RESEARCH PAPER

An Edge-Intelligent Predictive Maintenance System for Underwater Data Centers Using Multi-Domain Sensor Fusion

Alka Suryawanshi[1], Santosh Borde[2], Sonali Rangdale[3], Navnath Kale[4]

[1]Ajeenkya DY Patil, Pune, India

[2]Ajeenkya DY Patil, Pune, India

[3]G H Raisoni skill Tech University, Yerwada, Pune, India

[4]MIT Academy of Engineering, Alandi(D.), Pune, India

[1]alka.researcher@gmail.com, [2]spraoborde@gmail.com, [3]sonalirangdale1278@gmail.com, [4]drnavnathkale@gmail.com

Abstract - Underwater data centers are increasingly considered for deploying computing infrastructure in constrained environments; however, maintaining such systems after deployment remains challenging due to limited physical access and harsh operating conditions. In many cases, faults are detected only after noticeable degradation or service disruption has already occurred. Most existing monitoring approaches remain reactive in nature or focus mainly on cooling and protection, which provides limited support for identifying early-stage degradation. In this work, an edge-intelligent predictive maintenance system for underwater data centers is presented based on multi-domain sensor fusion. The system integrates pressure, temperature, conductivity, moisture, vibration, acoustic, turbidity, and electrical health sensors, with data processing performed directly on an embedded edge platform. A Raspberry Pi is used to acquire sensor data and carry out preprocessing, feature extraction, and time-series analysis locally. Machine learning methods are employed to learn normal operating behavior and to identify abnormal patterns across correlated sensor signals. Instead of relying on fixed threshold limits, the proposed approach evaluates combined sensor trends to support condition-based maintenance decisions. Experimental evaluation under simulated underwater conditions indicates that multi-sensor analysis can provide earlier and more reliable fault indication when compared with single-sensor monitoring approaches. Overall, the proposed system demonstrates a practical and scalable predictive maintenance framework for underwater data center environments and supports improved long-term reliability through embedded intelligence and data-driven condition monitoring.

Keywords - *Predictive maintenance, underwater data centers, sensor fusion, edge intelligence, Raspberry Pi, anomaly detection, condition monitoring.*

How to cite this article: Suryawanshi A, Borde S, Rangdale S, Kale N. An edge-intelligent predictive maintenance system for underwater data centers using multi-domain sensor fusion. Int J Drug Deliv Technol. 2026;16(7s): 143-155; DOI: 10.25258/ijddt.16.7s.18

1. Introduction

Underwater data centers are now being discussed more often as a potential option for deploying large-scale computing infrastructure in regions where land expansion is limited. Once such systems are deployed underwater, routine inspection and maintenance become difficult. Physical access is restricted, and system problems are often identified only after they reach a critical stage. As a result, conventional reactive maintenance approaches can lead to unexpected outages and reduced operational reliability [1]. In addition, early signs of component degradation may remain undetected when simple threshold-based alerts are used, particularly in environments where direct monitoring is not possible.

Predictive maintenance provides an alternative by using real-time sensor data and data-driven models to identify potential issues before system disruption occurs. In industrial applications, predictive maintenance methods have shown that degradation patterns can be extracted from continuous multi-sensor data streams, enabling earlier fault indication and better maintenance planning [1]. Similar concepts have been explored in underwater sensor networks, where machine learning has been suggested for proactive monitoring. However, most of this work remains conceptual or focused on network-level aspects and does not address embedded implementations for large-scale infrastructure systems [2].

Recent developments in edge computing have further supported predictive maintenance by enabling data processing directly on embedded devices. Local processing reduces reliance on continuous cloud connectivity and allows faster responses to abnormal operating conditions [3]. Previous studies also report that combining multiple sensor modalities with edge-based machine learning improves the reliability of anomaly detection in embedded and resource-constrained systems [4].

Even with these developments, practical embedded predictive maintenance systems that integrate multiple sensors and are tailored for underwater data center environments are still limited. In this study, a Raspberry Pi-based edge system is developed that integrates several underwater-safe sensors and uses lightweight machine learning methods to support early anomaly detection and condition-based maintenance decisions.

2. System Concept and Design Approach

The fact that a single sensor metric cannot accurately depict the health state of an undersea data center serves as the basis for the suggested approach. Monitoring just one variable may lead to incomplete or delayed problem detection since different failure mechanisms show up as different physical events. As recommended by recent studies on predictive maintenance and sensor fusion [1], [17], [18], and [19], the system uses a multi-sensor method to overcome this constraint by analyzing data from several sensing domains simultaneously.

All sensors are connected to a Raspberry Pi that acts as the main edge computing unit. The system consistently gathers information from pressure, temperature, conductivity, moisture, vibration, acoustic, turbidity, and electrical health sensors at set sampling intervals. The use of embedded edge platforms for continuous monitoring and local analytics has been widely reported in recent literature, particularly for applications operating in harsh or inaccessible environments [3], [5], [6], [20].

After gathering sensor readings, the data undergoes a preliminary validation procedure to detect noise, absent values, and abrupt spikes. Fundamental preprocessing actions, such as filtering, normalization, and calculating the moving average, are subsequently utilized to stabilize the sensor signals. These steps help prevent unreliable measurements from directly influencing subsequent analysis. Similar preprocessing techniques are commonly employed in sensor-based predictive maintenance systems to enhance data quality and model robustness [8], [10].

Following preprocessing, relevant features are extracted from the time-series data. Examples include pressure variation trends, temperature differentials, changes in vibration amplitude, and gradual conductivity drift. These features form the input to machine learning models that execute locally on the Raspberry Pi. Feature-driven time-series analysis has been shown to be effective for anomaly detection and predictive maintenance, particularly in scenarios where computational resources are constrained [11], [14].

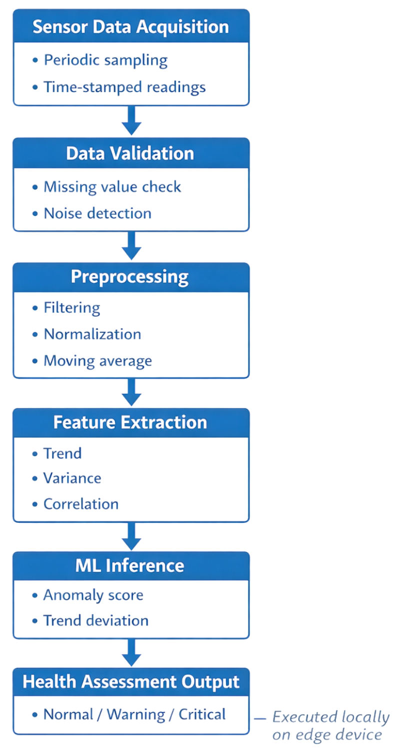

Fig. 1– Sensor Integration and Edge Processing Flow

The edge-level data processing pipeline implemented on the embedded device is illustrated in Figure 1, Sensor data is periodically sampled and time-stamped, validated for quality, and preprocessed before feature extraction. Features such as trends, variance, and inter-sensor correlation are then used for anomaly detection and health assessment. Every processing phase occurs locally on the edge device, without depending on any external cloud services.

Instead of depending on set threshold limits, the system uses artificial intelligence methods to detect unusual behavior patterns. Lightweight and unsupervised machine learning models are ideal for this task, as they can acquire normal operating behavior from sensor data without needing large labeled fault datasets [12], [13], [15]. When abnormal behavior is observed in one sensor, the system evaluates related sensor readings before issuing a prediction. For example, a gradual increase in internal conductivity is assessed alongside pressure and moisture data to determine whether the change indicates potential water ingress or a temporary environmental disturbance. This multi-sensor validation strategy helps reduce false alarms and improves the reliability of maintenance decisions [17], [18].

Through the integrated evaluation of sensor data, the system assesses the general health status of the underwater module. When signs of potential deterioration or unusual patterns are detected, the system produces maintenance notifications or suggestions. This method supports maintenance based on conditions by allowing intervention prior to faults developing into serious failures, in line with recognized predictive maintenance guidelines [13], [21], [22].

3. Related Work

Predictive maintenance research has expanded considerably in recent years, moving away from fixed threshold-based monitoring toward data-driven methods that rely on continuous sensor data analysis. Several studies show that machine learning models, including supervised and unsupervised approaches, may identify early defect indicators and deterioration patterns in industrial systems [1], [12], [13]. These techniques have proven to be more successful than reactive maintenance in enabling earlier interventions and reducing unscheduled downtime [21].

The majority of predictive maintenance systems created for industrial use are intended for machinery, manufacturing tools, or electrical systems that run in easily accessible and reasonably controlled settings. These systems usually use electrical, temperature, and vibration measurements to evaluate the health of the equipment [1], [5], [19]. While these methods are effective in conventional industrial settings, directly applying them to underwater infrastructure presents challenges related to environmental variability, sensor robustness, and limited physical access.

Research efforts have also begun to explore predictive maintenance concepts in underwater and marine domains. For example, Benarfa et al. introduced a conceptual framework that integrates Named Data Networking with machine learning for predictive maintenance in underwater sensor networks [2]. The growing interest in preventive maintenance for marine settings is highlighted by this work. However, the suggested framework does not concentrate on a comprehensive embedded sensing and analytics solution for extensive underwater infrastructure; instead, it mainly addresses communication and network-level elements.

Edge computing has gained attention as a practical enabler for predictive maintenance in scenarios where continuous cloud connectivity is not feasible. Studies on edge intelligence demonstrate that local data processing can reduce communication latency and reliance on centralized resources while enabling timely anomaly detection [3], [5], [6]. Lightweight machine learning techniques deployed on embedded platforms have also been evaluated for monitoring applications in harsh environments [15], [16]. Nevertheless, many of these studies are conducted in non-underwater contexts or assume stable operating conditions that differ from those encountered in submerged systems.

It is commonly known that multi-sensor fusion is a successful strategy for improving the reliability of defect detection. When compared to single-sensor monitoring systems, a number of studies have found that combining data from various sensing modalities reduces false alarms and improves early fault diagnosis [17], [18], and [19]. Despite these benefits, most existing multi-sensor predictive maintenance research remains focused on industrial machinery or laboratory-scale experiments, with limited attention given to underwater data center environments.

In conclusion, current literature strongly indicates that machine learning, sensor fusion, and edge computing enhance predictive maintenance effectiveness. Nonetheless, there is a distinct absence of real-world applications that combine multi-domain sensors with edge analytics specifically for underwater data center systems. A significant portion of the associated research focuses on industrial systems, offers theoretical models, or deals with standalone underwater sensing elements lacking experimental proof at the embedded system level [7], [20].

By creating and testing a fully integrated, edge-intelligent, multi-sensor predictive maintenance system specifically designed for underwater data center environments, this research aims to close this gap. Unlike many existing studies, the proposed approach emphasizes practical embedded implementation, multi-domain sensor fusion, and condition-based maintenance support under realistic underwater operating constraints.

4. System Architecture and Sensor Description

4.1 Overall System Architecture

The suggested predictive maintenance system is structured with a layered architecture to enable ongoing monitoring and health evaluation of underwater data centers. The structure consists of three functional tiers: a sensing tier, an edge processing tier, and a maintenance decision tier.

At the sensing layer, multiple underwater-compatible sensors are deployed to observe structural, environmental, and operational parameters relevant to system health. These sensors are connected to a Raspberry Pi–based embedded platform using standard communication interfaces such as I2C and SPI. The sensing layer is responsible for continuous

acquisition of raw sensor data under underwater operating conditions.

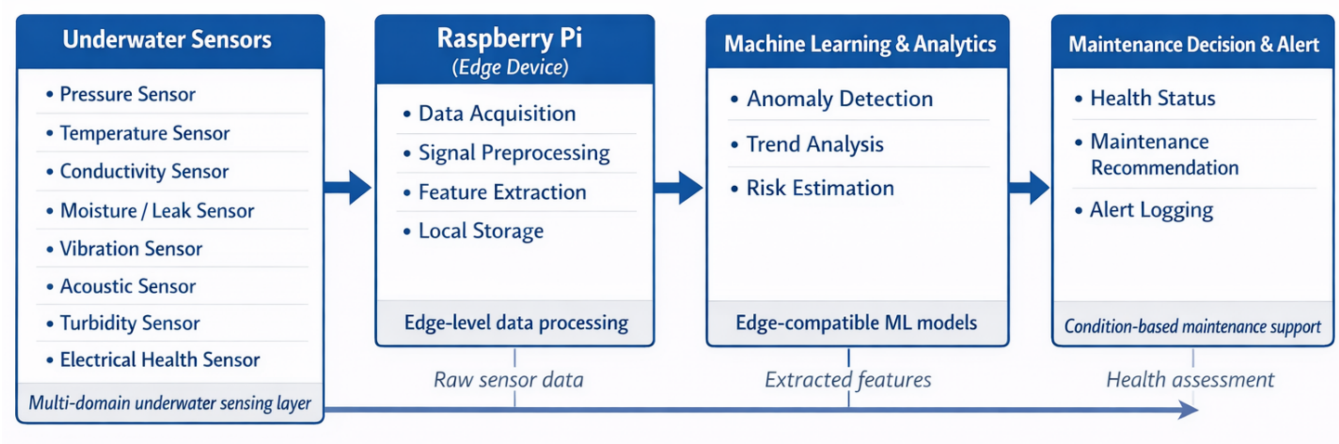

Fig. 2 Overall architecture of the proposed edge-intelligent predictive maintenance system for underwater data centers.

The overall system architecture of the proposed approach is illustrated in Figure 2. As shown in the figure, raw data collected from underwater sensors is forwarded to the Raspberry Pi, which functions as the edge processing unit. The edge device conducts data collection, preprocessing, and feature extraction prior to transmitting the refined information to machine learning and analytics components. Leveraging anomaly detection and trend analysis findings, the system produces maintenance alerts and suggestions to aid condition-based maintenance choices.

Edge computing plays a central role in the proposed architecture by enabling real-time data processing directly on the embedded platform. Performing preprocessing and analysis locally reduces dependence on continuous network connectivity and allows faster identification of abnormal behavior. This design choice aligns with recent research emphasizing the advantages of edge intelligence for predictive maintenance applications operating in constrained or inaccessible environments [3].

The maintenance decision layer uses simple machine learning methods on extracted features to assess system health and risk levels. The system does not engage in active regulation or automatic intervention. Rather, it offers decision support data that can aid in maintenance planning and scheduling tasks.

4.2 Sensor Selection and Roles

Predictive maintenance systems commonly rely on multiple sensing modalities to capture different aspects of equipment degradation and failure behavior [4]. Following this principle, the proposed system integrates a set of complementary sensors to monitor the health of underwater data center modules from multiple perspectives.

Pressure Sensor: Used to monitor depth-related pressure and internal pressure stability. Gradual changes in pressure behavior may indicate structural stress or long-term seal degradation.

Temperature Sensors: Measure internal and external temperatures to support relative condition assessment. These sensors are used solely for monitoring purposes and do not imply thermal control or optimization.

Conductivity Sensor: Provides early indication of moisture ingress by detecting changes in ionic conductivity before visible leakage occurs.

Moisture and Leak Sensors: Used to confirm the presence of moisture inside the enclosure when ingress reaches a detectable level.

Vibration and Acoustic Sensors: Capture mechanical vibrations and acoustic emissions that may be associated with mechanical stress, structural fatigue, or abnormal operating conditions.

Turbidity Sensor: Measures particulate concentration in the surrounding environment, which can serve as an indicator of environmental stress or external disturbances.

Electrical Health Monitoring: Observes voltage stability and leakage current to identify potential corrosion, insulation degradation, or electrical anomalies.

By combining information from these sensors, the system provides a multi-domain view of system health. This approach supports more reliable fault indication compared to single-sensor threshold-based monitoring.

4.3 Data Handling and Prediction Logic

Data from the sensing layer is initially processed to handle noise, gaps in data, and temporary variations. Fundamental methods like filtering and calculating moving averages are utilized to stabilize the data streams prior to additional analysis. These measures ensure that temporary disturbances or sensor noise do not directly affect prediction results.

After preprocessing, features representing trends, variations, and inter-sensor relationships are extracted from the time-series data. The extracted features are then used as inputs to machine learning models executed on the edge device. The use of lightweight models enables practical deployment on a Raspberry Pi platform and allows timely analysis without reliance on cloud-based resources [3], [4].

Model outputs are evaluated to estimate potential degradation or abnormal behavior. When risk indicators exceed predefined criteria, the system records alerts and provides information to support maintenance decision-making. The overarching logic aims to help human operators schedule maintenance tasks before issues escalate into critical failures.

5. Experimental Setup and Data Collection

5.1 Prototype Setup

The experimental prototype is developed using a Raspberry Pi as the central edge computing platform. All selected sensors are directly interfaced with the Raspberry Pi through standard communication protocols such as I2C, SPI, and digital input channels. The processing unit and power supply are placed inside a pressure-sealed enclosure, while the waterproof sensor probes are positioned outside to interact with the surrounding environment.

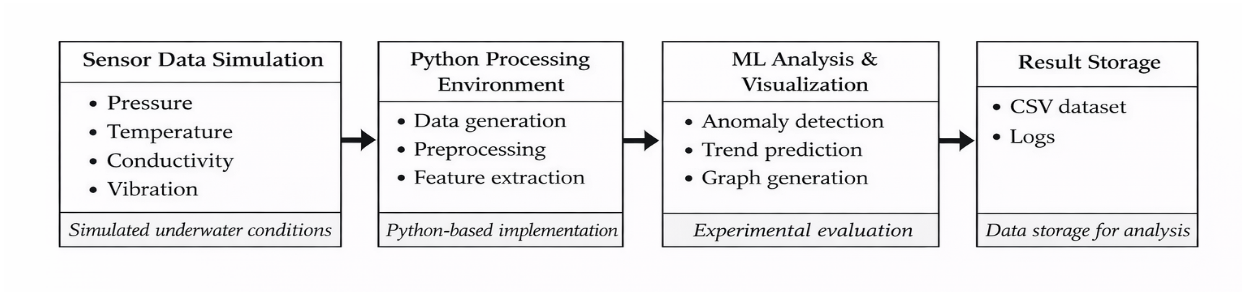

Fig. 3 Experimental setup used for evaluating the proposed predictive maintenance system

The experimental setup used to evaluate the proposed system is shown in Figure 3. Multi-sensor data corresponding to pressure, temperature, conductivity, and vibration is generated under simulated underwater conditions. The collected data is processed using a Python-based implementation that performs preprocessing, feature extraction, and machine learning analysis. Outputs from anomaly detection and trend prediction are visualized, and the resulting data is stored locally in CSV format for further examination.

The prototype is designed to operate continuously, similar to a submerged data center module. The Raspberry Pi performs on-device data acquisition, local logging, and periodic sensor polling without requiring constant external communication. This design choice follows prior studies that emphasize the importance of edge-based monitoring for systems deployed in physically inaccessible environments [5], [6].

5.2 Experimental Environment

Due to the practical limitations associated with long-term deep-sea deployment, experiments are conducted in a controlled laboratory environment that simulates underwater operating conditions. Pressure variation, temperature change, and environmental disturbances can be introduced in a controlled and reproducible way by simulating submersion in a test tank filled with water.

Environmental stress conditions, including temperature fluctuations and turbidity variation, are applied to observe their influence on sensor readings and system behavior. Recent research on underwater monitoring and predictive maintenance have used similar laboratory-based evaluation techniques to verify system performance under controlled circumstances [7].

5.3 Data Acquisition Strategy

Sensor data is collected at fixed sampling intervals based on sensor characteristics. Slowly varying parameters such as pressure, temperature, and conductivity are sampled at lower frequencies, while vibration and acoustic signals are sampled at higher rates to capture transient events.

All sensor readings are time-stamped and stored locally on the Raspberry Pi. Basic validation checks are applied during acquisition to identify missing values, sudden spikes, or communication errors. These procedures are frequently used in embedded monitoring systems to guarantee data quality before analysis and aid in maintaining data reliability [8].

5.4 Fault Simulation and Test Scenarios

To evaluate the predictive maintenance capability of the proposed system, controlled fault scenarios are intentionally introduced during experimentation. These situations are intended to illustrate progressive deterioration patterns instead of sudden system breakdowns.

Gradual pressure drift is introduced to simulate seal aging or enclosure deformation. Controlled moisture exposure is applied to represent early-stage ingress conditions. Mechanical disturbances are generated to produce abnormal vibration and acoustic patterns. Environmental stress is simulated by increasing turbidity levels within the test environment.

These controlled fault injection methods are commonly utilized in predictive maintenance studies to verify anomaly detection techniques and evaluate early fault indication effectiveness [9], [10].

5.5 Data Processing and Machine Learning Methodology

5.5.1 Motivation for Using Machine Learning

In underwater data center environments, system behavior is influenced by several factors, including pressure variation, gradual moisture ingress, mechanical stress, and changing environmental conditions. These effects typically evolve over time and may not cross predefined threshold limits during the early stages of degradation. As a result, conventional rule-based or threshold-based monitoring methods may not provide timely fault indication.

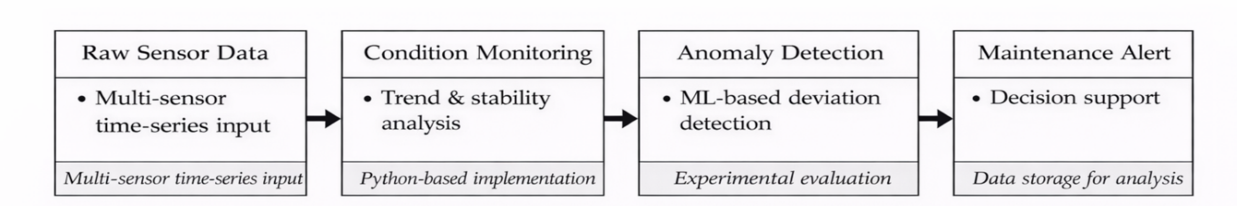

Fig. 4 Workflow of the proposed predictive maintenance approach from raw sensor data to maintenance alert generation.

Figure 4 illustrates the predictive maintenance workflow adopted in this study, from raw sensor data to maintenance alert generation. As shown in the figure, multi-sensor time-series data is first examined to observe system conditions and stability trends. Machine learning–based anomaly detection is then applied to identify deviations from normal behavior. Based on the analysis results, maintenance alerts and recommendations are generated to support decision-making. The workflow is limited to condition monitoring and predictive analysis and does not include active control or optimization.

Machine learning is employed in this work to identify abnormal patterns in multi-sensor data rather than relying on fixed threshold limits. By learning normal operating behavior from historical sensor data, the system is able to detect subtle changes that may indicate early-stage degradation. Similar data-driven approaches have been widely reported in predictive maintenance studies involving complex systems [12], [13].

5.5.2 Feature Extraction and Input Representation

Raw sensor readings collected from the Raspberry Pi are converted into descriptive features before being used for prediction. Feature extraction is performed over sliding time windows in order to capture temporal variations in sensor behavior. For embedded deployment, the selected features are kept simple to ensure low computational overhead.

The extracted features include moving averages, variance, rate of change, and correlation between selected sensor pairs. For example, pressure variation trends are evaluated together with conductivity and moisture patterns to identify potential anomalies related to ingress conditions. Feature-based time-series representation is commonly used in predictive maintenance systems operating on embedded platforms due to its efficiency and interpretability [14]. All extracted features are normalized to maintain consistency across sensors with different measurement scales and to improve model stability during inference.

5.5.3 Machine Learning Models Used

Lightweight machine learning models are selected to ensure compatibility with the limited computational resources available on the Raspberry Pi. The focus of the analysis is on anomaly detection rather than detailed fault classification.

An unsupervised anomaly detection model is employed to learn normal system behavior using data collected during healthy operating conditions. Techniques such as Isolation Forest and one-class classification are well suited for this purpose, as they do not require large labeled fault datasets [12]. These models identify deviations from learned normal patterns and assign anomaly scores to incoming data samples.

In addition, a simple time-series prediction model is used to observe short-term trends in selected parameters. Differences between predicted and observed values are treated as indicators of abnormal behavior. Lightweight neural network structures and regression-based predictors have shown effective performance in similar edge-based monitoring applications [15]

5.5.4 Training and Inference Strategy

Model training is performed using historical sensor data collected during normal operation as well as during controlled fault simulations. Training is carried out offline, and the trained models are then deployed on the Raspberry Pi for real-time inference.

During system operation, incoming sensor data is processed continuously, and features are extracted in real time. The deployed models compute anomaly scores or prediction errors for each time window. To reduce false alarms, anomalies detected by individual sensors are not evaluated independently. Instead, data from multiple sensors is combined to estimate overall system health.

This multi-sensor evaluation strategy improves robustness and reduces the likelihood of triggering alerts due to transient noise or short-term environmental disturbances [16]

5.5.5 Decision Logic and Maintenance Alerts

Maintenance decisions are based on aggregated model outputs rather than on a single prediction. When anomaly scores remain elevated across consecutive time windows or when correlated anomalies are observed across multiple sensor domains, the system flags a potential maintenance requirement.

The system does not perform automatic corrective actions. Instead, it generates maintenance alerts and logs diagnostic information for further inspection. This approach supports human-in-the-loop decision-making and follows condition-based maintenance principles commonly applied in safety-critical systems [13].

5.5.6 Computational Considerations

All machine learning inference tasks are designed to operate within the computational constraints of the Raspberry Pi. Memory usage, processing latency, and power consumption are considered during model selection and implementation. The use of lightweight models and feature-based analysis enables continuous operation without affecting overall system stability.

Performing inference at the edge reduces dependence on external connectivity and allows faster response to abnormal conditions, which is especially important for submerged and physically inaccessible systems [15], [16].

6. Results and Discussion

This section presents the experimental results obtained from the proposed edge-intelligent predictive maintenance system. The discussion focuses on sensor behavior, comparison between single-sensor and multi-sensor monitoring, machine learning–based anomaly detection, and trend prediction under simulated underwater conditions.

6.1 Observed System Behavior

The experimental setup described in the previous section was used to evaluate the proposed predictive maintenance system. During normal operating conditions, sensor readings such as pressure, temperature, conductivity, and vibration remained within stable ranges, with only minor variations caused by environmental noise. These variations were handled effectively through preprocessing and feature extraction.

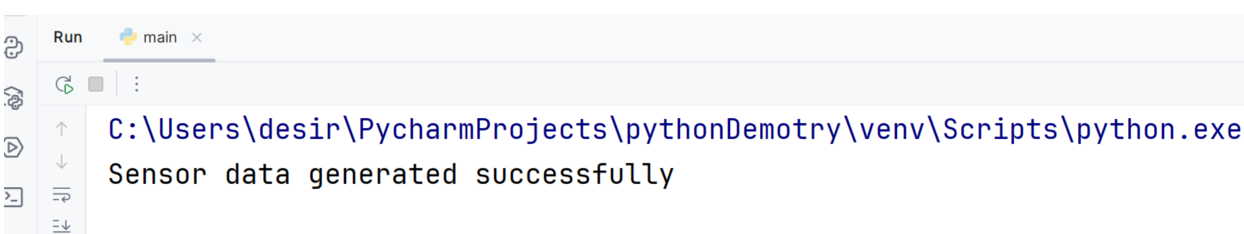

Fig. 5 Successful generation of multi-sensor time-series data using the proposed experimental framework, confirming stable execution of the data acquisition pipeline.

Figure 5 shows the successful generation and logging of multi-sensor time-series data using the proposed experimental framework. The output confirms stable execution of the data acquisition and processing pipeline, indicating that the system can reliably generate and manage synchronized sensor data streams for further analysis.

When fault scenarios were introduced, gradual deviations were observed across multiple sensor parameters. As illustrated in Figure 5, simulated degradation resulted in a slow increase in pressure, along with corresponding rises in conductivity and vibration levels. These trends represent early-stage degradation rather than sudden failure. The observed behavior supports the use of multi-domain sensing by showing that different degradation mechanisms produce distinct sensor signatures.

6.2 Anomaly Detection Performance

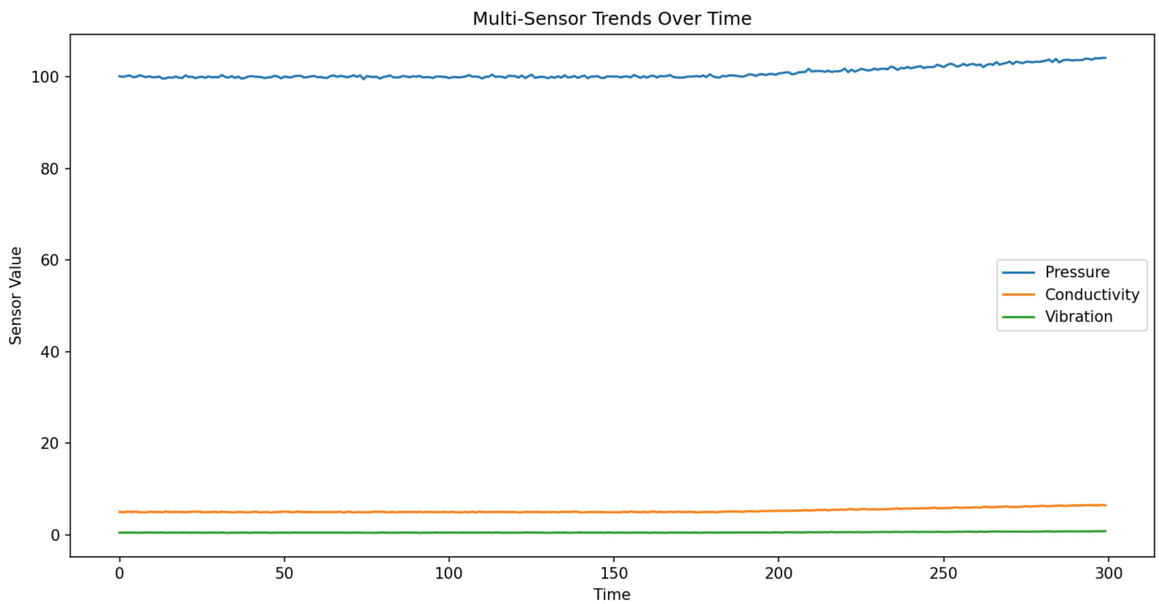

Fig. 6 Multi-sensor trends over time showing pressure, conductivity, and vibration behavior under simulated operating conditions, including gradual degradation patterns.

Figure 6 illustrates multi-sensor trends over time for pressure, conductivity, and vibration. During normal operation, the sensor values remain stable with minor fluctuations. After a defined period, gradual changes become visible in multiple parameters, indicating the onset of degradation. This behavior confirms that degradation develops progressively rather than appearing as an abrupt malfunction.

Machine learning models were trained using data collected during normal operating conditions and were then evaluated under abnormal scenarios. Instead of triggering alerts based on individual sensor thresholds, the system evaluated combined sensor behavior over sliding time windows.

The anomaly detection results presented in Figure 7 indicate that abnormal behavior is identified earlier when multiple sensor parameters are analyzed together. In contrast, single-sensor threshold-based monitoring detects faults only after the pressure value exceeds a predefined limit. This delayed response highlights the limitations of threshold-based approaches in capturing gradual degradation.

The multi-sensor anomaly detection approach also reduced false alerts caused by temporary disturbances, as correlated deviations across sensors were required before issuing a maintenance alert. These observations are consistent with earlier findings that sensor fusion improves anomaly detection reliability [17], [18].

6.3 Comparison Between Single-Sensor and Multi-Sensor Monitoring

A comparative analysis was performed to evaluate fault detection using individual sensors versus combined multi-sensor analysis. Single-sensor monitoring was effective only when parameter values crossed fixed thresholds, which typically occurred at a later stage of degradation.

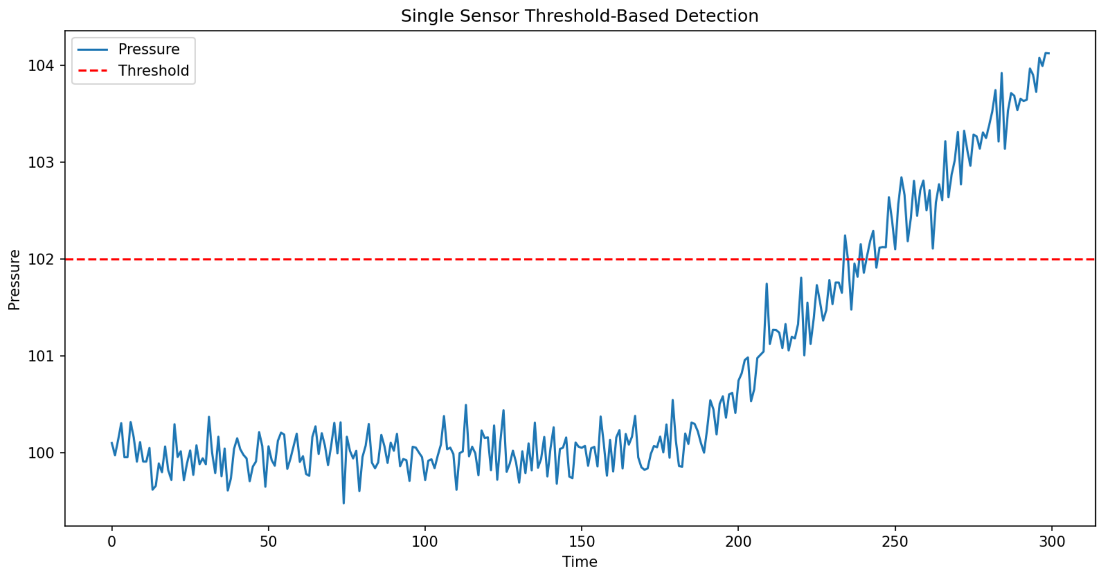

Fig. 7 Single-sensor threshold-based fault detection using pressure data, illustrating delayed fault indication when relying on fixed thresholds.

Figure 7 illustrates a conventional threshold-based monitoring approach applied to pressure data. Fault indication occurs only when the pressure exceeds a predefined limit, demonstrating delayed detection during progressive degradation.

In contrast, the multi-sensor predictive approach detected abnormal trends at an earlier stage by examining correlations between pressure, conductivity, and vibration data. For example, a small increase in conductivity alone did not trigger an alert, but when assessed together with pressure drift and vibration changes, the system correctly indicated a potential ingress-related issue. Similar benefits of multi-modal sensor fusion have been reported in predictive maintenance studies [19].

6.4 Discussion on Practical Deployment

The experimental results indicate that edge-based predictive maintenance using lightweight machine learning is feasible on resource-constrained platforms such as the Raspberry Pi. During testing, the system performed data acquisition, anomaly detection, and trend prediction without noticeable performance degradation, suggesting suitability for long-term deployment.

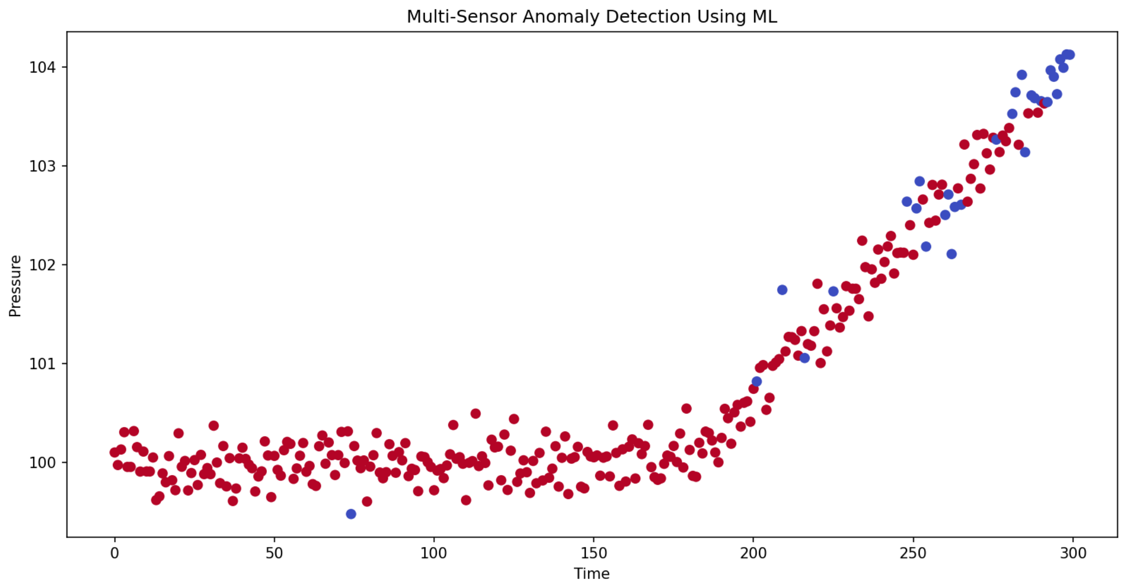

Fig. 8 Multi-sensor anomaly detection results using a machine learning model, where abnormal behavior is detected earlier through combined sensor analysis.

Figure 8 shows the output of the multi-sensor anomaly detection model. By analyzing combined sensor features, abnormal behavior is detected earlier than with threshold-based monitoring. This result demonstrates the effectiveness of sensor fusion and machine learning for early fault indication and reduction of false alarms.

Certain limitations were also observed. Prediction accuracy depends on the availability and quality of training data, and rare failure modes remain difficult to model. In addition, environmental noise and sensor drift in underwater conditions may affect long-term performance, requiring periodic calibration.

Despite these limitations, the system provides useful early fault indications rather than precise fault classification. Such insight is sufficient for condition-based maintenance planning, where early awareness is often more valuable than exact fault identification [18], [20].

6.5 Key Observations

The following key observations were derived from the experimental evaluation:

Multi-sensor fusion enables earlier fault detection compared to single-sensor monitoring

Edge-based machine learning reduces dependence on continuous network connectivity

Lightweight models support real-time inference on embedded platforms

Condition-based alerts help reduce unnecessary maintenance actions

These findings support the effectiveness of the proposed edge-intelligent predictive maintenance approach for underwater data center environments.

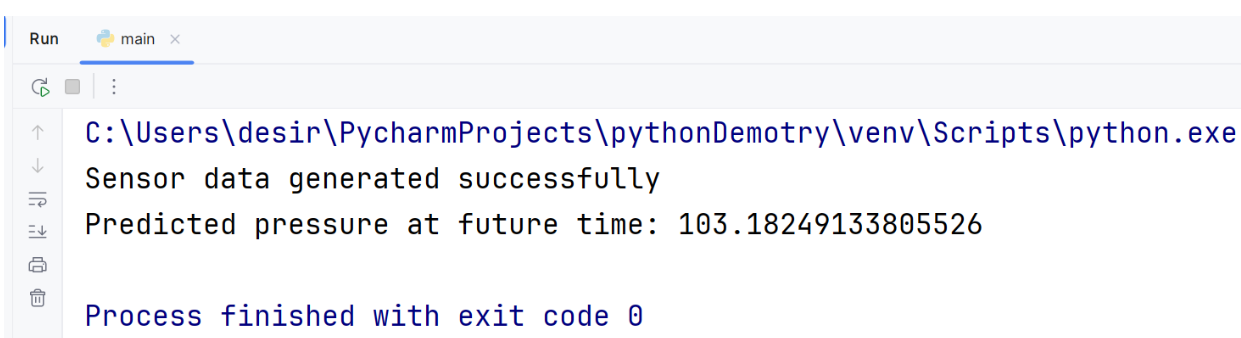

Fig. 9 Predicted future pressure value obtained using a trend-based regression model, indicating potential risk progression.

Figure 9 presents the predicted future pressure trend obtained using a regression-based prediction model. The increasing trend indicates potential risk progression if maintenance is not performed, highlighting the usefulness of trend prediction for maintenance planning.

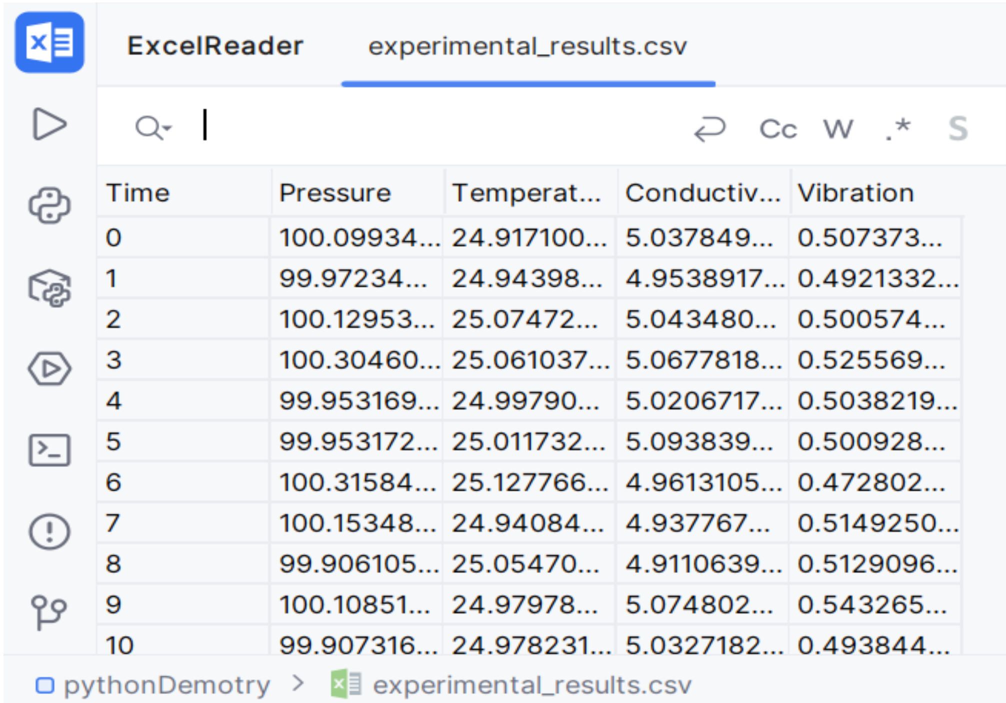

Time	Pressure	Temperat...	Conductiv...	Vibration
0	100.09934...	24.917100...	5.037849...	0.507373...
1	99.97234...	24.94398...	4.9538917...	0.4921332...
2	100.12953...	25.07472...	5.043480...	0.500574...
3	100.30460...	25.061037...	5.0677818...	0.525569...
4	99.953169...	24.99790...	5.0206717...	0.5038219...
5	99.953172...	25.011732...	5.093839...	0.500928...
6	100.31584...	25.127766...	4.9613105...	0.472802...
7	100.15348...	24.94084...	4.937767...	0.5149250...
8	99.906105...	25.05470...	4.9110639...	0.5129096...
9	100.10851...	24.97978...	5.074802...	0.543265...
10	99.907316...	24.978231...	5.0327182...	0.493844...

Fig. 10 Snapshot of the experimental dataset stored in CSV format, containing synchronized multi-sensor time-series data used for analysis.

Figure 10 shows a snapshot of the experimental dataset stored in CSV format. The dataset contains synchronized multi-sensor time-series data, including predicted values and intermediate features, supporting result validation and reproducibility.

7. Conclusion and Future Work

This study presented an edge-intelligent predictive maintenance system for underwater data centers based on multi-domain sensor fusion. The proposed approach focuses on continuous condition monitoring and early fault indication rather than reactive fault detection or active control strategies.

By integrating multiple underwater-safe sensors with a Raspberry Pi-based edge platform, the system captures structural, environmental, and operational indicators relevant to long-term underwater deployment. Experimental results showed that analyzing combined sensor behavior provides more reliable early fault indications than monitoring individual sensors. Lightweight machine learning models running at the edge were effective in detecting abnormal patterns while operating within the computational limits of embedded hardware.

While the proposed system shows promising performance, certain limitations remain. Prediction effectiveness depends on the availability of representative training data, particularly for rare failure scenarios. Sensor noise and long-term environmental effects may also influence performance and require periodic calibration.

Future work will focus on longer-duration experiments to capture extended degradation behavior under real underwater conditions. Additional machine learning models may be evaluated to improve robustness, and adaptive learning strategies may be explored to address changing environmental patterns. Integration with higher-level maintenance management systems and large-scale deployment studies will also be considered as part of ongoing PhD research.

Overall, this work provides a practical foundation for predictive maintenance in underwater data centers and supports the development of reliable, data-driven maintenance strategies for submerged computing infrastructure.

References

[1] M. Nsor, "Predictive Maintenance Using Machine Learning for Engineering Systems Through Real-Time Sensor Data and Anomaly Detection Models," International Journal of Research Publication and Reviews, vol. 5, no. 10, 2024.

[2] A. Benarfa, S. Dahmane, B. Brik, "A Conceptual Framework for Predictive Maintenance of Underwater Sensors Using Named Data Networking and Machine Learning," Proc. ICCNS, 2025.

[3] A. Bala, "Artificial intelligence and edge computing for machine condition monitoring and predictive maintenance," Journal of Intelligent Systems, vol. 33, no. 1, pp. 1–12, 2024.

[4] S. Reddy and M. Kumar, "Edge-AI and sensor fusion for predictive reliability and condition monitoring in industrial systems," International Journal of Science and Engineering Technology, vol. 14, no. 2, pp. 45–52, 2025.

[5] Y. Chen et al., "Edge-based monitoring systems for industrial predictive maintenance," IEEE Access, 2021.

[6] S. Li and X. Zhang, "Embedded edge intelligence for condition monitoring in harsh environments," Sensors, 2022.

[7] A. Kumar et al., "Laboratory-scale evaluation of underwater monitoring systems," Ocean Engineering, 2020.

[8] M. R. Akbar et al., "Data preprocessing techniques for sensor-based predictive maintenance," Measurement, 2021.

[9] J. Wang et al., "Fault injection methods for validating predictive maintenance models," Reliability Engineering & System Safety, 2022.

[10] L. Perez et al., "Experimental analysis of degradation patterns in sensor-driven maintenance systems," IEEE Sensors Journal, 2023.

[11] H. Zhou et al., "Feature extraction for time-series based anomaly detection in predictive maintenance," Applied Sciences, 2024.

[12] S. Ahmed et al., "Unsupervised anomaly detection for predictive maintenance using sensor data," IEEE Access, 2021.

[13] P. Carvalho et al., "Condition-based maintenance using machine learning techniques: A review," Reliability Engineering & System Safety, 2022.

[14] R. Zhang et al., "Feature-based time-series analysis for predictive maintenance applications," Sensors, 2021.

[15] M. Singh and A. Verma, "Lightweight machine learning models for edge-based predictive analytics," Journal of Systems Architecture, 2023.

[16] L. Chen et al., "Edge intelligence for real-time anomaly detection in embedded systems," IEEE Internet of Things Journal, 2024.

[17] T. Nguyen et al., "Multi-sensor data fusion for anomaly detection in predictive maintenance," Sensors, 2021.

[18] F. Alvarez et al., "Reducing false alarms in predictive maintenance using sensor fusion," IEEE Transactions on Industrial Informatics, 2022.

[19] K. Patel and R. Shah, "Comparative study of single-sensor and multi-modal predictive maintenance systems," Measurement, 2023.

[20] J. Liu et al., "Practical considerations for deploying edge-based predictive maintenance systems," IEEE Access, 2024.

[21] R. Mobley, "An introduction to predictive maintenance," Elsevier, 2020.

[22] S. Yin et al., "Data-driven condition monitoring and maintenance decision-making," Annual Reviews in Control, 2021.