

2.2. Data Preprocessing

Once acquired, the raw sensor data undergoes a structured preprocessing pipeline to ensure quality and consistency before entering the predictive modeling stage [10]. High-frequency noise is reduced using a five-point moving average filter, which smooths the signal while preserving transient behavior. Outliers caused by electrical interference or sensor jitter are identified using a z-score threshold of $|z| > 3$ and removed to prevent distortion of the learning process. All features are then normalized using min–max scaling to ensure compatibility across variables with different magnitudes. The dataset is subsequently divided into training, validation, and testing subsets using a 70–15–15 split. The held-out test set contains 5000 samples and is used exclusively for final evaluation. Fault and non-fault samples were balanced to avoid bias in the predictive model [11].

2.3. Predictive Models

Two lightweight machine learning models were selected based on their suitability for real-time deployment within a UNS-based architecture. Isolation Forest was used for anomaly detection, particularly for identifying abnormal temperature and vibration patterns. Its low computational overhead and ability to model normal behavior without requiring labeled fault data make it appropriate for edge-level execution. Logistic Regression was employed to estimate the probability of imminent failure using multivariate sensor inputs. Both models were trained using Scikit-learn and TensorFlow. The training process employed the Adam optimizer with a learning rate of 0.001, running for 50 epochs with a batch size of 64. The failure-score output was trained using a mean-squared error loss function. More complex models such as LSTM, GRU, and XGBoost were evaluated during preliminary experiments but were excluded due to higher inference latency and reduced suitability for real-time UNS integration [12].

2.4. Digital Twin

A digital twin of the test bench was developed using a discrete-time state-space representation. The model describes the system dynamics through the equations

$$x_{t+1} = Ax_t + Bu_t + w_t$$

and

$$y_t = Cx_t + v_t$$

where x_t denotes the internal system state, u_t represents the input vector derived from sensor measurements, and y_t is the predicted output. The matrices A , B , and C were identified from baseline operational data using system identification techniques, while w_t and v_t represent process and measurement noise respectively [13]. The fidelity of the digital twin was evaluated using the simulation error metric

$$E_{sim} = |y_{actual} - y_{simulated}|.$$

As reported in the manuscript, the simulation error remained below 0.05 across vibration, pressure, and temperature signals, demonstrating that the twin accurately replicates real system behavior and can safely validate prescriptive actions before they are applied to the physical system.

2.5. Prescriptive Optimization

Prescriptive maintenance decisions are generated by solving an optimization problem that balances intervention cost and failure risk. The decision rule is expressed as

$$\min_{a \in A} C(a) + \lambda R(a)$$

where $C(a)$ denotes the cost associated with a maintenance action, $R(a)$ represents the risk of failure if no action is taken, and λ is a weighting factor that determines the trade-off between cost and reliability [14]. The action set includes controlled shutdowns, speed reductions, and scheduled maintenance interventions. Each candidate action is evaluated within the digital-twin environment to ensure that it reduces risk without introducing instability or unintended consequences.

2.6. Definition of “4000 Avoided Failures”

The reported figure of 4000 avoided failures corresponds to simulated high-risk events detected during the test-bench experiments. A failure event is counted when the predictive model outputs a failure probability greater than 0.8, the digital twin confirms that the predicted state would lead to unsafe operation, and the prescriptive module identifies an action that reduces the risk below the safety threshold. As clarified in the manuscript, “These 4000 prevented failures refer to simulated high-risk events generated within the controlled test-bench environment.” They do not represent real industrial failures but rather demonstrate the system’s ability to detect and mitigate hazardous conditions in a controlled and repeatable manner.

2.7. Experimental Design Summary

The experimental campaign consisted of multiple controlled runs on the electromechanical test bench, each lasting approximately 330 s and generating ~33,000 samples per sensor (100 Hz sampling) per run, as described in the manuscript (“All sensors operate at a sampling rate of 100 Hz, resulting in approximately 33,000 samples per sensor for each experimental run.”). Fault conditions were injected systematically through mechanical imbalance, thermal overload, and pressure restriction to ensure balanced representation of normal and degraded states. The predictive models were trained using a 70/15/15 split, with 5000 held-out test samples (“The held-out test set contains 5000 samples and is used exclusively for final evaluation.”). Hyperparameters included a learning rate of 0.001, 50 epochs, and batch size of 64 for the Logistic Regression model, while the Isolation Forest used default contamination-based anomaly scoring. This consolidated design ensures reproducibility and clarifies the scope of the experimental evaluation.

3. System Design

3.1. Research Design

The system is designed to collect real-time operational data from industrial equipment and process it through a Unified Namespace (UNS). A controlled test bench is used to generate vibration, temperature, pressure, and power measurements. These measurements are streamed into the UNS, which acts as the central integration layer. Predictive models and a digital twin consume this data to support maintenance decisions [15].

Predictive modeling is formulated as a supervised learning task, where sensor inputs x_t are mapped to a failure score y_t :

$$\hat{y}_t = f_{\theta}(x_t)$$

where x_t represents the sensor input vector at time t , y_t is the predicted failure score, and $f_{\theta}(\cdot)$ denotes the predictive model parameterized by θ .

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n (y_i - f_{\theta}(x_i))^2$$

where $\mathcal{L}(\theta)$ is the mean-squared error loss, n is the number of training samples, and y_i and x_i are the true label and input vector for sample i .

3.2. Unified Namespace (UNS): Data Architecture for Industry 4.0

The UNS is implemented as a protocol-agnostic, hierarchical data layer that connects sensors, PLCs, SCADA, MES, ERP, and analytical applications [16]. Figure 1 shows the UNS architecture. All incoming data streams are normalized for timestamps, formats, and naming conventions to ensure consistency across the system.




Figure 1. Unified Namespace (UNS) architecture.

Communication performance is monitored using throughput:

$$T = \frac{N}{\Delta t}$$

where T is throughput, N is the number of packets received, and Δt is the measurement interval.

3.3. Implementation of UNS

The UNS is deployed on a layered architecture linking field devices, communication protocols, and analytics engines [17]. Figure 2 shows the workflow for publishing data from industrial devices. MQTT, OPC-UA, and Modbus are used for communication, and each data point is assigned a unique namespace identifier for traceability [18].

Latency is measured as

$$trvb\ cL = t_{\text{received}} - t_{\text{sent}}$$

where L is communication latency, t_{sent} is the timestamp when the packet was published, and t_{received} is the timestamp when it was consumed.




Figure 2. UNS implementation workflow.

3.4. Industrial Data Sources and Communication Protocols

The system collects data from vibration sensors, temperature probes, pressure transmitters, power meters, PLCs, and SCADA systems [19]. MQTT supports lightweight, low-latency messaging (Figure 3). OPC-UA provides structured and secure PLC/SCADA communication, while Modbus and BACnet support legacy devices [20].




Figure 3. MQTT communication protocol.

Protocol performance is evaluated using latency and packet integrity to ensure reliable data for analytics.

3.5. Data Processing and Analysis

After ingestion into the UNS, raw sensor data are standardized and cleaned. Noise is reduced using a moving average filter, and outliers are detected using the z-score:

$$z_t = \frac{x_t - \mu}{\sigma}$$

where z_t is the standardized value, x_t is the raw measurement, μ is the mean, and σ is the standard deviation.

Historical and real-time data are used for predictive modeling. Model stability is monitored using drift:

$$D_t = | \hat{y}_t - \hat{y}_{t-1} |$$

where D_t is the drift value at time t and y_t and y_{t-1} are consecutive model outputs.

This ensures consistent behavior under varying operating conditions.

The dataset was divided into 70% training, 15% validation, and 15% testing. The 5000 samples used for evaluation correspond to the held-out test set [21]. All results were computed using unseen data to ensure unbiased evaluation. Fault conditions were injected into real sensor streams to ensure balanced representation of failure and non-failure states.

3.6. Digital-Twin Development and Validation

A digital twin is implemented to simulate equipment behavior using a discrete-time state-space model:

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + w_t \\ y_t &= Cx_t + v_t \end{aligned}$$

where x_t is the system state, u_t is the input vector, y_t is the output, A , B , and C are system matrices, and w_t and v_t represent process and measurement noise. Figure 4 shows the digital-twin environment.



Figure 4. Digital-twin representation.

Simulation accuracy is evaluated using

$$E_{\text{sim}} = | y_{\text{actual}} - y_{\text{simulated}} |$$

where E_{sim} is the simulation error, y_{actual} is the real measurement, and $y_{\text{simulated}}$ is the digital-twin output.

The twin validates predictive outputs and tests prescriptive actions before applying them to the real system.

3.7. Core Technologies

The system integrates IoT sensors, machine learning models (TensorFlow, Scikit-learn), 3D visualization tools, distributed storage, and cloud services [22,23]. Table 1 lists the predictive models and their applications.

Table 1. Predictive models and application.

Model Type	Application	Algorithm Used	Data
Anomaly Detection	Temperature Abnormalities	Isolation Forest	Scikit-learn
Predictive Maintenance	Pressure Drop Prediction	Logistic Regression	Tensor Flow

These technologies form the foundation for real-time maintenance intelligence.

3.8. System Architecture

Figure 5 presents the full system architecture, integrating the UNS, predictive models, prescriptive logic, digital-twin engine, and dashboards. Data exchange uses structured APIs and message queues.




Figure 5. The real system architecture used in the test-bench environment.

End-to-end decision cycle time is measured as

$$C_{\text{cycle}} = L_{\text{UNS}} + L_{\text{model}} + L_{\text{prescriptive}}$$

where L_{UNS} denotes UNS communication latency, L_{model} model inference time, and $L_{\text{prescriptive}}$ prescriptive decision time. This ensured that the architecture created was able to accept real-time operations. In addition, Table 2 includes the most important technologies in question and the corresponding advantages.

Table 2. Key technological components and benefits.

Technology	Component	Benefit
IoT Sensors	Real-time data collection	Timely detection of anomalies
Machine Learning	Predictive failure models	Reduced downtime, proactive maintenance
3D Visualization	Digital-twin integration	Enhances monitoring and status updates
Data Management	Distributed data storage	Fast access and data reliability
Security Protocols	Encrypted data flow	Ensures secure, compliant system

3.9. Predictive and Prescriptive Support for Maintenance

The predictive model monitors equipment condition to detect anomalies and degradation trends [24]. The prescriptive module selects the best maintenance action using

$$\min_{a \in A} C(a) + \lambda R(a)$$

where $C(a)$ represents intervention cost, $R(a)$ the risk associated with failure, and λ a weighting factor balancing cost and reliability

3.10. Lifecycle Phases of the Prescriptive Maintenance Strategy

The maintenance strategy progresses through perceived, conceived, predictive, and adaptive phases. Sensor anomalies trigger transitions between these phases using an anomaly score:

$$S_t = \frac{|x_t - \mu|}{\sigma} \quad (1)$$

When the score exceeds a threshold, the system changes from historical analysis to predictive or adaptive decision-making. These phases ensure the maintenance strategy dynamically reflects real equipment conditions.

3.11. Data Security and Integrity

Security controls were used to provide confidentiality and integrity of the system over UNS channels of communication. Any information sent over MQTT and OPC-UA was encrypted with TLS [25]. Role-based access controls limited access to sensitive parameters of operations. The audit logs kept track of all the transactions, and thus traceability. The tampering probability was used to determine data integrity:

$$P_{\text{tamper}} = \frac{\text{Detected Anomalous Packets}}{\text{Total Packets}}$$

where P_{tamper} is the probability of data tampering.

4. Case Study Analysis

This case study shows how the proposed UNS system works in a real scenario. The test bench represents a small industrial machine. It behaves like a pump or a motor used in many factories. The setup includes vibration sensors, temperature sensors, pressure sensors, and a power meter. These sensors send data to the UNS in real time.

The UNS collects all sensor values under a single namespace. Each value has a timestamp. Each value has a clear name. This makes the data easy to use. The UNS sends

the data to the predictive model without delay. The average latency is about 18 ms. The packet loss is very low at 0.40 percent.

The predictive model checks the sensor data. It predicts if the equipment will fail. The model accuracy is about 94 percent. The model gives early warnings when vibration or temperature increases. The UNS makes the data stable and clean. This improves the model accuracy compared to point-to-point systems.

The digital twin receives the same inputs as the real machine. It simulates the behavior of the equipment. The digital-twin output matches the real output. The error stays below 0.05. The drift stays below 0.012. This shows that the digital twin is reliable. The twin tests the prescriptive actions before they are used on the real system.

The prescriptive module uses the prediction and the digital twin. It suggests the best action. It may suggest a shutdown. It may suggest a speed reduction. It may suggest a maintenance task. The test bench avoided more than 4000 possible failures during the study. This shows that the system can reduce downtime.

These 4000 prevented failures refer to simulated high-risk events generated within the controlled test-bench environment. They represent model-detected failure states validated through the digital twin, not real industrial failures.

This study demonstrates that a Unified Namespace can serve as an effective real-time integration layer for predictive and prescriptive maintenance within a controlled test-bench environment. The UNS reduced communication latency and packet loss, improved data consistency, and enabled higher predictive accuracy compared to a traditional P2P architecture. The digital twin provided a reliable mechanism for validating predictions and evaluating prescriptive actions before deployment.

These findings highlight the potential of UNS-based architectures; however, they are limited to the experimental setup used. Future work should evaluate scalability, multi-machine coordination, cybersecurity hardening, and deployment in real industrial environments with more complex operating conditions.

5. Benchmark Comparison

This section compares the UNS system with other maintenance systems. The first comparison is between UNS and a point-to-point (P2P) setup. The UNS system shows lower latency. The UNS system shows lower packet loss. The UNS system also shows higher throughput. As seen in Table 3 the predictive accuracy is also higher when the data comes through the UNS.

Table 3. Performance Comparison.

System	Accuracy	Latency	Packet Loss
UNS System	94.20%	~18 ms	0.40%
P2P System	87.51%	~31 ms	3.11%
Typical PdM (literature)	85–93%	Not reported	Not reported

The second comparison is with values reported in other predictive-maintenance studies. Many studies report accuracy between 85 percent and 93 percent. The UNS system reaches about 94 percent accuracy. This shows that the UNS improves the quality of the data. It also shows that the UNS helps the model make better predictions.

This comparison shows that the UNS system performs better. It also shows that the UNS system is more stable and more accurate.

6. Digital-Twin Validation

The digital twin is used to copy the behavior of the real equipment. It receives the same sensor inputs as the real system. It runs a simple state-space model. The goal is to match the real output as closely as possible. The digital twin helps to test actions before they are used on the real machine [26,27].

The digital twin shows high accuracy during the study. The vibration, pressure, and temperature signals match the real signals. The error stays below 0.05 for all three signals. This means the digital twin can follow the real system with good precision. The drift also stays low. The drift values remain below 0.012. This shows that the digital twin is stable over time [28].

The digital twin also supports prescriptive maintenance. It tests the recommended actions in a safe virtual space. It checks if the action reduces risk. It checks if the action prevents failure. It checks if the action keeps the system within safe limits. This helps to avoid wrong decisions. It also reduces the chance of unnecessary shutdowns.

The digital twin improves the reliability of the predictive model. It confirms the prediction before the action is taken. It also helps to tune the model when the behavior changes. This makes the full system more robust. It also makes the maintenance process safer.

7. Results

All results presented in this section refer specifically to the controlled test-bench environment described earlier. These findings should therefore be interpreted as evidence of feasibility under controlled conditions rather than as direct indicators of performance in full-scale industrial deployments. Performance metrics such as latency, drift, and predictive accuracy reflect the hardware configuration, injected fault scenarios, and UNS implementation used in this study. These results should therefore be interpreted within the scope of the experimental setup and not generalized to full-scale industrial systems without further validation.

7.1. Real-Time Data Performance and Communication Latency

The histograms shown in Figure 6 compare performance metrics in an UNS scenario. The top-left plot indicates that throughput using UNS is consistently high, with most measurements clustered tightly around 1600–1650 units, showing stable and predictable performance. The top-right plot reveals that MQTT latency is generally low, peaking around 16–18 ms, with a symmetric distribution indicating reliable response times. Likewise, the bottom-left histogram for OPC UA latency shows significantly higher values, centered around 25–27 ms, reflecting slower performance compared to MQTT. Finally, the bottom-right plot of maximum jitter demonstrates very low variation, mostly below 4–6 units, suggesting minimal timing instability across the tested system.




Figure 6. Cont.




Figure 6. Real-time data performance and latency.

7.2. Predictive Model Accuracy and Stability

Table 4 summarizes the predictive model's accuracy and stability across 5000 test samples. The model demonstrates strong overall performance, achieving a mean accuracy of 94.18%, with precision and recall averaging 91.79% and 92.69%, respectively. In addition to accuracy, the model demonstrated strong precision (mean 91.79%) and recall (mean 92.69%), with a low mean absolute error of 0.034, as shown in Table 4. These metrics confirm that the model performs consistently across both positive and negative classes, strengthening the reliability of the reported 94% accuracy. The low standard deviations ($\approx 1\%$ for accuracy, precision, and recall) indicate consistent performance across samples, while the MAE remains minimal at 0.034, reflecting precise predictions. The observed ranges show that even the lowest accuracy and precision values ($\approx 90.5\%$ and 87.7%) remain high.

Table 4. Predictive model accuracy and stability.

Statistic	Accuracy (%)	Precision (%)	Recall (%)	MAE
Count	5000.000	5000.000	5000.000	5000.000
Mean	94.182	91.792	92.688	0.034
Std	0.983	1.199	1.129	0.0049
Min	90.545	87.656	87.788	0.0168
25%	93.498	90.961	91.921	0.0306
50%	94.192	91.799	92.673	0.0340
75%	94.851	92.613	93.458	0.0373
Max	97.811	96.230	96.801	0.0505

Likewise, simulated prediction drift over 10 consecutive time windows remained very low, with all measured drift values remaining below 0.012, as shown in Figure 7. The drift fluctuates mildly between approximately 0.002 and 0.009, with minor peaks around time windows 3–4 and 8, but never exceeds the acceptable threshold of 0.012. This indicates excellent model stability over time; the data distribution and model behavior remain highly consistent with the original training conditions, with no significant concept or data drift occurring. The model does not require retraining or monitoring alerts based on this simulation, demonstrating robust real-world reliability even as new data arrives overtime.

7.3. UNS vs. Point-to-Point (P2P) Integration: Comparative Evaluation

The comparative analysis between the UNS and traditional P2P integration demonstrates the superiority of UNS in system performance and predictive reliability, as indicated in Figure 8. The mean latency for UNS is significantly lower at ≈ 18 ms compared to ≈ 31 ms for P2P, with packet loss reduced from 3.11% to 0.40%, indicating more efficient and reliable data transmission. UNS also achieves higher throughput (≈ 1620 messages/s) versus P2P

(≈ 1180 messages/s). In terms of predictive performance, UNS-processed data yields a mean accuracy of 94.20% and MAE of 0.034, outperforming P2P data, which achieves 87.51% accuracy and MAE of 0.061, as indicated in Figure 9.




Figure 7. Simulated drift over time.




Figure 8. UNS and P2P comparative performance metrics.




Figure 9. UNS and P2P mean error value. Box plot comparing MAE distributions for UNS and P2P. Circles represent statistical outliers, indicating data points that fall outside the typical range of each method's error distribution.

7.4. Digital-Twin Validation and Simulation Accuracy

Figure 10 consists of three time-series plots comparing real sensor measurements (blue) against predictions from a digital-twin model (orange) over approximately 7000 s, with error thresholds below 0.05 for all metrics.




Figure 10. Real vs. digital twin.

Vibration: The real and twin signals overlap almost perfectly, both showing highly similar oscillatory patterns with amplitudes ranging from about 3.5 to 6.5 units.

Pressure: The two lines are virtually indistinguishable, fluctuating tightly around 100 units with minor variations between roughly 99.7 and 100.1.

Temperature: Again, excellent alignment is observed, with both signals following the same gradual waves from around 44 to 56 degrees.

Across all three parameters, including vibration, pressure, and temperature, the digital twin demonstrates remarkable fidelity to real-world data, maintaining errors well under 0.05 throughout the entire period.

7.5. Evaluation of Prescriptive Maintenance Decisions

Figure 11 presents two histograms illustrating the outcomes of an optimization process. The left plot shows the distribution of parameter adjustment risk reduction, with most instances clustered around 20–25 units, forming a roughly normal distribution peaking near 23 and spanning from about 15 to 30. The right plot depicts optimized maintenance downtime reduction, similarly normally distributed but with a higher peak count, centered around 17–18 units and ranging from approximately 10 to 26. Both distributions indicate consistent and substantial benefits from optimization, with most cases achieving meaningful reductions in risk (around 22 units on average) and downtime (around 18 units), demonstrating reliable performance improvements across numerous simulations or scenarios [29].




Figure 11. Optimization process. The bars represent the frequency distribution of each metric, while the overlaid smooth lines depict the estimated probability density, highlighting the underlying distribution shape.

Moreover, as shown in Figure 12, which illustrates the impact of implementing controlled shutdowns on preventing system failures, for cases where no failure was prevented (labeled as 0), the count is zero, indicating no such occurrences. In contrast, when controlled shutdowns successfully prevented failures (labeled as 1), the count exceeds 4000 instances. This demonstrates that the predictive maintenance strategy, likely powered by the digital twin and monitoring system, effectively averted over 4000 potential failures through proactive interventions, resulting in zero unprevented critical failures and significantly enhancing system reliability and safety.

In addition, Figure 13 illustrates the impact of implemented prescriptive maintenance actions on risk reduction across three scenarios. Scenario 1, representing an optimal and

targeted maintenance strategy, shows the largest relative reduction, indicating that focused interventions are highly effective. Scenario 2, corresponding to a standard or conventional maintenance approach, achieves a moderate decrease in risk, demonstrating that even typical prescriptive actions contribute to mitigation. Scenario 3, reflecting a high-risk or reactive context with initially elevated risk levels, achieves only a moderate reduction, suggesting that while interventions are beneficial, high initial risk limits the overall impact.




Figure 12. Failure count.




Figure 13. Comparison of risk reduction across three scenarios, showing the effectiveness of targeted, standard, and high-risk maintenance interventions.

7.6. System-Level Decision Cycle Performance

Figure 14 presents the time distributions of key decision-making components. The total decision cycle is consistently centered around ~130 units, indicating stable overall performance. Communication time shows very low variability with a mean of ~22 units, reflecting efficient data exchange. Model inference time remains stable around ~40 units, while prescriptive computation time, averaging ~70 units, is the most time-intensive and variable component.




Figure 14. Distributions of decision-cycle components; bars show frequency and curves show estimated density.

8. Discussion

The empirical findings reveal that the suggested Unified Namespace-based predictive and prescriptive maintenance system provides high data communication performance, prediction accuracy, system integration, data fidelity in digital twins, and decision optimization. The live performance indicators prove that UNS offers stable, high-throughput communication with minimal jitters, whereas MQTT is always lower in latency than OPC UA. The outcome of the predictive model in terms of accuracy (>94), error, and drift is low, which means that the model will remain reliable in the long run without the need to be retrained regularly.

Traditional IoT and point-to-point (P2P) architectures often suffer from tightly coupled data flows, protocol heterogeneity, and limited scalability, leading to higher latency and inconsistent data quality [30]. As shown in the manuscript, the P2P baseline exhibited higher latency (~31 ms), greater packet loss (3.11%), and lower predictive accuracy (87.51%) compared to the UNS-based approach (“UNS System. . . 94.20% accuracy, ~18 ms latency, 0.40% packet loss”). The UNS architecture differs fundamentally by providing a protocol-agnostic, hierarchical namespace that decouples IT/OT layers, ensures consistent data semantics, and supports real-time analytics and digital-twin feedback loops. This integration of predictive, prescriptive, and simulation-based validation is not typically available

in conventional IoT platforms, highlighting the conceptual and operational novelty of the proposed framework.

Digital-twin validation has shown that the behavior of the simulated systems is like that of real sensor data, and there is almost ideal alignment of the simulated system behavior in vibration, pressure, and temperature, which supports its utilization in both monitoring and prediction applications [30]. Prescriptive maintenance assessments indicate that optimization-based decisions always minimize operational risk and downtime in a broad range of situations. The measures of controlled shutdown prevented all possible critical failures (>4000 cases), and risk reduction analysis by scenarios demonstrates that even in high-risk settings, prescriptive interventions have a significant impact on reducing risk.

The results are consistent with the available literature concerning Industry 4.0 data structures, predictive maintenance, and digital twins. Previous literature indicates the shortcomings of P2P integration, including data silos, accumulation of latency, and poor-quality data, that UNS-based architectures aim to overcome [31]. The findings also relate to prior work confirming that harmonized hierarchical namespaces enhance the responsiveness of a system and minimize the overhead of communication [32].

Moreover, deep learning and sensor fusion models have also proven to be successful in the prediction of maintenance in the industrial environment, with a range of accuracy between 85 and 93% in the predictive-maintenance literature. The higher-than-94% accuracy in this study compares to most reported values, which may be because the data streams were clean and consistent with UNS. The extremely low drift also reflects emerging evidence that stable data structures remove concept drift by reducing noise and inconsistency in the source data. Digital-twin accuracy in the previous literature generally records small but significant differences between the actual and simulated sensor values [33]. Conversely, the almost identical traces in this work (error < 0.05) imply greater fidelity. The prescriptive maintenance element, comprising optimization-based decision-making and scenario-based risk mitigation, is compatible with the current body of research prioritizing the importance of linking predictive and prescriptive layers. Most studies state a partial or moderate decrease in downtime, but this research has shown highly consistent improvements, such as over 4000 avoided failures, exceeding conventional maintenance using thresholds or rule-based decision logic.

The findings of this research have several implications both for research and practice in industries. The robust performance of the UNS communication with the high level of accuracy and predictability of the predictive models proves the idea that a harmonized data architecture can significantly enhance the solidity of maintenance systems based on analytics. Cloud Data flows keep data clean and consistent, thus improving the quality of predictions and minimizing drift, which implies that architecture decisions directly impact model robustness. In practice, a combination of a high-fidelity digital twin with prescriptive optimization allows for proactive and risk-conscious decision-making, which can be observed through the significant decrease in downtime and risk scores and the number of possible breakdowns prevented [34]. These results suggest that the presented framework can contribute to the improvement of operational safety, reliability, and maintenance efficiency considerably.

Although the outcomes of this study are encouraging, the study has several limitations. This testing was performed on a controlled test bench, and this might not be fully representative of large-scale industrial conditions. The small sensor set and the analysis of short-term simulated drift could limit generalizability and prescriptive maintenance cases, and the optimization strategy was not based on actual operating data but on simulations. Thus, it needs additional validation in the case of practical conditions. Likewise, large-scale, multi-site deployments need to be employed to determine scalability and robustness.

Implementation of more sensor modalities, long-term drift research, and investigation of more sophisticated prescriptive optimization methodologies would further optimize the performance of the system.

9. Limitations

This study has some limitations. The test bench is not a full industrial system. It uses a small setup that represents a real machine, but it does not include all the conditions found in a factory. The sensors used in the study are limited. More sensor types can give more detailed information. The study also only ran for a short period. A longer study could show how the system behaves over months or years.

The UNS system works well in the test bench. But large factories may have more devices and more data. The system must be tested in a larger environment to confirm its performance. The digital twin also uses a simple model. A more advanced model may give better results for complex machines.

The prescriptive module works with the digital twin. But it needs more real-world testing. It must be checked with real operators and real maintenance teams. This will show how well it supports human decisions.

These limitations do not reduce the value of the study. They show where future work can improve the system. They also show how the system can grow into a full industrial solution.

Although the UNS-based architecture shows strong performance in a controlled test bench, generalization to large-scale industrial systems requires further validation. Future work will evaluate scalability, multi-site UNS federation, and integration with more complex ML models.

10. Conclusions

This study demonstrates that a Unified Namespace can function as an effective real-time integration layer for predictive and prescriptive maintenance within a controlled test-bench environment. By reducing communication latency and packet loss, improving data consistency, and enabling higher predictive accuracy, the UNS-based architecture provides a practical foundation for closed-loop maintenance intelligence. The integration of a digital twin further enhances decision reliability by validating predictions and prescriptive actions before they are applied to the physical system.

While the results are promising, they remain specific to the experimental setup used in this work. The findings should therefore be interpreted within the scope of a single-machine test bench rather than full-scale industrial deployments. Future research will focus on evaluating scalability across multiple assets, strengthening cybersecurity measures, and validating the architecture in real industrial environments with more complex operational conditions.

Author Contributions: Conceptualization, R.S.P. and A.D.; methodology, R.S.P.; software, A.D., validation, P.D., E.O. and M.P.; formal analysis, R.S.P.; investigation, R.S.P.; resources, A.D.; data curation, R.S.P.; writing—original draft preparation, R.S.P.; writing—review and editing, P.D., E.O. and M.P.; visualization, A.D.; supervision, P.D., E.O. and M.P.; project administration, P.D. and E.O. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

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

References

1. Miah, M.T.; Erdei-Gally, S.; Dancs, A.; Fekete-Farkas, M. A Systematic Review of Industry 4.0 Technology on Workforce Employability and Skills: Driving Success Factors and Challenges in South Asia. *Economies* **2024**, *12*, 35. [CrossRef]
2. Alam, M.; Islam, M.R.; Shil, S.K. AI-Based Predictive Maintenance for US Manufacturing: Reducing Downtime and Increasing Productivity. *Int. J. Adv. Eng. Technol. Innov.* **2023**, *1*, 541–567.
3. Boya, V.R.; Rao, K.S. An Effective Time Management System Will Play a Vital Role in Achieving Operational Excellence in Pharmaceuticals. *Int. J. Recent Technol. Eng.* **2019**, *8*, 2572–2575.
4. Bele, K.; Munsamy, M.; Telukdarie, A. Case Study: Evaluation of Maintenance Practices on Road Infrastructure. In Proceedings of the 2022 IEEE 28th International Conference on Engineering, Technology and Innovation (ICE/ITMC) & 31st International Association For Management of Technology (IAMOT) Joint Conference, Nancy, France, 19–23 June 2022. [CrossRef]
5. Bousdekis, A.; Lepenioti, K.; Apostolou, D.; Mentzas, G. A Review of Data-Driven Decision-Making Methods for Industry 4.0 Maintenance Applications. *Electronics* **2021**, *10*, 828. [CrossRef]
6. Lambán, M.P.; Morella, P.; Royo, J.; Sánchez, J.C. Using Industry 4.0 to Face the Challenges of Predictive Maintenance: A Key Performance Indicators Development in a Cyber Physical System. *Comput. Ind. Eng.* **2022**, *171*, 108400. [CrossRef]
7. Pillai, R.K.S.; O'Connell, E.; Denny, P. Role of Unified Namespace (UNS) and Digital Twins in Predictive and Adaptive Industrial Systems. *Machines* **2026**, *14*, 252. [CrossRef]
8. Adrita, M.M.; Brem, A.; O'Sullivan, D.; Allen, E.; Bruton, K. Methodology for Data-Informed Process Improvement to Enable Automated Manufacturing in Current Manual Processes. *Appl. Sci.* **2021**, *11*, 3889. [CrossRef]
9. Akano, O.A.; Hanson, E.; Nwakile, C.; Esiri, A.E. Designing Real-Time Safety Monitoring Dashboards for Industrial Operations: A Data-Driven Approach. *Glob. J. Res. Sci. Technol.* **2024**, *2*, 1–9. [CrossRef]
10. Bettahi, A.; Belouadha, F.Z.; Harroud, H. AI and EDM: Revolutionizing Global Education and Crafting a Personalized Digitally Advanced Learning. In *Internationalization of Higher Education and Digital Transformation: Insights from Morocco and Beyond*; Springer: Berlin/Heidelberg, Germany, 2025; pp. 243–258.
11. Chang, S.-Y.; Saklecha, S.R.; Wu, Y. Intelligent Semantic Extraction and Transformation Pipeline for Large-Scale Multimedia Data Processing. In *Proceedings of the 2025 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*; IEEE: New York, NY, USA, 2025; pp. 1–6.
12. Chang, W.; Chen, X.; He, Z.; Zhou, S. A Prediction Hybrid Framework for Air Quality Integrated with W-BiLSTM (PSO)-GRU and XGBoost Methods. *Sustainability* **2023**, *15*, 16064. [CrossRef]
13. O'Connell, E.; O'Brien, W.; Bhattacharya, M.; Moore, D.; Penica, M. Digital Twins: Enabling Interoperability in Smart Manufacturing Networks. *Proc. Telecom* **2023**, *4*, 265–278. [CrossRef]
14. Surendran Pillai, R.; Denny, P.; O'Connell, E. Predictive and Adaptive Manufacturing. *Adv. Int. J. Multidiscip. Res.* **2026**, *4*. [CrossRef]
15. Manditereza, K. How Does a Unified Namespace (UNS) Work? Hivemq.Com. Available online: <https://www.hivemq.com/blog/how-does-unified-namespace-uns-work-iiot-industry-40/> (accessed on 12 February 2026).
16. Manditereza, K. What Is Unified Namespace (UNS) and Why Does It Matter? HiveMQ. 2022. Available online: <https://www.hivemq.com/blog/what-is-unified-namespace-uns-iiot-industry-40/> (accessed on 12 February 2026).
17. Péter, Á.; Werner, S. *The Impact of Unified Namespace in Industry 4.0*; Lund University Libraries: Lund, Sweden, 2024.
18. Freitas, L.; Pereira, F.; Lopes, H.; Avram, C.; Leal, N.; Morgado, T.; Machado, J. OPC-UA vs. MQTT (UNS): Evaluating Alignment with RAMI4. In *Proceedings of the 0 Through Literature Review. International Conference Innovation in Engineering*; Springer: Berlin/Heidelberg, Germany, 2025; pp. 434–445.
19. Quincozes, S.; Emilio, T.; Kazienko, J. MQTT Protocol: Fundamentals, Tools and Future Directions. *IEEE Lat. Am. Trans.* **2019**, *17*, 1439–1448. [CrossRef]
20. Arden, N.S.; Fisher, A.C.; Tyner, K.; Yu, L.X.; Lee, S.L.; Kopcha, M. Industry 4.0 for Pharmaceutical Manufacturing: Preparing for the Smart Factories of the Future. *Int. J. Pharm.* **2021**, *602*, 120554. [CrossRef]
21. Dhingra, P.; Gayathri, N.; Kumar, S.R.; Singanamalla, V.; Ramesh, C.; Balamurugan, B. Internet of Things–Based Pharmaceuticals Data Analysis. In *Emergence of Pharmaceutical Industry Growth with Industrial IoT Approach*; Elsevier: Amsterdam, The Netherlands, 2020.
22. Surendran Pillai, R.; Denny, P.; O'Connell, E. Predictive Maintenance and the Prevention of Batch Failures in Pharmaceutical Manufacturing Causes, Impacts, and Industry Evolution. *Adv. Int. J. Multidiscip. Res.* **2026**, *4*. [CrossRef]
23. Xiong, M.; Wang, H.; Fu, Q.; Xu, Y. Digital Twin–Driven Aero-Engine Intelligent Predictive Maintenance. *Int. J. Adv. Manuf. Technol.* **2021**, *114*, 3751–3761. [CrossRef]
24. Werner, A.; Mendez-Rial, R.; Salvo, P.; Charisi, V.; Piccini, J.; Mousavi, A.; Civardi, C.; Monios, N.; Espinosa, D.B.; Hildebrand, M.; et al. Architecture for Predictive Maintenance Based on Integrated Models, Methods and Technologies. In *Proceedings of the Intelligent and Transformative Production in Pandemic Times: Proceedings of the 26th International Conference on Production Research*; Springer: Berlin/Heidelberg, Germany, 2023; pp. 259–274.

25. Jaime, F.J.; Muñoz, A.; Rodríguez-Gómez, F.; Jerez-Calero, A. Strengthening Privacy and Data Security in Biomedical Microelectromechanical Systems by IoT Communication Security and Protection in Smart Healthcare. *Sensors* **2023**, *23*, 8944. [[CrossRef](#)] [[PubMed](#)]
26. Dooley, A.; Penica, M.; McGrath, S.; O'Brien, W.; Bhattacharya, M.; O'Connell, E. Digital Twin Bridging: Enabling Virtual Worlds to Manifest in the Physical. In *Proceedings of the 2024 35th Irish Signals and Systems Conference (ISSC)*; IEEE: New York, NY, USA, 2024; pp. 1–6.
27. Kumar, V.; Upmanu, V.; Chauhan, R.P.S.; Sharma, N.; Kumar, B.V.; Zadoo, M. Real-Time Environmental Monitoring and Control for Ensuring Optimal Temperature and Humidity Conditions for Patient Comfort and Medical Equipment Integrity. In *Proceedings of the 2024 3rd International Conference on Sentiment Analysis and Deep Learning (ICSADL)*, Bhimdatta, Nepal, 13–14 March 2024; IEEE: New York, NY, USA, 2024; pp. 558–562. [[CrossRef](#)]
28. Fuller, A.; Fan, Z.; Day, C.; Barlow, C. Digital Twin: Enabling Technologies, Challenges and Open Research. *IEEE Access* **2020**, *8*, 108952–108971. [[CrossRef](#)]
29. Pillai, R.S.; O'Connell, E.; Denny, P.; Shefeeque, M.B.K. Gazer 3d: A Robust, Data Driven, Digital Twin and Immersive Visualisation Platform. *IJSAT-Int. J. Sci. Technol.* **2026**, *17*. [[CrossRef](#)]
30. Elijah, O.; Ling, P.A.; Rahim, S.K.A.; Geok, T.K.; Arsad, A.; Kadir, E.A.; Abdurrahman, M.; Junin, R.; Agi, A.; Abdulfatah, M.Y. A Survey on Industry 4.0 for the Oil and Gas Industry: Upstream Sector. *IEEE Access* **2021**, *9*, 144438–144468. [[CrossRef](#)]
31. Harrison, R.; Vera, D.A.; Ahmad, B. A Connective Framework to Support the Lifecycle of Cyber–Physical Production Systems. *Proc. IEEE* **2021**, *109*, 568–581. [[CrossRef](#)]
32. Badihi, H.; Zhang, Y.; Jiang, B.; Pillay, P.; Rakheja, S. A Comprehensive Review on Signal-Based and Model-Based Condition Monitoring of Wind Turbines: Fault Diagnosis and Lifetime Prognosis. *Proc. IEEE* **2022**, *110*, 754–806. [[CrossRef](#)]
33. Katsikeas, S.; Fysarakis, K.; Miaoudakis, A.; Van Bemten, A.; Askoxylakis, I.; Papaefstathiou, I.; Plemenos, A. Lightweight & Secure Industrial IoT Communications via the MQ Telemetry Transport Protocol. In *Proceedings of the 2017 IEEE Symposium on Computers and Communications (ISCC)*; IEEE: New York, NY, USA, 2017; pp. 1193–1200.
34. Simić, M.; Dedeić, J.; Stojkov, M.; Prokić, I. A Hierarchical Namespace Approach for Multi-Tenancy in Distributed Clouds. *IEEE Access* **2024**, *12*, 32597–32617. [[CrossRef](#)]

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