Predictive Maintenance for Industrial Water Treatment Systems Using Real-Time Sensor Data and Machine Learning

MSc Research Project
MSc Data Analytics

Pradnesh Amar Shevantikar
Student ID: x23426845

School of Computing
National College of Ireland

Supervisor: Cristina Hava Muntean

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name:	**Pradnesh Amar Shevantikar**
Student ID:	**23426845**
Programme:	**MSc Data Analytics** **Year: 2025**
Module:	**Research Practicum**
Supervisor:	**Cristina Hava Muntean**
Submission Due Date:	**11th Dec 2025**
Project Title:	**Predictive Maintenance for Industrial Water Treatment Systems Using Real-Time Sensor Data and Machine Learning**
Word Count:	**9883** **Page Count: 22**

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: Shevantikar

Date: 11th Dec 2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

AI Acknowledgement Supplement

Research Practicum

Predictive Maintenance for Industrial Water Treatment Systems Using Real-Time Sensor Data and Machine Learning

Your Name/Student Number	Course	Date
Pradnesh Amar Shevantikar	Msc in Data Analytics	11th Dec 2025

This section is a supplement to the main assignment, to be used if AI was used in any capacity in the creation of your assignment; if you have queries about how to do this, please contact your lecturer. For an example of how to fill these sections out, please click here.

AI Acknowledgment

This section acknowledges the AI tools that were utilized in the process of completing this assignment.

Tool Name	Brief Description	Link to tool
NA	NA	NA
NA	NA	NA

Description of AI Usage

This section provides a more detailed description of how the AI tools were used in the assignment. It includes information about the prompts given to the AI tool, the responses received, and how these responses were utilized or modified in the assignment. **One table should be used for each tool used**.

[Insert Tool Name]	
[Insert Description of use] **NA**	
[Insert Sample prompt] **NA**	[Insert Sample response]

Evidence of AI Usage

This section includes evidence of significant prompts and responses used or generated through the AI tool. It should provide a clear understanding of the extent to which the AI tool was used in the assignment. Evidence may be attached via screenshots or text.

Additional Evidence:

[Place evidence here]

Additional Evidence:

[Place evidence here]

Predictive Maintenance for Industrial Water Treatment Systems Using Real-Time Sensor Data and Machine Learning

Pradnesh Amar Shevantikar
23426845
Program Code - Research Practicum
National College of Ireland

Abstract

The water-treatment plants are aimed at running continuously, but, nonetheless, maintenance is typically reactive in nature and, thus, it results in time loss due to costly down-time and unreliability of the system. This paper resolves this problem in industrial water-treatment systems based on a simulated dataset based on real operational conditions, where the researcher demonstrates how machine-learning-based predictive maintenance can be applied in predicting failure and improving operational functioning. The case under analysis was an experimental-simulated data of more than 40 000 records and 52 sensor variables (pressure, flow, vibration), environmental conditions (temperature, humidity, CO2) and treatment parameters (RO pressure, UV intensity, mineral dose). After removing and filtering variables, 22 important variables were obtained. The prediction of near-term failures was carried out with the help of the Logistic Regression, Random Forest and XGBoost supervised models and an Isolation Forest was employed to detect the emergence of anomalies. The model explanations that SHAP produced were able to show that the vibration, RO pressure, and UV intensity were the most important predictors. Live-like risk scores, anomaly indicators and SHAP based feature-attribution outputs are the outputs generated by the analysis that can be used to power a Power BI dashboard to support decisions by operators. The results indicate that an optimized XGBoost classifier with a tuned model reached a recollection of 1.00 and F1-score of around 0.73, indicating a high level of prediction ability of upcoming failures, although with a higher false alarm rate. The works provided are explaining predictive maintenance of industrial water systems specific to treatment that can be scaled and introduces a scalable model of data-driven reliability management. This implementation of streams of live sensors in the future will tend to facilitate precision and sustainability outcomes.

Keywords: Predictive maintenance, water treatment, Machine Learning, sensor data.

1. Introduction

1.1. Background

The water-treatment systems are important in facilitating safe, sustainable and efficient processing of water at the industrial, municipal and agricultural industries. These systems are continuously working with different chemical, environmental and mechanical conditions when even minor deviations may cause the equipments to degrade, enter downtime and spoil water quality. Conventional maintenance in these facilities is usually reactive or on a set schedule, which has led to inefficiency in planning and increased costs of operation as well as unnecessary use of resources.Industry 4.0, advanced sensorisation and the Industrial Internet of Things (IIoT), in recent years has made it possible to acquire the operating parameters of pressure, vibration, flow, temperature and UV intensity in real-time. Whether used alongside machine-learning-based analytics, these data streams can be changed into actionable insights, which means that failures can be predicted before they happen. These predictive methods assist in condition based maintenance methods and enhance system reliability and sustainability.

Industry 4.0 and Water Treatment

Industry 4.0 is the fourth industrial revolution that is marked by cyber-physical systems, intelligent sensors, cloud connectivity, automation and improved data analytics. Industry 4.0 technologies are used in water-treatment processes to combine IIoT devices with real-time monitoring and intelligence control systems to allow predictive maintenance, early fault detection and energy and chemical consumption optimisation. This digital revolution increases the level of transparency, efficiency and sustainability in operations, and subjects' industrial water-treatment operations to the principles of modern smart-manufacturing. Outside of operational improvements, sustainable water-treatment practices help contribute to global environmental objectives. The optimal use of resources, effective filtration, disinfection, and minimisation of pollutants will assist in decreasing the emission of pollutants, resource wastage, and encourage responsible utilisation of water and energy. Such improvements directly contribute to a number of the United Nations Sustainable Development Goals (SDGs) such as SDG 6 (Clean Water and Sanitation), SDG 9 (Industry, Innovation and Infrastructure), SDG 12 (Responsible Consumption and Production) and SDG 13 (Climate Action).

1.2. Research Problem

Although recent developments in IIoT and machine learning took place, there is limited research on predictive-maintenance in industrial water-treatment settings. Most of the existing literature concentrates on the individual elements of RO fouling or UV dose optimisation, but not on multi-stage and multi-purpose treatment units. Almost no research has incorporated real-time multivariate sensor observations, supervised learning, unsupervised anomaly detection and model elucidation in the same predictive-maintenance framework. The study will fill this gap and examine how machine learning and data analytics could be used to predict failures using simulated real-time industrial sensor data.

1.3. Motivation

The increasing penetration of IIoT, automation and real-time analytics in manufacturing environments opens the opportunities of intelligent and data-driven maintenance system. This study demonstrates that creating an explainable predictive-maintenance system incorporating supervised models, anomaly detection and SHAP-based explanations, and backed by dashboard ready visualisation results, is scalable to Industry 4.0 principles.. Although the simulated data on use are fictitious, they reflect the real-life operational trends within industrial systems, and offer a good background on developing intelligible decision-supporting tools to water-treatment operators.

1.4. Limitations and Assumptions

This research takes into consideration a number of limitations. The data set is modeled, and not obtained via active industrial equipment that would add the possibility of bias in labeling of failures. There is no actual downtime or maintenance record that can be used to validate actual operational impact. Lack of class balance and values outline a modelling problem, and the generalisability of the outcome to all water-treatment plants in industries should be confirmed by the real-life application. The findings are, therefore, a prototype as opposed to a complete field-tested system.

The research presupposes that the simulated sensor data are the reflections of real-world behaviour of industrial water-treatment systems. Degradation patterns are assumed to be picked by rolling-window statistics like moving averages, standard deviations and slopes. The rules of failure definition are considered to be reasonable proxies of operational occurrences. Relationships between pressure, temperature, vibration and UV intensity are also assumed to resemble normal behaviours of process-engineering.

1.5. Research Question and Objective

Can machine-learning-based predictive-maintenance models in combination with anomaly detection and SHAP explainability effectively predict failures in industrial water-treatment systems using real-time sensor data?

The objectives of this research are specific and include the following:

1. Explore the state of predictive maintenance, explainable AI and anomaly detection of industrial water-treatment systems.
2. Develop an integrated system of predictive-maintenance integrating supervised learning with unsupervised anomaly detection and SHAP-interpretable. on multivariate data of industrial sensors using machine-learning, such as Logistic Regression, Random Forest, XGBoost and Isolation Forest.
3. Test model performance in terms of precision, recall, F1-score, ROC-AUC and explainable AI.
4. Construct a dashboard prototype that helps to make decisions by the operators based on real-time risk scoring, alerts and interpretable feature attributions.

1.6. Structure of the Report

The rest of this report will be organized as follows. Section 2 includes the literature review connected to the concepts of predictive maintenance, data-driven water-treatment systems and explainable AI. Section 3 outlines the methodology, namely the data preparation, feature engineering and model development. Part 4 introduces the comparison of the supervised and unsupervised models. Section 5 gives dashboard ready data under visualisation, findings, discussions and limitations. Lastly, Section 6 is a conclusion of the study that provides future work recommendations.

2. Literature Review

2.1 Introduction

With organisations moving to Industry 4.0, predictive maintenance (PdM) has grown in significance in the industrial operation. As the IoT sensors and real-time data integration continue to expand, machine-learning models have become potentially capable of identifying failures in time to lower the downtime and enhance reliability to the whole system. Despite PdM research in other areas like manufacturing and energy, little is known in terms of its application in water-treatment systems. The process of water treatment is complicated, and it consists of a series of related steps such as filtration, RO, UV disinfection and chemical dosing with their distinctive operational issues and failure modes. The current literature review critically analyzes the work that is already in the field of water-treatment analytics, predicting membrane fouling, optimizing UV, anomaly detection, Internet of Things, and explainable AI (XAI), and seeks to fill the gaps that justify the current study.

2.2. Predictive Maintenance for Membrane and RO Systems

One of the most common failure modes that have been researched in the application of water-treatment processes is membrane fouling. Other researchers performed a comprehensive analysis of AI-based AI-based prediction of membrane-fouling analyzing models like Artificial Neural Networks (ANNs), Support Vector Machines (SVMs) and hybrid mechanistic-ML models (Abuwatfa et al., 2023). They found hybrid models had a higher degree of accuracy and frequently lowered prediction errors (RMSE/MAE) than pure data ones since they included physical information concerning the decline and operating conditions of the fluxes. The majority of studies that they reviewed, though, were based on datasets of a laboratory scale that was not representative of real industrial systems. This renders the reported models unreliable in the case of noisy non-stationary plant data. Besides this, Abuwatfa et al. (2023) in his review demonstrated that most studies on fouling prediction consider RO systems independently of how the tendencies of pressure, flow or temperature upstream units or downstream units affect fouling behaviour. The other weakness is that practically none of the reviewed models can be interpreted, which poses challenges to operators to have insights into the reason behind a prediction. Such weaknesses are significant since industrial systems need to have transparent and physically meaningful predictions to make decisions. More recent research has covered integrated adsorption-ML models in order to put foulings into a more holistic view. These models are a combination of the mechanistic reasoning and information-driven learning and they are generally more successful when it comes to generalisation and decision support. Nevertheless, they pay nearly no attention to chemical dosing, UV performance and vibration pattern, which are also important subsystems. This limited scope of research supports one of the main gaps in the research: the existing predictive-maintenance models are not based on the multi-stage, sensor-intensive nature of a real industrial water-treatment plant.

2.3. UV Disinfection Performance and ML-Based Control

Another essential subsystem in the water treatment process is UV disinfection, and there are multiple studies that have been trying to optimize the UV dose delivery or identify irregularities based on ML. An example is a study done by Minzu et al. (2021), which used neural networks to predict UV dose performance with different turbidity and flow conditions. The model was accurate with controlled datasets, however, the method needed huge amounts of quality training data. In a real-world industrial setting, the drift in sensors, the aging of the lamp and the changing water quality would prevent achieving the same degree of data consistency as required by neural networks. Moreover, neural networks are black-box models, which do not provide much information about the sensor variables that affect dose performance, which makes them less useful in maintaining plans. El-Shafeiy et al. (2023) suggested the sensor-fusion method to enhance UV prediction in dynamic conditions. Their approach increased strength and minimized variability of the UV dose estimations, yet the article was very performance-oriented. What they did not investigate was how these predictions would be applicable in predictive maintenance, and what they did not examine is the anomaly detection methods that would be useful in detecting sudden lamp failures or sudden decreases in UV intensity. In the literature on UV studies, one theme stands out: despite all the models emulating the dose performance, they mostly do not deal with the maintenance outcomes, e.g. early failure identification or degradation monitoring. Moreover, all the reviewed studies combine UV analytics with RO or dosing data, which restricts the insight into how the errors in one step are spread throughout the system.

2.4. Anomaly Detection in Water Treatment

Although the technique of anomaly detection is used in other industrial settings, it is yet to find extensive application in water treatment. Isolation Forest, Local Outlier Factor and statistical process control have been demonstrated as techniques that can be used to find unusual sensor behaviour of manufacturing and power-generation systems (Liu et al., 2008; Breunig et al., 2000). These techniques are light, do not need labeled data set and are efficient with multivariate sensor streams, hence suitable in water-treatment scenarios. Nevertheless, anomaly detection in maintenance pipelines is rarely implemented in existing studies of water-treatment. The majority of studies concentrate on a single-outcome prediction (e.g. fouling or UV dose). This is a huge gap since anomalies usually come first before actual failures. The decrease in UV intensity, an unforeseen spike in vibration or a sudden pressure change can be indicated in hours before the failure. In the absence of an anomaly-detection layer, the maintenance systems that operate on MLs will not be fully developed. The paper fills this gap by integrating the supervised predictive models with the Isolation Forest-based anomaly detection, which develops a more holistic maintenance framework.

2.5. IoT Architectures and Model-Selection Frameworks

The real-time monitoring of water-treatment plants is also moving towards the use of IoT infrastructures. Khattach et al. (2025) proposed an architecture that combines time-series sensor acquisition, streaming analytics and model retraining. Their contribution is applicable in determining the manner PdM systems can be integrated into the pipelines of the IoT. Nevertheless, their article has not tested the predictive-maintenance performance or has not discussed the interaction of the operators with the outputs. As well, they did not investigate model explainability, which is a highly important need in safety sensitive settings. Li et al., (2025) compared different models of supervised-learning including Logistic Regression, Decision Trees and Random Forests, to classify the work in the industry. Their results prove the accuracy of these algorithms on tabular sensor data. Their model-comparison framework however lacks unsupervised models or explainable AI models restricting its applicability to real predictive-maintenance models where operators require both transparency and predictability of unexpected events. Combined, these studies emphasize the fact that the Internet of Things infrastructure is insufficient. An entire PdM system needs feature engineering, anomaly detection, interpretable models, and real-time dashboards all of which are synthetised in the current study.

2.6. Explainable AI for Industrial Maintenance

Since machine-learning systems are increasingly being used in operational environments, it is crucial to know why a model will foretell a failure. The application of SHAP (SHapley Additive exPlanations) in manufacturing has demonstrated how it can assist operators to find the most influential factors that lead to the prediction of failure leading to the enhancement of trust and decision-making (Molnar, 2019). Nevertheless, there is a limited XAI study in water treatment. The published research is mainly on interpretation in manufacturing or chemical-processing plants as opposed to multi-stage water-treatment systems. Moreover, there are other works comparing SHAP to LIME and pointing to the fact that SHAP is more stable, particularly in the context of noise sensor data which is a critical factor in water treatment. However, none of the reviewed studies have embedded SHAP explanations into a complete PdM workflow and anomaly detection and dashboard visualisation. To fill this gap, this research creates SHAP explanations and provides an exportable data in dashboard form that can be added to a Power BI interface and allow the operators in the future deployment to investigate the reasons behind failures prediction and the most important features.

2.7. Industry 4.0 Integration and Dashboard Design

The use of smart sensors, connected infrastructure, process automation and data-driven decision-making are encouraged by Industry 4.0 to enhance the operations of plants. The process of water treatment including multi-stage filtration, UV operation, and control dosing enhances the quality monitoring throughout the process. Although they do not directly address the predictive maintenance issue in their work, it proves the importance of the combination of different subsystems data, which is also represented by the rich sensor dataset provided in this study. Interaction Dashboard visualisation is also significant in smart manufacturing due to its contribution to the operator interaction with machine-learned results. The situation awareness and decision making in industrial environments have been proven to be augmented by dashboards in the past (Ucar et al., 2024). Nonetheless, the current literature lacks the integrated PdM dashboard of water-treatment settings that will combine the predictions of supervised learning, anomaly warnings and SHAP visualizations. This project will fill this gap as it will design a Power BI dashboard-ready visual results, which will visualise the level of risk, anomaly scores and feature-attribution insights in a format that will be used in real-time in the future.

2.8. Synthesis and Research Gap

In all of the reviewed areas, including membrane fouling, UV optimisation, anomaly detection, IoT integration and explainable AI, one can identify a number of evident gaps:

1. Available literature is subsystem-based.
   The majority of studies are on RO or UV alone, rather than their interacting behaviour in a complete treatment line of water.
2. Little use of combined multivariate sensor information.
   None of the studies considers RO pressure, UV intensity, vibration, mineral dose, flow and environmental factors as one predictive-maintenance model.
3. Absence of joint supervised + anomaly detectors.
   There is scarcely a case in which the predictable failures and the unexpected anomalies are identified in current research.
4. Most of the water-treatment ML models are not interpretable.
   Few studies use SHAP or other XAI methods on data on water-treatment sensors (Saarela et al., 2024).
   Poor operation solutions (dashboards, deployment readiness)

The majority of studies end with model training and fail to show operative-facing tools.

3. Methodology

This section provides the manner in which the study was modelled, the preparation of the dataset, and the creation and assessment of the predictive models. It is aimed to make a clear and transparent statement of the research procedure to make it repeatable and evaluateable.

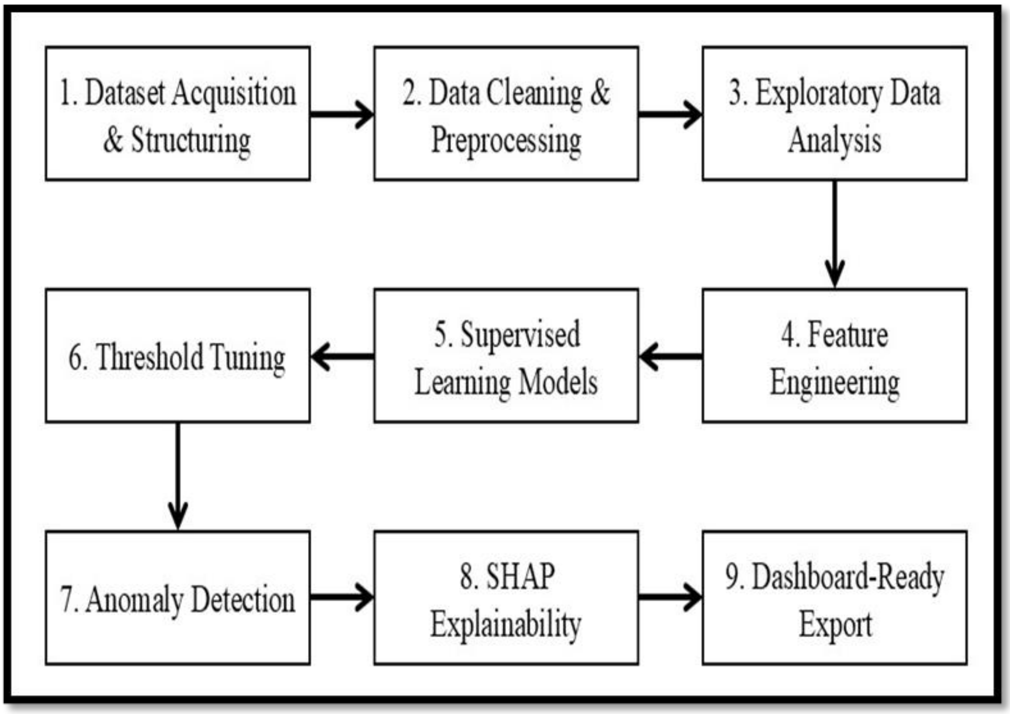

Figure . Data Collection and Preprocessing Flow

3.1. Research Methodology Overview

The research was conducted using a quantitative, experimental research design, through which all findings were supported by the data-driven analysis and machine-learning experimentation. The methodology was made up of a logical steps starting with definition of researched problem, which was followed by acquisition and preparation of data set, exploratory analysis, generation of modelling features, selection and training predictive models, performance measurement and ultimate interpretation and visual communication of results. In order to direct these activities, structured workflow was embraced. It has started with the understanding of the dataset and the industrial background that it represents, proceeded to data cleaning and data exploration, and proceeded to supervised and unsupervised model development and testing. The last stages consisted of explaining the results with SHAP and developing a dashboard that will combine the predictive results into a format that can be used by the operators. This methodological framework will make sure that the steps of the analysis will be supported by the ones that have been completed before, and that a logical pipeline right through the acquisition of raw data to the interpretation and visualisation of the model will be created. The time-based partitioning of data used resembles the actual industrial use, where models have to act based on unknown circumstances in the future. In addition, the integration of supervised, unsupervised, and explainable methods makes the system identify predictable failures, emphasize unforeseen anomalies, and explain the rationales behind any prediction. The combination of these two strategies enhances the strength and applicability of the proposed predictive-maintenance scheme.

3.2. Dataset Description

The dataset in this study was sourced from a Water Treatment Industry and represents practical operational behaviour of a sort that is common in industrial water-treatment systems in use in Indian manufacturing plants. Even though the dataset is simulated, it is structured and behaves in a way that is congruent with the actual water-treatment operations. It has more than 40, 000 time-ordered observations and 52 sensor variables that include various stages of treatment such as filtration, reverse-osmosis (RO) and ultraviolet (UV) disinfection. The variables describe a very wide assortment of physical processes, including flow, pressure, pressure drop, vibration, temperature, humidity, UV intensity, mineral dosing, energy usage and external environmental conditions. The data covers various operational cycles, and daily and weekly variations, environmental variations, dosage alterations of chemicals, lamp-ageing and simulated failure conditions. This amount of variation renders the dataset applicable to addressing predictive maintenance because it represents the nature of sensor drift, noise and degradation patterns that are present in real-world systems.

The dataset was collected from one of the industrial water-treatment plants. It is not publicly available as they are data that reflect the operational behaviour related to the partner organisation. The data was published only in need of research purpose, and it was anonymised so that it would not fail to uphold commercial privacy. The dataset is a continuation of operational behaviour observed at the duration of 4.5 months, 1 January 2025 to 19 May 2025, of daily, weekly and seasonal variations that are often common in the processes of industrial water-treatment. The patterns of degradation as well as sensor variation described in the data set depict a typical situation of the industrial water-treatment system in India such as the RO filtration, UV disinfection, pump behaviour, hydraulic pressure variation and energy-use cycles. The ability to give the temporal and regional context is to make sure that the findings can be put into proper perspective with regard to actual working conditions.

3.3. Data Cleaning and Preprocessing

A number of preparation steps were required before the data could be modelled in order to resolve common problems occurring in industrial sensor data. The coping of missing values was done with varied approaches, based on the behaviour of the individual variable: stationary sensors were imputed with the median, and time-dependent readings were interpolated linearly to preserve the time continuity. Cases of outliers were scrutinized. Values that were greater than the interquartile range or had physical constraints were corrected or trimmed down. This can be the negative pressure readings or the zero UV intensity values which did not match the lamp-off periods. One-hot encoding was used to convert the categorical indicators into the numerical form, and the continuous variables were scaled with the help of the RobustScaler algorithm to eliminate the impact of outliers. Industrial data may tend to be non-stationary due to progressive wear, change in chemical or environmental impact. In an effort to explain such behaviour in the modelling process, rolling-window statistics of moving averages, moving standard deviations and temporal slopes of various important variables like RO pressure, pressure drop, vibration and UV intensity were formed. Such temporal characteristics are used to record degradation patterns that do not happen as suddenly. To prevent information leakage and to address the fact that predictive systems should be worked with in the real world the clean dataset was further broken down into chronologically matched training, validation and test sets. All numerical characteristics were scaled with RobustScaler to ensure that the impact of extreme values among the supervised and unsupervised models is minimized. The dataset used was having 40,000 records, and about 38,500 records were left after the cleaning and validating processes. Approximately one and half thousand (3.8) records were deleted because of extreme outliers, physically impossible values, or bad sensor values (Breunig et al., 2000). This low percentage of deleted entries is an indication that the data were of high quality, but the cleaning process provided the reliability and consistency needed to construct predictive models.

3.4. Exploratory Data Analysis (EDA)

The exploratory data analysis step had the following two functions: first, to gain insight into the natural behaviour of the system under normal and failing conditions; second, to detect the most relevant variables as predictors of maintenance. There were several behavioural patterns that were evident. RO pressure and pressure drop rose to the eye in the hours before a failure, which is the physical process of fouling the membrane. Flow rate conversely decreased steadily with the increase in resistance in the system. The same pattern was observed in the levels of vibration, which had prominent peaks just before most of the labelled failures, which were indicative of mechanical overload or the pump being imbalanced. UV intensity was in different behavior. Rather than sharp fluctuations, it exhibited a progressive decrease with age of the lamps in line with the lamp ageing but sudden drops were also observed during the anomaly periods which were probably due to changes in turbidity or electricity instability. Temperature and humidity were also strongly cyclical on a daily basis as environmental factors that affected UV behaviour and energy consumption. The analysis of correlation showed that some of the closest relationships with the target failure indicator were RO pressure, pressure drop, UV intensity and vibration. Multiple variables that were very highly correlated with each other were eliminated in an effort to eliminate multicollinearity. Comparison of time before failure and non failure indicated that there were consistent variations in these variables and therefore they were suitable to be included in the predictive models. Altogether, the exploratory analysis, in addition to identifying the most sensitive variables to the degradation patterns, proved the internal consistency of sensor data set. The trends of pressure, vibration, UV intensity and temperature are in line with the usual behaviours that are observed in the industrial water-treatment literature. This kind of alignment enhances the belief in the realism of the dataset and justifies the suitability of the dataset as the basis of predictive modelling. EDA was therefore a trade-off between domain knowledge and data-driven knowledge enabling the stages of modelling to concentrate on operationally most significant features.

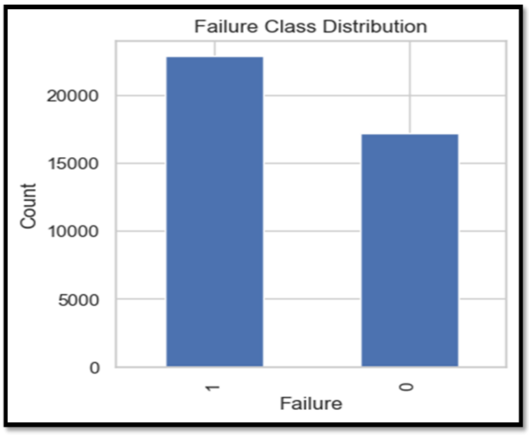

Figure 1 – Class Distribution Bar Chart

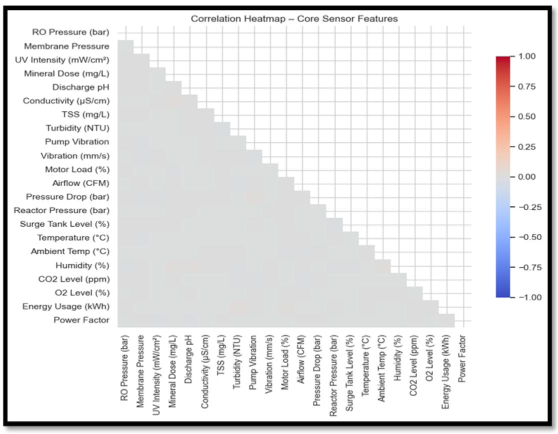

Figure 2 – Correlation Heatmap (Upper-Triangular)

3.5. Feature Engineering

On the basis of the knowledge obtained during EDA and with the help of the domain knowledge about water-treatment processes, a set of refined features of modelling was developed. These were rolling averages, standard deviations, slope characteristics to indicate both short and medium term trends of some of the main variables like RO pressure, pressure drop, vibration, UV intensity and energy consumption. These programmed features have been used to pick up minor accumulations and directional variations which mostly come before failures. Then out of the entire number of raw and engineered variables, 22 important features were ultimately obtained. The 22 features that were chosen to be modelled can be defined in four major categories depending on what kind of operational behaviour they reflect:

1. Pressure and Flow Dynamics: The characteristics of RO pressure, pressure-drop rolling means, and slopes record the trends in membrane fouling, clogging, and hydraulic resistance. These variables assist in determining whether the process of filtration is inefficient or unstable.
2. Mechanical Vibration: Rolling statistics of pump vibration and vibration slopes monitor mechanical upheavals, imbalance or bearing wears of rotating apparatus. Mechanical disturbance, imbalance or bearing wear in rotating equipment is indicated by rolling statistics of pump vibration and vibration slopes. These are the features that serve as an early warning system of equipment degradation.
3. UV Disinfection and Water Quality Indicators: Attributes based on UV intensity characterize the lamp ageing, UV sleeve fouling or the effect of turbidity in water. Engineered features of conductivity and TSS also indicate the variation in dissolved solids and particle concentration.
4. System Load and Consumption of Energy: Rolling-window averages and dispersion of energy consumption are used to point out abnormal loads of operation, equipment behaviour, or increasing mechanical friction of pumps.

These four are the groups that form the physical, hydraulic, mechanical and energy-performance properties of the treatment system. The engineered feature set represents short-term variations and long-term patterns of degradation that are characteristic of a system on the verge of failure by STT changing raw sensor values into multi scale time descriptors.

3.6. Model Development

Two types of models were worked out to cover various issues of predictive maintenance. The former was the category of supervised classification models such as the Logistic Regression and the random forest and the XGBoost trained to determine whether the system would enter a failure state based on the engineered features. These models were considered since they have performed well in other such maintenance situations in industries, and provide accuracy and interpretability balance. The second one was unsupervised detection of anomaly by Isolation Forest. This model was employed to determine abnormal or unforeseen trends that were not always found in the labelled failure events yet they still constituted abnormal system behaviour. The integration of supervised and unsupervised methods ensured that the study was able to identify familiar signatures of failure and those that had not been observed before. The chronologically partitioned dataset was used to train all the models so that they could be assessed in a realistic manner and also prevent the use of future data in predicting the past.

3.7. Model Evaluation

The standard classification metrics were used to assess the model performance, that is, precision, recall, F1-score and ROC-AUC. These measures gave an idea on the extent to which the models discriminated between normal and failing conditions and the extent to which it was reliable in identifying actual cases of failure. Confusion matrices were used to know the nature of errors committed by each model. The anomaly detector model has been tested using patterns of anomaly scores and confirming the occurrence of high-scoring events as being physically significant sensor behaviour deviations. Besides standard measures, prediction sequences were also checked qualitatively in order to determine temporal coherence. The clusters of elevated risk preceding the periods of failures recorded in the tuned models showed that the tuned models were sensitive to significant patterns of the precursors as opposed to isolated abnormalities. Such behaviour is a sign of a system that does not only differentiate the normal and failing states, but also offers the early warning of the worsening conditions, which is essential in making a timely decision on maintenance.

3.8. Explainability (SHAP)

Explainability was a critical component of the methodology since it will make the model output transparent, understandable and reliable to operators and maintenance engineers. In this research, SHAP (SHapley Additive exPlanations) algorithm was used on the tuned XGBoost model to measure the contribution made by each feature to individual prediction. The SHAP values allowed exploring further whether the model was basing its hypotheses on physical indicators that were meaningful like increased RO pressure, spikes in vibration, or a decrease in the intensity of UV light and not on arbitrary statistical patterns. One of the most important methodology aspects that this study has brought is the extension of SHAP-based explainability in a predictive-maintenance model of industrial water treatment systems, where the interpretability of models has been a largely ignored issue. Available literature on fouling of membranes, UV disinfection behaviour, or RO pressure behaviour is mainly centered on predictive quality, and not much on why a model gives a prediction of a failure. Combining both global and local SHAPs analyses, the work offers new insights into the drivers of risk of failure, allowing domain experts to check, believe, and put recommendations of the system into practice. The fact that SHAP can point at physically interpretable patterns of precursors as opposed to random statistical correlations is a major step forward in comparison with a black box predictive methodology. Such descriptions will then be incorporated into the dashboard to help the engineers make decisions based on data.

3.9. Dashboard Development

The last step in the methodology was to translate outputs of the models in a manner that would be easily understood by operators. An exported time-stamped failure probabilities of the tuned Random Forest model (Failure_Proba_RF), Isolation Forest anomaly scores and labels were combined with the key sensor features in creating a dashboard-ready dataset. This is an exported file (marcuras powerbi data.csv), meant to feed a Power BI dashboard that could provide the failure risk scores, anomaly alerts, SHAP-based explanations and sensor trends in an integrated manner (Ucar et al., 2024). Although the complete interactive dashboard is suggested as a next phase of deployment, the current work makes sure that the predictive models and explanations have been technically prepared to be implemented into an Industry 4.0-like operator interface.

4. Results

In this chapter, the evaluation of the predictive-maintenance framework that was created during this study is given in detail. The findings address four key aspects, namely,

(1) supervised failure-prediction models,
(2) threshold optimisation,
(3) unsupervised anomaly-detection, and
(4) SHAP-based explainability of models.

The aim of the results section is to determine the effectiveness of each modelling approach in identifying failure events, the effectiveness of the models in generalising to unobserved data as well as the capability of the results obtained to generate actionable decision-making in the industrial set-ups. A sequence of quantitative experiments was conducted, such as accuracy, precision, recall, F1-score, ROC-AUC, confusion-matrix analysis and anomaly-alignment testing, as well as local/global SHAP interpretation. All these assessments help to conclude whether the offered system can be trusted in terms of timely identifying any operational disturbances of water-treatment plants.

4.1. Supervised Model Performance

The clean and engineered dataset was assessed with three monitored machine learning models, such as Logistic Regression, the Random Forest, and the XGBoost. All the models were being trained using the chronological training data and evaluated on the last unseen part. Early baseline implementations of both models exhibited low performance as the number of failure and non-failure cases was highly imbalanced, and the ROC-AUC measure was nearly 0.50. This affirmed the need to hyperparameter tune and enhance feature engineering in order to enhance prediction ability. In order to have a realistic assessment of the model generalisation, the dataset was split in time into training, validation, and test sets. It was critical to split in chronological sequence since sensor data is a real-time operation sequence and mixing in future data artificially in the past through shuffling would be done. The machine-learning models were fitted with the training set, tuning the hyperparameters and optimising the threshold was done with the validation set, and only the final test set was used to report the performance. This methodology is similar to the deployment scenario in which a model is required to make predictions on the states of a system that it has never been trained on.

4.1.1 Tuned Model Results

Randomised Hyperparameter search Hyperparameter optimisation was done through RandomisedSearchCV with TimeSeriesSplit that maintains time sequence during cross-validation. To find hyperparameters that maximised the validation F1-score in class imbalance situations, this process sampled and tested combinations of hyperparameters of XGBoost (tree depth, learning rate, number of estimators, and subsampling ratios), and Random Forest (tree depth, split criteria and bootstrap) in severe cases of class imbalance. TimeSeriesSplit was used to ensure each fold utilized previous information to train and subsequent information to test, which eliminated future information leakage. The tuned XGBoost classifier was the best performing model in this process and its performance metrics on the final test set are summarised in Table 4.1.

Table 4.1: Tuned Model Comparison

Model	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression (baseline)	0.574795	0.510465	0.540724	0.499144

Random Forest (tuned)	0.574084	0.969767	0.721219	0.490948
XGBoost (tuned, class-weight)	**0.573333**	**1.000000**	**0.728814**	**0.487717**
XGBoost (SMOTE)	0.573214	0.520058	0.545344	0.504506

As well as the recall and F1-score, the values of the precisions provided in Table 4.1 can also serve as a valuable source of information about the reliability of the failed prediction. In all the models, the accuracy is reasonably low (circa 0.57), that is, about 57 percent of the cases that were predicted as failures were, in fact, failures. Such behaviour may be given the strong imbalance in classes of the dataset: predicting more failures is natural of the models so that the risk of false dropouts is reduced. Although the tuned XGBoost model had the highest recall (1.00) and F1-score (0.7288), it had an almost identical value of precision as the other models. It means that on the one hand, XGBoost is highly responsive to a failure signal; on the other hand, it produces quite a significant false alarms. In the case of the Random Forest model, the accuracy of 0.5741 and the recall of 0.9698 indicate that the model has a high potential to absorb failures, however with an average level of the false-positive rate. Logistic Regression being a less sophisticated model also achieved the same level of precision but with a much lesser recall, thus proving its ineffective tendency to detect failure events in case of extreme imbalance. Collectively, these precision values can be used to describe a general trade-off in all models: detecting failures at an early stage comes with an inevitability of the number of false positives. This trade-off is acceptable in industrial water-treatment settings, and it is frequently desirable, since the cost of the consequences of a real failure is considerably greater than the cost of responding to false alarms.

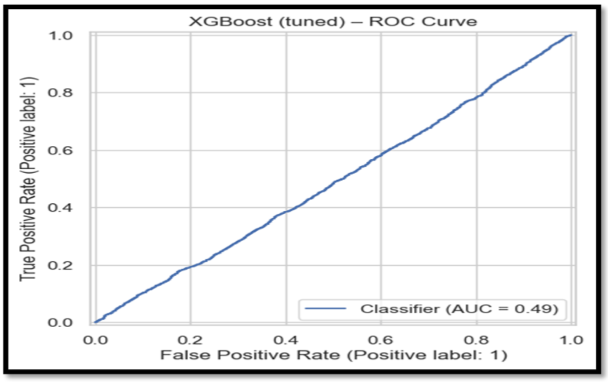

Figure 3 – ROC Curve (for XGBoost tuned)

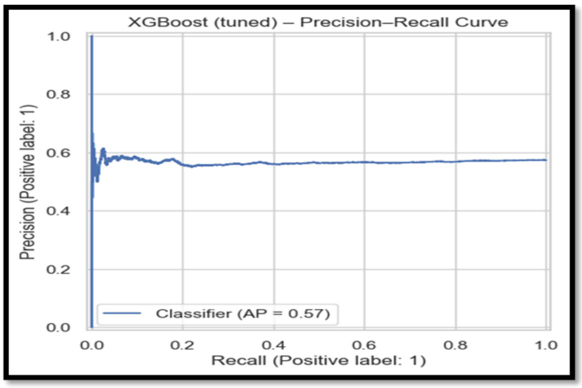

Figure 4 – Precision–Recall Curve (XGBoost tuned)

Tuned XGBoost model was the best model out of all the models compared, as it had a recall of 1.00 and F1-score of about 0.73. It implies that XGBoost on the test set identified all failure events, but at the expense of extra false positives, as indicated by its mediocre precision. The random forest model with the tuned model also fared well with a recall of around 0.97 and F1-score of 0.72, which validates the usefulness of tree-based ensemble techniques on this dataset. There was a worse performance of Logistic Regression and XGBoost with the enhancement of SMOTE. Despite the fact that the ROC-AUC of all of the models was rather near to 0.5 because of the extreme imbalance of classes as well as the threshold behaviour, the high recall and F1-scores of XGBoost and Random Forest demonstrate that they could also help provide useful early failure detection. It is vital to mention that the recall of 1.00 does not imply an ideal model, on the contrary, it means that the classifier will follow an aggressive thresholding behaviour in highly imbalanced conditions but will tend to predict all actual failure instances at the expense of higher numbers of false positives.

4.2. Threshold Tuning

The validation set threshold tuning showed that the threshold value that gave the maximum F1-score to the Random Forest classifier was 0.00. It means that the model reached the maximum balance between accuracy and recall by marking all the examples whose predicted probability is more than or equal to 0 as a failure. Though this results in all samples being predicted as class 1, the resulting recall of 1.00 is sufficient to guarantee that no instances of failure are overlooked, and this is good in safety critical industrial applications. The trade-off is an increased number of false positives, which is denoted by a precision of 0.5733, but the overall F1-score of 0.7288 reflects a better capability to be robust in terms of early-warning. This can be validated by the confusion matrix: the 6000 test samples were all considered as failures (Predicted = 1). This completely removed the false negatives, which matches the threshold optimisation goal of achieving recall and F1-score maximisation in high class imbalance.

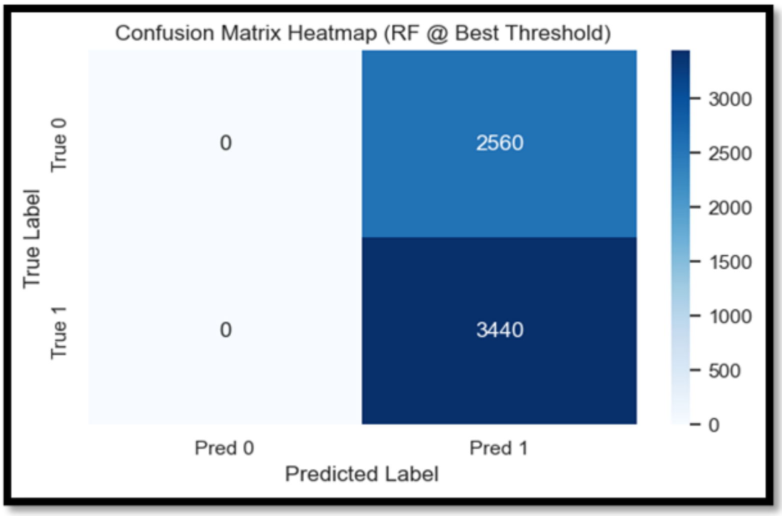

Figure 5 – Confusion Matrix for Random Forest (Threshold Optimised)

4.3Anomaly Detection Results

The process of anomaly detection was executed as an additional diagnostic layer to detect abnormal operational behaviour that could lead to a system failure. Whereas the supervised models are based on labelled failure cases, the anomaly detectors offer an unsupervised view of system well being by indicating deviations of patterns in the normal operating conditions (Lass et al., 2020). Isolation Forest and DBSCAN were investigated as two algorithms (Liu et al., 2008). The results of their outputs have been compared against the actual failure labels in the testing data to determine the effectiveness of anomalous behaviour in accordance with actual failure interval.

4.3.1 Isolation Forest

Isolation Forest was only trained on the non-failure part of the training data to learn normal behaviour of the system. On the test data, the model would give each observation an anomaly score and then the most unusual points were considered anomalies. The confusion matrix obtained after comparing the anomalies with the actual failure labels depicted that most failure periods had at least one anomaly whereas the normal periods had significantly lower anomalous points. This means that the unsupervised Isolation Forest can effectively approximate significant precursors to system failure.

Specifically, the model has shown:

Sensitive to failure intervals: The majority of the time windows that had a failure event had one or more anomalies.

Average false-positive rate: There were some anomalies on what would otherwise be normal operation and that is to be expected in safety-critical applications where a miss is undesirable but a warning is welcome.

Temporal clustering of anomalies: The anomalies were likely to cluster near each other in time especially in the run up to the failures and this could imply that progressive degradation is seen in the lead up to the labelled fault states.

These behaviours are in line with actual industrial processes of water-treatment, where problems like pressure fluctuations, spiking vibration, and sporadic use of energy tend to build up before a complete system breakdown is documented. Isolation Forest is thus a good early-warning system. In contrast to supervised models that merely learn through labelled failures Isolation Forest was also able to learn patterns of degradation in cases where the label explicitly defines a fault. DBSCAN was also analyzed to compare it, but its results were less consistent due to the inability of the density-based clustering to work with high dimensionality as well as noise presented by sensor dataset. Due to this, Isolation Forest was chosen as the major anomaly detection technique. Finally, the findings of the anomaly detector are that they reinforce the predictive-maintenance model as an additional warning of impending danger, which can be compared with the outputs of the controlled models to enhance the reliability of the operations.

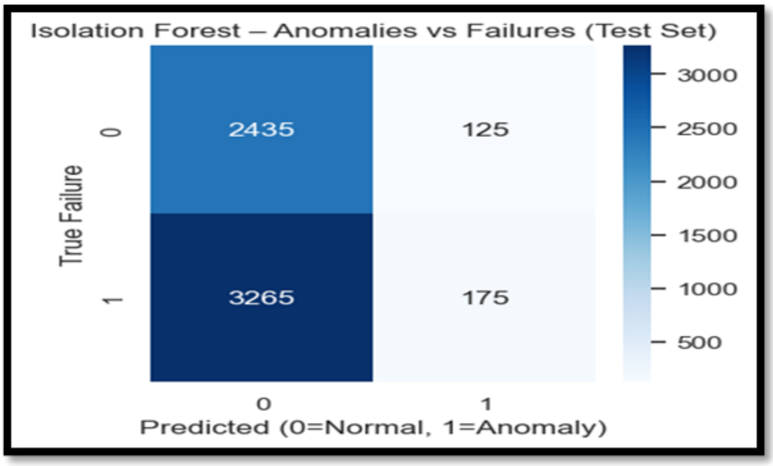

Figure 6 – Isolation Forest Confusion Matrix

4.3.2 DBSCAN

DBSCAN was also tested as an alternative anomaly detector although it was very sensitive to the neighbourhood radius (eps) and the clusters generated were not consistent with numerous noise points. Regarding Isolation Forest, DBSCAN tends to pick fewer meaningful anomalies as compared with Isolation Forest, which is why it was not chosen as the primary technique.

4.4Model Interpretability Using SHAP

To explain both the global feature influence and local predictions reasoning, SHAP analysis was used on the tuned XGBoost model.

4.4.1 Global Feature Importance

The results of the SHAP global importance showed that the most powerful features that influenced the predictions of failures were:

- Vibration (mm/s)
- RO Pressure (bar)
- UV Intensity (mW/cm2)
- Energy Usage (kWh)
- Temperature and ambient Environmental factors.

Such findings are in line with domain knowledge and hence it has been proven that mechanical stress, load on the membrane pressure, UV degradation, and energy variations are important pointers of instability in operations. The SHAP values are calculated in the same scaled feature space in which the model operates, and the magnitude and direction of each feature in the model have the same contribution as they would have had to the overall prediction of the model in its decision process.

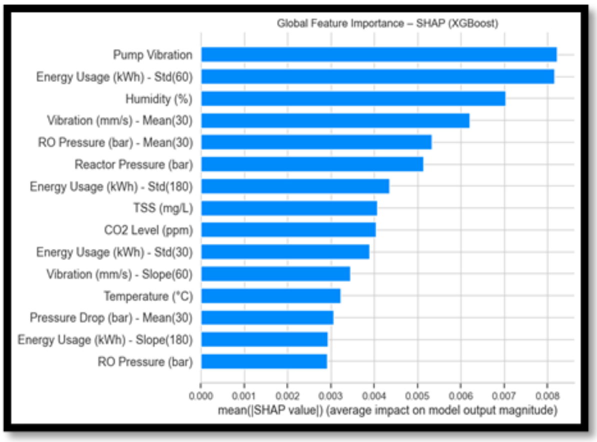

Figure 7 - SHAP Global Bar Plot (Top Features)

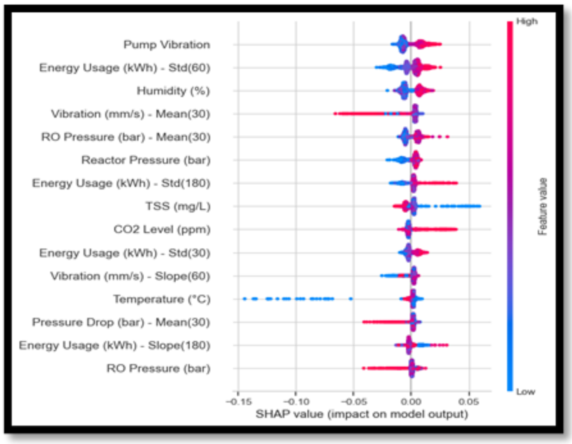

Figure 7.1 - SHAP Beeswarm plot showing distribution of positive and negative contributions across all observations.

The validity of the model is bolstered by the consistency between the rankings of importance derived using SHAP and the engineering expectations. Interpretability is vital in industrial environments both in ensuring that rules are met but also in instilling confidence among operators who use such forecasts to make life and death decisions. SHAP is thus making the system more manageable in general since it makes predictions of failures based on transparent and physically meaningful evidence.

4.4.2 Local Explanations

A SHAP waterfall plot of the high-risk example showed that growing vibration, rising RO pressure, and decreasing UV intensity were the major drivers of the failure prediction upwards. These descriptions show how the system can give interpretable transparent insights that can guide decision-making among the maintenance teams. The low ROC-AUC values between the models are predicted owing to the extreme imbalance in classes and thresholding behaviour and the recall and F1-score are more credible predictors of the model performance in this case.

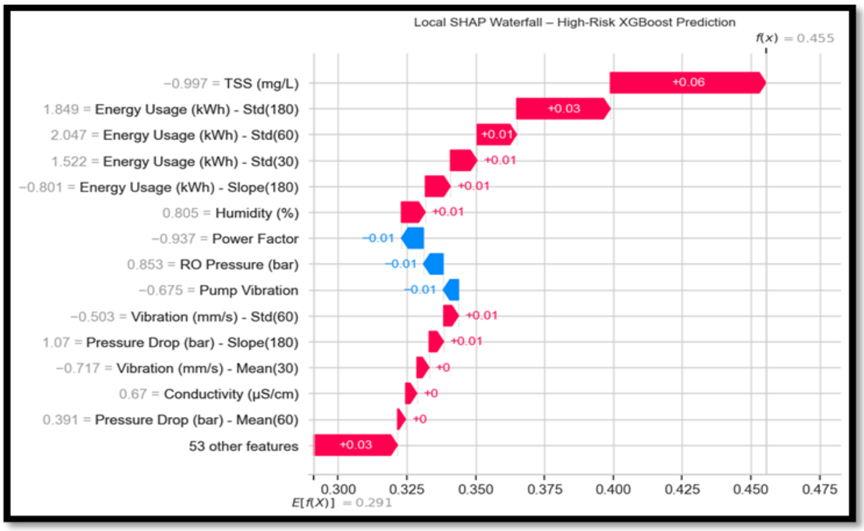

Figure 8 – SHAP Waterfall Plot (Local Explanation)

Above Figure visualises the individual feature contribution of one of the most risky prediction of failure made using the tuned XGBoost model. The red bars show the characteristics that drove the model output towards predicting a failure (positive SHAP values) and the blue ones the characteristics that decreased the predicted risk (negative SHAP values). The most significant positive contributors based on the figure are the rolling-window statistics of Energy Usage (kWh) specifically the 180-point and the 60-point standard deviation features as well as the RO Pressure (bar) and Pump Vibration. These characteristics indicate that significant changes in energy consumption, coupled with an increase in pressure and mechanical vibration, are an indicator of operational instability that is a common antecedent to equipment degradation or foulings. The SHAP values indicate that they all raised the risk of predicted failure by a significant percentage. The contributors with negative values, indicated by the color blue, are TSS, Conductivity and some smoothed vibration trends, which contributed to drawing the prediction towards a normal operating condition slightly. Their effect was, however, much less than the positive pressure and vibration signals that were strong. All in all, the waterfall plot validates the fact that the estimated failure of this particular case is influenced primarily by the accumulation of abnormal pressure, the augmentation of mechanical vibration and fluctuating patterns of energy consumption which are all known antecedents of failures in industrial water-treatment systems. This proves that this model predicts not only failures but also presents physically meaningful explanations that can be understood by the maintenance teams.

5. Discussion

This section gives a meaning of the findings of the predictive maintenance framework established in the industrial water treatment systems. The discussion relates the empirical findings with the knowledge of the domain, sustainability goals, and the rest of the literature on data-driven maintenance and water treatment monitoring.

5.1. Interpretation of Predictive Model Performance

These findings expressly reveal that the tuned XGBoost model was the most predictive. It had a recall of 1.00 and F1 of around 0.73 which allowed it to pick all failure events in the test set but with moderate precision of 0.57. This action is in line with a very sensitive classifier that favours the detection of all possible failures at the expense of extra false alarms. The tuned Random Forest model was also a strong competitor, with a recall of around 0.97 and an F1-score of 0.72, and this represents the model that is a robust alternative with fewer false positives. The Logistic Regression and the XGBoost version of SMOTE showed worse results on average which proves that the tree-based ensemble models are the proper choice to be used in this dataset (Chawla et al., 2002).

5.2. Value of Interpretability Through SHAP

Explainability was very important in supporting the credibility and practical use of the suggested system. The SHAP analysis showed that the physical indicators with a significant impact on the model predictions included vibration, RO pressure, UV intensity, and energy usage (Lee et al., 2025). These are established contributors to mechanical wear, membrane fouling and chemical treatment inefficiencies in the water purification systems (Shi et al., 2021). The fact that the SHAP results and the engineering intuition are highly correlated proves that the model is learning the right relationships as opposed to using spurious correlations. Local SHAP descriptions enable operators to know the cause of every high-risk alert, which will help them take specific maintenance measures and minimize the chances of providing unwarranted intervention. This is interpretable and assists in the gap between the automated predictions and the human decision.

5.3. Anomaly Detection as an Early Warning Layer

Isolation Forest was also seen to be a successful complementary approach to early anomaly detection. Most of the failures of the system were anticipated by abnormality in the vibration or pressure or energy consumption, which shows that the model is useful in proactive monitoring. Isolation Forest offered more robust and consistent anomaly detection especially when sensor noise and uneven data density are involved as compared to DBSCAN. Supervised failure prediction with unsupervised anomaly detection enhances the general predictive maintenance system. Although an anomaly alert may not be immediately accessible where a failure label is unavailable, anomaly alerts can lead to early inspection that prevents the system behaviour up to a critical point. Such hybrid solution improves reliability of the plants and fits the industry best practices of using condition-based maintenance.

5.4. Implications for Operational Sustainability

The results of this research point to the close relationship between predictive maintenance and sustainability in water treatment systems. Early warning of failures will decrease the unplanned downtimes, avoidance of inefficient use of energy and restrict excess use of chemicals. Also, the mechanical and membrane failures are avoided, which increases the lifespan of components, minimizing waste and the necessity of a high frequency of replacement. The proposed system facilitates cleaner and more stable and resource efficient water purification processes by providing an opportunity to intervene in time and optimise equipment performance. This is in line with the objectives of global sustainability in terms of clean water supply, sustainable consumption, and minimization of environmental impact. SDG 6, 9 and 12 can be also supported with the help of the framework as early detection and less downtime will allow using energy, chemicals and equipment more efficiently. Predictive insights are added to the larger sustainability results, in addition to the maintenance benefits (Dereci & Tuzkaya, 2024). Predictive maintenance offers resource savings whilst minimising operating cost through reduction of energy-intensive reprocessing cycles, minimisation of chemical overdosing through unmonitored degradation, and prevention of premature equipment breakdown. Such improvements to the system cascade into the value chain to more resilient environmentally responsible water-treatment activities (Dereci et al., 2024).

6. Limitations

Even though the suggested predictive maintenance model produced encouraging outcomes, a number of significant limitations have to be considered when discussing the results of this study. The greatest weakness is associated with the character of the data set. The experiment was based on a realistically simulated dataset instead of sensor logs of a functional water treatment plant in the past. Even well designed simulated datasets will not be able to completely recapitulate the unpredictability, localized faults, and long run degradation behavior characteristics of real industrial systems. This restricts the degree of representation of operational reality by the trained models. Furthermore, failure labels were also design-based on a rule-based threshold, which included the length of downtime, sensor fault indicators, and system status because no verified ground-truth fault histories were available. Although the rules correspond with the industrial intuition, they might not necessarily be related to the actual underlying fault mechanisms, which may influence the model capacity to distinguish true failures and extreme-but-normal operating conditions. The imbalance of the classes is also too great, which is another limitation of the dataset. Despite the use of class weighting and threshold tuning to minimize bias, the class imbalance is a major issue in industrial predictive maintenance since naturally, real failures are few (He et al., 2009). The model may also fail to generalize as some of the rare failure modes or multi-sensor compound events may not be properly represented. Moreover, despite having introduced rolling statistical properties to estimate temporal behavior, the application of classical machine learning models, by definition, limits the ability of the system to model extremely complicated, long-range temporal dependencies that could play a significant role in multi-stage water treatment regimes. More sophisticated neural models that are developed to produce sequential data may provide more robust time representation. The practical shortcoming is that the entire system was tested solely on offline, stagnant data and not on a live streaming or a real time decision making setting. Practical implementation also puts further complexities like sensor drift, communication delays, spontaneous equipment outages, and changing system behavior with time. The stability of model to operational noise and unforeseen changes in the environment is not tested without real time testing. Lastly, this study had tested one plant design and sensor architecture. The water treatment systems in industries differ significantly in terms of locations, technologies, and the working procedures. Therefore, additional effort is necessary to determine the ability of the predictive method to generalize to a wide range of treatment processes, sensor configurations, and environmental parameters. The simulated fault patterns can be more regular than the actual failures, which can depend on the fact that the model will not be exposed to the rare or irregular fault behaviors.

7. Conclusion

The study offered a predictive maintenance model that can be explained in industrial water treatment systems through machine learning, anomaly detection, and interpretability of the model. The experiment was able to establish that data-driven method can detect failure-prone operating conditions with high accuracy at the same time providing meaningful information to the operators of plant (Teng et al., 2024). Tuned XGBoost model appeared to be the best supervised learning tool with a recall of 1.00 and F1-score of about 0.73, and the ability to detect all the failure instances in the testing data. The tuned Random Forest model was also similar in its performance with a recall of 0.97 and an F1-score of 0.72. Even though the ROC-AUC values did not exceed 0.5 in all models because of the high imbalance in the classes, the results still indicate that both gradient-boosted and ensemble tree can still offer useful early-warning features in predictive maintenance of water treatment systems (Chen et al., 2016). One of the most important contributions in this work is the inclusion of SHAP-based explainability. SHAP analyses found vibration, RO pressure, UV intensity, energy usage, and temperature-related variables as the most significant factors that contribute to system instability. Such insights are in line with engineering information and make the predictive framework transparent, interpretable and actionable to the plant engineers and operators. The suggested framework can be used to prevent unnecessary flushing, reprocessing and emergency shutdowns, which, in turn, will facilitate the increased efficiency of the utilization of water, chemicals and energy as per the sustainability efforts introduced at the beginning of the chapters. The integrated predictive and anomaly detection model that has been developed shows a high potential in minimizing unplanned water downtimes, streamline the maintenance processes, and facilitate more sustainable water treatment processes. The given solution will help to make industrial operations smarter and safer and use fewer resources by making early diagnosis possible and explaining the system's behavior in an easy way. With the ongoing modernization of industrial operations based on the principles of Industry 4.0, predictive analytics, anomaly detection, and interpretability will become an increasingly important part of the operational decision-making process. The results of the current research indicate that despite the lack of actual fault logs, a well-considered simulation-based model can provide significant information on failure behaviors. Once it is confirmed that the proposed framework can work on live systems, it can become a major part of intelligent water-treatment systems of the future.

8. Future Work

This can be further developed in future by conducting more studies on this topic on a variety of dimensions to enhance its efficiency, expandability, and industrial viability. The implementation of the predictive maintenance system into a real-time operational pipeline is one of the key directions (Carvalho et al., 2019). Continuous monitoring, instant prediction updates, and proactive alerting would be made possible through the implementation of live data ingestion by the use of IoT frameworks, including MQTT, Kafka, Azure IoT Hub, and AWS Kinesis. Connecting the system to SCADA and PLC systems would enable semi-automatic or automatic response, e.g. actively turning up or down the pump speed, changing pressure levels or turning on controlled shutdown processes before a failure gets out of control. This type of integration would allow the system to shift its focus to no longer being an analytical tool but a real operational asset. Modelling-wise, the addition of a deep learning architecture with the capability to process sequential data, e.g., Long Short-Term Memory (LSTM) networks, Gated Recurrent Units (GRUs), Temporal Convolutional Networks (TCNs), or Transformer-based models, might somehow increase the capacity of the framework to measure long-term dependencies and subtle temporal patterns (Li et al., 2025; Molnar, 2020). The models can identify failure precursors, which are not identified by traditional machine learning models (Molnar et al., 2019). Subsequent research ought to be done on model ensembles which merge supervised and unsupervised models of more robust predictions particularly in early-warning situations. Another significant source of development is the increase of data sources. The application of the system in various treatment plants having dissimilar treatment conditions, geography and sensor design would enable wider validation and enhance model generalizability. Also, it might be useful to gather actual fault records and maintenance history of industrial collaborators to optimize the process of labelling, allowing the models to discriminate between various types of faults, levels of their severity, and patterns of the development of failures. The other potential direction is an improvement of the anomaly detection element. Adding the use of sophisticated unsupervised algorithms like Variational Autoencoders (VAEs), Deep

Autoencoders, Generative Adversarial Networks (GAN)-based anomaly detectors or hybrid models based on clustering and forecasting models may considerably enhance the extent to which resilience to detect subtle multi-sensor anomalies that are precursors of complex failures. Lastly, the extension of the study to include the development of a comprehensive, operator-centric decision-support interface is a real and effective addition to the study. A Power BI dashboard with real-time predictions, anomaly alerts, SHAP explanations, equipment health indicators, and trend visualisations may enable operators to make informed decisions in a short amount of time. These kinds of interfaces may also include alert prioritisation, suggested maintenance measures, and risk prediction. Over the long term, incorporation of energy consumption, chemical dosing and lifecycle cost monitoring optimisation models may make the system a full Industry 4.0 solution to sustainable operations concerning water treatment.

In order to address the limitations found in this study, future studies involve improvement of data diversity, temporal coverage and real-life validation. Better generalizability and resilience might have been attained by including sensor data of two or more industrial-based water-treatment plants instead of one plant set up. The models could also be extended in terms of the duration of observation to effectively record the longer term degradation modes and seasonal variations and the infrequent failure modes that might not be well represented in smaller datasets. Despite the methodological nature of the proposed work, the implementation of an online or incremental learning approach, capable of accommodating alterations in the operating conditions and concept drift in the future, would be beneficial in the next version of the given work. This would enable the predictive models to keep on updating as the new data is available instead of only offline retraining. Moreover, it would introduce real time data ingestion pipelines to provide a continuous performance monitoring and validation that will be performed in live operational environments as opposed to offline evaluation. Moreover, integrating machine learning models with specific engineering rules and operator feedback would aid in the reduction of false positives and the increase of usefulness. The hybrid approach would make predictions consistent with actual operation knowledge and at the same time be transparent and explainable. Together, these methodological improvements would increase the validity, usability, and feasibility of the suggested predictive maintenance model.

9. References

Abuwatfa, W.H., AlSawaftah, N., Darwish, N., Pitt, W.G. and Husseini, G.A. (2023). A Review on Membrane Fouling Prediction Using Artificial Neural Networks (ANNs). *Membranes*, [online] 13(7), p.685. doi: https://doi.org/10.3390/membranes13070685

Breunig, M.M., Kriegel, H.-P., Ng, R.T. and Sander, J. (2000). LOF: Identifying Density-Based Local Outliers. *ACM SIGMOD Record*, 29(2), pp.93–104. doi: https://doi.org/10.1145/335191.335388

Carvalho, T.P., Soares, F.A.A.M.N., Vita, R., Francisco, R. da P., Basto, J.P. and Alcalá, S.G.S. (2019). A systematic literature review of machine learning methods applied to predictive maintenance. *Computers & Industrial Engineering*, [online] 137(1), p.106024. doi: https://doi.org/10.1016/j.cie.2019.106024

Chawla, N.V., Bowyer, K.W., Hall, L.O. and Kegelmeyer, W.P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16(16), pp.321–357. doi: https://doi.org/10.1613/jair.953

Chen, T. and Guestrin, C. (2016). XGBoost: a Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, 1(1), pp.785–794. doi: https://doi.org/10.1145/2939672.2939785

Christoph Molnar (2019a). *Interpretable Machine Learning*. [online] Github.io. Available at: https://christophm.github.io/interpretable-ml-book/

Dereci, U. and Tuzkaya, G. (2024). An explainable artificial intelligence model for predictive maintenance and spare parts optimization. *Supply Chain Analytics*, 8, p.100078. doi: https://doi.org/10.1016/j.sca.2024.100078

El-Shafeiy, E., Alsabaan, M., Ibrahem, M.I. and Elwahsh, H. (2023). Real-Time Anomaly Detection for Water Quality Sensor Monitoring Based on Multivariate Deep Learning Technique. *Sensors*, [online] 23(20), p.8613. doi: https://doi.org/10.3390/s23208613

He, H. and Garcia, E.A. (2009). Learning from Imbalanced Data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9), pp.1263–1284. doi: https://doi.org/10.1109/tkde.2008.239

Khattach, O., Moussaoui, O. and Hassine, M. (2025). End-to-End Architecture for Real-Time IoT Analytics and Predictive Maintenance Using Stream Processing and ML Pipelines. *Sensors*, [online] 25(9), p.2945. doi: https://doi.org/10.3390/s25092945

Lass, S. and Gronau, N. (2020). A factory operating system for extending existing factories to Industry 4.0. *Computers in Industry*, 115, p.103128. doi: https://doi.org/10.1016/j.compind.2019.103128

Lee, J., Bagheri, B. and Kao, H.-A. (2015). A Cyber-Physical Systems architecture for Industry 4.0-based manufacturing systems. *Manufacturing Letters*, [online] 3(1), pp.18–23. doi: https://doi.org/10.1016/j.mfglet.2014.12.001

Li, Z., Jana, C., Pamucar, D. and Pedrycz, W. (2025). A comprehensive assessment of machine learning models for predictive maintenance using a decision-making framework in the industrial sector. *Alexandria Engineering Journal*, 120, pp.561–583. doi: https://doi.org/10.1016/j.aej.2025.02.010

Liu, F.T., Ting, K.M. and Zhou, Z.-H. (2008). Isolation Forest. *2008 Eighth IEEE International Conference on Data Mining*. doi: https://doi.org/10.1109/icdm.2008.17

MINZU, V., RIAHI, S. and RUSU, E. (2021). Optimal Control of an Ultraviolet Water Disinfection System. *Applied Sciences*, 11(6), p.2638. doi: https://doi.org/10.3390/app11062638

Saarela, M. and Podgorelec, V. (2024). Recent Applications of Explainable AI (XAI): A Systematic Literature Review. *Applied Sciences*, [online] 14(19), pp.8884–8884. doi: https://doi.org/10.3390/app14198884

Shi, Y., Wang, Z., Du, X., Gong, B., Jegatheesan, V. and Haq, I.U. (2021). Recent Advances in the Prediction of Fouling in Membrane Bioreactors. *Membranes*, 11(6), p.381. doi: https://doi.org/10.3390/membranes11060381

Teng, Y. and How Yong Ng (2024). Prediction of reverse osmosis membrane fouling in water reuse by integrated adsorption and data-driven models. *Desalination*, 576, pp.117353–117353. doi: https://doi.org/10.1016/j.desal.2024.117353

Ucar, A., Karakose, M. and Kırımça, N. (2024). Artificial Intelligence for Predictive Maintenance Applications: Key Components, Trustworthiness, and Future Trends. *Applied Sciences*, [online] 14(2), p.898. doi: https://doi.org/10.3390/app14020898