



Svedberg
DISSEMINATION OF KNOWLEDGE

Open

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Digital Twin-Enabled Predictive Maintenance Framework for Intelligent Fault Detection and Remaining Useful Life Prediction in Electrical Systems

Saurabh Kumar¹, Anand Ranjan², Shambhu Kumar³, Dr. Ramesh Khima Piprotar⁴, Dr. Amol Jagdish Shakadwipi⁵

¹Assistant Professor, Department of Electrical Engineering, Government Engineering College Sheikhpura, Bihar.
E-mail: saurabhk248@gmail.com

²Assistant Professor, Department of Electrical Engineering, Government Engineering College Buxar, Bihar.
E-mail: anandranjan9534@gmail.com

³Assistant Professor, Department of Electrical Engineering, Government Engineering College Sheikhpura, Bihar.
E-mail: kumarshambhu3011@gmail.com

⁴Assistant Professor of Balaji College of Commerce Science Technology & Management, Shri Govind Guru University –Godhra.
E-mail: rameshpiprotar24@gmail.com

⁵Assistant Professor, SNJB's K.B Jain college of Engineering, Chandwad. E-mail: amolshakadwipi@gmail.com

*Corresponding Author: E-mail: saurabhk248@gmail.com

Abstract

Failures in motors, transformers, switchgear, and converters can propagate from degradation to production interruption, safety exposure, revenue loss, and reputational damage. Conventional monitoring detects abnormalities but often lacks a synchronized asset model, quantified uncertainty, and a defensible link between technical evidence and managerial action. This paper develops a standards-informed digital twin-enabled prognostics and health management framework that integrates intelligent fault detection and probabilistic remaining useful life (RUL) prediction with lifecycle economics and strategic asset management. The framework couples an updated state-space twin with physics residuals, temporal learning, Bayesian estimation, explainability, risk-based maintenance optimization, and a managerial value layer. It specifies sensing, data quality, semantic context, hybrid model orchestration, uncertainty, cybersecurity, human authorization, asset criticality, spare-parts coordination, investment appraisal, and value-realization governance. Validation combines asset-disjoint chronological testing, fault-severity and regime stratification, calibration, warning lead time, decision utility, lifecycle cost, and organizational adoption indicators. Ten tables map failure modes, signals, twin layers, models, validation gates, economic metrics, stakeholders, and commercial options; three figures present the architecture, analytics pipeline, and deployment loop. Because no empirical dataset is supplied, no measured accuracy, cost saving, or return on investment is claimed. The contribution is a testable socio-technical framework for electrical assets that converts condition evidence into auditable engineering, financial, and managerial decisions without confusing a static model with an operational twin.

Keywords: digital twin, predictive maintenance, fault diagnosis, remaining useful life, strategic asset management, lifecycle economics, digital servitization, maintenance decision-making

1. Introduction

Failures happen in electrical systems due to combined thermal, electrical, mechanical, dielectric, and environmental phenomena. Motor faults impact vibration and current spectra; transformer faults influence gas, moisture, and temperature dynamics; switchgear faults manifest themselves in time, coil current, movement, or partial discharge; and converter mission profiles lead to cumulative failure of semiconductors and capacitors. Due to dependence on load and environmental variables, fixed alerts may erroneously report fault detection due to regime changes (Benbouzid, 2000; Nandi et al., 2005; Henao et al., 2014; Wang et al., 2012).

Condition-based maintenance is an analysis of diagnostic and intervention data, while prognostics and health management (PHM) involve additional forecasting of equipment health condition and RUL. Consistent reviews state that a PHM system needs not only a classifier but also data acquisition, construction of health indicators, detection of degradation stages, and maintenance planning (Jardine et al., 2006; Si et al., 2011; Lei et al., 2018). Machine learning extends capabilities for pattern recognition, but its performance in industry is limited by rare failures, label noise, domain shift, class imbalance, and insufficient integration with maintenance processes (Carvalho et al., 2019; Zonta et al., 2020; Theissler et al., 2021).

Digital twin can only be meaningful if its virtual model is linked to the physical asset and kept up to date through operational data. Static simulation and dashboard are the digital model, not the twin of operation. The research differentiates models, digital shadows one way, and connected twins bi-directionally by means of identity, synchronization, purpose, and traceability of their life cycle (Grieves & Vickers, 2017; Kritzinger et al., 2018; Tao et al., 2019; Jones et al., 2020). Predictive twin is about putting evidences in an asset state estimation and making projections of potential degradation paths (Fuller et al., 2020; Rasheed et al., 2020; Kapteyn et al., 2021).

The gap is an end-to-end solution that brings together physical electrical-asset properties, multirate telemetry, diagnostics, uncertainty, RUL, maintenance usefulness, lifecycle economics, organizational capability, and OT cybersecurity. Literature reviews identify computational cost, heterogeneous data, modeling complexity, and limited reference architectures as recurring barriers (Errandonea et al., 2020; van Dinter et al., 2022; Zhong et al., 2023). Management research adds a second gap: predictive information does not create value unless it is translated into budget priorities, work orders, spare-parts availability, workforce routines, supplier coordination, and accountable value-capture mechanisms. Electrical assets make this integration especially demanding because signals, failure consequences, intervention authority, and commercial exposure differ across asset classes and operating contexts.

From a commerce and management perspective, a digital twin is therefore both a technical capability and an organizational investment. Its business case must compare acquisition, integration, cybersecurity, data-governance, training, and recurring model-management costs with risk-adjusted benefits such as avoided downtime, deferred replacement, lower emergency procurement, improved service continuity, and new digitally enabled service offerings. ISO 55000:2024 emphasizes lifecycle value and alignment between asset-management and organizational objectives, while ISO/TS 55010:2024 addresses alignment between financial and non-financial functions. These principles motivate a framework in which technical predictions are evaluated through measurable value, risk, and accountability rather than accuracy alone (ISO, 2024a, 2024b; Meng et al., 2022).

Accordingly, this paper asks four questions: (RQ1) How should a digital twin-PHM architecture represent heterogeneous electrical assets while preserving a common decision interface? (RQ2) How can physics-based and data-driven models be combined to detect faults and estimate RUL under sparse failures and changing operating regimes? (RQ3) What validation and governance gates are required before predictions are used for maintenance? (RQ4) How can technical evidence be translated into lifecycle value, resource allocation, organizational adoption, and commercially viable service models? The paper contributes (i) a layered architecture aligned with digital-twin, PHM, and asset-management standards; (ii) a hybrid probabilistic formulation for diagnosis and RUL; (iii) asset-specific instantiation maps; (iv) a safety-gated evaluation and implementation protocol; and (v) a managerial value-realization layer connecting predictions with lifecycle cost, capital planning, supply-chain decisions, performance indicators, and digital servitization.

Table 1 Maintenance paradigms and the additional obligations created by a digital twin-PHM system

Paradigm	Primary question	Typical output	Critical limitation / added obligation
Reactive	What failed?	Repair after loss of function	High downtime; no advance warning
Preventive	When is the scheduled interval?	Time- or usage-based task	May replace healthy parts or miss accelerated degradation
Condition-based	Is present condition abnormal?	Alarm, diagnosis, condition index	Often stops at detection; regime dependence must be controlled
Predictive / PHM	When will a functional limit be reached?	RUL estimate and risk	Needs degradation evidence, uncertainty, and decision validation
Digital twin-enabled PHM	How will this identified asset evolve under candidate actions?	Synchronized state, fault posterior, RUL distribution, utility-ranked advice	Requires identity, bidirectional synchronization, model governance, security, and human authority

Note. Synthesized from Jardine et al. (2006), Lei et al. (2018), Kritzinger et al. (2018), ISO 13381-1 (2015), and van Dinter et al. (2023). The entries are conceptual distinctions, not comparative performance results.

Table 1 clarifies why attaching a prediction model to a dashboard is insufficient: the twin adds asset identity, synchronized state, counterfactual simulation, and governed decision feedback.

2. Technical, Economic, and Management Foundations

2.1 Digital twin, digital thread, and fidelity

The concept of digital twin has evolved from lifecycle engineering and simulation to a cyber-physical system that consists of the asset, virtual twin, data, and decision service (Glaessgen & Stargel, 2012; Qi & Tao, 2018; Tao et al., 2019). ISO 23247-1:2021 standardizes observable aspects and the twin itself, whereas the NIST guidance provides categorization of its applications and data flows (International Organization for Standardization [ISO], 2021; Shao, 2021). Fidelity should always be problem-specific; for example, motor twins to separate the rotor asymmetry from the load oscillation require different states and sampling rate than transformers with their insulation aging twins (Rasheed et al., 2020).

It is better to choose a probabilistic twin rather than a deterministic mirror due to such factors as noise of sensors, presence of damage unknown to the observer, uncertain parameters, and model-form error. The concept of predictive twin can be defined as two dynamic systems, i.e., physical and digital one in the framework of probabilistic graphical model by Kapteyn et al. (2021). Such approach allows for sequential state update, uncertainty inclusion, and action choosing. In addition, physics-based twin construction may be used to create residuals and synthetic degradation paths. However, synthetic data cannot be regarded as field evidence (Aivaliotis et al., 2019).

2.2 Fault diagnosis and RUL as related but distinct tasks

Detection determines if there is deviation from the permissible healthy envelope; isolation pinpoints the faulty subsystem or mechanism; severity assessment gauges the extent of the damage; and prognosis estimates the time before reaching a certain level of failure. A particular model can identify the fault type correctly but yield useless RUL estimates if the fault is non-monotonic in nature. Conversely, an RUL model may track degradation without isolating its physical cause. The framework therefore keeps diagnosis and prognosis as separate services linked by a shared latent health state. ISO 13374-1 provides a layered view from data acquisition to advisory generation, ISO 17359 sets condition-monitoring program guidance, and ISO 13381-1 frames the prognosis process (ISO, 2003, 2015, 2018).

Table 2 Electrical asset failure modes, observables, and twin state requirements

Asset	Representative failure mechanisms	Useful observables	Twin states / context
Induction / synchronous motor	Bearing damage; broken rotor bars; eccentricity; stator inter-turn fault; cooling loss	Three-phase current and voltage; vibration; speed; flux; temperature; acoustic emission	Slip, torque, load, speed, thermal network, electromagnetic parameters, bearing health
Oil-filled transformer	Partial discharge; arcing; oil/cellulose thermal fault; moisture; bushing or tap-changer degradation	Dissolved gases; moisture; top-oil and winding temperature; load; partial discharge; dielectric response	Thermal hot spot, insulation age, gas-generation trend, moisture equilibrium, loading history
Circuit breaker / switchgear	Mechanism friction; coil or latch defect; contact wear; insulation defect; gas leakage	Trip/close coil current; travel; operating time; motor current; vibration; contact resistance; partial discharge; gas pressure	Mechanism timing, stroke, contact erosion, operation count, ambient and control-voltage context
Converter / inverter	Bond-wire or solder fatigue; gate-driver fault; DC-link capacitor aging; thermal-interface degradation	Junction/case temperature; current; voltage; switching transients; ESR/capacitance; cooling conditions	Mission profile, thermal cycles, loss model, damage accumulation, control mode

Note. Fault-signal mappings are grounded in Benbouzid (2000), Nandi et al. (2005), Henao et al. (2014), Duval (2003), IEC 60599 (2022), IEEE C57.104-2019, Hosseini et al. (2018), and Wang et al. (2012, 2014). Sensor availability and thresholds must be adapted to the actual installation.

As summarized in Table 2, a common software architecture can be reused, but the physical states, observables, sampling rates, and failure thresholds cannot be copied blindly between asset classes.

2.3 Strategic asset management, lifecycle economics, and digital servitization

Strategic asset management frames maintenance as a value decision across cost, risk, performance, sustainability, and organizational objectives. Under ISO 55000:2024, the relevant outcome is not simply high technical availability but value realized from assets over their life cycles. Consequently, the digital twin should connect health evidence to asset criticality, budget cycles, replacement planning, service obligations, and financial and non-financial performance indicators. ISO/TS 55010:2024 further supports a shared decision language between engineering, operations, finance, and senior management (ISO, 2024a, 2024b).

Lifecycle economics requires incremental analysis: the expected benefits of earlier detection and better intervention timing must be compared with sensor, platform, integration, cloud or edge infrastructure, cybersecurity, training, validation, and model-governance costs. Economic evaluation should account for adoption dynamics and interactions among technology providers, asset owners, maintainers, and users rather than assuming that model accuracy automatically produces savings (Meng et al., 2022). Decision metrics therefore include net present value (NPV), return on investment (ROI), discounted payback period, avoided-loss value, cost per avoided failure, and risk-adjusted maintenance cost.

The same connected capability can support digital servitization. Equipment suppliers and service organizations may offer remote condition monitoring, availability guarantees, subscription analytics, or outcome-based maintenance contracts. Such models alter the value proposition, delivery ecosystem, revenue logic, data rights, liability, and customer relationship. Research on industrial digital servitization shows that IoT-enabled information can support business-model innovation, but value creation depends on co-creation, ecosystem coordination, and an explicit value-capture design (Paiola & Gebauer, 2020; Sjödin et al., 2020).

2.4 Literature synthesis and research gap

Digital-twin maintenance reviews report rapid growth in monitoring, diagnosis, and RUL use cases, yet many studies remain prototype-specific and provide limited evidence about uncertainty calibration, cross-regime generalization, maintenance utility, secure deployment, organizational adoption, or financial value realization (Errandonea et al., 2020; van Dinter et al., 2022). Reference architectures improve consistency, but domain-level adaptation remains necessary (van Dinter et al., 2023). Electrical-machine literature provides mature signal-processing and fault-signature knowledge, while power-equipment standards provide diagnostic interpretation rules. Asset-management and digital-servitization research explains lifecycle value, cross-functional alignment, capability development, and commercial value capture. The missing bridge is a stateful hybrid architecture that preserves engineering mechanisms while linking predictions to portfolio priorities, resources, stakeholder accountability, and business-model choices.

Table 3 Evidence streams, strengths, and unresolved integration gaps

Evidence stream	Established strength	Unresolved gap addressed here
Digital twin foundations	Identity, synchronization, lifecycle models, simulation and decision services	Translate generic twin concepts into maintenance states and promotion gates for electrical assets
Predictive maintenance / PHM	Condition indicators, prognosis methods, and maintenance decision logic	Make asset state and uncertainty continuously updateable through the twin
Electrical fault diagnosis	Physics-linked signatures and standards for motors, transformers, switchgear, and converters	Fuse heterogeneous signatures without discarding operating-regime context
Machine and deep learning	Nonlinear feature learning and temporal pattern recognition	Control imbalance, domain shift, calibration, explainability, and rare-failure evidence
OT and AI governance	Risk controls, security safeguards, traceability and human oversight	Tie model promotion and maintenance authority to measurable assurance gates
Strategic asset management and lifecycle economics	Lifecycle value, criticality, financial/non-financial alignment and investment appraisal	Connect predictive evidence to portfolio prioritization, budget decisions, and value-realization baselines
Digital servitization and business models	Connected services, co-creation, ecosystem delivery, subscriptions and outcome-based contracts	Specify data rights, liability, pricing logic, customer value and evidence needed for commercial claims

Note. Synthesis informed by Fuller et al. (2020), Carvalho et al. (2019), Liu et al. (2018), Arrieta et al. (2020), Stouffer et al. (2023), and Tabassi (2023).

Table 3 converts the literature review into design obligations rather than treating the sources as a list of algorithms.

3. Framework Development Method

This conceptual framework uses a standards-informed integrative synthesis; it claims neither a PRISMA review nor experimental accuracy. Functional requirements came from digital-twin foundations, maintenance reviews, and reference architectures. PHM functions were aligned with ISO 13374-1, ISO 13381-1, and ISO 17359; managerial value requirements were aligned with ISO 55000:2024 and ISO/TS 55010:2024. Asset states and observables were mapped from electrical-machine, transformer, switchgear, and power-electronics sources. Economic evaluation, cross-functional responsibility, organizational capability, digital-servitization, assurance, uncertainty, explainability, cybersecurity, and human-authorization gates were then added so that a technically functional prototype is not mistaken for a deployable or commercially justified system.

The design unit is an identified asset nested within an asset portfolio and an organizational value system. Each twin has an identifier, configuration, sensor map, model lineage, histories, state posterior, decision authority, cost center, criticality class, service obligation, and value baseline. Class models may be shared, but calibration and state beliefs remain asset-specific. Field observations, inspection labels, maintenance observations, laboratory observations, simulations, and expert rule observations retain provenance and confidence tags; however, simulation-based faults may increase coverage but cannot determine field failure rates, cash savings, or ROI. Business-case assumptions are versioned alongside model assumptions so that technical and financial claims can be audited together.

4. Proposed Digital Twin-PHM Architecture

The proposed four-layer architecture and its technical, governance, and managerial-value overlays are depicted in Figure 1. Figures and tables used in this paper have been developed by the authors from the cited literature; no fabricated empirical or financial results are included.




Figure 1. Proposed digital twin-enabled PHM architecture for heterogeneous electrical assets.

4.1 Physical asset and sensing layer

This layer includes the asset, protection and control equipment, monitoring sensors, and contextual information. Sensors are selected according to failure-mode observability rather than information richness.

High rate information like current, voltage, vibration, acoustic, and switching measurements is processed close to the asset; lower rate variables include temperatures, DGA, moisture, counts, inspections, and weather information. All observations have timestamp, units, sensor ID, calibration, quality checks, and operational modes. Protection systems take precedence – the mirror should not delay or supersede the primary protection system.

4.2 Edge, data quality, and semantic layer

It executes filtering, resampling, timing adjustment, missing data verification, plausibility checking, and feature extraction for specific events. The storage of raw waveforms is selective, but feature extractions preserve provenance. A semantic graph connects measurements with components, failure modes, rating, topology,

maintenance, and limitations. Processing blocks from ISO 13374 provide functionality, whereas ISO 23247 provides digital representations and exchange capabilities (ISO, 2003, 2021).

4.3 Hybrid digital twin and state estimation layer

The hybrid twin consists of a reduced-order model in combination with a learned residual model or a learned temporal model: motor electromagnetic or thermal models, transformer hot-spot or aging models, breaker time or wear models, converter electrothermal damage. The physics limits the extrapolation and generates interpretable residuals, while the learning part covers nonlinearities not captured by the models. Theory-driven and physics-driven learning incorporate the relationships or penalties in the governing equations (Karpatne et al., 2017; Raissi et al., 2019).

4.4 Fault, RUL, and decision-service layer

Fault services calculate anomaly scores, fault probabilities, severity, and evidence explanations. RUL services project a distribution of time to one or more functional thresholds under specified future-load scenarios. Decision services combine failure risk with criticality, maintenance duration, spare-parts lead time, workforce capacity, production and customer-service constraints, budget availability, sustainability targets, and safety rules. The output is an advisory package, not an autonomous command: predicted fault, supporting signals, counterevidence, uncertainty interval, expected lead time, model-validity status, ranked actions, resource implications, and estimated risk-adjusted value. Explainability may combine physics residuals with local feature attributions such as SHAP or LIME, but explanation stability, engineering plausibility, and managerial actionability must be tested (Ribeiro et al., 2016; Lundberg & Lee, 2017; Arrieta et al., 2020).

Table 4 Architecture layers, data products, and assurance controls

Layer	Principal functions	Data products	Mandatory controls
Physical and sensing	Measure condition and operational context	Waveforms, trends, events, inspections	Calibration, time synchronization, sensor-health monitoring, protection independence
Edge and semantic data	Quality control, alignment, feature extraction, contextualization	Quality-tagged windows, features, asset graph	Schema/version control, unit checks, provenance, retention policy
Hybrid twin	State estimation, simulation, parameter update, scenario projection	State posterior, residuals, model-validity status	Validity domain, parameter bounds, model registry, drift detection
PHM services	Anomaly detection, diagnosis, severity, RUL and uncertainty	Fault posterior, health index, RUL distribution	Calibration, class imbalance, asset-disjoint testing, explanation checks
Decision and integration	Risk ranking, scenario comparison, CMMS/EAM integration, portfolio prioritization and value tracking	Advisory, work-order recommendation, resource plan, business-case update and evidence bundle	Human approval, financial/non-financial alignment, audit trail, rollback and least privilege

Note. The layer separation adapts ISO 13374-1 (2003), ISO 23247-1 (2021), Kapteyn et al. (2021), and van Dinter et al. (2023) to electrical-asset PHM.

Table 4 makes assurance controls first-class architecture elements instead of afterthoughts added after model training.

Table 5 Twin state, synchronization rate, and fit-for-purpose fidelity by asset class

Asset	Fast state / update	Slow degradation state / update	Fidelity test
Motor	Current, voltage, speed, torque proxy, vibration features; sub-second to seconds	Bearing, rotor, stator and thermal health; minutes to days	Separates fault signatures from load/speed variation across the declared operating envelope
Transformer	Load, temperatures, cooling status; seconds to minutes	Gas-generation, moisture and insulation-age states; hours to months	Reproduces thermal behavior and diagnostic trends within uncertainty, not just single DGA labels

Asset	Fast state / update	Slow degradation state / update	Fidelity test
Circuit breaker	Coil current, travel, vibration and timing per operation	Contact wear, friction, gas/insulation condition; per event to months	Detects timing/mechanism deviations while accounting for control voltage and ambient effects
Converter	Electrical and thermal states; microseconds at control level, aggregated at edge	Thermal-cycle damage and capacitor health; minutes to weeks	Preserves mission-profile stress and predicts condition indicators within validated load/temperature ranges

Note. Update rates are design ranges, not universal requirements. Fidelity is evaluated against the intended maintenance decision, consistent with Rasheed et al. (2020) and ISO 23247-1 (2021).

Table 5 shows why the fastest sensor rate should not dictate the synchronization rate of every twin state.

5. Mathematical Formulation

5.1 Coupled asset-twin state estimation

Let x_t denote hidden physical health and operating states, u_t control and load inputs, θ_t asset parameters, y_t measured signals, and c_t contextual variables. A stochastic state-space representation is:

$$x_{t+1} = f_{\text{phys}}(x_t, u_t, c_t, \theta_t) + r_{\phi}(z_t) + w_t \quad (1)$$

$$y_t = h(x_t, u_t, c_t) + v_t \quad (2)$$

Here, f_{phys} is a mechanistic transition model, r_{ϕ} is a learned residual driven by recent observations z_t , and w_t and v_t represent process and measurement uncertainty. Sequential Bayesian filtering updates $p(x_t, \theta_t | y_{1:t}, u_{1:t}, c_{1:t})$. This is consistent with the coupled probabilistic view of predictive twins proposed by Kapteyn et al. (2021). If the system operates outside the model validity domain, the twin must raise an abstention or low-confidence state rather than extrapolate silently.

5.2 Intelligent fault detection and isolation

A regime-conditioned residual $e_t = y_t - \hat{y}_t$ is calculated from the twin. An anomaly score can combine normalized physics residuals, learned reconstruction error, and sensor-quality penalties:

$$A_t = \alpha \|\Sigma_e^{-1/2} e_t\|_2^2 + \beta L_{\text{rec}}(z_t) + \gamma Q_t \quad (3)$$

A persistent alarm is generated only when A_t crosses a regime-specific threshold for a defined duration and data quality is acceptable. A temporal encoder then estimates $p(F, =, k | z_{1:t}, x_t, c_t)$, where k includes healthy, known fault classes, and unknown/other. Long short-term memory networks can model temporal relationships, whereas attention models can model long sequences; but the former should not be employed without simple baselines and ablation studies (Hochreiter & Schmidhuber, 1997; Vaswani et al., 2017; Zhao et al., 2019). When it comes to rare faults, precision and recall provide more information than the ROC area (Saito & Rehmsmeier, 2015).

5.3 Probabilistic RUL prediction

For failure mode k , let D_t be a monotonic or piecewise degradation state and $D_{\text{crit},k}$ the functional threshold. RUL is the first-passage time:

$$\text{RUL}_{t,k} = \inf\{\tau > 0: D_{t+\tau} \geq D_{\text{crit},k} | I_t\} \quad (4)$$

The information set I_t includes the state posterior, maintenance history, intended mission profile, and model version. The framework outputs quantiles or a survival function, not only a point estimate. Epistemic uncertainty can be approximated by ensembles or Monte Carlo dropout, while aleatory uncertainty arises from measurement noise and future-use variability (Gal & Ghahramani, 2016). Calibration must be tested using prediction-interval coverage, interval width, Brier score for horizon-based failure risk, and proper scoring rules (Brier, 1950; Gneiting & Raftery, 2007).

5.4 Risk-based maintenance and resource-allocation decision

For candidate maintenance action a at decision time t , the decision layer minimizes expected discounted lifecycle loss subject to safety, workforce, spare-parts, production, budget, and regulatory constraints. Candidate actions include continued operation, targeted inspection, load derating, planned repair, replacement, or acquisition of additional evidence. The loss function includes direct maintenance expenditure, production or service interruption, failure consequences, safety and compliance penalties, environmental impact, and the opportunity cost of scarce resources.

$$a^* = \mathbb{E}[C_m(a) + C_d(a) + C_f \mathbf{1}(T_f < T_a) + C_s(a) \mid I_t] \quad (5)$$

The optimization does not convert uncertain predictions into false precision. Failure probabilities and RUL distributions must be propagated through cost and consequence scenarios, and the recommended action should remain robust under plausible changes in demand, energy price, spare lead time, and intervention duration. If uncertainty is high, targeted inspection or additional sensing may dominate immediate replacement. Critical protection decisions remain outside the twin's authority, and financially attractive actions cannot override safety constraints.

5.5 Business-case and value-realization model

The digital-twin investment is assessed incrementally against an agreed baseline such as corrective, preventive, or conventional condition-based maintenance. For evaluation horizon T and discount rate r , NPV equals the discounted stream of avoided failure, downtime, emergency procurement, energy, quality, and premature replacement costs plus additional service revenue, less implementation and recurring operating costs. ROI is the net discounted benefit divided by discounted investment cost. These measures are reported as distributions or scenarios when benefits depend on uncertain failure incidence and intervention effectiveness (Meng et al., 2022). A benefits register assigns each value mechanism to an owner, baseline, measurement frequency, evidence source, and realization date. Engineering metrics such as warning lead time and calibrated failure probability are therefore paired with management metrics such as avoided downtime value, maintenance-plan compliance, inventory turnover for critical spares, emergency purchase frequency, asset-life extension, user adoption, and benefit realization. No commercial claim is promoted unless both the technical causal pathway and the accounting treatment are documented.




Figure 2. Hybrid analytics pipeline from quality-tagged sensor windows to uncertainty-aware maintenance utility.

Figure 2 separates mechanistic and learned evidence before they are fused into fault and RUL estimates, making ablation and failure analysis possible.

Table 6 Recommended analytics by task, evidence condition, and failure mode

Task	Candidate method	When appropriate	Required challenge test
Healthy-envelope detection	Robust statistics, one-class SVM, autoencoder, physics residual	Few labeled faults; stable healthy coverage	Unseen loads, sensor drift, missing channels, false alarms per operating hour
Fault isolation	Gradient-boosted trees, CNN, temporal network, probabilistic classifier	Fault labels or credible laboratory evidence exist	Asset-disjoint split, severity stratification, unknown-fault rejection, probability calibration
Health indicator	Physics-derived indicator, monotonic representation learning, latent state filter	Degradation trend is observable	Monotonicity, trendability, prognosability and reset after maintenance

Task	Candidate method	When appropriate	Required challenge test
RUL	Wiener/gamma process, particle filter, survival model, LSTM/transformer, hybrid first-passage model	Threshold and degradation trajectory are defined	Right censoring, late-life scarcity, scenario shift, interval coverage
Explanation	Residual decomposition, SHAP/LIME, counterfactual and rule trace	Operator must understand evidence and counterevidence	Stability, physical plausibility, actionability, no leakage of post-failure features
Uncertainty	Bayesian filter, ensemble, MC dropout, conformal interval	Decision cost is asymmetric or domain shift likely	Calibration, sharpness, out-of-distribution abstention, coverage by regime

Note. Model families are options, not prescriptions. Selection should follow baseline comparison and ablation. See Si et al. (2011), Liu et al. (2018), Lei et al. (2018), Gal and Ghahramani (2016), and Arrieta et al. (2020).

Table 6 ties each model family to an evidence condition and a failure-oriented challenge test, preventing algorithm selection by novelty alone.

6. Asset-Specific Instantiation

6.1 Rotating electrical machines

For motors and generators, the twin should estimate operating regime before interpreting signatures. Motor-current signature analysis, negative-sequence components, Park-vector features, vibration envelopes, temperature, and flux can support rotor, stator, bearing, and eccentricity diagnosis (Benbouzid, 2000; Nandi et al., 2005; Henao et al., 2014). One such effective combination would involve the use of electromagnetic and thermal models along with a time-based classifier. Fault probabilities would be contingent upon speed, slip, torque, power supply, and operating conditions of the drive system. The RUL estimation of the bearing can have an entirely different degradation model since the signatures differ.

6.2 Oil-filled transformers

A transformer diagnosis must not limit the DGA analysis to a simplistic multi-class output. The twin sustains gas levels and rates, sampling periods, load, thermal conditions, oil handling and maintenance history, moisture, and analytical uncertainty. Interpretation guidelines for IEC 60599:2022 and IEEE C57.104-2019 are offered, while graph plotting techniques of Duval allow interpreting the fault zones (Duval, 2003). The PHM framework combines the standard-based knowledge with thermal aging models and trend-based classifier. The RUL is evaluated based on specific criteria like the life loss limit for insulation, bushing risk level, or tap changer decision rule; the RUL for the transformer as a whole does not make sense.

6.3 Circuit breakers and switchgear

Twin circuit breakers are event-based. Each event triggers updates to timing, coil current segment, movement, speed, vibration, contact wear equivalent, control voltage, gas pressure, and environmental factors. The trip coil signal may show problems in the mechanisms or control circuits and may be used for machine learning diagnosis where traditional time-based feature extraction fails to capture abnormal shapes (Hosseini et al., 2018). In addition, the twin should also regard periods with no events as uncertainty increase, not proof of good condition. Insulation threats have their own health states with individual thresholds.

6.4 Power-electronic converters

Converter twin models generate loss, temperature cycling, and fatigue from mission profiles. Condition indicators encompass on-voltage, thermal impedance, switch dynamics, capacitor ESR and capacitance, cooling efficiency, and control residuals. Physics of failure is essential since accelerated testing and actual mission profiles seldom correspond. Mission-profile-based approach has been suggested by reliability research (Wang et al., 2014). Residuals may be used to update loss and thermal models, but state update should keep switching regime, environmental temperature, and cooling regime the same.

7. Validation Protocol and Evaluation Metrics

The system should be assessed in stages. The first stage assesses sensors and the data pipeline. The second assesses the mechanistic model and state estimator through controlled tests. The third trains and evaluates diagnostic and prognostic models using asset-disjoint, time-ordered datasets. The fourth uses shadow mode, where predictions cannot initiate maintenance. The fifth conducts a prospective human-in-the-loop advisory pilot. In parallel, a managerial evaluation compares costs and benefits with the pre-registered maintenance baseline, tests workflow

adoption, and confirms whether spare, workforce, and budget decisions actually change. Commercial evaluation is permitted only after technical safety and validity gates are satisfied.

Table 7 Proposed multi-stage validation design

Stage	Evidence and split	Primary tests	Exit condition
1. Data pipeline	Recorded and injected sensor/data faults	Time alignment, units, missingness, calibration, latency, failover	Quality faults detected; provenance and replay verified
2. Twin model	Bench/lab tests plus normal field operation	Parameter identifiability, residual whiteness, thermal/electrical state error, validity envelope	Errors bounded for intended decision and invalid regions flagged
3. Offline PHM	Asset-disjoint and chronological splits; faults stratified by type/severity/regime	Detection lead time, PR-AUC, macro-F1, calibration, RUL error and interval coverage	Predefined criteria met without test-set tuning
4. Shadow deployment	Live data; no maintenance authority	False alarms per asset-month, drift, latency, abstention, operator review	Stable performance across seasons and operational regimes
5. Human-in-loop pilot	Prospective advisory trial	Utility, avoided downtime, unnecessary actions, explanation use, safety events	Independent review approves bounded use and rollback

Note. The protocol adapts PHM evaluation principles from Lei et al. (2018), industrial ML deployment concerns from Theissler et al. (2021), and risk governance from Tabassi (2023).

Table 7 makes shadow operation mandatory before maintenance authority is expanded.

Table 8 Evaluation metrics and proposed acceptance logic

Capability	Metrics	Acceptance logic (pre-registered, asset-specific)
Data quality	Completeness; timestamp error; sensor-health detection rate; latency	No silent unit/time mismatch; degraded inputs trigger quality state or abstention
Fault detection	Event precision/recall; PR-AUC; macro-F1; false alarms per asset-month; detection delay	Meet minimum recall for critical faults while staying within an operational false-alarm budget
Fault probability	Brier score; expected calibration error; reliability diagram	Predicted probabilities remain calibrated overall and by operating regime
RUL point accuracy	MAE/RMSE; relative error; asymmetric score; horizon-specific error	Beat persistence, population median, and single-model baselines on untouched assets
RUL uncertainty	Coverage; mean interval width; weighted interval score	Coverage near nominal with intervals sharp enough to affect decisions
Decision value	Expected cost; avoided downtime; unnecessary maintenance; lead-time utility	Positive prospective utility with no increase in safety incidents
Robustness	Performance by asset, severity, regime, season, missing channel and OOD case	No hidden subgroup collapse; invalid regimes produce abstention
Economic value	NPV; ROI; payback; avoided-loss value; cost per avoided failure	Positive risk-adjusted value relative to the agreed baseline, with assumptions and sensitivity ranges disclosed
Operational management	Schedule compliance; emergency work ratio; critical-spares service level; planning lead time	Predictions improve resource coordination without shifting hidden costs or degrading service
Organizational adoption	Advisory use; override rate and rationale; training completion; decision-cycle time	Qualified users understand, contest, document and appropriately act on recommendations

Note. Numerical cutoffs are intentionally not fabricated. They must be pre-registered from asset criticality and maintenance economics. Precision-recall guidance follows Saito and Rehmsmeier (2015); probabilistic evaluation follows Brier (1950) and Gneiting and Raftery (2007).

Table 8 deliberately avoids universal accuracy thresholds; an acceptable false-alarm rate or RUL error depends on asset criticality, failure consequence, inspection cost, and available lead time.




Figure 3. Safety-gated validation and maintenance decision loop with evidence returned after intervention.

As shown in Figure 3, maintenance findings return to the evidence store and may trigger recalibration, but retraining does not automatically promote a new model.

8. Cybersecurity, Governance, Organizational Adoption, and Commercial Control

A networked twin expands the OT attack surface since any compromise to the sensor, broker, model artifact, or CMMS API may lead to erroneous state estimates and work orders. As such, the architecture employs network segmentation, device identity authentication, encryption where possible, least privilege, signing of model artifacts, use of reliable time sources, logging, backups and recovery, vulnerability management, and testing of the fallback manual option. NIST Special Publication 800-82 Rev. 3 stresses that OT security must maintain performance, reliability, and safety (Stouffer et al., 2023). IEC 62443-3-3 offers system security requirements and security levels for structuring the controls choice (International Electrotechnical Commission [IEC], 2013).

Governance of AI is no less important. Each model card shall include intended use, unintended use, training evidence, class balance, validity range, subgroup performance, uncertainty approach, explanation method, known failure cases, economic assumptions, accountable owner, and rollback plan. The NIST AI Risk Management Framework organizes trustworthy AI work around govern, map, measure, and manage functions (Tabassi, 2023). For safety-significant electrical assets, the responsible engineer remains accountable; a confidence score or financial benefit estimate cannot substitute for a qualified decision. Operators should be trained to interpret supporting and contradictory evidence and to record reasons for accepting or rejecting an advisory.

Organizational implementation requires sponsorship, cross-functional process ownership, role clarity, competence development, and incentives that do not reward unsafe deferral or unnecessary maintenance. ISO 55012:2024 emphasizes people involvement and competence in asset management. A RACI structure should assign responsibility for data quality, model approval, maintenance execution, financial validation, cybersecurity, procurement, supplier management, and benefit reporting. Worker consultation is essential because experienced maintainers provide contextual knowledge and can identify false explanations or impractical recommendations (ISO, 2024c).

Table 9 Implementation roadmap, deliverables, and promotion gates

Phase	Core deliverable	Promotion gate	Failure if skipped
1. Scope and criticality	Asset register, failure modes, stakeholders, decision rights, baseline costs and value hypothesis	Named technical and business owners; endpoints, costs and benefits defined	A technically impressive model solves the wrong maintenance problem
2. Instrument and contextualize	Sensor map, time sync, quality rules, asset graph	Replayable, quality-tagged data with provenance	False alarms and data leakage remain invisible
3. Build minimum viable twin	Reduced-order physics, state estimator, baseline alarms	Validated within a declared operating envelope	Complexity grows without identifiable states
4. Add fault and RUL models	Calibrated diagnosis, uncertainty-aware RUL, explanations	Asset-disjoint validation beats simple baselines	Novel model appears accurate because of leakage or imbalance
5. Shadow and pilot	Live evidence bundles, human-reviewed advisories and benefits register	Operational, adoption, utility, economic and safety criteria met	Offline metrics are mistaken for field value
6. Scale and govern	Registry, monitoring, drift response, cyber controls, training, portfolio governance and service model	Independent approval, sustainable funding, accountable ownership, rollback and auditability	Model decay or compromise becomes an operational hazard

Note. The roadmap integrates van Dinter et al. (2023), Stouffer et al. (2023), and Tabassi (2023). It is intentionally gate-based rather than calendar-based.

Table 9 provides a practical sequence: start with decision scope and evidence quality, not with a preferred AI algorithm.

9. Discussion

9.1 Theoretical contribution

In the proposed framework, the digital twin is treated as a governed probabilistic state and decision system rather than a visualization. Its theoretical significance lies in the explicit separation and controlled coupling of (a) asset state estimation, (b) fault hypotheses, (c) degradation and RUL, (d) maintenance utility, and (e) organizational value realization. This structure prevents anomaly detection, fault diagnosis, prognosis, and business benefit from being treated as equivalent outcomes. It also connects engineering PHM with strategic asset management by making cost, risk, performance, and stakeholder accountability part of the decision architecture rather than downstream reporting.

9.2 Practical and managerial implications

This approach discourages four weak practices: acquiring signals without a failure-mode map; improving aggregate accuracy while hiding rare but critical failures; issuing a single RUL estimate without a threshold, mission scenario, or interval; and claiming savings without an agreed baseline or auditable benefit pathway. Managers should begin with asset criticality and the decision to be improved, then select the minimum twin fidelity and data infrastructure needed for that decision. A maintainable low-fidelity twin linked to work planning and finance can create more value than an unmanageable high-fidelity model isolated from operations.

Data services, identity, provenance, model registries, evaluation, cybersecurity, training, and governance are reusable organizational capabilities, while physical models and failure states remain asset-specific. A portfolio approach can therefore standardize controls and interfaces without forcing one model across all equipment. Maintenance, operations, finance, procurement, information technology, OT security, and suppliers should jointly review recommendations that affect shutdown timing, critical spares, capital replacement, warranty, or service obligations. Power-system twin research similarly links situational awareness, diagnosis, security, and resilience through model-data interaction rather than a universal model (Wu et al., 2024).

9.3 Commercialization and business-model implications

For asset owners, the primary value proposition is risk-adjusted operational continuity: fewer disruptive failures, better planned outages, longer useful asset life, and more defensible capital allocation. For original-equipment manufacturers and service providers, the same infrastructure can enable tiered remote monitoring, analytics

subscriptions, availability guarantees, shared-savings contracts, and outcome-based maintenance. Digital servitization research indicates that these offerings require simultaneous redesign of value proposition, value delivery, and value capture, with customers and ecosystem partners participating in co-creation (Paiola & Gebauer, 2020; Sjödin et al., 2020).

Commercial deployment also creates governance questions. Contracts should specify data ownership and portability, model-update rights, cybersecurity duties, service-level metrics, responsibility for false alarms and missed detections, audit access, termination assistance, and limits on automated authority. Pricing should reflect verified customer value and risk allocation rather than model sophistication alone. A staged offering can begin with monitoring and advisory services, progress to planning integration, and reach outcome-based contracts only after calibrated field performance and benefit attribution are demonstrated.

Table 10 Managerial value-realization and commercial deployment matrix

Dimension	Value mechanism	Core metric	Accountable stakeholder	Promotion condition
Operational continuity	Earlier, better-timed intervention	Avoided downtime value; service availability	Operations and maintenance	Prospective utility demonstrated without safety degradation
Lifecycle finance	Deferred replacement and lower total ownership cost	NPV; ROI; payback; risk-adjusted lifecycle cost	Finance and asset manager	Baseline, attribution and sensitivity independently reviewed
Supply chain	Forecasted demand for critical spares and planned procurement	Stockout risk; emergency purchase rate; inventory turnover	Procurement and stores	Forecast horizon exceeds actionable replenishment lead time
Workforce and process	Better job planning and knowledge-supported decisions	Adoption; override rationale; schedule compliance	Maintenance manager and HR/training	Competence, consultation and documented decision rights in place
Customer value	Improved uptime, transparency and service assurance	Service-level attainment; verified customer benefit	Commercial/service manager	Customer-accepted metric and data-governance terms agreed
Digital service revenue	Subscription, availability or outcome-based offering	Recurring revenue; margin; renewal; shared savings	Provider and customer contract owners	Field calibration, auditable attribution, liability and exit terms established

Note. The matrix operationalizes lifecycle-value principles in ISO 55000:2024 and financial/non-financial alignment in ISO/TS 55010:2024. Commercial options draw on digital-servitization research (Paiola & Gebauer, 2020; Sjödin et al., 2020). Metrics require organization-specific baselines; no financial performance is claimed here. Table 10 makes commercial scale-up conditional on validated technical performance, credible benefit attribution, organizational readiness, and explicit allocation of data and operational risk.

9.4 Limitations and research agenda

The framework has five limitations. First, it is conceptual and has not been validated on a specific dataset or deployment; therefore, it makes no claims about accuracy, cost reduction, ROI, or revenue. Second, its multi-asset scope requires local specification of failure modes, sensors, thresholds, model validity, criticality, cost structure, and service obligations. Third, RUL may not be identifiable when preventive interventions censor degradation, thresholds change, or damage is non-monotonic. Fourth, security, organizational readiness, incentives, supplier dependence, and workforce acceptance may affect outcomes more than algorithmic performance. Fifth, financial attribution is difficult when maintenance policy, demand, prices, and asset utilization change simultaneously.

Future research requires prospective multi-site validation with asset splits, censoring, lead time, calibration, false alarms per exposure, managerial utility, and risk-adjusted economic outcomes. Studies should ablate physics-only,

data-only, and hybrid twins; test transfer, unknown-fault rejection, sensor attacks, and post-maintenance state reset; and report context and maintenance history rather than waveforms alone. Management research should examine cross-functional decision quality, user adoption, procurement effects, data-governance arrangements, value-capture mechanisms, and the causal attribution of benefits under alternative contracts. Comparative pilots should pre-register both technical and commercial baselines.

10. Conclusion

This study developed a digital twin-enabled predictive maintenance framework for intelligent fault detection, RUL forecasting, lifecycle economics, and strategic management of electrical assets. The framework integrates failure-mode-based sensing, quality-tagged data, semantic asset context, hybrid physics and machine-learning models, Bayesian state updating, calibrated fault and RUL predictions, risk-based maintenance, resource allocation, OT cybersecurity, organizational adoption, and human authorization. Its managerial value layer connects technical evidence with asset criticality, lifecycle cost, capital planning, spare-parts coordination, workforce decisions, benefit realization, and potential digital-service offerings. The architecture is reusable across motors, transformers, switchgear, and converters without assuming identical states, signals, thresholds, sampling rates, or commercial consequences. Evidence discipline remains central: simulated observations are labeled, predictions carry uncertainty and validity information, financial claims require baselines and attribution, and maintenance authority expands only after offline, shadow, prospective, and governance gates. The next step is to implement the framework for one asset class, pre-register engineering and business criteria, and test whether the twin improves safety, maintenance effectiveness, and risk-adjusted organizational value in real operations.

Declarations

Funding: This research received no external funding.

Conflict of interest: The authors declare no conflict of interest.

Data availability: No empirical dataset was generated or analyzed because this paper develops a conceptual framework and validation protocol.

Author contributions: To be completed by the authors according to the target journal's required contribution taxonomy before submission.

References

1. Aivaliotis, P., Georgoulas, K., Arkouli, Z., & Makris, S. (2019). Methodology for enabling digital twin using advanced physics-based modelling in predictive maintenance. *Procedia CIRP*, 81, 417–422. <https://doi.org/10.1016/j.procir.2019.03.072>
2. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
3. Benbouzid, M. E. H. (2000). A review of induction motors signature analysis as a medium for faults detection. *IEEE Transactions on Industrial Electronics*, 47(5), 984–993. <https://doi.org/10.1109/41.873206>
4. Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1), 1–3. [https://doi.org/10.1175/1520-0493\(1950\)078<0001:VOFEIT>2.0.CO;2](https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2)
5. Carvalho, T. P., Soares, F. A. A. M. N., Vita, R., Francisco, R. P., Basto, J. P., & Alcalá, S. G. S. (2019). A systematic literature review of machine learning methods applied to predictive maintenance. *Computers & Industrial Engineering*, 137, 106024. <https://doi.org/10.1016/j.cie.2019.106024>
6. Duval, M. (2003). New techniques for dissolved gas-in-oil analysis. *IEEE Electrical Insulation Magazine*, 19(2), 6–15. <https://doi.org/10.1109/MEI.2003.1192031>
7. Errandonea, I., Beltrán, S., & Arrizabalaga, S. (2020). Digital twin for maintenance: A literature review. *Computers in Industry*, 123, 103316. <https://doi.org/10.1016/j.compind.2020.103316>
8. Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8, 108952–108971. <https://doi.org/10.1109/ACCESS.2020.2998358>
9. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning* (Vol. 48, pp. 1050–1059). PMLR.

10. Glaessgen, E., & Stargel, D. (2012). The digital twin paradigm for future NASA and U.S. Air Force vehicles. In 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2012-1818>
11. Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359–378. <https://doi.org/10.1198/016214506000001437>
12. Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85–113). Springer. https://doi.org/10.1007/978-3-319-38756-7_4
13. Henao, H., Capolino, G.-A., Fernandez-Cabanas, M., Filippetti, F., Bruzzese, C., Strangas, E., Pusca, R., Estima, J., Riera-Guasp, M., & Hedayati-Kia, S. (2014). Trends in fault diagnosis for electrical machines: A review of diagnostic techniques. *IEEE Industrial Electronics Magazine*, 8(2), 31–42. <https://doi.org/10.1109/MIE.2013.2287651>
14. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
15. Hosseini, M., Stephen, B., McArthur, S. D. J., & Helm, J. (2018). Current-based trip coil analysis of circuit breakers for fault diagnosis. In 2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe). IEEE. <https://doi.org/10.1109/ISGTEurope.2018.8571478>
16. IEEE Standards Association. (2019). IEEE guide for the interpretation of gases generated in mineral oil-immersed transformers (IEEE Std C57.104-2019). <https://doi.org/10.1109/IEEESTD.2019.8890040>
17. International Electrotechnical Commission. (2013). Industrial communication networks—Network and system security—Part 3-3: System security requirements and security levels (IEC 62443-3-3:2013).
18. International Electrotechnical Commission. (2022). Mineral oil-filled electrical equipment in service—Guidance on the interpretation of dissolved and free gases analysis (IEC 60599:2022). <https://webstore.iec.ch/en/publication/66491>
19. International Organization for Standardization. (2003). Condition monitoring and diagnostics of machines—Data processing, communication and presentation—Part 1: General guidelines (ISO 13374-1:2003).
20. International Organization for Standardization. (2015). Condition monitoring and diagnostics of machines—Prognostics—Part 1: General guidelines (ISO 13381-1:2015). <https://www.iso.org/standard/51436.html>
21. International Organization for Standardization. (2018). Condition monitoring and diagnostics of machines—General guidelines (ISO 17359:2018).
22. International Organization for Standardization. (2021). Automation systems and integration—Digital twin framework for manufacturing—Part 1: Overview and general principles (ISO 23247-1:2021).
23. International Organization for Standardization. (2024a). Asset management—Vocabulary, overview and principles (ISO 55000:2024). <https://www.iso.org/standard/83053.html>
24. International Organization for Standardization. (2024b). Asset management—Guidance on the alignment of financial and non-financial functions in asset management (ISO/TS 55010:2024).
25. International Organization for Standardization. (2024c). Asset management—Guidance on people involvement and competence (ISO 55012:2024).
26. Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7), 1483–1510. <https://doi.org/10.1016/j.ymssp.2005.09.012>
27. Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020). Characterising the digital twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29, 36–52. <https://doi.org/10.1016/j.cirpj.2020.02.002>
28. Kapteyn, M. G., Pretorius, J. V. R., & Willcox, K. E. (2021). A probabilistic graphical model foundation for enabling predictive digital twins at scale. *Nature Computational Science*, 1(5), 337–347. <https://doi.org/10.1038/s43588-021-00069-0>
29. Karpatne, A., Atluri, G., Faghmous, J. H., Steinbach, M., Banerjee, A., Ganguly, A., Shekhar, S., Samatova, N., & Kumar, V. (2017). Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10), 2318–2331. <https://doi.org/10.1109/TKDE.2017.2720168>
30. Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11), 1016–1022. <https://doi.org/10.1016/j.ifacol.2018.08.474>

31. Lei, Y., Li, N., Guo, L., Li, N., Yan, T., & Lin, J. (2018). Machinery health prognostics: A systematic review from data acquisition to RUL prediction. *Mechanical Systems and Signal Processing*, 104, 799–834. <https://doi.org/10.1016/j.ymssp.2017.11.016>
32. Liu, R., Yang, B., Zio, E., & Chen, X. (2018). Artificial intelligence for fault diagnosis of rotating machinery: A review. *Mechanical Systems and Signal Processing*, 108, 33–47. <https://doi.org/10.1016/j.ymssp.2018.02.016>
33. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30* (pp. 4765–4774). Curran Associates.
34. Meng, H., Liu, X., Xing, J., & Zio, E. (2022). A method for economic evaluation of predictive maintenance technologies by integrating system dynamics and evolutionary game modelling. *Reliability Engineering & System Safety*, 222, 108424. <https://doi.org/10.1016/j.res.2022.108424>
35. Nandi, S., Toliyat, H. A., & Li, X. (2005). Condition monitoring and fault diagnosis of electrical motors—A review. *IEEE Transactions on Energy Conversion*, 20(4), 719–729. <https://doi.org/10.1109/TEC.2005.847955>
36. Paiola, M., & Gebauer, H. (2020). Internet of things technologies, digital servitization and business model innovation in BtoB manufacturing firms. *Industrial Marketing Management*, 89, 245–264. <https://doi.org/10.1016/j.indmarman.2020.03.009>
37. Qi, Q., & Tao, F. (2018). Digital twin and big data towards smart manufacturing and Industry 4.0: 360 degree comparison. *IEEE Access*, 6, 3585–3593. <https://doi.org/10.1109/ACCESS.2018.2793265>
38. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
39. Rasheed, A., San, O., & Kvamsdal, T. (2020). Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE Access*, 8, 21980–22012. <https://doi.org/10.1109/ACCESS.2020.2970143>
40. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). ACM. <https://doi.org/10.1145/2939672.2939778>
41. Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3), e0118432. <https://doi.org/10.1371/journal.pone.0118432>
42. Shao, G. (2021). Use case scenarios for digital twin implementation based on ISO 23247 (NIST AMS 400-2). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AMS.400-2>
43. Si, X.-S., Wang, W., Hu, C.-H., & Zhou, D.-H. (2011). Remaining useful life estimation—A review on the statistical data driven approaches. *European Journal of Operational Research*, 213(1), 1–14. <https://doi.org/10.1016/j.ejor.2010.11.018>
44. Sjödin, D., Parida, V., Kohtamäki, M., & Wincent, J. (2020). An agile co-creation process for digital servitization: A micro-service innovation approach. *Journal of Business Research*, 112, 478–491. <https://doi.org/10.1016/j.jbusres.2020.01.009>
45. Stouffer, K., Pease, M., Tang, C. Y., Zimmerman, T., Pillitteri, V., Lightman, S., Hahn, A., Saravia, S., Sherule, A., & Thompson, M. (2023). Guide to operational technology (OT) security (NIST SP 800-82 Rev. 3). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-82r3>
46. Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
47. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415. <https://doi.org/10.1109/TII.2018.2873186>
48. Theissler, A., Pérez-Velázquez, J., Kettelgerdes, M., & Elger, G. (2021). Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry. *Reliability Engineering & System Safety*, 215, 107864. <https://doi.org/10.1016/j.res.2021.107864>
49. van Dinter, R., Tekinerdogan, B., & Catal, C. (2022). Predictive maintenance using digital twins: A systematic literature review. *Information and Software Technology*, 151, 107008. <https://doi.org/10.1016/j.infsof.2022.107008>
50. van Dinter, R., Tekinerdogan, B., & Catal, C. (2023). Reference architecture for digital twin-based predictive maintenance systems. *Computers & Industrial Engineering*, 177, 109099. <https://doi.org/10.1016/j.cie.2023.109099>
51. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems 30* (pp. 5998–6008). Curran Associates.

52. Wang, H., Liserre, M., & Blaabjerg, F. (2014). Transitioning to physics-of-failure as a reliability driver in power electronics. *IEEE Journal of Emerging and Selected Topics in Power Electronics*, 2(1), 97–114. <https://doi.org/10.1109/JESTPE.2013.2290282>
53. Wang, H., Liserre, M., Blaabjerg, F., de Place Rikken, P., Jacobsen, J. B., Kvisgaard, T., & Landkildehus, J. (2012). Transition from component-level to system-level reliability in power electronic systems. *IEEE Transactions on Industry Applications*, 48(1), 180–191. <https://doi.org/10.1109/TIA.2011.2175880>
54. Wu, Y., Guerrero, J. M., Wu, Y., Bazmohammadi, N., Vasquez, J. C., Botero Cabrera, A., & Lu, N. (2024). Digital twins for microgrids: Opening a new dimension in the power system. *IEEE Power & Energy Magazine*, 22(1), 35–42. <https://doi.org/10.1109/MPE.2023.3324296>
55. Zhao, R., Yan, R., Chen, Z., Mao, K., Wang, P., & Gao, R. X. (2019). Deep learning and its applications to machine health monitoring. *Mechanical Systems and Signal Processing*, 115, 213–237. <https://doi.org/10.1016/j.ymssp.2018.05.050>
56. Zhong, D., Xia, Z., Zhu, Y., & Duan, J. (2023). Overview of predictive maintenance based on digital twin technology. *Heliyon*, 9(4), e14534. <https://doi.org/10.1016/j.heliyon.2023.e14534>
57. Zonta, T., da Costa, C. A., da Rosa Righi, R., de Lima, M. J., da Trindade, E. S., & Li, G. P. (2020). Predictive maintenance in the Industry 4.0: A systematic literature review. *Computers & Industrial Engineering*, 150, 106889. <https://doi.org/10.1016/j.cie.2020.106889>