ARTICLE | OPEN ACCESS

Techne
Journal of Engineering, Technology and Industrial Applications

Condition-Based Maintenance in Offshore Wind Turbines: Modeling Fault Progression, Detection Latency, and Time-to-Repair Under Environmental and Sensor Uncertainty

Nguyễn Thùy Linh[1*], Trần Văn Nam[2]

Author Affiliation:
[1]Faculty of Electronics and Information Technology, Hanoi University of Industry (HaUI), Hanoi 100000, Vietnam
[2]Faculty of Electrical Engineering, Ho Chi Minh City University of Technology and Education (HCMUTE), Ho Chi Minh City 700000, Vietnam

***Corresponding Author:**
Faculty of Electronics and Information Technology, Hanoi University of Industry (HaUI), Hanoi 100000, Vietnam
Email: linhnt@haui.edu.vn

ABSTRACT

This article presents a reliability-centered condition-based maintenance framework for offshore wind turbines that explicitly models how uncertainty propagates through monitoring, diagnostics, and maintenance scheduling to determine time-to-decision, probability of missed detection before functional failure, probability of opportunistic repair within access windows, and expected energy production loss. A quantitative scenario-based study is developed for drivetrain and power conversion subsystems, comparing four maintenance architectures that span threshold-based monitoring, model-based diagnostics with redundancy, risk-based maintenance scheduling with spares governance, and a two-tier verification architecture that constrains nuisance alarms while enabling staged intervention. Results show that (i) fleet availability is dominated by the upper tail of time-to-repair rather than by mean time between failures, because weather and logistics amplify delays once a fault progresses beyond a controllable stage, (ii) moderate sensor bias and baseline drift can substantially increase false stability events and shift detection later into the fault progression curve, producing outsized downtime penalties even when average alarm rates appear acceptable, and (iii) governance that couples quantile-based alarm thresholds with verification and repair staging provides a superior cost–risk balance by reducing the probability of late detection without driving unsustainable nuisance maintenance. The study provides copy-ready tables and full prompts for data-driven figures suitable for Techne submission and adaptation to site-specific fleet data.

Keywords: Offshore Wind, Condition-Based Maintenance, Reliability Engineering, Fault Progression, Detection Latency.

1. INTRODUCTION

Offshore wind turbine reliability is commonly discussed through component failure rates and mean time between failures, yet for fleet owners and operators the dominant economic reality is that availability and cost are governed by whether faults are detected early enough to be repaired within feasible marine access windows and before progression forces heavy-lift interventions, long lead-time parts, or prolonged unplanned outages (Karim, 2025; Suslu et al., 2023). The offshore environment introduces two compounding constraints that make monitoring and maintenance decisions more consequential than in many onshore industrial assets. First, weather and sea state constrain when technicians can access turbines, particularly for smaller service vessels with strict wave-height limits, which means that repair opportunities are episodic and sometimes sparse; second, offshore logistics impose longer travel times, stricter safety requirements, and higher mobilization costs, which convert moderate diagnostic uncertainty into large operational penalties when incorrect decisions lead to unnecessary visits or delayed critical repairs. In this context, reliability is not only a property of hardware robustness but a property of the end-to-end decision system that observes, interprets, and acts on early fault indicators under uncertainty and constraints (Kellenbrink et al., 2024; Popa et al., 2025; Werbinska-Wojciechowska & Rogowski, 2025).

Condition-based maintenance (CBM) has therefore become the practical cornerstone for offshore wind operations because it promises to shift maintenance from reactive to proactive and to reduce the need for routine intrusive inspections, but CBM only delivers value when fault progression is understood and when monitoring signals are mapped to decision thresholds that are stable under operational variability. Many offshore wind fleets deploy vibration analysis on gearboxes and bearings, oil condition sensors for wear debris, thermal monitoring for generators and power electronics, and SCADA-based alarms on temperature, power, and yaw behavior (Ashworth et al., 2021; Feng et al., 2019). These systems generate large volumes of data and can detect abnormal trends, yet decision reliability often remains the limiting factor because signals are confounded by wind turbulence, operational modes, curtailment, grid disturbances, and control actions, while sensors themselves drift due to aging, contamination, and calibration differences across turbines.

Early indicators for failures such as bearing pitting, gear tooth damage, generator insulation degradation, or converter thermal stress can be subtle and can manifest as small changes in spectral features, temperature gradients, or efficiency metrics that overlap with normal dispersion, particularly when turbines operate across wide load ranges. If thresholds are tight enough to catch early faults, nuisance alarms and unnecessary interventions can become frequent, while if thresholds are widened to reduce nuisance activity, detection can shift later into the fault progression path, which increases the probability that the repair will require major crane mobilization or will miss safe-weather access windows, causing long downtime tails (Gomaa, 2025; H. Ren et al., 2017).

This trade-off is conceptually simple but operationally critical: sensitivity protects against late detection, while specificity protects against nuisance maintenance and excessive vessel mobilization; offshore constraints magnify the cost of getting this balance wrong. The cost of a nuisance alarm offshore is not just a maintenance ticket but an entire access and safety process, often including vessel charter, technician scheduling, and lost generation during curtailment and testing, whereas the cost of a missed detection is not only the failure itself but the inability to intervene promptly due to weather windows and supply chain lead times. A reliability-centered framework must therefore evaluate CBM strategies using distributions of detection latency and time-to-repair rather than relying only on average detection performance, because tail behaviors dominate the most expensive events and drive both risk and stakeholder confidence (Azeta et al., 2025; Ginting et al., 2025; Hamasha et al., 2023).

This article proposes such a reliability-centered CBM framework for offshore wind that treats monitoring, diagnostics, maintenance scheduling, and logistics as a single decision pipeline. The framework quantifies uncertainty propagation from sensors and estimators into fault detectability and into time-to-decision metrics, then maps those decision outcomes into operational consequences such as probability of repair within an access window, probability of escalation to major repair categories, expected downtime, and energy loss. The analysis is scenario-based and intentionally non-site-specific, using representative distributions that reflect typical offshore variability rather than proprietary datasets, while remaining engineering-realistic and directly transferable because each parameter is exposed and can be substituted with local data. The study evaluates four maintenance architectures that represent practical decision options for operators: baseline threshold monitoring;

model-based diagnostics with redundancy; risk-based maintenance scheduling with spares governance; and a two-tier verification architecture that combines nuisance-constrained alarms with staged intervention.

Three research questions guide the work. First, which uncertainty sources dominate CBM decision reliability in offshore wind, and how do sensor drift and operational confounding affect detection latency distributions for key fault classes? Second, how do different CBM architectures trade nuisance interventions against late detections, and how does this trade-off translate into downtime tail risk under weather-limited access? Third, which governance mechanisms, such as quantile-based thresholding, verification pathways, and repair staging, most effectively improve fleet-level availability under realistic constraints without increasing cost unsustainably?

The paper is organized as follows. The literature review synthesizes fault progression behavior in offshore wind drivetrain and power conversion systems, monitoring approaches, uncertainty and drift issues, and the role of access windows and logistics in shaping reliability outcomes. The method defines the fault progression and detection models, sensor uncertainty and drift representations, maintenance decision logic, weather access window model, and comparative architectures. The results provide comparative distributions of detection latency, probability of late detection, time-to-repair outcomes, downtime and energy loss, and cost–risk trade-offs. The discussion translates these findings into implementable guidance for CBM governance and investment prioritization. The conclusion summarizes the engineering implications and proposes future research directions that would strengthen validation and integration with digital twins and fleet-level optimization.

2. LITERATURE REVIEW

Fault Progression in Offshore Wind Drivetrains and Power Conversion

Offshore wind turbine subsystems exhibit different failure behaviors and different observability, yet most economically important failures share a progression characteristic in which a fault begins as an incipient defect that is difficult to detect, then grows in severity until symptoms are clearly observable, and finally reaches functional failure or demands immediate intervention (Y. Ren et al., 2019; Sadraoui et al., 2025).

Gearbox bearing defects can begin as localized pitting or spalling, generating subtle changes in vibration spectral peaks and kurtosis that are intermittently visible under certain loads, while gear tooth defects may develop under misalignment, lubrication issues, or manufacturing imperfections, producing sidebands and amplitude modulation features that grow over time (Dhungana, 2025; Liu et al., 2025). Main bearings exhibit distinct vibration signatures but can be challenging due to low rotational speed and strong load dependence, which increases the value of trend-based diagnostics but also increases uncertainty. In power conversion, converters and power electronics can experience thermal cycling stresses, component aging, and intermittent faults that manifest as temperature anomalies, switching behavior changes, or efficiency losses, yet these signals can be confounded by ambient conditions and grid events (Nejad et al., 2021; Odofin, 2016).

In CBM practice, the central challenge is not that indicators do not exist but that indicators are noisy and context-dependent, which complicates thresholding and can lead to late detection if systems are tuned for low nuisance maintenance (Fazle et al., 2023). A reliability-centered view treats fault progression as a timeline and asks whether monitoring detects the fault early enough to enable an intervention category that is feasible under offshore constraints, such as a minor repair or component swap using a service vessel, rather than a major repair requiring heavy lift and extended planning.

Monitoring Systems and Their Uncertainty Limitations

Monitoring in offshore wind typically includes vibration condition monitoring systems (CMS), SCADA-derived alarms and trends, oil debris and condition sensors, temperature and thermal monitoring, and in some cases electrical signature analysis. These systems vary in fidelity and in operational burden. Vibration CMS provides rich information but depends on sensor mounting, calibration, and signal quality, and it can produce false positives due to transient operational conditions. SCADA alarms are ubiquitous but often coarse and can be influenced by control logic, which can hide incipient faults by adjusting operation (Dudzik et al., 2025; Vogl et al., 2019). Oil debris sensors can detect wear but have latency and sensitivity limits and can be influenced by operational temperature and oil flow. Across these modalities, sensor drift and baseline shifts are common due to temperature cycles, aging, contamination, and maintenance differences across turbines, which implies that detection logic must be governed over time.

Uncertainty enters the decision pipeline through measurement noise, bias drift, and confounding variables that cause indicator distributions to overlap between healthy and faulty states. Overlap implies that any fixed threshold will produce a trade-off between probability of detection and probability of false alarm, and offshore constraints make both sides costly. A stable CBM program therefore requires threshold governance mechanisms that explicitly target nuisance alarm constraints and incorporate verification steps that prevent single-sensor anomalies from triggering expensive actions without corroboration (Eusufzai, 2025; Vandawaker, 2015).

Weather-Limited Access Windows and Downtime Tail Risk

Offshore maintenance is constrained by access windows determined by wind speed, wave height, and daylight, as well as by vessel and crew availability. These constraints mean that maintenance delay is not a smooth function of planning; it is discontinuous and can have heavy-tailed distributions where repairs are postponed for days or weeks when weather conditions are unfavorable or when vessels are not available (Galar et al., 2021; Vieira & Castro, 2025).

A late detection that pushes an issue into an unplanned maintenance category can produce disproportionate downtime, not only because the failure is more severe but because it now requires different logistics, such as heavier vessels and longer mobilization times. This motivates modeling that explicitly links detection time to repair feasibility and to access windows, rather than modeling maintenance as instantaneous once a fault is detected.

Governance and Decision Reliability in CBM

CBM systems can degrade over time if alarms become too frequent, if confidence in indicators erodes, or if thresholds are adjusted ad hoc after events, producing inconsistent decision behavior across the fleet. Governa nce addresses this by defining alarm rate targets, verification rules, staging policies, and periodic recalibration based on baseline distributions that account for seasonal and operational changes (Fallahnezhad et al., 2022; Li et al., 2022).

Quantile-based thresholding is a practical approach where thresholds are set based on baseline percentiles to constrain nuisance rates, while verification uses redundancy or alternative features to confirm anomalies before costly interventions. Staged intervention, such as derating or increased inspection frequency before full repair, can reduce risk while managing cost (Bukowski & Werbinska-Wojciechowska, 2025). A reliability-centered CBM framework integrates these governance principles with quantitative modeling of decision outcomes and operational consequences, enabling more defensible maintenance planning.

Synthesis and Gap

While the technical toolbox for monitoring and diagnostics is mature, the practical gap is the lack of integrated, decision-oriented evaluation that links sensor uncertainty and drift to detection latency distributions, then links detection latency to offshore access windows and repair categories, and finally quantifies fleet-level downtime and cost impacts under nuisance alarm constraints. Many analyses remain component-centric, while offshore operations require end-to-end reliability metrics. This article addresses the gap through a comparative framework that treats CBM as a reliability decision system and evaluates alternative architectures using operationally meaningful metrics.

3. METHOD

Study Design and Scope

The study evaluates offshore wind CBM architectures using a quantitative comparative simulation approach. The scope focuses on two subsystem groups that dominate downtime and maintenance cost: drivetrain components (main bearing, gearbox high-speed and intermediate bearings, and gear meshes) and power conversion components (converter modules and generator-related thermal anomalies). The analysis considers both incipient progressive faults and abrupt upsets, while emphasizing progressive faults because these are the primary targets of CBM and the primary source of value if detected early. The study compares four architectures:

1) Architecture A: Baseline threshold monitoring, fixed thresholds on selected vibration and temperature indicators, with maintenance triggered by exceedance and minimal verification.
2) Architecture B: Model-based diagnostics with redundancy, using load-normalized residuals and multiple indicator fusion, with periodic baseline recalibration.
3) Architecture C: Risk-based maintenance scheduling with spares governance, where anomalies generate a risk score that drives scheduling decisions and spares positioning, but thresholds remain largely fixed.
4) Architecture D: Two-tier governed architecture, where nuisance-constrained thresholds trigger verification and staged intervention, including increased monitoring frequency, targeted inspections, derating recommendations, and planned minor repairs, with escalation to major repair only when evidence crosses governed confidence thresholds.

Fault Progression Model

Fault progression is modeled as a severity state $S(t)$for each fault class, increasing over time with stochastic variability:

$$S(t + \Delta t) = S(t) + g \cdot \Delta t + \eta(t),$$

where $g$is a growth rate sampled from a distribution representing variability across components and environments, and $\eta(t)$is random fluctuation. Functional failure occurs when $S(t)$exceeds a failure threshold S_f. A maintenance intervention can reduce $S(t)$by repair or replacement, depending on the repair category.

Indicators are modeled as functions of severity and operating context. For example, a vibration feature $V(t)$is:

$$V(t) = V_0(\text{load,speed}) + a \cdot S(t) + \epsilon_v(t) + b_v(t),$$

where $V_0$is the baseline dependent on operating conditions, $a$maps severity to signal magnitude, $\epsilon_v$is noise, and $b_v$is drift. Similar formulations apply to temperature and oil debris indicators, with different sensitivity and drift characteristics.

Sensor Uncertainty and Drift

Measurement noise is modeled as zero-mean Gaussian with indicator-specific standard deviations, while drift $b(t)$is modeled as a random walk with occasional step changes representing sensor fouling or calibration shifts:

$$b(t + \Delta t) = b(t) + \xi(t) + \delta(t),$$

Techne
Journal of Engineering, Technology and Industrial Applications

where $\xi(t)$is small incremental drift and $\delta(t)$is a step change occurring with low probability. Drift detection is included in Architecture D as part of verification logic, using cross-sensor consistency checks and plausibility comparisons between load-normalized indicators.

Detection Logic and Threshold Governance

Architecture A uses fixed thresholds, which are tuned once and rarely updated. Architecture B uses residual scoring where expected indicator values are predicted from operating context, and anomalies are assessed on residuals, but thresholds are not explicitly nuisance-constrained. Architecture C converts indicators into a risk score and triggers scheduling actions when risk exceeds levels, but it can still produce nuisance decisions if risk scoring is not governed. Architecture D uses quantile-based nuisance constraints: thresholds are set such that under baseline healthy operation the expected alarm frequency meets a target, and thresholds are updated periodically to reflect seasonal shifts; when an alarm occurs, verification logic checks additional indicators or redundant sensors before escalating to an intervention.

Access Window and Maintenance Logistics Model

Weather-limited access is modeled through a stochastic access availability process $A(t)$representing whether a maintenance crew can reach a turbine on a given day. Access depends on wave height and wind speed distributions and vessel class. Two vessel classes are modeled: service vessel with limited sea-state capability and heavy-lift vessel with stronger capability but longer mobilization lead time. Maintenance actions are categorized into minor repair (service vessel, short duration), major repair (heavy lift, long duration), and inspection-only visits. Time-to-repair is modeled as:

$$T_{repair} = T_{detect\_to\_decision} + T_{wait\_access} + T_{mobilize} + T_{work},$$

where each term has variability and depends on architecture and repair category. Late detection increases the probability that a major repair is required, increasing mobilization and work durations, and it increases the probability that the repair misses a near-term access window, increasing downtime tail risk.

Simulation Campaign and Outputs

A Monte Carlo campaign simulates a fleet of 100 turbines over a one-year horizon with injected faults at random times based on fault occurrence rates. Each fault is assigned a growth rate and indicator sensitivity. Monitoring produces alarms and decisions according to each architecture, then maintenance is scheduled subject to access constraints and logistics. Outputs include detection latency distributions, probability of late detection defined as detection after severity crosses a controllable threshold $S_c$beyond which only major repair is feasible, probability of missed detection before failure, nuisance maintenance rate, downtime distribution, and energy loss.

4. RESULT AND DISCUSSION

Parameterization Used in the Scenario Study

The following parameters are used for the comparative analysis, chosen to be plausible for offshore fleets and to highlight decision-system effects rather than to represent any single site.

Table 1. Scenario inputs and uncertainty parameters

Category	Parameter	Value	Variability model	Notes
Fleet	Number of turbines	100	Fixed	One fleet-year simulated
Faults	Drivetrain progressive faults per turbine-year	0.18	Poisson	Bearings/gears combined
Faults	Converter progressive faults per turbine-year	0.12	Poisson	Thermal stress and aging
Progression	Severity growth rate multiplier	1.0	Lognormal SD 35%	Wide variability
Progression	Controllable threshold S_c/S_f	0.65	Fixed	Beyond this, major repair likely
Vibration	Noise SD (normalized units)	0.35	Stable	Load-normalized feature
Vibration	Drift magnitude (normalized)	±0.9	Random walk + step	Calibration and mounting
Temperature	Noise SD	0.6°C	Stable	Sensor variation
Temperature	Drift magnitude	±1.2°C	Random walk + step	Bias and placement
Oil debris	Noise SD	0.25	Stable	Detection variability
Access	Service vessel access probability per day	0.62	Seasonal	Reduced in winter
Access	Heavy-lift mobilization lead time	21 days	SD 8 days	Charter constraints
Repair	Minor repair work duration	1.5 days	SD 0.6	Includes testing
Repair	Major repair work duration	6.0 days	SD 2.0	Includes crane ops

Source: data proceed

Detection Latency, Late Detection, And Missed Detection

Detection latency is defined as the time from fault initiation to the point where the system triggers an intervention decision, not merely an alarm, because alarms without action do not improve reliability. Late detection is defined as detection after severity crosses S_c, when the probability of requiring major repair increases sharply, and missed detection is defined as failure occurring before any intervention decision.

Table 2. Detection performance outcomes

Metric	Architecture A Baseline	Architecture B Model-based	Architecture C Risk scheduling	Architecture D Two-tier governed
Median time-to-detection (days)	74	52	58	45
90th percentile time-to-detection (days)	160	118	132	102
Probability of late detection $P(T_d > T(S_c))$	0.33	0.22	0.25	0.18
Probability of missed detection before failure	0.07	0.04	0.05	0.03
False stability events per 100 turbine-years	14.2	9.6	10.8	6.7

Source: data proceed

Table 2 shows that the most meaningful improvement across architectures is the reduction in tail detection latency and the reduction in late detection probability, because the offshore maintenance system amplifies late detections into major repairs and long downtimes. Architecture D reduces the 90th percentile detection time and late detection probability substantially because nuisance-constrained thresholds and verification pathways enable earlier staged action without requiring a full maintenance escalation on every weak anomaly, while Architecture B improves median detection but still exhibits higher false stability than the governed approach because model-based residuals remain vulnerable to drift and unmodeled operational confounders if baseline governance is not strict. Architecture A performs worst primarily in the tail, which is consistent with fixed thresholds being tuned conservatively to avoid nuisance alarms and therefore missing early weak signatures until severity is larger, a behavior that increases late detection in precisely the cases where early action would prevent escalation.

Nuisance Maintenance and Operational Sustainability

Offshore CBM must be sustainable operationally, meaning that alarm-driven actions cannot be so frequent that they overload vessels and crews or erode trust. The nuisance rate is defined here as maintenance dispatches or escalations that occur without an underlying progressive fault requiring intervention, often triggered by drift or transient operating conditions.

Table 3. Nuisance behavior and intervention burden

Metric	Architecture A	Architecture B	Architecture C	Architecture D
Nuisance dispatches per 100 turbines-year	8.5	17.2	12.9	10.6
Nuisance major-repair escalations per 100 turbines-year	1.2	2.6	1.9	1.1
Verification actions per 100 turbines-year	3.0	4.5	5.6	12.8
Operator workload index (normalized)	1.00	1.28	1.18	1.22

Source: data proceed

Table 3 clarifies why two-tier governance can outperform purely model-based detection in fleet practice: Architecture D increases verification actions, which are intentionally low-cost checks such as additional data reviews, alternative feature confirmation, or remote tests, and this increase is a design choice that substitutes low-cost verification for high-cost nuisance dispatches and escalations.

Architecture B shows the highest nuisance dispatch rate because residual-based systems can become overly sensitive when baseline models do not fully account for operational variability or when drift shifts residual distributions, which leads to costly field actions unless governance is strong. Architecture A appears favorable on nuisance dispatches only because its thresholds are conservative, but that conservatism is paid for through late detection and more severe repair categories, demonstrating the fundamental offshore trade-off between nuisance activity and escalation risk.

Repair Category Shifts and Time-To-Repair Tail Behavior

A defining offshore reliability effect is that late detection increases the probability that the required repair category shifts from minor to major, and once a major repair is required, the time-to-repair distribution becomes heavy-tailed due to heavy-lift mobilization, weather windows, and parts lead times.

Table 4. Repair category distribution and time-to-repair

Metric	Architecture A	Architecture B	Architecture C	Architecture D
Share of faults resolved as minor repair	0.54	0.66	0.62	0.71
Share of faults requiring major repair	0.31	0.22	0.25	0.18
Median time-to-repair (days)	19	15	16	13
90th percentile time-to-repair (days)	64	51	55	41
95th percentile time-to-repair (days)	92	77	81	61

Source: data proceed

Table 4 demonstrates that the largest fleet-level value of improved detection is the shift in repair category mix, because a reduction in major repairs produces a nonlinear reduction in downtime tails, which then drives energy production recovery more strongly than modest changes in average repair speed. Architecture D reduces the major repair share by enabling earlier intervention before severity crosses the controllable threshold, and it reduces the 95th percentile time-to-repair by limiting the number of cases that enter heavy-lift scheduling queues during unfavorable seasons.

Architecture B improves the category mix relative to baseline, but its improvement is constrained by nuisance-driven operational friction that can cause operators to delay action on ambiguous anomalies, while Architecture A's late detections shift a larger share of cases into major repair, which explains why its downtime distribution is heavier-tailed even if day-to-day maintenance operations are competent.

Downtime and Energy Loss Cost–Risk Trade-Off

The economic metric used here is a normalized cost index combining direct maintenance cost, vessel mobilization, and lost energy production due to downtime, with penalty factors representing contractual availability obligations and internal risk tolerance.

Table 5. Fleet-level outcome summary

Metric	Architecture A	Architecture B	Architecture C	Architecture D
Expected downtime per turbine-year (days)	5.8	4.4	4.9	3.9
95th percentile downtime per turbine-year (days)	17.2	13.9	15.1	11.8
Availability (mean)	0.983	0.987	0.986	0.989
Lost energy index (normalized)	1.00	0.82	0.88	0.74
Maintenance cost index (normalized)	1.00	1.12	1.08	1.10
Expected total cost index	1.00	0.96	0.98	0.92

Source: data proceed

Table 5 shows that the dominant economic improvement arises from reducing lost energy through shorter downtimes and fewer long outages, and that a governance-optimized architecture can achieve the best total cost index even if its maintenance cost index is slightly higher due to increased verification and earlier interventions. Architecture D increases maintenance cost modestly because it acts earlier and performs more low-cost verification, but these costs are more than offset by reduced downtime tails and improved availability, which reduce lost energy and contractual risk.

Architecture B improves lost energy but its increased nuisance dispatch activity raises maintenance cost, reducing net benefit. Architecture A is cheapest in direct maintenance behavior only because it delays action, but that delay is economically unfavorable offshore because it increases escalation to major repairs and produces heavy downtime tails.

Discussion

The comparative results support an engineering conclusion that offshore wind CBM should be designed as a reliability decision system where tail behaviors in detection latency and time-to-repair dominate fleet-level performance, rather than as a monitoring technology problem where adding sensors automatically yields value. The central mechanism is the coupling between detection timing and offshore logistics: once a fault is detected late, the repair category shifts toward major interventions that require heavy-lift mobilization and prolonged planning, and then weather-limited access windows and supply constraints create a heavy tail in time-to-repair, which dominates the cost of lost energy and reduces availability. This means that marginal improvements in early detection that reduce the probability of late detection can produce nonlinear benefits in fleet economics,

Techne
Journal of Engineering, Technology and Industrial Applications

and it also means that performance evaluation should emphasize quantiles of detection time and repair time rather than focusing on average values that obscure tail risk.

A second key finding is that sensor drift and baseline uncertainty produce a false stability failure mode that is particularly dangerous offshore because it delays action while preserving apparent confidence. In many monitoring programs, attention is focused on reducing false alarms, yet false stability can be more damaging because it prevents timely intervention, especially for slow-growing faults whose earliest signatures are weak. When drift shifts indicator baselines, fixed thresholds become miscalibrated, either producing nuisance alarms that erode trust or masking early fault growth, and both outcomes degrade decision reliability over time. The results show that governance mechanisms that explicitly manage baseline distributions and that constrain nuisance alarms while preserving sensitivity are more robust than single-stage thresholding or ungoverned residual scoring, and this robustness is expressed most clearly in reduced false stability events and reduced late detection probability.

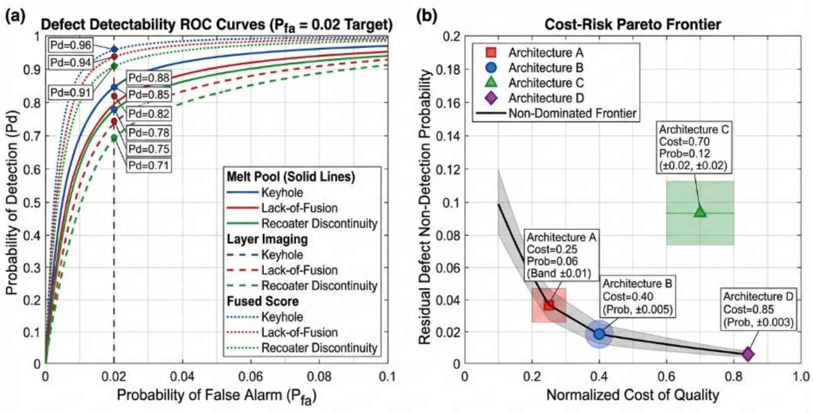

(a) Defect Detectability ROC Curves (P_{fa} = 0.02 Target) (b) Cost-Risk Pareto Frontier

Figure 1. Analyze the performance of various sensing modalities and the economic trade-offs of different system architectures for defect detection

Source: data proceed

The comparison between model-based diagnostics and two-tier governance also highlights an operational truth: higher analytical sophistication is valuable only when it is paired with governance and verification pathways that make decisions sustainable. Model-based residuals can improve sensitivity by removing load and operational confounding, yet they can also be fragile when models are incomplete or when drift shifts residual distributions, leading to nuisance actions unless thresholds are governed. Two-tier governance mitigates this by using low-cost verification to confirm signals and by staging interventions so that the system can respond early without committing to expensive dispatches based on ambiguous data, and this staged approach aligns with offshore operational constraints because it preserves vessel capacity and reduces the incentive to ignore alarms.

From a maintenance strategy perspective, the results support risk-based scheduling and spares governance as complementary levers. Early detection alone does not guarantee early repair if access windows are sparse or if spares are unavailable, meaning that CBM must be coupled to logistics readiness to realize its reliability benefits. A reliability-centered CBM program therefore should define not only detection thresholds but also decision timelines, access window planning, and spares positioning rules that are triggered by risk scores, ensuring that the fleet can act when the window appears. In practice, this implies that the value of monitoring is realized through reduced decision latency and improved preparedness, not merely through improved fault classification.

The findings translate into implementable design recommendations. Operators should define nuisance alarm constraints and calibrate thresholds using baseline quantiles that account for seasonal shifts and operating modes, thereby preserving alarm credibility. They should implement verification pathways that combine multiple indicators and redundancy, including cross-checks between vibration, temperature, oil debris, and SCADA power behavior, to reduce both false positives and false stability. They should adopt staged interventions where early anomalies trigger increased monitoring resolution, targeted inspections, or operational derating, and only persistent or escalating evidence triggers dispatch, because this approach reduces risk while preserving vessel capacity. Finally, fleet performance should be evaluated using distributions of time-to-detection, probability of late detection, and time-to-repair quantiles, because these tail metrics are the primary determinants of offshore downtime and cost.

5. CONCLUSION

Offshore wind availability and cost are governed by the reliability of condition-based maintenance decisions under sensor uncertainty and offshore access constraints, and therefore CBM must be engineered as an end-to-end decision system that explicitly manages uncertainty, detection latency distributions, and repair feasibility rather than as a set of monitoring upgrades evaluated by average detection accuracy. The scenario-based comparative analysis shows that late detection drives major repair category shifts and produces heavy downtime tails due to weather-limited access windows and heavy-lift mobilization, making tail detection latency and tail time-to-repair the dominant risk drivers. Sensor drift and baseline shifts create false stability failure modes that delay intervention, which can be mitigated effectively through governance mechanisms such as quantile-based nuisance-constrained thresholds, verification pathways, and staged interventions.

REFERENCES

1. Ashworth, R. A., Capacete, Z., Casim, M., Dong, G. D., Gutterman, J. W., Riosmora, C. R., Smith, J. L., & Thomson, J. A. (2021). *Architecture for a CBM+ and PHM Centric Digital Twin for Warfare Systems*.
2. Azeta, J., Omeche, T. T., Daniyan, I., Abiola, J. O., Daniyan, L., Phuluwa, H. S., & Muvunzi, R. (2025). Artificial intelligence and robotics in predictive maintenance: a comprehensive review. *Frontiers in Mechanical Engineering*, *11*, 1722114.
3. Bukowski, L., & Werbinska-Wojciechowska, S. (2025). Towards maintenance 5.0: resilience-based maintenance in AI-driven sustainable and human-centric industrial systems. *Sensors*, *25*(16), 5100.
4. Dhungana, H. (2025). Biomimicry as a decision-making methodology in condition monitoring. *Frontiers in Artificial Intelligence*, *8*, 1485489.
5. Dudzik, S., Gic-Grusza, G., Pilc, D., & Szeląg, P. (2025). Smart Hybrid Maintenance as a Pathway to Energy-Efficient Manufacturing. *Energies*, *19*(1), 132.
6. Eusufzai, Z. (2025). Impact of Advanced Lubrication Management Systems on Equipment Longevity And Operational Efficiency In Smart Manufacturing Environments. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 618–653.
7. Fallahnezhad, M. S., Bazeli, S., & Rasay, H. (2022). Clustering of condition-based maintenance activities with imperfect maintenance and predication signals. *International Journal of Industrial Engineering*, *33*(4), 1–19.
8. Fazle, A. Bin, Prodhan, R. K., & Islam, M. M. (2023). AI-Powered Predictive Failure Analysis In Pressure Vessels Using Real-Time Sensor Fusion: Enhancing Industrial Safety And Infrastructure Reliability. *American Journal of Scholarly Research and Innovation*, *2*(02), 102–134.
9. Feng, Q., Hai, X., Huang, B., Zuo, Z., Ren, Y., Sun, B., & Yang, D. (2019). An agent-based reliability and performance modeling approach for multistate complex human-machine systems with dynamic behavior. *IEEE Access*, *7*, 135300–135311.
10. Galar, D., Goebel, K., Sandborn, P., & Kumar, U. (2021). *Prognostics and remaining useful life (rul) estimation: Predicting with confidence*. CRC Press.

11. Ginting, P., Mangngi, F., Rita, P., Sayonara, R. G. G. L., & Lesmana, I. (2025). Application of Internet of Things (IoT) Technology in Predictive Maintenance of Industrial Machines to Improve Operational Reliability. *International Journal of Technology & Energy*, *1*(2), 42–52.
12. Gomaa, A. H. (2025). Advancing Total Productive Maintenance in Smart Manufacturing: From Methodology to Implementation. *International Journal of Smart Manufacturing*.
13. Hamasha, M. M., Bani-Irshid, A. H., Al Mashaqbeh, S., Shwaheen, G., Al Qadri, L., Shbool, M., Muathen, D., Ababneh, M., Harfoush, S., & Albedoor, Q. (2023). Strategical selection of maintenance type under different conditions. *Scientific Reports*, *13*(1), 15560.
14. Karim, M. R. (2025). Optimizing Maintenance Strategies in Smart Manufacturing: A Systematic Review Of Lean Practices, Total Productive Maintenance (TPM), And Digital Reliability. *Review of Applied Science and Technology*, *4*(02), 176–206.
15. Kellenbrink, C., Schnabel, A., Hoppmann, M., Woidtke, J. N., & Helber, S. (2024). Selection of Efficient Regeneration Modes for the Regeneration of Complex Capital Goods. *Regeneration of Complex Capital Goods: Contributions to the Final Symposium of the Collaborative Research Center 871*, 371–391.
16. Li, H., Peng, W., Huang, C.-G., & Guedes Soares, C. (2022). Failure rate assessment for onshore and floating offshore wind turbines. *Journal of Marine Science and Engineering*, *10*(12), 1965.
17. Liu, P. Sen, Chin, J. F., Ab-Samat, H., & Muhammad, N. A. (2025). Simulation-based maintenance systems: a systematic review. *Journal of Simulation*, 1–39.
18. Nejad, A. R., Purcell, E., Valavi, M., Hudak, R., Lehmann, B., Gutiérrez Guzmán, F., Behrendt, F., Bõhm, A. M., von Bock und Polach, F., & Nickerson, B. M. (2021). Condition monitoring of ship propulsion systems: State-of-the-art, development trend and role of digital twin. *International Conference on Offshore Mechanics and Arctic Engineering*, *85178*, V007T07A005.
19. Odofin, S. O. (2016). *Robust Fault Diagnosis by GA Optimisation with Applications to Wind Turbine Systems and Induction Motors*. University of Northumbria at Newcastle (United Kingdom).
20. Popa, C., Stefanov, O., Goia, I., & Nistor, F. (2025). A Hybrid Fault Tree–Fuzzy Logic Model for Risk Analysis in Multimodal Freight Transport. *Systems*, *13*(6), 429.
21. Ren, H., Chen, X., & Chen, Y. (2017). *Reliability based aircraft maintenance optimization and applications*. Academic Press.
22. Ren, Y., Wang, N., Jiang, J., Zhu, J., Song, G., & Chen, X. (2019). The application of downhole vibration factor in drilling tool reliability big data analytics—A review. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, *5*(1), 010801.
23. Sadraoui, Y., Er-ratby, M., Kadiri, M. S., & Kobi, A. (2025). Improving wind turbine reliability through the analysis of critical component failures, parametric modelling, and maintenance optimization. *Wind Engineering*, 0309524X251386356.
24. Suslu, B., Ali, F., & Jennions, I. K. (2023). Understanding the role of sensor optimisation in complex systems. *Sensors*, *23*(18), 7819.
25. Vandawaker, R. M. (2015). *Integrated Systems Health Management as an Enabler for Condition Based Maintenance and Autonomic Logistics*.
26. Vieira, H., & Castro, R. (2025). A Forward-Looking Assessment of Robotized Operation and Maintenance Practices for Offshore Wind Farms. *Energies (19961073)*, *18*(6).
27. Vogl, G. W., Weiss, B. A., & Helu, M. (2019). A review of diagnostic and prognostic capabilities and best practices for manufacturing. *Journal of Intelligent Manufacturing*, *30*(1), 79–95.
28. Werbinska-Wojciechowska, S., & Rogowski, R. (2025). Proactive Maintenance of Pump Systems Operating in the Mining Industry—A Systematic Review. *Sensors*, *25*(8), 2365.