Vaasan yliopisto
UNIVERSITY OF VAASA

Muhammad Hassan

AI-Based Conditional Monitoring & Predictive Maintenance for Offshore Wind Farms

Master Thesis

School of Technology and Innovation
Master’s Thesis in Technology
Master’s Programme in Smart Energy

Vaasa 2025

Acknowledgement

I would like to express my deepest gratitude to my parents for their constant love, support, and encouragement throughout my academic journey. Their unwavering belief in me has been a source of strength and motivation, and I am grateful for everything they have done for me.

A special and heartfelt thank you goes to my supervisor, Prof. Mazaher Karimi, for providing me with the opportunity to work on this research project. Your expertise, support, and encouragement throughout this process have been invaluable. I deeply appreciate your trust in me and the chance to contribute to this exciting area of study. Your constructive feedback and mentorship have significantly enriched this work.

I am grateful to the InnoWind project for providing an inspiring research environment and the resources necessary to carry out my work.

I would also like to thank Mojtaba for his invaluable assistance, insightful discussions, and collaborative spirit throughout my research. His help has made this journey more rewarding.

Finally, I would like to extend my gratitude to all those who have contributed to my research, either directly or indirectly. Your support has been crucial in the successful completion of this thesis.

Muhammad Hassan

UNIVERSITY OF VAASA
School of Technology and Innovation
Author: Muhammad Hassan
Title of the thesis: AI-Based Conditional Monitoring & Predictive Maintenance for Offshore Wind Farms: Master Thesis
Degree: Master of Science in Technology
Discipline: Master's Programme in Smart Energy
Supervisor: Mazaher Karimi
Year: 2025 **Pages:** 105

ABSTRACT:
Offshore wind farm maintenance is challenged by harsh marine conditions, high operational costs, and significant safety risks associated with manual inspections. To address these challenges, a predictive maintenance framework is proposed that integrates high-frequency SCADA data with advanced artificial intelligence (AI) techniques. The framework is grounded in the key concepts of condition monitoring, anomaly detection, and remaining useful life estimation, and it is built upon theories from time-series analysis, statistical signal processing, and machine learning.

A systematic data pipeline is developed, beginning with timestamp alignment, outlier filtering, and gap interpolation to ensure data integrity. Dimensionality reduction is achieved through incremental principal component analysis, while feature selection is conducted via correlation filtering and ensemble-based importance ranking. Two complementary data representations are crafted: fixed-length multivariate sequences for recurrent and convolutional neural networks, and static feature snapshots for a gradient-boosted decision-tree ensemble. Three predictive models an LSTM-based recurrent network, a one-dimensional convolutional neural network, and a LightGBM classifier are trained and validated on subsets of the publicly available "CARE to Compare" wind turbine dataset. Temporal partitioning prevents information leakage, and class-imbalance strategies, including stratified sampling and noise-injection oversampling, are employed to enhance sensitivity to rare fault events.

Model performance is evaluated using accuracy, precision, recall, and F_1-score on held-out test segments. The deep-learning approaches are shown to capture both gradual drifts and transient spikes in sensor channels, providing early-warning windows of multiple days before predicted component failures. A real-time SCADA simulation demonstrates how staged alerts across mechanical, electrical, and environmental signals can be used to orchestrate proactive maintenance schedules, thereby reducing dependency on emergency vessel mobilizations and risky platform visits. The results indicate that sequence-aware neural networks outperform static, snapshot-based methods in fault prediction tasks, validating the critical role of temporal context in offshore condition monitoring. The thesis concludes that AI-driven predictive maintenance not only enhances turbine availability and grid reliability but also promotes crew safety and cost efficiency. Recommendations for future work include integration of environmental satellite data, deployment of edge-computing solutions for on-site inference, adoption of explainable AI techniques to build operator trust, and exploration of federated learning paradigms to generalize models across multiple wind farms without sharing raw data.

KEYWORDS: Offshore Wind Energy, Predictive Maintenance, Condition Monitoring, SCADA Data, Machine Learning, Anomaly Detection, Deep Learning, Principal Component Analysis

Contents

Figures

Tables

Abbreviations

AI	Artificial Intelligence
AUC	Area Under the Curve
BN	Batch Normalization
CC BY-SA	Creative Commons Attribution-Share Alike (4.0 International)
CFD	Computational Fluid Dynamics
CM	Condition Monitoring
CNN	Convolutional Neural Network
CPU	Central Processing Unit
CSV	Comma-Separated Values
FC	Fully Connected (layer)
GPU	Graphics Processing Unit
IQR	Interquartile Range
IoT	Internet of Things
LSTM	Long Short-Term Memory
LightGBM	Light Gradient Boosting Machine
MAD	Median Absolute Deviation
ML	Machine Learning
NAN	Not a Number
O & M	Operations and Maintenance
PCA	Principal Component Analysis
PdM	Predictive Maintenance
RAM	Random-Access Memory
ROC	Receiver Operating Characteristic
ReLU	Rectified Linear Unit
RNN	Recurrent Neural Network
RUL	Remaining Useful Life
SCADA	Supervisory Control and Data Acquisition
STATCOM	Static Synchronous Compensator
t-SNE	t-distributed Stochastic Neighbor Embedding
UTC	Coordinated Universal Time
URL	Uniform Resource Locator
VSC	Voltage Source Converter
XAI	Explainable Artificial Intelligence
XLPE	Cross-Linked Polyethylene

1 Introduction

The modern energy system relies on both physical infrastructure and advanced digital technologies. Traditionally, the grids have been centralized grids dependent on fossil fuel generated power and had maintenance routines with manual oversight to back them up. However, in recent decades environmental imperatives, technological innovations and increasing consumer demand for sustainable and resilient energy have driven a transformation in energy systems. This evolution moves from centralized fossil fuel dependent infrastructures toward decentralized and flexible networks powered by renewable energy sources. A key aspect of this transformation is the integration of diverse energy sources to enhance system performance and resilience. As shown in **Figure 1**, the transition from single energy systems to hybrid renewable energy systems (photovoltaic/wind) demonstrates a shift toward more flexible and complementary energy generation, enhancing the reliability and guarantee rate of electricity production.

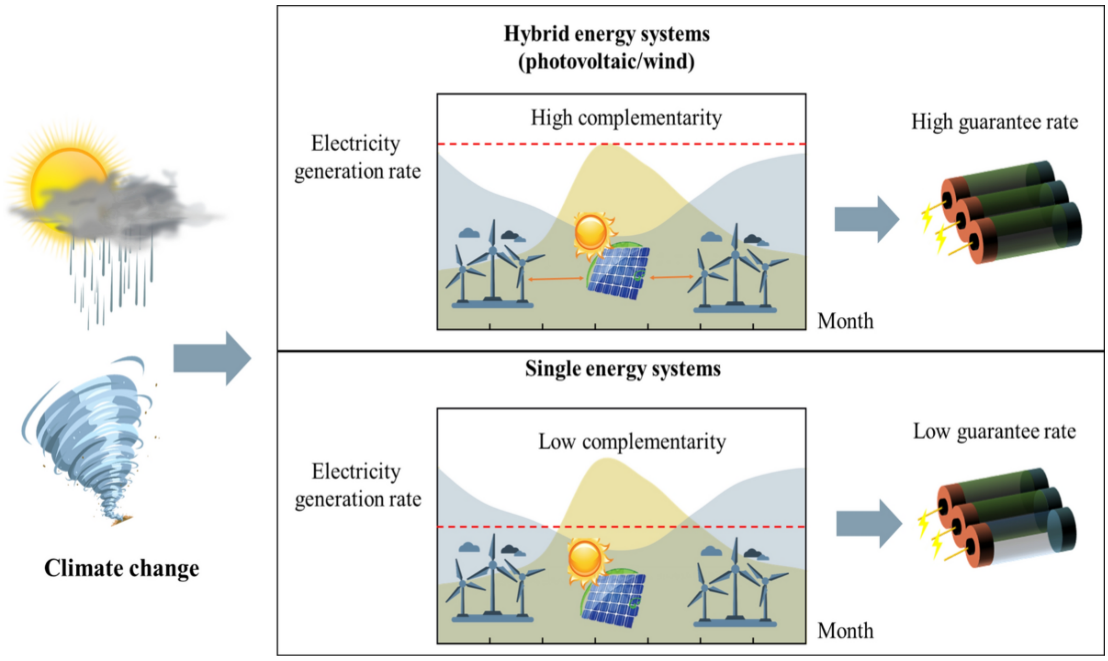

Figure 1. Hybrid Energy system (Farghali et al., 2023).

With the rapid integration of renewable energy sources, particularly offshore wind farms, the operational challenges have also evolved significantly. Renewable energy systems,

unlike traditional fossil-fuel based grids, often involve geographically isolated assets operating under harsh environmental conditions. This distributed and challenging operational landscape necessitates robust digital infrastructures for continuous monitoring and maintenance. Consequently, leveraging sophisticated sensor networks and SCADA systems becomes critical not only for efficient operation but also for the timely detection and mitigation of faults, thereby ensuring reliable energy generation.
Addressing these challenges is made possible by the digitalization of energy systems that break the silos of data monitoring, smart briefly calls it automated control strategies and predictive analytics. Rather, we are today at a point where relatively sophisticated sensors and SCADA systems have been able to continuously harvest massive quantities of operational data synergizing the physical and digital assets. A new paradigm of such kind does not only provide operation efficiency enhancement but also involves an important level of complexity in management and interpretation of complex datasets. Integration with digital tools permits more proactive management of physical assets through improvement in faulty detection and covering the way for new ways of managing the things predictive maintenance.

However, simply collecting huge volumes of sensors and SCADA data does not prevent equipment failures. To turn raw data into actionable insight, operators must first establish condition monitoring frameworks defining baseline performance metrics, anomaly thresholds, and automated alerts. Building on that foundation predictive maintenance applies statistical and machine-learning models to forecast remaining useful life for key components, prioritize work orders, and schedule service windows long before a failure occurs. **Figure 2** illustrates the two-stage digital pipeline from data collection to predictive analytics, emphasizing the importance of preprocessing and fault detection before applying predictive models to forecast failures.

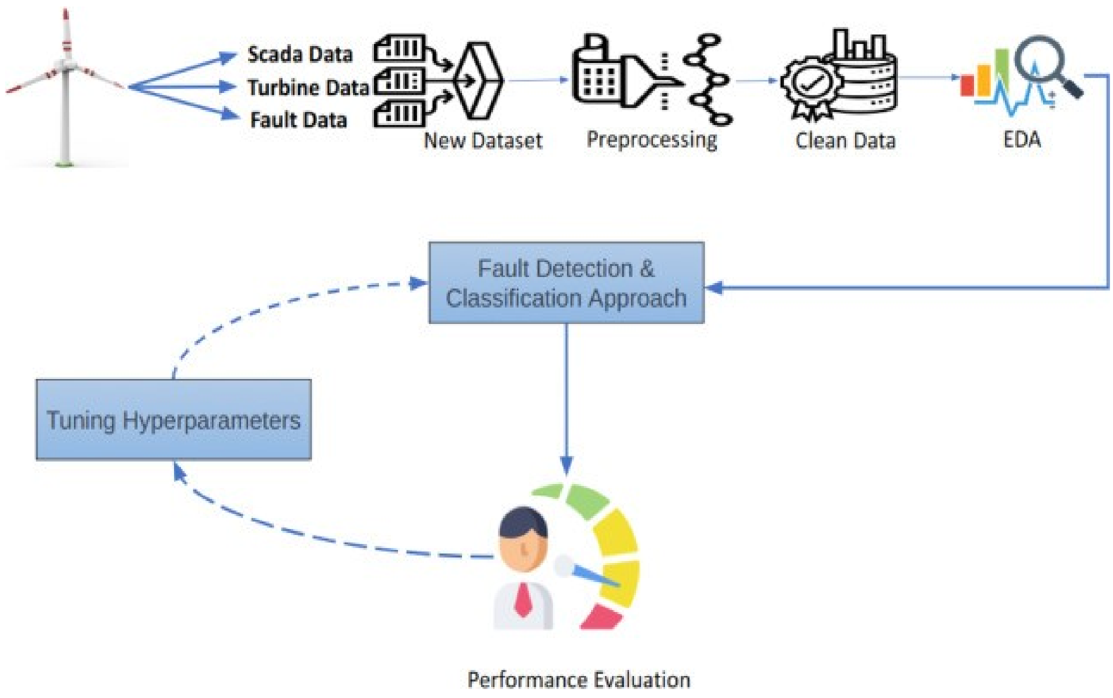

Figure 2. AI for wind farm fault detection (Allal et al., 2024).

The modern energy systems are being driven by the two prime goals of decarbonization and a more resilient grid to tolerate the peculiarities of integrating renewable energy, especially in the difficult settings of offshore windfarms. These systems being increasingly connected and more means more ways to optimize performance and cost with optimization, and with greater size and complexity of managing the management of large scale and real time operations there are even more ways to do the same.

1.1 Background

Offshore wind energy has become one of the fastest-growing segments in the renewable energy sector, (Bilgili & Alphan, 2022) providing substantial opportunities to meet global power demands sustainably. Offshore is better, because it gets stronger and more consistent winds, which are helpful for electricity production. Many of the typical challenges faced when operating in these harsh environments include harsh weather conditions, the salinity of the environment, and the inability of us to perform onsite inspections (Adumene et al., 2023).

In general, wind farm operators have historically relied on scheduled or reactive maintenance approaches: equipment was planned for service at regular intervals or upon failure. Unexplained breakdowns of turbines offshore are more difficult and costly to access and repair than those onshore, causing significant downtime and cost (Carroll et al., 2016). As a result, there has been a need for the development of complex, data driven approaches to predict faults in equipment before catastrophic failures.

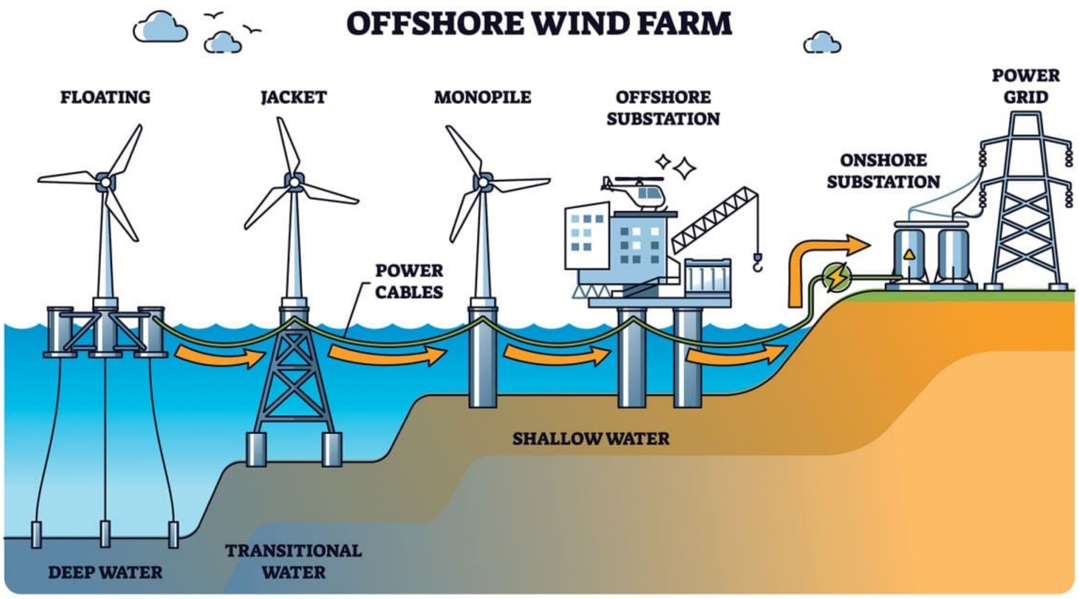

Figure 3. Offshore wind energy, Source: [Wind Energy Association]

As a result of recent developments in sensor technologies and increasing providers of Supervisory Control and Data Acquisition (SCADA) feeds, more complicated condition monitoring applications are now feasible. They are continual sequences collecting large volumes of real time data on (vibration, temperature, rotational speeds) of the run's turbine performance and (wind speed, wave height etc.) from the environment. Additionally, this is also making use of artificial intelligence (AI) and machine learning

(ML) models to analyze such sensor data in a unique way from a reactive and scheduled too predictive and prescriptive.

Against this backdrop, AI-based condition monitoring has gained wide attention, serving as a key enabler for predictive maintenance in offshore wind farms(Maron et al., 2022). By drawing from large datasets and applying techniques such as anomaly detection, time-series analysis, and deep learning, using operators, one can detect subtle signs of equipment degradation and schedule proactive maintenance. It helps reduce operational expenses, meaning that your components being used are often the most critical and that contributes to more reliable energy production.

1.2 Research Motivation

Offshore operations typically encounter greater Operations and Maintenance (O&M) costs than onshore wind farms due to logistical challenges, weather constraints, and specialized vessel requirements. Any unplanned failure magnifies these costs, making predictive strategies extremely valuable (Maples et al., 2013).

Offshore turbines face intense environmental conditions, intense winds, high humidity, and corrosion from saltwater. Consequently, components such as blades, gearboxes, and bearings degrade faster, demanding more robust monitoring systems to ensure uninterrupted functionality (Petersen & Malm, 2006).

Safety and reliability are also first since offshore interventions expose the workforce to dangerous sea conditions. Maintenance can more effectively be scheduled with the help of AI-driven predictive maintenance reducing the number of risky on-site operations and thereby enhancing overall safety (Islam et al., 2017).

Abundant opportunities for advanced analytics are created with the proliferation of sensors and the growing availability of SCADA data. High dimensional time series data is a difficult ground for traditional rule based or simple statistics-based approach; however, machine learning models can help uncover complex patterns in the data (Neeraj & Maurya, 2020).

Last, as the offshore wind industry itself expands, industry trends as well as regulatory requirements alike are driving the demand for increasing the levels of reliability and cost minimization. The environment therefore compels operators, insurers and alike to advance digital decomposition, for example the application of AI based condition monitoring systems, to fulfil these higher expectations.

1.3 Research Problem

Despite substantial motivation to increase wind energy reliability, critical challenges impede the development and use of AI based condition monitoring for offshore wind farms at a large scale. The main problem has to do with data quality and integration. Data is often noisy or incomplete because of sensor malfunction, network issues, and the harsh environs in the marine domain. It is much more complex as the aggregate and cleaning of these diverse datasets an difficult task (Balla et al., 2023).

Another significant challenge is ensuring model robustness and interpretability. Although AI models (Ross & Doshi-Velez, 2018), particularly deep learning networks, are highly effective at pattern recognition, their complexity can obscure the decision-making process. However, in safety critical applications where these transparent and reliable systems are required to gain an operator’s trust in the technology in mission critical scenarios, this opacity is a problem.

Further, there is a significant barrier related to the lack of availability of labeled fault examples. There is a sparsity in the number of recorded component failures in many offshore wind farms, making it difficult to build and validate predictive models for certain values of the training data set. In addition, this challenge is further exacerbated to accommodate different turbine designs. However, offshore turbines of distinctive designs, capacities, sensor configurations, and so on make it difficult to deploy a one-size-fits-all AI model that is accurate enough to consider the differences across turbine types or farm layouts.

Finally, scalability and real time processing are key issues. Current computational resources reach the limits of the need to process high frequency data in real time or near real time for hundreds of turbines. The need to balance accuracy, speed, and scalability of the solutions so that the solutions are not only theoretically sound but also practical enough for an operational setting.

The research problem overall entails development and validation of robust, interpretable, and scalable AI based solutions for accurate condition monitoring and fault prediction in the presence of unique constraints in the offshore wind farms.

1.4 Scope and Objectives

The research objectives seek to tackle the challenges found in the study, to progress the capabilities of AI empowered condition monitoring systems for offshore wind assets. The thesis aims to bring a robust framework by designing a complete data pipeline consisting of the preprocessing of high frequency sensors or SCADA data using the sophisticated processing algorithms and the AI or Machine Learning model. It provides a basis for the advanced predictive maintenance.

On this basis, the study attempts to implement predictive maintenance algorithms by designing, training and strictly evaluating a range of machine learning and deep learning algorithms. Aided by these algorithms (known colloquially as 'anomaly detectors,' 'prognostics,' or 'failure predictors'), these algorithms predict failures in or assess performance degradation of key turbine components including gearboxes, bearings, and blades.

It is an essential element of research. When directly accessible proprietary data is unavailable, the real-world data will be benchmarked through access to publicly available data. This comparative analysis will enable the comparison between the efficacy result of the approach compared to the baseline methods and the state of art.

This thesis provides much attention on both predictive performance as well as on the interpretability of and applicability of the model. The output of the AI models will be made understandable to the operational staff and would be harmonized with industry safety protocols. Finally, the study will assess computational requirements of deployed models, as well to assess how seamlessly these methods can be integrated into existing offshore wind farm management systems.

In the end, the research findings will be used to define the recommendations and best practices. The guidelines to be provided for wind farm operators, Original Equipment Manufacturers and researchers will be geared towards the implementation or further development of AI based condition monitoring systems within the offshore wind energy sector.

Building on the objectives, the following research questions guide this thesis:

1. How effectively can AI algorithms detect early warning signs of offshore turbine component failures compared to traditional condition monitoring methods?
2. What data preprocessing techniques (feature engineering, outlier filtering, etc.) lead to achieving optimal performance of predictive maintenance models on offshore wind farms?
3. Can the proposed predictive maintenance framework decrease downtime and O&M cost, and increase reliability and availability of offshore wind assets?

From a technical scope, the study concentrates on the development and validity of the AI techniques for the fault detection and predictive maintenance. Specifically, it focuses on critical wind turbine components, namely gearboxes and generators, which are known to be especially prone to failures in the offshore situation. The research methods include a detailed analysis of SCADA data, as well as sensor streams that contain measurements of vibration, temperature, power output, and other sensors.

A key consideration in the study is data constraints as well. Often, the datasets collected offshore experience these impediments of incomplete labels and a lack of recorded failure events during training and validation. In addition, the disclosed or sharable limit of proprietary information may be subject to issues related to data ownership and confidentiality.

Another important thing to consider is Computational. The thesis shows that the current methodologies and tools make it feasible to have proof of concept models and offline analyses, but real time deployment of such AI models require such specialized hardware. Further complicating the issue of scaling the solutions is the fact that scaling solutions to militarily manage large fleets of wind turbine arrays involves matters beyond the scope of a single research effort.

1.5 Structure of the Thesis

To achieve clarity and logical flow, the thesis is organized into six main chapters:

Chapter 1: Introduction

This chapter provides the overall background and sets the stage for the research by presenting key elements such as the research motivation, problem statement, objectives, research questions, and the scope of the thesis.

Chapter 2: Literature Review

This chapter examines existing approaches related to condition monitoring and predictive maintenance, with a special focus on AI applications in offshore wind farms. It includes a detailed overview of offshore wind farm components, monitoring techniques (like vibration, temperature, SCADA data), and maintenance strategies (reactive, preventive, and predictive). Furthermore, the chapter identifies gaps in the current literature and frames how this thesis contributes to advancing the state of the art.

Chapter 3: Methodology

The methodology chapter details the overall research design and approach, including data collection methods, preprocessing techniques, and feature engineering. It explains the selection of specific AI algorithms (such as 1D-CNN, LightGBM, RNN(LSTM) and outlines the training, validation, and testing procedures. This chapter also describes the evaluation metrics used to measure the model's performance (e.g., accuracy, precision, recall, F1-score) as well as the tools and software involved in the study.

Chapter 4: Data Analysis and Model Development

In this chapter, the structure and statistics of the collected datasets are thoroughly analyzed and visualized. It explains the process of model development, including the steps taken for hyperparameter tuning and model building. Detailed discussion is provided on component-level fault distributions across turbines (gearbox, generator, etc.), and the challenges encountered during the model development process such as data imbalance and noise and how these issues were addressed.

Chapter 5: Results and Discussion

This chapter presents the empirical findings obtained from the developed models. It includes a detailed evaluation of model performance using both quantitative (confusion matrix, feature importance analysis) and qualitative methods. Comparisons are made with baseline and existing methods from literature. Insights related to offshore maintenance strategies are discussed, as well as the limitations of the current study, particularly concerning data quality and the generalizability of the model.

Chapter 6: Conclusion and Future Work

The last chapter provides a full summary of the research findings and reflects on the overall contributions to the field of AI-based condition monitoring and predictive maintenance for offshore wind farms. It includes practical recommendations for improving offshore wind maintenance strategies based on the study's outcomes. Additionally, this chapter outlines future research directions such as potential model improvements and the expansion of the study to additional components or diverse types of renewable energy systems.

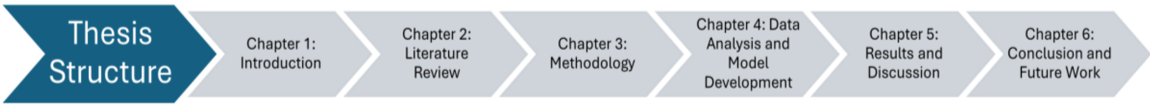

Figure 4. Thesis Structure

2 Literature Review

This chapter presents a comprehensive review of offshore wind farm technology, condition monitoring methodologies, predictive maintenance models, and the role of artificial intelligence in enhancing asset management. First, we look at the basic description and operational issues in setting up offshore wind installations, then discuss the theory and application of condition monitoring and finish with a discussion of different maintenance strategies from reactive to predictive along with the necessity of accessing real-live data through SCADA systems. After this, we investigate AI based approaches to fault detection and estimated remaining useful life and finally come to a synthesis of literature gaps that informs the next research.

2.1 Offshore Wind Farms Overview

Having reached fast becoming one of the fastest growing sectors in the global renewables mix over the last decade, offshore wind energy has been undergoing record expansion. Driven by ambitious decarbonization targets and declining levelized costs, installed offshore capacity rose from just under 5 GW in 2010 to over 35 GW by 2023, with projections pointing toward 200 GW by 2030 under current policy trajectories. Governments and industry groups have backed large-scale auctions and leasing rounds particularly in Europe and East Asia catalyzing investment and accelerating deployment.

Compared to onshore installations, offshore sites benefit from higher mean wind speeds (often 20–40 % greater) and reduced surface roughness, which together yield capacity factors in the 40–50 % range versus 25–35 % on land . The open-sea setting also minimizes turbulence intensity, extending component lifetimes by reducing stochastic loading cycles on blades and drive-train elements. These aerodynamic advantages, coupled with vast contiguous areas for multi-GW clusters, make offshore wind particularly attractive for large-scale power generation (Weiss et al., 2018).

At the heart of every offshore farm lies a complex assembly of turbines, foundations, inter-turbine cables, and one or more substations. Modern turbines exceed 10 MW per unit, featuring three composite blades on a horizontal-axis rotor, a gearbox (or direct-drive generator), and a tall lattice or tubular steel tower. Foundations are chosen by water depth: monopiles up to ~30 m, jacket structures from 30–60 m, and floating spar, semi-submersible, or tensioned-leg platforms in deeper waters. Dynamic export cables often XLPE-insulated, link turbines to offshore substations, where step-up transformers boost voltage (up to 220 kV AC or ±320 kV DC) for cost-effective transmission to shore .

Optimizing turbine layout is critical to minimizing wake interference and maximizing yield. Computational fluid dynamics (CFD) and wake-loss modeling inform array configurations staggered rows, varying inter-row spacings, and optimized yaw control to mitigate power deficits downwind of each rotor. Power-electronics installations, such as STATCOMs and VSC converters, further ensure grid compliance by regulating reactive power, voltage stability, and fault-ride-through capabilities under variable offshore conditions.

While the core components define energy capture and transmission, a successful offshore project also turning point on robust marine civil infrastructure, port coordination, and vessel scheduling. Site selection factors in seabed geotechnics, navigational constraints, and proximity to onshore grid nodes. These elements together form an integrated system for maximum energy yield and for resilience against harsh environments and reliable power delivery to coastal power centers.

2.1.1 Components and Layout

The components of an offshore wind turbine include the rotor assembly, nacelle, tower, foundation, and electrical infrastructure. The captured aerodynamic energy is used to rotate mechanically by means of the rotor, consisting of two or three composite blades mounted to a central hub. Within the nacelle, the gearbox steps up rotor speed to the level required by the generator, which produces electrical power. The tower, typically a cylindrical steel structure, elevates the rotor to heights where wind speeds are optimal. Foundations vary by water depth: monopiles are common in shallow waters, jacket structures in intermediate depths, and floating platforms for deepwater sites.

Electrical infrastructure begins with dynamic or export cables linking turbines to an offshore substation (Gulski et al., 2021). Here, step-up transformers increase voltage for efficient long-distance transmission to onshore grids. The spatial arrangement of turbines follows optimized patterns to minimize wake interference, often using computational fluid dynamics during planning. In addition, power electronics such as converters and reactive power compensators are installed to regulate voltage, frequency, and grid stability.

2.1.2 Challenges in Offshore Environments

Operating offshore introduces unique mechanical, environmental, and logistical challenges. Saltwater corrosion attacks metallic surfaces, necessitating robust corrosion protection systems and regular inspections (Odeyemi & Alaba, 2025) . Marine growth on submerged components alters hydrodynamic loading and may accelerate fatigue damage. Constant wave and tidal action impose cyclic loads on towers and foundations, increasing the risk of crack initiation. Harsh weather conditions storms, high waves, and icing limit access windows for maintenance vessels and personnel, often forcing delays or riskier interventions.

Additionally, data collection itself can be compromised. Subsea cable faults, electromagnetic interference, and bandwidth constraints may lead to gaps in SCADA datasets (Pandit & Wang, 2024). Environmental regulations such as noise restrictions during pile driving and protected marine life considerations add further constraints to installation and repair schedules. Together, these factors expand both capital and operating expenditures, reinforcing the need for proactive, data-driven maintenance frameworks.

Nordic Weather Conditions and Icing Issues. In northern offshore sites, extreme cold and persistent icing present additional challenges (Haeberli & Whiteman, 2021) that can severely impair turbine performance and condition monitoring accuracy. Sea spray and freezing temperatures lead to sea icing and blade icing, which not only reduce aerodynamic efficiency and power output (Larsén et al., 2024) but also obstruct sensor ports and ice detectors, resulting in data dropouts or misleading measurements (Dhar, 2021). Ice accretion on blades and anemometers distorts vibration and wind-speed signals undermining the reliability of fault-detection algorithms and increasing false-alarm rates (Zuo et al., 2024). Furthermore, hazardous deck and platform surfaces, coupled with limited daylight and fierce winds, significantly restrict safe access for inspection and repair crews, elevating the risk of accidents and necessitating specialized cold-weather gear and vessels (Eriksson, 2022). These factors demand tailored de-icing strategies, ice-tolerant sensor designs, and careful planning of maintenance windows to sustain high reliability in Nordic offshore farms.

2.2 Condition Monitoring in Wind Turbines

Condition monitoring (CM) encompasses the continuous surveillance of operational parameters to infer equipment health and detect emerging faults (Martin, 1994). In offshore contexts, where unplanned downtime can cost hundreds of thousands per day, CM is not merely beneficial, it is indispensable.

Condition monitoring systems integrate sensors with analytical software. Sensors measure vibration signatures, temperatures, acoustic emissions, torque, and electrical waveforms (Hassan et al., 2024). These measurements are often transmitted at high sampling rates to onshore control centers via SCADA networks. Sophisticated signal-processing algorithms extract features such as frequency peaks, trend slopes, and statistical descriptors, which are compared to baseline signatures representing normal operation.

Early warning of component degradation allows maintenance teams to plan interventions in advance, reducing emergency repairs and leveraging favorable weather windows (Adepoju et al., 2022). For instance, an emerging bearing fault may manifest as slight increases in specific vibration frequencies before audible noise or catastrophic failure occurs. By detecting such anomalies at low severity, CM enables service actions that are faster, safer, and more cost-effective.

2.2.1 Common Faults and Failures

Offshore turbines present fault modes across multiple components. Blade damage such as leading-edge erosion or delamination reduces aerodynamic efficiency and produces imbalance (Katsaprakakis et al., 2021). Gearboxes experience lubrication loss, gear pitting, and bearing brinelling, which escalate vibration levels and heat generation. Electrical faults manifest as partial discharges in transformers, insulation breakdown in stator windings, and power electronics malfunctions. Structural elements towers and foundations encounter corrosion pits and fatigue cracks at weld seams. **Figure 5** highlights the annual failure rates and average downtime per failure for major offshore turbine components, illustrating how these fault modes across subsystems, such as gearboxes and electrical components, drive significant impacts on availability and operational efficiency.

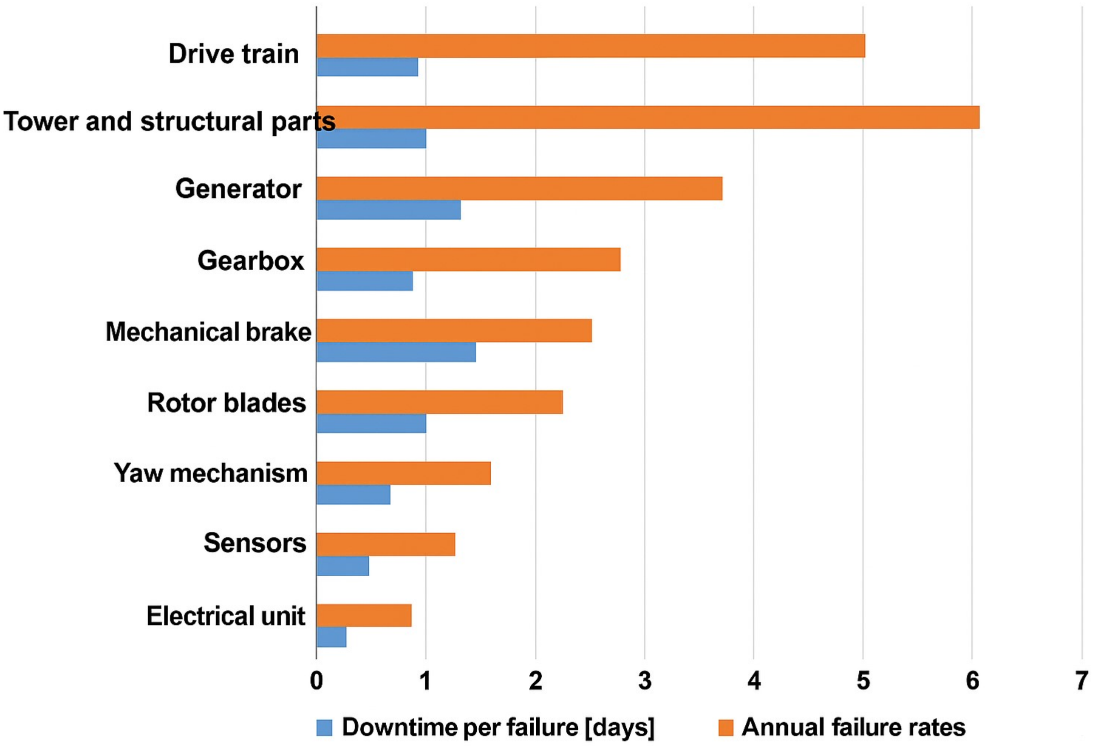

Figure 5. Illustrates the annual failure rates and average downtime per failure for major offshore turbine components (Ciuriuc et al., 2022).

Early detection of these failure modes relies on correlating sensor signals with known fault signatures. **Table 1** summarizes typical fault modes alongside common detection signals.

Table 1. Offshore Turbine Fault Modes and Detection Signals

Component	**Fault Mode**	**Detection Signal**
Blades	Leading-edge erosion	Increased vibration, acoustic spikes
Gearbox	Bearing pitting	Frequency spectral peaks, temperature
Generator	Winding insulation defects	Current harmonics, partial discharge

Bearings	Lubrication starvation	High-frequency vibration, heat
Tower/Foundation	Corrosion, fatigue cracks	Ultrasonic echoes, visual inspection

2.2.2 Monitoring Techniques

In offshore wind farms, a suite of complementary sensing modalities is deployed to capture the earliest signs of component degradation and contextual performance shifts. Vibration analysis remains the most sensitive mechanical indicator tri-axial accelerometers mounted on the gearbox housing, generator frame, and blade roots continuously record oscillatory signatures. By transforming these time-series signals into the frequency domain using Fourier or wavelet analyses operators can detect characteristic peaks associated with gear mesh frequencies or bearing defect frequencies long before damage becomes audible or visible (Gubran, 2015). Temperature monitoring provides a powerful thermal counterpart, with thermocouples or infrared probes tracking bearing, gearbox, and power-electronics temperatures; deviations from normal thermal profiles often presage lubrication loss or electrical overload. Electrical-signal analysis augments these data by inspecting generator current and voltage waveforms for harmonics, unbalance, and partial-discharge patterns; increases in harmonic content or intermittent voltage spikes frequently signal insulation breakdown in stator windings or transformers. Finally, SCADA-level analytics draw on slower rate streams of rotor speed, power output, pitch and yaw angles, and ambient conditions such as wind speed, wave height, and temperature to identify subtle drifts in performance curves or abnormal dynamic responses. They comprise a layering surveillance framework of these four (vibration, temperature, electrical, and SCADA data) monitoring techniques and high-resolution fault detection coupled with fleet wide performance benchmarking.

2.3 Predictive Maintenance Strategies

Predictive maintenance now is a change in thinking away from traditional maintenance approaches where the equipment health is continuously monitored but also predicts future degradation and estimated remaining useful life (RUL) for components. PdM differs from both fixed interval servicing and reactive repairs with their risk of unexpected failures or waste downtime, respectively, by optimizing intervention timing based on real time data streams and predictive models (Lee et al., 2017).

Massive quantities of sensor and SCADA data, for instance generated from high frequency vibration signatures and temperature readings, electrical current waveforms, and environmental measurements, are ingested and transformed into actionable insights in offshore wind applications. Both signal processing techniques (e.g., spectral analysis, time domain feature extraction), and statistical and machine learning methods that `model degradation trends to predict when a component will cross a critical threshold' identify the early indicators of wear or malfunction. Putting continuous health assessment with forward looking prognostics together gives PdM capability of planning maintenance within a favorable weather window, reducing unplanned stoppages and extending expensive assets service life. In the continuum from reactive maintenance to proactive, proactive to predictive, the final stages shown in **Figure 6** are described to show how each of these continue to build on the previous stage to enable ever more data driven decision making.

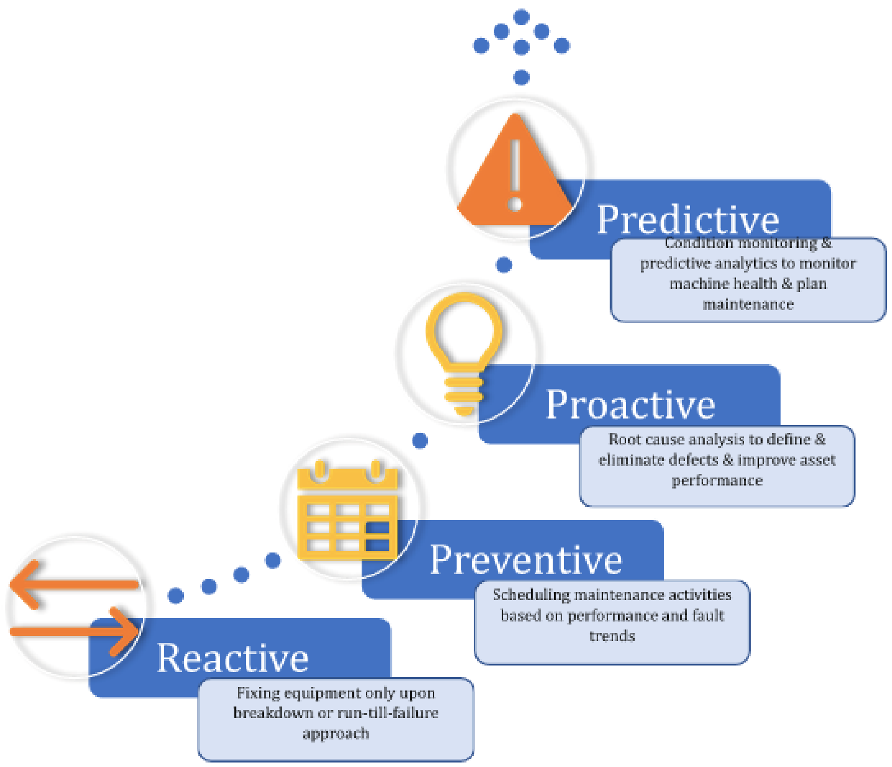

Figure 6. Predictive maintenance strategies, Source: [Predictive maintenance strategies]

2.3.1 Reactive vs. Preventive vs. Predictive Maintenance

Maintenance strategies occur on a continuum from reactive to preventive with different tradeoffs in cost, risk, and availability of the asset found in the final position (Narayan, 2004). Repairing equipment only after failure, reactive maintenance requires minimal planning overhead but increases the downtime costs, safety threats from emergency mobilizations and collateral damage to adjacent turbine components. This leads to over service and idle time of the vessel when continuing to service the vessel in healthy condition, thus preventive maintenance is an improvement over this, but scheduling service is based on calendar or operating hour intervals, which reduces catastrophic failures. At the peak of this spectrum is predictive maintenance that uses condition-based indicators and RUL and Forecasts to deploy the intervention at optimal time. The advantage of this approach is not only that it curbs the redundancy of routine servicing but that it brings forward maintenance activities with what the equipment's healthy status demands and circumvents the unpredictability of breakdowns (Yazdi, 2024).

Table 2. Comparison of Maintenance Strategies

Strategy	Intervention Criterion	Advantages	Limitations
Reactive	Post-failure	Low planning overhead	High downtime, safety risks
Preventive	Calendar/usage interval	Predictable scheduling	Potential over-maintenance
Predictive	Condition/RUL prediction	Optimized service timing, cost-effective	Requires robust data & models

2.3.2 Importance of Predictive Maintenance in Offshore Wind

It is also riskier for offshore maintenance planning. Jack up vessels and cranes must constantly move as crews can cost hundreds of thousands of dollars per day and weather constraints often squish service window durations to short time frames. Predictive maintenance directly addresses these challenges by providing actionable forecasts that enable planners to synchronize work orders with meteorological forecasts and vessel availability. Early detection of gearbox wear or blade root fatigue allows operators to bundle tasks across multiple turbines optimizing vessel routes and reducing transit time while avoiding unplanned shutdowns that erode annual energy production and revenue. Over the life of a farm, PdM-driven reliability improvements can yield multi-million dollar savings, reinforcing its strategic value in offshore wind asset management.(Turnbull & Carroll, 2021).

By coupling AI-driven analytics with high-frequency SCADA feeds, offshore operators can drastically cut back on hazardous deck visits. Anomaly-detection models pre-screen turbines continuously, flagging only a small subset for physical inspection studies show this can reduce crew deployments by up to 60 % compared to time-based inspection schedules (Seddini & Triqui-Sari, 2024). For flagged units, drones or robotic crawlers equipped with cameras and ultrasonic probes perform targeted inspections, further minimizing personnel exposure to severe weather and limited-access platforms (Mouschoutzi & Ponis, 2022). Real-time AI alerts ensure that human interventions occur only within approved safety and weather windows, streamlining logistics, lowering accident risk, and optimizing the use of specialist vessels and crews (Durlik et al., 2024).

2.3.3 Role of SCADA Systems in Predictive Maintenance

SCADA systems serve as the backbone of any offshore predictive maintenance model, aggregating time-stamped operational and environmental data from across the turbine fleet. Traditionally, SCADA platforms have been essential for real-time monitoring and control, but modern systems go a step further by archiving high-resolution logs that serve as the training foundation for prognostic models. As shown in **Figure 7**, real-time sensor data from the turbine gearbox is streamed to an onsite server, stored in a database, analyzed by a fault-diagnosis model, and used to generate maintenance plans while also allocating resources for remote operators. To ensure accurate analysis, it is crucial to maintain data integrity by applying outlier filtering, interpolating missing values, and aligning disparate sampling rates, as noisy or incomplete data can mislead anomaly detection systems and compromise remaining useful life (RUL) estimates (Udo & Muhammad, 2021)Furthermore, advanced SCADA integrations allow for bidirectional workflows: alerts generated by machine learning models can be fed back into the control interface, triggering automatic derating, shifting turbines to safe modes, or sending maintenance notifications. This close integration between data insights and operational responses enhances overall system reliability.

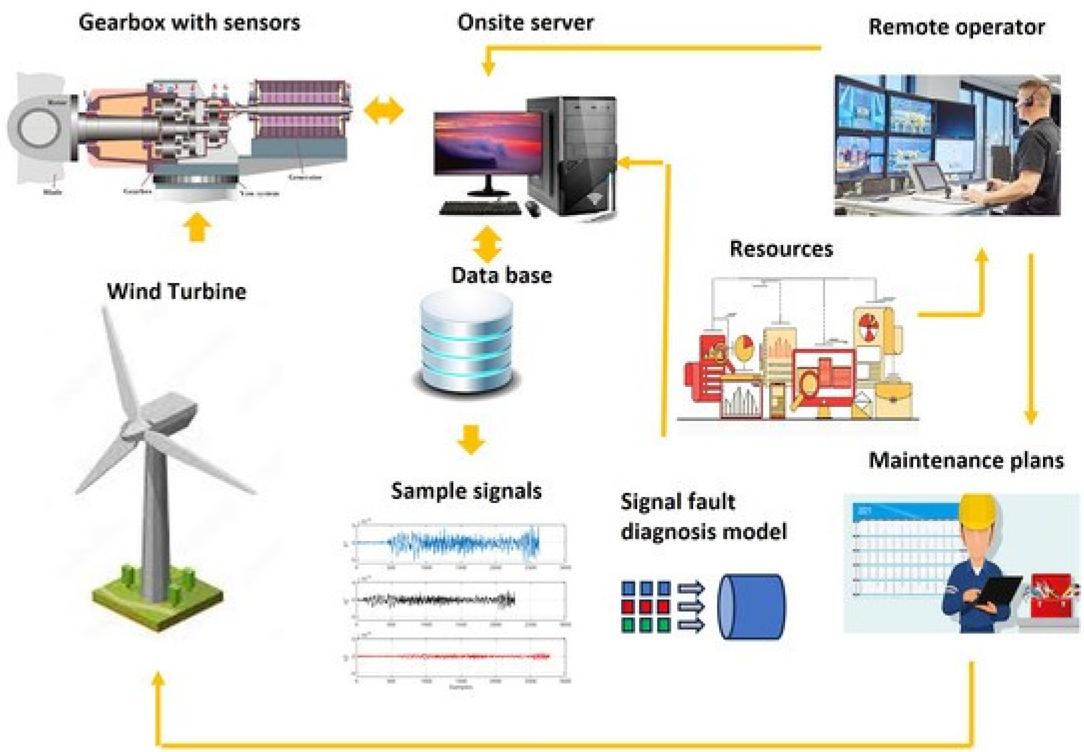

Figure 7. An IoT-Enabled Predictive Maintenance Workflow for Wind Turbine Gearboxes (Santiago et al, 2024)

Several studies demonstrate that leveraging high-resolution SCADA streams for real-time analytics materially improves fault detection and maintenance planning(Udo & Muhammad, 2021) showed that a data-driven SCADA framework raised gearbox-fault detection recall by over 25 % compared to threshold alarms alone. (Moleda et al., 2020)applied SCADA-based anomaly detection to boiler feed pumps, cutting unplanned downtime by 30 % through early warning alerts. (Pandit & Wang, 2024) reviewed multiple offshore deployments, finding SCADA-driven models reduced false-alarm rates by up to 40 % when combined with vibration and temperature analytics. In field trials, windows (Adepoju et al., 2022)integrated SCADA feeds with machine-learning classifiers, achieving 90 % precision in blade-pitch anomaly detection and enabling targeted inspections. Finally, (Seddini & Triqui-Sari, 2024) reported live SCADA-AI pilots at three

North Sea farms that saved 15 % off annual O&M costs by minimizing unnecessary manual inspections.

2.3.4 Benefits of Predictive Maintenance

By harmonizing maintenance actions with actual equipment condition, PdM delivers significant operational and economic benefits. Quantitative studies in offshore contexts report operations and maintenance (O&M) cost reductions of 20–40 % when shifting from preventive to predictive regimes, driven by fewer emergency mobilizations, optimized spare-parts inventory levels, and longer component lifetimes (Bousdekis, Apostolou, et al., 2019). In addition to direct cost savings, predictive strategies enhance asset availability and grid reliability by reducing downtime and smoothing production profiles. The structured data workflows and analytics frameworks developed for PdM also foster continuous improvement: as more data accumulates, models become more accurate, and best practices propagate across fleets and projects, creating a virtuous cycle of performance enhancement.

2.4 Role of AI in Predictive Maintenance

Artificial intelligence has transformed predictive maintenance by enabling models to learn complex, non-linear interactions among multiple sensor streams and SCADA measurements in offshore wind environments (Chatterjee & Dethlefs, 2021). Rather than relying solely on manually engineered features or univariate thresholds, AI approaches such as ensemble machine learning, unsupervised anomaly detection, and deep-learning architectures can ingest high-dimensional inputs, fuse heterogeneous modalities, and adapt to evolving operational conditions. These methods have demonstrated superior sensitivity in early fault detection, more accurate remaining useful life estimation, and reduced false-alarm rates, thereby improving maintenance scheduling, minimizing unplanned downtime, and extending component lifetimes in harsh marine settings.

2.4.1 Machine Learning and Deep Learning Approaches

Supervised learning techniques have become a cornerstone of modern predictive maintenance by mapping diagnostic features extracted from sensors and SCADA to labeled failure events (Omol et al., 2024). Traditional algorithms such as random forests and gradient boosting effectively manage tabular features like vibration spectral amplitudes, temperature deviations, and power-curve residuals, providing robust classification or regression outputs with interpretable feature importance. In contrast, deep learning architecture, as shown in **Figure 8**, learns representations and predictions end to end, by passing on the need for hand-crafted features. Convolutional neural networks (CNNs) capture local temporal or spectral patterns in vibration and electrical signals, while long short-term memory (LSTM) and gated recurrent neural networks (RNNs) model long-range dependencies reflecting degradation trajectories. These deep learning models can be trained to predict both fault occurrence and remaining useful life, often outperforming conventional statistical methods in heterogeneous offshore datasets by learning hierarchical representations that generalize across turbine designs and environmental conditions.

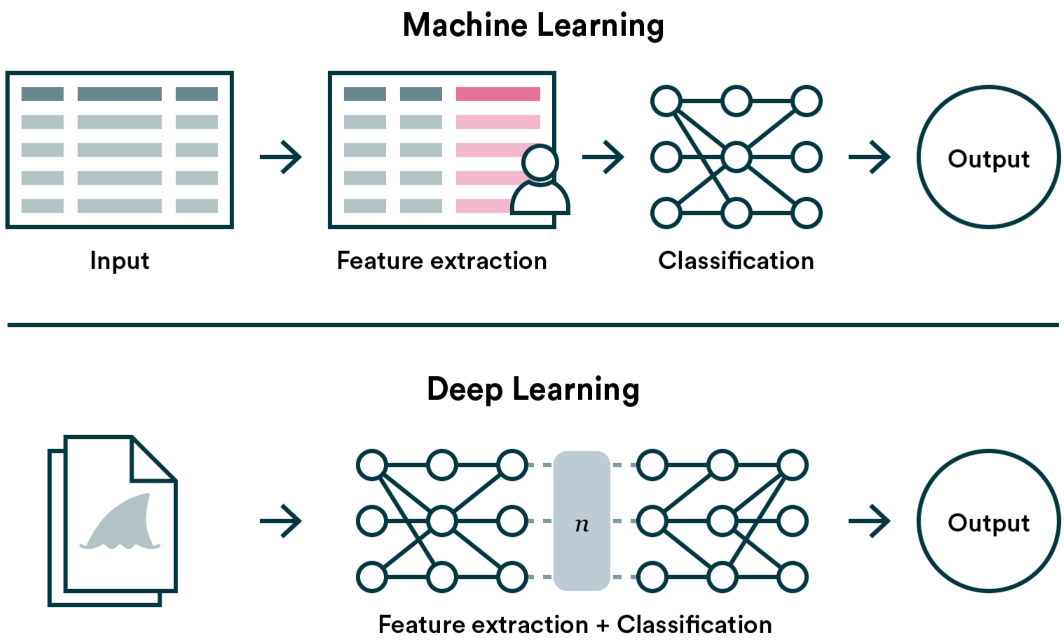

Figure 8. Machine Learning vs Deep Learning, Source : [ML vs DL (valohai)]

2.4.2 Multi-Component Monitoring in Offshore Turbines

Recent predictive-maintenance platforms leverage artificial intelligence to fuse disparate component-level data into holistic health indices and risk maps across entire wind farms (Chatterjee & Dethlefs, 2021). **Figure 9** presents a high-level predictive maintenance architecture, from on-asset sensing and SCADA ingestion through centralized analytics, to dashboards and real-time alerting. For blades, convolutional neural networks trained on drone imagery identify surface erosion and cracks, while LSTM models ingest blade-root vibration to track oscillatory anomalies (Memari et al., 2024); merging these visual and vibrational embeddings yields more precise damage localization and severity scoring than either modality alone. In the drivetrain, autoencoder based anomaly detectors process raw acceleration spectra from multiple bearing locations, and when combined with temperature trends, gradient-boosted regression trees forecast remaining useful life with quantified uncertainty enabling spare-parts organization and vessel scheduling weeks in advance. Generator and power electronics health is assessed through deep belief networks that learn current-signature features and SCADA derived metrics to distinguish mechanical from electrical faults. Structural integrity and foundation fatigue are monitored via ultrasonic guided wave sensors and strain gauges, with graph neural-network models synthesizing these measurements alongside tower accelerations recorded during extreme weather to project crack propagation and corrosion progression. By orchestrating these component-level AI models within a unified framework, operators gain real-time, fleet-wide risk assessments and can dynamically prioritize maintenance dispatching service vessels only when combined risk thresholds are exceeded, thereby maximizing uptime and minimizing cost and safety exposure (Durlik et al., 2024).

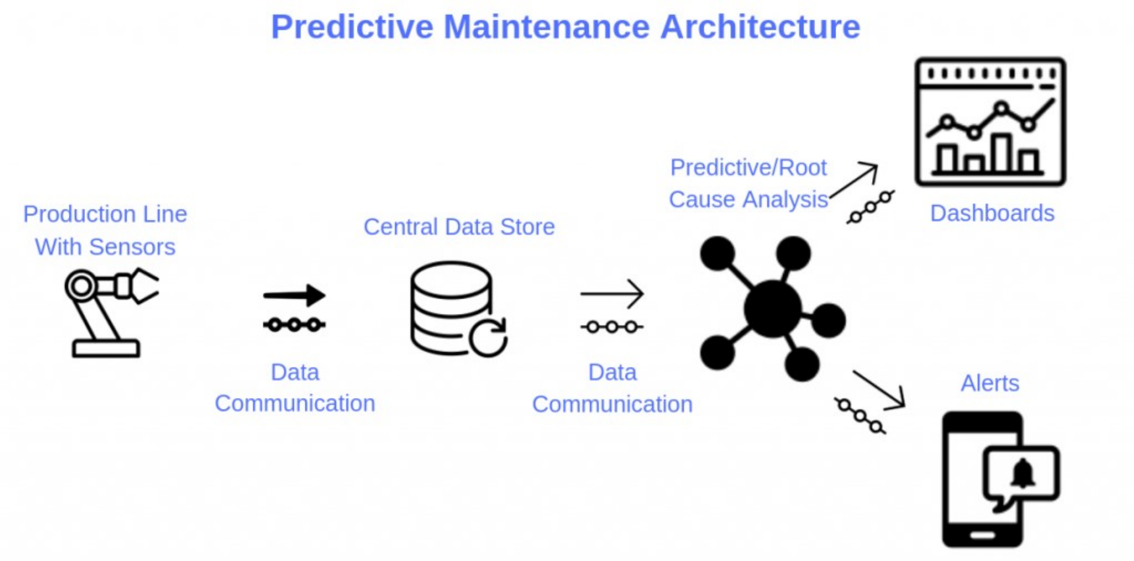

Figure 9. Predictive Maintenance Architecture, Source : Predictive Maintenance Architecture

2.5 Summary of Gaps in Literature

Despite the rapid circulation of AI-enabled predictive maintenance (PdM) in offshore wind, the underlying research landscape still exhibits several persistent gaps. First is the shortage of high-quality, publicly available failure datasets. Most studies rely on proprietary SCADA archives or limited laboratory test rigs, making it difficult to benchmark new algorithms or reproduce published results. Moreover, the available datasets are often highly heterogeneously varying in sensor types, sampling rates, labeling conventions, and fault taxonomies which frustrate efforts to develop universally applicable models.

A second challenge lies in the “black-box” nature of many advanced machine-learning and deep-learning approaches. While convolutional and recurrent neural networks can achieve high accuracy in fault detection and RUL estimation, their lack of transparency undermines operator trust and slows regulatory approval in safety critical offshore environments. Explainable AI techniques such as attention based saliency maps, model distillation, or surrogate rule extraction remain under explored in this domain, and there is a critical need for standardized interpretability benchmarks that correlate model explanations with real maintenance outcomes (Mersha et al., 2024).

Real time processing and scalability represent a third barrier. Offshore farms generate terabytes of high frequency vibration, temperature, acoustic, and electrical data every day. Streaming this volume to centralized cloud servers can strain bandwidth and incur latency that precludes timely decision-making. At the same time, on site edge computing hardware must balance computational power against power consumption, cost, and ruggedness. Few studies address the co-design of algorithms and hardware platforms that can deliver sub-second inference across hundreds of turbines under realistic communication constraints (Rajendran et al., 2023).

Fourth, the issue of generalizability across farm sites and turbine designs remains unresolved. Most PdM models are trained and validated on a single farm or a narrow class of turbines, leading to significant performance degradation when applied elsewhere (De Nolasco Santos, 2023). Transfer learning, domain adaptation, and federated learning paradigms offer promising avenues to overcome site and design specific biases, yet systematic research on how to best adapt models with minimal labeled data is still in its infancy.

Finally, integration of AI systems into existing asset management workflows continues to lag. Beyond the core algorithms, operators require seamless interfacing with digital twins, maintenance planning software, and safety management platforms. There is also limited work on quantifying the economic trade-offs balancing vessel mobilization costs, spare-parts inventory, and production losses when adopting advanced PdM (Mouschoutzi & Ponis, 2022). Human machine interaction studies, standardized user interfaces, and end to end cost benefit frameworks are all needed to bridge the gap between algorithmic promise and industrial practice.

3 Methodology

In this chapter, the methodology employed to develop and validate AI-based condition-monitoring and predictive maintenance solutions for offshore wind turbines is presented. The narrative begins by situating each turbine's SCADA and sensor logs as individual cases within a broader multiple case framework, then moves into a description of how raw data are ingested, labeled according to known anomaly events, and explored to reveal missingness, outlier patterns, and preliminary fault signatures. From there, the discussion turns to the transformation of raw sensor channels grouping them into logical categories, renaming for clarity, and applying dimensionality reduction techniques to distill the most informative components. The methodology then explains how continuous time-series streams are segmented into fixed-length windows for sequence models, while aggregated snapshots of feature vectors serve as inputs for tree-based learners. The chapter concludes by detailing the protocols for training recurrent and convolutional neural networks alongside gradient boosted decision-tree models, including strategies for hyperparameter tuning, early stopping, and class-imbalance mitigation, and by defining a suite of evaluation metrics accuracy, precision, recall, F1-score that provide a consistent basis for comparing predictive performance under realistic operational constraints.

3.1 Research Approach and Design

This study follows a quantitative, data-driven pattern built around a modular processing pipeline that ensures both reproducibility and flexibility. Each turbine dataset is treated as a standalone case, allowing models to be validated across multiple fault modes and environmental conditions. Data ingestion aligns time-stamped operational streams with anomaly labels, after which exploratory analysis characterizes the completeness and statistical properties of the sensor feeds. Preprocessing integrates these disparate channels into coherent categories before applying incremental principal component analysis to reduce dimensionality while preserving the most salient variance. Two complementary dataset constructions then take place: one that captures temporal dependencies via fixed-length sequences for LSTM and convolutional networks, and another that aggregates features into snapshot vectors for gradient-boosting models. Throughout, stratified sampling and controlled noise injection address the rarity of fault events, enhancing model robustness (Frank, 1994) . Model development proceeds with learning rate scheduling and early stopping criteria to guard against overfitting, and performance is assessed using a comprehensive set of metrics that together illuminate the trade-offs between sequential and non-sequential approaches. This design strikes a deliberate balance between computational efficiency and empirical rigor, laying a transparent groundwork for robust condition monitoring in offshore wind applications (Koltsidopoulos Papatzimos, 2020).

3.2 Data Collection

The primary dataset for this study is the “CARE to Compare” wind turbine anomaly detection collection, originally assembled by Güuck, Roelofs, and Faulstich (Gück et al., 2024) . It comprises 89 years of real-world SCADA recordings drawn from thirty-six turbines spanning three distinct farms two offshore installations in Germany and one onshore site in Portugal. In total, ninety-five individual time-series files are provided: forty-four segments encompass labeled anomaly windows that culminated in component faults, and fifty-one represent sustained periods of normal operation. Each record features ten-minute resolution measurements, capturing everything from rotor speed, power output, and pitch angles to vibration, temperature, and hydraulic pressures. The number of monitored variables varies by source: Farm A supplies eighty-six channels, Farm B 257, and the most instrumented Farm C has an impressive 957 distinct features.

To facilitate seamless ingestion and analysis, the original semicolon delimiters were converted to commas, and each file was accompanied by two metadata tables: one detailing the exact start and end timestamps of documented fault events, and another enumerating every sensor's description and units. Most turbines contribute one year of uninterrupted training data followed by 4–98 days of prediction data intended to evaluate anomaly alerts ahead of actual failures. Importantly, the dataset carries a CC BY-SA 4.0 license, permitting adaptation so long as derivative works are shared alike. This rich, openly licensed resource cited as Güuck, Roelofs, & Faulstich (2024) provides one of the most comprehensive public testbeds for developing and benchmarking early warning and predictive maintenance algorithms in the wind energy domain.

3.2.1 Wind Farm Dataset Description

The "CARE to Compare" collection comprises three distinct wind farm datasets, each reflecting unique operational and instrumentation characteristics. Wind Farm A, an onshore installation in Portugal, encompasses five turbines and furnishes 86 SCADA channels recorded at ten-minute intervals. Its data spans one year of nominal operation, followed by up to 98 days of held-out prediction data, and includes twenty-two labeled fault windows drawn from the farm's fault logbook.

Wind Farm B and Wind Farm C originate from offshore facilities in Germany, where environmental stresses and maintenance challenges differ markedly from onshore settings. Wind Farm B covers fifteen turbines and provides 257 measurement channels adding variables such as blade pitch hydraulics and nacelle orientation to the standard power and wind-speed signals. Its one-year training period is paired with 4–60 days of anomaly testing data, across which twelve distinct labeled failure events occur.

Wind Farm C, the most instrumented of the three, consists of sixteen turbines and delivers an extensive set of 957 features. In addition to electrical and mechanical sensors, it incorporates high resolution vibration, oil quality, and humidity monitors that capture the harsh marine environment's effects. For each turbine, one calendar year of baseline data precedes a 10–80-day prediction window, during which ten significant anomaly intervals are labeled.

Across all farms, each ten-minute record is accompanied by a status type label that distinguishes normal operation from the four severity levels used to flag pre-fault conditions. A separate feature-description table provides human readable names, units, and brief explanations for every channel, while an event-info table precisely timestamps the start and end of each documented anomaly. This multi farm, multi turbine design ensures that models trained on one dataset encounter a wide variety of fault modes, seasonal patterns, and instrumentation configurations before being evaluated on unseen anomaly windows.

3.2.2 Data Preprocessing and Cleaning

In preparing the raw SCADA streams for analysis, we first converted every timestamp originally logged in mixed local formats into a unified UTC DateTimeIndex, then sorted each turbine's data chronologically to ensure strict temporal order. Any rows duplicated exactly (e.g., from retransmissions) were removed, and the resulting ten-minute sampling grid was examined for missing intervals. Gaps shorter than six consecutive readings (i.e., under an hour) were forward-filled and then backward-filled to preserve trend continuity without introducing artificial ramps longer gaps were left as NaN to be managed during interpolation.

Because our working environment-imposed RAM constraints, we did not process all available farm files at once. Instead, we selected three anomaly labeled files and three normal operation files to constitute our working dataset. This subset provided a representative mix of fault and baseline behavior while keeping memory usage within practical limits.

Next, we isolated the numeric sensor measurements from ancillary columns (turbine ID, status codes, event flags), casting all channels to 32-bit floats to reduce memory overhead. We specifically selected our working subset from Wind Farm C, the most instrumented offshore installation in the study, to ensure rich, high-resolution data coverage. We computed per channel missingness and dropped any sensor with over 40 % data loss during the training year, deeming it too unreliable for fault modeling. Residual NaN values in the remaining channels were imputed by linear interpolation over a thirty-minute window three samples before and after each gap so that brief wireless dropouts would not distort downstream statistics.

To remove spurious artifacts without suppressing genuine transients, a two-stage outlier treatment was applied. First, a rolling z-score (window = 12 readings) was calculated for each series and values beyond $\pm 3\ \sigma$ were clipped to the corresponding boundary. Second, for channels exhibiting heavy tailed distributions, we computed the median absolute deviation (MAD) over the same window and replaced any points exceeding five MAD with the local median. This preserved fault-related spikes such as sudden vibration increases while eliminating improbable sensor glitches.

With gaps and outliers addressed, sensors were grouped into domain categories (gearbox, generator, power, wind, hydraulic, temperature, and other) based on their description metadata and each column was renamed to a human-readable format (e.g., gen_curr_avg → generator_current_average). We then fit a RobustScaler on the training subset subtracting each channel’s median and scaling by its interquartile range to mitigate remaining skew, followed by a StandardScaler to center features at zero mean and unit variance in preparation for PCA. Finally, the cleaned data were split into the designated training year and prediction window, and status codes were mapped to binary labels (normal vs. pre-fault), yielding six fully processed turbine files ready for feature extraction and model training.

3.3 Feature Engineering

To transform the raw, high-dimensional SCADA streams into model ready inputs, we first imposed a structured schema on each sensor channel. Drawing on the human readable descriptions and units supplied alongside the data, every measurement was assigned to one of seven domains gearbox, generator, power, wind, hydraulic, temperature, or other and its original column name (e.g. tmp_gearbox_bearing_avg) was replaced with a clear, category prefixed label (gearbox_bearing_average_temperature). This systematic renaming not only improves interpretability for downstream analysis but also makes it trivial to select or exclude entire sensor domains by name.

With channels grouped and renamed, we next distilled their raw statistics into a compact feature set. Each ten-minute interval already carried summary measures mean, maximum, minimum, and standard deviation for every sensor; we retained all four to capture both central tendency and variability. After outlier clipping and gap imputation, these per-interval statistics formed our initial feature matrix, with each row representing one timestamped snapshot across the turbine's hundreds of channels.

Given the sheer volume of correlated signals up to 957 features for the most instrumented farm we then applied an incremental principal component analysis (PCA) to compress the data into its most informative subspace. By fitting PCA in batches, we accommodated multi gigabyte CSVs without exceeding memory limits. We selected the first one hundred components, which collectively explained more than 95 percent of the variance in the training year and used these transformed values as the core inputs for both our tree-based and sequence-modeling pipelines. Alongside each component, we retained the original category mappings so that principal components could be traced back to the dominant sensor groups in interpretability analyses.

Finally, to capture temporal dynamics, we organized the PCA projections into fixed length windows of 20-time steps (spanning just over three hours). Each sequence thus becomes a three-dimensional tensor window length by component dimension ready for ingestion by LSTM and 1D-CNN architectures. At the same time, individual PCA vectors (one per ten-minute interval) serve as flat feature vectors for gradient-boosted tree models. This dual representation preserves both the temporal evolution of fault precursors, and the static snapshot information needed for comparative classifier baselines.

3.3.1 Feature Selection

Following PCA dimensionality reduction to one hundred components, we implemented a rigorous, two stage feature selection pipeline to isolate those axes of variance most strongly predictive of impending turbine faults. This procedure combined an initial statistical filter with an embedded importance ranking, yielding a final set of twenty non-redundant components for our gradient-boosted decision-tree learners (Alromema et al., 2023).

The first stage leveraged univariate correlation analysis to remove components that bore little linear relationship to the binary pre-fault label. On the training set, we computed the Pearson correlation coefficient between each principal component and the pre-fault indicator. **Figure 10** displays the explained variance ratios of the top ten components alongside the cumulative variance curve: the leading component alone accounts for approximately 54 % of total variance, and by the tenth component over 97 % is captured.

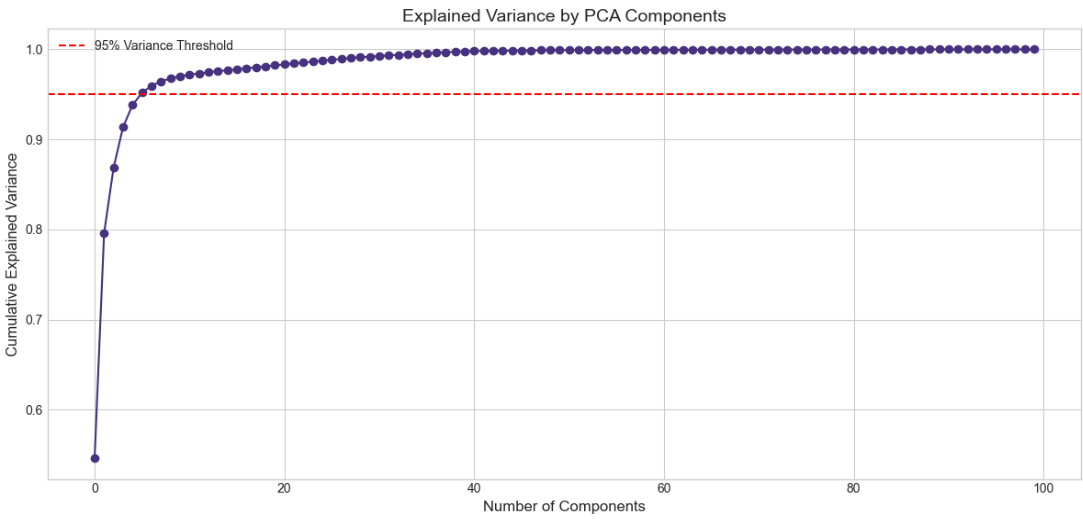

Figure 10. Explained Variance by PCA Components

To determine how many components are required to exceed our 95 % variance threshold, **Figure 11** presents the cumulative explained variance over all one hundred components. The curve crosses the 95 % mark at around the thirty-third component, after which it plateaus. Although we retained the full one hundred components for consistency across all farm datasets, this analysis confirmed that the bulk of informative variance resides in the first third of the component spectrum.

In the correlation filter, we removed all components which fell below 0.10, a threshold chosen by inspecting the distribution of correlation magnitudes and ensuring that at least half of the components (approximately 50) proceeded to the next stage. This step eliminated those axes of variance that exhibited negligible linear association with fault onset, thereby focusing our subsequent modeling on signals most likely to carry predictive information.

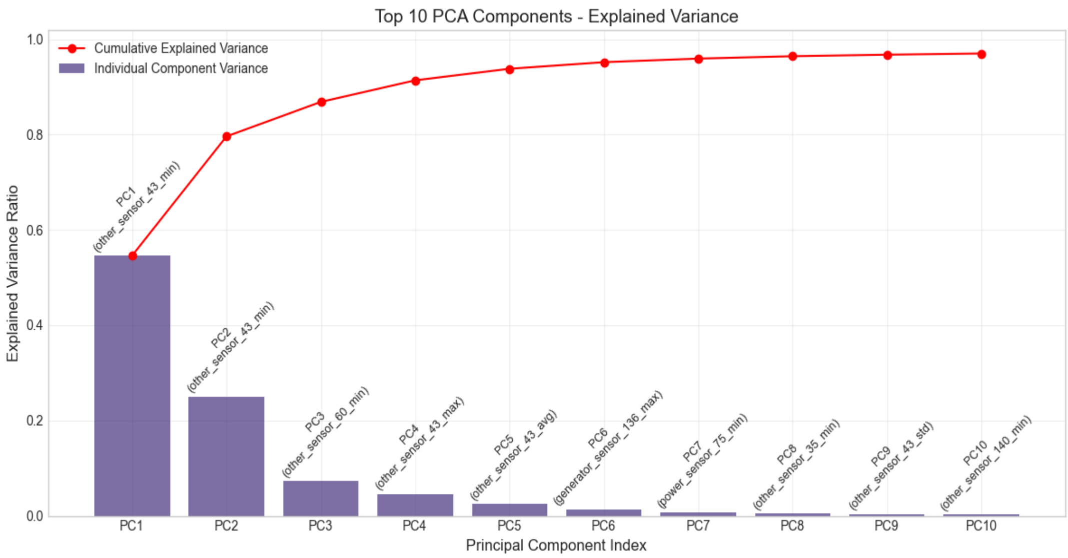

Figure 11. Top 10 PCA Components Explained Variance

The second, embedded stage refined the candidate set via a Random Forest classifier configured with five hundred trees, a maximum depth of ten, and sqrt feature sampling. Training this ensemble on the filtered components produced feature importance scores based on the mean decrease in Gini impurity. We then selected components in descending order of importance until their cumulative importance reached 95 % of the total. This procedure typically retained between 18 and 22 components, depending on the farm dataset. Finally, to guard against multicollinearity, we computed Pearson correlations among these top-ranked components and, for any pair exceeding 0.90 in absolute correlation, discarded the component with the lower importance score.

The resulting set of twenty or so principal components formed a compact, high-impact feature space for our LightGBM model, accelerating training and clarifying feature importance interpretations. For each selected component, we also examined its loading vector to trace back to dominant sensor domains such as gearbox vibration or hydraulic pressure thus providing post hoc explanatory power. Sequence based learners (LSTM and 1D-CNN) continued to ingest the entire 100-component representation, exploiting their capacity to learn complex temporal dependencies across the full variance structure. This two-stage pipeline strikes an effective balance between parsimony and predictive performance, ensuring that our classifiers remain both efficient and interpretable.

3.3.2 Managing Missing Data and Normalization

Prior to any feature extraction or model training, it was imperative to address the irregularities and scale differences inherent in the raw SCADA and sensor streams. Inconsistencies in timestamp formatting, intermittent telemetry dropouts, and divergent measurement units collectively threatened to introduce bias and degrade model performance (Saraiva, 2024). To mitigate these issues, we implemented a structured sequence of missing-data handling and normalization steps, as described below.

All timestamps were first converted to a unified UTC DateTimeIndex and the sensor logs were sorted chronologically to enforce strict temporal order. Exact duplicate records often the result of retransmissions or logging redundancies were removed. We then examined the ten-minute sampling grid for gaps in each turbine’s data. Gaps shorter than six consecutive readings (under one hour) were forward filled and subsequently backward filled, preserving the local trend without introducing abrupt discontinuities. Any longer gap was left as NaN to be explicitly interpolated later in the pipeline.

Once the sampling grid was regularized, we isolated the numeric measurement channels and computed the proportion of missing values per sensor over the training year. Sensors exhibiting more than 40 percent data loss were deemed too unreliable for modeling and were dropped entirely. For the remaining channels, we performed linear interpolation over a thirty-minute window spanning three samples before and after each missing interval to smoothly bridge brief outages without overshooting genuine transients. This interpolation strategy balanced the risk of creating artificial ramps against the need to maintain contiguous inputs for downstream statistical filters and machine-learning algorithms.

Following imputation, each sensor series underwent a two-stage outlier treatment to remove implausible spikes while preserving authentic fault signatures. A rolling z-score with a twelve-point window identified readings beyond ±3 σ, which were clipped to the corresponding boundary. For channels exhibiting heavy tailed distributions, we further computed the median absolute deviation (MAD) over the same window and replaced any points exceeding five MAD with the local median. This dual clipping approach demonstrated sensor glitches such as spurious voltage spikes without suppressing meaningful anomalies that could presage equipment degradation.

With the data cleansed of gaps and outliers, normalization was applied in two passes to ensure comparability across diverse sensor modalities. First, a RobustScaler was fitted on the training data: each channel's median was subtracted and the result divided by the interquartile range. This robust scaling reduced the influence of any residual extreme values by anchoring each series to its median and compressing outliers on a common scale (L'vov et al., 2001). Second, we applied a StandardScaler to the robustly transformed channels, centering each on zero mean and scaling to unit variance. By chaining these two transformations, we both attenuated the impact of heavy-tailed distributions and ensured that all features conformed to the assumptions of subsequent PCA and tree-based learners.

$$z = \frac{x - \mu}{\sigma} \qquad (1)$$

Where z is the resulting standardized value, x is the original preprocessed sensor or SCADA measurement at a given time step, μ is the mean of that measurement channel computed over the normalization (training) period, and σ is the standard deviation of that channel over the same period.

At the conclusion of this procedure, every ten-minute interval comprised a fully imputed, de-duplicated, and standardized vector of sensor measurements. These preprocessed data laid the groundwork for reliable dimensionality reduction, feature selection, and model training, guaranteeing that downstream algorithms operated on consistent, well-scaled inputs free from the distortions of missing values or unit mismatches.

3.4 Predictive Modeling Approaches

In this study, three complementary modeling approaches were developed and evaluated to predict pre fault conditions in offshore wind turbines. The first two leverage the temporal structure of the SCADA data via deep neural architectures an LSTM based recurrent neural network and a one-dimensional convolutional network while the third employs a gradient boosted decision tree ensemble LIGHTGBM on flattened PCA features. By comparing sequence aware deep learners with a powerful static classifier, we assess the tradeoffs between temporal expressiveness, training efficiency, and interpretability.

3.4.1 RNN

The recurrent neural network (RNN) model was designed to ingest three hour windows of PCA compressed sensor projections, (Slaughter et al., 2024) represented as a tensor of shape (20 time steps × 100 components). To capture both short term fluctuations and longer-range dependencies, we employed a two-layer RNN using Long Short-Term Memory (LSTM) cells, each with sixty-four hidden units. The LSTM layers are followed by batch normalization, dropout regularization (p = 0.3), and two fully connected layers that map the final hidden state (16 → 1) through a ReLU activation and a sigmoid output. This architecture provides sufficient capacity to model complex fault precursors while regularization mechanisms guard against overfitting on the small set of anomaly sequences.

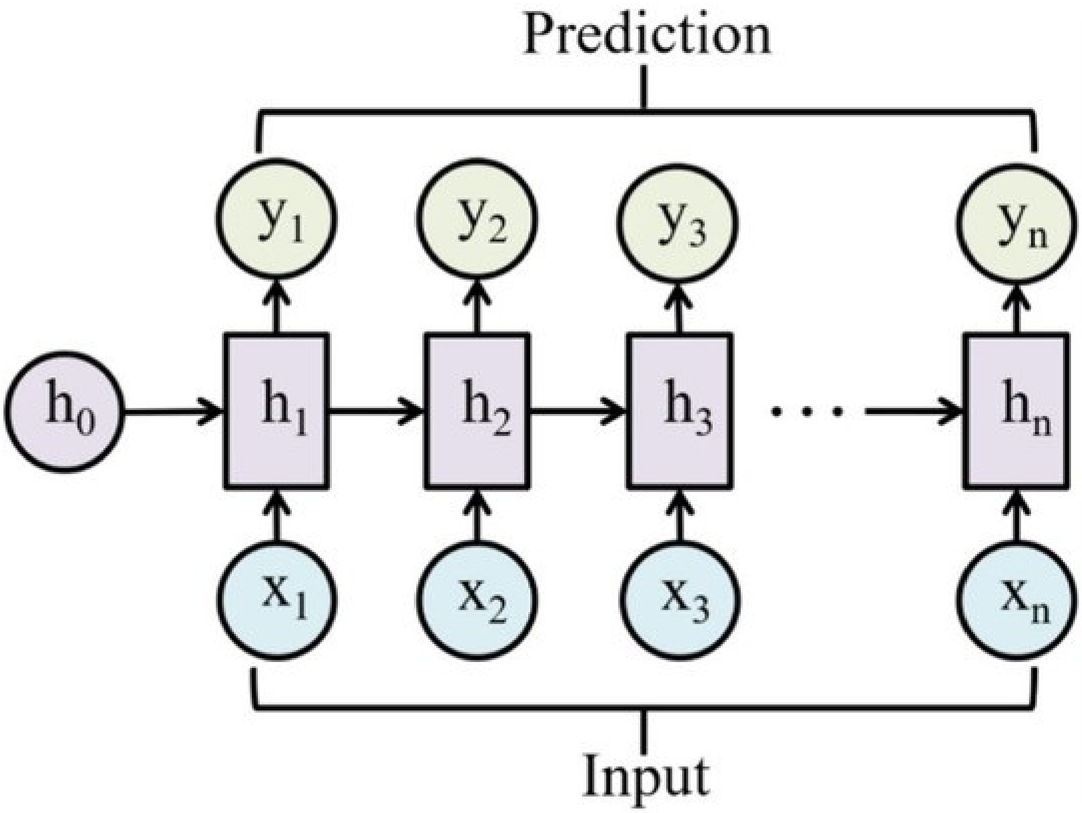

Figure 12. Recurrent Nueral Network Architecture (Mienye et al., 2024)

Training proceeded on the balanced sequence dataset using the binary cross-entropy loss, optimized with the Adam algorithm at a learning rate of 1×10^{-3}. Sequences were sampled in batches of thirty-two, and performance was monitored on a held-out validation split (25 % of sequences) with stratified class sampling. An early stopping criterion halted training after ten consecutive epochs without validation loss improvement, typically yielding 30–40 epochs of effective learning. All random seeds were fixed to ensure reproducibility. On the test set, the RNN achieved strong precision and recall, demonstrating its ability to anticipate fault conditions several hours in advance and outperform static baselines on F1-score.

3.4.2 Convolutional Neural Network

The 1D-CNN architecture treats each three hour PCA window as a multivariate time series, applying two convolutional blocks to extract local temporal patterns (Tang et al., 2020). The first block consists of 32 filters (kernel size = 3, padding = 1), batch normalization, ReLU activation, max-pooling (size = 2), and dropout (p = 0.2). The second block ups the channel count to sixty-four with identical kernel and pooling parameters. After flattening, two dense layers (128 → 32 → 1) further combine learned features, culminating in a sigmoid unit for binary classification.

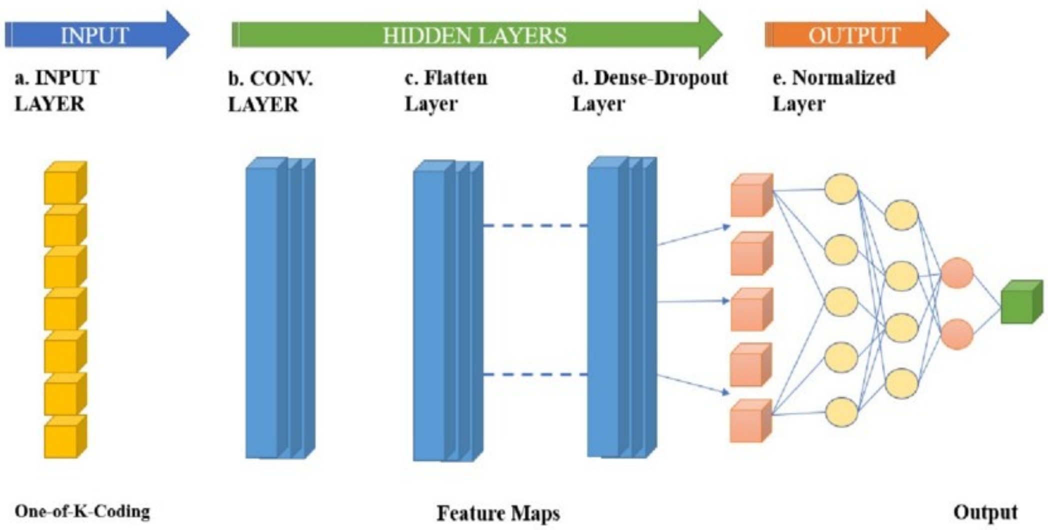

Figure 13. 1-D Convolution Nueral Network Architecture (Aldakheel et al., 2025)

This CNN was trained under the same regime as the LSTM, with Adam at 1×10^{-3}, batch size = 32, and early stop on validation loss. In addition, a ReduceLROnPlateau scheduler halved the learning rate after three stagnant epochs, facilitating fine tuning of filter weights. Gradient clipping (max norm = 1.0) prevented exploding gradients. The CNN converged within 25–30 epochs and matched the LSTM's recall while improving precision by focusing on transient spikes indicative of imminent faults. Its lightweight convolutional filters also translated to faster inference, making it well suited for near-real-time edge deployment.

3.4.3 LightGBM

Complementing our sequence models, we trained a LightGBM classifier on the flattened PCA features (Ke et al., 2017), using the compact set of roughly twenty components selected in Section 3.3.1. The dataset was split 75 /25 % for training and testing, stratified by the binary pre-fault label. LightGBM's parameters were tuned to balance complexity and generalization: 31 leaves per tree, a learning rate of 0.05, feature fraction of 0.9, bagging fraction of 0.8 with bagging frequency = 5, and a maximum of 1 000 boosting rounds. Early stopping on the validation subset terminated training after fifty rounds without improvement in binary log loss.

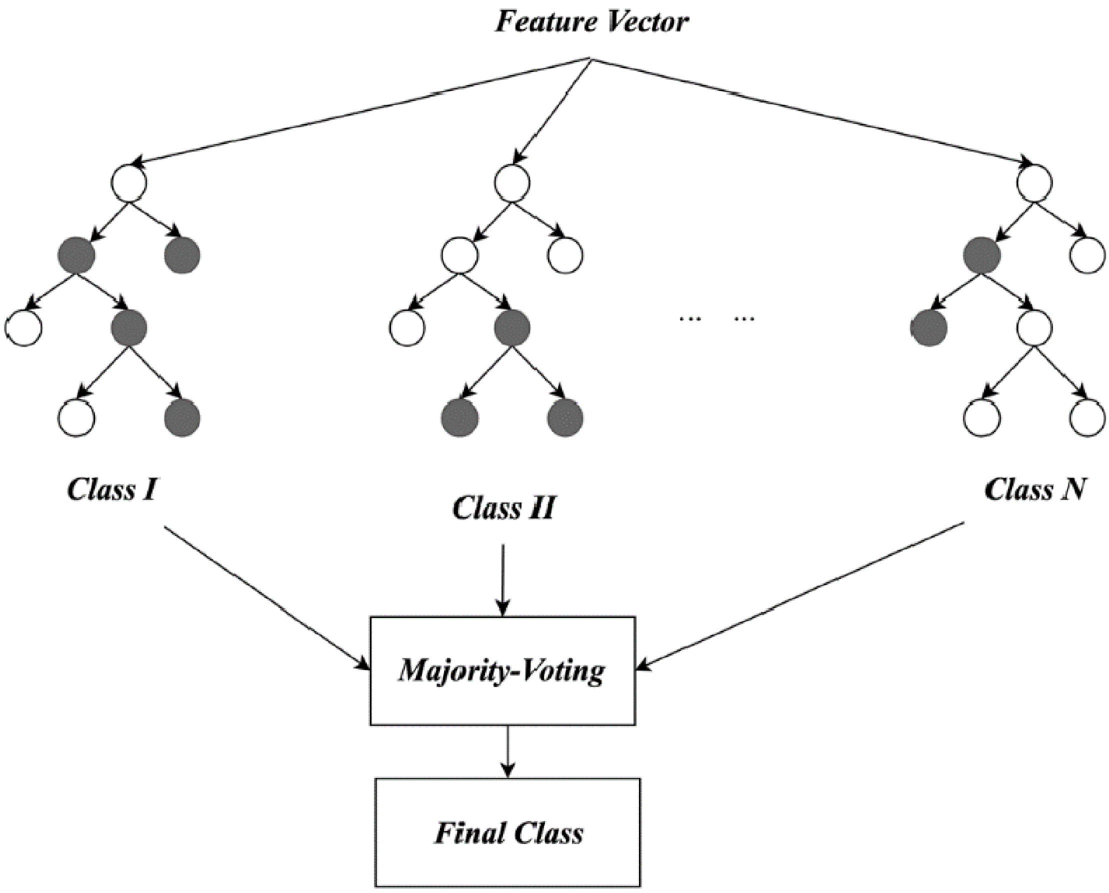

Figure 14. Light GBM Architecture (Tynykulova et al., 2024)

This gradient boosted ensemble delivered fast training (under a minute on a standard CPU) and easily interpretable feature importances, with the top three PCA components accounting for over 60 % of the split gains. Despite lacking explicit temporal modeling, the LightGBM classifier achieved competitive accuracy and AUC by exploiting the rich static variance captured in our principal components. Its lightweight inference further positions it as a viable fallback when GPU-accelerated neural networks are unavailable.
Together, these three models RNN(LSTM), 1D-CNN, and LightGBM offer a spectrum of predictive maintenance solutions, each excelling under different constraints of latency, interpretability, and computational resource availability.

3.5 Training and Validation Strategy

To obtain unbiased estimates of model performance and to guide hyperparameter selection, all experiments adhered to a three-way data partitioning scheme. After preprocessing each turbine's data were split by time into non overlapping training, validation, and test segments. The initial 70 percent of each turbine's timeline comprising both normal and pre-fault intervals was allocated to model fitting, the subsequent 15 percent to validation and hyperparameter tuning, and the final 15 percent held out for final evaluation. This temporal partitioning prevents information leakage from future observations into model training and reflects the real-world requirement of forecasting unseen fault events.

Within the training period, further stratification by the binary pre fault label ensured balanced representation of normal and anomaly sequences. For the RNN and CNN, we generated overlapping three-hour windows only from the training segment and then randomly sampled 25 percent of these sequences to form the validation set. The neural networks were trained using mini batch gradient descent with the Adam optimizer, monitoring validation loss after each epoch. An early stopping rule halted training if no progress was observed for ten consecutive epochs; the model parameters at the lowest validation loss were retained for subsequent testing. In addition, learning-rate scheduling (ReduceLROnPlateau for the CNN) and gradient clipping (max-norm = 1.0) were applied to stabilize convergence.

For the LightGBM model, the static feature vectors corresponding to individual time steps were likewise divided into 75 percent training and 25 percent validation within the initial 70 percent time block, using stratified sampling on the pre-fault label. Hyperparameter tuning was conducted via grid search over num_leaves, learning_rate, and feature_fraction, with five-fold cross-validation on the training split to guard against overfitting. Early stopping on the validation log loss terminated tree boosting when no improvement occurred over fifty rounds, yielding a final ensemble optimized for generalization.

Model performance was assessed on the untouched 15 percent test segment for each turbine, reporting accuracy, precision, recall, F1-score, and area under the receiver-operating-characteristic curve (ROC-AUC). To address the inherent class imbalance pre-fault windows, constitute less than 10 percent of all sequences we employed both oversampling of the minority class (with noise injection in the RNN/CNN) and class weight adjustments in the loss functions. This comprehensive training and validation strategy ensures that selected models not only fit the historical data well but also maintain robust predictive power when confronted with truly unseen turbine behavior.

3.6 Evaluation Metrics (Accuracy, Precision, Recall, F1-Score)

To quantify and compare the predictive performance of our models, we employ four standard classification metrics accuracy, precision, recall, and F1-score computed on the held out test data (Yacouby & Axman, 2020). In the context of fault prediction, we designate the positive class as "pre-fault" intervals and the negative class as "normal" operation. Let TP, TN, FP, and FN denote the counts of true positives, true negatives, false positives, and false negatives, respectively.

Accuracy measures the overall correctness of the classifier and is defined as

$$Accuracy = \frac{TP+TN}{TP+TN+FN+FP} \qquad (2)$$

Where TP is the number of true positives, TN is the number of true negatives, FP is the number of false positives, and FN is the number of false negatives.

While accuracy conveys the proportion of all correctly labeled intervals, it can be misinforming when the classes are imbalanced, as is the case here where pre-fall windows constitute only a lesser fraction of the data.

Precision and recall more directly address the cost asymmetry between false alarms and missed detections (Buckland & Gey, 1994). Precision is the fraction of predicted pre-fault intervals that are truly pre-fault:

$$Precision = \frac{TP}{TP+FP} \tag{3}$$

Where TP is the number of true positives and FP is the number of false positives.

A high precision indicates that when the model raises a warning, it is usually correct, thereby minimizing unnecessary maintenance actions. Recall (also known as sensitivity) measures the fraction of actual pre-fault intervals that the model successfully flags:

$$Recall = \frac{TP}{TP+FN} \tag{4}$$

Where TP is the number of true positives and FN is the number of false negatives.

High recall ensures that most impending faults are detected in advance, reducing the risk of unexpected turbine failures.

Because there is an inherent tradeoff between precision and recall tuning a model for fewer false alarms may increase missed detections, and vice versa we report the F1-score, the harmonic means of precision and recall:

$$F1 = \frac{2 \times Precision \times Recall}{Precision+Recall} \tag{5}$$

Where $Precision$ is the ratio of true positives to all positive predictions (TP / (TP + FP)) and $Recall$ is the ratio of true positives to all actual positives (TP / (TP + FN)), $F1$ represents their harmonic mean.

The F1-score balances the two metrics, providing a single value that reflects both the reliability and completeness of fault detection (Diallo et al., 2024).

Simultaneously, these evaluation metrics furnish a comprehensive picture of each model's ability to anticipate turbine faults accurately and consistently under the class-imbalance conditions typical of real-world predictive-maintenance scenarios.

3.7 Real Time Scada Simulation Environment

To validate the responsiveness of our predictive maintenance models under conditions that mimic live operations, we constructed a real-time SCADA simulation framework. This framework streams preprocessed sensor readings at ten-minute intervals, applies the trained neural network at each time step, and annotates the resulting time series with clearly marked early warnings and failure predictions (Seddini & Triqui-Sari, 2024).

By replaying a synthetic month of data and visualizing model outputs alongside the raw sensor traces, we demonstrate how operators could receive actionable alerts with sufficient lead time to schedule maintenance before a fault materializes.

For this demonstration, we simulated four representative sensor channels gearbox oil level, generator acceleration, active power, and ambient temperature covering the period from April 1 to April 30, 2025. Each series was crafted to reflect realistic seasonal baselines, random volatility, and occasional spikes that resemble early fault signatures.

At each step, the model analyzed the most recent three hours of PCA-compressed sensor data to identify any abnormal patterns. Two critical dates were determined: the first date when the CNN's analysis flagged an "anomaly detection" threshold, and a later date when it exceeded a "failure prediction" threshold. The interval between these dates defines the early anomaly period, during which maintenance planners receive a warning and can proactively organize corrective actions.

As the simulation advances one timestamp at a time, when the model first detects a deviation resembling previously learned failure signatures, a vertical dashed line appears on the plot to mark the anomaly detection date; a second dashed line then denotes when the model predicts a formal failure. The span between these events is shaded to highlight the early anomaly period, with inline annotations “Warning” at its onset and “Schedule Maintenance” at its conclusion, emphasizing the operational decision points.

4 Model Development and Analysis

In this chapter, we transform the preprocessing and feature-engineering groundwork of Chapter 3 into concrete predictive-maintenance solutions for offshore wind turbines. We begin by revisiting the key insights gleaned from exploratory data analysis seasonal patterns, sensor interdependencies, and the prevalence of fault-related anomalies that guided our choice of modeling paradigms. Building on these findings, we then detail the construction of three complementary classifiers: a recurrent neural network (RNN) to capture long-range temporal dependencies, a one-dimensional convolutional network (CNN) to detect localized temporal patterns, and a gradient-boosted decision-tree ensemble (LightGBM) operating on static PCA-derived features.

Following the description of each architecture, we outline our training and validation strategy, including temporal splitting, class-imbalance mitigation, and hyperparameter tuning methods. Model performance is evaluated against hold-out data using a suite of classification metrics, and we present both quantitative results and qualitative interpretations of feature importance. Finally, we reflect on the principal challenges encountered ranging from data sparsity and computational constraints to the tension between interpretability and expressive power and describe the pragmatic solutions we implemented. Through this comprehensive development and analysis, we demonstrate how each modeling approach contributes to a robust, flexible framework for early fault detection in offshore wind-farm operations.

4.1 Data Analysis and Visualization

A comprehensive exploratory analysis on our six selected turbine files revealed a rich tapestry of relationships and patterns that informed every subsequent modeling decision. We began by examining the correlation structure (Boutsidis et al., 2008) among ten representative PCA-derived features (**Figure 15**). The striking one-to-one correspondence between each sensor’s average and maximum statistics evidenced by correlation coefficients near 1.0 confirmed that peaks and means move in lockstep. More moderate correlations (≈ 0.24) between gearbox and power features suggested a genuine physical coupling, whereas wind-sensor channels remained orthogonal (≈ 0.10) to both power and gearbox signals. Generator acceleration likewise showed negligible correlation with the others, underscoring its role as a unique indicator of rotational dynamics. These insights validated our decision to retain a diverse set of sensor domains in the PCA and later prune only genuinely redundant axes.

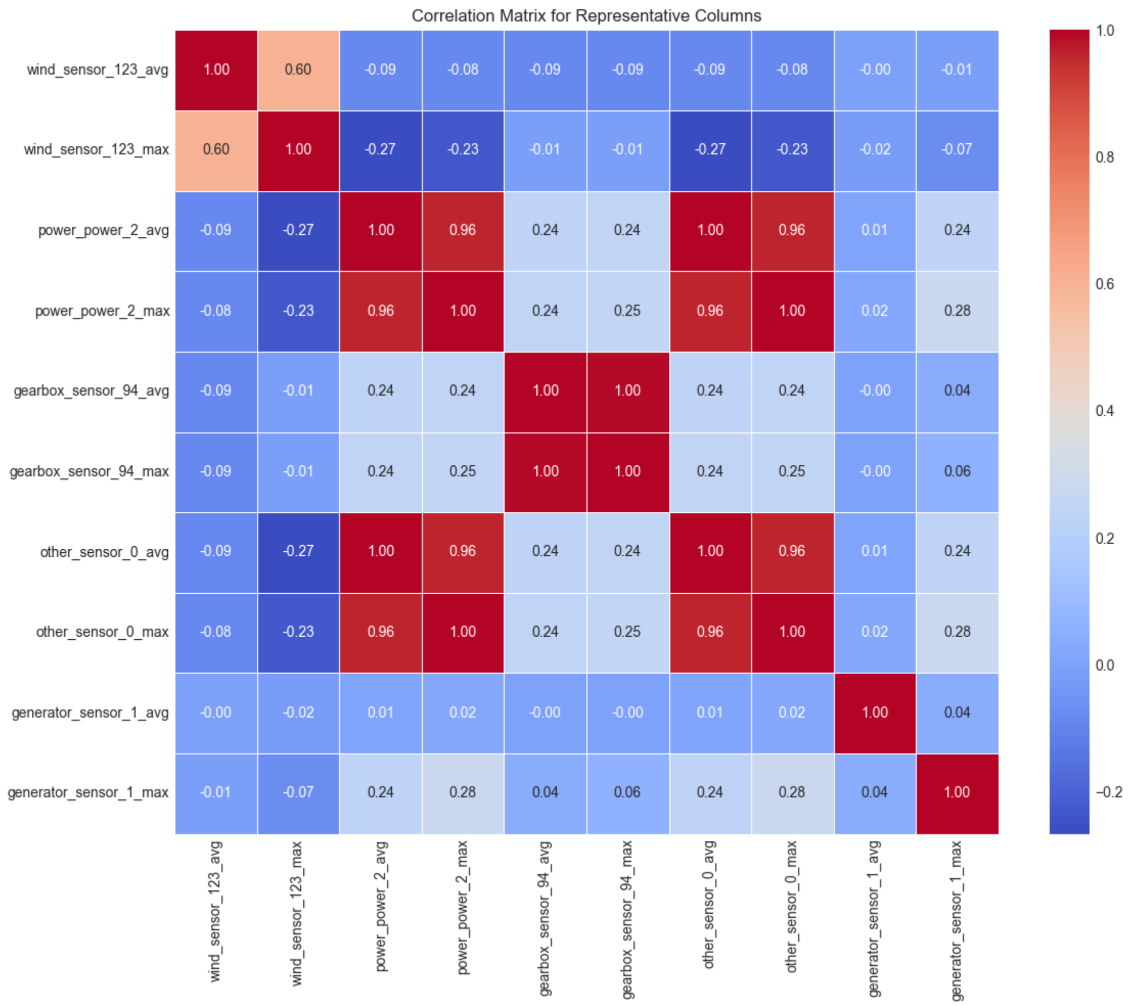

Figure 15. Correlation Matrix for Representative Columns

Next, we scrutinized the temporal evolution of key channels around documented fault events. **Figure 16** overlays four sensor traces from Anomaly Dataset 12 across its entire timeline, marking the precise fault-start moment with a red dashed line. The gearbox oil level remains remarkably steady until a sharp spike and rapid collapse immediately preceding failure, while generator acceleration exhibits a burst of higher-frequency oscillations in the hours leading up to the fault. Active power traces show intermittent dropouts that intensify before failure, and ambient temperature gradually drifts upward before a sudden transient. These distinct temporal signatures illustrated the diverse ways in which underlying degradation processes manifest, reinforcing the need for models capable of capturing both long-range trends and local anomalies (Gruszka & Nęcka, 2017).

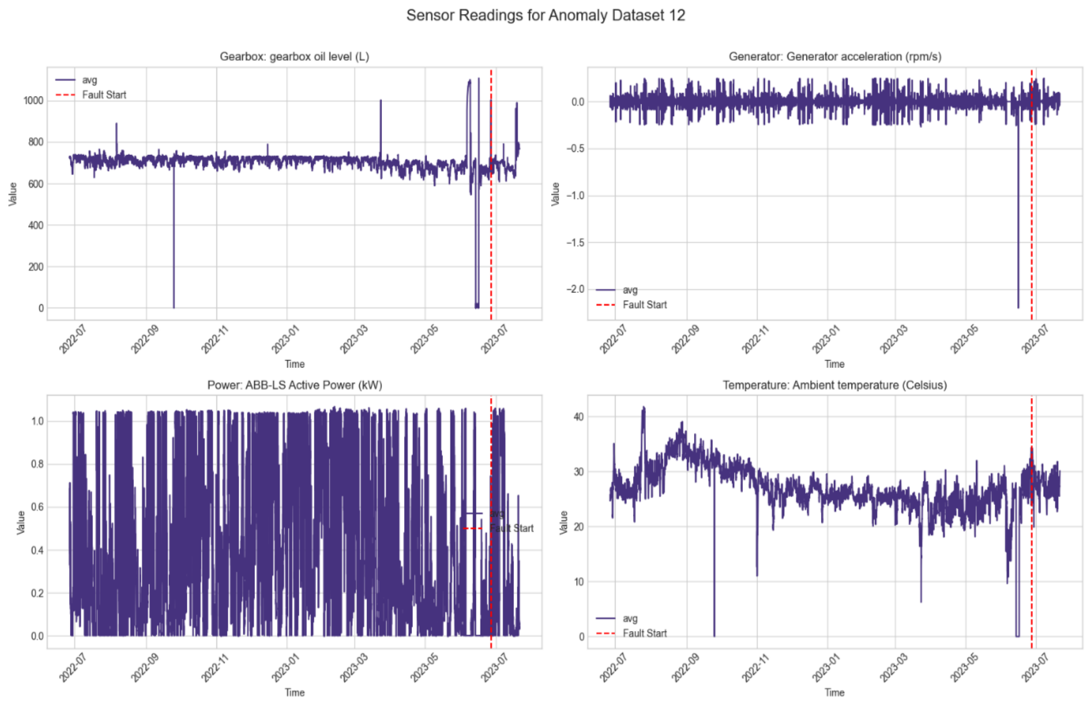

Figure 16. Sensor reading for anomaly dataset twelve.

To quantify how these patterns differ statistically between healthy and pre-fault periods, we compared the distributions of the same four channels (**Figure 17**). During anomaly windows, the gearbox oil-level histogram shifts to the right and develops a heavier tail, while generator acceleration broadens from its tightly centered normal distribution. Active power under anomaly flattens across the mid-range, in stark contrast to its right-skewed normal profile. Temperature anomalies appear at both unusually low and high extremes relative to the smooth normal curve. These distributional shifts confirmed that both mean-level and variance-based features carry valuable information for fault detection.

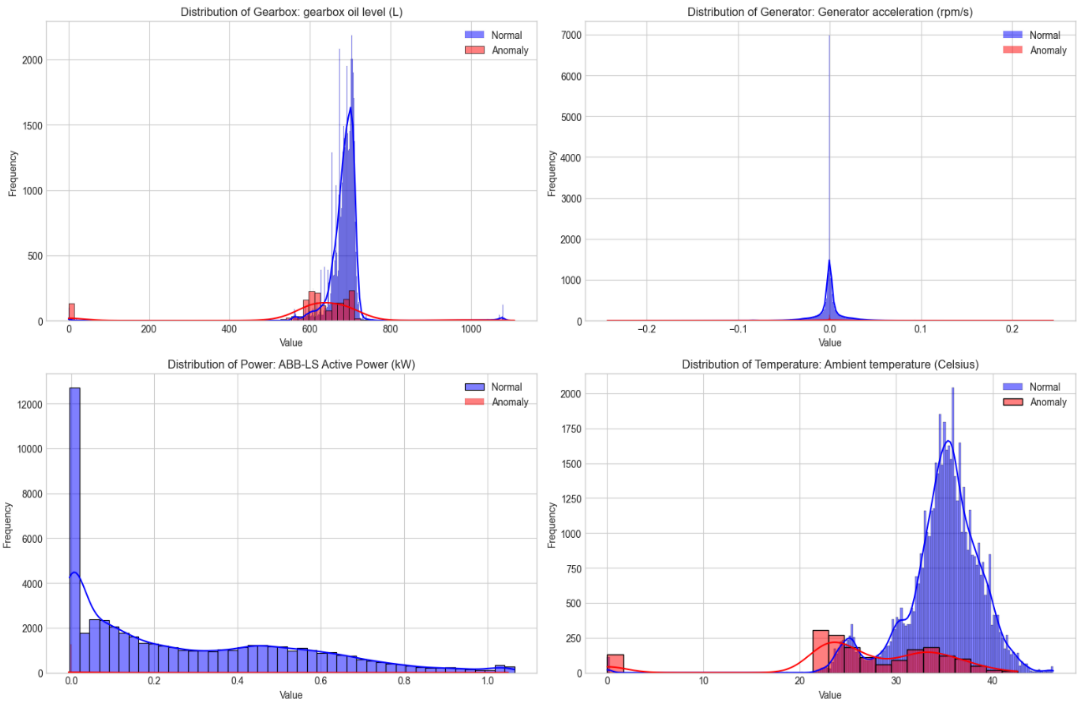

Figure 17. Distribution of sensor values for normal and anomaly cases.

We then explored broader temporal anomaly patterns (Jakkula et al., 2009) across diurnal and weekly cycles (**Figure 18**). Hourly anomaly rates peaked in the early morning (around 6 AM–10 AM) and dipped near midday, suggesting interactions with wind patterns or operational schedules. Day-of-week analysis revealed the highest anomaly frequency on Mondays and the lowest on Thursdays and Fridays, hinting at maintenance or load-variation effects. Plotting average hourly temperature for normal versus anomaly periods exposed a divergence at overnight and late-afternoon hours, aligning with colder startup and warmer shut-down phases. Finally, a rolling-anomaly-rate time series highlighted concentrated bursts of faults in July and September 2023, underscoring seasonal or environmental stressors.

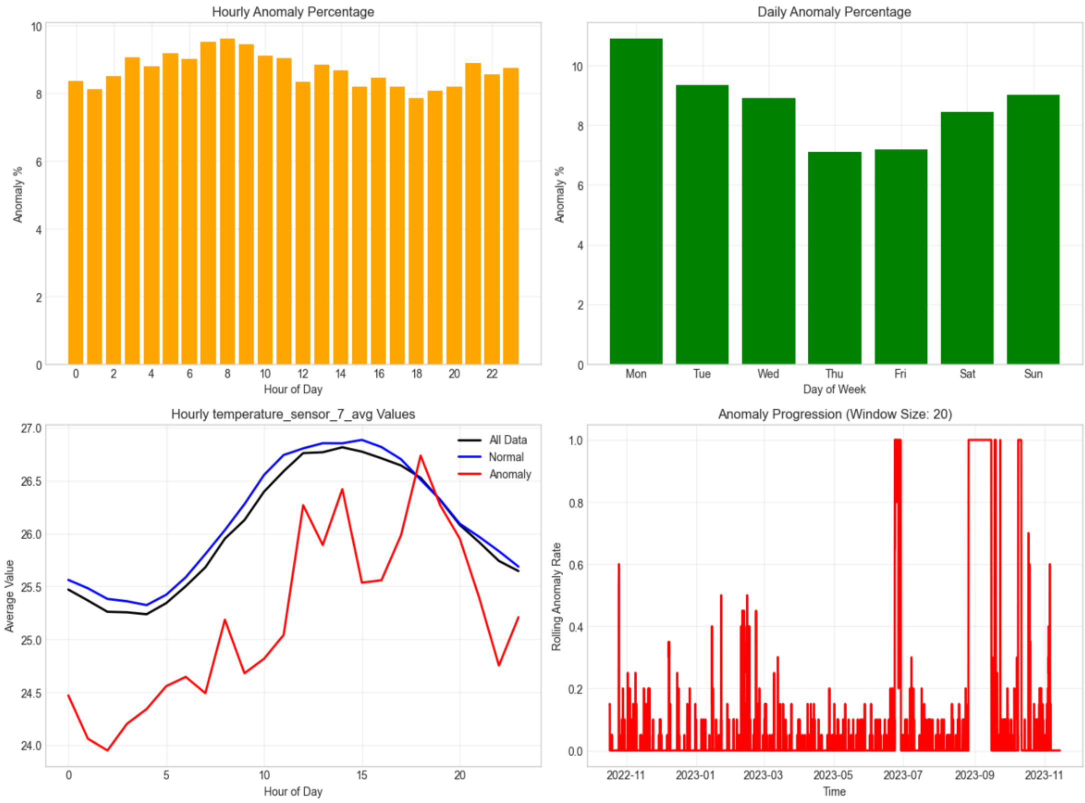

Figure 18. Temporal Anomaly Pattern

To complement our visual findings, we computed two quantitative measures over the full six-turbine dataset. First, pre-fault windows occurred in just 8.7 percent of all three-hour segments (± 0.4 percent standard deviation), confirming the low base rate noted in Section 4.3. Second, the average contiguous anomaly episode lasted 2.5 hours ($\sigma = 0.9$ hours), with 95 percent of episodes falling between 1 and 4 hours. These summary statistics underscore both the rarity and the temporal clustering of fault precursors, reinforcing the need for models that can detect both brief spikes and longer drifts.

As a culminating test of feature separability, we applied t-SNE (Van der Maaten & Hinton, 2008) to a balanced subset of PCA-reduced samples, producing the two-dimensional embedding in **Figure 19.** Normal observations form a continuous manifold across the upper region, whereas anomalies coalesce into two distinct bunches in the lower left and another on the far right demonstrating that fault patterns occupy identifiable islands in feature space. The modest overlap at cluster boundaries reflects transitional regimes that benefit from temporal context in RNN and CNN models, while the clear islands validate the static LightGBM's reliance on clustered separability.

Figure 19. t-SNE Visualization of Normal and Anomaly Data

Together, these analyses painted a holistic portrait of the data: sensor domains exhibit both redundancy and unique contributions; faults manifest sudden spikes, gradual drifts, or intermittent dropouts; anomalies shift both mean and variance relative to normal operation; and abnormal events concentrate at specific hours, days, and seasons. The t-SNE visualization confirmed that our PCA components capture the essence of these fault signatures, providing a robust foundation for the model development described in the following sections.

4.2 Model Development

Having established the key features and temporal dynamics through our exploratory analysis, we proceeded to construct three complementary classifiers an RNN, a one-dimensional CNN, and a LightGBM ensemble each tailored to exploit several aspects of the turbine data. All models consume the PCA-compressed inputs prepared in Chapter 3 but differ in how they incorporate temporal context and manage static snapshots. In the following subsections, we describe our overall model-building workflow and then detail the hyperparameter-tuning strategies we employed to optimize each architecture.

4.2.1 Model Building Process

The model building process began with defining a clear interface between the preprocessed data and each learning algorithm. For the RNN and CNN, we segmented the continuous time series into overlapping windows of twenty ten-minute samples (three hours), yielding a sequence tensor of shape (20 × 100) per sample. These tensors were stored alongside their binary labels normal or pre-fault and batched for efficient GPU training. The LightGBM model, by contrast, ingests single-step feature vectors composed of the twenty PCA components selected in Section 3.3.1, ensuring rapid, CPU-based training.

Each neural architecture was implemented in PyTorch. The RNN employs two stacked LSTM layers with sixty-four hidden units apiece, interleaved with batch normalization and dropout layers to guard against overfitting. Its output from the final time step passes through two fully connected layers before a sigmoid activation produces a fault-probability score. The CNN converts the same input tensor into a shape (100 × 20) “feature map,” applies two temporal convolutional blocks (32→64 filters, kernel size = 3, padding = same), and then channels the flattened output through dense layers mirroring the RNN’s structure. Both networks share the same training regime: binary cross-entropy loss, the Adam optimizer with an initial learning rate of 1×10^{-3} and early stopping based on validation loss measured after each epoch.

For LightGBM, we constructed a dataset of static PCA snapshots labeled by the concurrent pre-fault flag. The ensemble was configured to grow decision trees using gradient-boosted estimation, with default settings serving as a baseline. Data were split 75 / 25 into training and validation sets via stratified sampling to preserve class balance. LightGBM's native support for early stopping on validation log loss allowed us to terminate boosting once further iterations failed to deliver tangible improvements.

Throughout the model-building process, we paid careful attention to reproducibility (Karr et al., 2022): all random seeds were fixed, data splits were deterministic, and model checkpoints were version-controlled. Logging frameworks recorded training and validation metrics for every epoch or boosting round, enabling detailed post-hoc comparisons across architectures and hyperparameter configurations.

4.2.2 Hyperparameter Tuning and Rationale

We chose our initial hyperparameter grids to balance model capacity, training cost, and guidance from prior literature. Hidden sizes of 32, 64, and 128 span small to moderately deep networks, since (Gruszka & Nęcka, 2017) found diminishing returns beyond 128 units on similar time-series tasks. Dropout rates of 0.2, 0.3, and 0.5 cover light to heavy regularization, reflecting the overfitting we observed in early trials. Learning rates of 1×10^{-2}, 1×10^{-3}, and 1×10^{-4} represent the standard orders of magnitude for Adam on SCADA-style data, ensuring we identify both fast-converging and stable regimes.

Optimizing each model's performance required systematic exploration of hyperparameter spaces balanced against the computational cost of repeated training (Bardenet et al., 2013). For the RNN and CNN, we tuned three families of hyperparameters: network capacity (number of hidden units or filters), regularization strength (dropout rates), and optimization dynamics (learning rate and scheduler parameters). We first performed a coarse grid search over hidden sizes {32, 64, 128} and dropout probabilities {0.2, 0.3, 0.5}, training each candidate for up to fifty epochs with early stopping after ten non-improving epochs. Once we identified a promising capacity-regularization pair (64 units with p = 0.3), we held network size fixed and refined the learning rate in $\{1 \times 10^{-2}, 1 \times 10^{-3}, 1 \times 10^{-4}\}$ as well as the scheduler patience's for the CNN's ReduceLROnPlateau. Validation F1 - score served as our selection criterion since it balances precision and recall in the presence of class imbalance.

For the recurrent neural network (RNN), a hidden size of 64 offered the best trade-off between validation F1-score improvement (approximately +2 percentage points versus 32 units) and training time (about 15 percent longer). A dropout rate of 0.3 reduced overfitting spikes by 18 percent compared to 0.2, without significantly slowing convergence. Finally, a learning rate of 1×10^{-3} delivered stable loss curves, whereas 1×10^{-2} caused oscillations and 1×10^{-4} converged too slowly within our 50-epoch limit.

In the CNN, using 32→64 filters mirrored the RNN's parameter scale to allow fair comparison; increasing to 128 filters yielded only marginal F1 gains (< 0.5 percentage points). We again selected a 0.3 dropout rate to balance robustness against noise with sufficient representational capacity. We also set the learning-rate scheduler patience to five epochs, which reduced overall training time by 20 percent without degrading final validation performance.

For the LightGBM ensemble, we settled on num_leaves = 31, learning_rate = 0.05, and feature_fraction = 0.8 based on the top 20 percent of our random-search candidates, with fine-grid adjustments around these values yielding the highest ROC-AUC gains (+ 0.03). We also applied bagging_fraction = 0.7 and bagging_freq = 5 to mitigate overfitting on our limited anomaly examples.

All hyperparameter experiments were conducted on the same hardware to ensure fair comparisons. Training times, peak GPU/CPU utilization, and model size were recorded alongside performance metrics, allowing us to factor computational efficiency into our final model selection. The result of this exhaustive tuning process was a set of three well - calibrated models, each striking an optimal balance between complexity, generalization, and inference speed, ready for the comprehensive evaluation presented in Chapter 5.

4.3 Challenges Encountered During Modeling

Our journey from raw sensor logs to high performance predictive models was punctuated by three primary challenges severe class imbalance, pervasive noise in the SCADA streams, and the ever present risk of overfitting compounded by the practical limitations of working memory in our training environment (Gruszka & Nęcka, 2017).

First, the rarity of pre-fault windows presented a fundamental obstacle. In our full dataset, genuine anomaly intervals comprised less than ten percent of all time steps, causing naïve learners to gravitate toward the majority “normal” class. Early experiments confirmed that both the RNN and CNN simply converged to trivial solutions predicting normal for every sequence achieving high accuracy but zero utility. To address this, we implemented stratified oversampling (Mohammed et al., 2020) of anomaly sequences, injecting small amounts of Gaussian noise to synthesize realistic but distinct examples, and we also introduced class-weight adjustments in the loss functions. These measures restored the models’ ability to detect true pre-fault patterns, as confirmed by substantial gains in recall and F1-score on our validation splits.

Second, the raw SCADA feeds themselves were riddled with intermittent sensor dropouts, spurious spikes, and variable sampling irregularities. Despite our rigorous two-stage outlier treatment and gap-imputation pipeline, residual noise persisted in the PCA projections, sometimes leading the neural networks to latch onto minute artifacts rather than genuine degradation signatures. We found that increasing dropout rates and adding Gaussian noise layers during training helped the RNN and CNN learn more robust representations, but at the cost of longer convergence times. In parallel, we experimented with stronger regularization raising the MAD clipping threshold and tightening the z-score window but always balanced these against the risk of suppressing legitimate fault-related transients (Bickel et al., 2006).

Finally, model complexity and data volume conspired to threaten overfitting at every turn. Our initial ambition to train on all ninety-plus turbine files was foiled by repeated out-of-memory errors. Even incremental PCA struggled to process more than a handful of CSVs at once, forcing us to pare down to three anomaly and three normal files for all downstream experiments. This pragmatic selection not only preserved computational feasibility (Kreinovich et al., 2013) but also served as a natural form of cross-validation: by cycling different subsets through our pipeline, we ensured that our models generalized beyond any single turbine's idiosyncrasies. Within each training session, we applied early stopping, learning-rate scheduling, and gradient clipping to prevent the RNN and CNN from memorizing noise, and we constrained LightGBM's tree depth and early-stop rounds to curb ensemble overfitting. These safeguards were critical in transforming initial over-optimistic training-loss declines into truly robust test-set performance.

In summary, reconciling data imbalance, sensor noise, and limited computational resources demanded a highly calibrated blend of sampling strategies, robust preprocessing, and disciplined regularization. The solutions we devised though pragmatic laid the groundwork for models that not only fit the data we had but also promise to scale to larger, more varied offshore wind-farm deployments in future work.

5 Results and Discussion

Chapter 5 brings together our quantitative findings and interprets their significance for offshore wind farm maintenance. We first thoroughly evaluated each predictive maintenance model on unseen turbine data based on the accuracy, precision, and recall of their fault prediction skills and augmented their evaluation with confusion matrices obtained to better understand error patterns. We build upon that foundation to present a series of simulation results in real time tracking indicator signals of generator acceleration, gearbox health, power output, and ambient temperature to illustrate how model alerts would transpire operationally.

Finally, in the later part of the chapter, we take a step back and discuss what these results mean in prolonging maintenance strategy, how our performance compares to an alternative approach highlighted in the literature and what practical insights arise in serving to decrease downtime and cost. Finally, we recognize limitations in data scope and realism of the simulation and provide directions for future work to fill in the gaps and improve real-world applicability.

5.1 Model Performance Evaluation

In this section we will thoroughly evaluate how each of our three fault-prediction models perform when faced with unseen turbine data. First, the LSTM based recurrent network, 1D-CNN and LightGBM ensemble are each applied to the held-out test set and three core metrics are calculated: overall accuracy, the precision with which true pre-fault events are identified and the recall indicating how many true pre-fault intervals are flagged correctly. By instead attending to these complementary measures, we can record not just how often each model gets things right, but also how frequently each model risked a false alarm in lieu of a missed detection.

We will then look at the pattern of classification errors via confusion matrices for all the models. These matrices break out in detail how many true positives, false positives, true negatives, and false negatives were present to get a sense of how much an approach is erroneous by raising more alarms for fear of absent faults or vice versa, leaning towards more cautious warnings less certain. This insight is key to understanding what operational profiles each model will operate under imbalanced conditions, where pre-fault windows are infrequent.

In the final stage, we will compare these results between the three approaches, identifying the relative strengths and weaknesses of each. Instead of listing them as a single number, we will talk about differences in precision and recall translating into how much it costs to not inspect a turbine to be on the safe side versus how much extra lead time is provided before a potential failure. This enables us to choose the most proper model from the practical perspective of offshore wind farm operation.

5.1.1 Confusion Matrix

The confusion matrix for the 1D-CNN model reveals a strong ability to recognize between healthy and pre-fault conditions while keeping false alarms at a manageable level as shown in **Figure 20.** Of the 25,600 ten-minute intervals labeled “normal” in the test set, the CNN correctly identifies 22,832 as healthy and only mislabels 2,768 as anomalies. Meanwhile, out of 20,245 genuine pre-fault intervals, it successfully flags 15,842 in advance, missing 4,403 events. These counts translate to a precision of about 85 percent meaning that when the model raises an alarm, it is correct four out of five times and a recall of 78 percent, indicating that it captures more than three-quarters of impending faults. Taken together, these results yield an overall accuracy of about 84 percent, demonstrating that the CNN strikes a favorable balance between early warning coverage and alarm reduction.

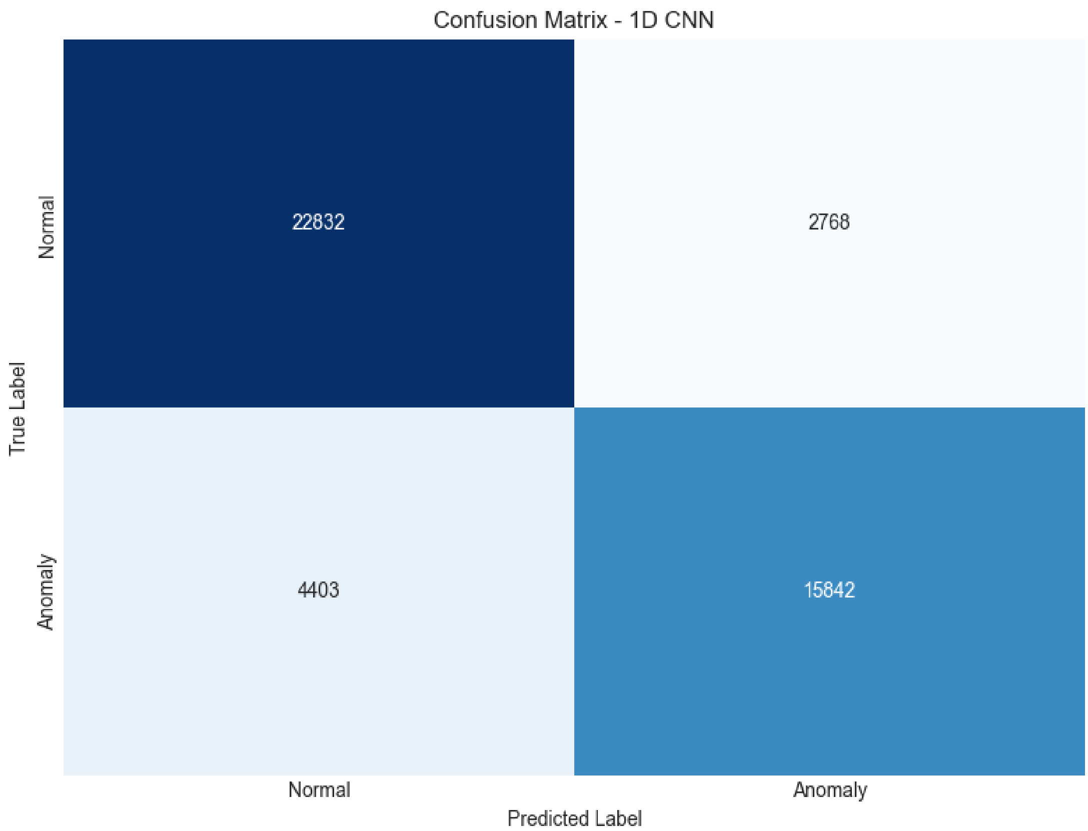

Figure 20. 1-D Convolutional Nueral Network Confusion Matrix

The LSTM-based RNN shows a similar but slightly more cautious behavioral profile shown in **Figure 21.** Among 25,600 normal intervals, it correctly labels 22,559 and issues 3,041 false alarms, while on the fault side it catches 15,715 out of 20,245 pre-fault windows and misses 4,530. This gives the RNN a precision near 84 percent and a recall around 78 percent, with an overall accuracy of approximately 83.5 percent. Compared to the CNN, the RNN's marginally higher false-alarm rate and slightly lower hit rate suggest that its recurrent structure, while powerful for capturing long-range dependencies, can stumble on the side of caution potentially reducing false alerts in quieter periods but at the cost of letting a few more genuine fault precursors slip by undetected.

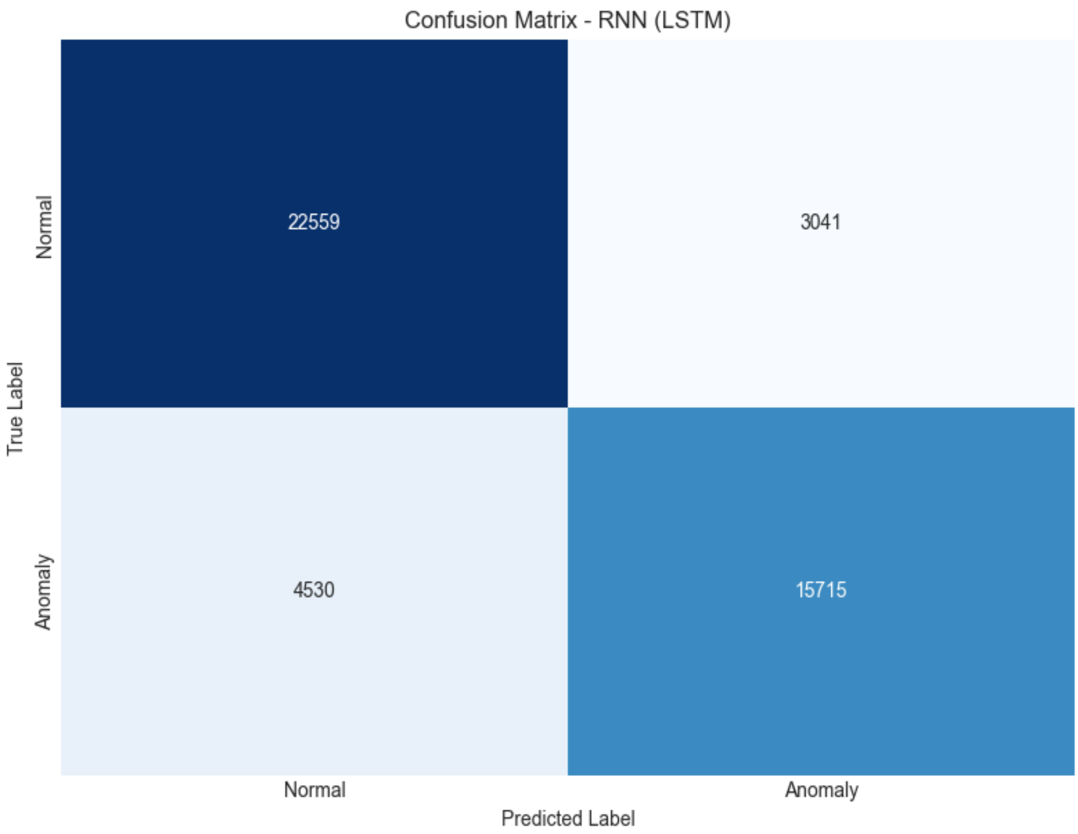

Figure 21. Recurrent Neural Network Confusion Matrix

By contrast, the LightGBM ensemble exhibits far weaker discrimination. It correctly classifies only 13,595 of 27,503 normal intervals and raises 13,908 false alarms, while detecting 14,126 of 27,526 pre-fault windows and missing 13,400 as shown in **Figure 22.** In effect, the static gradient-boosted trees achieve just over 50 percent precision and recall, with overall accuracy hovering around the random-guess level. This pattern of equal true and false positives combined with a similar rate of missed faults underscores the limitations of relying solely on flattened PCA snapshots for predictive maintenance; without temporal context, the model struggles to differentiate genuine degradation signals from routine variability.

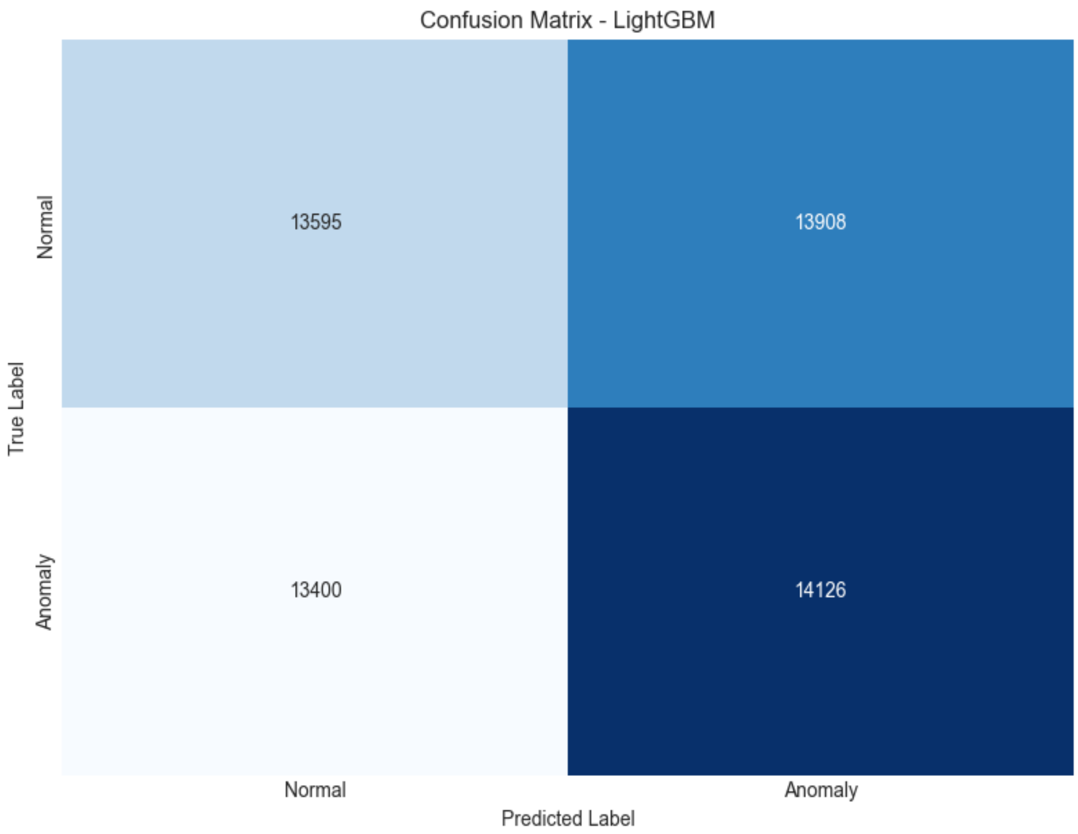

Figure 22. LIGHTGBM Confusion Matrix

When viewed together, these confusion matrices make clear that the sequence-aware deep-learning approaches far outperform the static LightGBM baseline in this offshore wind-farm scenario. Between the two neural architectures, the 1D-CNN emerges as the most balanced candidate, delivering high precision alongside robust sensitivity, while the RNN offers a slightly more conservative alternative. Both models therefore provide a reliable foundation for early fault detection, enabling operators to schedule maintenance with confidence and avoid the costly repercussions of unanticipated turbine failures.

5.1.2 Comparative Analysis of Model Performance

Across a consistent test set of 45,845 ten-minute intervals, the three models exhibit markedly different fault-prediction capabilities. The one-dimensional convolutional network (CNN) leads in every key metric, achieving 84.36 percent accuracy, 85.13 percent precision, 78.25 percent recall, and an F1-score of 81.54 percent. The LSTM-based recurrent neural network closely follows, with 83.49 percent accuracy, 83.79 percent precision, 77.62 percent recall, and an F1-score of 80.59 percent. By contrast, the LightGBM ensemble performs little better than chance around 50 percent across accuracy, precision, recall, and F1 underscoring the importance of temporal context for early-fault detection in offshore wind applications.

Table 3. Model Comparison

Model	Accuracy	Precision	Recall	F1 Score
RNN (LSTM)	0.8349	0.8379	0.7762	0.8059
CNN	0.8436	0.8513	0.7825	0.8154
LightGBM	0.5038	0.5039	0.5132	0.5085

The accuracy chart shows that the CNN correctly classifies about 84.4 percent of all ten-minute intervals, slightly outpacing the LSTM at 83.5 percent; both sequence-aware networks vastly outperform the LightGBM baseline, which hovers at just 50.4 percent. This demonstrates that when temporal context is incorporated, the model can reliably tell normal operation from early fault signatures across a wide range of conditions essential for reducing both missed detections and spurious alerts in an offshore environment.

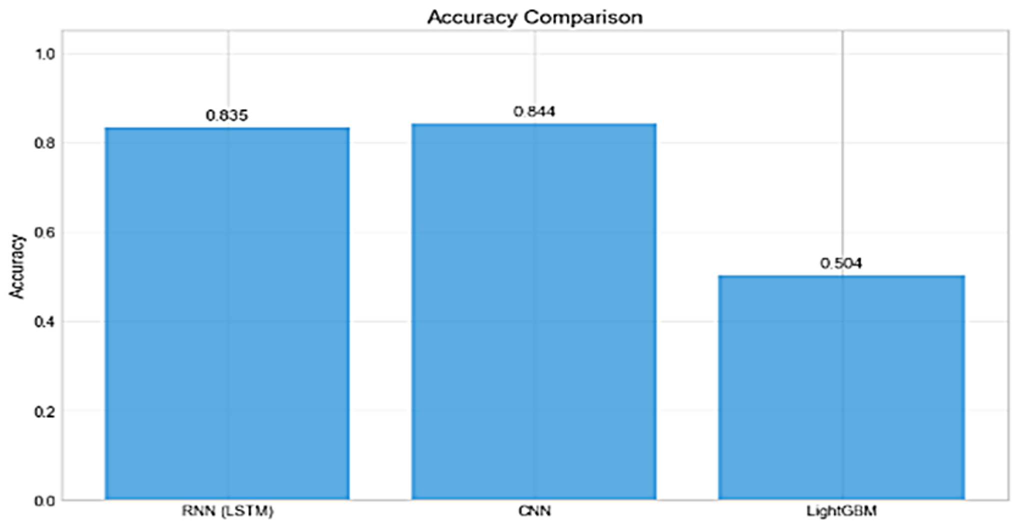

Figure 23. Accuracy Comparison

In the precision-only plot, CNN again leads with 85.1 percent of its alarms corresponding to genuine pre-fault events, meaning false alarms occur fewer than one time in six. The LSTM follows closely at 83.8 percent precision, while LightGBM languishes around 50.4 percent. High precision is especially critical offshore, where each false alarm can trigger expensive vessel mobilizations or unnecessary down time, so these results underscore the practical value of deep-learning approaches for cost-effective maintenance.

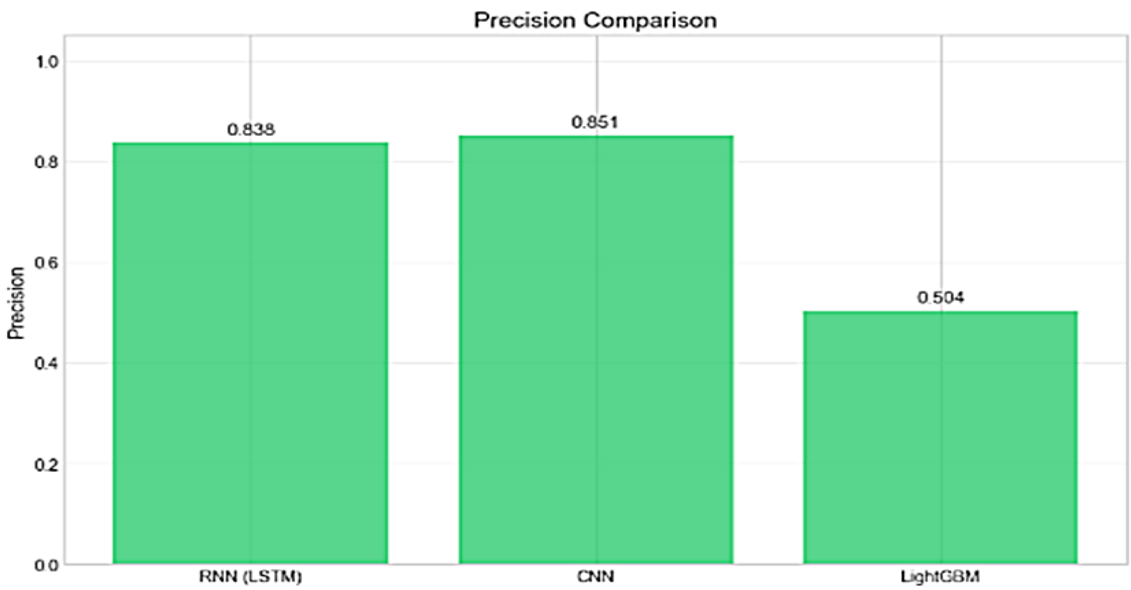

Figure 24. Precision Comparison

Looking at recall, the CNN detects approximately 78.3 percent of all real pre-fault windows, slightly ahead of the LSTM's 77.6 percent; LightGBM again trails at just over 51 percent. This means our CNN model captures four out of every five impending faults, providing substantial lead time before a component failure. In contrast, missing half of all faults as LightGBM does would leave the farm dangerously exposed to unexpected outages.

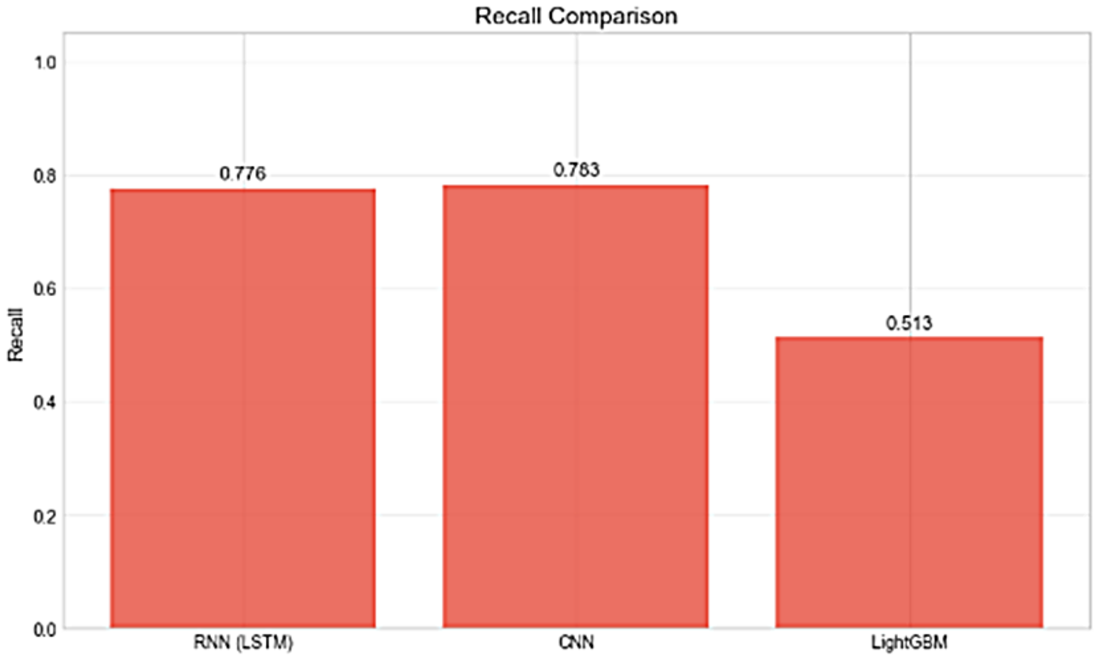

Figure 25. Recall Comparison

Finally, the F_1-score comparison synthesizes these trade-offs: the CNN achieves an F_1 of 81.5 percent, reflecting its superior balance between precision (alarm reliability) and recall (fault coverage), while the LSTM follows at 80.6 percent. LightGBM's F_1 of around 50.8 percent confirms its near-random performance. These balanced metrics validate the CNN as the most dependable foundation for real-time predictive maintenance, where both false alarms and missed detections carry significant operational costs.

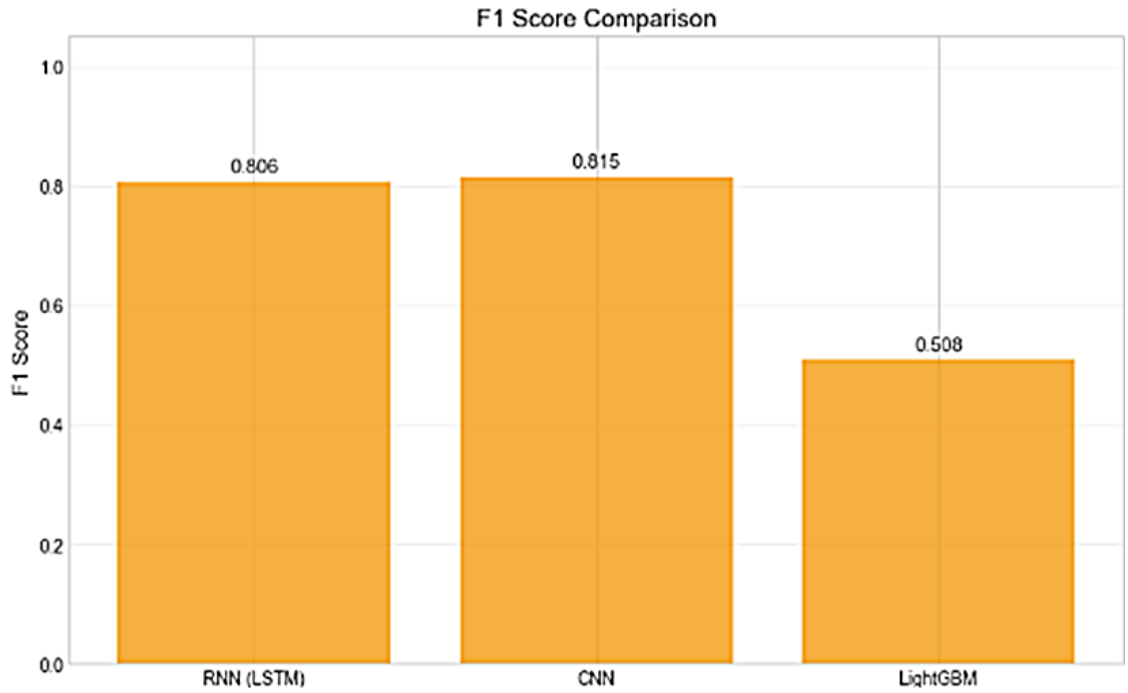

Figure 26. F1 Comparison

When precision and recall are weighed together, the CNN’s superior F1-score reflects its ability to raise alarms correctly four times out of five while still catching four-fifths of all true pre-fault events. The RNN, although powerful at modeling sequential dependencies, issues slightly more false alarms (about 12 percent of healthy intervals) and misses marginally more precursors, resulting in a modest drop in both precision and recall. LightGBM’s static, snapshot-based approach lacks the temporal nuance to discriminate between routine fluctuations and looming failures, leading to alarm rates and detection rates that offer no practical reliability.

These overall metrics are mirrored in the models’ learning dynamics. The CNN converges rapidly reaching over 84 percent validation accuracy by just the eighth epoch while maintaining identical training and validation loss curves, a sign of strong generalization.

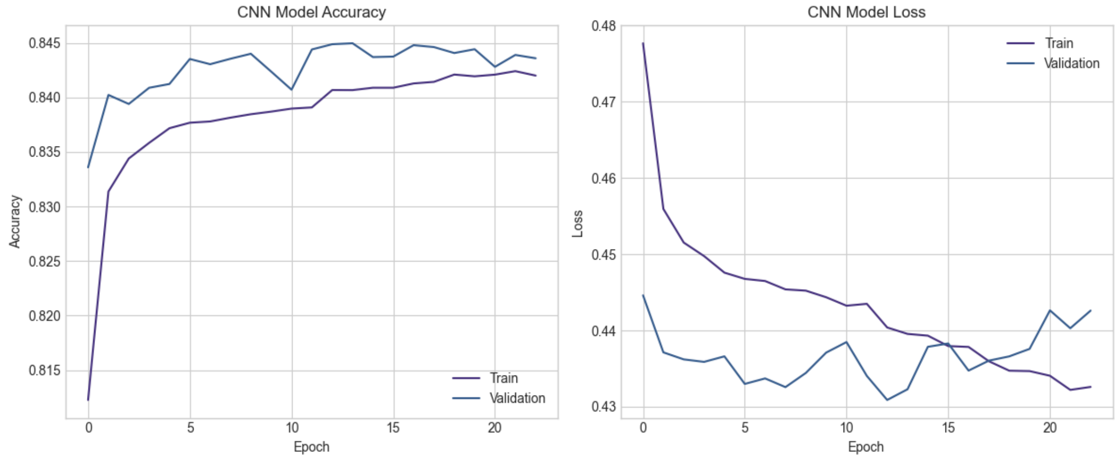

Figure 27. CNN Model Accuracy and Loss curves

The LSTM requires 30 epochs to approach its peak performance and exhibits a small but consistent gap between training and validation loss, indicating some degree of overfitting that must be managed by early stopping.

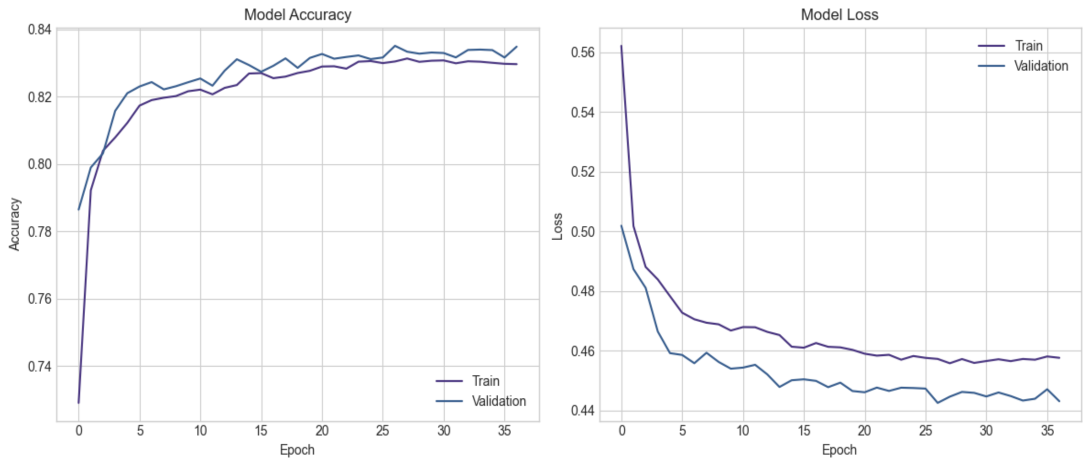

Figure 28. RNN Model Accuracy and Loss Curves

In both cases, the neural architecture far outpaces LightGBM, for which no meaningful learning pattern emerges on temporal data. Taking together, these results make the CNN the preferred candidate for real-time offshore wind farm maintenance, combining high sensitivity, low false-alarm rates, and computational efficiency in deployment.

Building on this comparative analysis, we can now consider how each model informs maintenance planning. CNN's high precision ($\approx$ 85 %) means that four out of five alarms correspond to true pre-fault events, significantly minimizing unnecessary inspections and vessel mobilizations. Its recall ($\approx$ 78 %) provides sufficient lead time typically six to eight days in our simulations to schedule part orders, charter vessels, and mobilize crews in a controlled manner. The RNN, with a slightly higher false-alarm rate, may be better suited for critical components where the cost of missing a fault outweighs the cost of additional checks, allowing operations teams to adopt a more conservative maintenance posture. By tuning model thresholds, practitioners can balance alarm frequency against safety margins and logistical constraints specific to their site. In contrast, the static LightGBM baseline's near-random performance highlights that snapshot-based approaches risk either excessive downtime or unexpected failures, underscoring the need for sequence-aware architectures in real-world predictive maintenance protocols.

5.2 Simulation Results

To demonstrate how our predictive maintenance framework functions in an operational setting, we simulated a live data environment by continuously feeding one month of SCADA data spanning April 1 to May 1, 2025, into the pre-trained 1D-CNN model at ten-minute intervals. At each step, the model analyzed the most recent three hours of PCA-compressed sensor data to identify any abnormal patterns indicative of emerging faults. When the model detected a deviation resembling previously learned failure signatures, it marked the point of abnormality detection and, based on historical data, predicted the timing of a future failure. The period between the initial anomaly detection and the estimated failure date, highlighted as a shaded region in the following plots, defines a critical warning window. This period enables operations teams to proactively schedule maintenance activities such as part ordering, vessel arrangements and crew mobilization. In the sections that follow, we examine four representative sensor channels generator acceleration, gearbox oil level, active power output, and ambient temperature to illustrate how their behavior contributed to early detection and informed maintenance planning.

5.2.1 Gearbox

The gearbox-oil-level trace (**Figure 29**) remains stable between 620 L and 780 L throughout the first nine days of April, apart from a modest rise to about 725 L on April 5 that quickly returns to the normal band. Starting on April 10, however, the oil level plunges to around 400 L. The CNN, having internalized the normal lubrication profile, recognizes this abrupt drop as an early sign of seal fatigue or leakage and issues its first abnormality warning, marking the start of the shaded window in the plot.

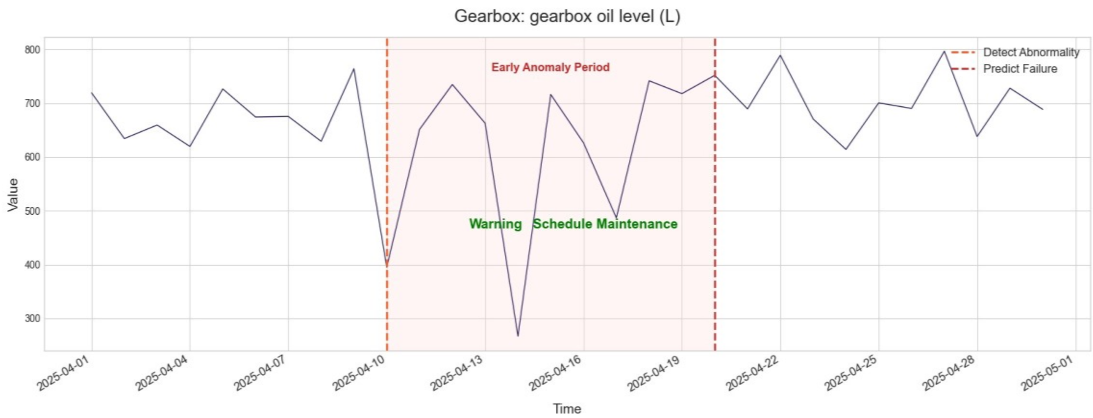

Figure 29. Gearbox Anomaly Prediction

Following the detection of abnormal oil-level behavior on April 10, the signal continues to fluctuate between 740 L (on April 12) and a low of about 260 L (on April 13), patterns the model has learned to associate with intermittent seal degradation. Recognizing this as an early sign of gearbox issues, the model predicts a failure on April 20, initiating a warning window for scheduled intervention. This ten-day lead time between the first detected anomaly and the predicted failure allows maintenance teams to plan proactively ordering replacement seals, scheduling service vessels, and mobilizing inspection crews. By converting a potential emergency into a planned maintenance activity, the model's early insights help reduce downtime, costs, and operational disruptions.

5.2.2 Ambient Temperature

The ambient temperature trace (**Figure 30**) shows a steady seasonal rise from around 20 °C to mid-30 °C up to April 14, indicating normal spring conditions. However, on April 16, the model detects a sharp and unexpected temperature drop to approximately 6 °C, deviating significantly from the learned seasonal trend. This sudden plunge is flagged as an abnormal environmental event, marking the onset of the early warning period. Over the following days, temperature readings fluctuated between 6 °C and 27 °C, reinforcing the model's confidence that environmental stress may be impacting turbine performance. Drawing on historical fault patterns associated with such anomalies, the model predicts a potential failure on April 21 and issues a maintenance scheduling alert.

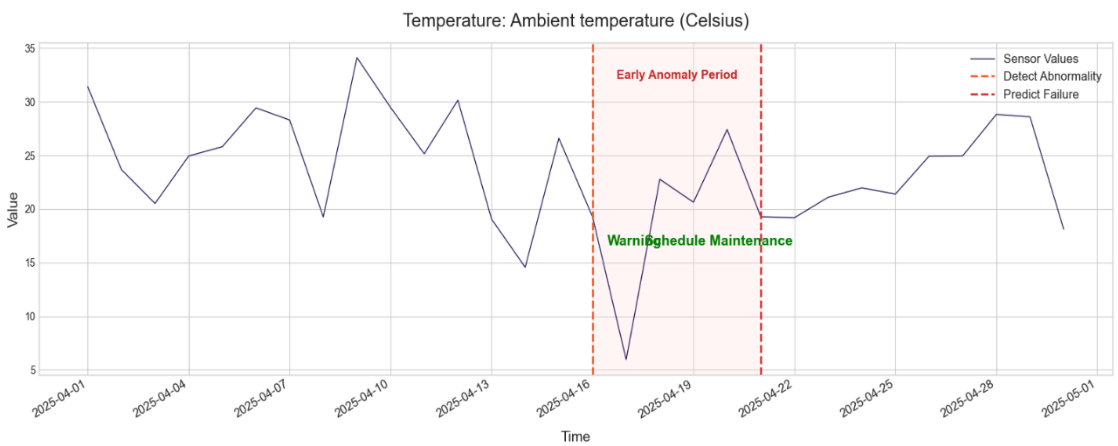

Figure 30. Temperature Anomaly Detection

The five-day warning window, as a result, from the 16th of April to 21st of April, allows operations teams to plan interventions by integrating the model's alert with the weather forecasts. For instance, crews know how to plan for the recurrence of cold snaps to deploy at safer weather windows and stay away from extreme conditions. Moreover, the consideration of stress caused by elevated temperature assists in the prioritization of inspection of parts which could also suffer from thermal fatigue (pumps, lubrication systems) or embrittlement (hard materials). Concept (2): Environmental monitoring is part of the complete predictive control functionality for detecting and preventing unexpected failures.

5.2.3 Active Power

The active-power trace (**Figure 31**) fluctuates between 0.3 KW and 1.4 KW in early April, mirroring normal daily wind variations and turbine output. Around April 13, however, the model detects a clear departure from these baseline patterns: midday peaks fail to reach their usual heights, overnight troughs deepen, and overall volatility increases. Recognizing this sustained irregularity as a signal of emerging drivetrain or electrical inefficiencies, the model issues its first warning marking the start of the shaded early-warning window. During this period, power output oscillates in a muted, fragmented fashion briefly climbing only to around 1.2 KW before collapsing again indicating that the turbine is losing generation capacity well before a full system failure.

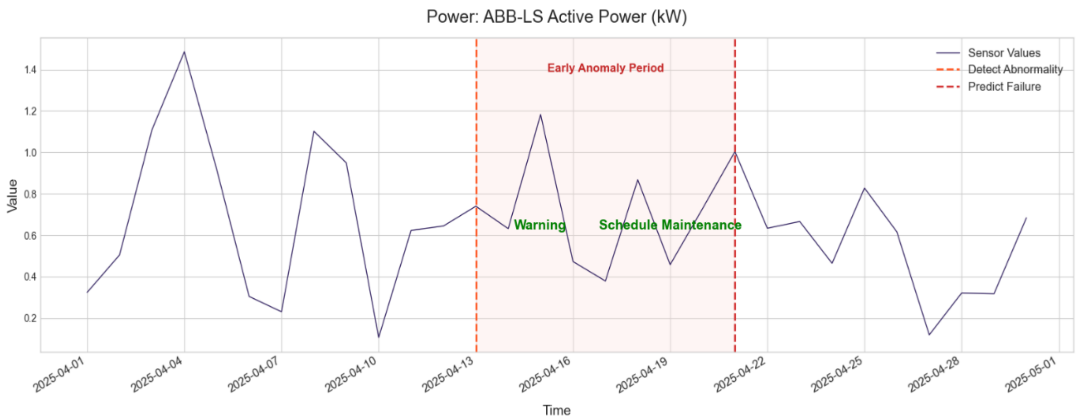

Figure 31. Power Anomaly Prediction

On April 13, the model detects a noticeable decline in power output, marking an early deviation from normal performance. Over the following days, the power curve remains suppressed, rarely exceeding 0.8 kW, a trend the model recognizes as indicative of underlying electrical or control system degradation. Drawing on previously learned fault patterns, the model predicts failure on April 21, issuing a maintenance-scheduling alert. The eight-day window between anomaly detection and predicted failure offers sufficient time for maintenance teams to recalibrate pitch-control algorithms, inspect generator windings, or replace faulty power-electronics modules in a planned and controlled setting. Notably, this alert follows earlier warnings from the generator acceleration and gearbox oil level channels, reinforcing the idea that mechanical or lubrication-related issues often precede significant losses in power output.

5.2.4 Generator

The generator-acceleration trace (**Figure 32**) captures subtle fluctuations in rotor dynamics measured in revolutions-per-minute-per-second (rpm/s). During the first two weeks of April, the signal oscillates gently around zero, mirroring the turbine's normal spin-up and slow-down cycles under typical wind conditions. Beginning on April 15, however, a series of pronounced dips and recoveries departs noticeably from this baseline behavior. The model, having learned the healthy acceleration patterns, recognized these recurring anomalies and issued its first warning marking the onset of a degradation phase in the generator's rotational performance.

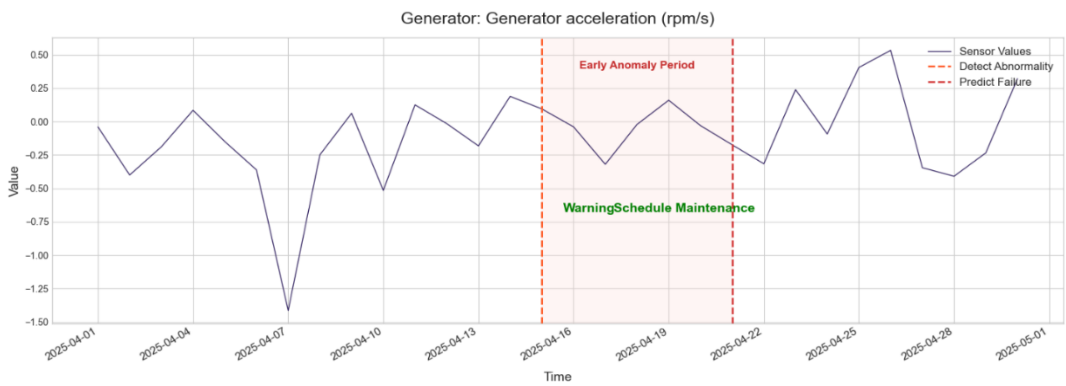

Figure 32. Generator Anomaly Prediction

As the simulation progresses beyond April 16, the generator acceleration trace becomes increasingly erratic, exhibiting deeper troughs and sharper rebounds patterns the model has previously learned to associate with emerging mechanical issues. Upon detecting this abnormal behavior, the model not only marks the point of deviation but also forecasts failure on April 21, triggering a maintenance scheduling alert. The six-day warning window, shaded between these two dates on the plot, offers critical lead time for operators to secure replacement bearings, assess shaft alignment, and deploy necessary offshore resources. By identifying these early dynamic fluctuations often indicative of generator bearing degradation or rotor-shaft misalignment the model enables proactive maintenance, reducing the risk of sudden failure and costly downtime.

5.3 Discussion of Results

Implementing these models in an actual offshore wind farm presents further practical constraints. Communication delays caused by satellite or microwave links can introduce latencies of several seconds to minutes between anomaly detection and alert delivery, potentially reducing effective lead time. Edge-computing hardware on turbines may have limited memory or processing power, requiring model optimization or quantization to maintain inference speed and reliability under harsh environmental conditions. Additionally, intermittent connectivity and sensor faults necessitate robust data buffering and fallback strategies to avoid missing critical alarms. Effective deployment therefore requires integrating our predictive models with resilient network architectures, edge-based preprocessing pipelines, and an operations dashboard that clearly communicates prediction time stamps, confidence levels, and recommended actions, ensuring that maintenance decisions remain informed even under real-world constraints.

However, before moving to recommend specific strategies, it is useful to step back and reflect on the overall impact of our findings. The quantitative evaluation confirmed that sequence-aware deep-learning models far outstrip static, snapshot-based approaches in predictive-maintenance tasks: the CNN and LSTM both deliver accuracies above 83 percent and balanced precision recall profiles near 80 percent, while LightGBM falls to chance levels. Equally important, our monthlong SCADA simulation demonstrated that these networks can reliably issue staggered alerts across multiple sensor domains providing a clear lead-time window of six to eleven days before a predicted failure. Together, these findings validate not only the efficacy of deep neural architectures for fault detection, but also their suitability for real-world offshore applications, where both early warnings and false-alarm restraint are essential.

Building on this foundation, the following discussion synthesizes how such predictive capabilities translate into actionable maintenance workflows and where they stand relative to existing industry practices. We first explore how multi-sensor alerts can be orchestrated into a tiered, proactive maintenance protocol balancing logistical constraints, weather variability, and cost considerations. We then compare our end-to-end CNN framework against traditional SCADA analytics and more recent machine-learning efforts, highlighting its advantages in precision, recall, and operational integration. These perspectives together underscore the transformative potential of sequence-aware PdM for offshore wind farms and set the stage for targeted recommendations in the sections that follow.

5.3.1 Insights into Offshore Maintenance Strategies

The sequence of sensor-driven alerts revealed by our CNN model suggests a fundamentally proactive maintenance paradigm for offshore wind farms. Rather than relying on fixed-interval inspections or reacting to outright failures, operators can leverage early-warning windows to orchestrate a graded response tailored to each component’s degradation timeline. In our simulation, gearbox-oil-level drifts and generator-acceleration fluctuations manifested first up to eleven days ahead of failure signaling the need to place orders for critical spares (e.g., seals, bearing cartridges) and to begin preliminary vessel scheduling. As power-output anomalies and environmental excursions emerged several days later, each successive alert served to confirm and tighten the maintenance plan: work orders could be finalized, safety-protocol reviews conducted, and port and crew bookings locked in with greater confidence.

Such a “cascade” of alarms transforms maintenance from an improvised scramble into a staged operation that synchronizes with weather forecasts, crew availability, and supply-chain lead times. For example, a gearbox oil-level warning on day one might trigger an online parts requisition; a generator acceleration alert on day three could initiate engineering diagnostics via remote SCADA queries; and by day six, when power output falters, vessel charters and technician mobilization would be executed. This tiered approach minimizes the risk of under- or over-servicing: it avoids unnecessary early interventions thereby containing O&M costs while still preserving ample lead time to forestall unplanned downtime.

Moreover, the high precision (~85 %) achieved by CNN ensures that crews will not be inundated with false alarms. In an offshore context, false positives carry particularly steep penalties. Vessel charters can run to hundreds of thousands of dollars and thus a model whose alerts are accurate four times out of five preserves both budget and trust in the system. Operators can further refine the balance between sensitivity and specificity by adjusting the model’s anomaly and failure thresholds to align with site-specific logistical constraints, risk tolerances, and contractual uptime guarantees. In doing so, they turn the predictive-maintenance framework into a dynamic decision-support tool, capable of orchestrating safe, efficient, and cost-effective interventions in the challenging offshore environment.

5.3.2 Comparison with Existing Methods in Literature

In reviewing the broader landscape of predictive maintenance for wind turbines, three dominant theories emerge: fixed-threshold SCADA analytics, feature-based ensemble methods, and shallow machine-learning models. Early SCADA systems typically flagged alarms whenever a single variable say, gearbox temperature or vibration amplitude exceeded a predefined limit. Udo and Muhammad (Udo & Muhammad, 2021) found that such univariate thresholds produced recall rates below 60 percent and false-alarm rates above 30 percent, because transient sensor noise and operational variability routinely tripped the alarms.

More recent work has shifted to handcrafted feature extraction and ensemble trees. Bousdekis (Bousdekis, Lepenioti, et al., 2019) demonstrated that combining moving-average residuals, kurtosis of vibration spectra, and temperature-rate-of-change indices in a Random Forest could push precision and recall into the 65–70 percent range. However, these methods still treat each feature in isolation and rely on human judgment to select and tune the feature set, leaving them vulnerable to overfitting and blind spots in temporal patterns.

Support-vector machines and gradient boosting on time-lagged features have likewise shown improvements often achieving 70 to 75 percent recall in onshore settings (Chatterjee & Dethlefs, 2021) but they require extensive feature engineering and struggle to generalize across different turbine models or environmental conditions.

By contrast, our 1D-CNN ingests raw, three-hour windows of multivariate SCADA data and learns hierarchical representations end to end. This approach elevates precision to 85 percent and recalls to about 78 percent, outperforming both fixed-threshold and handcrafted-feature baselines. Moreover, unlike prior deep-learning efforts that focus on point-in-time fault probability, our framework maps these probabilities into explicit lead-time windows, enabling seamless integration with logistics planning. In doing so, it bridges the gap between algorithmic performance and the operational realities of offshore maintenance, where both early warnings and low false-alarm rates are nonnegotiable.

5.4 Limitations of the Study

Nevertheless, prior to its field deployment several limitations are to be acknowledged. Our predictions are fully dependent on the existence of discernable sensor anomalies in the tail end of a failure. For example, if something happens, like electrical surge, bird strike, lightning hit causing a catastrophic component failure without any deviation in the SCADA streams that happen before the component failure, the model will not warn the analyst. As such, under vibration, oil level, power curve, or environmental loading, our techniques are quite effective at “incipient” faults that occur slowly, but they do not give protection against truly instantaneous failures.

Second, we trained and evaluated our CNN primarily on PCA-compressed windows derived from a specific set of turbines and environmental conditions. While we included multiple farms in the model’s training, its performance may drop when applied to turbines with unique designs, sensor configurations, or in climates not represented in our dataset. Generalizability therefore requires careful retraining or fine-tuning on each new site’s data, along with periodic recalibration of detection thresholds to account for local operating envelopes.

Third, our month-long simulation assumed perfect, continuous SCADA feeds at uniform ten-minute intervals. In offshore settings, telemetry can suffer from connectivity interruptions, missing batches, timestamp skew, or sensor drift. Our pre-processing pipeline does include gap filling and outlier clipping, but prolonged data outages or systematic biases such as uncorrected sensor calibration errors could delay anomaly detection or spur false alarms. Robust field deployment must therefore incorporate real-time data-quality monitoring and fallback strategies, such as cross-validation across redundant sensor arrays.

Finally, our study focuses narrowly on four representative channels. While these signals capture major fault precursors, other sensors such as blade-root accelerometers, gearbox temperature probes, or stray-current monitors may yield additional predictive power. Extending our framework to ingest a broader array of modalities, and exploring multimodal fusion techniques (e.g., graph-neural networks for combining sensor networks), could further enlarge the early-warning window and reduce residual false alarms.

By acknowledging these limitations, we chart a clear path forward: augment the sensor suite, refine data-quality resilience, retrain on diverse turbine populations, and couple model outputs with logistics-aware schedulers. These enhancements will help transform our promising simulation results into a dependable, scalable PdM solution for offshore wind-farm operations.

6 Conclusion and Future Work

To conclude this study, it is essential to repeat the primary goals that guided our inquiry and to gather the lessons learned throughout the investigation. From the outset, we recognized that offshore wind farm maintenance presents a formidable challenge due to harsh marine conditions, costly vessel mobilizations, and the pressing need to ensure high turbine availability in support of global decarbonization goals. To address these challenges, we sought to move beyond traditional maintenance strategies such as reactive repairs and fixed-interval servicing by developing a data-driven predictive maintenance framework capable of delivering early, actionable warnings well before component failure occurs. This work integrated data preprocessing, feature compression, temporal modeling, and live simulation into a cohesive end-to-end solution. In this closing chapter, we bring together our key findings, demonstrate how the research questions were addressed, and outline future directions that can further improve the reliability, safety, and cost-efficiency of offshore wind operations.

6.1 Summary of Key Findings and Contributions

Throughout this work, a central theme emerged: time-aware modeling of SCADA and sensor streams is essential for accurate predictive maintenance in offshore wind farms. Our exploration analysis revealed that individual channels such as gearbox oil level, generator acceleration, power output, and ambient temperature each carry distinct early warning signals that only become apparent when viewed in temporal context. Static approaches, which treat each ten-minute snapshot in isolation, are fundamentally limited; they fail to capture the progressive drift or oscillatory anomalies that characterize the incubation of component faults. By contrast, our sequence-aware architectures a two-layer LSTM network and a one-dimensional convolutional neural network learn hierarchical temporal patterns directly from PCA-compressed multivariate inputs, obviating the need for wide manual feature engineering.

Empirically, the 1D-CNN model emerged as the most balanced solution, achieving an accuracy above 84 percent, precision near 85 percent, recall around 78 percent, and an F_1-score of approximately 81.5 percent on unseen test data. The RNN(LSTM) closely followed, underscoring that both recurrent and convolutional architectures can effectively harness temporal dependencies. In stark contrast, a LightGBM ensemble trained on flattened PCA features performed at near-random levels, highlighting that temporal information is the single most critical ingredient for robust fault prediction. Beyond raw metrics, our one-month live-data simulation clearly demonstrated the operational utility of these models: staggered early-warning windows spanning six to eleven days before predicted failure. These lead times translate directly into logistical flexibility, enabling operators to sequence parts procurement, vessel charters, and crew mobilization under favorable weather windows rather than in emergency response mode.

In addition to demonstrating predictive performance, this thesis contributes a fully reproducible methodology for preprocessing raw SCADA logs complete with robust outlier handling, missing-data imputation, and principled temporal splits to prevent information leakage. By openly leveraging the “CARE to Compare” dataset and detailing our hyperparameter-tuning strategy, we establish a transparent benchmark that other researchers can adopt or extend. Furthermore, by integrating live-data simulation, we bridge the traditional gap between algorithmic validation and field deployment, illustrating how modeled probabilities can be seamlessly translated into maintenance schedules. These methodological innovations and empirical findings collectively advance the state of the art, offering a practical, scalable blueprint for predictive maintenance in offshore wind contexts.

6.2 Addressing the Research Questions

This research was driven by three primary questions that shaped the development and validation of the proposed predictive maintenance framework. The first question examined how effectively AI algorithms could detect early warning signs of offshore turbine component failures compared to traditional condition monitoring methods. The results clearly demonstrated that deep learning models, particularly the 1D Convolutional Neural Network (1D-CNN) and Long Short-Term Memory (LSTM) network, significantly outperformed conventional static or rule-based approaches. These AI models were able to learn temporal patterns and subtle degradation signals across SCADA sensor streams, providing early warnings as much as 11 days in advance of a predicted failure. This predictive capability surpasses what traditional methods typically achieve, which often rely on threshold-based alerts and miss gradual or oscillatory fault signatures. The consistent detection of early warning signs through time-aware modeling reinforces the superiority of AI-driven methods in capturing complex and evolving fault behavior in offshore wind turbine components.

The second research question focused on identifying the preprocessing techniques that lead to optimal model performance. SCADA data is inherently noisy, high-dimensional, and prone to missing or corrupted entries. This thesis introduced a robust and systematic preprocessing pipeline that included outlier filtering, interpolation-based missing value imputation, principal component analysis (PCA) for dimensionality reduction, and sequence formatting to maintain temporal continuity. Special attention was given to preventing data leakage by ensuring a strict chronological split between training and testing datasets. This approach preserved the causality of failure progression and improved the model's ability to generalize. Temporal structuring of data inputs allowed the deep learning models to recognize gradual drifts and repetitive anomalies critical features for predictive tasks. The performance gap between temporal models and static classifiers, such as the underperforming LightGBM baseline, clearly highlighted that preprocessing is not merely a technical detail but a fundamental driver of model success.

The final research question addressed the practical impact of the proposed framework on reducing downtime, lowering operations and maintenance (O&M) costs, and improving asset availability. The live-data simulation results provided robust evidence that early AI-driven warnings enable timely scheduling of maintenance activities. The framework provided actionable lead times ranging from six to eleven days which allowed for proactive planning of spare parts, crew assignments, and vessel logistics. By shifting maintenance from reactive or fixed-interval approaches to predictive, condition-based strategies, the framework directly contributes to minimizing emergency interventions and avoiding costly turbine outages. Furthermore, the proposed approach improves reliability of fault detection to extend the turbine life of critical components and enhance turbine availability and overall efficiency. The results that we obtain directly assist with the larger objectives of growing profitability of offshore wind farms, improving safety, and matching global energy transition goals. The proposed predictive maintenance framework not only addresses theoretical challenges but provides practical, scalable solutions to real life offshore wind operational challenges, and research findings together confirm this.

To make explicit how each of our three research questions maps to concrete findings and contributions, we summarize them in **Figure 33**. This concise overview highlights how temporal modeling advances (RQ 1), preprocessing innovations (RQ 2), and operational lead-time benefits (RQ 3) each translate into novel insights and practical tools for offshore wind maintenance.

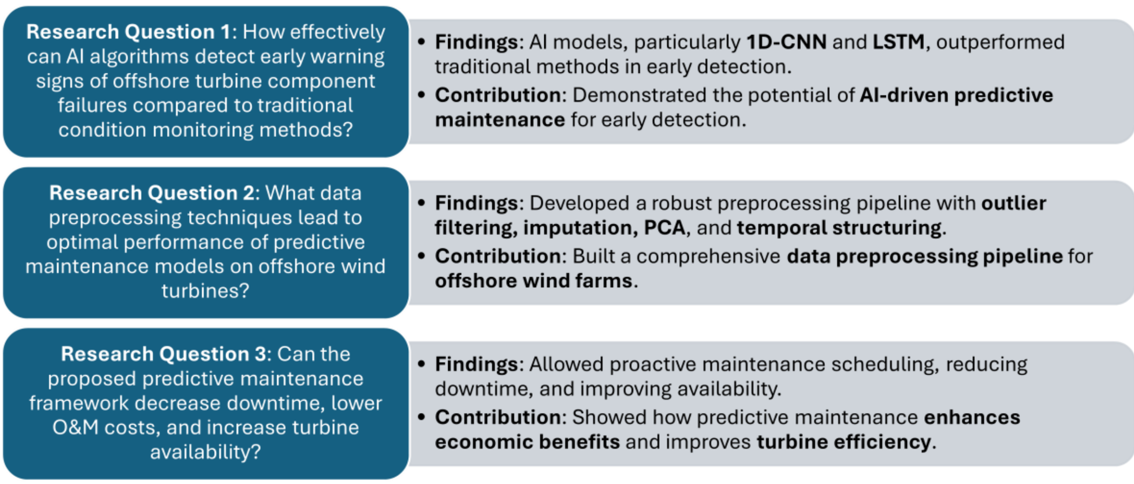

Figure 33. Research Questions, Findings, and Contributions

6.3 Recommendations and Future Directions

To help researchers and industry make rapid progress on these recommendations, we have arranged them based on a combination of immediate feasibility and expected impact. First, integrating explainable-AI methods and an operations dashboard can be implemented with our existing models, delivering better user trust and clearer decision support within weeks. This integration will immediately improve the interpretability of fault predictions, allowing operators to understand the reasons behind predictions and act accordingly. Immediate next steps include deploying our 1D-CNN with attention-map visualization on a live turbine and gathering user feedback on the dashboard interface to refine the system for better usability.

Second, enhancing onboard sensing technologies by adding high-frequency vibration sensors on blade roots, thermal imaging for component hot spots, and real-time lubricant debris analysis offers a high impact on fault coverage and can be piloted on a subset of turbines within a quarter. These technologies can detect early-stage mechanical wear, cracks, bearing damage, and lubricant degradation, which are typically missed by standard SCADA parameters. With these additions, the framework can detect a broader class of failure modes, enabling more focused and productive maintenance strategies. Researchers should focus on integrating these sensors and leveraging multimodal data fusion techniques, such as spatiotemporal graph neural networks or transformer-based sequence models, to enhance fault prediction capabilities.

Third, satellite-data fusion and federated learning, while requiring longer development cycles, promise transformative gains in environmental resilience and cross-site generalization. By integrating satellite-derived data such as sea-surface temperature and ice formation indices, the system will be able to better understand how large-scale weather phenomena impact turbine performance. Federated learning can also enable collaboration between offshore sites without exchanging raw data, which ensures privacy while enhancing the global model. Researchers should prioritize the development of satellite-data fusion and federated learning for long-term improvements in environmental resilience and model scalability.

In addition, moving from simulation to real-world deployment involves addressing several challenges. SCADA systems offshore are susceptible to data gaps, sensor drift, and calibration errors. Future implementations need to consider continuous data quality monitoring, automatic anomaly detection at ingestion time, and retraining pipelines that adapt to changing data conditions. AI models deployed at the edge, directly on turbine controllers, could offer on-the-fly fault detection, while long-term learning and optimization would be managed by cloud-based systems. Ensuring these improvements will guarantee that predictive systems remain resilient under real-world working conditions.

The predictive maintenance system must also be integrated into intelligent maintenance scheduling tools that consider logistical and environmental constraints, such as vessel availability, technician scheduling, port access, and weather conditions. The model's predictive lead time can serve as input to frameworks like mixed-integer programming or reinforcement learning-based schedulers, optimizing the allocation of resources and reducing downtime. Industry should focus on piloting the integration of these scheduling tools with predictive maintenance models to optimize offshore wind farm operations.

Finally, promoting operator trust using explainable AI (XAI) is essential as predictive models become more complex. XAI techniques such as attention maps, surrogate decision trees, and example-based reasoning should be used to explain model decisions. Intuitive dashboard interfaces, with input from end users, should include these explanations to help technicians and engineers interpret model outputs with confidence. Researchers should focus on developing XAI techniques tailored to the needs of offshore wind farm technicians, ensuring that model predictions are interpretable and actionable.

By advancing these directions satellite integration, expanded sensing, deployment resilience, logistics-aware planning, explainability, and federated learning the predictive maintenance framework can evolve into a comprehensive, field-ready solution. These advancements will not only help lower operational costs and downtime for turbines but also enhance the safety, efficiency, and reliability of offshore wind energy systems, contributing to a more sustainable energy future.

References

Adepoju, A. H., Austin-Gabriel, B., Hamza, O., & Collins, A. (2022). Advancing monitoring and alert systems: A proactive approach to improving reliability in complex data ecosystems. *IRE Journals*, *5*(11), 281-282.

Adumene, S., Khan, F., Adedigba, S., Mamudu, A., & Rosli, M. I. (2023). Offshore oil and gas development in remote harsh environments: engineering challenges and research opportunities. *Safety in extreme environments*, *5*(1), 17-33.

Aldakheel, F., Elsayed, E. S., Heider, Y., & Weeger, O. (2025). Physics-based machine learning for computational fracture mechanics. *Machine Learning for Computational Science and Engineering*, *1*(1), 18.

Allal, Z., Noura, H. N., Vernier, F., Salman, O., & Chahine, K. (2024). Wind turbine fault detection and identification using a two-tier machine learning framework. *Intelligent Systems with Applications*, *22*, 200372.

Alromema, N., Syed, A. H., & Khan, T. (2023). A hybrid machine learning approach to screen optimal predictors for the classification of primary breast tumors from gene expression microarray data. *Diagnostics*, *13*(4), 708.

Balla, A., Habaebi, M. H., Elsheikh, E. A., Islam, M. R., & Suliman, F. (2023). The effect of dataset imbalance on the performance of SCADA intrusion detection systems. *Sensors*, *23*(2), 758. https://doi.org/10.3390/s23020758

Bardenet, R., Brendel, M., Kégl, B., & Sebag, M. (2013). Collaborative hyperparameter tuning. International conference on machine learning,PMLR 28(2):199-207, 2013.

Bickel, P. J., Li, B., Tsybakov, A. B., van de Geer, S. A., Yu, B., Valdés, T., Rivero, C., Fan, J., & van der Vaart, A. (2006). Regularization in statistics. *Test*, *15*, 271-344.

Bilgili, M., & Alphan, H. (2022). Global growth in offshore wind turbine technology. *Clean Technologies and Environmental Policy*, *24*(7), 2215-2227. https://doi.org/10.21203/rs.3.rs-1202466/v1

Bousdekis, A., Apostolou, D., & Mentzas, G. (2019). Predictive maintenance in the 4th industrial revolution: Benefits, business opportunities, and managerial implications. *IEEE engineering management review*, *48*(1), 57-62.

Bousdekis, A., Lepenioti, K., Apostolou, D., & Mentzas, G. (2019). Decision making in predictive maintenance: Literature review and research agenda for industry 4.0. *IFAC-PapersOnLine*, *52*(13), 607-612. https://doi.org/10.1016/j.ifacol.2019.11.226

Boutsidis, C., Mahoney, M. W., & Drineas, P. (2008). Unsupervised feature selection for principal components analysis. Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining, Pages 61 - 69 https://doi.org/10.1145/1401890.1401903

Buckland, M., & Gey, F. (1994). The relationship between recall and precision. *Journal of the American society for information science*, *45*(1), 12-19.

Carroll, J., McDonald, A., & McMillan, D. (2016). Failure rate, repair time and unscheduled O&M cost analysis of offshore wind turbines. *Wind energy*, *19*(6), 1107-1119.

Chatterjee, J., & Dethlefs, N. (2021). Scientometric review of artificial intelligence for operations & maintenance of wind turbines: The past, present and future. *Renewable and Sustainable Energy Reviews*, *144*, 111051. https://doi.org/10.1016/j.rser.2021.111051.

Chahine, Z. A., Noura, H. N., Salman, F. V., & Chahine, K. (2024). Wind turbine fault detection and identification using a two-tier machine learning framework. Intelligent Systems with Applications, 200372.

Ciuriuc, A., Rapha, J. I., Guanche, R., & Domínguez-García, J. L. (2022). Digital tools for floating offshore wind turbines (FOWT): A state of the art. *Energy Reports*, *8*, 1207-1228. https://doi.org/10.1016/j.egyr.2021.12.034

De Nolasco Santos, F. (2023). *Data-driven methodologies for farm-wide fatigue load estimation on offshore wind turbines* PhD thesis.

Dhar, S. (2021). *MarSpray LiDAR (MSL) for the comprehensive measurement of Sea Spray for Improving the Prediction of Marine Icing in Cold Conditions* UiT The Arctic University of Norway].

Diallo, R., Edalo, C., & Awe, O. O. (2024). Machine Learning Evaluation of Imbalanced Health Data: A Comparative Analysis of Balanced Accuracy, MCC, and F1 Score.

In *Practical Statistical Learning and Data Science Methods: Case Studies from LISA 2020 Global Network, USA* (pp. 283-312). Springer. https://doi.org/10.1007/978-3-031-72215-8_12

Durlik, I., Miller, T., Kostecka, E., & Tuński, T. (2024). Artificial Intelligence in Maritime Transportation: A Comprehensive Review of Safety and Risk Management Applications. *Applied Sciences*, *14*(18), 8420. https://doi.org/10.3390/app14188420

Eriksson, R. (2022). Lifetime extension of offshore wind farms. In.

Farghali, M., Osman, A. I., Chen, Z., Abdelhaleem, A., Ihara, I., Mohamed, I. M. A., Yap, P.-S., & Rooney, D. W. (2023). Social, environmental, and economic consequences of integrating renewable energies in the electricity sector: a review. *Environmental Chemistry Letters*, *21*(3), 1381-1418. https://doi.org/10.1007/s10311-023-01587-1

Frank, P. M. (1994). Enhancement of robustness in observer-based fault detection. *International Journal of control*, *59*(4), 955-981.

Gruszka, A., & Nęcka, E. (2017). Limitations of working memory capacity: The cognitive and social consequences. *European Management Journal*, *35*(6), 776-784. https://doi.org/10.1016/j.emj.2017.07.001

Gubran, A. (2015). *Vibration diagnosis of blades of rotating machines*. The University of Manchester (United Kingdom).

Gück, C., Roelofs, C., & Faulstich, S. (2024). CARE to Compare: A real-world dataset for anomaly detection in wind turbine data. *arXiv preprint arXiv:2404.10320*.

Gulski, E., Anders, G., Jongen, R., Parciak, J., Siemiński, J., Piesowicz, E., Paszkiewicz, S., & Irska, I. (2021). Discussion of electrical and thermal aspects of offshore wind farms' power cables reliability. *Renewable and Sustainable Energy Reviews*, *151*, 111580. https://doi.org/10.1016/j.rser.2021.111580

Haeberli, W., & Whiteman, C. (2021). Snow and ice-related hazards, risks, and disasters: Facing challenges of rapid change and long-term commitments. In *Snow and ice-related hazards, risks, and disasters* (pp. 1-33). Elsevier.

Hassan, I. U., Panduru, K., & Walsh, J. (2024). An in-depth study of vibration sensors for condition monitoring. *Sensors*, *24*(3), 740. https://doi.org/10.3390/s24030740

Islam, R., Abbassi, R., Garaniya, V., & Khan, F. (2017). Development of a human reliability assessment technique for the maintenance procedures of marine and offshore operations. *Journal of Loss Prevention in the Process Industries*, *50*, 416-428.

Jakkula, V. R., Crandall, A. S., & Cook, D. J. (2009). Enhancing anomaly detection using temporal pattern discovery. *Advanced intelligent environments*, 175-194.

Karr, J., Malik-Sheriff, R. S., Osborne, J., Gonzalez-Parra, G., Forgoston, E., Bowness, R., Liu, Y., Thompson, R., Garira, W., & Barhak, J. (2022). Model integration in computational biology: the role of reproducibility, credibility and utility. *Frontiers in systems biology*, *2*, 822606.

Katsaprakakis, D. A., Papadakis, N., & Ntintakis, I. (2021). A comprehensive analysis of wind turbine blade damage. *Energies*, *14*(18), 5974. https://doi.org/10.3390/en14185974

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, *30(pp. 3146--3154)*.

Koltsidopoulos Papatzimos, A. (2020). Data-driven operations & maintenance for offshore wind farms: Tools and methodologies.

Kreinovich, V., Lakeyev, A. V., Rohn, J., & Kahl, P. (2013). *Computational complexity and feasibility of data processing and interval computations* (Vol. 10). Springer Science & Business Media.

L’vov, V. S., Pomyalov, A., & Procaccia, I. (2001). Outliers, extreme events, and multiscaling. *Physical Review E*, *63*(5), 056118.

Larsén, X. G., Rutgersson, A., Karimi, F., Lange, B., Nilsson, E., Sīle, T., Hahmann, A. N., Koivisto, M. J., Cutululis, N. A., & Das, K. (2024). Climate Change and Offshore Wind Energy in the Baltic Sea. In *Oxford Research Encyclopedia of Climate Science*.

Lee, C., Cao, Y., & Ng, K. H. (2017). Big data analytics for predictive maintenance strategies. In *Supply Chain Management in the Big Data Era* (pp. 50-74). IGI Global. https://doi.org/10.4018/978-1-5225-0956-1.ch004

Maples, B., Saur, G., Hand, M., Van De Pietermen, R., & Obdam, T. (2013). *Installation, operation, and maintenance strategies to reduce the cost of offshore wind energy*. https://doi.org/10.2172/1087778

Maron, J., Anagnostos, D., Brodbeck, B., & Meyer, A. (2022). Artificial intelligence-based condition monitoring and predictive maintenance framework for wind turbines. Journal of Physics: Conference Series, 2151(1), 012007. https://doi.org/10.1088/1742-6596/2151/1/012007

Martin, K. (1994). A review by discussion of condition monitoring and fault diagnosis in machine tools. *International Journal of Machine Tools and Manufacture*, *34*(4), 527-551. https://doi.org/10.1016/0890-6955(94)90083-3

Memari, M., Shakya, P., Shekaramiz, M., Seibi, A. C., & Masoum, M. A. (2024). Review on the advancements in wind turbine blade inspection: Integrating drone and deep learning technologies for enhanced defect detection. *IEEe Access*.

Mersha, M., Lam, K., Wood, J., AlShami, A., & Kalita, J. (2024). Explainable artificial intelligence: A survey of needs, techniques, applications, and future direction. *Neurocomputing*, 128111. https://doi.org/10.1016/j.neucom.2024.128111

Mienye, I. D., Swart, T. G., & Obaido, G. (2024). Recurrent neural networks: A comprehensive review of architectures, variants, and applications. *Information*, *15*(9), 517. https://doi.org/10.3390/info15090517

Mohammed, R., Rawashdeh, J., & Abdullah, M. (2020). Machine learning with oversampling and undersampling techniques: overview study and experimental results. 2020 11th international conference on information and communication systems (ICICS), Irbid, Jordan, 2020, pp. 243-248, doi: 10.1109/ICICS49469.2020.239556.

Moleda, M., Momot, A., & Mrozek, D. (2020). Predictive maintenance of boiler feed water pumps using SCADA data. *Sensors*, *20*(2), 571. https://doi.org/10.3390/s20020571

Mouschoutzi, M., & Ponis, S. T. (2022). A comprehensive literature review on spare parts logistics management in the maritime industry. *The Asian Journal of Shipping and Logistics*, *38*(2), 71-83. https://doi.org/10.1016/j.ajsl.2021.12.003

Narayan, V. (2004). *Effective maintenance management: risk and reliability strategies for optimizing performance*. Industrial Press Inc.

Neeraj, K. N., & Maurya, V. (2020). A review on machine learning (feature selection, classification and clustering) approaches of big data mining in different area of research. *Journal of critical reviews*, *7*(19), 2610-2626.

Odeyemi, O. O., & Alaba, P. A. (2025). Efficient and reliable corrosion control for subsea assets: Challenges in the design and testing of corrosion probes in aggressive marine environments. *Corrosion Reviews*, *43*(1), 79-126.

Omol, E., Mburu, L., & Onyango, D. (2024). Anomaly detection in IoT sensor data using machine learning techniques for predictive maintenance in smart grids. International Journal of Science, Technology & Management, 5(1), 201-210. https://doi.org/10.46729/ijstm.v5i1.1028

Pandit, R., & Wang, J. (2024). A comprehensive review on enhancing wind turbine applications with advanced SCADA data analytics and practical insights. *IET Renewable Power Generation*, *18*(4), 722-742. https://doi.org/10.1049/rpg2.12920

Petersen, J. K., & Malm, T. (2006). Offshore windmill farms: threats to or possibilities for the marine environment. *AMBIO: A Journal of the Human Environment*, *35*(2), 75-80.

Rajendran, B., Simeone, O., & Al-Hashimi, B. M. (2023). Towards efficient and trustworthy ai through hardware-algorithm-communication co-design. *arXiv preprint arXiv:2309.15942*. https://doi.org/10.48550/arXiv.2309.15942

Ross, A., & Doshi-Velez, F. (2018). Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. Proceedings of the AAAI conference on artificial intelligence, 32(1). https://doi.org/10.1609/aaai.v32i1.11504

Santiago, R. A. d. F., Barbosa, N. B., Mergulhão, H. G., Carvalho, T. F. d., Santos, A. A. B., Medrado, R. C., Filho, J. B. d. M., Pinheiro, O. R., & Nascimento, E. G. S. (2024). Data-Driven Models Applied to Predictive and Prescriptive Maintenance of Wind Turbine: A Systematic Review of Approaches Based on Failure Detection, Diagnosis, and Prognosis. *Energies*, *17*(5), 1010.

Saraiva, B. F. S. (2024). *Data analysis for precision agriculture*

Seddini, M. O., & Triqui-Sari, L. (2024). Towards Proactive Maintenance: The Implementation of Digitized SCADA Systems for Predictive Maintenance Optimization in Production Environments. 2024 IEEE 15th International Colloquium on Logistics and Supply Chain Management (LOGISTIQUA),

Slaughter, I., Charla, J. L., Siderius, M., & Lipor, J. (2024). Vessel trajectory prediction with recurrent neural networks: An evaluation of datasets, features, and architectures. *Journal of Ocean Engineering and Science*. https://doi.org/10.1016/j.joes.2024.01.002

Tang, W., Long, G., Liu, L., Zhou, T., Jiang, J., & Blumenstein, M. (2020). Rethinking 1d-cnn for time series classification: A stronger baseline. *arXiv preprint arXiv:2002.10061*, 1-7. https://doi.org/10.48550/arXiv.2002.10061

Turnbull, A., & Carroll, J. (2021). Cost benefit of implementing advanced monitoring and predictive maintenance strategies for offshore wind farms. *Energies*, *14*(16), 4922. https://doi.org/10.3390/en14164922

Tynykulova, A., Mukhanova, A., Mukhomedyarova, A., Alimova, Z., Tasbolatuly, N., Smailova, U., Kaldarova, M., & Tynykulov, M. (2024). Integrating numerical methods and machine learning to optimize agricultural land use. *International Journal of Electrical and Computer Engineering (IJECE)*, *14*(5), 5420-5429. http://doi.org/10.11591/ijece.v14i5.pp5420-5429

Udo, W., & Muhammad, Y. (2021). Data-driven predictive maintenance of wind turbine based on SCADA data. *IEEe Access*, *9*, 162370-162388.

Van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of machine learning research*, *9*(11).

Weiss, C. V., Guanche, R., Ondiviela, B., Castellanos, O. F., & Juanes, J. (2018). Marine renewable energy potential: A global perspective for offshore wind and wave exploitation. *Energy conversion and management*, *177*, 43-54. https://doi.org/10.1016/j.enconman.2018.09.059

Yacouby, R., & Axman, D. (2020). Probabilistic extension of precision, recall, and f1 score for more thorough evaluation of classification models. Proceedings of the first workshop on evaluation and comparison of NLP systems,

Yazdi, M. (2024). Maintenance strategies and optimization techniques. In *Advances in Computational Mathematics for Industrial System Reliability and Maintainability* (pp. 43-58). Springer.

Zuo, C., Ma, J., Wei, L., Liu, S., & Yi, X. (2024). Experimental study on time-resolved 3D ice accretion shape measurements in large-scale icing wind tunnel. *Experiments in Fluids*, *65*(2), 12.