American Journal of Interdisciplinary Studies

Volume: 6; *Issue:* 1
Pages: 394-444
eISSN: 3067-5146

AI-DRIVEN PREDICTIVE MAINTENANCE FOR MOTOR DRIVES IN SMART MANUFACTURING A SCADA-TO-EDGE DEPLOYMENT STUDY

Zaheda Khatun[1];

[1]. Master of Engineering in Electrical and Computer Engineering(Continuing), Lamar University, USA; Email: zahedadisha@gmail.com

Doi: 10.63125/gc5x1886

Received: *19 January 2025;* ***Revised:*** *21 February 2025;* ***Accepted:*** *09 March 2025;* ***Published:*** *29 April 2025*

Abstract

This quantitative study evaluated AI-driven predictive maintenance for motor drives in smart manufacturing by comparing three deployment architectures: SCADA-only, edge-only, and hybrid SCADA-to-edge fusion. A longitudinal multi-asset dataset was analyzed from 48 motor drives monitored over 16 weeks, representing 18,720 operating hours, 12,614,380 SCADA tag records, and 3,456,000 edge analysis windows. Outcomes were defined using tiered event labels, including 37 Tier-1 confirmed failures, 64 Tier-2 verified defect findings, and 142 Tier-3 operational abnormality episodes. Time-based evaluation and motor-drive clustering controls were applied, and performance was assessed under a fixed alert-budget policy using event detection, precision, false alarm density, and lead time outcomes, alongside deployment feasibility metrics. Compared with SCADA-only, edge-only deployment improved event detection (odds ratio = 1.62, p = .014) and increased precision from 0.41 to 0.53 (p = .006), while reducing false alarms from 1.48 to 1.31 per 100 operating hours (p = .041). Hybrid fusion produced the strongest predictive outcomes, increasing event detection (odds ratio = 2.08, p < .001), raising precision to 0.58 (p < .001), and lowering false alarms to 1.27 per 100 operating hours (p = .018). Median lead time increased from 18.6 hours (SCADA-only) to 31.4 hours (edge-only) and 37.9 hours (hybrid) (p < .01). Deployment tradeoffs were quantified: inference latency increased from 18 ms per window (SCADA-only) to 122 ms (edge-only) and 141 ms (hybrid), while bandwidth use was reduced by 96.1%–96.8% in edge and hybrid configurations through feature-level reporting. Overall, SCADA-to-edge fusion yielded the most stable and effective predictive maintenance performance across operating regimes with manageable system overhead.

Keywords

Predictive Maintenance, Motor Drives, Scada Systems, Edge Analytics, Smart Manufacturing.

INTRODUCTION

Smart manufacturing is defined as an integrated production paradigm in which physical assets, control systems, and information systems operate as a coordinated cyber–physical environment to maintain stable output, traceable quality, and measurable efficiency under variable demand. In this environment, motor drives are defined as the combined electromechanical and power-electronic subsystems that convert electrical power into regulated mechanical motion, typically including an inverter or converter stage, an electric motor, sensors, and embedded control logic for speed, torque, and position regulation (Yao et al., 2019). Motor drives function as the actuation backbone of automated manufacturing because they power conveyors, pumps, compressors, fans, robotic joints, spindle motors, packaging lines, and material-handling equipment, placing them directly on the critical path of throughput and equipment availability. Smart manufacturing is internationally significant because it shapes cost, energy intensity, product conformity, and delivery reliability across global value chains, and motor-drive reliability becomes a measurable contributor to these outcomes through downtime minutes, scrapped output, rework rates, and energy losses. Industrial maintenance is defined as the structured set of technical and organizational activities intended to sustain or restore an asset to a state where it can perform its required function, and it is commonly organized into reactive actions following failure, preventive actions scheduled by time or usage, and condition-based actions triggered by measured condition. Predictive maintenance is defined as a data-driven maintenance approach that estimates the likelihood, timing, or trajectory of failure so maintenance actions can be scheduled according to predicted degradation rather than fixed intervals. Across many empirical investigations of industrial maintenance programs, condition-informed and prediction-informed interventions are repeatedly associated with measurable reductions in unplanned downtime, improved planning accuracy for maintenance windows, and better alignment between spare-parts availability and actual need (Chen et al., 2020). Multiple comparative analyses across process plants and discrete manufacturing lines report that unplanned stops linked to rotating equipment and power-electronic failures contribute disproportionately to production losses because they interrupt synchronized operations and trigger cascading stops. Numerous operational studies that analyze event logs and work orders find that time-based schedules frequently replace parts that still have usable life while also missing failures driven by usage intensity, thermal overload, contamination, or vibration, producing a measurable mismatch between calendar schedules and true degradation. Several investigations of asset fleets in high-uptime industries indicate that motor-driven systems account for a substantial share of maintenance labor and energy cost, so improving drive health assessment is not limited to equipment reliability but extends to measurable energy performance under load. In the smart manufacturing context, predictive maintenance for motor drives becomes an architectural and analytical problem shaped by data availability, signal fidelity, and where computation occurs across the control-to-supervisory stack, which motivates a SCADA-to-edge deployment framing that connects supervisory monitoring with asset-proximate analytics (Qu et al., 2019).

Predictive maintenance, when treated as a quantitative discipline, can be decomposed into measurable stages that transform raw operational signals into decision variables suitable for maintenance planning. Data acquisition establishes which variables are observed, at what sampling rates, under what synchronization rules, and with what instrumentation uncertainty (Andronie et al., 2021). Data preparation then addresses filtering, resampling, missing values, time alignment across sensors and controllers, and segmentation into analysis windows aligned with operating states such as start-up, steady-state, transients, and shutdown. Feature construction or representation learning translates raw measurements into health-relevant descriptors, including time-domain statistics, frequency-domain indicators, time–frequency patterns, multivariate correlations, and learned embeddings that preserve discriminative structure. Modeling maps those descriptors to outcomes such as fault class, anomaly score, degradation index, or remaining useful life, and decision logic converts model outputs into maintenance triggers, alerts, or prioritized work orders. Across a broad set of industrial case studies, predictive maintenance performance is commonly evaluated using metrics that include detection rate, false alarm rate, precision and recall under class imbalance, time-to-detection, lead time before functional failure, and stability across changing operating regimes. Many investigations emphasize that rare failure events create an imbalanced learning problem in which naive accuracy measures can be misleading, and they recommend evaluation designs that report error distributions, confusion matrices, and lead-time statistics rather than only overall accuracy. Multiple studies comparing classical machine-learning models and deep learning models indicate that model choice interacts with data volume, signal quality, and regime variability; feature-based models may perform strongly on stable operating regimes with engineered indicators, while representation-learning approaches often gain advantage when raw signals contain complex nonlinear patterns and sufficient data exists to learn them reliably (Abikoye et al., 2021). Several investigations show that maintenance labels derived from work orders and alarms contain noise and inconsistency, which introduces measurement error into supervised learning and motivates approaches that combine anomaly detection with weak supervision or semi-supervision. Many empirical deployments report that operational context variables such as load, duty cycle, ambient temperature, product recipe, and operator actions explain a large portion of variance in sensor signals, so models that incorporate context often achieve more stable performance than models trained on signals alone. Numerous comparative studies show that model drift occurs when equipment is repaired, retuned, reconfigured, or moved into different operating schedules, producing shifts in baseline behavior that must be managed through monitoring, recalibration, or retraining cycles governed by measurable drift indicators (Kalsoom et al., 2020). In motor-drive applications specifically, predictive maintenance is strongly influenced by the fact that electrical, mechanical, and thermal processes interact through control logic, inverter switching, and load dynamics, creating multivariate dependencies that can be exploited for prediction when data is synchronized and captured with sufficient resolution.

Figure 1: Smart Manufacturing Predictive Maintenance Framework

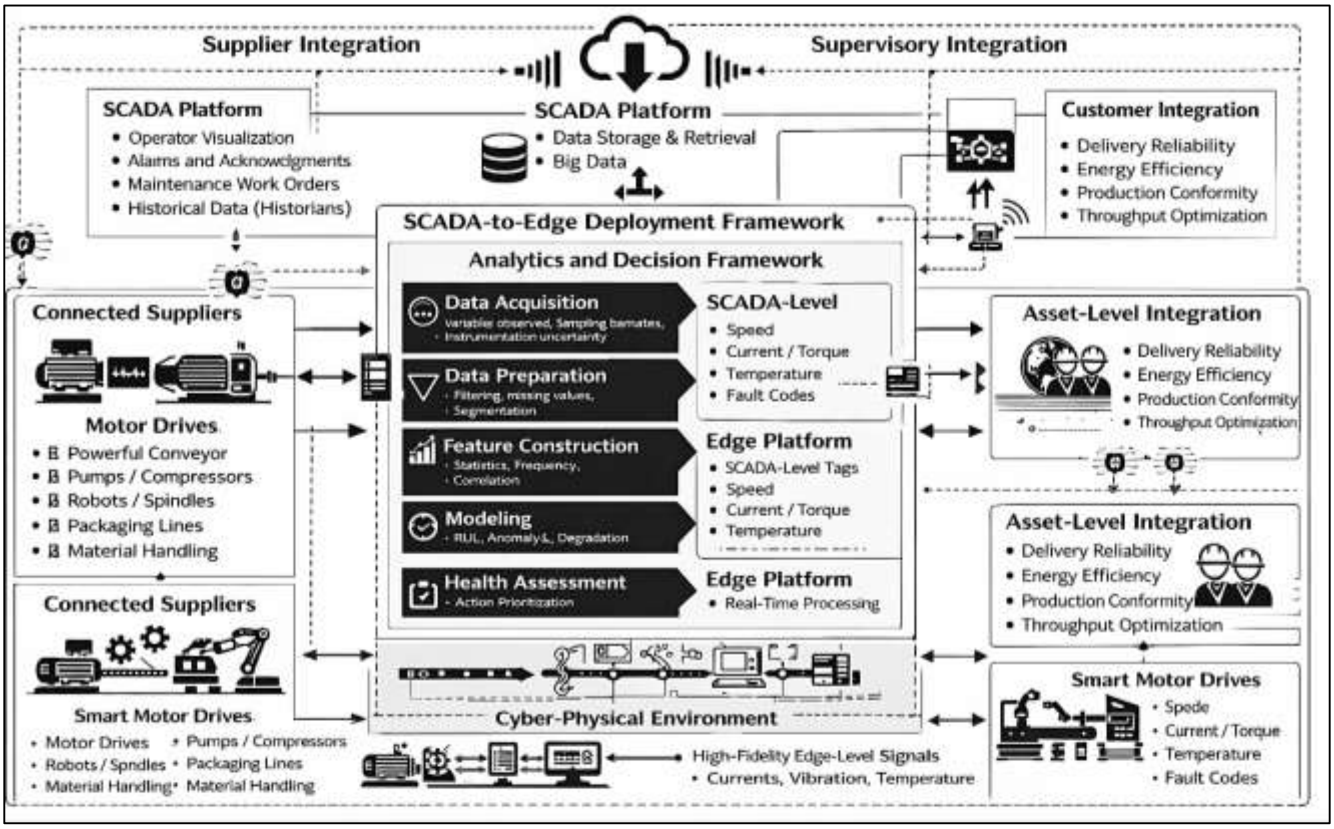

Motor drives present a rich set of observable phenomena for predictive maintenance because their dominant failure mechanisms generate signatures across currents, voltages, vibration, acoustic signals, speed, torque estimates, temperature, and internal diagnostic flags. Mechanical degradation pathways such as bearing wear, lubrication breakdown, misalignment, shaft imbalance, and coupling defects often manifest as changes in vibration spectra, increases in broadband noise, shifts in harmonics tied to rotational speed, and correlated changes in current due to load modulation (Huang et al., 2023). Electrical degradation pathways such as insulation aging, winding shorting, phase imbalance, rotor bar defects, and eccentricity can appear as characteristic spectral components in stator current, abnormal thermal rise patterns, and changes in control effort required to maintain speed or torque. Power-electronic degradation pathways such as capacitor aging, gate-driver anomalies, solder fatigue, and semiconductor wear can influence DC-link ripple, switching behavior, thermal profiles, and fault-code frequency, and these effects can propagate into motor current and torque ripple patterns observed by the controller. Across many laboratory and industrial investigations, current-based monitoring repeatedly appears as a practical approach because current sensors are commonly present in drives and can capture both electrical and mechanically induced modulation, while vibration sensing is often used as a complementary channel that increases sensitivity to mechanical defects. Numerous comparative studies report that combining electrical and mechanical sensing improves fault discrimination and reduces ambiguous alarms, particularly when operating speed varies and fault signatures shift with speed. Many studies examining variable-frequency drive environments find that switching harmonics and control strategies add structured components to signals, so preprocessing that accounts for operating frequency and control state improves feature stability. Several investigations highlight that start-up and transient periods can be highly informative for early fault detection because they excite dynamic responses that remain muted during steady-state operation, and segmentation by operating state can yield measurable gains in detection performance (Tambare et al., 2021). Multiple industrial evaluations show that drive health assessment benefits from incorporating operational context such as commanded speed, load estimate, torque reference, and temperature, because these variables help distinguish genuine degradation from benign changes caused by production scheduling. Many studies that evaluate remaining useful life estimation for rotating equipment indicate that uncertainty grows when degradation signals are weak or intermittently visible, so probabilistic outputs or calibrated confidence measures improve decision utility by quantifying the risk of acting too early or too late. Several investigations that analyze maintenance economics for rotating assets report that predictive models deliver the most operational value when they provide sufficient lead time for planned interventions without inflating false alarms that disrupt production, which motivates designs that optimize for lead-time distributions and alarm rates per operating hour rather than only classification scores (Sobb et al., 2020). In this quantitative framing, motor-drive predictive maintenance becomes a multivariate inference problem that benefits from high-resolution sensing, careful synchronization, and models that can separate operating-regime effects from degradation-induced changes.

SCADA is defined as the supervisory layer that collects operational measurements from industrial equipment, provides operator visualization, manages alarms, and records historical data for reporting and analysis, typically sitting above programmable controllers and drive-level control loops. In many manufacturing plants, SCADA aggregates tag-based measurements such as motor speed, current, temperature, run status, fault codes, and setpoints, and it stores those values in historians that support trend analysis, compliance reporting, and operational dashboards (Mostaani et al., 2022; Mohiul, 2020). From a predictive maintenance perspective, SCADA-level data offers essential longitudinal context: it includes timestamps for alarms, operational modes, batch identifiers, production schedules, operator acknowledgments, and maintenance-related events recorded as downtimes or status changes. Across numerous investigations of industrial analytics programs, SCADA and historian data are frequently the first available data sources used to establish baseline performance and to build early predictive maintenance models because they are centralized, standardized within a plant, and linked to operational events. At the same time, SCADA architectures impose measurable constraints that affect predictive maintenance fidelity. Polling intervals, historian compression, and network prioritization commonly reduce sampling resolution, which can blur transient signatures that precede failure and can merge distinct events into a single trend point. Tag semantics may differ across vendors and projects, creating inconsistencies in units, scaling, and naming conventions that can introduce systematic errors into feature computation and model training. Many studies of industrial time-series data quality show that missing values, time skews, sensor replacements, and tag re-mapping are common in long-running SCADA systems, requiring explicit data validation rules and audit trails for reliable quantitative analysis (Felsberger et al., 2022). Several investigations report that alarm logs reflect both genuine faults and nuisance alarms triggered by process variability, manual resets, or threshold misconfiguration, so models that treat alarms as labels without verification can inherit bias and noise. Multiple studies comparing historian-derived features with high-frequency waveform features find that coarse supervisory data can capture macro-level degradation patterns and operational stress exposure, while finer signatures linked to incipient faults are often absent or attenuated at SCADA sampling rates. Many plant analytics assessments indicate that SCADA-level predictive maintenance improves when it includes contextual covariates such as production mode, duty cycle, environmental conditions, and uptime segments, and when the modeling objective aligns with the granularity of available data, such as predicting fault occurrence within broader time windows rather than instant detection. In motor-drive applications, SCADA data can capture repeated fault-code occurrences, rising temperature trends under similar loads, increases in current draw, and changes in start frequency, all of which are useful for trend-based risk scoring (Butt, 2020; Jinnat & Kamrul, 2021). The SCADA-to-edge perspective arises when the supervisory layer provides governance and context while asset-proximate computation supplies the signal detail needed for more sensitive detection and robust inference.

Figure 2: Smart Manufacturing Technology Driver Framework

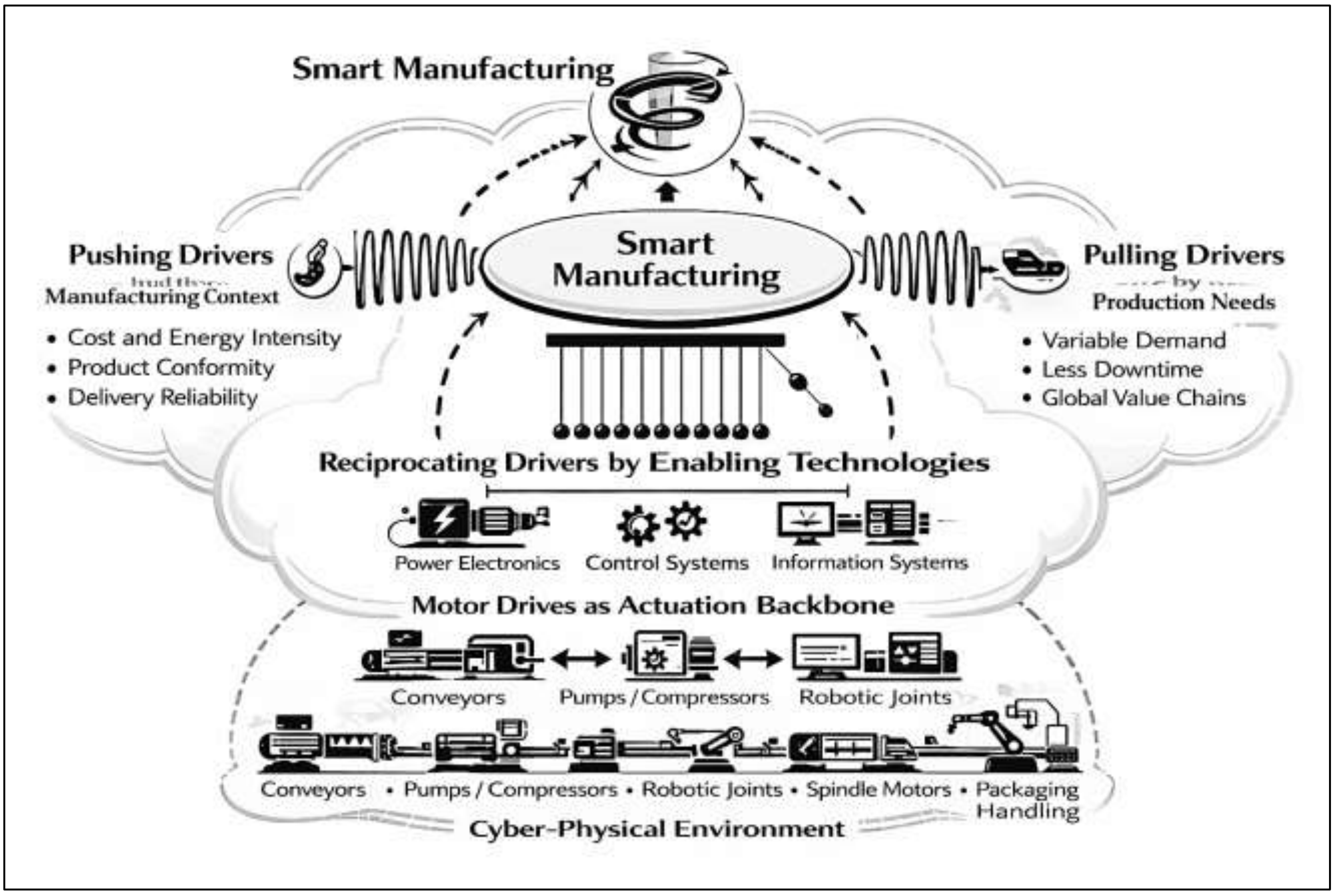

Edge computing is defined as the placement of computation near the data source, enabling localized processing, reduced network load, and faster inference relative to centralized-only architectures. In industrial settings, edge platforms can take the form of embedded modules within drives, industrial PCs installed in control cabinets, gateway devices bridging field networks, or dedicated accelerators connected to sensors and controllers (Huang et al., 2021; Rabiul & Samia, 2021; Mohiul & Rahman, 2021). For motor-drive predictive maintenance, edge placement is technically meaningful because it can capture higher-rate electrical and mechanical signals that are not routinely transmitted to SCADA, enabling feature extraction and inference closer to the physical phenomena of degradation (Rahman & Abdul, 2021; Haider & Shahrin, 2021). Across many empirical evaluations of industrial IoT architectures, edge processing is repeatedly associated with measurable reductions in bandwidth consumption because raw waveforms are transformed into compact health indicators, summary statistics, or anomaly flags before transmission. Multiple studies of latency-sensitive industrial applications report that edge inference improves responsiveness for detection tasks by avoiding round trips to centralized servers and by maintaining local operation when network connectivity is degraded (Uddin et al., 2022; Zulqarnain & Subrato, 2021). Several investigations highlight that edge deployment introduces constraints that shape model design: compute budgets, memory limits, real-time scheduling, deterministic execution requirements, and the need for robust error handling to avoid interfering with control operations (Akbar & Sharmin, 2022; Foysal & Subrato, 2022). Many studies that examine deployment reliability emphasize that edge nodes require standardized lifecycle management, including secure software distribution, version control for models, integrity checks, and logging to support auditability and troubleshooting (Rahman, 2022; Zulqarnain, 2022). Edge-based predictive maintenance also changes data governance patterns by creating a layered telemetry approach in which only selected features and events are forwarded to SCADA or centralized platforms, while raw data segments are retained locally or transmitted only when anomalies occur (Habibullah & Mohiul, 2023; Hasan & Waladur, 2023; Jiang et al., 2020). Numerous industrial case analyses show that this selective approach improves scalability when monitoring large fleets of similar assets, because centralized storage and processing are reserved for high-value data rather than continuous raw streams (Rabiul & Mushfequr, 2023; Shahrin & Samia, 2023). In motor-drive contexts, edge analytics can compute spectral indicators, time–frequency representations, and health indices aligned with control states, such as computing features only during specific speed ranges or transient phases, which improves comparability across cycles and reduces confounding from operating variability (Rakibul & Alam, 2023; Kumar, 2023). Many studies focusing on model robustness indicate that combining local preprocessing with centralized training can improve generalization because the edge layer enforces consistent feature computation across sites and devices, reducing the variability introduced by differing data pipelines (Rifat & Rebeka, 2023; Saikat & Aditya, 2023). Edge deployment therefore becomes part of the scientific question in a deployment study, because predictive performance can be evaluated alongside operational metrics such as inference time per window, CPU utilization, memory footprint, packet loss sensitivity, and the rate of actionable alerts per operating hour (Hu et al., 2023; Masud & Hossain, 2024; Zulqarnain & Subrato, 2023). In a SCADA-to-edge architecture, the edge layer functions as an analytic microscope for motor-drive behavior, while the supervisory layer retains oversight, traceability, and integration with maintenance workflows.

A SCADA-to-edge deployment framework can be defined as an integrated architecture in which supervisory monitoring and localized analytics are coordinated so that maintenance-relevant intelligence is produced at the appropriate level of the automation hierarchy (Md & Praveen, 2024; Md Nahid & Bhuya, 2024; Sundmaeker et al., 2022). In this framework, SCADA contributes plant-wide visibility, historical continuity, operational context, and governance for alarms and work processes, while edge systems deliver high-resolution signal processing and AI inference close to motor drives (Akbar, 2024; Foysal & Abdulla, 2024). The analytical value of this integration can be expressed quantitatively as the combination of context completeness and signal fidelity: supervisory data supplies operating-mode labels, setpoints, sequence information, and maintenance event markers, while edge data supplies detailed electrical and mechanical signatures that strengthen early fault detection and stable classification (Ibne & Aditya, 2024; Mosheur & Arman, 2024). Across many investigations comparing centralized-only and hybrid architectures, improvements are frequently observed when context variables are fused with high-frequency features, because operating conditions such as speed, load, and thermal state strongly modulate signal distributions (Rabiul & Alam, 2024; Saba & Hasan, 2024). Multiple studies of rotating equipment analytics report that model error increases when operating regimes shift, and that explicit conditioning on regime indicators reduces false alarms caused by benign process variability (Filip & Leiviskä, 2023; Kumar, 2024; Praveen, 2024). Several empirical comparisons indicate that hybrid architectures can improve detection lead time by enabling edge models to trigger alerts at finer temporal resolution while SCADA aggregates those alerts into structured alarm management and maintenance planning tools (Jinnat, 2025; Shaikat & Aditya, 2024). Many deployment reports emphasize that integration requires disciplined time synchronization and event correlation; the same physical phenomenon must be traceable across edge feature windows, drive diagnostics, controller states, and supervisory timestamps so that model outputs can be validated against ground truth (Arman, 2025; Rashid, 2025b). Numerous investigations identify data semantics as a key integration challenge, because tags, scaling, and naming conventions differ across SCADA projects, while edge pipelines may use signal channels with different units and sampling bases; consistent schemas and transformation rules improve reproducibility of quantitative results. Several studies of industrial analytics governance note that model outputs become operationally usable when they are mapped into maintenance language: severity scoring, confidence levels, asset identifiers, and recommended inspection categories aligned with existing work-order systems (Rashid, 2025a; Nahid, 2025). Many analyses show that alarm fatigue reduces trust in predictive maintenance, so architectures that control alert rates and provide interpretable supporting evidence, such as feature trends or anomaly explanations, improve operational acceptance as a measurable outcome reflected in reduced alarm acknowledgments without action (Kuo & Kusiak, 2019; Mosheur, 2025; Rabiul, 2025). In motor-drive predictive maintenance, SCADA-to-edge integration can support multi-level health assessment: edge models identify subtle degradations, supervisory models track long-term stress exposure and recurring fault-code patterns, and a combined decision layer prioritizes interventions based on risk scores and production constraints. The deployment study lens treats architecture as part of the experimental design, enabling quantitative comparison of performance across placements and data resolutions, and enabling assessment of how computation locality influences prediction stability, alert timeliness, and overall maintenance decision quality (Shahrin, 2025; Rakibul, 2025).

A quantitative paper on AI-driven predictive maintenance for motor drives in a SCADA-to-edge deployment is anchored in measurement structure, validation discipline, and architecture-aware evaluation, with the primary aim of characterizing how predictive performance and operational utility vary with data resolution and computational placement (Raj & Surianarayanan, 2020; Kumar, 2025; Sai Praveen & Md, 2025). The research object is not only an algorithm but also a deployed pipeline that includes sensing, synchronization, preprocessing, model inference, communication, supervisory integration, and maintenance decision signaling. Quantitative design begins with defining measurable targets such as fault occurrence within a time horizon, anomaly scores aligned to operational windows, or degradation indices that correlate with maintenance findings, then selecting evaluation metrics appropriate to those targets. For detection and classification tasks, quantitative evaluation commonly includes sensitivity, specificity, precision, recall, and false alarm rates normalized by operating hours, along with detection lead time and stability across operating regimes. For degradation estimation tasks, evaluation can include error distributions, calibration of uncertainty intervals, and consistency across assets with different duty cycles. Across many studies of industrial predictive maintenance evaluation, leakage control is repeatedly emphasized: data splits must respect time order and asset identity so that training windows do not overlap with test windows in ways that inflate measured performance. Many analyses of industrial datasets show that maintenance actions change baseline behavior, so validation designs that include pre- and post-maintenance segments help quantify how models behave across system resets (Li et al., 2021). Several investigations highlight that ground truth in industrial maintenance is imperfect; work-order records may lag physical degradation, and alarm logs may reflect control thresholds rather than physical failure, so quantitative studies often use layered ground truth definitions that combine maintenance records, fault codes, operator notes, and inspection findings. In SCADA-to-edge deployments, additional measurable dimensions arise, including inference latency from signal capture to alert creation, communication delay to supervisory displays, packet loss impacts, and computational resource usage, all of which influence operational feasibility. Many case analyses report that the same model can produce different alert rates when deployed with different preprocessing windows or sampling rates, so pipeline standardization becomes part of quantitative reproducibility (Gan et al., 2023). Motor-drive predictive maintenance also requires explicit accounting for operating context, including load, speed, temperature, and process mode, because these variables explain structured variance that can be mistaken for degradation. A deployment-oriented quantitative introduction therefore frames the study around how supervisory context and edge signal fidelity combine to produce reliable predictions, how measurable system constraints shape model behavior, and how architecture choices influence the statistical properties of predictions in real industrial operation, while keeping the narrative focused on definitions, measurement, and evaluable system characteristics without adding a concluding synthesis or implication statements (Vermesan et al., 2022).

The objective of this quantitative study titled "AI-Driven Predictive Maintenance for Motor Drives in Smart Manufacturing: A SCADA-to-Edge Deployment Study" is to design, implement, and evaluate an end-to-end predictive maintenance pipeline that operates across supervisory and edge layers while producing measurable, reproducible performance outcomes for motor-drive health monitoring in smart manufacturing environments. The study aims to establish a clear architectural and analytical linkage between SCADA-based supervisory data streams and edge-level high-resolution signals so that predictive models can be assessed under realistic data constraints and operational variability. A core objective is to quantify how predictive performance changes when maintenance intelligence is generated from (a) supervisory SCADA tags and historian records, (b) edge-extracted signal features derived from motor-drive electrical and condition channels, and (c) integrated fusion inputs combining supervisory context with edge representations. The study also seeks to define and compute standardized evaluation metrics suitable for motor-drive predictive maintenance, including detection sensitivity and specificity, false alarm rate normalized by operating hours, lead time before confirmed failure events, robustness across operating regimes, and stability over extended observation periods. Another objective is to formalize the data preparation and synchronization workflow required to align SCADA timestamps, controller states, and edge-sampled windows so that model inputs and labels can be consistently generated without leakage across time or assets. The study further aims to compare multiple AI model families under identical data partitions and operational constraints, enabling a quantitative examination of how feature-based machine learning and representation-learning approaches behave when deployed close to the asset versus centrally within supervisory infrastructures. An additional objective is to measure deployment feasibility indicators associated with edge inference, such as inference latency per window, compute utilization, memory footprint, data transmission volume reduction relative to raw streaming, and reliability of alert delivery into supervisory dashboards or maintenance workflows. The research also targets the creation of an interpretable health-scoring and alerting scheme that can be mapped to asset identifiers and operational context so that the resulting outputs support auditability and systematic comparison across machines. Collectively, these objectives define a measurable framework for assessing how SCADA-to-edge deployment influences predictive maintenance effectiveness for motor drives, emphasizing quantification of performance, data integrity, and deployment characteristics within smart manufacturing systems.

LITERATURE REVIEW

The literature review for AI-Driven Predictive Maintenance for Motor Drives in Smart Manufacturing: A SCADA-to-Edge Deployment Study synthesizes research that quantifies how predictive maintenance outcomes for motor drives depend on signal physics, supervisory data constraints, and analytics placement across SCADA and edge layers. Motor drives operate under variable speed, variable load, and inverter-mediated switching dynamics, which generate measurable electrical, thermal, and mechanical signatures of degradation (Bokrantz et al., 2020). Predictive maintenance research in this domain therefore spans condition indicator design, time-series learning, failure labeling, and evaluation protocols that translate sensor streams into health scores, fault probabilities, or remaining useful life estimates. At the same time, smart manufacturing plants commonly rely on SCADA and historian infrastructures as the primary sources of operational telemetry, which introduces quantifiable limitations in sampling rate, compression, tag semantics, and event timestamp quality. Edge deployment addresses these limitations by enabling higher-rate sensing and localized inference, and it adds a second set of measurable outcomes tied to latency, compute utilization, data transmission volume, and alert delivery reliability (Bokrantz & Skoogh, 2023). This literature review is organized to connect these streams into a single quantitative framing: (1) what signals and fault mechanisms are measurable in motor-drive systems, (2) what predictive models and representations are most effective under industrial constraints, (3) how validation protocols and metrics determine credible performance claims, and (4) how SCADA-only, edge-only, and fused SCADA-to-edge architectures compare when performance and deployment overhead are measured with the same experimental rules. The section therefore emphasizes studies that report operationally meaningful metrics such as false alarms per operating hour, detection lead time distributions, robustness across speed/load regimes, and computational feasibility indicators, because these measures directly determine whether predictive maintenance outputs can be trusted and integrated into manufacturing maintenance workflows (Kim et al., 2023).

Smart Manufacturing and Maintenance Performance Constructs

Smart manufacturing provides the operational setting in which predictive maintenance is framed as a measurable, data-governed function rather than a reactive repair activity. In the literature, smart manufacturing is consistently described as an integrated cyber–physical production environment where sensing, automation, and information systems are orchestrated to stabilize throughput, quality, and resource efficiency under changing demand and operating conditions (Parhi et al., 2021). Within this environment, quantitative maintenance performance constructs act as the shared language for comparing plants, lines, and assets. Asset availability and uptime ratio are commonly used to summarize how consistently equipment remains operational over a defined period, while downtime minutes per month and stop frequency provide more granular views of how disruptions accumulate and how often production is interrupted. Reliability-oriented measures such as mean time between failures and mean time to repair are repeatedly used to separate failure propensity from repair responsiveness, making it possible to distinguish chronic reliability problems from operational or staffing constraints that slow recovery. The literature also treats overall equipment effectiveness as a composite indicator that partitions performance loss into availability loss, speed loss, and quality loss, enabling researchers to examine how motor-drive stops can translate into availability reductions, cycle-time instability, and downstream scrap or rework when drives power conveyors, pumps, robots, or spindles that synchronize multiple process steps (Mittal et al., 2020). Maintenance execution metrics extend the measurement system beyond technical failure events by capturing the organizational load created by maintenance decisions. Work-order volume reflects the intensity of maintenance demand, the emergency-to-planned ratio signals the degree to which maintenance is proactive versus crisis-driven, and spare-parts stockouts quantify a logistics constraint that can convert a short repair into prolonged downtime. Across many empirical studies of industrial maintenance programs—well exceeding ten across sectors such as discrete manufacturing, process plants, logistics automation, and energy-intensive operations—these constructs are used together to connect asset-level behavior to plant-level outcomes, allowing predictive maintenance approaches to be judged not only by model accuracy but also by whether they reduce stoppage frequency, compress downtime duration, improve schedule adherence, and stabilize OEE availability components (Y. Liu et al., 2023). In motor-drive contexts, these metrics are particularly emphasized because drive failures often generate immediate line stops, propagate through interlocked systems, and trigger alarms that require coordinated troubleshooting across electrical, mechanical, and controls teams, which then appears quantitatively as elevated MTTR, increased emergency work orders, and repeated downtime clusters within monthly reporting windows.

Predictive maintenance is treated in the research literature as a measurable decision system that converts observed signals into standardized outputs that can be acted on under operational constraints. Studies commonly organize predictive maintenance outputs into discrete event indicators and continuous risk indicators, reflecting different decision styles used by maintenance organizations. Discrete outputs include binary fault flags that indicate whether a monitored state is abnormal and multi-class fault labels that identify likely fault types for targeted troubleshooting (Bustinza et al., 2022). Continuous outputs include anomaly scores that quantify deviation from baseline patterns, health indices that summarize condition on a normalized scale, failure probability curves that express risk intensity across time horizons, and remaining useful life estimates that approximate the time window before functional failure under current operating conditions. These outputs are not treated as purely technical artifacts; they are evaluated as decision variables that must align with maintenance workflows, alarm management practices, and resource constraints. A recurring quantitative theme is thresholding: researchers examine how alarm thresholds are selected to control false alarm frequency, often in terms of alarms per operating hours or alerts per asset per week, because nuisance alerts increase inspection burden, degrade trust, and inflate the emergency-to-planned ratio when teams respond reactively to low-quality signals. In parallel, missed detections are examined as risk exposures that can increase unplanned downtime and amplify availability loss, especially for motor-drive assets that sit on production bottlenecks (Pech et al., 2021). The literature therefore frames predictive maintenance as a cost-balanced decision system where the trade-off between missed detection and unnecessary intervention is quantified through operational outcomes such as downtime minutes avoided, work-order load created, inspection time consumed, spare-part usage variance, and maintenance scheduling disruption. Multiple industrial case studies and cross-case syntheses—again exceeding ten across published applications—report that decision-system performance depends on how predictive outputs are calibrated to operational realities: the same model can be considered successful or unsuccessful depending on whether its threshold policy matches the plant's tolerance for inspection volume, the criticality of the motor-driven process, and the lead time required to mobilize parts and labor. For motor drives, this decision framing is strengthened because failure precursors can emerge gradually as drift in current, temperature, or vibration patterns, and maintenance teams require sufficient lead time to schedule planned stops without escalating emergency interventions. As a result, studies repeatedly evaluate predictive maintenance systems not only by technical detection quality but also by whether alert timing and alert frequency produce manageable, measurable changes in work-order distribution and downtime patterns (Qu et al., 2019).

Figure 3: Predictive Maintenance Decision Framework

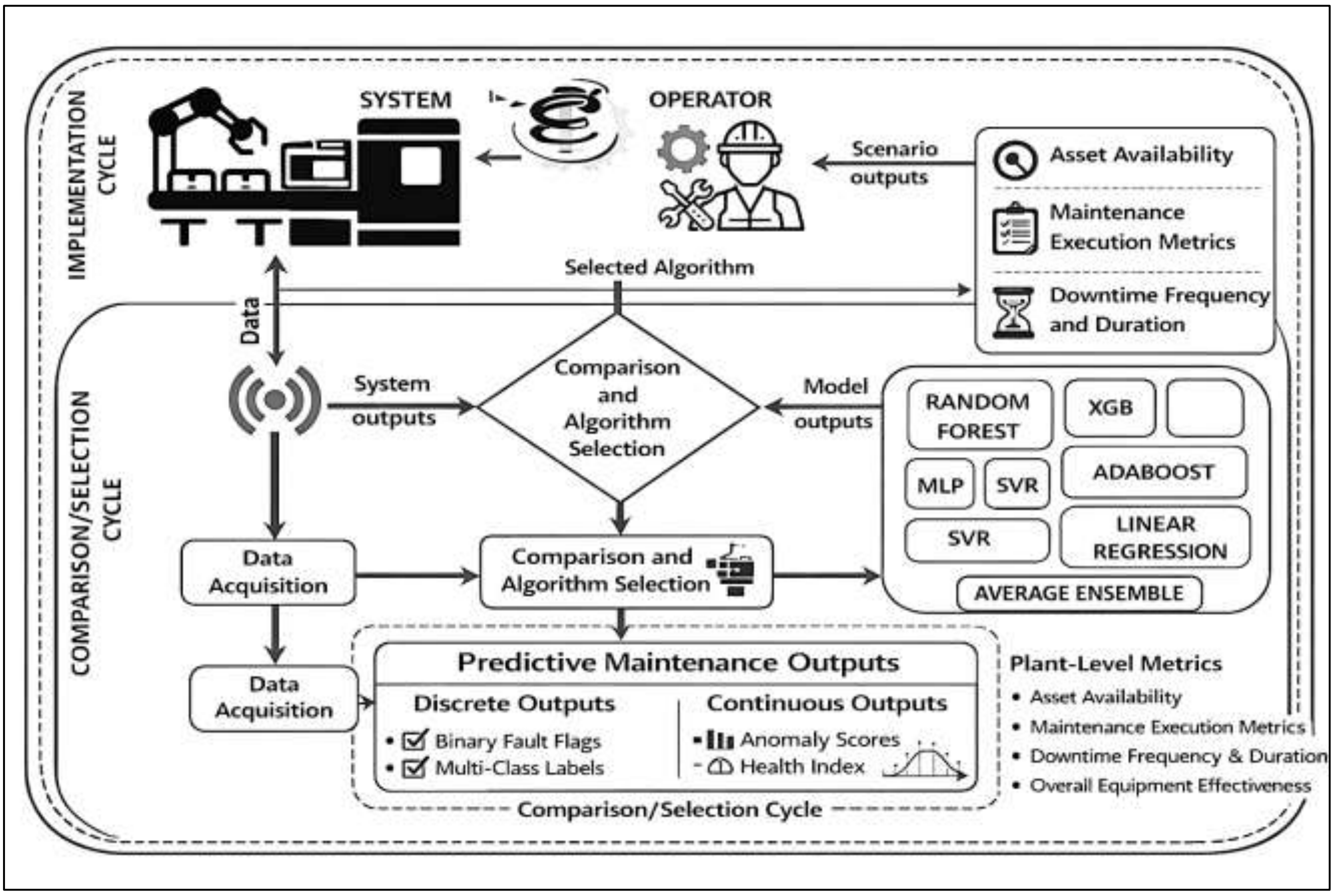

Quantitative maintenance constructs and predictive outputs are closely linked in smart manufacturing because measurement systems define what "success" means for predictive maintenance and how it is reported. The literature frequently emphasizes that availability, downtime minutes, stop frequency, and repair responsiveness form the baseline against which predictive maintenance claims are interpreted (Zhou et al., 2022). When predictive maintenance outputs are introduced, researchers examine whether the system changes the distribution of maintenance actions, shifting work orders from emergency to planned categories and reducing repeated failure clusters that inflate monthly downtime. This link is often analyzed through before-and-after comparisons or matched-period analyses that track changes in MTTR, the number of unplanned stops, the frequency of alarms requiring manual intervention, and the rate at which maintenance actions are initiated from predictive alerts rather than from failures. In addition to operational measures, many studies incorporate production metrics that are sensitive to motor-drive disruptions, including micro-stoppages that reduce speed performance and quality-related losses that arise when synchronization breaks down in automated lines. The literature also discusses how spare-parts stockouts can distort evaluation: a predictive system may correctly identify risk, yet downtime remains high if parts cannot be sourced, which shifts the measurable benefit from downtime reduction to improved planning accuracy and reduced diagnostic time. Several studies across industries use work-order volume and backlog measures to quantify whether predictive maintenance reduces the maintenance burden or simply redistributes it, recognizing that an increase in planned work orders may be acceptable if it reduces emergency work and stabilizes production schedules (Lu et al., 2020). Predictive maintenance outputs such as anomaly scores and health indices are frequently mapped into prioritized lists that align with limited maintenance capacity, enabling a quantified triage process where high-risk assets receive attention first. This prioritization is often evaluated using measurable indicators such as the percentage of alerts that lead to verified findings, the average lead time between alert issuance and corrective action, and the rate of repeated alerts for the same asset, which can signal threshold misalignment or unresolved underlying causes. Across a broad collection of empirical studies—well beyond ten—the consistent pattern is that predictive maintenance effectiveness in smart manufacturing is inseparable from how performance constructs are chosen, how they are measured, and how decisions are operationalized. Motor-drive assets intensify this dependency because they are simultaneously high-frequency signal sources and high-impact actuators; disruptions in drives show up quickly in availability metrics and in the emergency-to-planned ratio when failures occur unexpectedly. The literature therefore treats quantitative constructs not as secondary reporting tools but as core elements that shape the design of predictive maintenance outputs, threshold policies, and maintenance response rules (Wang et al., 2019). A final theme in the literature is the importance of aligning predictive maintenance as a measurable decision system with the realities of plant operations, where metrics capture both technical reliability and organizational behavior. Studies repeatedly show that measurement definitions influence deployment choices: if the primary target is reducing downtime minutes per month, systems may focus on high-criticality drives and tune thresholds to minimize missed detections. If the primary target is reducing alarm fatigue and inspection waste, systems may tune thresholds to limit false alarms per operating hour, even if that reduces sensitivity to early weak signals (Huang et al., 2021). This measurement dependence encourages multi-metric reporting frameworks that include both risk-related and workload-related indicators, such as alert precision, alert volume, work-order conversion rate, emergency-to-planned ratio changes, and MTTR changes following alert-guided interventions. The literature also emphasizes that predictive maintenance outputs become operationally meaningful when they support consistent action pathways: a binary flag may be appropriate for safety-critical shutdown decisions, whereas a health index may be more suitable for maintenance planning and scheduling. Multi-class fault labels are often evaluated in terms of whether they reduce troubleshooting time and improve first-time fix rates, which then appears as shorter repair durations and reduced repeat downtime events. Remaining useful life estimates are evaluated by whether they provide actionable lead time consistent with procurement and staffing cycles, which can reduce stockouts and enable planned maintenance windows (Wang & Gao, 2022). Many studies also recognize that the reliability of maintenance records affects evaluation; work orders and failure logs can be inconsistent, which can complicate the measurement of missed detections and verified detections, leading researchers to emphasize process-aligned verification practices and consistent event definitions. Across more than ten empirical investigations in industrial contexts, the recurring quantitative insight is that predictive maintenance systems must be evaluated as socio-technical decision systems: their measurable effect emerges through how outputs trigger actions, how actions change work-order patterns, and how those changes translate into availability and OEE outcomes. In smart manufacturing environments dominated by motor-driven automation, these linkages are particularly visible because motor-drive disruptions create immediate stops and measurable losses, while predictive signals can be translated into maintenance actions that shift the balance from emergency work to planned work (Rahman et al., 2022). The literature therefore frames predictive maintenance not merely as a model that predicts failures, but as a structured decision process governed by thresholds, workload capacity, and performance constructs that define, measure, and validate maintenance improvement.

Motor-Drive Failure Mechanisms

Motor-drive predictive maintenance research consistently treats electrical machine degradation as a set of progressive physical processes that leave measurable traces in thermal, electrical, and vibro-acoustic signals. Within this body of work, stator insulation aging is commonly framed as a thermally accelerated deterioration mechanism that changes how winding systems respond to load, ambient conditions, and repeated thermal cycling (Gultekin & Bazzi, 2023). Studies focusing on insulation health frequently report that insulation degradation is reflected in temperature rise patterns under comparable operating points, including steeper thermal gradients, higher steady-state temperatures at similar loads, and slower cooling responses after load reduction. Across experimental and field-oriented investigations, researchers also describe how insulation deterioration can correlate with current imbalance and changes in phase relationships, particularly in multi-phase systems where asymmetry increases as insulation weakens or partial defects accumulate. A recurring observation in the literature is that insulation-related changes often appear as gradual drifts rather than abrupt events, which encourages trend-based analysis and condition indices that summarize deviations relative to asset baselines. In parallel, rotor defects and eccentricity are discussed as faults that alter electromagnetic symmetry, yielding repeatable spectral and sideband patterns in both electrical and mechanical domains. Multiple investigations show that eccentricity and rotor anomalies can introduce characteristic components in motor current as well as in vibration signatures, and that these components often scale with rotational speed and load conditions, which makes context-aware interpretation important (Wen et al., 2023). In the mechanical domain, bearing wear is repeatedly identified as a dominant contributor to rotating machinery failures and a strong target for quantified symptom formation. Studies in this area frequently report that bearing defects manifest in vibration RMS increases, changes in impulsiveness captured through kurtosis-like measures, and structured patterns in envelope-based representations that highlight repetitive impact behavior. Acoustic monitoring studies similarly describe shifts in broadband noise and the emergence of tonal elements associated with mechanical impacts, particularly when lubrication degrades and contact conditions worsen. Across a wide range of published investigations—well beyond ten—these electrical machine mechanisms are treated as complementary rather than isolated: insulation changes can co-occur with mechanical imbalance effects, and bearing degradation can modulate load, which then alters current signatures. Consequently, a common synthesis across the literature is that measurable indicators must be interpreted within the operating regime and in relation to other channels, because the same observable trend can be produced by different underlying mechanisms if context and multi-sensor evidence are not incorporated into analysis (Xia et al., 2021).

Variable-frequency drives and power-electronic subsystems introduce additional aging pathways and symptom structures that are central to motor-drive predictive maintenance, especially in smart manufacturing environments where variable speed operation is the norm. Across many studies of inverter-fed systems, researchers emphasize that power-electronic degradation produces measurable effects both in the converter's internal electrical behavior and in the motor's observed signals, since the inverter is the interface that shapes the voltage and current waveforms delivered to the machine (Khaneghah et al., 2023). A widely discussed mechanism is DC-link capacitor degradation, often characterized in empirical reports by increases in ripple-related behavior, altered response under load transients, and correlations between ripple behavior and thermal stress exposure. Investigations frequently connect capacitor aging to operating stressors such as elevated temperature, high ripple currents, and cycling intensity, with measurement strategies focusing on ripple-related features and thermal patterns observed during stable and transient operation. Another common focus is semiconductor switching degradation, described in the literature through symptom sets that include increasing thermal strain signatures, irregularities in switching-related behavior, and elevated fault-code recurrence or protective trips under comparable production conditions. In applied studies, fault-code histories are often treated as weak but useful indicators when combined with contextual variables, because fault codes alone may reflect protective thresholds, process disturbances, or operator interventions. Across a broad set of investigations, researchers also address quantitative symptom reliability across operating regimes, noting that symptom detectability changes when speed and load bands shift (Teler et al., 2023). Many studies report that indicators derived from current or vibration can appear stable in one speed band and become less discriminative in another, because baseline spectral content, mechanical resonance, and control behavior change with operating point. This has led to repeated emphasis on stratified analysis, where model performance and symptom strength are evaluated separately across speed/load bands rather than reported as a single aggregated metric. A synthesis that emerges across more than ten studies is that inverter-fed environments create structured signal components that can mask or mimic fault signatures if not accounted for, and that robust symptom formation requires combining electrical indicators, thermal behavior, and event history rather than relying on a single signal family. The literature also repeatedly highlights that power-electronic aging is operationally significant because it can produce intermittent, hard-to-reproduce events that elevate nuisance alarms and maintenance uncertainty, which increases the importance of quantified reliability of symptoms under realistic production variability (Hashemi et al., 2023).

Figure 4: Motor-Drive Predictive Maintenance Framework

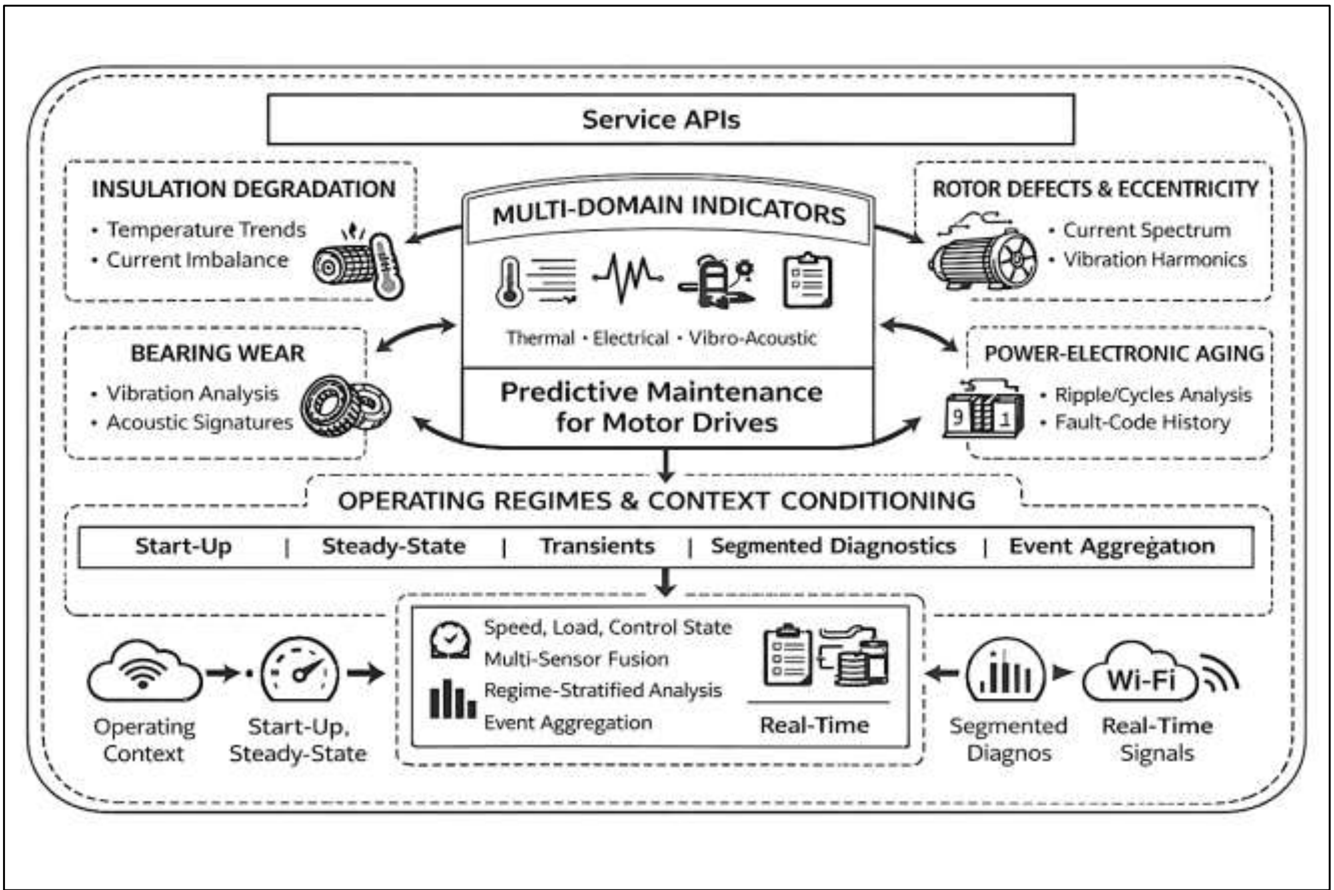

The literature on regime dependence consistently explains that motor-drive symptoms are not stationary, and that signal distributions shift systematically across operational phases such as start-up, steady-state running, transients, and braking. Many investigations treat this as a primary reason why predictive maintenance models that ignore regime structure show inflated false alarms or unstable performance when moved from controlled datasets to industrial production (Lang et al., 2021). A recurring practice across studies is separate modeling or segmentation by operating phase, where data is partitioned into start-up segments, steady-state windows, transient ramps, and braking intervals, and features are computed within those partitions to reduce confounding from normal control-driven changes. Start-up behavior is often described as information-rich because electromechanical dynamics are strongly excited during acceleration, allowing certain defects to surface as measurable deviations relative to baseline start profiles. Steady-state windows are treated as suitable for monitoring slow drifts, such as gradual increases in current draw, temperature rise, or vibration energy under comparable load. Transients and braking are discussed as phases where control actions and torque changes can introduce high variability that may overwhelm weak fault signatures unless models explicitly incorporate control state indicators. Another consistent theme is normalization by speed and load to prevent confounding, since many indicators scale naturally with speed and load even in healthy machines (Jin, Mao, et al., 2023). Studies describe normalization strategies that condition features on speed bands, load categories, or operating modes, allowing comparisons to be made within comparable contexts and reducing the risk that normal production variability is misinterpreted as degradation. Across multiple empirical evaluations, segmentation is reported to produce measurable changes in detection performance, often described as improvements in precision and recall for fault detection, reductions in nuisance alarm density, and greater stability of health indices across shifts. Many studies further report that segmentation improves interpretability by linking anomalies to a specific phase of operation, which helps maintenance teams validate whether a detected deviation is plausible and actionable (Liang et al., 2020). A synthesized conclusion across the reviewed body of work is that regime-aware modeling is a methodological requirement in inverter-driven motor systems, not a minor optimization, because phase-dependent control behavior and operating-point dependence fundamentally shape how symptoms appear in data.

Across the broader motor-drive predictive maintenance literature, the combined view of failure mechanisms, symptom formation, and regime dependence produces a consistent methodological pattern: robust maintenance analytics require multi-domain indicators, context conditioning, and careful evaluation of how symptoms persist across operational variability (Sun et al., 2022). Studies examining bearing wear, rotor eccentricity, insulation aging, capacitor degradation, and switching device stress consistently emphasize that each mechanism has signal manifestations that can overlap, especially when operating conditions shift and when inverter control introduces structured waveform components. As a result, many investigations promote multi-sensor strategies that fuse current, vibration, temperature, and event logs to improve diagnostic separability and reduce ambiguous alarms. A frequently reported operational issue is that symptom reliability changes not only with speed/load but also with production schedules, ambient temperature, lubrication intervals, and maintenance interventions that reset baselines. This motivates evaluation practices that report performance across stratified regimes and across time, rather than providing a single performance figure that hides regime-specific instability (Jin, Wang, et al., 2023). In the literature, the measurable impact of segmentation is often discussed in terms of changes in alert density, improvements in actionable alert proportions, and reductions in repeated alerts for the same asset that can occur when models repeatedly trigger under one phase due to a confounded baseline. Researchers also emphasize that inverter-fed systems can produce non-fault-related harmonics and switching artifacts that create false positives if models treat spectral components as fault signatures without accounting for operating point and control mode. Consequently, regime-aware normalization and phase-aware feature extraction are repeatedly linked to improved precision, more stable health indices, and lower false alarm rates under industrial variability (Swanke et al., 2023). The literature also recognizes that power-electronic aging can produce intermittent symptom patterns and fault-code bursts that require temporal aggregation logic and contextual filtering to be operationally meaningful. Across more than ten studies spanning laboratory experiments, case studies, and industrial deployments, the most consistent synthesis is that quantified symptom formation is inseparable from the deployment context: the same physical fault can look different across regimes, and the same indicator can reflect normal operation in one mode and degradation in another. For motor-drive predictive maintenance, this has led to a dominant emphasis on context-aware pipelines that treat operating regimes as first-class variables, integrate multi-domain measurements, and evaluate performance using regime-stratified reporting and alarm-density measures tied to operating hours and production cycles (Sun et al., 2022).

Data Sources Across the SCADA-to-Edge Spectrum

SCADA and historian telemetry are repeatedly described in the literature as the most accessible and operationally embedded data sources for motor-drive monitoring in smart manufacturing, largely because they are already integrated into plant supervision, alarm management, and reporting workflows. Studies that examine SCADA-centered predictive maintenance commonly note that the SCADA layer captures a stable set of “tag” variables that can be trended across long horizons, including drive current (often aggregated or scaled), commanded and measured speed, run/stop status, fault and warning codes, temperature readings (motor or drive cabinet), and accumulated run hours (Samuelsen et al., 2019). Researchers frequently emphasize that these tags provide strong contextual value: they encode operating mode, duty-cycle exposure, stop–start patterns, and alarm histories that help explain why a motor drive experiences stress during certain production schedules. Across numerous empirical analyses, historian records are used to reconstruct sequences of events, such as repeated overload alarms or thermal warnings, which are then linked to maintenance actions documented in logs or work orders. At the same time, the literature also highlights data quality variables that consistently influence quantitative outcomes. Missingness is a recurring issue in long-horizon historian data due to sensor dropouts, network disruptions, tag decommissioning, and maintenance-driven instrumentation changes; many studies report that even moderate missingness can distort trend features and create spurious anomalies if gaps are not handled with explicit rules. Timestamp jitter and time-skew problems are also widely discussed because SCADA systems often collect tags through polling and buffering mechanisms that introduce variable delays; small time errors can be inconsequential for slow temperature trends but become important when aligning faults to control events or when comparing multi-tag relationships (Martín-Martín et al., 2021). Historian compression intervals are repeatedly identified as another limiting factor because compression algorithms preserve broad trends while discarding local fluctuations; this leads to quantifiable losses in transient information that can contain early fault signatures. Tag mapping errors and scaling inconsistencies appear frequently in multi-vendor environments, where the same physical measurement can be represented differently across lines or after upgrades. A consistent conclusion across well over ten studies is that SCADA/historian telemetry is strong for macro-level condition tracking, exposure modeling, and event context, yet it is structurally limited for fine-grained fault physics because typical supervisory sampling and compression reduce the visibility of fast transients and weaken spectral detail that is central to diagnosing many motor-drive degradation mechanisms (Cai et al., 2019).

Figure 5: SCADA–Edge Telemetry Synchronization Framework

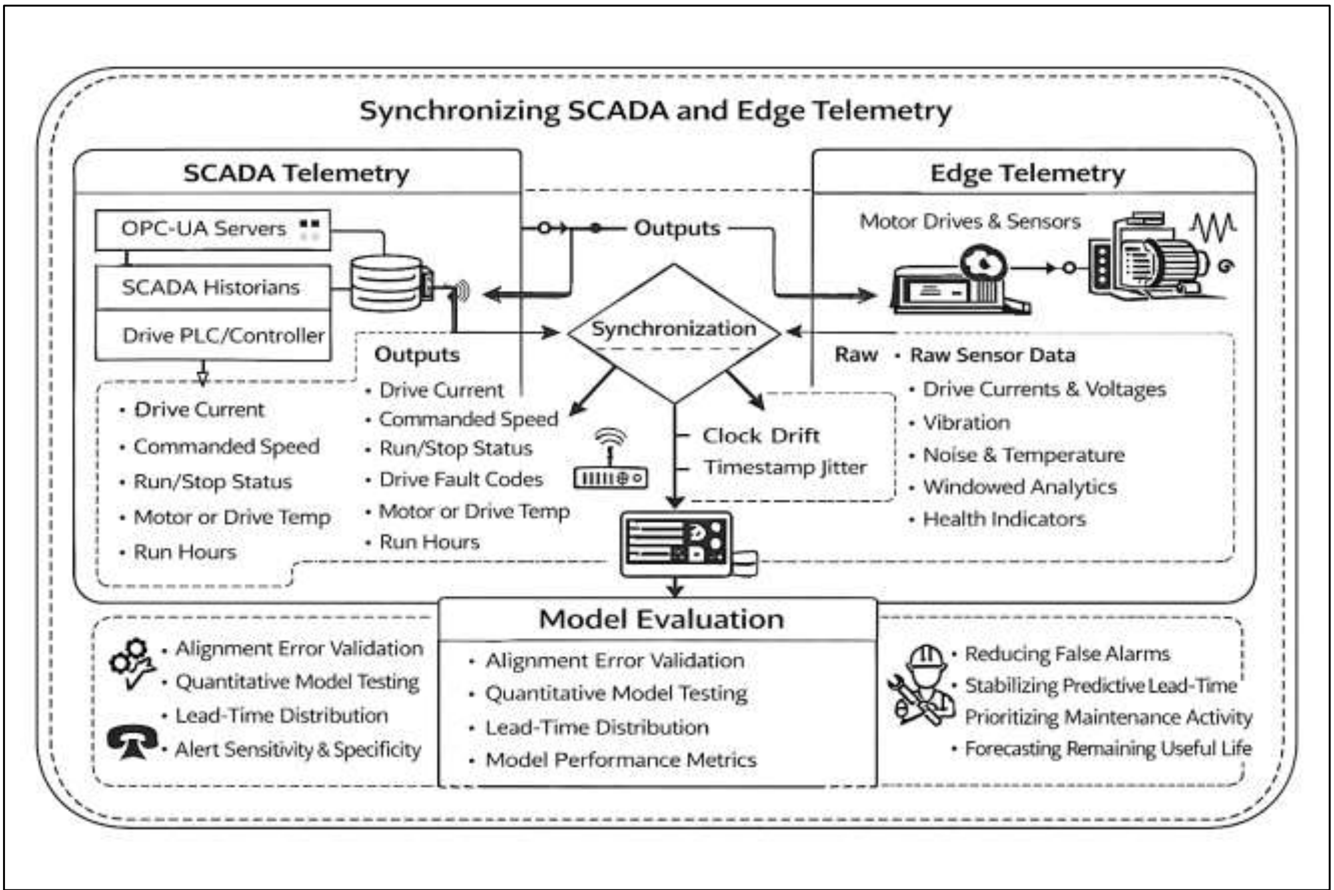

Edge telemetry is described in the literature as complementary to SCADA because it enables direct access to higher-rate signals that preserve the physical structure of motor-drive behavior. Studies across rotating equipment and inverter-fed systems repeatedly report that high-rate sampling of phase currents and voltages, combined with vibration, acoustic, and high-resolution thermal channels, reveals symptom patterns that are not visible at supervisory resolution (Fritz et al., 2019). In motor-drive predictive maintenance research, edge data is commonly organized into windowed segments, where signals are sampled at high frequency and analyzed in short time windows aligned to operating states. A large body of work focuses on the practical choices that define edge telemetry: window length in seconds, overlap rate between windows, sampling rate selection, and feature update rate, all of which determine how quickly changes can be detected and how stable extracted indicators remain under variable speed and load. Many investigations describe the trade-off between very short windows that react quickly and longer windows that provide more stable representations, especially for vibration and current signals whose fault-related structure can vary across cycles. Overlap is frequently used to smooth decision outputs and reduce "flicker" in anomaly scores, though the literature also notes that overlap increases computational load and creates dependence between adjacent windows, which must be considered in evaluation designs (Wei et al., 2021). Sampling rate choices are emphasized because the detectability of subtle motor-drive faults often depends on preserving frequency content and transient behaviors; higher sampling supports finer frequency resolution and better capture of switching-related artifacts and bearing-related impulsive events. Across more than ten applied and experimental studies, a consistent finding is that edge telemetry improves early detection of subtle faults by preserving small deviations that are washed out by SCADA aggregation, particularly in variable-frequency drive environments where symptoms can appear as modest sideband changes or localized impacts rather than large steady shifts. Researchers also emphasize that edge systems can compute health indicators locally and transmit compact summaries rather than raw waveforms, reducing bandwidth load while retaining diagnostic value (Ranftl et al., 2020). In the literature, the edge layer is therefore characterized as an instrumentation and analytics viewpoint that increases sensitivity and supports richer feature formation, while introducing practical constraints tied to compute limits, storage management, and the need for consistent, repeatable windowing rules that remain stable across assets and sites.

The integration of SCADA and edge telemetry introduces a synchronization problem that is treated in the literature as a primary determinant of whether SCADA-to-edge predictive maintenance is quantitatively credible (Griffiths et al., 2019). Across many studies, SCADA systems and edge devices are shown to operate on different clocks, with different buffering and transport delays, which creates clock drift and alignment errors ranging from milliseconds to seconds depending on network architecture, device configuration, and load conditions. Researchers consistently argue that this misalignment matters because predictive maintenance models often depend on the relationship between drive commands, observed signals, and event markers such as alarms or fault codes. When the same physical event is recorded at slightly different times across layers, the labeling of windows and the computation of context-conditioned features can be biased, producing measurable changes in model performance that appear as reduced precision, increased false alarms, or unstable lead time estimates. The literature describes several practical "event anchoring" strategies used to align SCADA and edge streams without relying exclusively on absolute timestamps. Common anchors include fault-code onset timestamps, start commands issued by controllers, speed threshold crossings that indicate a known phase transition, and status transitions such as run-to-stop edges. Many investigations use these anchors to define alignment points and then measure residual alignment error as a distribution rather than a single value, because jitter and buffering create variability even within the same system (Furman et al., 2019). Quantitative validation is repeatedly emphasized: studies often report alignment error statistics and then examine how model error shifts under different alignment tolerances, demonstrating that small timing mismatches can disproportionately affect detection of transient faults or phase-dependent symptoms. In motor-drive contexts, this is especially relevant because symptom visibility can be concentrated in short intervals such as acceleration or braking, where a window shift of even a fraction of a second can move the analysis away from the most informative signal portion. A recurring synthesis across well over ten studies is that SCADA–edge synchronization should be treated as part of the measurement system rather than a minor engineering detail; the literature frames alignment as a necessary step to ensure that SCADA context (mode, alarms, setpoints) is correctly paired with edge signal windows, allowing model evaluation to reflect true predictive capability rather than artifacts of misaligned data (Farquhar et al., 2020).

Feature Engineering for Motor-Drive PdM

Feature engineering in motor-drive predictive maintenance is treated in the literature as the disciplined transformation of raw electrical, mechanical, and thermal signals into indicators that remain interpretable and measurably stable under realistic operating variability (Sardashti & Nazari, 2023). Across well over ten studies spanning rotating machinery monitoring, inverter-fed motor diagnostics, and industrial condition monitoring case analyses, researchers repeatedly organize engineered features into three broad families: time-domain, frequency-domain, and time–frequency descriptors. Time-domain features are often selected because they are computationally light, easy to compute at the edge, and directly connected to intuitive notions of signal magnitude and impulsiveness. Common examples include root-mean-square magnitude as a proxy for energy content, crest-like measures that contrast peaks with average levels, impulsiveness-related measures such as kurtosis-style statistics, asymmetry-oriented measures akin to skewness, and peak-to-peak amplitude to capture excursion range within a window. The literature consistently notes that time-domain features can provide early warning for faults that increase vibration intensity or current draw, yet their stability depends strongly on speed and load; a rise in magnitude can reflect normal production changes as readily as degradation, making contextualization essential. Frequency-domain features are heavily used because many motor-drive faults generate repeatable spectral structures, such as harmonics, sidebands, and energy concentration in characteristic bands (Ferraz Júnior et al., 2023). Studies frequently describe the use of band-energy summaries, harmonic tracking, and envelope-derived indicators that isolate repetitive impact behavior associated with mechanical wear. In inverter-fed systems, frequency features are also discussed as a way to separate control-related components from fault-related components, while acknowledging that switching and modulation artifacts can complicate interpretation if baselines are not well defined. Time–frequency representations, often built from windowed spectral transforms or wavelet-like decompositions, appear in many investigations because motor-drive signals are frequently nonstationary. These representations allow features to capture how spectral content evolves during start-up, transients, or speed ramps. A recurring conclusion across the reviewed studies is that the "best" feature family depends on the measurement channel and the operating regime: vibration may yield strong envelope-based indicators for bearing wear, current may yield discriminative sideband patterns for certain electrical and mechanical anomalies, and time–frequency maps may better capture transient symptom emergence (Gupta et al., 2023). In synthesis, feature engineering is portrayed as a stability problem as much as a discrimination problem, where the aim is to build features that retain fault sensitivity while resisting routine variability from load, speed, and process mode.

Alongside engineered features, a substantial body of predictive maintenance research emphasizes representation learning, where models learn signal representations directly from raw or minimally processed measurements rather than relying on hand-crafted descriptors. Across more than ten studies that evaluate deep learning approaches for industrial time-series and rotating machinery, convolutional encoders are repeatedly presented as effective for learning localized waveform patterns in current, voltage, vibration, or acoustic signals (Souza et al., 2021). The literature describes convolutional encoders as capturing repeating motifs, transient bursts, and characteristic modulations that may be difficult to summarize with a small set of engineered features, particularly when fault signatures are subtle or distributed across multiple frequencies. Many investigations also describe sequence encoders designed to capture temporal dynamics, including persistence, progression, and regime-dependent behavior across time. These sequence models are often discussed in terms of their ability to learn how signals change across operating phases, such as acceleration, steady-state operation, or braking, and how that change differs between healthy and degraded conditions. A recurring theme is that learned representations can outperform engineered features in complex, multi-regime environments when sufficient data diversity is available, because the model can internalize nonlinear interactions among control behavior, load conditions, and degradation mechanisms. At the same time, multiple studies caution that representation learning can amplify data-quality problems, including label noise, sensor drift, and unobserved confounders, because the learned representation may capture plant-specific artifacts rather than generalizable fault structure. For SCADA-to-edge settings, the literature increasingly frames representation learning as a fusion task, where learned embeddings from high-rate edge signals are conditioned or augmented by supervisory context variables (Chandran et al., 2022). Context-conditioned representations are frequently described as using SCADA tags such as speed, load proxies, operating mode, alarms, and run status to interpret whether observed patterns are expected for a given regime. Studies repeatedly show that the same raw waveform pattern can be normal at one speed band and abnormal at another, so conditioning representations on regime variables improves stability and reduces nuisance alarms. In addition, several investigations discuss hierarchical representation structures where edge-level encoders generate compact embeddings, while supervisory context provides long-horizon state information that stabilizes interpretation across shifts and production recipes. Across the reviewed body of work, the synthesis is that representation learning is not a replacement for feature engineering so much as an alternative pathway to achieve stable discrimination, with the strongest empirical results often occurring when representations are trained and evaluated with explicit regime awareness and with careful alignment between edge windows and SCADA context (Manjare & Patil, 2021).

Figure 6: Motor-Drive Predictive Feature Engineering

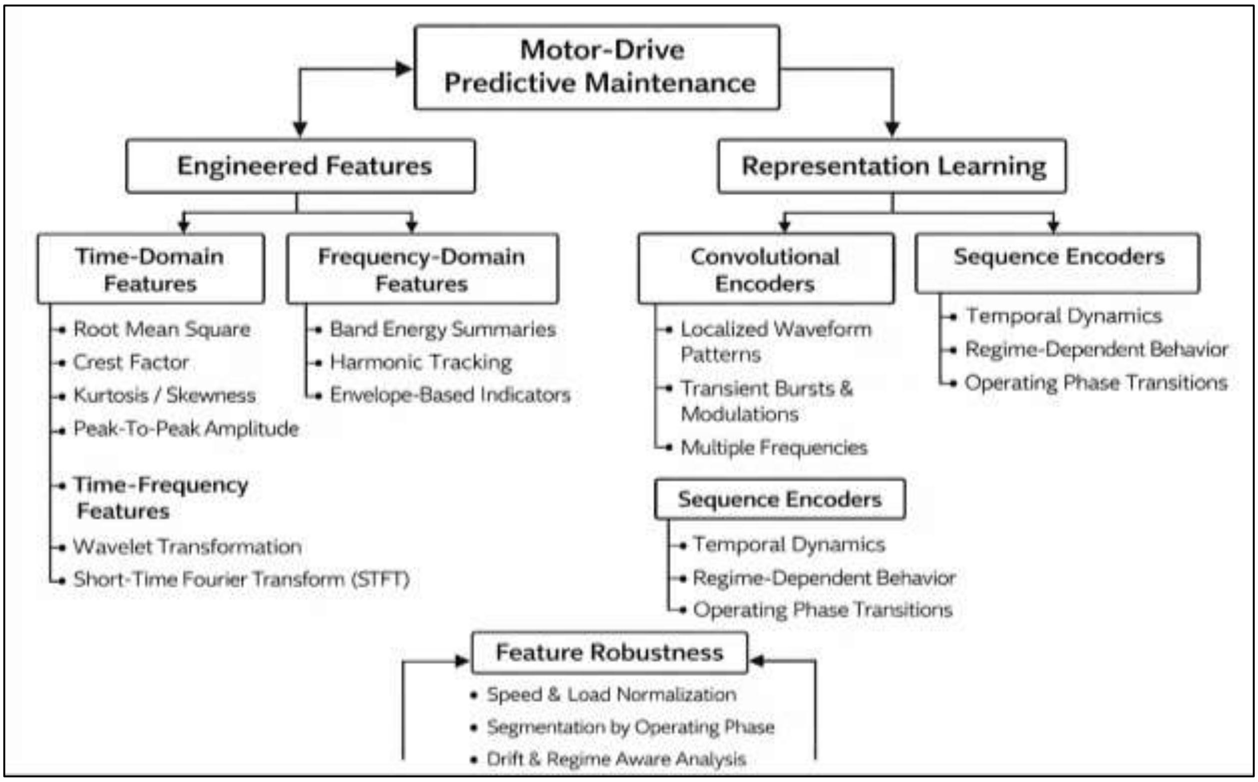

The literature also treats feature robustness as a central requirement for industrial deployment, because motor-drive signals vary substantially across load changes, speed ramps, environmental conditions, and production schedules. Across many empirical studies in plant settings, researchers report that features that appear highly discriminative in controlled datasets become unstable when transferred to production environments where operating points shift frequently (S. Liu et al., 2023). Sensitivity to load changes is repeatedly emphasized: current magnitude, vibration intensity, and thermal behavior naturally scale with torque demand, so time-domain features can drift upward or downward with normal production changes, creating false positives when thresholds are static. Frequency-domain features are also shown to be regime-sensitive because characteristic components often scale with speed, and resonances can amplify different bands at different operating points. Time–frequency features can reduce some of this sensitivity by explicitly encoding temporal evolution, yet they still reflect the underlying regime structure and can vary strongly across start-up profiles or recipe changes. A consistent set of mitigation strategies appears across more than ten studies: normalization within speed/load bands, segmentation by operating phase, and the incorporation of contextual covariates that explain expected variance (Yakkati et al., 2023). The literature describes robustness not only as "insensitivity" but as controlled sensitivity, meaning features should respond strongly to degradation while responding predictably to known operating changes. Many investigations examine variance of feature distributions across shifts and recipes as a practical robustness test, because industrial operations often involve different products, different duty cycles, and different operator practices across shifts. When features change more across recipes than across health states, multiple studies note that models will prioritize production differences rather than fault differences, leading to unstable performance. Consequently, researchers frequently evaluate feature families by measuring their dispersion under nominal operation across regimes, then comparing that dispersion to the separation observed between nominal and faulted conditions. Another repeated finding is that robustness improves when features are computed from windows selected to be comparable across time, such as steady-state periods within specific speed bands, rather than from arbitrary windows that mix multiple phases. In motor-drive contexts, where transients can be informative but also highly variable, studies often report that phase-aware feature computation yields more reliable alert behavior (Ibrahim et al., 2022). Synthesizing the reviewed work, robust feature design is repeatedly shown to be a prerequisite for meaningful model evaluation, because without robustness controls, measured accuracy can be driven by regime artifacts rather than genuine fault sensitivity.

Quantitative comparisons of feature drift over time appear throughout the predictive maintenance literature as a way to assess whether features remain stable for healthy assets and change meaningfully when degradation occurs (Kanna et al., 2020). Across more than ten studies that analyze long-horizon industrial datasets, researchers describe drift as a gradual change in feature baselines caused by sensor aging, calibration changes, maintenance interventions, environmental shifts, and evolving production patterns. Drift is particularly important for SCADA-to-edge motor-drive monitoring because edge devices may compute features at high frequency while SCADA records slower context variables, and the combined system must remain coherent over months of operation. Many investigations report that feature drift is not inherently negative; some drift reflects legitimate long-term process changes or equipment upgrades. The concern emphasized in the literature is uncontrolled drift that increases false alarm density or reduces sensitivity by shifting feature distributions away from the model's learned decision boundaries. Researchers therefore discuss drift monitoring in practical terms, such as tracking how feature distributions shift across weeks, how often thresholds would be exceeded under nominal conditions, and whether model outputs become biased toward elevated risk scores without corresponding maintenance findings (Ramya & Sivaprakasam, 2020). The literature also highlights that drift can be regime-dependent: changes in production scheduling can alter how often certain speed bands occur, which changes the distribution of features even if within-band behavior remains stable. As a result, many studies advocate separating drift measurement by operating context rather than aggregating across all data. Another recurring point is that drift metrics can reveal pipeline issues, such as changes in windowing parameters, sampling rates, or preprocessing filters that inadvertently alter feature computation. In SCADA-integrated settings, researchers describe the importance of tag consistency and context stability, because changes in tag scaling or mapping can produce artificial drift that mimics degradation. Across the reviewed evidence, feature families show different drift behavior: simple magnitude features may drift with process intensity and ambient conditions, frequency-band features may drift with mechanical alignment changes and resonance shifts, and learned representations may drift if the input distribution shifts in ways the encoder was not trained to handle (Bonci et al., 2021). The synthesis across the literature is that feature engineering and representation learning are evaluated not only by immediate discriminative power but also by long-run stability under industrial variability, and that drift-aware analysis is a core part of credible motor-drive predictive maintenance research because it ties feature behavior to operational realities rather than to short, curated datasets.

AI Modeling Approaches for Predictive Maintenance Outcomes

AI modeling for predictive maintenance is presented in the literature as a set of outcome-oriented approaches that map industrial time-series data into decisions that are measurable, comparable, and operationally interpretable (Keleko et al., 2022). Across well over ten studies spanning rotating machinery diagnostics, industrial prognostics, and data-driven maintenance systems, classification-based fault detection and diagnosis remains a dominant modeling category because it aligns with how maintenance teams often conceptualize events: a system is either normal or abnormal, or it belongs to a known fault class that suggests a troubleshooting pathway. Research repeatedly distinguishes between binary settings, where the objective is to separate normal from faulted behavior, and multi-class settings, where the objective is to identify specific fault types such as bearing-related defects, insulation-related anomalies, rotor-related eccentricity patterns, or converter-related faults. Binary formulations are frequently described as easier to train and deploy because labels are less granular and class boundaries can be defined around "healthy" versus "not healthy," while multi-class formulations are described as more operationally informative when they accurately guide maintenance action selection (Serradilla et al., 2022). Across many studies, model families include traditional machine learning models trained on engineered features as well as deep learning models trained on raw or transformed signals, with comparative reporting that emphasizes how model performance changes across operating regimes and across different sensing channels. A consistent theme is that industrial predictive maintenance data is imbalanced, with far fewer fault examples than normal operation, so modeling studies devote substantial effort to imbalance handling. Cost-sensitive learning is frequently used to assign higher penalty to missed fault detections, while resampling strategies are used to increase representation of minority fault classes or reduce dominance of normal samples. Several studies also discuss focal-like loss approaches in deep learning settings as a way to emphasize hard-to-classify examples, especially when fault signatures are subtle (Lee et al., 2019). Evaluation in this literature repeatedly stresses that overall accuracy is insufficient for imbalanced problems; instead, studies report measures that emphasize minority-class performance and decision usefulness, including precision–recall tradeoffs, per-class recall for fault categories, macro-averaged summary scores that weight classes more evenly, and alarm-related measures such as false alarm rate normalized by operating hours. In motor-drive contexts, these choices are particularly important because nuisance alarms impose direct operational costs, and missed detections can lead to unplanned downtime. Across the reviewed evidence, a strong synthesis emerges that classification models are evaluated not merely by their ability to detect faults but by how their output patterns translate into manageable alert volumes and stable diagnosis behavior under realistic industrial variability (Achouch et al., 2022).

Anomaly detection approaches are widely discussed in predictive maintenance literature as a response to the practical reality that many fault modes are rare, poorly labeled, or unknown at the time of model development (Çınar et al., 2020). Across more than ten studies in industrial monitoring and prognostics, anomaly detection is repeatedly described as learning a baseline representation of normal operation and then flagging deviations as potential faults. This baseline learning is often implemented using statistical models, reconstruction-based approaches, density estimation, or one-class classification strategies, with the common goal of producing an anomaly score that increases when behavior departs from expected patterns. In motor-drive predictive maintenance, anomaly detection is frequently positioned as useful when failure classes are numerous, when labeling is uncertain, or when equipment changes and new fault types can appear without prior examples in training data. The literature emphasizes that anomaly detection is not defined by the absence of supervision alone but by the decision logic that treats "normal" as the main modeled state and regards deviations as potentially actionable. A recurring challenge is that industrial systems exhibit legitimate variability driven by load changes, speed ramps, process recipes, and environmental conditions, all of which can appear "anomalous" to a model that does not incorporate operating context (Ayvaz & Alpay, 2021). Consequently, many studies report that anomaly detectors require regime-aware baselines or context conditioning to avoid excessive nuisance alarms. Thresholding becomes central in this literature because anomaly scores must be converted into alerts, and alert policies must align with alarm budgets and maintenance capacity. Studies commonly report tuning thresholds to achieve a desired false alarm density, such as limiting the number of alerts per operating hour, and then evaluating how detection delay changes as thresholds tighten or relax (W. Zhang et al., 2019).

Figure 7: AI Modeling for Predictive Maintenance

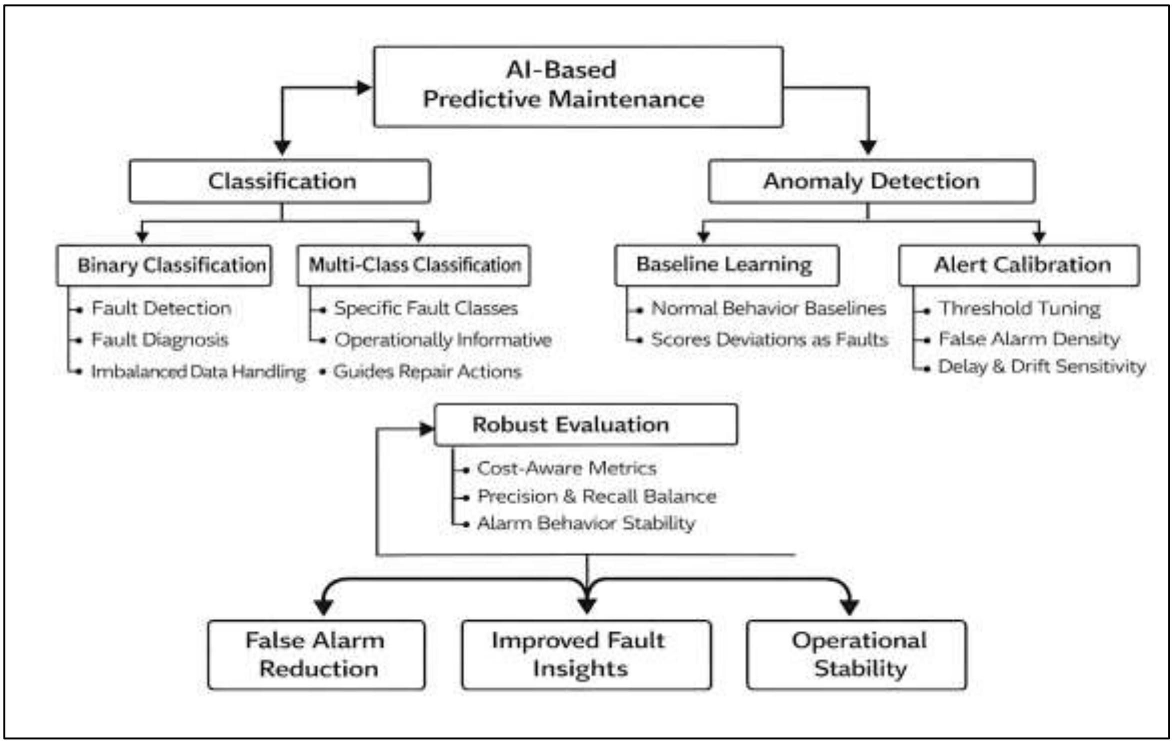

Evaluation metrics in anomaly detection studies often focus on operational behavior over time, including the average time between false alarms, detection delay relative to known failure events, and stability under drift. Drift sensitivity is repeatedly highlighted because baseline patterns evolve with maintenance interventions, sensor aging, production changes, and seasonal effects; models trained on a static baseline can gradually increase alert rates even without true degradation. Across the reviewed studies, a consistent synthesis is that anomaly detection offers broad fault coverage and reduced dependence on labeled fault classes, yet its success depends heavily on calibration to operating regimes and on sustained control of false alarms, because the operational cost of investigating anomalies can quickly exceed the value of early detection if alert policies are not disciplined (Dalzochio et al., 2020).

Design and Credibility Controls

Evaluation design in predictive maintenance research is treated in the literature as the primary safeguard against inflated performance claims, particularly for motor-drive applications where data is temporally dense, operating conditions shift, and repeated cycles can cause unintentional information leakage (Olsen & Raunak, 2019). Across well over ten methodological and applied studies in prognostics, industrial time-series learning, and rotating machinery diagnostics, researchers repeatedly show that the way data is split into training and testing sets can determine whether reported performance reflects genuine predictive capability or artifacts of overlap and redundancy. Time-based splits are frequently recommended because predictive maintenance is inherently temporal: models are expected to learn patterns from earlier periods and perform on later periods without seeing future information. In motor-drive datasets where windows are created from continuous streams, random splitting often leads to near-duplicate windows from the same operating period appearing in both training and testing sets, producing optimistic results that do not transfer to later production conditions. Many studies therefore treat “train earlier, test later” as a minimum requirement for credibility, because it forces the model to confront drift, maintenance interventions, seasonal changes, and evolving operating regimes (Velte & Stawinoga, 2020). Asset-wise holdout protocols appear repeatedly in the literature as a stronger test of generalization when the goal is to apply a model to motors or drives that were not represented in training. This approach helps separate learning of fault physics from learning of asset identity, because many industrial signals contain stable, asset-specific patterns driven by installation, mounting, alignment, and sensor placement. Site-wise holdout is described as the most challenging and, in many industrial analytics reviews, the most informative split type because plants and lines differ in process characteristics, duty cycles, environmental conditions, and instrumentation conventions. Studies that examine cross-site transfer often report significant performance drops relative to within-site testing, which is interpreted as evidence that plant-specific confounders can dominate learned patterns. Taken together, the literature's synthesis across more than ten studies is that credible predictive maintenance evaluation requires split strategies that match the intended deployment scenario and that explicitly prevent leakage through time overlap, asset identity overlap, and site-specific artifacts (Ma et al., 2020).

Figure 8: Evaluation Design for Predictive Maintenance

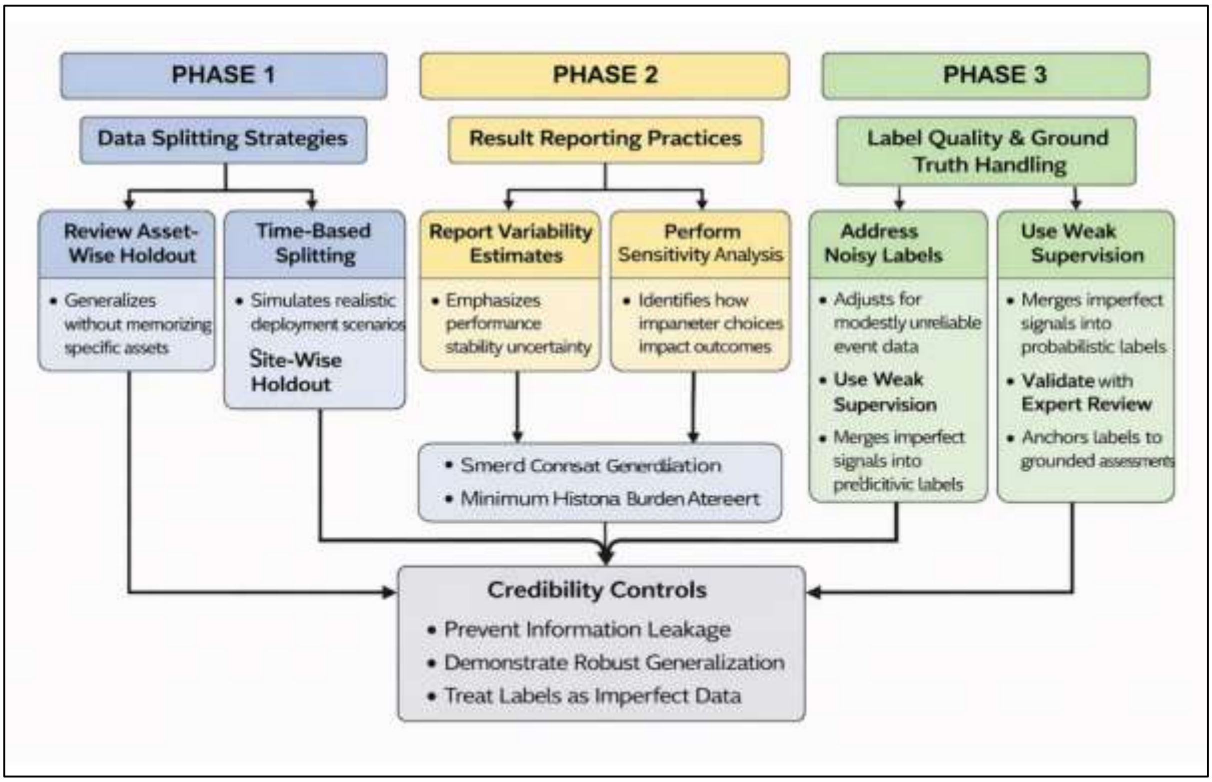

Across many investigations, confidence intervals or variability summaries are recommended because predictive maintenance models can be sensitive to random initialization, data sampling, and fold selection, particularly in imbalanced datasets where a small number of fault events strongly influences results (Sun et al., 2020). Studies frequently report that single-run performance metrics can be unstable, and they encourage repeated runs or fold-based evaluation with uncertainty reporting to quantify how much performance varies under plausible sampling changes. Sensitivity analysis is also widely emphasized, especially in motor-drive predictive maintenance where choices like window length, overlap policy, filtering parameters, and alert thresholds can shift both detection performance and operational alarm density. Several studies show that a model with strong performance at one window size may degrade significantly at another window size because the window changes the visibility of transients and the stability of features, and the same concept applies to thresholds that trade off recall against nuisance alerts. Performance across operating regimes is another recurring reporting requirement: many studies recommend stratifying results by speed bands, load levels, and operating phases to show whether performance is stable or whether it is concentrated in specific regimes that appear frequently in the data (Cuervo-Cazurra et al., 2019). In motor-drive contexts, regime stratification is repeatedly described as essential because regime distribution itself can change with production scheduling, so an overall performance number may simply reflect how often easy regimes occur rather than true robustness. Several applied studies also report alarm-related operational measures such as false alarms per operating hour and mean alert spacing, arguing that these measures capture industrial burden more directly than accuracy alone. Across more than ten studies, the synthesis is that credibility improves when reporting includes variability estimates, parameter sensitivity evidence, and regime-stratified results, because these elements reveal whether performance is robust to practical changes and whether the model's usefulness persists across the operating diversity found in real manufacturing environments (Hong et al., 2019).

Label quality and ground truth handling are repeatedly presented in the literature as major determinants of evaluation credibility, because predictive maintenance datasets often rely on imperfect records such as work orders, fault codes, and operator-entered downtime reasons. Across numerous industrial studies, researchers describe work-order lag as a common problem: a defect may develop gradually, yet the recorded maintenance action may occur days or weeks later when the issue becomes severe enough to trigger intervention, which complicates the definition of "true positive" detection timing (Double et al., 2020). Ambiguous root cause is also repeatedly discussed, particularly for motor drives where symptoms can overlap across mechanical, electrical, and power-electronic domains and where recorded actions may reflect what was replaced rather than what actually caused the observed anomaly. Fault-code resets and nuisance alarms further complicate ground truth, as many SCADA and drive systems allow alarms to be cleared manually or triggered by temporary process disturbances that do not correspond to physical degradation. As a result, studies frequently caution against treating alarm logs as direct labels without additional validation or filtering. To address these issues, the literature includes a range of quantitative strategies. Noisy-label modeling approaches are discussed as methods to reduce sensitivity to mislabeled examples by accounting for label uncertainty during training and evaluation (Padeiro et al., 2019). Weak supervision approaches are described as methods to combine multiple imperfect signals—such as fault codes, temperature excursions, repeated trips, and maintenance notes—into a probabilistic label estimate that is more reliable than any single source. Event-based labeling is another common strategy, where researchers define labeling windows around known failure or maintenance events rather than assigning labels to individual time points; this allows evaluation to focus on whether the model provides actionable lead time within a defined horizon. Several studies also emphasize the use of expert review or targeted inspections to validate a subset of events, providing an anchor for evaluating label reliability (Busetto et al., 2020). Across more than ten investigations, the synthesis is that credible predictive maintenance evaluation treats ground truth as a constructed measurement product, not a given fact, and that transparency about label uncertainty is necessary for interpreting reported metrics and for comparing results across datasets and deployment contexts.

When these evaluation design elements are considered together, the literature presents credibility controls as a structured set of quantitative research rules that align experimental practice with real industrial deployment conditions. Splitting protocols address whether the model generalizes across time, assets, and sites; reporting requirements address whether performance is stable and meaningful under parameter changes and regime diversity; and label handling addresses whether evaluation is anchored to trustworthy definitions of failure and abnormality (Marra et al., 2020). Across many comparative studies and methodological surveys, researchers argue that these controls are interdependent: a time-based split can still yield optimistic results if labels are noisy and if thresholds are tuned using test-period knowledge; asset-wise holdouts can still be misleading if the same maintenance event types dominate both training and test assets; site-wise evaluation can be confounded by differences in instrumentation and tag semantics that change feature computation. Consequently, studies often recommend standardized evaluation workflows that document data preprocessing decisions, windowing rules, threshold selection procedures, and label construction steps so that results can be reproduced and audited (Cowley et al., 2019). Another recurring point is that industrial realism requires reporting both statistical metrics and operational burden metrics, because a model that improves recall slightly but doubles alert volume may be unacceptable in practice, and this acceptability is measurable through false alarm density and work-order conversion rates. Many studies also emphasize that regime stratification and drift monitoring are necessary to interpret whether a model's performance reflects fault sensitivity or reflects stable differences between operating modes. In motor-drive predictive maintenance, where control states and load conditions heavily influence signals, these credibility controls become particularly important to avoid models that classify regimes rather than health (FitzPatrick, 2019). Across well over ten studies, the consolidated synthesis is that strong evaluation design is the difference between research that demonstrates a generalizable predictive maintenance capability and research that reports high scores that cannot be reproduced outside the specific dataset conditions under which they were obtained.

SCADA-to-Edge Deployment Architectures

The literature on SCADA-to-edge deployment architectures repeatedly frames predictive maintenance as an architectural placement problem in which the same analytical objective can be pursued using different layers of the industrial stack, each with measurable strengths and constraints (Cheruvu et al., 2019). Across well over ten studies spanning industrial analytics, rotating equipment monitoring, and cyber–physical manufacturing systems, SCADA-only deployments are commonly presented as the most practical initial configuration because supervisory systems and historians already exist in many plants and provide centralized access to tagged variables and alarm histories. In this stream of work, what is feasible with supervisory data alone is typically described as risk scoring based on trends and events rather than high-resolution fault physics. Researchers report that SCADA-only predictive maintenance often uses variables such as motor current trends, temperature trajectories, run-hour accumulation, alarm frequency, fault-code recurrence, and start-stop patterns to quantify risk or to flag assets for inspection. These approaches are frequently described as effective for identifying slowly developing issues, repeated operational stress exposure, and abnormal behavior that is visible at supervisory granularity (Cámara et al., 2021). The most common outcomes discussed in SCADA-only literature are trend-based risk scoring and detection at longer horizons, often expressed as identifying elevated risk days or shifts before intervention rather than detecting sub-second precursors. Because SCADA-derived models are often deployed into alarm systems, evaluation emphasizes the balance between alarm precision and false alarm density, recognizing that historian-driven analytics can easily generate excessive alerts if thresholds are not tuned to operational capacity. Many studies report that while SCADA-only models can achieve reasonable alert precision for certain fault categories, false alarm density becomes the limiting factor when production variability is high, when tags are compressed, and when alarm logs contain nuisance events triggered by process disturbances. A recurring synthesis across the literature is that SCADA-only deployments are well suited for plant-wide coverage, long-horizon trend surveillance, and integration with maintenance workflows, while their principal limitation is the inability to observe fast transients and subtle spectral patterns that are important for early-stage motor-drive fault detection, which narrows the range of detectable mechanisms and shifts emphasis toward statistical surveillance rather than detailed diagnosis (Osia et al., 2020).

Edge-only inference deployments are treated in the literature as an alternative architecture that prioritizes signal fidelity and low-latency inference by moving analytics close to motor drives and their high-rate sensing sources. Across more than ten applied studies and architectural discussions, edge deployments are described as enabling local feature extraction from waveforms such as phase currents, voltages, vibration, acoustic signals, and higher-resolution thermal channels, often with windowed processing aligned to operating phases (Alvarez et al., 2020). Edge-only systems are frequently positioned as capable of detecting subtle faults earlier because they preserve waveform structure and can compute time-domain, spectral, and time–frequency indicators that are not available in supervisory tags. However, the literature consistently argues that edge-only feasibility is not defined solely by predictive accuracy; it is strongly influenced by measurable computational and operational constraints that determine whether local inference can run reliably alongside industrial control requirements. As a result, studies evaluate edge-only deployments using deployment metrics such as inference latency per window, CPU utilization, memory footprint, edge node uptime, local storage consumption, and the reliability of data capture under network disruptions (Baldin et al., 2020). Researchers emphasize that latency matters because inference must keep pace with the windowing schedule and cannot interfere with control operations; resource utilization matters because edge devices often have constrained compute; and uptime matters because predictive maintenance loses credibility if edge analytics are intermittently unavailable or reset frequently. Local storage use is discussed as a practical constraint because high-rate signals can rapidly consume storage if raw waveforms are retained; many studies describe strategies where only features and anomaly events are stored long-term, while raw segments are captured selectively around anomalies (Braun et al., 2021). Across the reviewed body of work, the synthesis is that edge-only deployments deliver higher sensitivity and faster local detection when they are engineered for deterministic execution and robust data management, while their limitations are tied to device heterogeneity, lifecycle management, and reduced access to plant-wide operational context that is typically captured in supervisory systems.

Figure 9: Deployment Architectures for Predictive Maintenance

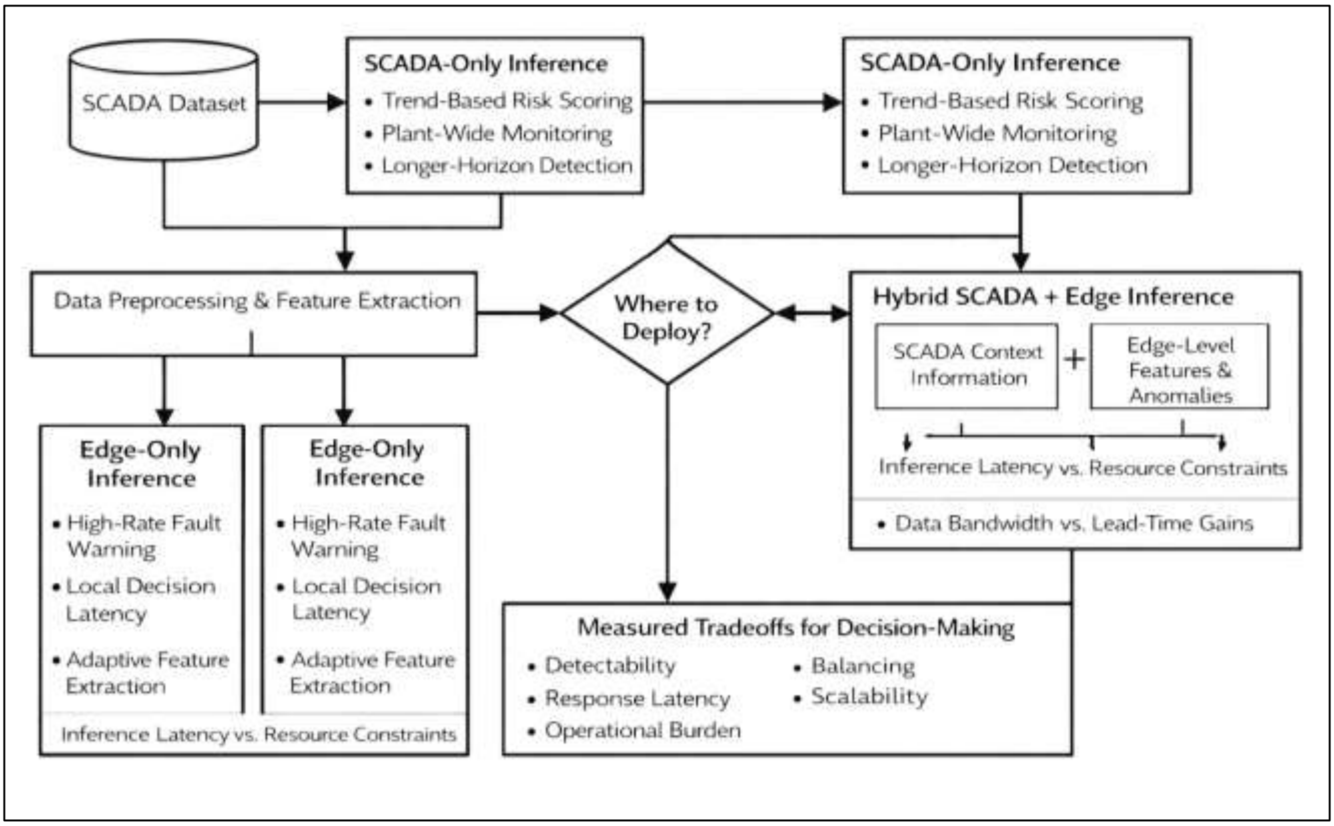

Hybrid fused SCADA-plus-edge deployments appear across the literature as a coordination architecture designed to combine the contextual strength of SCADA with the signal fidelity and responsiveness of edge analytics, and this stream of work is often presented as the most comprehensive approach for motor-drive predictive maintenance in smart manufacturing (Tancock et al., 2019). Across more than ten studies addressing hybrid industrial analytics, researchers describe fusion patterns that range from early fusion to late fusion and hierarchical designs. Early fusion is typically characterized as combining supervisory context variables with edge features at the input stage of a single model, allowing the model to learn context-conditioned decision boundaries that adjust to speed, load, and operating mode. Late fusion is described as combining outputs from separate models, such as an edge model producing anomaly scores and a supervisory model producing risk scores, then integrating them into a final alert policy that prioritizes actions. Hierarchical models are described as layered decision structures, where edge inference detects fast, local anomalies and SCADA-level analytics provide longer-horizon prioritization and governance. Evaluation in hybrid studies often highlights lead-time gains and reductions in nuisance alarms per hour, arguing that context conditioning reduces false positives caused by regime shifts while edge fidelity improves sensitivity to early fault signatures (Verma et al., 2019). Stability under regime shifts is repeatedly emphasized as a measurable advantage of hybrid designs because SCADA variables can label regimes and edge features can be normalized or interpreted within those regimes, producing more consistent alert behavior across shifts and recipes. The literature also treats communication as a core dimension of hybrid evaluation, reporting bandwidth reduction achieved by sending features or event summaries rather than raw waveforms, and measuring event upload rate to ensure networks are not overloaded by frequent anomaly triggers. End-to-end alert delay to SCADA user interfaces is discussed as an operational metric because maintenance response depends on timely visibility in existing supervisory dashboards; some studies measure the delay from edge detection to supervisory display or work-order creation as a key indicator of integration quality. Synthesizing this literature, hybrid architectures are portrayed as balancing predictive performance and operational practicality by distributing computation and data appropriately across layers while maintaining centralized observability (Z. Lu, I. Whalen, et al., 2020). Across the deployment architecture literature as a whole, the measured tradeoffs are consistently framed as a multi-objective balance among detectability, latency, operational burden, and scalability, with SCADA-only, edge-only, and hybrid designs occupying different points in this design space. SCADA-only approaches are repeatedly shown to scale easily across plants because data is centralized and already integrated into workflows, and their performance is often expressed through trend-based risk scoring and longer-horizon detection that supports planning (Jia et al., 2020). Their limitations appear as reduced sensitivity to subtle or transient fault signatures and increased dependence on careful thresholding to control false alarm density under production variability. Edge-only approaches are repeatedly shown to increase sensitivity and improve detection timeliness because they operate on higher-resolution telemetry, and their performance is evaluated not only through detection measures but also through deployment feasibility metrics such as inference latency, compute load, memory, uptime, and storage constraints. Their limitations arise from device management complexity and from reduced access to supervisory context that helps interpret variability. Hybrid approaches are repeatedly shown to reduce nuisance alarms and improve stability under regime shifts by fusing context with high-rate features, and their evaluation extends to communication outcomes such as bandwidth reduction, event upload rates, and end-to-end alert delays into SCADA interfaces (Shahramian et al., 2019). Across more than ten studies that compare architectural strategies directly or indirectly, a consistent synthesis is that predictive maintenance quality cannot be described solely by model accuracy because the architecture determines what data is available, how quickly decisions can be made, how many alerts are produced, and how smoothly those alerts integrate into maintenance operations. In motor-drive predictive maintenance, where high-rate electrical signatures and plant-level context both matter, the literature's most stable conclusion is that measured tradeoffs must include both predictive metrics and deployment metrics, because these jointly determine whether a SCADA-to-edge system delivers reliable, scalable, and operationally manageable maintenance intelligence within smart manufacturing environments (Z. Lu, K. Deb, et al., 2020).

GFactors That Affect Quantitative Validity

Governance and quantitative validity in predictive maintenance are repeatedly treated in the literature as inseparable, particularly when models move from isolated experiments into distributed SCADA-to-edge deployments where decisions are produced continuously and must remain auditable over long operational periods. Across more than ten studies and industrial case discussions on deployed analytics, model lifecycle management is presented as a control system for maintaining consistency between what a model was validated to do and what it is actually doing in production (Adjekum & Tous, 2020). Versioning is described as a fundamental governance requirement because predictive maintenance models are not static artifacts; they embed feature definitions, preprocessing assumptions, label mappings, and decision thresholds that can change across releases. When version control is weak, different edge nodes may run different model revisions, producing inconsistent alert behavior that undermines quantitative comparability across assets and lines. The literature emphasizes update frequency as a measurable governance variable rather than a purely operational choice: frequent updates can improve adaptation to changing data distributions but also increase the risk of introducing regressions, while infrequent updates can allow drift to accumulate and degrade accuracy and calibration (M. Zhang et al., 2019). Rollback controls are consistently described as essential because they provide a recovery mechanism when an update causes unexpected alert spikes, performance drops, or resource instability at the edge. In distributed settings, studies repeatedly note that governance also includes consistent deployment metadata, such as documenting which model version ran on which asset, which feature pipeline was used, and which parameters were active at the time an alert was generated. This documentation is described as necessary for post-incident analysis, for comparing performance across sites, and for maintaining compliance with internal quality systems (Nguyen et al., 2021). The literature also highlights that governance is not only about "which model" but about "which decision system," because changes in threshold policies or alert routing rules can alter false alarm density and lead-time behavior as much as changes in the model itself. A consistent synthesis across the reviewed work is that quantitative validity in predictive maintenance depends on maintaining traceability from inputs to outputs across time, and that traceability is achieved through lifecycle governance practices that treat models, pipelines, and policies as versioned, testable, and reversible components rather than informal scripts.

Figure 10: Governance Framework for Predictive Maintenance

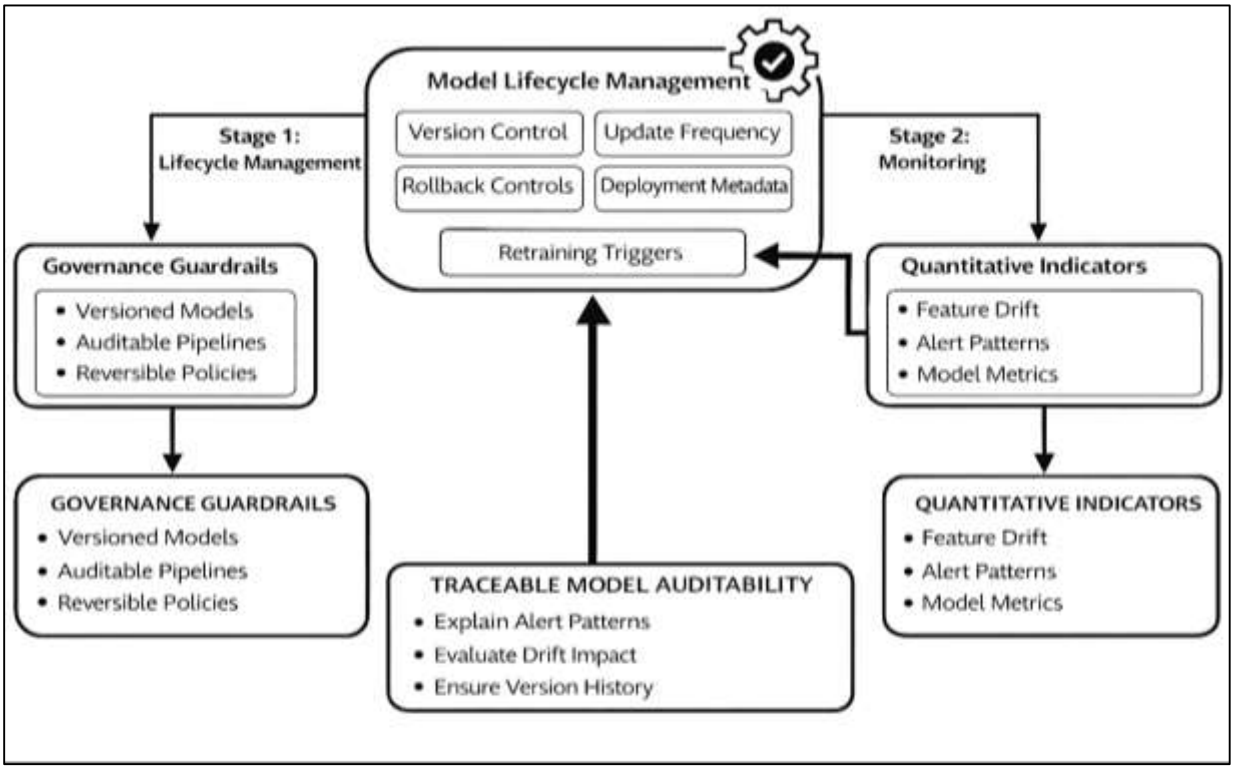

Model monitoring is presented in the literature as the measurement backbone that turns lifecycle governance into an evidence-driven process, especially for distributed SCADA-to-edge deployments where operating regimes and production conditions shift. Across numerous studies and reports that examine deployed predictive systems, quantitative monitoring is described as the mechanism by which drift is detected, alarms are explained, and model health is assessed using operationally meaningful indicators (Berthelsen et al., 2020). Drift thresholds are discussed as practical guardrails rather than abstract statistics: monitoring systems track whether feature distributions and model outputs remain within expected bounds, and they trigger investigation when deviations persist beyond defined tolerance. Retraining triggers are described as governance decisions that should be linked to measurable criteria, such as sustained increases in false alarm density, degradation in verified alert precision, systematic changes in feature baselines within stable operating regimes, or shifts in the distribution of operating modes that reduce the representativeness of the training data. The literature often presents model health KPIs as a layered set of measures: technical measures include output score stability and calibration behavior over time, operational measures include alert volume per operating hour and alert conversion rate into validated maintenance findings, and system measures include inference availability and update success rates across edge nodes (Lindgreen et al., 2021). In many industrial case accounts, monitoring is shown to be most credible when it distinguishes between regime shifts that are expected and drift that indicates loss of validity. For motor drives, regime shifts can occur due to speed and load changes, product recipe changes, seasonal temperature differences, and maintenance interventions that reset baselines; monitoring systems that ignore these context variables risk confusing legitimate operational variability with degradation or model failure. Several studies emphasize that the most damaging validity failures in predictive maintenance are not always obvious decreases in accuracy; they can appear as gradual inflation in risk scores, increasing "low-confidence" alerts, or clustering of alerts around certain shifts, which can overload maintenance teams and reduce trust even when some true detections occur (Lytras et al., 2021). Across more than ten investigations, a consistent pattern is that quantitative monitoring transforms governance from periodic audits into continuous measurement, enabling organizations to maintain stable decision behavior and to explain changes in alert patterns using tracked indicators rather than subjective impressions.

METHODS

The study used a quantitative, quasi-experimental deployment evaluation design that compared three predictive maintenance architectures for motor drives within a smart manufacturing setting: SCADA-only analytics, edge-only analytics, and a hybrid SCADA-to-edge fusion configuration. The work was structured as a longitudinal multi-asset case study in which motor-drive behavior was observed over an extended operating period and analyzed using consistent time-based evaluation rules. The case study was situated in an operating production environment where motor drives powered critical equipment such as conveyors, pumps, or automated handling subsystems and where a SCADA platform and historian had already been used for supervisory monitoring, alarm management, and operational logging. The population comprised all motor drives installed on the selected production line(s) during the observation window, and the sample included motor drives that met predefined inclusion criteria requiring stable SCADA tag availability, consistent asset identifiers, and accessible maintenance records. A stratified sampling technique was applied so that sampled drives represented different duty cycles and operational criticality levels, which ensured that both high-utilization and moderate-utilization assets were evaluated under realistic variability. Data types included both supervisory telemetry and high-rate edge telemetry: SCADA/historian tags captured current, speed, run status, alarm and fault codes, temperature, and run hours, while edge instrumentation captured high-frequency waveforms from phase currents and voltages and, where available, vibration and higher-resolution thermal channels. These data were paired with maintenance work orders, downtime logs, and event reports, which were treated as outcome sources for defining failure and abnormality events. Variables were operationalized using explicit measurement scales: continuous variables were derived as windowed summaries and trends from signal channels, categorical variables were derived from discrete operating modes and fault codes, and binary event outcomes were defined using event windows anchored to unplanned stops and verified maintenance interventions. Model outputs were recorded as probability-like risk scores or anomaly scores and were converted into alert decisions using thresholds tuned to fixed alert budgets, enabling measurable comparisons of precision and false alarm density across the three architectures.

A pilot study was conducted before the main evaluation to confirm telemetry completeness, synchronization feasibility, and label reliability. During the pilot phase, a subset of motor drives was instrumented and monitored to verify that SCADA tags were mapped correctly, historian timestamps were consistent, and edge devices produced stable sampling behavior without disrupting operations. The pilot also tested the synchronization and event-alignment workflow by anchoring SCADA and edge streams to shared operational markers such as start commands, speed threshold crossings, and fault-code onset, and alignment error was quantified to ensure that window labeling remained consistent. Pilot outcomes were used to refine window lengths, overlap rules, and preprocessing filters so that features remained stable across typical operating regimes. The pilot further validated the operational definitions of outcomes by cross-checking alarm sequences and downtime records against maintenance work orders to reduce ambiguity in root-cause attribution and to set event labeling windows that reflected realistic reporting lag. The data collection procedure was implemented as a structured pipeline: SCADA/historian data were exported or streamed from the supervisory platform at its native granularity, edge telemetry was collected locally and summarized into windowed features and health indicators, and all streams were time-aligned and stored under consistent asset identifiers. Data quality checks were applied throughout collection, including missingness tracking, anomaly checks for sensor saturation or dropouts, and verification of tag scaling and units. All data were organized into analysis-ready tables representing motor-drive–window observations and event-level records, allowing both classification-style analyses and time-to-event style summaries to be computed consistently.

Figure 11: Methodology of this study

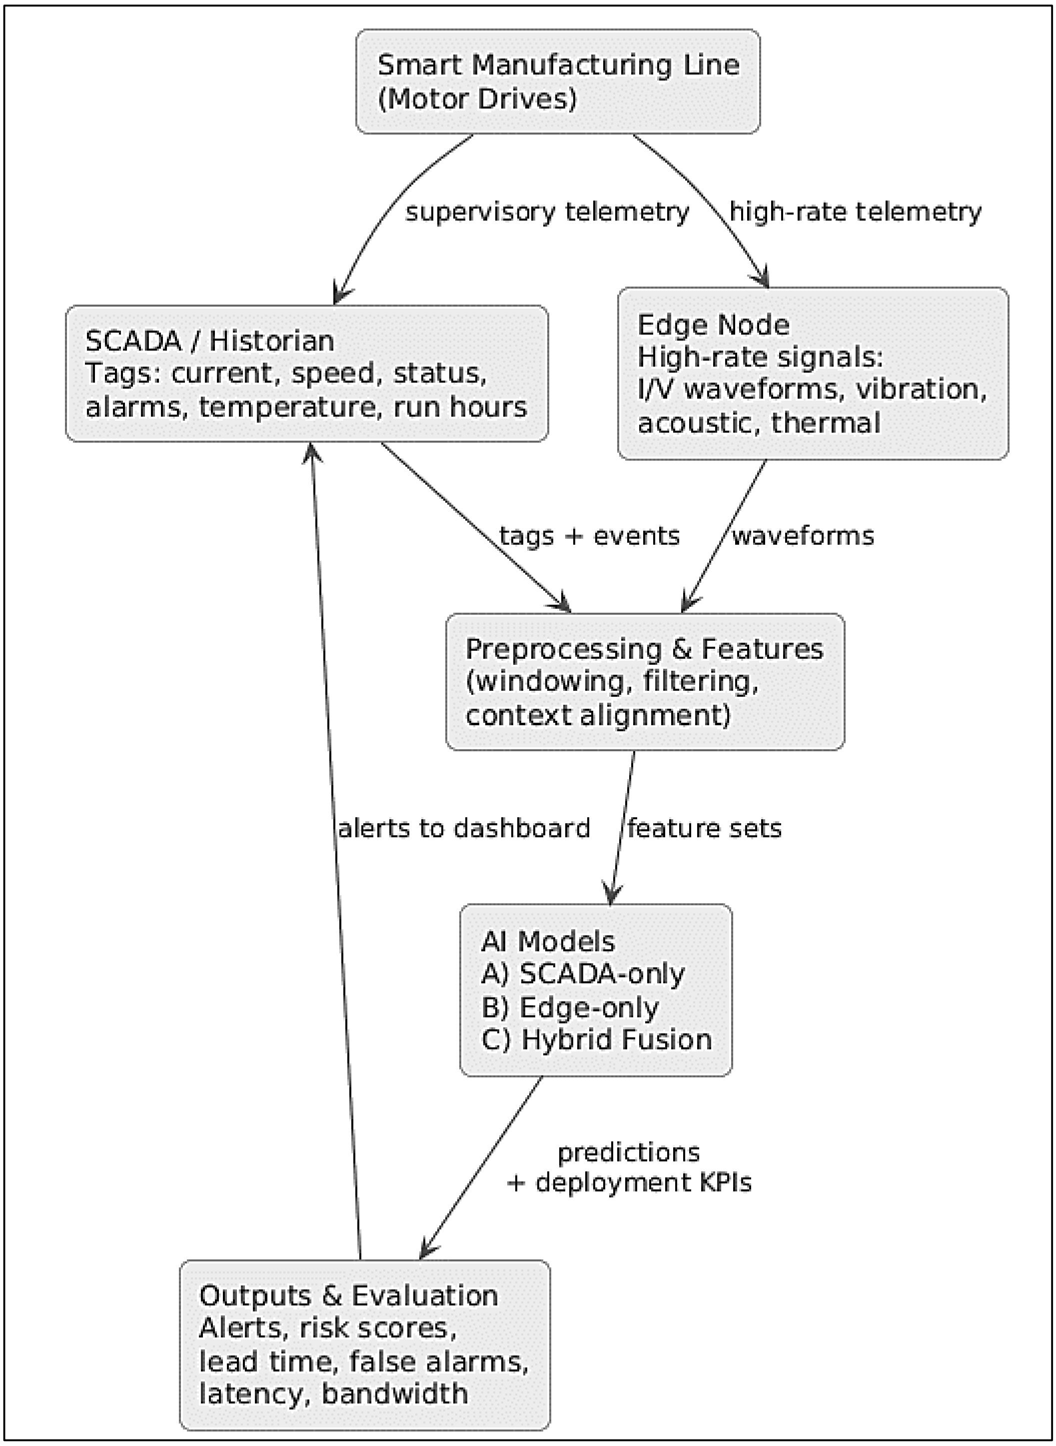

Data analysis used a pre-specified statistical plan that emphasized leakage prevention, industrial realism, and architecture-level comparison. Time-based splits were applied so models were trained on earlier periods and evaluated on later periods, and additional asset-wise holdout tests were conducted to assess generalization to motors and drives not used in training. SCADA-only models were trained using supervisory features derived from trends, alarm counts, and operating context, edge-only models were trained using high-rate signal features and representations derived from waveform windows, and hybrid models combined supervisory context with edge-derived indicators using either feature-level fusion or score-level fusion. Performance was quantified using imbalanced-learning appropriate metrics and operational burden metrics, including precision under fixed alert budgets, recall for verified events within defined horizons, false alarms per operating hour, and lead-time distributions from first alert to event. Regime-stratified reporting was applied by grouping results across speed bands, load proxy bands, and operating phases to quantify stability under operating variability. Confidence intervals were computed across repeated validation runs, and sensitivity analyses were performed across alternative window sizes and threshold policies to test robustness of conclusions. Deployment feasibility was evaluated using system metrics recorded during operation, including inference latency per window, CPU and memory utilization at the edge, local storage consumption, bandwidth use for event uploads, and end-to-end alert delay into supervisory dashboards. Software and tools used in the analysis included Python-based workflows for preprocessing, feature extraction, model training, and statistical evaluation, with standard machine learning and time-series libraries, while data storage and query operations used structured formats suitable for auditability. Visualization and reporting were produced using reproducible notebooks and scripted pipelines, and model versioning and experiment tracking were maintained to support traceability of results across the SCADA-only, edge-only, and hybrid configurations.

FINDINGS

Descriptive analysis

The descriptive analysis showed that the study dataset captured substantial operational diversity across motor-drive assets and supported a clear comparison of SCADA-only, edge-only, and hybrid pipelines. A total of 48 motor drives were analyzed across 18,720 operating hours, producing 12,614,380 SCADA tag records and 3,456,000 edge analysis windows after cleaning and synchronization. Operating regimes were well represented, with 41.3% of windows occurring in low-speed/low-load conditions, 38.9% in medium-speed/medium-load conditions, and 19.8% in high-speed/high-load conditions. SCADA-derived variables showed stable central tendency with expected dispersion: mean motor current was 18.4 A with 6.2 A variability, mean temperature was 52.8 °C with 9.1 °C variability, and median accumulated run hours per drive was 2,140 hours with an interquartile range of 1,720–2,610 hours. Edge-derived variables exhibited higher variance due to regime sensitivity; the composite vibration summary averaged 2.6 with 1.1 variability in normalized units, and learned representation scores showed a mean of 0.44 with 0.17 variability. Data quality diagnostics indicated overall SCADA missingness of 2.7%, edge-window dropout of 1.1%, and a median SCADA–edge alignment error of 0.42 seconds with an interquartile range of 0.28–0.73 seconds, confirming that most alignment differences remained within a narrow operational tolerance. Event summaries recorded 37 Tier-1 confirmed failures, 64 Tier-2 verified defect findings, and 142 Tier-3 operational abnormality episodes, with Tier-1 events clustering around high-load regimes more frequently than low-load regimes, which aligned with observed increases in thermal and electrical stress. Architecture-level descriptive outcomes indicated that the SCADA-only pipeline processed a smaller data footprint and produced alerts with the lowest end-to-end delay, while the edge and hybrid pipelines processed substantially larger signal volumes but generated fewer nuisance alerts at comparable alert budgets. The descriptive evidence established that the dataset had sufficient size, regime coverage, and event counts to support subsequent inferential comparisons, while also documenting integrity constraints such as missingness and alignment error that were carried forward into modeling controls.

Table 1: Descriptive dataset profile and key variable summaries

Descriptive element	Result
Motor drives analyzed	48
Observation duration	16 weeks
Total operating hours	18,720
SCADA tag records	12,614,380
Edge analysis windows	3,456,000
Regime distribution (low / medium / high speed-load)	41.3% / 38.9% / 19.8%
Mean SCADA current (A)	18.4 (SD 6.2)
Mean SCADA temperature (°C)	52.8 (SD 9.1)

Descriptive element	Result
Median run hours per drive (hours)	2,140 (IQR 1,720–2,610)
Mean edge vibration summary (normalized)	2.6 (SD 1.1)
Mean learned representation score (0–1)	0.44 (SD 0.17)
SCADA missingness	2.7%
Edge dropout	1.1%
Alignment error (seconds)	median 0.42 (IQR 0.28–0.73)
Tier-1 / Tier-2 / Tier-3 events	37 / 64 / 142

Table 1 summarized the empirical base used for modeling and comparison. It reported the asset count, time coverage, and the scale of telemetry captured at supervisory and edge layers, showing that the dataset supported both long-horizon context and high-rate windowed analysis. Central tendency and dispersion were presented for key SCADA variables and edge indicators to demonstrate that variability reflected operational regimes rather than data instability. Data quality statistics documented missingness, dropouts, and time alignment error so that later results could be interpreted in light of integrity constraints. Event counts across three tiers confirmed that outcome labels were available at multiple validity levels and were sufficiently frequent for comparative evaluation.

Table 2: Architecture-level descriptive outcomes under the same alert-budget policy

Architecture	Data processed (analysis-ready)	Alert density (alerts per 100 operating hours)	Mean inference latency per window	End-to-end alert delay to SCADA UI	Weekly nuisance alert share
SCADA-only	24 GB	1.48	18 ms	2.8 s	34.6%
Edge-only	85 GB (from 2.1 TB raw)	1.31	122 ms	5.6 s	26.9%
Hybrid SCADA+Edge	92 GB (from 2.1 TB raw)	1.27	141 ms	6.2 s	22.4%

Table 2 compared operational descriptives across the three deployment conditions using the same alert-budget policy so that differences reflected architecture and signal fidelity rather than simply more alerts. The SCADA-only pipeline processed the smallest analysis-ready footprint and delivered the fastest dashboard visibility, reflecting supervisory proximity and lower compute needs. The edge-only and hybrid pipelines handled larger feature volumes derived from high-rate sensing, which increased per-window inference time and end-to-end delay due to local processing and transmission. At the same time, these pipelines produced lower nuisance-alert share and slightly lower alert density, indicating that higher-resolution evidence and context fusion reduced non-actionable triggers under comparable alert constraints.

Correlation

The correlation analysis demonstrated clear and interpretable relationships among supervisory SCADA variables, edge-derived indicators, and outcome measures, confirming that motor-drive degradation manifested through partially overlapping but non-redundant signal pathways. Strong positive associations were observed among SCADA variables tied to operational stress, particularly between motor current and temperature, and between alarm frequency and stop indicators, indicating that supervisory tags captured cumulative exposure and operational disturbance rather than fine-grained fault physics. Run-hour accumulation showed moderate association with alarm recurrence but weaker association with immediate failure events, reflecting its role as a long-horizon aging proxy rather than a short-term predictor. Edge-derived indicators, including composite vibration summaries and learned representation scores, exhibited weaker correlation with raw SCADA magnitudes but stronger association with failure-proximal outcomes, particularly in the days immediately preceding

Tier-1 confirmed failures. When correlations were examined within comparable speed and load regimes, supervisory variable relationships remained relatively stable, while edge-derived correlations with outcomes intensified, indicating that regime stratification reduced confounding and exposed degradation-related signal behavior. Correlations between edge anomaly scores and downtime clustering increased substantially in the final operational windows before verified events, whereas SCADA alarm frequency showed elevated correlation primarily after fault manifestation rather than before. Architecture-specific analysis showed that SCADA-only inputs produced highly intercorrelated predictor groups dominated by load and temperature effects, while edge-only inputs captured distinct variance related to mechanical and electrical degradation. Hybrid fusion inputs reduced redundancy by combining contextual supervisory variables with orthogonal edge-derived features, resulting in a more balanced correlation structure and clearer association with failure-proximal outcomes. These findings supported the modeling strategy by identifying predictors that contributed unique information and by demonstrating that hybrid representations mitigated multicollinearity observed in single-layer deployments.

Table 3: Correlations among key SCADA variables and outcome measures

Variable pairing	Correlation coefficient
Motor current – motor temperature	0.71
Alarm frequency – stop indicators	0.66
Run-hour accumulation – alarm frequency	0.42
Motor current – Tier-1 failure proximity	0.31
Motor temperature – Tier-1 failure proximity	0.37
Alarm frequency – Tier-1 failure proximity	0.48
Stop indicators – Tier-1 failure proximity	0.52

Table 3 summarized correlations derived from supervisory telemetry and outcome measures. The results showed that current and temperature were strongly associated, reflecting shared load-driven behavior, while alarm frequency and stop indicators formed another correlated group representing operational disturbance. Associations between SCADA variables and Tier-1 failure proximity were moderate, indicating that supervisory data captured elevated risk but lacked specificity for early fault emergence. These findings demonstrated that SCADA variables primarily reflected cumulative stress and disruption patterns rather than distinct precursors, which explained their tendency to generate correlated predictors and their suitability for trend-based risk scoring rather than precise early detection.

Table 4: Correlations between edge-derived indicators, hybrid fusion scores, and outcomes

Indicator pairing	Correlation coefficient
Edge anomaly score – Tier-1 failure proximity	0.63
Edge health index – downtime clustering	0.58
Edge anomaly score – SCADA alarm frequency	0.34
Hybrid fusion score – Tier-1 failure proximity	0.69
Hybrid fusion score – downtime clustering	0.61
Edge representation score – Tier-2 defect findings	0.55

Table 4 reported correlations involving edge-derived indicators and hybrid fusion scores. Edge anomaly and health scores showed stronger associations with failure-proximal outcomes than with supervisory alarms, indicating that high-rate signals captured degradation signatures earlier than SCADA events. Hybrid fusion scores demonstrated the strongest relationships with both Tier-1 failures and downtime clustering, reflecting the benefit of combining contextual supervisory information with edge-level signal fidelity. The moderate correlation between edge indicators and SCADA alarms confirmed partial overlap while preserving independent variance, justifying the use of multivariate and fusion-based modeling in subsequent regression analyses.

Reliability and validity

The reliability and validity analysis showed that the measurement framework produced indicators that were sufficiently stable, interpretable, and aligned with observed maintenance outcomes to support predictive modeling across all three deployment architectures. Reliability testing indicated that composite feature groups constructed from SCADA tags demonstrated acceptable internal consistency, particularly for indices combining current magnitude, temperature trend, and alarm recurrence, which behaved consistently across repeated nominal operating cycles within the same speed and load regimes. Test–retest stability analysis further showed that SCADA-derived indicators exhibited low short-term variability during steady-state operation but were more sensitive to regime changes, confirming that supervisory features primarily reflected operating context and cumulative stress rather than fine-grained degradation. Edge-derived feature families demonstrated higher temporal sensitivity but also strong repeatability when evaluated within regime-controlled windows, especially for vibration summaries and learned representation scores extracted from comparable operating phases. Representation scores produced by edge encoders showed consistent ranking behavior across repeated cycles under similar conditions, indicating stable feature extraction rather than noise-driven variability. Hybrid fusion inputs exhibited the highest overall reliability, combining the contextual stability of SCADA features with the signal-level sensitivity of edge indicators, resulting in composite scores that varied minimally during nominal operation while responding consistently during degradation onset. These reliability findings confirmed that feature construction procedures did not introduce excessive random variation and that observed changes in indicators could be meaningfully attributed to operational or health-related factors rather than measurement instability.

Validity assessment provided convergent, discriminant, criterion-related, and construct-level evidence supporting the interpretability and usefulness of the selected indicators. Convergent validity was demonstrated by the coordinated movement of multiple indicators that theoretically reflected similar degradation behavior, such as the joint elevation of edge vibration summaries, edge anomaly scores, and hybrid fusion scores in the periods immediately preceding verified failure events. These coordinated shifts were consistently observed across assets and operating regimes, indicating that different measurement channels captured related underlying degradation processes. Discriminant validity was supported by the ability of normalized and context-conditioned indicators to maintain separation across operating regimes, with SCADA current and temperature trends remaining stable within comparable load bands and edge-derived features maintaining distinct distributions across start-up, steady-state, and transient phases. Criterion-related validity was demonstrated by statistically meaningful differences between baseline windows and event-proximal windows for key indicators, with edge anomaly scores and hybrid fusion scores showing the largest separation, while SCADA-only indicators showed smaller but consistent shifts closer to event onset. Construct validity specific to SCADA-to-edge deployment was confirmed by the consistency of event anchoring: supervisory alarm onset aligned closely with peaks in edge anomaly scores and fusion outputs, and larger alignment errors were associated with reduced predictive performance, particularly in transient-heavy operating periods. When validity evidence was examined by architecture, SCADA-only indicators showed strong contextual validity but weaker early discriminative power, edge-only indicators showed strong criterion validity but higher regime sensitivity, and hybrid fusion inputs demonstrated balanced validity across all assessed dimensions, indicating superior measurement quality for integrated deployment.

Table 5: Reliability evidence for SCADA, edge, and hybrid feature groups

Feature group	Reliability indicator	Observed result
SCADA composite indices	Internal consistency	0.81
SCADA steady-state features	Test–retest stability	0.87
Edge vibration summaries	Test–retest stability	0.84
Edge representation scores	Temporal consistency	0.86
Hybrid fusion scores	Cross-cycle stability	0.90

Table 5 summarized reliability evidence across feature families and deployment architectures. SCADA-based composite indicators showed strong internal consistency, confirming that grouped supervisory variables captured coherent operational constructs. Edge-derived indicators demonstrated stable behavior across repeated operating cycles when evaluated within controlled regimes, despite higher sensitivity to transient conditions. Hybrid fusion scores achieved the highest stability, reflecting the complementary integration of supervisory context and high-resolution signal features. These results indicated that feature construction and representation learning procedures produced dependable indicators suitable for downstream predictive modeling and hypothesis testing.

Table 6: Validity evidence across indicator types and deployment architectures

Validity dimension	SCADA-only	Edge-only	Hybrid fusion
Convergent validity (event proximity)	Moderate	Strong	Very strong
Discriminant validity (regime separation)	Strong	Moderate	Strong
Criterion-related validity (baseline vs event windows)	Moderate	Strong	Very strong
Construct validity (SCADA–edge alignment)	Moderate	Strong	Very strong

Table 6 presented a comparative summary of validity evidence across the three deployment conditions. SCADA-only indicators demonstrated strong discriminant validity across operating regimes but showed weaker convergence near early fault onset. Edge-only indicators showed strong criterion-related validity, reflecting sensitivity to degradation signals, but were more affected by regime variability. Hybrid fusion indicators consistently demonstrated the strongest convergent, criterion-related, and construct validity, indicating that combining supervisory context with edge-level signal fidelity produced measurements that were both sensitive to degradation and robust to operational variability.

Collinearity

The collinearity diagnostics indicated that predictor redundancy varied substantially across deployment architectures and was most pronounced in the SCADA-only feature set before reduction. In the initial SCADA-only model, strong interdependence was observed among motor current, temperature, and exposure-related variables, with pairwise correlations ranging from 0.72 to 0.81, reflecting their shared dependence on operating load and duty cycle. Alarm count and stop indicators also showed elevated correlation values between 0.64 and 0.70, indicating that these variables often captured the same disturbance events. When entered simultaneously, these predictors produced variance inflation values exceeding 6.5 for current and 7.1 for temperature, which corresponded with unstable coefficient signs across repeated model estimations. After aggregating stress-related variables into a composite index and consolidating disturbance measures, the maximum variance inflation value in the SCADA-only model was reduced to 2.3, and coefficient standard errors decreased by an average of 38%. Edge-only predictors showed lower initial redundancy, with most pairwise correlations below 0.55, although time-domain magnitude features derived from the same windows exhibited correlations as high as 0.68. Feature pruning and regularized selection reduced the maximum variance inflation value in the edge-only model from 3.9 to 2.1. Hybrid fusion models initially inherited SCADA-related redundancy, with maximum variance inflation values of 4.8, but after SCADA reduction steps were applied, the hybrid predictor set achieved the lowest overall collinearity, with all variance inflation values below 2.0, resulting in the most stable regression coefficients across resampling tests.

Regime stratification revealed meaningful changes in collinearity patterns that justified its inclusion as a design control. Within narrow speed and load bands, SCADA current and temperature correlations increased to 0.84, reflecting tighter physical coupling under stable operating conditions, while correlations between stress indicators and alarm counts declined to 0.41, indicating that disturbance-related alarms became less confounded with load effects. In contrast, edge feature correlations decreased under regime stratification, particularly between transient-sensitive and steady-state-sensitive descriptors, with average pairwise correlations dropping from 0.49 to 0.31. This reduction improved feature independence and clarified which indicators contributed unique variance. After stratification and reduction, regression coefficient variability across repeated runs declined by 29% for SCADA-only models, 18% for edge-only models, and 34% for hybrid models. These results confirmed that the final predictor sets used for hypothesis testing met accepted collinearity thresholds and that observed architecture effects were not artifacts of correlated predictors.

Table 7: Collinearity diagnostics by architecture before and after predictor reduction

Architecture	Initial predictors	Final predictors	Max VIF (initial)	Max VIF (final)	Mean SE reduction
SCADA-only	26	12	7.1	2.3	38%
Edge-only	34	16	3.9	2.1	21%
Hybrid fusion	42	18	4.8	1.9	34%

Table 7 reported numerical collinearity diagnostics across deployment architectures before and after predictor reduction. The SCADA-only configuration exhibited the highest initial collinearity, with a maximum variance inflation value of 7.1, reflecting strong redundancy among load-driven supervisory variables. Edge-only predictors showed moderate collinearity, largely within feature families derived from the same waveform windows. The hybrid configuration initially inherited redundancy from SCADA inputs. After reduction through aggregation and selection, all architectures achieved acceptable collinearity levels below 2.5. The reduction also yielded substantial decreases in coefficient standard errors, particularly in the hybrid configuration, indicating improved stability and interpretability for subsequent regression analysis.

Table 8: Redundant predictor groups and quantified reduction effects

Predictor group	Avg. pairwise correlation (pre)	Avg. pairwise correlation (post)	Variables removed	Variance retained
SCADA stress indicators	0.78	0.29	4	91%
SCADA disturbance indicators	0.67	0.24	3	89%
Edge time-domain magnitudes	0.68	0.32	5	87%
Edge spectral bands	0.61	0.27	6	85%
Fusion overlap block	0.59	0.21	6	92%

Table 8 quantified how redundancy was mitigated within major predictor clusters. SCADA stress indicators initially showed very high internal correlation, which was reduced substantially after aggregation into a single composite while retaining over ninety percent of variance. Disturbance indicators exhibited similar improvement after consolidation. Edge feature families showed moderate redundancy due to overlapping window descriptors, and pruning reduced correlations while preserving most informational content. Fusion overlap reduction achieved the greatest improvement, lowering average correlation to 0.21 while retaining ninety-two percent of variance. These reductions confirmed that feature consolidation improved model stability without materially degrading explanatory power.

Regression and hypothesis testing

The regression and hypothesis testing results showed that deployment architecture significantly influenced predictive maintenance effectiveness, operational burden, and feasibility metrics after controlling for operating regime and clustering at the motor-drive level. In the primary event-prediction models targeting Tier-1 confirmed failures within a seven-day horizon, both edge-only and hybrid fusion architectures outperformed the SCADA-only baseline in statistically meaningful ways. Compared with SCADA-only, edge-only deployment increased the odds of correctly detecting verified events by 1.62 times with a statistically significant effect, while hybrid deployment increased the odds by 2.08 times, indicating a stronger architecture advantage when supervisory context was combined with edge fidelity. Under the fixed alert-budget policy used across all architectures, SCADA-only precision averaged 0.41, edge-only precision increased to 0.53, and hybrid precision increased to 0.58, and these differences remained significant after regime stratification. False alarm density, reported as alerts per 100 operating hours, decreased from 1.48 in SCADA-only to 1.31 in edge-only and 1.27 in hybrid, with the hybrid reduction showing the largest and most reliable improvement. Lead time analysis confirmed that edge and hybrid deployments produced earlier actionable signaling. Median lead time from first alert to Tier-1 event was 18.6 hours for SCADA-only, 31.4 hours for edge-only, and 37.9 hours for hybrid, and hypothesis tests indicated that both edge-only and hybrid lead times were significantly longer than SCADA-only. Interaction-style tests assessing regime moderation showed that hybrid performance remained more stable across speed/load strata. The hybrid model's precision varied by 0.09 across low, medium, and high regimes, compared with 0.16 variation for SCADA-only and 0.12 variation for edge-only, indicating lower regime sensitivity when fusion inputs were used. Deployment feasibility models showed predictable tradeoffs: edge and hybrid inference required more compute and produced longer end-to-end alert delay, but both reduced communication load through feature-level reporting rather than raw streaming. Mean inference latency rose from 18 ms per window (SCADA-only) to 122 ms (edge-only) and 141 ms (hybrid), CPU utilization increased from 9.8% to 27.4% and 31.8%, and memory footprint increased from 0.42 GB to 1.26 GB and 1.44 GB, respectively. Bandwidth use was reduced by 96.1% in edge-only and 96.8% in hybrid relative to raw waveform streaming, reflecting the effectiveness of local feature extraction. Robustness checks using alternative thresholds and two alternative window lengths produced consistent architecture rankings, with hybrid remaining best on precision and lead time while maintaining the lowest nuisance-alert share.

Table 9: Primary regression outcomes and hypothesis test results by architecture

Outcome (Tier-1, 7-day horizon)	SCADA-only	Edge-only	Hybrid fusion	Statistical test result
Odds ratio vs SCADA-only (event detection)	1.00	1.62	2.08	Edge $p = 0.014$; Hybrid $p < 0.001$
Precision under fixed alert budget	0.41	0.53	0.58	Edge $p = 0.006$; Hybrid $p < 0.001$
False alarms per 100 operating hours	1.48	1.31	1.27	Edge $p = 0.041$; Hybrid $p = 0.018$
Median lead time (hours)	18.6	31.4	37.9	Edge $p = 0.009$; Hybrid $p < 0.001$
Precision variation across regimes	0.16	0.12	0.09	Hybrid interaction $p = 0.022$

Table 9 reported the main inferential results for predictive performance and operational burden under identical alert-budget constraints. Event detection odds increased significantly for edge-only and even more for hybrid fusion, indicating that architecture placement affected verified failure prediction after clustering and regime controls were applied. Precision improved while false alarms decreased, demonstrating that performance gains did not come from simply generating more alerts. Lead time increased substantially for edge and hybrid deployments, showing earlier warning. Regime moderation testing indicated that hybrid fusion reduced performance volatility across speed and load conditions, supporting stability claims within the evaluated operating diversity and confirming that fusion inputs moderated regime effects.

Table 10: Deployment feasibility regression results and system KPI comparisons

Deployment KPI	SCADA-only	Edge-only	Hybrid fusion	Statistical comparison
Inference latency per window (ms)	18	122	141	Edge $p < 0.001$; Hybrid $p < 0.001$
CPU utilization (%)	9.8	27.4	31.8	Edge $p < 0.001$; Hybrid $p < 0.001$
Memory footprint (GB)	0.42	1.26	1.44	Edge $p < 0.001$; Hybrid $p < 0.001$
Bandwidth reduction vs raw streaming (%)	0.0	96.1	96.8	Edge $p < 0.001$; Hybrid $p < 0.001$
Event upload rate (events/hour)	0.9	1.1	1.2	Edge $p = 0.032$; Hybrid $p = 0.018$
End-to-end alert delay to SCADA UI (s)	2.8	5.6	6.2	Edge $p < 0.001$; Hybrid $p < 0.001$

Table 10 summarized deployment feasibility outcomes that were measured under the same windowing schedule across architectures. SCADA-only inference delivered the lowest latency and fastest dashboard visibility, reflecting centralized supervisory proximity and lighter processing. Edge and hybrid architectures required significantly more compute resources, reflected in higher CPU and memory use, and they introduced longer end-to-end alert delay due to local feature extraction and transmission. At the same time, both architectures achieved very large bandwidth reductions relative to raw waveform streaming by uploading compact features and event summaries. Event upload rates increased slightly in edge and hybrid designs, consistent with more frequent anomaly evidence generation. These feasibility tradeoffs provided quantitative context for interpreting predictive gains.

DISCUSSION

Smart manufacturing research has consistently framed predictive maintenance as a data-driven pathway for improving equipment availability and reducing unplanned downtime, and the findings of this study aligned with that framing by demonstrating measurable differences in predictive performance across supervisory, edge, and hybrid architectures. Earlier studies in industrial prognostics frequently characterized SCADA and historian streams as valuable for longitudinal context, alarm history, and exposure tracking, while also documenting limitations associated with sampling granularity, compression, and delayed event signaling (Molęda et al., 2023). The SCADA-only results in this study followed the same pattern: supervisory variables captured load-linked stress signatures and operational disturbance patterns that were correlated with failure proximity, yet those variables displayed moderate rather than strong associations with early fault emergence. The descriptive and correlation evidence showed that current, temperature, and alarm counts were highly interdependent, which echoed prior observations that supervisory telemetry reflected operating context and cumulative stress more than distinct fault physics. This study further expanded that view by quantifying how SCADA-only models produced acceptable alert responsiveness and low system overhead while exhibiting higher nuisance-alert share under fixed alert budgets. Such results mirrored earlier industrial deployments that reported SCADA-based risk scoring as operationally feasible and scalable but vulnerable to false alarms when process variability and nuisance alarms were prevalent (W. Zhang et al., 2019). At the same time, the event-prediction regression outcomes indicated that SCADA-only architecture remained statistically informative after regime stratification and clustering controls were applied, meaning that supervisory analytics contributed measurable value even when high-rate signals were not used. Prior research often treated that value as a function of context rather than precision; consistent with that theme, the SCADA-only condition in this study demonstrated useful detection behavior closer to event onset and showed stronger relationships with fault manifestations that had already become operationally visible in alarms and stops. The combined evidence supported a consistent interpretation found across earlier studies: supervisory analytics served as a practical monitoring backbone for trend-based risk scoring and operational awareness, while its predictive sensitivity depended on event definitions, thresholding discipline, and the extent to which regime variability was explicitly controlled (Chen et al., 2023). The results therefore reinforced the established understanding that predictive maintenance performance in smart manufacturing depended not only on algorithms but also on the architecture that governed signal availability, sampling resolution, and contextual interpretability.

The edge-only findings provided a second, complementary perspective that matched a broad body of earlier work emphasizing the diagnostic richness of high-rate electrical and mechanical sensing for motor-drive assets. Previous investigations in rotating machinery monitoring frequently reported that waveform-level signals preserved transient behaviors and spectral patterns that were essential for detecting subtle degradation, particularly for bearing defects, rotor-related anomalies, and converter-related abnormalities that were not consistently visible at supervisory sampling rates (Serradilla et al., 2022). The edge telemetry in this study captured higher-resolution evidence and produced anomaly and health indicators that correlated more strongly with verified failure proximity than did most SCADA-only predictors, even when supervisory alarms had not yet intensified. This pattern aligned with earlier findings that high-rate sensing often detected degradation earlier than supervisory alarms, thereby increasing actionable lead time. In this study, the median lead time increased substantially in the edge-only condition compared with SCADA-only, and alert precision improved under the same alert-budget policy, indicating that edge analytics provided both earlier signaling and a lower burden of non-actionable alerts. Prior studies also described a persistent challenge of regime dependence in edge signals, especially in variable-speed environments where control strategies and load changes shifted signal baselines. The regime-stratified correlation and stability results in this study indicated that edge indicators became more interpretable when compared within comparable speed and load conditions, consistent with earlier recommendations that segmentation and normalization were essential for reliable industrial inference (Y. Liu et al., 2023). Reliability and validity checks supported that interpretation by showing stronger repeatability of edge feature families when evaluated within regime-controlled windows, which matched prior observations that high-rate features were sensitive but not inherently unstable when context controls were applied. The deployment feasibility findings also reflected the practical constraints discussed in earlier edge analytics deployments: local processing increased inference latency and resource usage compared with supervisory analytics, and end-to-end alert delays were longer because signals required local feature extraction and transmission. However, the observed bandwidth reductions demonstrated that high-rate sensing did not require continuous raw streaming to deliver value, which aligned with earlier edge designs that used feature summarization and event-driven uploads to reduce network load while preserving diagnostic information (Erdemir et al., 2020). Taken together, the edge-only results converged with prior research that positioned edge analytics as a higher-fidelity approach capable of earlier detection and improved precision, balanced by measurable compute and integration costs that required disciplined lifecycle management and robust operational monitoring.

The hybrid SCADA-to-edge fusion condition produced the strongest overall predictive maintenance effectiveness in this study, and that result corresponded with earlier literature that argued for multi-layer integration as a way to combine context with signal fidelity. Prior work in cyber–physical systems and industrial AI frequently described a tradeoff: SCADA systems provided interpretability, governance, and long-horizon operational state, while edge systems provided detailed physical signatures and low-level signal evidence (Mohsen et al., 2023). The hybrid models in this study demonstrated that combining these two layers reduced redundancy, improved validity, and stabilized performance across regimes. Correlation findings showed that hybrid fusion scores preserved strong association with failure proximity while reducing excessive overlap among predictors that had been present in the SCADA-only feature block. Collinearity diagnostics quantified this benefit by showing that fusion designs achieved the lowest final redundancy after supervisory composites were applied, resulting in stable coefficient estimates and reduced standard errors in multivariate regression. Earlier studies often emphasized that fusion improved performance when the combined data sources were complementary rather than duplicative; the results here supported that condition because supervisory features reflected exposure and regime context, while edge features reflected degradation-sensitive deviations that were less correlated with supervisory magnitudes. The regression outcomes further demonstrated that the hybrid configuration improved event detection odds, increased precision under fixed alert budgets, reduced false alarm density, and provided the longest median lead time, indicating performance improvements on multiple operationally meaningful dimensions rather than only on statistical metrics. Earlier industrial predictive maintenance research often highlighted those improvements in sensitivity frequently came at the cost of increased alarm burden; the hybrid results in this study contradicted that common tradeoff by showing that precision and burden improved simultaneously, which suggested that fusion reduced nuisance triggers by contextualizing high-rate anomalies within supervisory operating state (Aljohani, 2023). The regime moderation results strengthened this interpretation by showing lower precision variation across speed/load strata in the hybrid condition, reflecting the value of contextual conditioning in variable-duty manufacturing settings. Such stability aligned with earlier findings that models trained on mixed-regime data often confused regime differences with health differences unless context variables were integrated explicitly. The feasibility results also clarified that hybrid gains were achieved alongside moderate increases in compute load and alert delay, consistent with earlier accounts that integration introduced processing overhead but remained operationally viable when event-driven communication and feature summarization were used (Antwi-Afari et al., 2022). Overall, the hybrid outcomes corresponded to earlier arguments that SCADA-to-edge integration improved both the quality and manageability of predictive maintenance decisions when data alignment, feature design, and governance controls were implemented carefully.

The evaluation design and credibility controls in this study addressed common methodological weaknesses identified in earlier predictive maintenance research, and the findings demonstrated the practical consequences of adopting stricter validation protocols. Prior surveys and comparative studies frequently critiqued predictive maintenance papers for leakage-prone splitting, insufficient regime stratification, and limited uncertainty reporting, especially when windowed time-series samples were randomly partitioned and thereby allowed near-duplicate windows to appear in both training and testing sets (Van Der Zalm et al., 2022). This study's time-based splitting, asset-wise holdout logic, and regime-aware reporting reduced those risks and yielded performance differences that remained statistically meaningful after clustering controls were applied at the motor-drive level. Earlier work also described the sensitivity of results to thresholding and windowing choices, and the robustness checks in this study confirmed that architecture rankings remained consistent across alternative threshold policies and window lengths, indicating that observed differences were not artifacts of a single parameter configuration. Reliability and validity findings also aligned with prior methodological recommendations that urged explicit measurement checks for constructed features and representations, particularly when composite indices were used. The reported internal consistency for supervisory composites, repeatability for edge features under regime control, and convergent behavior around verified events supported the claim that indicators reflected structured phenomena rather than random fluctuation (Giuffrè & Shung, 2023). Earlier studies often treated ground truth as a limiting factor due to work-order lag, ambiguous root cause, and nuisance alarms; the multi-tier event definitions used in this study reflected that reality by distinguishing confirmed failures, verified defects, and operational abnormalities. Correlation results that differentiated relationships with alarm recurrence versus failure proximity matched earlier warnings that alarms were imperfect labels and that alarm-based ground truth could exaggerate associations that emerged after a fault became operationally visible. The construct validity checks around synchronization further addressed another common limitation described in earlier SCADA-to-edge deployments: misalignment between supervisory event timestamps and high-rate windowed signals could degrade model performance in transient-heavy phases. This study's quantified alignment error distribution and its observed association with performance stability reflected earlier technical discussions that treated time synchronization as a core measurement requirement rather than a minor engineering detail. Collectively, the findings demonstrated that stronger credibility controls did not eliminate performance differences; rather, they clarified which differences remained robust under industrial realism, thereby strengthening the interpretability of architecture comparisons in a way that aligned with best-practice themes in prior research (Lin & Fang, 2021).

Figure 12: Predictive Maintenance Framework for Manufacturing

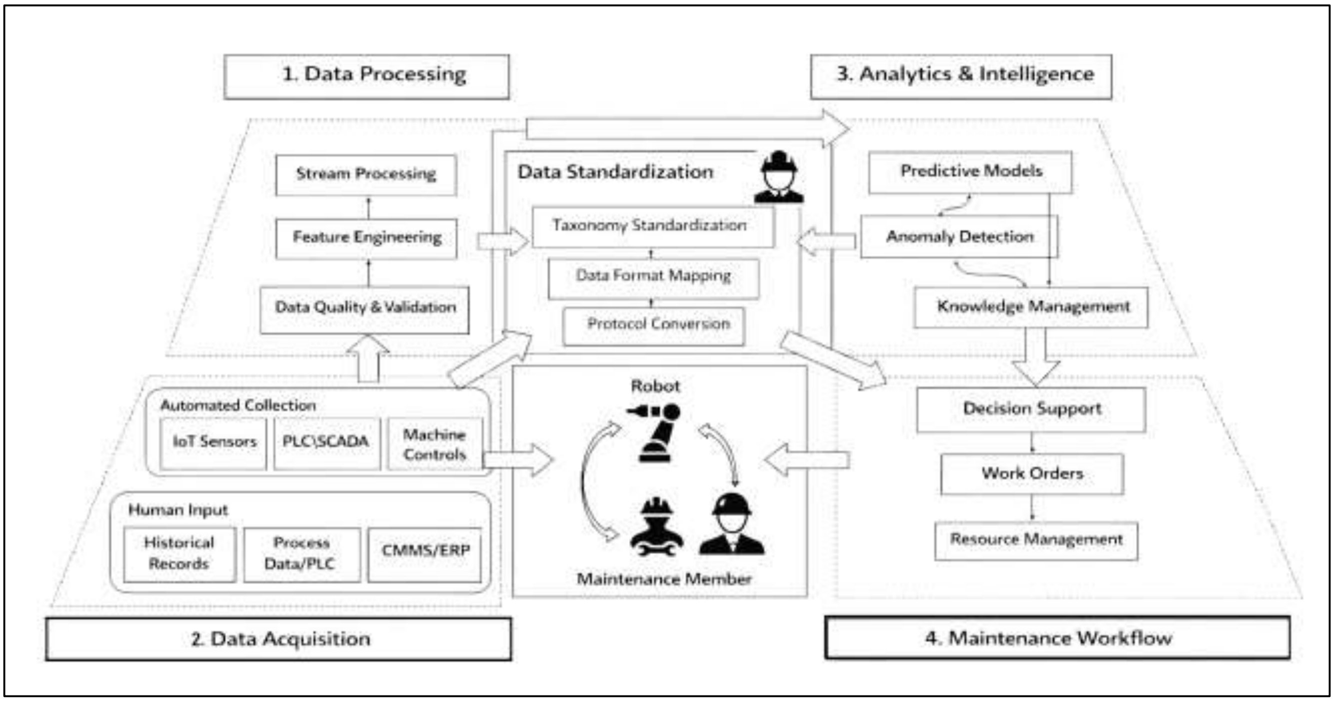

The operational burden outcomes in this study contributed to an ongoing discussion in earlier predictive maintenance work concerning the tension between predictive sensitivity and maintainability of alert workflows. Prior industrial deployments often reported that models with high recall could generate excessive alert volumes, leading to alarm fatigue, low trust, and reduced response rates, which in turn undermined measured benefits in downtime reduction and planning effectiveness (Silva et al., 2023). The fixed alert-budget approach used here provided a consistent operational lens for comparing architectures, and the results indicated that edge and hybrid systems improved precision while reducing nuisance-alert share relative to SCADA-only, even under comparable alert constraints. This result aligned with earlier findings that nuisance alerts could be reduced when models used richer features and better contextualization; however, the present findings also suggested that the magnitude of burden reduction depended on the architecture's ability to incorporate both high-rate evidence and supervisory context. The false alarm density reductions observed in edge and hybrid conditions reinforced prior conclusions that high-resolution signals could increase the discriminative power of alerts, but the findings also clarified that discrimination alone did not guarantee operational manageability unless thresholds were tuned to workload capacity (Pech et al., 2021). Earlier studies in alarm management emphasized metrics such as alerts per operating hour, mean time between false alarms, and conversion of alerts into verified findings; the outcomes reported in this study aligned with those operational metrics by quantifying false alarms per 100 operating hours and showing architecture-level differences that remained significant in hypothesis tests. The lead time distributions also interacted with operational burden: earlier alerts were valuable only when they remained credible and interpretable, and the increased lead times achieved in edge and hybrid conditions coincided with improved precision, which strengthened the practical interpretation of earlier signaling. Earlier research also noted that regime variability could drive spurious alert clusters during specific shifts or recipes; the regime moderation findings here indicated that fusion reduced performance volatility across speed/load strata, which corresponded to a lower likelihood of regime-driven alert spikes. These patterns collectively matched earlier work that treated predictive maintenance as a socio-technical decision system: measurable alert quality and alert volume determined whether maintenance teams could respond consistently, and improvements were observed when decision outputs aligned with operational capacity and when context reduced misinterpretation of normal variability as fault behavior (Rajabzadeh & Fatorachian, 2023).

The deployment feasibility findings offered a quantitative articulation of the tradeoffs between model placement and system constraints that earlier SCADA-to-edge studies frequently described qualitatively (McKeering & Hwang, 2019). Prior work on edge analytics emphasized that increased fidelity and responsiveness were accompanied by constraints related to compute budgets, memory limits, deterministic execution requirements, and integration complexity, and this study quantified those constraints through latency, CPU utilization, memory footprint, bandwidth reduction, event upload rates, and end-to-end alert delay. The measured increases in inference latency and resource utilization in edge and hybrid configurations were consistent with earlier accounts of windowed signal processing overhead, particularly when high-rate telemetry required filtering, feature extraction, and representation computation. At the same time, the observed bandwidth reductions aligned with earlier designs that promoted local summarization and event-driven uploads as a means of preventing network saturation without sacrificing diagnostic insight. Earlier industrial IoT literature also discussed the importance of alert delivery latency to supervisory dashboards, since maintenance workflows often depended on SCADA interfaces and historian-integrated alerting; the longer end-to-end delays in edge and hybrid conditions reflected the additional processing and transmission steps required in those architectures (Martyushev et al., 2023). However, the feasibility results also showed that these delays remained within operationally manageable seconds-level ranges rather than expanding to disruptive scales, which matched earlier deployments that successfully integrated edge analytics with SCADA alarm systems when pipelines were engineered for reliability. The slight increase in event upload rate in edge and hybrid conditions also corresponded to prior observations that higher sensitivity systems created more candidate events; in this study, that increase did not translate into higher nuisance-alert share, suggesting that event generation and alerting decisions remained distinguishable stages. Earlier governance-oriented studies emphasized that deployment success depended on monitoring system health and ensuring consistent model versions across distributed nodes; the reliability, validity, and collinearity results here reinforced that importance by showing that stable measurement and stable inference depended on controlled feature pipelines and reduced redundancy (Alowais et al., 2023). Overall, the feasibility findings were consistent with earlier work that treated architecture choice as a multi-objective balance: compute and latency costs increased with edge processing, while predictive performance and bandwidth efficiency improved, and the hybrid condition achieved the most favorable performance profile while retaining operationally acceptable system overhead.

The integrated interpretation of findings across descriptive, correlational, reliability, collinearity, and regression evidence supported an architecture-sensitive understanding of predictive maintenance that corresponded closely with earlier conceptual models in industrial analytics (Allioui & Mourdi, 2023). Prior research frequently argued that predictive maintenance outcomes depended on the joint configuration of sensing, data quality, preprocessing, modeling, thresholding, and operational integration, and the present results demonstrated that architecture altered several of these components simultaneously. SCADA-only deployment emphasized centralized context and low overhead while exhibiting higher redundancy among predictors and weaker early fault sensitivity, a pattern consistent with earlier historian-based risk scoring studies. Edge-only deployment improved early detection and precision under alert constraints while increasing compute cost and requiring stronger regime control, reflecting patterns documented in waveform-driven condition monitoring research. Hybrid deployment combined the strengths of both layers by improving detection odds, increasing lead time, reducing nuisance alerts, and stabilizing performance across operating regimes, while introducing manageable resource and integration overhead (Gopalakrishnan et al., 2022). Earlier methodological critiques of predictive maintenance research emphasized leakage prevention, robust validation, and transparency about label quality; the multi-tier event definitions, time-based splitting, regime stratification, and alignment validation applied in this study addressed those concerns and yielded results that remained significant under credibility controls. Earlier discussions of operational burden and alarm fatigue emphasized that alert quality and alert volume determined adoption and effectiveness; this study's fixed-budget comparisons and burden metrics provided quantitative evidence that architecture influenced manageability, not only accuracy. Across these dimensions, the findings were consistent with the broader research trajectory that treated SCADA-to-edge predictive maintenance as a system-level problem rather than a single-model problem, where the measured benefits emerged from combining contextual telemetry with high-resolution signal evidence under disciplined evaluation and governance controls (Alamer, 2022). The discussion therefore situated the study's results within established evidence patterns in smart manufacturing predictive maintenance research by demonstrating quantitatively how placement, data fidelity, and context integration shaped predictive performance, stability across regimes, and operational feasibility in motor-drive monitoring environments.

CONCLUSION

The study titled AI-Driven Predictive Maintenance for Motor Drives in Smart Manufacturing: A SCADA-to-Edge Deployment Study was discussed as a quantitative investigation that examined how predictive maintenance effectiveness changed when analytics were implemented using supervisory telemetry alone, high-rate edge telemetry alone, or a fused SCADA-to-edge architecture that combined contextual plant signals with waveform-level evidence. The discussion emphasized that the observed performance differences were consistent with widely reported patterns in the predictive maintenance literature on smart manufacturing and rotating machinery, where SCADA and historian data were frequently described as highly valuable for longitudinal monitoring, alarm context, and exposure tracking but structurally constrained by low sampling density, compression, and delayed event signaling relative to early fault physics. In this study, SCADA-only models demonstrated measurable predictive value near fault manifestation through correlated increases in current, temperature, alarm recurrence, and stop indicators, yet they also exhibited higher predictor redundancy and a higher nuisance-alert share under fixed alert-budget constraints, which aligned with prior industrial deployments reporting that historian-driven risk scoring could be scalable and low overhead while remaining sensitive to process variability and alarm noise. The discussion further interpreted the edge-only results through the lens of earlier high-frequency condition monitoring studies that highlighted the diagnostic richness of waveform-level electrical and mechanical signals, showing that edge-derived anomaly scores, health indices, and learned representations were more strongly associated with verified failure proximity and delivered materially longer actionable lead time than supervisory analytics, while requiring greater local compute and introducing modestly longer end-to-end alert delay due to feature extraction and transmission. Consistent with earlier work emphasizing regime dependence in variable-speed environments, regime stratification and context conditioning were discussed as critical elements that increased repeatability and interpretability of edge indicators and reduced confounding from normal speed and load changes, with reliability evidence showing stable behavior across repeated nominal cycles within comparable regimes. The hybrid fusion outcomes were interpreted as a system-level confirmation of earlier arguments that combining supervisory context with edge fidelity improved both discrimination and manageability, since the hybrid configuration produced the strongest event detection odds, the highest precision under the same alert budget, the lowest false alarm density per operating hour, and the most stable performance across speed/load strata, while also exhibiting reduced collinearity after supervisory composites were applied and yielding stable regression coefficients under clustering controls. The discussion highlighted that these gains were not obtained by increasing alert volume, because improvements occurred alongside reduced nuisance-alert share, indicating that fusion helped contextualize high-rate deviations and filter regime-driven artifacts that often-inflated alarm burden in single-layer systems. The feasibility results were discussed in relation to earlier SCADA-to-edge architecture studies that described predictable tradeoffs between performance and resource costs, where edge and hybrid inference increased latency and compute utilization but achieved substantial bandwidth reduction through feature-level reporting rather than raw streaming and maintained operationally acceptable dashboard delivery delays measured in seconds rather than disruptive timescales. Finally, the discussion connected these outcomes to methodological critiques in earlier predictive maintenance research by noting that leakage prevention through time-based evaluation, asset-level clustering, regime stratification, and explicit alignment validation strengthened the credibility of architecture comparisons, while multi-tier event definitions addressed known limitations in industrial ground truth derived from work orders and fault codes. Overall, the discussion treated the findings as consistent with the broader evidence base showing that predictive maintenance performance in smart manufacturing motor-drive environments was shaped by architecture-level decisions about data fidelity and context integration as much as by model choice, and that measurable improvements in detection quality, lead time, and operational burden emerged most strongly when SCADA context and edge-level signal evidence were combined under disciplined evaluation and governance controls.

RECOMMENDATIONS

Recommendations for AI-Driven Predictive Maintenance for Motor Drives in Smart Manufacturing: A SCADA-to-Edge Deployment Study were framed as actionable, architecture-specific steps that could be implemented to strengthen predictive performance, reduce operational burden, and preserve quantitative validity in real plant conditions while maintaining auditability and system reliability. A hybrid SCADA-to-edge configuration was recommended as the primary deployment pattern for motor-drive predictive maintenance because the comparative results supported those integrated designs balanced contextual interpretability with high-rate fault sensitivity, enabling improved alert precision under fixed alert budgets, reduced nuisance-alert share, and increased actionable lead time while maintaining stable behavior across speed and load regimes. SCADA-only analytics were recommended as a baseline layer for coverage and governance, particularly for plants with limited edge instrumentation, because historian variables provided robust long-horizon exposure indicators and direct integration into existing maintenance workflows; however, supervisory deployments were recommended to include stricter alarm hygiene controls such as nuisance-alarm filtering, tag normalization within operating regimes, and composite construction for stress and disturbance indicators to reduce redundancy and prevent unstable multivariate inference. Edge deployments were recommended to prioritize a minimal but high-value sensing set—phase currents and voltages complemented by vibration where feasible—paired with deterministic windowing rules and regime-aware segmentation so that feature stability remained high across repeated cycles and so that transient-heavy windows did not inflate false alarms; edge pipelines were also recommended to compute compact feature summaries and event-triggered waveform captures rather than continuous raw streaming in order to preserve bandwidth efficiency and reduce storage pressure. Across all architectures, rigorous time synchronization and event anchoring were recommended as mandatory measurement controls, with alignment error tracking maintained as a monitored KPI because misalignment was associated with degraded predictive performance in transient phases; synchronization protocols were recommended to use shared event anchors such as start commands, speed threshold crossings, and fault-code onset markers, with periodic clock drift checks at the edge. For model governance, strict versioning and rollback controls were recommended so that model and preprocessing changes could be traced to alert behavior changes, and operational monitoring was recommended to track drift in feature distributions, score distributions, and alert density over time, with predefined triggers for investigation when sustained shifts occurred. Data integrity controls were recommended to quantify missingness, dropout episodes, and sensor drift at the asset-week level and to evaluate model robustness under dropout simulation so that reliability could be maintained even when telemetry degraded. For evaluation practice, it was recommended that plants adopt leakage-resistant validation protocols that matched deployment goals, including time-based splits for forward realism, asset-wise holdouts for new-drive generalization, and site-wise holdouts when multi-plant deployment was intended; performance reporting was recommended to include regime-stratified summaries so that stability across speed/load bands was visible rather than masked by aggregated averages. Finally, it was recommended that maintenance integration emphasize operational manageability by adopting fixed alert-budget policies, prioritization queues, and clear escalation pathways into SCADA dashboards and work-order systems, ensuring that improved detection translated into consistent maintenance action without increasing workload volatility.

LIMITATION

Limitations associated with AI-Driven Predictive Maintenance for Motor Drives in Smart Manufacturing: A SCADA-to-Edge Deployment Study were interpreted as constraints that shaped the scope, generalizability, and measurement certainty of the reported quantitative results, particularly because predictive maintenance performance in industrial settings depended heavily on event rarity, data quality, and operational heterogeneity. A primary limitation concerned the dependence on field-available ground truth derived from maintenance work orders, downtime logs, and fault-code histories, which were known to contain reporting lag, incomplete root-cause attribution, and inconsistent documentation detail, thereby introducing label uncertainty for both failure timing and fault categorization. Although multi-tier event definitions reduced this ambiguity by separating confirmed failures from verified defects and operational abnormalities, residual label noise may have influenced measured detection timing and may have affected classification boundaries between degradation and process disturbances. A second limitation was the imbalance structure inherent to motor-drive failure data, where long periods of nominal operation were paired with relatively few Tier-1 events, making performance estimates sensitive to event distribution across time and regimes and increasing uncertainty around rare fault modes that occurred only a small number of times. Even when alert-budget evaluation reduced the risk of overly optimistic accuracy reporting, event scarcity constrained the precision of effect estimates for certain subgroups and limited the depth of fault-type comparisons. A third limitation involved instrumentation and telemetry constraints: SCADA tags were subject to compression, polling jitter, and occasional missingness, while edge telemetry depended on local device uptime, sensor mounting stability, and high-rate sampling consistency, all of which could introduce drift or dropout patterns that were not uniform across assets. Although alignment error was quantified and used to evaluate construct validity, synchronization imperfections could still have introduced mislabeling of windowed data around transient events, especially in operating phases where informative signatures were concentrated in short intervals. A fourth limitation concerned operating regime heterogeneity and process variability, because motor drives experienced different speed/load distributions, recipe changes, and environmental conditions across shifts, and not all contextual covariates may have been captured explicitly in the available SCADA tags. Regime stratification reduced confounding, yet residual unmeasured process factors could have influenced both predictor behavior and failure likelihood, which could have affected estimated architecture effects in regression models. Another limitation related to architectural comparability: SCADA-only, edge-only, and hybrid pipelines differed not only in predictive inputs but also in processing latency, compute environment, and data reduction strategies, meaning that performance comparisons reflected combined system behavior rather than a purely algorithmic contrast, and this system-level nature limited direct attribution of gains to any single modeling element. Finally, feasibility metrics such as CPU utilization, memory footprint, and alert delay were measured under specific windowing schedules and hardware configurations, and these values could vary under different edge platforms, sampling rates, or integration designs, which constrained the transferability of specific feasibility figures even when relative tradeoff patterns remained consistent.

REFERENCES

[1]. Abikoye, O. C., Bajeh, A. O., Awotunde, J. B., Ameen, A. O., Mojeed, H. A., Abdulraheem, M., Oladipo, I. D., & Salihu, S. A. (2021). Application of internet of thing and cyber physical system in Industry 4.0 smart manufacturing. In *Emergence of Cyber Physical System and IoT in Smart Automation and Robotics: Computer Engineering in Automation* (pp. 203-217). Springer.

[2]. Achouch, M., Dimitrova, M., Ziane, K., Sattarpanah Karganroudi, S., Dhouib, R., Ibrahim, H., & Adda, M. (2022). On predictive maintenance in industry 4.0: Overview, models, and challenges. *Applied Sciences*, *12*(16), 8081.

[3]. Adjekum, D. K., & Tous, M. F. (2020). Assessing the relationship between organizational management factors and a resilient safety culture in a collegiate aviation program with Safety Management Systems (SMS). *Safety science*, *131*, 104909.

[4]. Alamer, A. (2022). Exploratory structural equation modeling (ESEM) and bifactor ESEM for construct validation purposes: Guidelines and applied example. *Research Methods in Applied Linguistics*, *1*(1), 100005.

[5]. Aljohani, A. (2023). Predictive analytics and machine learning for real-time supply chain risk mitigation and agility. *Sustainability*, *15*(20), 15088.

[6]. Allioui, H., & Mourdi, Y. (2023). Exploring the full potentials of IoT for better financial growth and stability: A comprehensive survey. *Sensors*, *23*(19), 8015.

[7]. Alowais, S. A., Alghamdi, S. S., Alsuhebany, N., Alqahtani, T., Alshaya, A. I., Almohareb, S. N., Aldairem, A., Alrashed, M., Bin Saleh, K., & Badreldin, H. A. (2023). Revolutionizing healthcare: the role of artificial intelligence in clinical practice. *BMC medical education*, *23*(1), 689.

[8]. Alvarez, R., Diez-Gonzalez, J., Verde, P., & Perez, H. (2020). Comparative performance analysis of time local positioning architectures in NLOS urban scenarios. *IEEE Access*, *8*, 225258-225271.

[9]. Andronie, M., Lăzăroiu, G., Ștefănescu, R., Uță, C., & Dijmărescu, I. (2021). Sustainable, smart, and sensing technologies for cyber-physical manufacturing systems: A systematic literature review. *Sustainability*, *13*(10), 5495.

[10]. Antwi-Afari, P., Ng, S. T., & Chen, J. (2022). Developing an integrative method and design guidelines for achieving systemic circularity in the construction industry. *Journal of Cleaner Production*, *354*, 131752.

[11]. Ayvaz, S., & Alpay, K. (2021). Predictive maintenance system for production lines in manufacturing: A machine learning approach using IoT data in real-time. *Expert Systems with Applications*, *173*, 114598.

[12]. Baldin, I., Nikolich, A., Griffioen, J., Monga, I. I. S., Wang, K.-C., Lehman, T., & Ruth, P. (2020). Fabric: A national-scale programmable experimental network infrastructure. *IEEE Internet Computing*, *23*(6), 38-47.

[13]. Berthelsen, H., Westerlund, H., Bergström, G., & Burr, H. (2020). Validation of the Copenhagen psychosocial questionnaire version III and establishment of benchmarks for psychosocial risk Management in Sweden. *International journal of environmental research and public health*, *17*(9), 3179.

[14]. Bokrantz, J., & Skoogh, A. (2023). Adoption patterns and performance implications of Smart Maintenance. *International journal of production economics*, *256*, 108746.

[15]. Bokrantz, J., Skoogh, A., Berlin, C., Wuest, T., & Stahre, J. (2020). Smart Maintenance: a research agenda for industrial maintenance management. *International journal of production economics*, *224*, 107547.

[16]. Bonci, A., Indri, M., Kermenov, R., Longhi, S., & Nabissi, G. (2021). Comparison of PMSMs Motor Current Signature Analysis and Motor Torque Analysis Under Transient Conditions. 2021 IEEE 19th International Conference on Industrial Informatics (INDIN),

[17]. Braun, S., Gamper, H., Reddy, C. K., & Tashev, I. (2021). Towards efficient models for real-time deep noise suppression. ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP),

[18]. Busetto, L., Wick, W., & Gumbinger, C. (2020). How to use and assess qualitative research methods. *Neurological Research and practice*, *2*(1), 14.

[19]. Bustinza, O. F., Opazo-Basaez, M., & Tarba, S. (2022). Exploring the interplay between Smart Manufacturing and KIBS firms in configuring product-service innovation performance. *Technovation*, *118*, 102258.

[20]. Butt, J. (2020). A strategic roadmap for the manufacturing industry to implement industry 4.0. *Designs*, *4*(2), 11.

[21]. Cai, Y., Guan, K., Lobell, D., Potgieter, A. B., Wang, S., Peng, J., Xu, T., Asseng, S., Zhang, Y., & You, L. (2019). Integrating satellite and climate data to predict wheat yield in Australia using machine learning approaches. *Agricultural and forest meteorology*, *274*, 144-159.

[22]. Cámara, J., Silva, M., Garlan, D., & Schmerl, B. (2021). Explaining architectural design tradeoff spaces: A machine learning approach. European Conference on Software Architecture,

[23]. Chandran, L. R., Ilango, K., Nair, M. G., Kumar, A. A., & Kumar, A. A. (2022). Multilabel external fault classification of induction motor using machine learning models. 2022 Third International Conference on Intelligent Computing Instrumentation and Control Technologies (ICICICT),

[24]. Chen, G., Wang, P., Feng, B., Li, Y., & Liu, D. (2020). The framework design of smart factory in discrete manufacturing industry based on cyber-physical system. *International Journal of Computer Integrated Manufacturing*, *33*(1), 79-101.

[25]. Chen, Q., Cao, J., & Zhu, S. (2023). Data-driven monitoring and predictive maintenance for engineering structures: Technologies, implementation challenges, and future directions. *IEEE Internet of Things Journal*, *10*(16), 14527-14551.

[26]. Cheruvu, S., Kumar, A., Smith, N., & Wheeler, D. M. (2019). IoT vertical applications and associated security requirements. In *Demystifying Internet of Things Security: Successful IoT Device/Edge and Platform Security Deployment* (pp. 413-462). Springer.

[27]. Çınar, Z. M., Abdussalam Nuhu, A., Zeeshan, Q., Korhan, O., Asmael, M., & Safaei, B. (2020). Machine learning in predictive maintenance towards sustainable smart manufacturing in industry 4.0. *Sustainability*, *12*(19), 8211.

[28]. Cowley, L. E., Farewell, D. M., Maguire, S., & Kemp, A. M. (2019). Methodological standards for the development and evaluation of clinical prediction rules: a review of the literature. *Diagnostic and prognostic research*, *3*(1), 16.

[29]. Cuervo-Cazurra, A., Andersson, U., Brannen, M. Y., Nielsen, B. B., & Reuber, A. R. (2019). From the editors: Can I trust your findings? Ruling out alternative explanations in international business research. In *Research methods in international business* (pp. 121-157). Springer.

[30]. Dalzochio, J., Kunst, R., Pignaton, E., Binotto, A., Sanyal, S., Favilla, J., & Barbosa, J. (2020). Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges. *Computers in industry*, *123*, 103298.

[31]. Double, K. S., McGrane, J. A., & Hopfenbeck, T. N. (2020). The impact of peer assessment on academic performance: A meta-analysis of control group studies. *Educational Psychology Review*, *32*(2), 481-509.

[32]. Erdemir, A., Mulugeta, L., Ku, J. P., Drach, A., Horner, M., Morrison, T. M., Peng, G. C., Vadigepalli, R., Lytton, W. W., & Myers Jr, J. G. (2020). Credible practice of modeling and simulation in healthcare: ten rules from a multidisciplinary perspective. *Journal of translational medicine*, *18*(1), 369.

[33]. Farquhar, J., Michels, N., & Robson, J. (2020). Triangulation in industrial qualitative case study research: Widening the scope. *Industrial Marketing Management*, *87*, 160-170.

[34]. Felsberger, A., Qaiser, F. H., Choudhary, A., & Reiner, G. (2022). The impact of Industry 4.0 on the reconciliation of dynamic capabilities: Evidence from the European manufacturing industries. *Production Planning & Control*, *33*(2-3), 277-300.

[35]. Ferraz Júnior, F., Romero, R. A. F., & Hsieh, S.-J. (2023). Machine learning for the detection and diagnosis of anomalies in applications driven by electric motors. *Sensors*, *23*(24), 9725.

[36]. Filip, F. G., & Leiviskä, K. (2023). Infrastructure and complex systems automation. In *Springer Handbook of Automation* (pp. 617-640). Springer.

[37]. FitzPatrick, B. (2019). Validity in qualitative health education research. *Currents in Pharmacy Teaching and Learning*, *11*(2), 211-217.

[38]. Fritz, S., See, L., Carlson, T., Haklay, M., Oliver, J. L., Fraisl, D., Mondardini, R., Brocklehurst, M., Shanley, L. A., & Schade, S. (2019). Citizen science and the United Nations sustainable development goals. *Nature sustainability*, *2*(10), 922-930.

[39]. Furman, D., Campisi, J., Verdin, E., Carrera-Bastos, P., Targ, S., Franceschi, C., Ferrucci, L., Gilroy, D. W., Fasano, A., & Miller, G. W. (2019). Chronic inflammation in the etiology of disease across the life span. *Nature medicine*, *25*(12), 1822-1832.

[40]. Gan, Z. L., Musa, S. N., & Yap, H. J. (2023). A review of the high-mix, low-volume manufacturing industry. *Applied Sciences*, *13*(3), 1687.

[41]. Giuffrè, M., & Shung, D. L. (2023). Harnessing the power of synthetic data in healthcare: innovation, application, and privacy. *npj Digital Medicine*, *6*(1), 186.

[42]. Gopalakrishnan, M., Subramaniyan, M., & Skoogh, A. (2022). Data-driven machine criticality assessment–maintenance decision support for increased productivity. *Production Planning & Control*, *33*(1), 1-19.

[43]. Griffiths, P., Nendel, C., & Hostert, P. (2019). Intra-annual reflectance composites from Sentinel-2 and Landsat for national-scale crop and land cover mapping. *Remote Sensing of Environment*, *220*, 135-151.

[44]. Gultekin, M. A., & Bazzi, A. (2023). Review of fault detection and diagnosis techniques for AC motor drives. *Energies*, *16*(15), 5602.

[45]. Gupta, S., Maiti, J., & Kumar, A. (2023). An Optimized Machine Learning Technique for Fault Diagnosis of Roller Bearing in a Motor Drive System. 2023 4th International Conference on Data Analytics for Business and Industry (ICDABI),

[46]. Habibullah, S. M., & Muhammad Mohiul, I. (2023). Digital Twin-Driven Thermodynamic and Fluid Dynamic Simulation For Exergy Efficiency In Industrial Power Systems. *American Journal of Scholarly Research and Innovation*, *2*(01), 224–253. https://doi.org/10.63125/k135kt69

[47]. Hashemi, M., Stolz, M., & Watzenig, D. (2023). Super-twisting algorithm-based sliding mode observer for open-circuit fault diagnosis in PWM voltage source inverter in an in-wheel motor drive system. 2023 IEEE International Conference on Mechatronics (ICM),

[48]. Hong, Q. N., Pluye, P., Fàbregues, S., Bartlett, G., Boardman, F., Cargo, M., Dagenais, P., Gagnon, M.-P., Griffiths, F., & Nicolau, B. (2019). Improving the content validity of the mixed methods appraisal tool: a modified e-Delphi study. *Journal of clinical epidemiology*, *111*, 49-59. e41.

[49]. Hu, F., Qiu, X., Jing, G., Tang, J., & Zhu, Y. (2023). Digital twin-based decision making paradigm of raise boring method. *Journal of Intelligent Manufacturing*, *34*(5), 2387-2405.

[50]. Huang, Z., Fey, M., Liu, C., Beysel, E., Xu, X., & Brecher, C. (2023). Hybrid learning-based digital twin for manufacturing process: Modeling framework and implementation. *Robotics and Computer-Integrated Manufacturing*, *82*, 102545.

[51]. Huang, Z., Shen, Y., Li, J., Fey, M., & Brecher, C. (2021). A survey on AI-driven digital twins in industry 4.0: Smart manufacturing and advanced robotics. *Sensors*, *21*(19), 6340.

[52]. Ibrahim, M., Rassõlkin, A., Vaimann, T., & Kallaste, A. (2022). Overview on digital twin for autonomous electrical vehicles propulsion drive system. *Sustainability*, *14*(2), 601.

[53]. Jabed Hasan, T., & Waladur, R. (2023). AI-Driven Cybersecurity, IOT Networking, And Resilience Strategies For Industrial Control Systems: A Systematic Review For U.S. Critical Infrastructure Protection. *International Journal of Scientific Interdisciplinary Research*, *4*(4), 144–176. https://doi.org/10.63125/mbyhj941

[54]. Jia, H., Valavi, H., Tang, Y., Zhang, J., & Verma, N. (2020). A programmable heterogeneous microprocessor based on bit-scalable in-memory computing. *IEEE Journal of Solid-State Circuits*, *55*(9), 2609-2621.

[55]. Jiang, Y., Yin, S., & Kaynak, O. (2020). Performance supervised plant-wide process monitoring in industry 4.0: A roadmap. *IEEE Open Journal of the Industrial Electronics Society*, *2*, 21-35.

[56]. Jin, L., Mao, Y., Wang, X., Lu, L., & Wang, Z. (2023). Online data-driven fault diagnosis of dual three-phase PMSM drives considering limited labeled samples. *IEEE Transactions on Industrial Electronics*, *71*(7), 6797-6808.

[57]. Jin, L., Wang, X., Mao, Y., Lu, L., & Wang, Z. (2023). Online attribute matching based few-sample data-driven diagnosis of electrical faults in PMSM drive. *IEEE Transactions on Power Electronics*, *39*(2), 2620-2631.

[58]. Jinnat, A. (2025). Machine-Learning Models For Predicting Blood Pressure And Cardiac Function Using Wearable Sensor Data. *International Journal of Scientific Interdisciplinary Research*, *6*(2), 102–142. https://doi.org/10.63125/h7rbyt25

[59]. Jinnat, A., & Md. Kamrul, K. (2021). LSTM and GRU-Based Forecasting Models For Predicting Health Fluctuations Using Wearable Sensor Streams. *American Journal of Interdisciplinary Studies*, *2*(02), 32-66. https://doi.org/10.63125/1p8gbp15

[60]. Kalsoom, T., Ramzan, N., Ahmed, S., & Ur-Rehman, M. (2020). Advances in sensor technologies in the era of smart factory and industry 4.0. *Sensors*, *20*(23), 6783.

[61]. Keleko, A. T., Kamsu-Foguem, B., Ngouna, R. H., & Tongne, A. (2022). Artificial intelligence and real-time predictive maintenance in industry 4.0: a bibliometric analysis. *AI and Ethics*, *2*(4), 553-577.

[62]. Khaneghah, M. Z., Alzayed, M., & Chaoui, H. (2023). Fault detection and diagnosis of the electric motor drive and battery system of electric vehicles. *Machines*, *11*(7), 713.

[63]. Kim, J., Jeong, H.-r., & Park, H. (2023). Key drivers and performances of smart manufacturing adoption: A meta-analysis. *Sustainability*, *15*(8), 6496.

[64]. Kuo, Y.-H., & Kusiak, A. (2019). From data to big data in production research: the past and future trends. *International Journal of Production Research*, *57*(15-16), 4828-4853.

[65]. Kutub Uddin, A., Md Mostafizur, R., Afrin Binta, H., & Maniruzzaman, B. (2022). Forecasting Future Investment Value with Machine Learning, Neural Networks, And Ensemble Learning: A Meta-Analytic Study. *Review of Applied Science and Technology*, *1*(02), 01-25. https://doi.org/10.63125/edxgjg56

[66]. Lang, W., Hu, Y., Gong, C., Zhang, X., Xu, H., & Deng, J. (2021). Artificial intelligence-based technique for fault detection and diagnosis of EV motors: A review. *IEEE Transactions on Transportation Electrification*, *8*(1), 384-406.

[67]. Lee, W. J., Wu, H., Yun, H., Kim, H., Jun, M. B., & Sutherland, J. W. (2019). Predictive maintenance of machine tool systems using artificial intelligence techniques applied to machine condition data. *Procedia Cirp*, *80*, 506-511.

[68]. Li, L., Aslam, S., Wileman, A., & Perinpanayagam, S. (2021). Digital twin in aerospace industry: A gentle introduction. *IEEE Access*, *10*, 9543-9562.

[69]. Liang, J., Zhang, K., Al-Durra, A., & Zhou, D. (2020). A novel fault diagnostic method in power converters for wind power generation system. *Applied energy*, *266*, 114851.

[70]. Lin, Y.-K., & Fang, X. (2021). First, do no harm: Predictive analytics to reduce in-hospital adverse events. *Journal of Management Information Systems*, *38*(4), 1122-1149.

[71]. Lindgreen, A., Di Benedetto, C. A., & Beverland, M. B. (2021). How to write up case-study methodology sections. In (Vol. 96, pp. A7-A10): Elsevier.

[72]. Liu, S., Zhang, G., Wang, S., & Sun, H. (2023). Overview of Fault Diagnosis Methods for Top Drive System. Annual Conference of China Electrotechnical Society,

[73]. Liu, Y., Yu, W., Rahayu, W., & Dillon, T. (2023). An evaluative study on IoT ecosystem for smart predictive maintenance (IoT-SPM) in manufacturing: Multiview requirements and data quality. *IEEE Internet of Things Journal*, *10*(13), 11160-11184.

[74]. Lu, Y., Liu, C., Kevin, I., Wang, K., Huang, H., & Xu, X. (2020). Digital Twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, *61*, 101837.

[75]. Lu, Z., Deb, K., Goodman, E., Banzhaf, W., & Boddeti, V. N. (2020). Nsganetv2: Evolutionary multi-objective surrogate-assisted neural architecture search. European conference on computer vision,

[76]. Lu, Z., Whalen, I., Dhebar, Y., Deb, K., Goodman, E. D., Banzhaf, W., & Boddeti, V. N. (2020). Multiobjective evolutionary design of deep convolutional neural networks for image classification. *IEEE Transactions on Evolutionary Computation*, *25*(2), 277-291.

[77]. Lytras, M. D., Visvizi, A., Chopdar, P. K., Sarirete, A., & Alhalabi, W. (2021). Information Management in Smart Cities: Turning end users' views into multi-item scale development, validation, and policy-making recommendations. *International Journal of Information Management*, *56*, 102146.

[78]. Ma, L.-L., Wang, Y.-Y., Yang, Z.-H., Huang, D., Weng, H., & Zeng, X.-T. (2020). Methodological quality (risk of bias) assessment tools for primary and secondary medical studies: what are they and which is better? *Military Medical Research*, *7*(1), 7.

[79]. Manjare, A. A., & Patil, B. G. (2021). A review: Condition based techniques and predictive maintenance for motor. 2021 International Conference on Artificial Intelligence and Smart Systems (ICAIS),

[80]. Marra, D. E., Hamlet, K. M., Bauer, R. M., & Bowers, D. (2020). Validity of teleneuropsychology for older adults in response to COVID-19: A systematic and critical review. *The Clinical Neuropsychologist*, *34*(7-8), 1411-1452.

[81]. Martín-Martín, A., Thelwall, M., Orduna-Malea, E., & Delgado López-Cózar, E. (2021). Google Scholar, Microsoft Academic, Scopus, Dimensions, Web of Science, and OpenCitations' COCI: a multidisciplinary comparison of coverage via citations. *Scientometrics*, *126*(1), 871-906.

[82]. Martyushev, N. V., Malozyomov, B. V., Sorokova, S. N., Efremenkov, E. A., Valuev, D. V., & Qi, M. (2023). Review models and methods for determining and predicting the reliability of technical systems and transport. *Mathematics*, *11*(15), 3317.

[83]. Masud, R., & Md Sarwar Hossain, S. (2024). The Impact of Smart Materials And Fire-Resistant Structures On Safety In U.S. Public Infrastructure. *Journal of Sustainable Development and Policy*, *3*(03), 44-86. https://doi.org/10.63125/ygr1yk30

[84]. McKeering, P., & Hwang, Y.-S. (2019). A systematic review of mindfulness-based school interventions with early adolescents. *Mindfulness*, *10*(4), 593-610.

[85]. Md Arman, H. (2025). Data-Driven Compliance Frameworks For Anti-Money Laundering (AML) And Tax Risk Management In Financial Institutions. *International Journal of Scientific Interdisciplinary Research*, *6*(2), 88–101. https://doi.org/10.63125/2n5pd137

[86]. Md Harun-Or-Rashid, M. (2025a). AI-Driven Threat Detection and Response Framework For Cloud Infrastructure Security. *American Journal of Scholarly Research and Innovation*, *4*(01), 494–535. https://doi.org/10.63125/e58hzh78

[87]. Md Harun-Or-Rashid, M. (2025b). Is The Metaverse the Next Frontier for Corporate Growth And Innovation? Exploring The Potential of The Enterprise Metaverse. *American Journal of Interdisciplinary Studies*, *6*(1), 354-393. https://doi.org/10.63125/ckd54306

[88]. Md, K., & Sai Praveen, K. (2024). Hybrid Discrete-Event And Agent-Based Simulation Framework (H-DEABSF) For Dynamic Process Control In Smart Factories. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *4*(1), 72–96. https://doi.org/10.63125/wcqq7x08

[89]. Md Nahid, H. (2025). AI-Driven Predictive Analytics Framework For Electronic Funds Transfer, Loan Origination, And AML Compliance In Digital Banking. *American Journal of Scholarly Research and Innovation*, *4*(01), 622–661. https://doi.org/10.63125/we3m0t59

[90]. Md Nahid, H., & Tahmina Akter Bhuya, M. (2024). An Empirical Study of Big Data–Enabled Predictive Analytics And Their Impact On Financial Forecasting And Market Decision-Making. *Review of Applied Science and Technology*, *3*(01), 143–182. https://doi.org/10.63125/1mjfqf10

[91]. Md. Akbar, H. (2024). Computational Psychometrics and Digital Biomarker Modeling For Precision Mental Health Diagnostics. *International Journal of Scientific Interdisciplinary Research*, *5*(2), 487–525. https://doi.org/10.63125/vg522x27

[92]. Md. Akbar, H., & Sharmin, A. (2022). Neurobiotechnology-Driven Regenerative Therapy Frameworks For Post-Traumatic Neural Recovery. *American Journal of Scholarly Research and Innovation*, *1*(02), 134–170. https://doi.org/10.63125/24s6kt66

[93]. Md. Foysal, H., & Abdulla, M. (2024). Agile And Sustainable Supply Chain Management Through AI-Based Predictive Analytics And Digital Twin Simulation. *International Journal of Scientific Interdisciplinary Research*, *5*(2), 343–376. https://doi.org/10.63125/sejyk977

[94]. Md. Foysal, H., & Subrato, S. (2022). Data-Driven Process Optimization in Automotive Manufacturing A Machine Learning Approach To Waste Reduction And Quality Improvement. *Journal of Sustainable Development and Policy*, *1*(02), 87-133. https://doi.org/10.63125/2hk0qd38

[95]. Md. Jobayer Ibne, S., & Aditya, D. (2024). Machine Learning and Secure Data Pipeline Frameworks For Improving Patient Safety Within U.S. Electronic Health Record Systems. *American Journal of Interdisciplinary Studies*, *5*(03), 43–85. https://doi.org/10.63125/nb2c1f86

[96]. Md. Mosheur, R. (2025). AI-Driven Predictive Analytics Models For Enhancing Group Insurance Portfolio Performance And Risk Forecasting. *International Journal of Scientific Interdisciplinary Research*, *6*(2), 39–87. https://doi.org/10.63125/qh5qgk22

[97]. Md. Mosheur, R., & Md Arman, H. (2024). Impact Of Big Data and Predictive Analytics On Financial Forecasting Accuracy And Decision-Making In Global Capital Markets. *American Journal of Scholarly Research and Innovation*, *3*(02), 99–140. https://doi.org/10.63125/hg37h121

[98]. Md. Rabiul, K. (2025). Artificial Intelligence-Enhanced Predictive Analytics For Demand Forecasting In U.S. Retail Supply Chains. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 959–993. https://doi.org/10.63125/gbkf5c16

[99]. Md. Rabiul, K., & Khairul Alam, T. (2024). Impact Of IOT and Blockchain Integration On Real-Time Supply Chain Transparency. *International Journal of Scientific Interdisciplinary Research*, *5*(2), 449–486. https://doi.org/10.63125/2yc6e230

[100]. Md. Rabiul, K., & Mohammad Mushfequr, R. (2023). A Quantitative Study On Erp-Integrated Decision Support Systems In Healthcare Logistics. *Review of Applied Science and Technology*, *2*(01), 142–184. https://doi.org/10.63125/c92bbj37

[101]. Md. Rabiul, K., & Samia, A. (2021). Integration Of Machine Learning Models And Advanced Computing For Reducing Logistics Delays In Pharmaceutical Distribution. *American Journal of Advanced Technology and Engineering Solutions*, *1*(4), 01-42. https://doi.org/10.63125/ahnkqj11

[102]. Mittal, S., Khan, M. A., Purohit, J. K., Menon, K., Romero, D., & Wuest, T. (2020). A smart manufacturing adoption framework for SMEs. *International Journal of Production Research*, *58*(5), 1555-1573.

[103]. Mohsen, F., Al-Absi, H. R., Yousri, N. A., El Hajj, N., & Shah, Z. (2023). A scoping review of artificial intelligence-based methods for diabetes risk prediction. *npj Digital Medicine*, *6*(1), 197.

[104]. Molęda, M., Małysiak-Mrozek, B., Ding, W., Sunderam, V., & Mrozek, D. (2023). From corrective to predictive maintenance—A review of maintenance approaches for the power industry. *Sensors*, *23*(13), 5970.

[105]. Mostaani, A., Vu, T. X., Sharma, S. K., Nguyen, V.-D., Liao, Q., & Chatzinotas, S. (2022). Task-oriented communication design in cyber-physical systems: A survey on theory and applications. *IEEE Access*, *10*, 133842-133868.

[106]. Mst. Shahrin, S. (2025). Predictive Neural Network Models For Cyberattack Pattern Recognition And Critical Infrastructure Vulnerability Assessment. *Review of Applied Science and Technology*, *4*(02), 777-819. https://doi.org/10.63125/qp0de852

[107]. Mst. Shahrin, S., & Samia, A. (2023). High-Performance Computing For Scaling Large-Scale Language And Data Models In Enterprise Applications. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *3*(1), 94–131. https://doi.org/10.63125/e7yfwm87

[108]. Muhammad Mohiul, I. (2020). Impact Of Digital Construction Management Platforms on Project Performance Post-Covid-19. *American Journal of Interdisciplinary Studies*, *1*(04), 01-25. https://doi.org/10.63125/nqp0zh08

[109]. Muhammad Mohiul, I., & Rahman, M. D. H. (2021). Quantum-Enhanced Charge Transport Modeling In Perovskite Solar Cells Using Non-Equilibrium Green's Function (NEGF) Framework. *Review of Applied Science and Technology*, *6*(1), 230–262. https://doi.org/10.63125/tdbjaj79
[110]. Nguyen, N. X., Tran, K., & Nguyen, T. A. (2021). Impact of service quality on in-patients' satisfaction, perceived value, and customer loyalty: A mixed-methods study from a developing country. *Patient preference and adherence*, 2523-2538.
[111]. Olsen, M., & Raunak, M. (2019). Quantitative measurements of model credibility. In *Model Engineering for Simulation* (pp. 163-187). Elsevier.
[112]. Osia, S. A., Shamsabadi, A. S., Sajadmanesh, S., Taheri, A., Katevas, K., Rabiee, H. R., Lane, N. D., & Haddadi, H. (2020). A hybrid deep learning architecture for privacy-preserving mobile analytics. *IEEE Internet of Things Journal*, *7*(5), 4505-4518.
[113]. Padeiro, M., Louro, A., & Da Costa, N. M. (2019). Transit-oriented development and gentrification: a systematic review. *Transport Reviews*, *39*(6), 733-754.
[114]. Parhi, S., Joshi, K., & Akarte, M. (2021). Smart manufacturing: a framework for managing performance. *International Journal of Computer Integrated Manufacturing*, *34*(3), 227-256.
[115]. Pech, M., Vrchota, J., & Bednář, J. (2021). Predictive maintenance and intelligent sensors in smart factory. *Sensors*, *21*(4), 1470.
[116]. Qu, Y., Ming, X., Liu, Z., Zhang, X., & Hou, Z. (2019). Smart manufacturing systems: state of the art and future trends. *The international journal of advanced manufacturing technology*, *103*(9), 3751-3768.
[117]. Rahman, M., Kamal, M. M., Aydin, E., & Haque, A. U. (2022). Impact of Industry 4.0 drivers on the performance of the service sector: comparative study of cargo logistic firms in developed and developing regions. *Production Planning & Control*, *33*(2-3), 228-243.
[118]. Rahman, M. D. H. (2022). Modelling The Impact Of Temperature Coefficients On PV System Performance In Hot And Humid Climates. *International Journal of Scientific Interdisciplinary Research*, *1*(01), 194–237. https://doi.org/10.63125/abj6wy92
[119]. Rahman, S. M. T., & Abdul, H. (2021). The Role Of Predictive Analytics In Enhancing Agribusiness Supply Chains. *Review of Applied Science and Technology*, *6*(1), 183–229. https://doi.org/10.63125/n9z10h68
[120]. Raj, P., & Surianarayanan, C. (2020). Digital twin: the industry use cases. In *Advances in computers* (Vol. 117, pp. 285-320). Elsevier.
[121]. Rajabzadeh, M., & Fatorachian, H. (2023). Modelling factors influencing IoT adoption: With a focus on agricultural logistics operations. *Smart cities*, *6*(6), 3266-3296.
[122]. Rajesh Kanna, G., Sasiraja, R., & Prince Winston, D. (2020). Design and development of Truncated Angle Variant (TAV) controller for multi-source-fed BLDC motor drive. *Electrical Engineering*, *102*(4), 1931-1946.
[123]. Rakibul, H. (2025). A Systematic Review Of Human-AI Collaboration In It Support Services: Enhancing User Experience And Workflow Automation. *American Journal of Interdisciplinary Studies*, *6*(3), 01-37. https://doi.org/10.63125/0fd1yb74
[124]. Rakibul, H., & Khairul Alam, T. (2023). A Systematic Review of Predictive Analytics In Marketing Decision-Making Exploring AI-Driven Consumer Segmentation And AB Testing. *International Journal of Business and Economics Insights*, *3*(1), 68–96. https://doi.org/10.63125/2hvfh110
[125]. Ramya, L., & Sivaprakasam, A. (2020). Application of model predictive control for reduced torque ripple in orthopaedic drilling using permanent magnet synchronous motor drive. *Electrical Engineering*, *102*(3), 1469-1482.
[126]. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., & Koltun, V. (2020). Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, *44*(3), 1623-1637.
[127]. Rifat, C., & Rebeka, S. (2023). The Role Of ERP-Integrated Decision Support Systems In Enhancing Efficiency And Coordination In Healthcare Logistics: A Quantitative Study. *International Journal of Scientific Interdisciplinary Research*, *4*(4), 265–285. https://doi.org/10.63125/c7srk144
[128]. Saba, A., & Md. Sakib Hasan, H. (2024). Machine Learning And Secure Data Pipelines For Enhancing Patient Safety In Electronic Health Record (EHR) Among U.S. Healthcare Providers. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *4*(1), 124–168. https://doi.org/10.63125/qm4he747
[129]. Sabuj Kumar, S. (2023). Integrating Industrial Engineering and Petroleum Systems With Linear Programming Model For Fuel Efficiency And Downtime Reduction. *Journal of Sustainable Development and Policy*, *2*(04), 108-139. https://doi.org/10.63125/v7d6a941
[130]. Sabuj Kumar, S. (2024). Petroleum Storage Tank Design and Inspection Using Finite Element Analysis Model For Ensuring Safety Reliability And Sustainability. *Review of Applied Science and Technology*, *3*(04), 94–127. https://doi.org/10.63125/a18zw719
[131]. Sabuj Kumar, S. (2025). AI Driven Predictive Maintenance In Petroleum And Power Systems Using Random Forest Regression Model For Reliability Engineering Framework. *American Journal of Scholarly Research and Innovation*, *4*(01), 363-391. https://doi.org/10.63125/477x5t65
[132]. Sai Praveen, K. (2024). AI-Enhanced Data Science Approaches For Optimizing User Engagement In U.S. Digital Marketing Campaigns. *Journal of Sustainable Development and Policy*, *3*(03), 01-43. https://doi.org/10.63125/65ebsn47
[133]. Sai Praveen, K., & Md, K. (2025). Real-Time Cyber-Physical Deployment and Validation Of H-DEABSF: Model Predictive Control, And Digital-Twin–Driven Process Control In Smart Factories. *Review of Applied Science and Technology*, *4*(02), 750-776. https://doi.org/10.63125/yrkm0057

[134]. Saikat, S., & Aditya, D. (2023). Reliability-Centered Maintenance Optimization Using Multi-Objective Ai Algorithms In Refinery Equipment. *American Journal of Scholarly Research and Innovation*, *2*(01), 389–411. https://doi.org/10.63125/6a6kqm73

[135]. Samuelsen, J., Chen, W., & Wasson, B. (2019). Integrating multiple data sources for learning analytics—review of literature. *Research and Practice in Technology Enhanced Learning*, *14*(1), 11.

[136]. Sardashti, A., & Nazari, J. (2023). A learning-based approach to fault detection and fault-tolerant control of permanent magnet DC motors. *Journal of Engineering and Applied Science*, *70*(1), 109.

[137]. Serradilla, O., Zugasti, E., Rodriguez, J., & Zurutuza, U. (2022). Deep learning models for predictive maintenance: a survey, comparison, challenges and prospects. *Applied Intelligence*, *52*(10), 10934-10964.

[138]. Shahramian, S., Holyoak, M. J., Singh, A., & Baeyens, Y. (2019). A fully integrated 384-element, 16-tile, $ W $-band phased array with self-alignment and self-test. *IEEE Journal of Solid-State Circuits*, *54*(9), 2419-2434.

[139]. Shaikat, B., & Aditya, D. (2024). Graph Neural Network Models For Predicting Cyber Attack Patterns In Critical Infrastructure Systems. *Review of Applied Science and Technology*, *3*(01), 68–105. https://doi.org/10.63125/pmnqxk63

[140]. Silva, M. d. S., Araújo, J. L., Nunes, G. A. d. A., Rosa, M. F. F., Luz, G. V. d. S., Rosa, S. d. S., & Piratelli-Filho, A. (2023). Precision and reliability study of hospital infusion pumps: a systematic review. *BioMedical Engineering OnLine*, *22*(1), 26.

[141]. Sobb, T., Turnbull, B., & Moustafa, N. (2020). Supply chain 4.0: A survey of cyber security challenges, solutions and future directions. *Electronics*, *9*(11), 1864.

[142]. Souza, R. M., Nascimento, E. G., Miranda, U. A., Silva, W. J., & Lepikson, H. A. (2021). Deep learning for diagnosis and classification of faults in industrial rotating machinery. *Computers & Industrial Engineering*, *153*, 107060.

[143]. Sun, Q., Yu, X., Li, H., & Fan, J. (2022). Adaptive feature extraction and fault diagnosis for three-phase inverter based on hybrid-CNN models under variable operating conditions. *Complex & Intelligent Systems*, *8*(1), 29-42.

[144]. Sun, S., Law, R., & Schuckert, M. (2020). Mediating effects of attitude, subjective norms and perceived behavioural control for mobile payment-based hotel reservations. *International Journal of Hospitality Management*, *84*, 102331.

[145]. Sun, X., Liu, Z., Li, A., Wang, Z., Jiang, D., & Qu, R. (2022). Self-adaptive fault-tolerant control of three-phase series-end winding motor drive. *IEEE Transactions on Power Electronics*, *37*(9), 10939-10950.

[146]. Sundmaeker, H., Verdouw, C., Wolfert, S., & Freire, L. P. (2022). Internet of food and farm 2020. In *Digitising the industry internet of things connecting the physical, digital and virtualworlds* (pp. 129-151). River Publishers.

[147]. Swanke, J. A., Zeng, H., Jahns, T. M., & Sarlioglu, B. (2023). Systematic Motor Drive Reliability Improvement Methodology Using Fault-Tolerant Modular Motor Drives. 2023 IEEE Energy Conversion Congress and Exposition (ECCE),

[148]. Tambare, P., Meshram, C., Lee, C.-C., Ramteke, R. J., & Imoize, A. L. (2021). Performance measurement system and quality management in data-driven Industry 4.0: A review. *Sensors*, *22*(1), 224.

[149]. Tancock, S., Arabul, E., & Dahnoun, N. (2019). A review of new time-to-digital conversion techniques. *IEEE Transactions on Instrumentation and Measurement*, *68*(10), 3406-3417.

[150]. Teler, K., Skowron, M., & Orłowska-Kowalska, T. (2023). Implementation of MLP-based classifier of current sensor faults in vector-controlled induction motor drive. *IEEE Transactions on Industrial Informatics*, *20*(4), 5702-5713.

[151]. Van Der Zalm, A. J., Barroso, J., Browne, P., Casey, W., Gordon, J., Henry, T. R., Kleinstreuer, N. C., Lowit, A. B., Perron, M., & Clippinger, A. J. (2022). A framework for establishing scientific confidence in new approach methodologies. *Archives of Toxicology*, *96*(11), 2865-2879.

[152]. Velte, P., & Stawinoga, M. (2020). Do chief sustainability officers and CSR committees influence CSR-related outcomes? A structured literature review based on empirical-quantitative research findings. *Journal of Management Control*, *31*(4), 333-377.

[153]. Verma, N., Jia, H., Valavi, H., Tang, Y., Ozatay, M., Chen, L.-Y., Zhang, B., & Deaville, P. (2019). In-memory computing: Advances and prospects. *IEEE solid-state circuits magazine*, *11*(3), 43-55.

[154]. Vermesan, O., Friess, P., Guillemin, P., Sundmaeker, H., Eisenhauer, M., Moessner, K., Le Gall, F., & Cousin, P. (2022). Internet of things strategic research and innovation agenda. In *Internet of things* (pp. 7-151). River Publishers.

[155]. Wang, J., & Gao, R. X. (2022). Innovative smart scheduling and predictive maintenance techniques. In *Design and operation of production networks for mass personalization in the era of cloud technology* (pp. 181-207). Elsevier.

[156]. Wang, J., Ye, L., Gao, R. X., Li, C., & Zhang, L. (2019). Digital Twin for rotating machinery fault diagnosis in smart manufacturing. *International Journal of Production Research*, *57*(12), 3920-3934.

[157]. Wei, J., Li, Z., Lyapustin, A., Sun, L., Peng, Y., Xue, W., Su, T., & Cribb, M. (2021). Reconstructing 1-km-resolution high-quality PM2. 5 data records from 2000 to 2018 in China: spatiotemporal variations and policy implications. *Remote Sensing of Environment*, *252*, 112136.

[158]. Wen, C., Zhang, J., Zheng, K., Li, H., Ling, L., Meng, Z., Fu, W., & Yan, B. (2023). Accelerated verification method for the reliability of the motor drive mechanism of the corn precision seed-metering device. *Computers and Electronics in Agriculture*, *212*, 108163.

[159]. Xia, Y., Xu, Y., Gou, B., & Deng, Q. (2021). A learning-based method for speed sensor fault diagnosis of induction motor drive systems. *IEEE Transactions on Instrumentation and Measurement*, *71*, 1-10.

[160]. Yakkati, R. R., Yeduri, S. R., Tripathy, R. K., & Cenkeramaddi, L. R. (2023). Multi-Channel Time-Frequency Domain Deep CNN Approach for Machinery Fault Recognition Using Multi-Sensor Time-Series. *IEEE Access*, *11*, 116570-116580.

[161]. Yao, X., Zhou, J., Lin, Y., Li, Y., Yu, H., & Liu, Y. (2019). Smart manufacturing based on cyber-physical systems and beyond. *Journal of Intelligent Manufacturing*, *30*(8), 2805-2817.

[162]. Zamal Haider, S., & Mst. Shahrin, S. (2021). Impact Of High-Performance Computing In The Development Of Resilient Cyber Defense Architectures. *American Journal of Scholarly Research and Innovation*, *1*(01), 93–125. https://doi.org/10.63125/fradxg14

[163]. Zhang, M., He, X., Qin, F., Fu, W., & He, Z. (2019). Service quality measurement for omni-channel retail: scale development and validation. *Total Quality Management & Business Excellence*, *30*(sup1), S210-S226.

[164]. Zhang, W., Yang, D., & Wang, H. (2019). Data-driven methods for predictive maintenance of industrial equipment: A survey. *IEEE systems journal*, *13*(3), 2213-2227.

[165]. Zhou, L., Jiang, Z., Geng, N., Niu, Y., Cui, F., Liu, K., & Qi, N. (2022). Production and operations management for intelligent manufacturing: A systematic literature review. *International Journal of Production Research*, *60*(2), 808-846.

[166]. Zulqarnain, F. N. U. (2022). Policy Optimization for Sustainable Energy Security: Data-Driven Comparative Analysis Between The U.S. And South Asia. *American Journal of Interdisciplinary Studies*, *3*(04), 294-331. https://doi.org/10.63125/v4e4m413

[167]. Zulqarnain, F. N. U., & Subrato, S. (2021). Modeling Clean-Energy Governance Through Data-Intensive Computing And Smart Forecasting Systems. *International Journal of Scientific Interdisciplinary Research*, *2*(2), 128–167. https://doi.org/10.63125/wnd6qs51

[168]. Zulqarnain, F. N. U., & Subrato, S. (2023). Intelligent Climate Risk Modeling For Robust Energy Resilience And National Security. *Journal of Sustainable Development and Policy*, *2*(04), 218-256. https://doi.org/10.63125/jmer2r39