American Journal of Scholarly Research and Innovation

Volume: 2; **Issue:** 1
Pages: 389– 411
eISSN: 3067-2163

RELIABILITY-CENTERED MAINTENANCE OPTIMIZATION USING MULTI-OBJECTIVE AI ALGORITHMS IN REFINERY EQUIPMENT

Saikat Sarkar[1]; Aditya Dhanekula[2]

[1]. *Master of Engineering in Mechanical Engineering (Continuing), Lamar University, Texas, USA.*
Email: saikat.mac@gmail.com

[2]. *Bachelor of Technology, Mechanical Engineering , VNR VJIET, India*
Email: dhanekulaaditya1@gmail.com

Doi: 10.63125/6a6kqm73

Received: 21 September 2023; *Revised:* 25 October 2023; *Accepted:* 27 November 2023; *Published:* 28 December 2023

Abstract

This quantitative study on Reliability-Centered Maintenance Optimization Using Multi-Objective AI Algorithms in Refinery Equipment had examined how reliability-centered maintenance (RCM) logic could be translated into measurable decision variables and optimized under refinery constraints using multi-objective AI. A structured review of 62 peer-reviewed papers had informed objective selection, constraint definition, and algorithm benchmarking practices. The empirical analysis had used a retrospective refinery case dataset comprising 420 assets observed over 60 months (190 rotating, 170 static, 60 instrumentation/control) with 17,980,000 operating exposure hours. The dataset had included 612 functional failure events, 9,284 work orders (59.0% preventive, 41.0% corrective), and 4,116 inspection observations. Downtime severity had been right-skewed, with median downtime per event of 6.2 hours (IQR 2.1–18.4) and a 95th percentile of 92.0 hours, with static equipment exhibiting longer restoration burdens. Composite indices for operating severity, monitoring coverage, execution burden, degradation evidence, and risk exposure had demonstrated acceptable internal consistency (0.77–0.85). Collinearity screening reduced model predictors and improved stability, lowering maximum VIF from 9.4 to 2.8 in rotating equipment models. Regression results showed that operating severity increased functional failure hazard (HR 1.29, 95% CI 1.18–1.41, $p < 0.001$) and degradation evidence increased hazard (HR 1.33, 95% CI 1.22–1.46, $p < 0.001$), while monitoring coverage reduced hazard (HR 0.83, 95% CI 0.75–0.92, $p = 0.001$). Preventive intensity reduced corrective recurrence (IRR 0.79, 95% CI 0.70–0.90, $p < 0.001$) but increased direct cost (+5.4%, $p = 0.001$). Turnaround proximity reduced unavailability (RR 0.86, $p = 0.001$). Safety-critical assets showed higher inspection adherence (96.1% vs 88.7%; OR 2.72, $p < 0.001$). The findings had clarified the trade-off structure motivating Pareto-efficient maintenance portfolios that remained interpretable and constraint-feasible in refinery operations.

Keywords

Reliability-centered maintenance; Multi-objective AI; Refinery equipment; Maintenance optimization; Risk constraints.

INTRODUCTION

Reliability-centered maintenance is defined as a structured and systematic approach to maintenance decision-making that focuses on preserving the required functions of physical assets within a specified operational context (Piechnicki et al., 2021). In refinery equipment systems, required functions are formally defined in measurable operational terms such as throughput stability, containment integrity, pressure and temperature control, and compliance with safety limits. Reliability, within this framework, is quantified as the probability that equipment performs its intended function for a defined period under stated conditions, making maintenance inherently a probabilistic and data-driven discipline. Refinery environments amplify the importance of this definition because equipment operates under high thermal stress, pressure fluctuations, corrosive chemical exposure, and continuous process demands. Maintenance decisions therefore directly influence operational continuity, safety performance, and economic stability. Reliability-centered maintenance differs from traditional maintenance philosophies by shifting the analytical focus from asset condition alone to the consequences of functional failure, thereby linking maintenance tasks to safety, operational, environmental, and economic outcomes. This consequence-oriented structure provides a logical foundation for quantitative modeling, as failure effects can be translated into measurable loss functions (Arfan et al., 2021; Sifonte & Reyes-Picknell, 2017). Internationally, refinery reliability is a shared industrial concern due to the interconnected nature of energy supply chains, global fuel markets, and cross-border environmental exposure. Equipment failures in refineries contribute to production instability, energy price volatility, and large-scale safety incidents with transnational implications. As a result, maintenance optimization in refineries has evolved into a discipline that integrates reliability engineering, risk analysis, and asset management under a common quantitative structure. Reliability-centered maintenance establishes this structure by formally identifying functional failures, dominant degradation mechanisms, and allowable risk thresholds. These definitions are essential because they determine which maintenance actions are admissible within an optimization framework and which performance metrics are prioritized. In refinery contexts, maintenance is not an isolated technical activity but a system-level control mechanism that regulates failure probability distributions and operational risk profiles (Jahid, 2021; Ma et al., 2020). Quantitative research in this domain therefore begins with explicit definitions to ensure that optimization outcomes remain aligned with functional performance requirements and safety boundaries that govern refinery operations on an international scale.

Refinery equipment systems consist of interconnected mechanical, thermal, hydraulic, and control components whose failure behaviors are governed by multiple interacting degradation processes. Equipment such as pumps, compressors, heat exchangers, furnaces, pressure vessels, piping networks, and instrumentation systems experience failure mechanisms including corrosion, erosion, fatigue, fouling, vibration-induced wear, thermal cracking, and control drift (Souza et al., 2021; Akbar & Farzana, 2021). These mechanisms operate across different time horizons and respond unevenly to inspection, repair, and replacement actions. Reliability-centered maintenance provides a structured means of classifying these failure modes according to their functional consequences, allowing maintenance tasks to be selected based on risk significance rather than uniform scheduling. Quantitatively, this classification enables the assignment of failure probabilities, detection likelihoods, and consequence magnitudes to each failure mode. Refinery maintenance decisions also interact with operational constraints such as continuous processing requirements, turnaround schedules, and regulatory inspection intervals, which introduce temporal and resource-based dependencies across equipment groups (Kefalidou et al., 2018; Reza et al., 2021). These dependencies transform maintenance planning into a high-dimensional decision problem involving discrete task selection, continuous interval optimization, and constrained scheduling. Empirical observations across refinery operations show that maintenance policies focused solely on corrective actions produce higher variability in availability and greater exposure to cascading failures. Preventive and condition-based strategies reduce uncertainty in failure occurrence but introduce trade-offs involving downtime, inspection cost, and execution risk. Quantitative optimization becomes necessary because these trade-offs cannot be resolved intuitively when asset populations are large and constraints are tightly coupled. Maintenance optimization in refineries therefore requires formal objective functions representing reliability, availability, cost, and risk, along with constraints reflecting safety limits, manpower availability, spare parts logistics, and outage windows. The international relevance of this challenge is reinforced by the standardization of refinery designs and operating practices across regions, which creates comparable reliability problems in different jurisdictions. Quantitative optimization frameworks enable consistent evaluation of maintenance policies across assets and sites, supporting comparability and governance transparency (Saikat, 2021; Velmurugan et al., 2021). Reliability-centered maintenance defines the permissible decision space for such optimization by ensuring that maintenance actions remain functionally justified and risk-informed. This alignment between engineering logic and quantitative modeling is critical for refinery environments where maintenance errors can have high-consequence outcomes.

Figure 1: Reliability-Centered Maintenance Optimization Framework

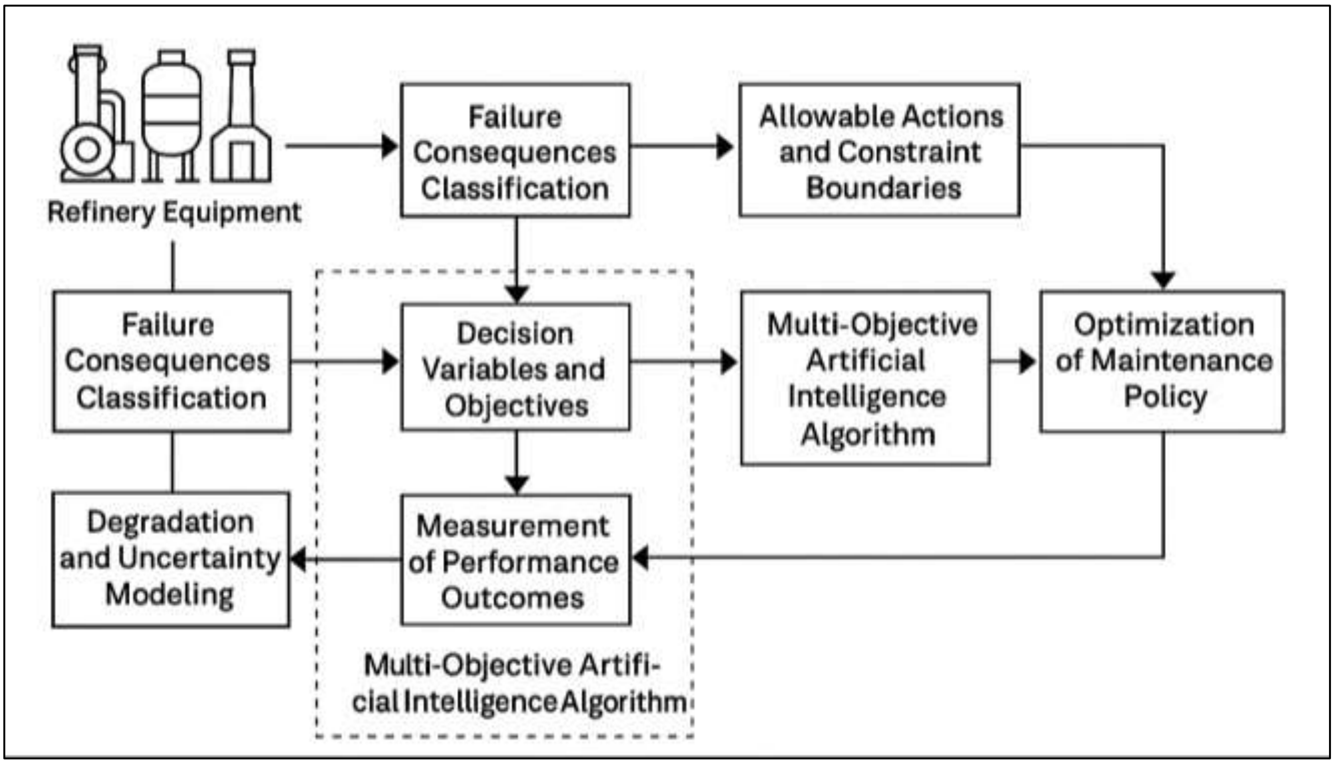

Reliability-centered maintenance decisions in refinery equipment inherently involve multiple objectives that cannot be reduced to a single performance measure without loss of critical information. Maintenance policies simultaneously influence equipment reliability, operational availability, maintenance cost, safety risk, and environmental exposure (Mehairjan, 2017; Shaikh & Aditya, 2021). These objectives are often conflicting, as actions that improve reliability may increase downtime or cost, while cost minimization strategies may elevate failure probability and risk. Quantitative modeling of refinery maintenance therefore requires a multi-objective formulation in which each objective is explicitly represented and evaluated. Reliability objectives are commonly expressed through failure probability reduction, hazard rate control, or availability maximization. Economic objectives include direct maintenance expenditure, indirect downtime losses, energy efficiency penalties, and secondary damage costs. Risk objectives capture the likelihood and severity of loss-of-containment events, fires, explosions, and process upsets. Reliability-centered maintenance supports this multi-objective structure by linking each maintenance task to specific functional failures and consequence categories. This linkage allows objective functions to be derived directly from failure effects rather than abstract performance targets (Bousdekis et al., 2021; Kanti & Shaikat, 2021). Maintenance intervals, inspection coverage, and task selection become decision variables whose effects on each objective can be quantitatively estimated. In refinery contexts, safety-critical functions impose hard constraints rather than soft trade-offs, meaning certain failure probabilities or degradation levels are not acceptable regardless of cost implications. Multi-objective optimization frameworks accommodate this structure by distinguishing between objectives to be optimized and constraints that must be satisfied. This distinction is essential for maintaining alignment with refinery safety governance and regulatory compliance. Internationally, asset management standards emphasize balanced consideration of performance, cost, and risk, reinforcing the relevance of multi-objective approaches. Quantitative studies in refinery maintenance therefore adopt optimization formulations that generate sets of non-dominated solutions representing alternative maintenance policies, each reflecting a different balance among objectives (Zeinalnezhad et al., 2020; Zobayer, 2021a, 2021b). These solution sets enable structured evaluation and selection by decision-makers without collapsing complex trade-offs into a single scalar metric. Reliability-centered maintenance ensures that all candidate solutions remain grounded in functional necessity and consequence awareness, preserving engineering validity within the optimization process.

Artificial intelligence algorithms are defined in this context as computational search methods capable of exploring large, complex decision spaces through adaptive and iterative improvement mechanisms. In refinery maintenance optimization, the decision space is characterized by nonlinear relationships, mixed discrete and continuous variables, and tightly coupled constraints (Hall et al., 2021; Arman & Kamrul, 2022; Nahid, 2022). Traditional analytical optimization methods are limited in such environments due to nonconvexity, discontinuity, and combinatorial explosion. Multi-objective artificial intelligence algorithms address these challenges by maintaining populations of candidate solutions and evolving them toward improved trade-offs across objectives. Evolutionary algorithms, swarm-based methods, and metaheuristic search techniques are particularly suited to maintenance optimization because they do not require gradient information and can accommodate arbitrary objective functions (Hossain & Milon, 2022; Abdur &Haider, 2022). These algorithms generate diverse sets of feasible maintenance policies, allowing exploration of alternative task schedules, inspection intervals, and resource allocations. In multi-objective settings, algorithm performance is evaluated based on convergence toward optimal trade-offs and diversity across the solution set. For refinery applications, this capability is essential because decision-makers require visibility into multiple policy options that satisfy safety and operational constraints (Bellinello et al., 2020; Mushfequr & Praveen, 2022; Mortuza & Rauf, 2022). Artificial intelligence algorithms function as optimization engines rather than decision authorities, operating within the decision space defined by reliability-centered maintenance logic. The maintenance model specifies allowable actions, failure relationships, and constraint boundaries, while the algorithm searches for combinations that optimize performance measures. This separation preserves explain ability and accountability, which are critical in refinery environments where maintenance decisions are subject to audit and regulatory scrutiny. International application of such algorithms is facilitated by their adaptability to different refinery configurations and operating conditions, as long as the underlying reliability and cost models are properly parameterized (Netto et al., 2020; Rakibul & Samia, 2022; Rony & Ashraful, 2022). Quantitative research leverages artificial intelligence to handle computational complexity while maintaining fidelity to engineering principles embedded in reliability-centered maintenance.

Quantitative optimization of reliability-centered maintenance in refinery equipment requires explicit modeling of degradation processes, uncertainty sources, and maintenance effectiveness. Degradation mechanisms in refineries often interact, producing nonstationary failure behavior that challenges simple lifetime assumptions (Abdul, 2023; Saikat, 2022). Quantitative models must therefore represent how operating conditions, load variations, and environmental factors influence degradation rates (Abdulla & Zaman, 2023; Arfan et al., 2023; Teixeira et al., 2020). Failure data in refinery contexts are frequently sparse, censored, or aggregated, introducing statistical uncertainty into reliability estimates. Maintenance actions themselves introduce variability, as repairs and inspections may partially restore equipment condition rather than returning it to an idealized state. This imperfect maintenance behavior affects subsequent failure probability and must be reflected in reliability models. Uncertainty also arises from measurement error in condition monitoring data, inspection findings, and reporting practices. Quantitative frameworks address these uncertainties through probabilistic modeling, stochastic simulation, or Bayesian updating mechanisms. Reliability-centered maintenance provides a structured means of integrating expert judgment and historical evidence by explicitly linking degradation modes to observable indicators and maintenance responses (Amin & Md Mesbaul, 2023; Foysal & Aditya, 2023; Nunez & Borsato, 2017). When these models are embedded within multi-objective optimization, uncertainty propagates into objective evaluations, influencing the relative ranking of candidate maintenance policies. Artificial intelligence algorithms can accommodate this variability by evaluating expected performance across simulated scenarios rather than relying on deterministic estimates (Hamidur, 2023; Rashid et al., 2023). In refinery environments, uncertainty modeling is particularly important for low-frequency, high-consequence failures where empirical data are limited but consequences are severe. Maintenance optimization models therefore incorporate conservative constraints for safety-critical equipment while allowing more flexible trade-offs for lower-consequence assets (Galar et al., 2021; Musfiqur & Kamrul, 2023; Muzahidul & Mohaiminul, 2023). This differentiated treatment aligns with barrier management principles that prioritize containment integrity and hazard control. Quantitative treatment of degradation and uncertainty enhances the robustness of optimized maintenance policies and supports defensible decision-making across refinery operations.

Figure 2: Reliability-centred Multi-Objective Maintenance Framework

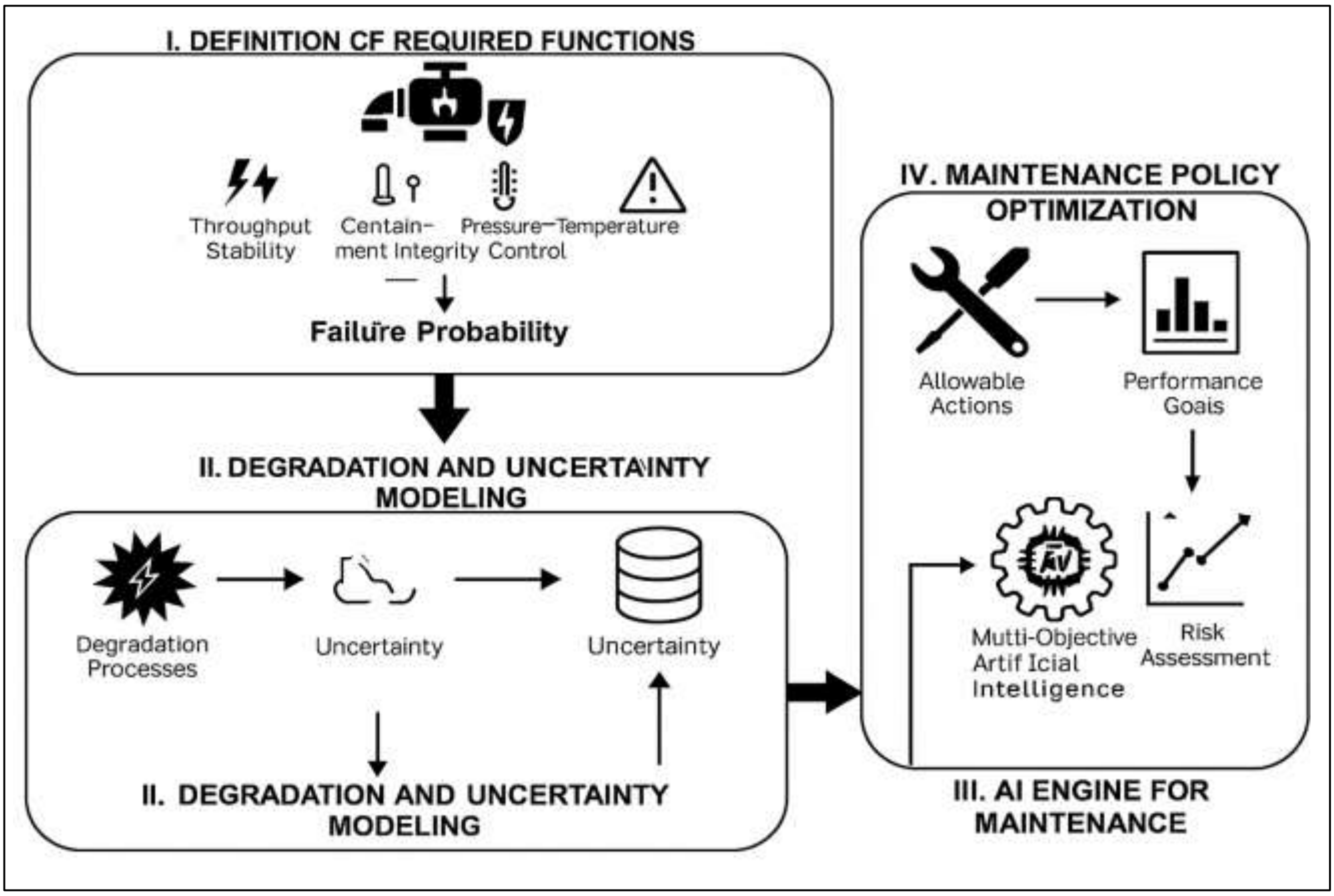

A quantitative study of reliability-centered maintenance optimization requires a clearly defined measurement architecture that links decision variables to performance outcomes. Reliability measures include failure rates, hazard functions, and survival probabilities, while maintainability is represented through repair time distributions and restoration effectiveness (Amin & Sai Praveen, 2023; Hasan & Ashraful, 2023; Moraes et al., 2021). Availability combines reliability and maintainability into a single operational metric reflecting the proportion of time equipment is capable of performing its function. Economic measures encompass preventive and corrective maintenance costs, inspection expenditures, spares inventory costs, and downtime losses. Risk measures quantify the likelihood and severity of adverse events, often incorporating consequence modeling for safety and environmental impacts. Decision variables include maintenance task type, inspection method, intervention interval, and resource allocation. Constraints represent safety requirements, inspection mandates, manpower limits, and outage scheduling restrictions. Reliability-centered maintenance ensures that all decision variables correspond to valid functional preservation actions rather than arbitrary interventions (Galar et al., 2017; Ibne &Kamrul, 2023; Mushfequr & Ashraful, 2023). In multi-objective optimization, these variables are encoded in a form suitable for algorithmic search, while constraints are enforced through feasibility rules or penalty mechanisms (Roy & Kamrul, 2023; Saba et al., 2023). Quantitative evaluation of candidate solutions relies on consistent time horizons and comparable metrics to ensure meaningful trade-off analysis. Refinery maintenance optimization models often incorporate execution-related parameters such as task completion probability and scheduling delays, recognizing that planned actions do not always translate directly into realized outcomes. This comprehensive measurement structure supports rigorous quantitative analysis and enables reproducible comparison of alternative maintenance policies across equipment populations and operating contexts (Nowakowski et al., 2019; Saba & Kanti, 2023; Shaikh & Farabe, 2023).

Reliability-centered maintenance optimization in refinery equipment carries international significance due to the global integration of energy production, refining capacity, and distribution networks. Refinery reliability affects fuel availability, industrial feedstock supply, and economic stability across regions. Maintenance-related failures can propagate through supply chains, influencing markets beyond national boundaries. Safety incidents in refineries also have transnational implications through regulatory response, industry learning, and public trust in energy infrastructure (Marquez et al., 2020; Haider & Hozyfa, 2023; Zobayer, 2023). The convergence of asset management practices across countries has created shared expectations for risk-informed and performance-based maintenance decision-making. Quantitative optimization frameworks support this convergence by providing standardized methods for evaluating maintenance policies using common performance metrics. Multi-objective artificial intelligence algorithms enhance this capability by generating transparent sets of alternative solutions that reflect different balances among reliability, cost, and risk. Reliability-centered maintenance anchors these solutions in functional necessity and consequence awareness, ensuring that optimization outputs remain aligned with engineering intent and safety governance (Vogl et al., 2019). The refinery sector provides a rigorous application domain for such methods because of its asset complexity, high consequence exposure, and data-rich operating environment. Quantitative maintenance optimization in this context contributes to consistent decision logic across refineries while allowing adaptation to local conditions through parameterization. The integration of reliability-centered maintenance with multi-objective artificial intelligence algorithms therefore represents a structured analytical approach for managing refinery equipment performance within internationally relevant operational and safety constraints (Bakri & Januddi, 2020).

The objective of the quantitative study titled "Reliability-Centered Maintenance Optimization Using Multi-Objective AI Algorithms in Refinery Equipment" is to develop and test a structured optimization framework that translates reliability-centered maintenance decision logic into measurable variables and then identifies maintenance policies that are simultaneously efficient, feasible, and performance-balanced across refinery asset systems. The study aims to formalize refinery equipment maintenance planning as a multi-objective optimization problem in which key outcomes—equipment reliability, operational availability, maintenance cost burden, and quantified risk exposure—are represented as explicit objective functions rather than implicit managerial preferences. A central objective is to construct a refinery-relevant decision space by defining admissible maintenance actions using reliability-centered maintenance principles, including task selection categories (condition-based, time-based, failure-finding, corrective, and redesign-equivalent substitutes) and task timing variables such as inspection intervals, preventive replacement frequencies, and condition thresholds. Another objective is to encode the principal operational constraints that determine feasibility in real refinery contexts, including turnaround windows, manpower and contractor capacity, spare parts availability, inspection resource limits, and safety-critical barrier requirements that must be satisfied as hard constraints. The study also seeks to implement multi-objective AI algorithms capable of searching this constrained space efficiently and producing a diverse set of non-dominated maintenance plans, allowing comparative evaluation of trade-offs rather than forcing a single maintenance policy. A further objective is to evaluate the resulting solution sets using quantitative indicators of both optimization quality and operational interpretability, including solution feasibility rate, Pareto diversity, convergence stability, and scenario robustness under uncertainty in failure behavior and maintenance effectiveness. The study additionally aims to compare optimized policies against benchmark maintenance strategies commonly used in refinery practice, such as fixed-interval preventive maintenance, reactive dominant maintenance, and condition-triggered maintenance with static thresholds, using consistent metrics across equivalent time horizons. Finally, the study's objective includes establishing a reproducible methodological pipeline that integrates failure mode structures, reliability modeling outputs, and AI optimization results into a coherent quantitative analysis, enabling maintenance decision-makers to select policies that align with functional requirements and risk acceptability while maintaining transparent trade-off visibility across competing refinery performance priorities.

LITERATURE REVIEW

The literature review for the quantitative study titled "Reliability-Centered Maintenance Optimization Using Multi-Objective AI Algorithms in Refinery Equipment" synthesizes research streams that jointly explain how refinery maintenance decisions can be formulated as a measurable, constrained, multi-objective optimization problem and solved using artificial intelligence algorithms while remaining consistent with reliability-centered maintenance logic (Sifonte & Reyes-Picknell, 2017). This section surveys foundational theories and empirical findings in reliability engineering, reliability-centered maintenance, maintenance optimization, and refinery asset integrity management to establish how failure behavior, degradation mechanisms, inspection effectiveness, and maintenance task selection are represented quantitatively. It then reviews multi-objective optimization theory and AI-based solution approaches that are commonly used to search complex decision spaces involving mixed discrete–continuous variables and nonlinear constraints, which match the structure of refinery maintenance planning. In addition, the literature review examines how refinery-specific operational realities—turnaround scheduling, resource limits, safety-critical barrier constraints, and consequence severity—shape objective function design and feasibility modeling (Karajagikar & Sonawane, 2020). The section also discusses how uncertainty is treated in maintenance optimization research, including statistical life modeling, imperfect maintenance representation, and scenario-based evaluation, because refinery data often include censoring, sparse failure events for critical assets, and measurement variability from inspection and condition monitoring. Finally, the literature review consolidates evidence on evaluation practices for multi-objective AI outputs, including Pareto-front quality metrics and operational interpretability criteria, to support rigorous comparison between optimized maintenance policies and benchmark strategies (Yang et al., 2020). Overall, the literature review builds the conceptual and quantitative basis for the study's framework by connecting refinery maintenance decision logic, reliability modeling, and multi-objective AI optimization into a unified research foundation.

Conceptual and Quantitative Foundations of Reliability-Centered Maintenance (RCM)

Reliability-centered maintenance (RCM) is commonly framed as a decision-oriented maintenance methodology that preserves required system functions by selecting tasks that are justified by failure consequences and supported by evidence from operating context. In quantitative decision terms, the "required function" is not a general statement of purpose; it is a set of measurable outputs that can be observed and audited in a refinery environment (Alrifaey et al., 2020). These outputs typically include availability at the equipment and unit level, throughput stability expressed through sustained production rates and reduced process variability, and containment integrity represented through sustained pressure boundary performance and controlled emissions. Within this framing, functional failure is defined as a measurable deviation from allowable operating limits rather than a vague notion of "breakdown." A pump that still rotates but cannot maintain discharge pressure, a heat exchanger that maintains flow but loses thermal duty due to fouling, or a valve that cycles but exhibits stiction leading to unstable control can each constitute functional failure because performance moves outside defined limits. RCM translates these concepts into an explicit taxonomy of failure modes that can be mapped to measurable event types and measurable indicators (Tang et al., 2017). For rotating equipment, event types such as seal leakage, bearing distress, and misalignment-related vibration escalation are often associated with observable signals and maintenance logs. For static equipment, event types such as tube rupture, corrosion-driven wall loss exceedance, and fouling-induced thermal performance degradation are captured through inspection findings, process data patterns, and integrity reports. The quantitative value of defining failure modes this way is that it creates traceable relationships between operational performance requirements and the maintenance actions that can be justified to preserve them. RCM also clarifies that "reliability" and "availability" are not interchangeable; reliability is tied to failure occurrence behavior over time, while availability is a combined reflection of failure occurrence and restoration effectiveness. This distinction matters because refinery decisions regularly require separating the frequency of functional failure from the duration of downtime and the quality of restoration (Rosita & Rada, 2021). Studies across maintenance engineering, reliability analytics, and process industries repeatedly emphasize that consistent definitions of required function, functional failure, and failure mode taxonomy are the foundation for defensible quantitative modeling, since objectives and constraints cannot be meaningfully specified when failure is ambiguously defined. In refinery equipment, where high-energy processes and integrity requirements dominate the operating context, this definitional discipline becomes the base layer for every subsequent optimization step.

Figure 3: Reliability-centred Maintenance Decision Framework

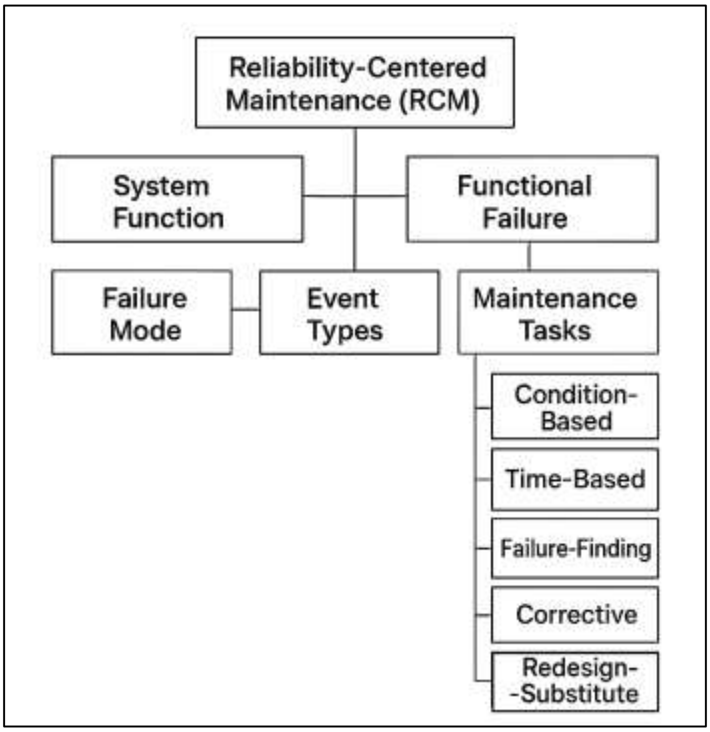

RCM becomes operationally useful when its decision structure is translated into variables that can be manipulated, evaluated, and compared in a quantitative framework. The classic RCM logic organizes maintenance actions into categories that correspond to different control philosophies: condition-based actions triggered by observed degradation, time-based actions scheduled by interval, failure-finding actions designed to reveal hidden failures, corrective actions that restore function after failure, and redesign-substitute actions that modify conditions or configurations to reduce failure likelihood or consequence (Yavuz et al., 2019). In quantitative maintenance planning, these categories are not merely labels; they define admissible policy types for each asset-mode pair. Decision variables then specify which task category is assigned to each failure mode, the interval length for time-directed tasks, the trigger threshold for condition-based tasks, and the inspection coverage level that defines how much of the relevant failure space is being monitored. Coverage is especially important in refinery settings because inspection and monitoring rarely observe the entire failure mechanism directly; instead, coverage represents the practical reach of methods such as vibration monitoring points, thickness measurement locations, or sampling frequency for oil analysis. RCM-based variables also include sequencing and packaging variables when tasks are constrained by shutdown windows or turnaround schedules. A key element of translating RCM into quantitative form is consequence classification

(Afzali et al., 2019). Refinery equipment failures differ sharply in consequence severity, so RCM distinguishes safety-critical functions from production-critical functions and assigns different decision rules to each. In quantitative models, this distinction is represented either through weights that penalize safety-impacting failure outcomes more heavily, or through constraint classes that prohibit trading safety performance for cost savings. This structure supports consistency with barrier-management logic, where certain functions must be preserved within tight tolerance regardless of economic pressure. The decision structure also supports mixed-variable optimization because some choices are discrete (task type selection, inspection method selection) and some are continuous (interval lengths, thresholds), mirroring the real maintenance planning environment. A recurring finding across diverse maintenance research is that refinery maintenance problems become computationally complex not because the concepts are unclear, but because the decision space is large and constrained: there are many equipment items, many failure modes, and many feasible combinations of tasks and intervals, each with measurable effects on downtime, risk exposure, and cost (Prasetyo & Rosita, 2020). Translating RCM into explicit variables is therefore not a cosmetic exercise; it is the step that converts engineering logic into a mathematically searchable decision space while keeping the search anchored to functional necessity and consequence awareness.

RCM output artifacts provide the structured information needed to populate optimization models with traceable inputs, and the literature consistently treats these artifacts as the bridge between qualitative reasoning and quantitative computation. The most commonly used artifacts include FMEA and FMECA outputs that list failure modes, their causes, detection opportunities, local and system effects, and severity characterizations (Ross et al., 2018). These artifacts support quantitative modeling because they create a standardized inventory of what can go wrong and how it matters, which then supports mapping each failure mode to measurable event definitions and measurable indicators. Detection methods and detectability assessments inform inspection and monitoring variables, while effects and severity assessments inform objective priorities and constraint strictness. When criticality indices are produced—whether through scoring systems, ranking matrices, or more elaborate risk indices—these values serve as prioritization rules that shape how optimization allocates limited resources such as inspection capacity, skilled labor hours, and spare parts. In quantitative optimization models, criticality can appear as constraints that require minimum attention levels for high-criticality assets, as penalty terms that discourage resource diversion away from barrier functions, or as selection rules that restrict the action set to those considered technically justified for each mode. Another major role of RCM artifacts is defining admissible actions. Not every maintenance action is appropriate for every failure mode; condition monitoring is meaningful for some progressive degradation patterns, time-based replacement may be meaningful for wear-driven modes, and failure-finding is relevant for hidden protective functions (Okwuobi et al., 2018). The RCM logic encoded in FMEA/FMECA outputs therefore defines the allowable pairing between failure modes and maintenance task categories. This pairing is essential when AI algorithms are used for search because unconstrained search can generate solutions that appear numerically attractive while violating basic engineering validity. By embedding RCM artifacts into the optimization, the action set becomes a filtered set of technically admissible choices rather than an unrestricted menu. The literature also emphasizes the importance of maintaining traceability from optimized decisions back to RCM artifacts for auditability and organizational learning, especially in high-consequence sectors such as refining. Traceability supports validation of results by reliability engineers and integrity specialists, and it supports governance requirements that demand evidence-based maintenance justification. Across multiple streams of research on maintenance strategy design and industrial reliability management, studies repeatedly show that optimization outcomes become more implementable when they directly reference structured artifacts already used in maintenance programs, because implementation teams can connect outputs to familiar failure modes, inspection plans, and work management processes (Alizadeh et al., 2021). This role of RCM artifacts positions them as both technical inputs and organizational alignment mechanisms in quantitative maintenance optimization.

A final strand in the literature on conceptual and quantitative foundations of RCM concerns how the structured logic of RCM constrains and enables AI-based search for maintenance optimization without sacrificing interpretability (Gupta & Mishra, 2018). Multi-objective AI algorithms are powerful because they can explore large decision spaces, handle mixed-variable problems, and generate sets of trade-off solutions across competing objectives such as reliability performance, downtime burden, cost, and risk exposure. However, refinery maintenance research consistently indicates that algorithmic power alone does not produce valid maintenance plans unless the search space is engineered to reflect permissible actions, consequence boundaries, and operational constraints. RCM provides this engineering of the search space through its outputs and decision logic. When RCM defines required functions and functional failures as measurable outcomes, AI evaluation can score candidate plans using consistent performance measures grounded in real operational limits. When RCM classifies failure modes into event types such as seal leakage, bearing degradation, fouling-related duty loss, and tube rupture, AI evaluation can connect maintenance decisions to measurable event likelihoods and measurable consequence categories recorded in plant data systems (Fuentes-Huerta et al., 2021). When RCM decision variables are formalized—task categories, intervals, thresholds, and coverage—AI search can manipulate those variables in controlled ways that correspond to actual maintenance levers. When consequence classification is embedded as safety-critical constraint classes, AI-generated solutions must satisfy barrier-preservation requirements as feasibility conditions, which aligns algorithmic outputs with refinery safety governance. The literature also supports the use of RCM artifacts as a means to constrain AI search to admissible action sets, reducing the risk of producing plans that are numerically optimized but operationally illegitimate (Fang et al., 2019). Importantly, interpretability is addressed not by simplifying the optimization problem, but by preserving the mapping between optimized choices and RCM constructs those engineers already use: each decision can be traced to a failure mode, a consequence rationale, and a task justification. Studies across maintenance analytics emphasize that such traceability increases the likelihood of implementation, improves stakeholder acceptance, and supports continuous improvement because realized outcomes can be compared back to the assumptions encoded in the RCM artifacts. In refinery environments, where maintenance decisions must remain auditable and aligned with integrity programs, the synthesis of RCM logic and AI search is presented in the literature as a disciplined integration: RCM structures what is allowed and what matters, and AI explores how to allocate tasks and intervals within that structured space to balance multiple quantitative objectives under real constraints (Catelani et al., 2020).

Refinery Equipment Context

Refinery equipment reliability literature consistently treats the refinery as a heterogeneous asset ecosystem in which degradation signatures vary by equipment class, duty severity, and process environment, requiring differentiated maintenance analytics rather than uniform scheduling (Khodabakhsh et al., 2018). Rotating equipment—particularly pumps, compressors, turbines, fans, and associated drivetrains—exhibits degradation patterns that are strongly expressed through dynamic and lubrication-related indicators. Common signatures include changes in vibration amplitude and frequency content associated with imbalance, misalignment, looseness, bearing defects, and resonance; these signals often co-occur with temperature rise, abnormal acoustic patterns, and changes in power draw. Lubrication degradation is repeatedly characterized as both a root contributor and a diagnostic window, with contamination, viscosity shift, oxidation, and wear debris serving as measurable proxies for frictional and fatigue mechanisms in bearings and gears. Cavitation in pumps is frequently described as a combined hydraulic and mechanical stressor that can produce noise, vibration, performance loss, and localized erosion, while seal wear and seal system instability can manifest as leakage events, pressure deviations, or repeated corrective work orders. Refinery literature also separates rotating equipment degradation into progressive modes, which support condition-based monitoring, and rapid-onset modes, which require protective logic and robust operational controls (Xie et al., 2019). Static equipment—heat exchangers, pressure vessels, piping circuits, fired heater tubes, and storage systems—shows a different family of signatures dominated by material loss, cracking, and performance drift. Corrosion is treated as a principal mechanism with multiple subtypes tied to chemistry, temperature, velocity, and metallurgy, often captured through thickness measurements, corrosion rate estimates, and localized inspection findings. Erosion and erosion–corrosion is discussed as flow-accelerated mechanisms that concentrate damage at elbows, tees, impingement zones, and high-velocity regions. Fouling in heat exchangers is presented as a degradation pathway with direct thermodynamic consequences, measurable through loss of heat duty, increased pressure drops, changes in approach temperatures, and increased energy consumption, making process historians valuable sources of early warning data. Cracking mechanisms are linked to cyclic stress, thermal gradients, and corrosive environments, with inspection modalities providing intermittent observations that must be interpreted in the context of operating history (Ishola et al., 2020). Instrumentation and control loops form a third class of refinery assets in which degradation often appears as drift, stiction, hysteresis, intermittent faults, and control instability, measurable through calibration records, process variability, oscillation metrics, and valve signature behavior. Across these asset classes, the literature emphasizes that degradation signatures are not isolated indicators but interacting patterns embedded in process context, and refinery maintenance research repeatedly uses this classification to justify asset-specific monitoring, inspection selection, and task interval differentiation (Flath & Stein, 2018).

Quantitative refinery maintenance research commonly represents failure behavior using reliability models that translate observed events and condition evidence into probabilistic descriptions of time-dependent failure propensity, enabling maintenance decisions to be evaluated using comparable metrics across asset populations (Yuan et al., 2017). Time-to-failure representations are used to summarize the elapsed operating time between functional failures, while hazard-oriented reasoning is used to express how failure propensity changes with age, stress exposure, or accumulated degradation. Survival-curve thinking is widely used to compare alternative maintenance policies by examining how the probability of continued function changes over operating time under different intervention strategies. Within refinery datasets, the literature repeatedly highlights that failure observations often include right-censoring, where assets have not failed by the end of the observation window; left-censoring, where degradation began before tracking started; and interval censoring, where inspection finds damage within a range of possible onset times rather than at a precise moment. Missingness is also common because work order records may omit root cause detail, sensor streams may contain gaps, and integrity inspection results may be sparse in time relative to degradation dynamics (Teixeira et al., 2020). High-integrity systems create an additional challenge: some high-consequence failure events are rare, which produces sparse data that complicates direct frequency estimation and increases uncertainty around model parameters. Refinery literature addresses these realities by emphasizing careful event definition, consistent failure coding, and the separation of failure events from maintenance-induced outages, because downtime does not always correspond to functional failure and repair actions can be preventive, corrective, or opportunistic. Condition indicators are treated as crucial complements to event data, particularly for rotating equipment and exchanger performance, where continuous or high-frequency measurements can reveal degradation progression before failure occurs. These indicators include vibration features, lubrication and wear-debris markers, temperature and pressure trends, process variability signatures, and heat-transfer efficiency proxies (Gertler, 2017). Remaining useful life representations are discussed as operationally meaningful summaries that convert condition evidence into an estimate of how much functional time remains before a failure threshold is crossed, often expressed with uncertainty bounds rather than single-point certainty. Research also acknowledges that remaining useful life estimates depend on the stability of operating context, the representativeness of training data, and the definition of failure thresholds that reflect refinery functional requirements. The literature frequently stresses that quantitative reliability representations become more actionable when they are aligned with specific functional failure definitions, because the same physical degradation can be tolerable for one service and unacceptable for another depending on safety, environmental, and production constraints (Jiang et al., 2018). As a result, refinery maintenance research tends to integrate reliability models with equipment criticality and consequence categories, ensuring that probabilistic outputs are interpreted through refinery risk priorities rather than used as isolated statistical artifacts.

Figure 4: Refinery Reliability Analytics Framework Diagram

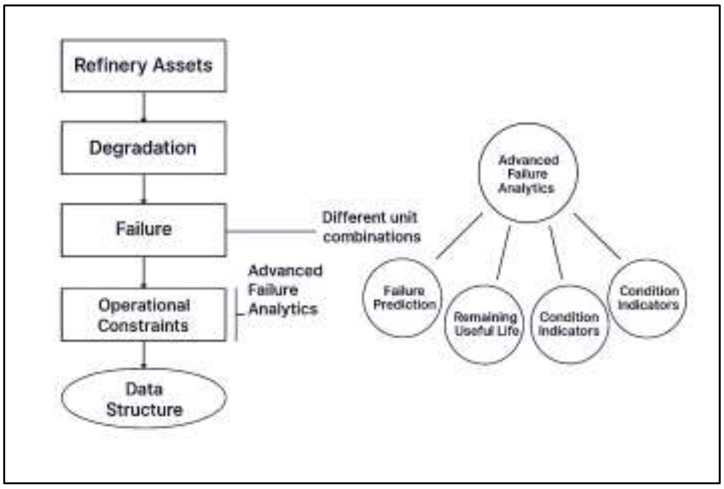

Refinery maintenance optimization studies regularly treat operational constraints as first-class modeling elements because they shape which maintenance decisions are feasible, when they can be executed, and how they interact across equipment groups. Turnaround windows are typically represented as discrete feasibility periods during which intrusive inspections, major overhauls, and internal vessel work can be performed, while on-stream windows limit interventions to those compatible with continuous operation and safe isolation practices (Milazzo & Bragatto, 2019). This structure creates timing dependencies that link maintenance decision variables to unit shutdown schedules, scaffolding availability, and work-package sequencing. Resource capacity constraints are repeatedly emphasized as binding limitations: manpower availability, contractor hours, specialty craft access, and nondestructive testing capacity often determine whether an optimized plan can be executed within the available calendar time. Nondestructive testing availability is treated as particularly constraining for integrity-focused work because inspection queues, certification requirements, and equipment access planning can bottleneck execution. Spare parts lead time is also modeled as a coupling constraint across assets because the procurement timeline for seals, bearings, exchanger bundles, valves, and specialized rotating components can exceed the time between detection of degradation and required intervention. This produces coordination requirements between reliability predictions, materials management, and maintenance scheduling, and it can cause one asset's demand to crowd out another if inventory is limited (Niu, 2017). Permit-to-work systems introduce additional feasibility and risk constraints by controlling when high-risk activities can occur and by requiring isolation, gas testing, confined space controls, and simultaneous operations coordination; these procedural requirements translate into measurable time costs, staffing requirements, and exposure hours. Safety exposure hours are treated in the literature as quantifiable proxies for maintenance-induced risk, since intrusive work increases human exposure to hazardous environments even when it reduces in-service failure risk. Consequently, refinery maintenance modeling often recognizes that maintenance is not purely risk-reducing; it redistributes risk across operating and maintenance phases, requiring careful balancing between intervention frequency and exposure burden. Constraints are also discussed as interacting rather than independent: a turnaround window increases access but compresses work into limited time, which intensifies manpower constraints; long-lead spares constraints can force schedule shifts that alter permit loads and simultaneous operations risks; inspection capacity limits can reduce coverage and increase uncertainty in degradation assessment (Iqbal et al., 2017). Quantitative studies frequently incorporate these constraints to avoid producing maintenance policies that look optimal in reliability or cost terms but fail under practical execution limits. The literature presents constraint modeling as essential for refinery realism because it preserves the operational context in which reliability-centered decisions are implemented and evaluated.

Across refinery maintenance research, the structure and quality of data are treated as decisive factors that influence the validity of reliability modeling and optimization outputs, and the literature repeatedly emphasizes data integration across maintenance, operations, and integrity systems. Work order management systems provide event histories, repair actions, labor hours, and parts usage, while process historians supply continuous streams of operating conditions and performance indicators that often contain early signatures of degradation (P. Wang et al., 2021). Inspection databases contribute intermittent integrity observations—thickness readings, defect characterizations, crack indications, and inspection coverage notes—while condition monitoring platforms provide high-frequency vibration, lubrication, and performance features for rotating equipment (Bakhtiari et al., 2017). These datasets differ in granularity, temporal resolution, and semantic consistency, which introduces challenges in aligning timestamps, matching equipment identifiers, and ensuring that a "failure event" means the same thing across sources. Missingness and inconsistent coding are frequently identified as sources of bias, particularly when failure cause fields are incomplete or when preventive interventions are recorded in ways that resemble corrective repairs. Sparse failure events for high-integrity assets also lead researchers to rely more heavily on condition indicators, inspection findings, and expert-informed categorizations to characterize risk and degradation progression. Data structure issues become especially important for instrumentation and control loops because faults may appear as subtle control instability, oscillation, or drift rather than clear breakdown, which requires extracting performance signals from process data and associating them with maintenance records. The literature also notes that refinery data often reflect the maintenance policy already in place, meaning observed failure patterns are shaped by prior interventions; this dependence influences how time-to-failure is interpreted and how model calibration is performed (Yao et al., 2021). For optimization purposes, researchers frequently emphasize the need for consistent definitions of failure thresholds and functional performance limits so that reliability estimates and remaining useful life indicators correspond to the same operational meaning used by maintenance planners. When operational constraints are included—turnarounds, staffing limits, inspection capacity, spares lead time, and permit-to-work loads—the data required expand to include scheduling calendars, resource rosters, procurement timelines, and safety systems information, creating a multi-source modeling environment. Studies in this area repeatedly show that optimization frameworks gain credibility when they explicitly document how data sources are reconciled, how censoring and missingness are handled, and how condition indicators are linked to functional failure definitions. This integrated data structure perspective is central in refinery maintenance literature because it supports traceability: decision variables and constraints can be mapped to real records, predicted outcomes can be compared to observed performance, and maintenance plans can be evaluated using metrics that stakeholders recognize from existing reliability and integrity governance systems (Kim et al., 2020).

Maintenance Optimization Models

Maintenance optimization literature in industrial engineering commonly begins with single-objective formulations because they provide clear measurable outputs and tractable benchmarking, especially when maintenance programs are evaluated against one dominant performance target (Bányai, 2021). Studies in this stream frequently represent optimization as minimizing a maintenance-related cost rate, minimizing downtime burden, or maximizing reliability performance, with outputs expressed in directly observable measures such as total maintenance expenditure over a fixed horizon, cumulative downtime hours, average production loss, failure counts, or mean operating time between functional failures. Cost-focused models typically aggregate direct maintenance costs, including labor and parts, and may incorporate indirect losses associated with stoppages, quality deviations, or energy inefficiency caused by degraded equipment. Downtime-focused models treat lost availability as the key penalty, often evaluating schedules based on how frequently assets are removed from service and how long restoration takes. Reliability-maximization models emphasize reducing failure occurrence through intervention timing, inspection scheduling, or proactive replacement actions, and they evaluate outcomes using counts of functional failures, reliability levels over time, or stability of performance within operating limits. In refinery contexts, single-objective models are frequently used to isolate the impact of a specific policy lever—such as preventive replacement interval or inspection frequency—under controlled assumptions (Pisacane et al., 2021). The literature also identifies persistent limitations of single-objective framing when applied to refineries, where maintenance decisions simultaneously affect safety risk, environmental exposure, production continuity, and cost. Refinery operations introduce competing priorities that single-objective optimization cannot represent without embedding hidden trade-offs into assumptions or arbitrarily converting dissimilar outcomes into a single score. For example, a policy that lowers direct cost can increase unplanned downtime by shifting work from preventive to corrective actions; a policy that lowers downtime through fewer interventions can increase failure probability by extending operating periods beyond tolerable degradation thresholds. Another recurring limitation is that safety-critical constraints in refineries cannot be appropriately treated as simple cost terms, because compliance and barrier integrity requirements impose non-negotiable thresholds on inspection and maintenance actions. As a result, refinery-focused maintenance studies often use single-objective models as baseline references or preliminary screens, while acknowledging that single-objective outputs provide an incomplete representation of decision quality when objectives conflict. This literature emphasizes that measurable outputs remain valuable, but they must be situated in a framework that preserves the multi-dimensional nature of refinery maintenance performance (Faddoul et al., 2018). That shift forms the analytical bridge from single-objective optimization to multi-objective maintenance optimization, where trade-offs are explicitly represented and evaluated rather than indirectly embedded in a single scalar objective.

Multi-objective maintenance optimization literature expands the modeling structure to represent refinery and industrial decisions as simultaneous attempts to improve several performance dimensions under realistic constraints. In this body of work, objective sets commonly include minimizing total expected cost, minimizing expected unavailability or maximizing availability, minimizing quantified risk exposure, and minimizing the number of intrusive interventions as a proxy for workforce exposure, disruption, and maintenance-induced risk (Altman, 2021). Total cost objectives are typically defined to include both direct maintenance expenditures and indirect costs associated with downtime losses, production instability, and secondary damage from cascading failures.

Figure 5: Maintenance Optimization Policy Decision Framework

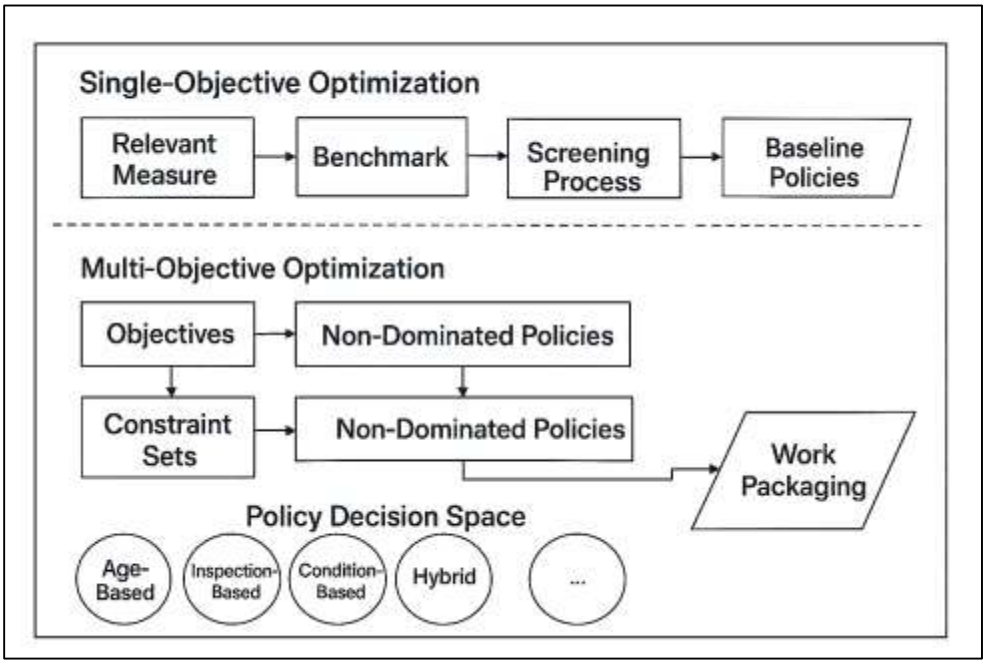

Availability objectives capture the combined effect of failure occurrence and restoration duration, making them central for continuous-process industries where unit interruptions can propagate across interconnected systems. Risk objectives are represented through probability–consequence reasoning, where probability reflects failure propensity and consequence reflects safety, environmental, and operational severity; in refinery studies, high-consequence failure modes are often handled as constraint-dominant elements rather than being treated as a simple weighted preference. The intervention-count objective appears in the literature as a practical way to capture the operational and safety burden of maintenance, reflecting that each intrusive action creates execution risk, isolation complexity, permit workload, and staffing exposure. Constraint sets are treated as essential in multi-objective formulations because feasible maintenance decisions in refineries are bounded by safety minima, inspection frequency requirements, access constraints, and turnaround feasibility (Nemati et al., 2018). Safety minima constrain how far inspection intervals can be extended for integrity-critical components and require verification tasks for protective functions. Inspection frequency requirements reflect regulatory, standard-based, or programmatic rules that must be satisfied regardless of economic pressures. Turnaround feasibility constraints limit when intrusive work can be performed and introduce coupling across assets, because many tasks compete for the same outage window, scaffolding access, and specialty testing capacity. The multi-objective literature commonly frames the outcome as a set of non-dominated solutions rather than a single "best" policy, supporting structured comparison among trade-offs. This approach aligns with refinery governance environments in which stakeholders from operations, maintenance, integrity, and safety evaluate maintenance programs using different success criteria. Studies repeatedly stress that a major advantage of multi-objective modeling is transparency: trade-offs between cost, availability, and risk are visible rather than hidden (Bousdekis et al., 2021). Within refinery maintenance research, multi-objective frameworks are also valued because they allow safety-critical constraints to be enforced while still optimizing production and cost outcomes for less critical assets, producing decision sets that reflect the differentiated consequence structure of refinery equipment systems.

The literature on maintenance policy decision spaces emphasizes that optimization is ultimately an exercise in selecting and tuning policy types that correspond to implementable maintenance levers, and it repeatedly categorizes these levers into age-based, condition-based, inspection-based, and hybrid portfolios. Age-based replacement intervals are among the most frequently optimized variables because they correspond to simple preventive replacement policies, are relatively easy to implement, and can be evaluated using historical failure and repair records (Deng et al., 2017). Studies optimizing age-based intervals examine how different interval lengths change failure occurrence patterns, downtime frequency, and total cost outcomes across equipment populations. Condition threshold optimization is emphasized in research that leverages condition monitoring and process performance signals, particularly for rotating equipment vibration indicators, lubrication degradation markers, corrosion rates derived from thickness readings, and fouling indicators derived from thermal performance degradation. Thresholds are treated as decision points that trigger maintenance actions when measured indicators cross predefined limits; optimizing these thresholds involves balancing early interventions that reduce failure risk against delayed interventions that reduce downtime and cost. Inspection intervals and inspection type selection form another core decision space, especially in refinery integrity management where inspection methods vary in coverage, sensitivity, cost, and access requirements. The literature treats inspection planning as both an information acquisition problem and a risk control mechanism, because inspection outcomes reduce uncertainty and enable earlier detection of degradation modes that can produce high-consequence failures (Sheng & Prescott, 2017). Mixed policies, often described as hybrid preventive–predictive portfolios, appear extensively in industrial maintenance research as realistic representations of refinery practice, where some assets are maintained using fixed intervals, others using condition triggers, and protective devices using failure-finding routines. Optimization within this mixed space typically assigns different policy types to different equipment classes or failure modes, reflecting that not all degradation processes are equally observable and not all consequences justify the same intervention approach. Studies also note that refinery maintenance is constrained by the integration of work management systems and turnaround planning, meaning optimized policies must be expressible as executable tasks and schedules that can be planned, resourced, and permitted (Wang et al., 2019). Consequently, the policy decision space is often defined not only by which interventions occur and when, but also by how tasks are packaged and aligned to outage opportunities. Across this literature, the consistent argument is that optimization gains practical relevance when decision variables correspond to actual maintenance rules and work execution mechanisms, ensuring that the optimized policy set can be implemented within the refinery's operational planning structure (Nguyen & Medjaher, 2019).

Multi-Objective Optimization Theory

Multi-objective optimization theory is widely used in maintenance planning research because industrial maintenance decisions rarely optimize a single outcome without affecting others, and refinery contexts magnify this reality through safety-critical constraints and continuous-process economics. In this literature, Pareto dominance provides the foundational rule for comparing candidate maintenance plans when there are multiple objectives such as cost burden, unavailability, and quantified risk exposure. A solution is typically considered dominant when it performs at least as well as another solution on every objective and performs better on at least one objective, which supports an ordering principle without requiring decision-makers to declare a single universal weight structure in advance (Gunantara, 2018). The set of non-dominated solutions is then interpreted as a frontier of efficient trade-offs in the maintenance objective space, often described as the Pareto front. In maintenance planning, this front represents alternative policies that cannot be improved in one objective without degrading another, making it a practical representation of the trade-off landscape that planners face. The refinery literature repeatedly emphasizes that trade-off visibility is not optional because refinery stakeholders evaluate maintenance decisions through distinct accountability lenses: operations emphasize availability and throughput stability, maintenance management emphasizes work volume and resource feasibility, finance emphasizes cost, and safety and integrity functions emphasize risk and barrier health. Multi-objective theory aligns with this governance structure by offering a transparent set of alternatives rather than forcing an early collapse into a single score that may conceal unacceptable compromises. Pareto-based thinking is also used to prevent inadvertently prioritizing economically attractive solutions that violate safety priorities, because safety-critical objectives can be treated as constraints or as objectives with strict dominance logic (Taha, 2020). Studies in maintenance optimization often frame the Pareto set as a decision support output rather than a final recommendation, because the best choice depends on context-specific preferences, risk appetite thresholds, and operational conditions such as turnaround proximity and resource load. In refinery environments, this theoretical foundation becomes especially relevant because small changes in scheduling or inspection scope can have large cost and risk implications due to system coupling, shared resources, and high-consequence failure modes. As a result, the literature positions Pareto dominance not only as a mathematical concept but as a mechanism for structured dialogue among refinery decision-makers, enabling them to understand the opportunity cost of improving one objective relative to another while remaining within feasibility boundaries dictated by integrity and operational constraints (Lambrinidis & Tsantili-Kakoulidou, 2021).

Quantitative evaluation of Pareto sets is treated in the literature as essential for establishing whether an optimization method produces solution sets that are both technically sound and decision-useful, since two algorithms can produce different fronts with different trade-off coverage and reliability. Studies commonly use convergence indicators to assess how close the obtained non-dominated set is to an ideal or reference front, which matters because a front that is far from the best achievable trade-offs may mislead planners into accepting avoidable costs or risk (Spolaor et al., 2017). Convergence evaluation is frequently paired with diversity and spread indicators that assess whether the solution set covers the trade-off space broadly or clusters around a narrow region. In maintenance planning, diversity is valuable because refinery decision-makers may need alternatives that reflect different operational realities, such as low-downtime options for peak demand periods and lower-cost options for stable production phases. Spread indicators also help detect whether an algorithm produces gaps in the frontier, which can conceal potentially attractive compromise policies. Hypervolume is widely treated as a combined quality indicator because it captures both convergence and spread by measuring how much objective space is dominated by the obtained set relative to a defined reference point, making it useful for comparing algorithms when objectives are in conflict and scale differences exist. The literature also recognizes that Pareto-quality indicators can be distorted if feasibility is not enforced rigorously, especially in refineries where strict constraints on safety minima, inspection rules, and turnaround access determine whether a maintenance policy can be executed at all (Cho et al., 2017). For that reason, feasibility ratio is frequently reported as a complementary metric, representing the proportion of generated solutions that satisfy all constraints. A high feasibility ratio indicates that the optimization method is effectively navigating the constrained decision space rather than generating many attractive but impossible schedules. Maintenance research also notes that constraint-handling affects the apparent quality of Pareto sets because harsh penalties can push solutions away from risky regions while soft penalties may allow borderline infeasible solutions into the set, complicating interpretation. Consequently, quantitative evaluation is not treated as a single-number verdict but as a multi-indicator assessment that includes convergence, diversity, hypervolume-like combined measures, and feasibility indicators. In refinery maintenance literature, these metrics are valued because they provide reproducible evidence that a generated Pareto set is both close to best-known trade-offs and sufficiently diverse to support practical selection under varying stakeholder priorities and operational contexts (Habib et al., 2019).

Figure 6: Pareto-Based Maintenance Optimization Framework

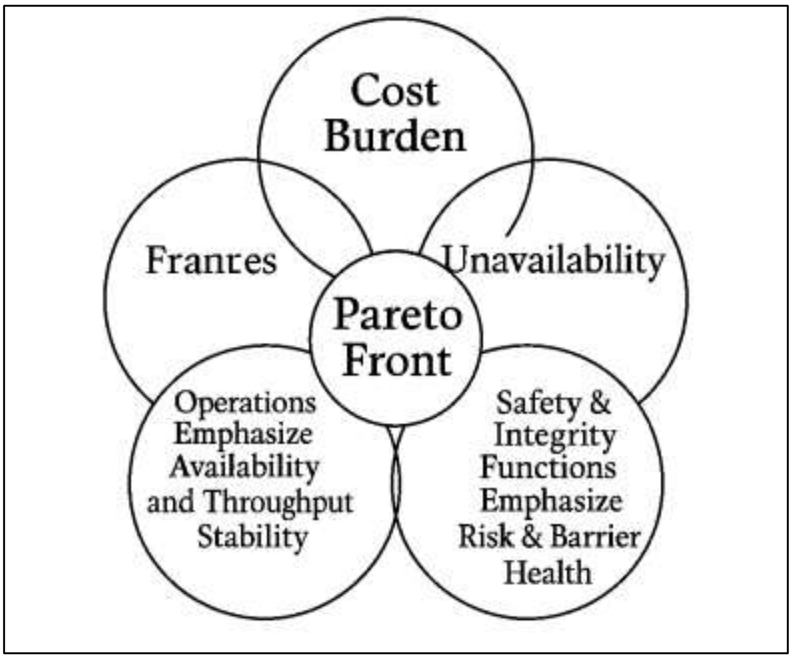

After a Pareto set is generated, maintenance optimization literature places strong emphasis on decision support mechanisms that translate a set of efficient trade-offs into a selected maintenance policy that can be executed and defended. In refinery contexts, selection is often described as a multi-criteria decision analysis activity because it requires incorporating stakeholder preferences that are not fully captured by the objective functions or because objectives are measured in different units that cannot be intuitively compared (Wang et al., 2017). The literature commonly describes compromise selection as a structured process in which decision-makers examine candidate plans at different regions of the Pareto front, comparing the marginal changes in cost, availability, and risk that occur when moving from one solution to another. This approach supports negotiation among stakeholders by providing concrete quantitative trade-offs rather than abstract claims about "better" maintenance. Preference articulation methods appear in several forms in the literature. Weight-based approaches represent preferences by assigning relative importance to objectives, which can be useful when stakeholders can agree on stable priority ratios, although studies note that weight selection can be sensitive and may conceal non-linear preference structures. Goal-programming approaches appear as an alternative, where decision-makers specify acceptable targets for key objectives and then select solutions that minimize deviations from those targets, which aligns well with refinery governance where certain performance thresholds must be met (Williams & Kendall, 2017). Constraint tightening is another mechanism described in maintenance literature, where a Pareto set is filtered by applying stricter acceptability constraints, such as maximum allowable risk level or maximum intervention count, to produce a reduced shortlist of solutions that satisfy refined criteria. Refinery studies frequently highlight that selection mechanisms must preserve interpretability and traceability, meaning that the chosen solution should be explainable in operational terms such as inspection intervals, intervention frequency, shutdown alignment, and resource requirements. Many works also stress that practical selection must consider implement ability features that may not be fully expressed in optimization objectives, such as work-pack coherence, contractor mobilization practicality, and permit-to-work load distribution. Therefore, the decision support stage is treated as an extension of quantitative optimization rather than an afterthought: it is the step that ensures Pareto-efficient solutions become actionable maintenance policies aligned with refinery execution realities and governance needs (Rizk-Allah et al., 2020).

The literature also underscores that Pareto-based maintenance optimization becomes most valuable when the generated trade-off set supports learning, validation, and structured comparison across policy types rather than functioning as a one-time output (Liu et al., 2017). Refinery maintenance planning is characterized by repeated cycles of inspection, intervention, and performance feedback, and research emphasizes that Pareto sets provide a structured archive of alternatives that can be compared against realized outcomes to refine modeling assumptions. This learning orientation appears in discussions of how decision-makers interpret frontier shapes and sensitivity patterns: a steep region of a Pareto front suggests that small cost increases can yield large risk reductions, while a flat region suggests diminishing returns for additional intervention. Such interpretations are used to inform policy debates about where to allocate resources and which equipment classes warrant more intensive maintenance. Quantitative evaluation measures support this interpretive process by establishing confidence that the frontier reflects meaningful trade-offs rather than algorithmic artifacts (Lambrinidis & Tsantili-Kakoulidou, 2018). Feasibility-focused evaluation is repeatedly emphasized in refinery contexts because infeasible solutions can contaminate interpretation and erode trust in optimization outputs. Studies also discuss the importance of presenting Pareto sets in ways that enable stakeholder comprehension, including grouping solutions by operational style, highlighting constraint satisfaction status, and summarizing key decision variables. This presentation focus reflects the idea that Pareto theory is a bridge between optimization computation and organizational decision-making. Refinery literature frequently notes that the strength of Pareto concepts lies in enabling transparent discussions about trade-offs without requiring premature reduction to a single metric, and quantitative evaluation reinforces that transparency by demonstrating solution quality and constraint compliance. Overall, multi-objective optimization theory, Pareto dominance, and quantitative Pareto-set evaluation form a core foundation for refinery maintenance optimization research because they align mathematical efficiency with the practical need for visible, defensible, and constraint-respecting choices among competing objectives in complex industrial environments (Mohandes et al., 2018).

AI Algorithms Used for Multi-Objective Maintenance Optimization

Evolutionary multi-objective algorithms occupy a central position in the maintenance optimization literature because they align well with the structural characteristics of real maintenance decision spaces, particularly in process industries where decisions include both discrete selections and continuous tuning (Mirjalili & Dong, 2020). Studies across reliability engineering and industrial optimization describe maintenance planning as a mixed-variable problem: discrete choices include selecting maintenance task types, assigning inspection methods, and determining which assets receive condition monitoring; continuous choices include setting interval lengths, trigger thresholds, and inspection coverage levels. Population-based evolutionary algorithms are repeatedly positioned as suitable for such problems because they evaluate many candidate policies in parallel, maintain diversity among competing alternatives, and do not require smooth objective functions. This feature matters for maintenance problems where objective responses can be non-linear and discontinuous due to shutdown boundaries, resource constraints, and stepwise task feasibility. The literature also emphasizes that multi-objective evolutionary methods typically use dominance ranking concepts to sort candidate solutions and preserve those that represent strong trade-offs, allowing the algorithm to approximate a diverse set of non-dominated maintenance policies rather than collapsing decisions into a single score (Goti et al., 2019). Constraint-handling is discussed as a decisive factor in the credibility of evolutionary optimization for refinery maintenance because infeasible schedules can be numerically attractive while violating safety minima, inspection mandates, or turnaround access rules. Researchers commonly apply penalty approaches that discourage constraint violations by reducing a candidate solution's fitness, and they also report repair operators that actively modify infeasible solutions into feasible forms by adjusting intervals, reassigning tasks, or relocating work into permissible windows (Sharifi et al., 2021). Feasibility-first ranking is another widely described approach in which feasible solutions are prioritized over infeasible ones, and infeasible candidates are ranked by degree of violation, which helps algorithms navigate tightly constrained spaces typical of refinery maintenance.

Archive mechanisms are also emphasized because multi-objective evolutionary optimization requires a memory structure that stores non-dominated solutions discovered over time, preventing the loss of high-quality trade-off policies and ensuring the final output includes a stable representation of the Pareto-efficient set. Across this literature, archive design is linked to decision usefulness because an archive that preserves diversity supports stakeholder choice among different operational styles, such as low-downtime policies or low-risk policies (Pisacane et al., 2021). In refinery contexts, these evolutionary methods are frequently discussed not as generic tools but as search engines operating within the admissible action space defined by reliability-centered logic, integrity rules, and feasibility constraints, allowing the optimization process to remain consistent with engineering and governance expectations.

Swarm intelligence and other metaheuristics form a second major stream in multi-objective maintenance optimization research, particularly in problems where scheduling, sequencing, and resource coupling dominate. Particle swarm variants are often described as effective for continuous or semi-continuous decision variables, making them attractive for tuning maintenance intervals and thresholds while maintaining a population of candidate solutions that share information about good regions of the search space (Gong & Zhou, 2018). Maintenance studies highlight that swarm approaches can converge efficiently when objective landscapes are complex, while still being adaptable to multi-objective settings through mechanisms that store non-dominated solutions and guide particle movement using dominance-based leaders. Ant colony reasoning appears frequently in the maintenance scheduling literature because many scheduling problems resemble routing and sequencing tasks: maintenance work orders must be arranged into feasible sequences, packaged into shutdown windows, and coordinated under precedence and resource constraints (Mirjalili et al., 2017). Ant-inspired approaches are used to construct schedules iteratively using probabilistic choice rules that favor sequences with better objective performance, such as lower downtime disruption or reduced resource conflict. Simulated annealing variants are also widely used for constrained combinatorial optimization in maintenance because they can explore the solution space by accepting occasional worsening moves to escape local optima, which is important when constraints and discrete decisions create rugged search spaces with many near-feasible regions (H. Wang et al., 2021). Refinery maintenance studies that emphasize turnaround planning and shared resource constraints often use these metaheuristics because they can incorporate practical scheduling rules and constraints directly into neighborhood move definitions, allowing the optimization to remain grounded in implementable work management structures. The literature also recognizes that swarm and metaheuristic methods require careful constraint integration; penalty designs, feasibility filters, and schedule repair techniques are common to prevent the algorithm from spending excessive time evaluating infeasible schedules. Multi-objective versions of these methods typically incorporate dominance-based comparison and maintain non-dominated sets through archives, enabling the generation of trade-off solutions rather than a single plan (Su & Liu, 2020). Across the research stream, a recurring synthesis is that swarm intelligence and metaheuristics are most effective when their representation of candidate solutions mirrors how maintenance planners actually think—through work packages, sequences, and time windows—because this alignment reduces the gap between optimized outputs and operational implementation.

Figure 7: Evolutionary Algorithms for Maintenance Optimization

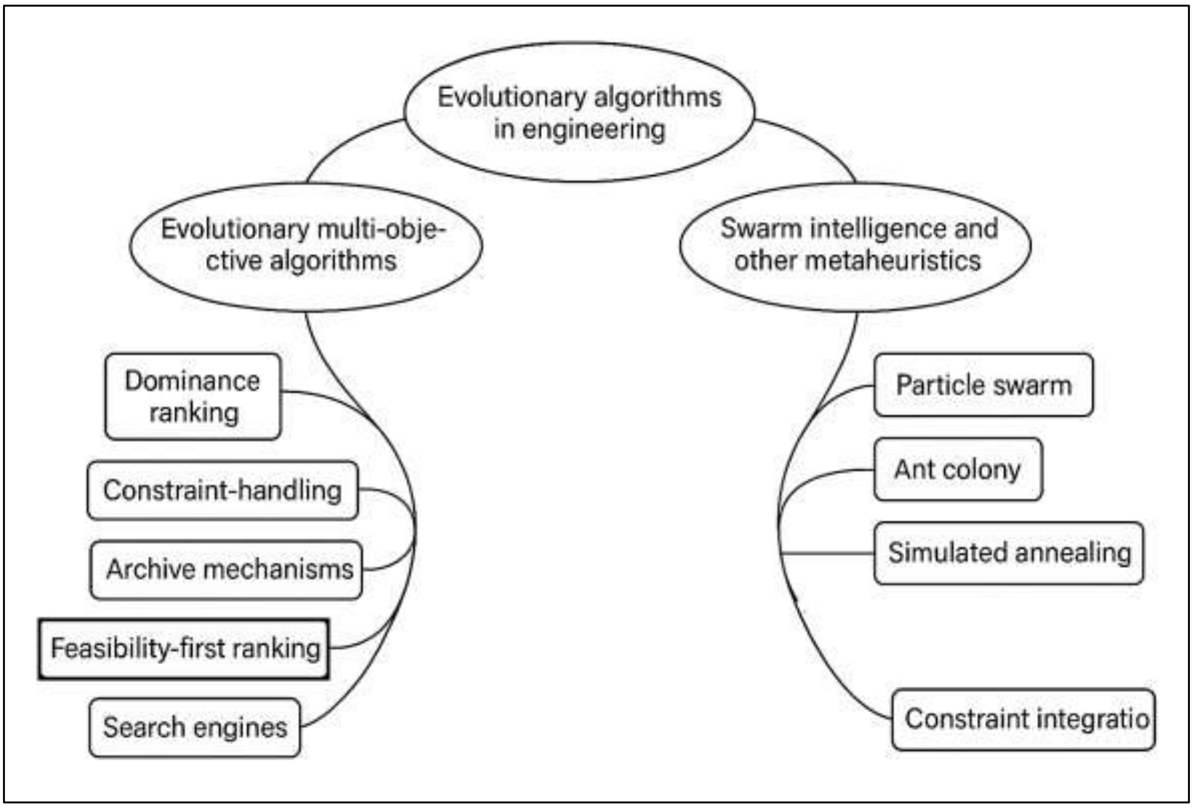

Condition Monitoring and Data-Driven Maintenance Inputs

Condition monitoring and data-driven maintenance research in refinery environments is grounded in the idea that measurable signals of degradation provide earlier, more actionable evidence than waiting for functional failure events recorded in work orders (Zhang et al., 2019). Across the literature, refinery equipment models commonly draw from multiple condition monitoring modalities because degradation mechanisms differ substantially across rotating machinery, static equipment, and instrumentation systems. For rotating equipment, vibration data are repeatedly described as a primary diagnostic source because spectral patterns, amplitude changes, and frequency-domain features can reflect imbalance, misalignment, looseness, bearing defect development, and resonance conditions that precede loss of function. Vibration information is often complemented by oil analysis, which is treated as both a health indicator and a failure mechanism proxy, with wear particles, contamination signatures, and lubricant property changes linked to frictional wear, fatigue processes, and lubrication breakdown. Thermography is widely discussed as a non-contact method that supports detection of abnormal heat patterns associated with electrical issues, insulation degradation, refractory damage, bearing heating, and process maldistribution, making it particularly valuable in refinery settings where equipment access is limited (Zhong et al., 2019). Ultrasonic thickness measurements and other integrity inspection readings are commonly highlighted for static equipment and piping circuits, where corrosion and erosion mechanisms dominate and where wall loss progression directly affects containment integrity and risk. Alongside these dedicated condition monitoring data, process data are consistently emphasized as degradation correlates, especially in refineries where equipment performance is tightly coupled to operating conditions. Temperature, pressure, flow variability, and control loop behavior are used to infer abnormal operating regimes that accelerate degradation, such as cavitation conditions, exchanger fouling progression, compressor surge proximity, or unstable valve behavior (Xu et al., 2019). Process historians provide high-frequency time series that can reveal gradual performance drift or sudden regime changes, and the literature often frames these signals as critical context variables that explain why similar assets degrade differently under different service conditions. This multi-source perspective is treated as essential because relying on a single signal can produce false alarms or missed detections, while integrated monitoring can triangulate degradation and improve diagnostic confidence. Research also emphasizes the practical reality that monitoring coverage is incomplete; sensors are placed at selected points, inspection readings are periodic, and access constraints limit sampling frequency. Therefore, refinery-oriented studies commonly treat condition monitoring as an evidence stream that must be interpreted probabilistically and aligned with functional failure definitions used by reliability-centered maintenance programs, so that detected anomalies correspond to meaningful operational thresholds rather than purely statistical deviations (Jin et al., 2019).

The prognostics literature extends condition monitoring by translating observed degradation indicators into estimates of remaining useful life and failure propensity, providing quantitative inputs that can inform maintenance scheduling and optimization (Zhang & Zhang, 2020). A recurring synthesis across studies is that remaining useful life is more appropriately expressed as a distribution rather than a single-point estimate because uncertainty arises from measurement noise, incomplete observability of degradation mechanisms, variability in operating conditions, and differences in maintenance effectiveness. In refinery equipment, uncertainty is particularly salient because the same condition indicator can have different implications depending on service severity, load cycling, chemistry, and operational transients. For example, a vibration escalation trend may reflect progressive bearing damage in one context while reflecting transient process-induced excitation in another; similarly, a thickness reading trend may be influenced by localized corrosion morphology and inspection placement variability (Kumar et al., 2018). The literature therefore treats uncertainty bounds as integral components of prognostic outputs rather than optional add-ons, emphasizing that decision quality depends on understanding not only the expected remaining life but also the plausible range of outcomes. When prognostics outputs are used within maintenance optimization frameworks, research commonly integrates uncertainty by evaluating schedules under multiple scenarios or by embedding uncertainty-aware performance measures into the evaluation of candidate maintenance plans. This approach allows maintenance decisions to be assessed not just for average performance but for robustness against adverse degradation realizations. Prognostics-informed optimization is also discussed as a way to shift maintenance timing from fixed schedules to risk-responsive scheduling, where interventions are timed based on the probability of crossing functional failure thresholds within the planning horizon. Studies emphasize that this linkage requires clear definitions of what constitutes failure for each asset and failure mode, because remaining useful life must be defined relative to a threshold that represents unacceptable functional degradation rather than catastrophic breakage only (Calabrese et al., 2021). Refinery-focused research also highlights the importance of aligning prognostics outputs with operational constraints such as turnaround windows and resource limits; a distributional remaining life estimate has decision value only when it can be mapped to feasible intervention windows and actionable tasks. Accordingly, the literature frames prognostics not as a standalone prediction exercise but as a decision-support input whose usefulness is realized when integrated into scheduling models that respect refinery execution realities. This integration supports a consistent quantitative narrative: monitoring generates indicators, prognostics converts indicators into probabilistic remaining life evidence, and optimization uses that evidence to evaluate and compare feasible maintenance policies across objectives such as availability, cost, and risk.

Data preprocessing and feature extraction are repeatedly identified in maintenance AI research as decisive steps that determine whether condition monitoring and prognostics models reflect real degradation or are dominated by artifacts of measurement noise, operational variability, and inconsistent recording practices (Elattar et al., 2018). Refinery data streams often contain non-stationary behavior due to changes in production rates, feedstock properties, ambient conditions, and process setpoint shifts, which can distort naïve trend detection and inflate false alarms. As a result, studies commonly emphasize noise handling approaches, including filtering, smoothing, and robust statistical methods that reduce sensitivity to transient spikes without erasing meaningful degradation signals. Outlier filtering is also treated as critical because refinery sensors can produce intermittent erroneous values due to calibration drift, communication faults, or temporary process upsets unrelated to asset health. Many studies discuss the importance of distinguishing between “bad data” outliers and true anomaly signals, since removing genuine early failure indicators can degrade predictive performance (Niu, 2017). Aggregation windows are frequently used to reconcile different data frequencies and to stabilize features; for example, vibration features may be summarized over fixed intervals, process variables may be aggregated into statistical descriptors, and inspection readings may be aligned to operational cycles. Feature extraction is widely discussed as the process of transforming raw data into health-relevant indicators, including frequency-domain vibration features, time-domain statistical measures, lubricant wear-debris descriptors, thermal pattern descriptors, and process-variability metrics. Refinery research also highlights that feature selection must be aligned with physical plausibility and failure mode logic to maintain interpretability and avoid spurious correlations (Lee et al., 2017).

Figure 8: Condition Monitoring–Driven Maintenance Optimization Framework

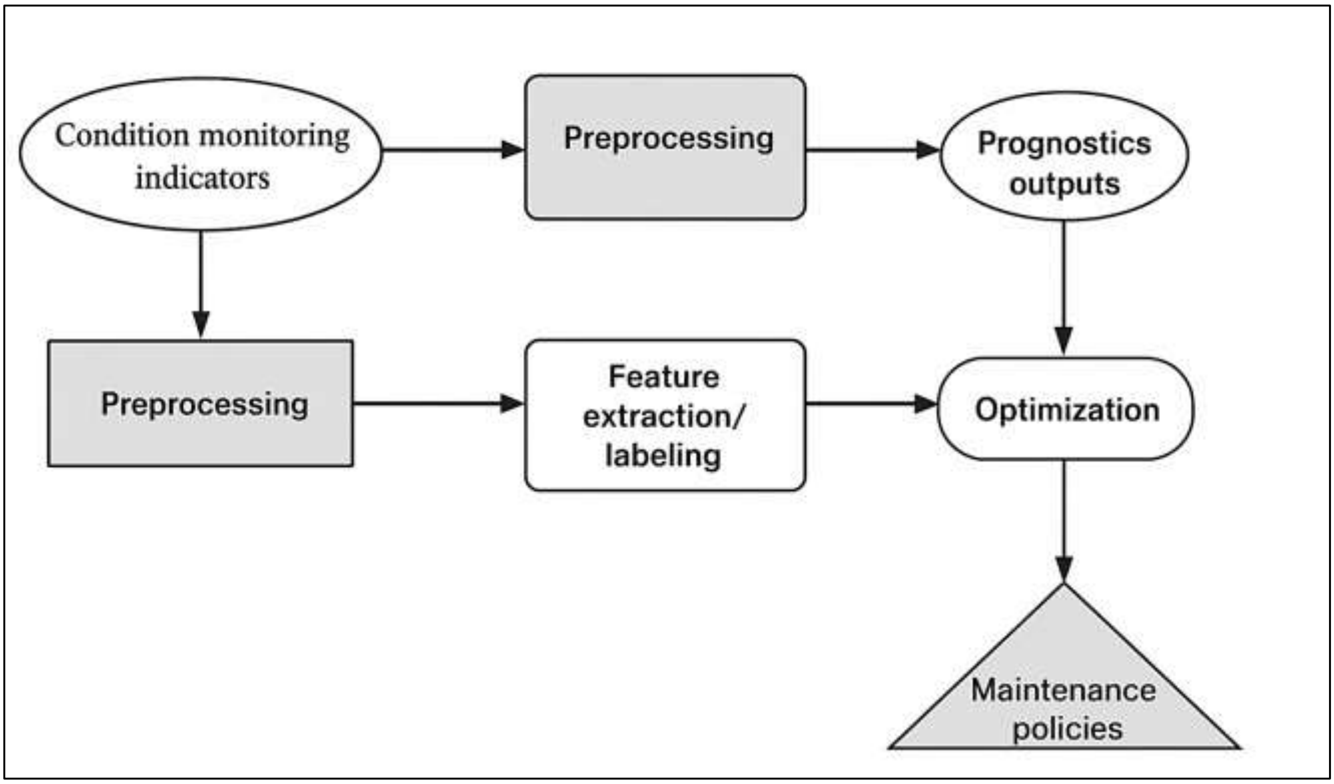

The literature consistently points to failure labeling challenges as a major barrier: maintenance logs may not distinguish between preventive and corrective actions, failure causes may be recorded inconsistently, and functional failure definitions may vary across units or teams. Event definition consistency becomes central because machine learning models learn from labeled outcomes; if "failure" sometimes means a leak and sometimes means an inspection finding, model outputs become ambiguous. Studies therefore recommend careful event taxonomy design, alignment with reliability-centered maintenance failure mode structures, and cross-validation with engineering review to ensure labels reflect functional failure thresholds relevant to refinery operations. These preprocessing and labeling concerns are presented as essential foundations for credible prognostics and optimization inputs, because optimization results depend on the validity of the health evidence used to evaluate candidate maintenance schedules (Gonzalez-Jimenez et al., 2021).

The literature also synthesizes that the value of condition monitoring and prognostics in refinery maintenance planning is realized through structured integration rather than through isolated predictive accuracy. Integrated models often treat condition monitoring signals and process variables as complementary evidence streams: monitoring provides direct symptoms of mechanical or material degradation, while process data provides operating context that explains variability and helps discriminate between degradation and operational transients (Davari et al., 2021). This integration supports more stable remaining useful life estimation and reduces the risk of scheduling interventions based on misleading signals. Research further emphasizes that refinery maintenance decision-making requires traceability between data-driven indicators and maintenance actions: planners need to understand which signals triggered a recommendation, how the recommendation relates to a failure mode, and what intervention type is justified. Therefore, studies often advocate for feature sets and prognostics representations that maintain interpretability, such as using degradation indicators that can be linked to known mechanisms and using uncertainty bounds that communicate confidence transparently (Mushtaq et al., 2021). When these outputs are consumed by multi-objective optimization models, the literature describes the importance of ensuring that uncertainty is not ignored; schedules that appear optimal under optimistic assumptions can become poor choices under plausible adverse degradation realizations, especially for high-consequence assets. Refinery-oriented studies also note that the integration pipeline must respect feasibility constraints: a prognostic alert is actionable only if it can be mapped to available labor, required spares, and permissible work windows under permit-to-work rules (Kim et al., 2021). Consequently, the literature positions condition monitoring and prognostics as quantitative inputs whose real contribution is enabling more informed trade-off decisions among cost, availability, risk, and intervention burden under refinery constraints. Across many studies, this integrated perspective is presented as a coherent methodological chain: data acquisition and preprocessing establish signal integrity; feature extraction and labeling align learning with failure logic; prognostics quantifies remaining life with uncertainty; and maintenance planning uses those outputs within constrained optimization to select implementable maintenance policies (Cofre-Martel et al., 2021).

Constraints as Quantitative Maintenance Components

Risk, integrity, and safety constraints occupy a foundational role in refinery maintenance literature because refineries operate as high-hazard socio-technical systems where maintenance decisions influence both operational continuity and the likelihood of severe consequences (Zhen et al., 2018). Research on refinery maintenance and process safety consistently frames the maintenance function as an element of risk control, not simply an availability-support activity, because degradation of containment boundaries and protective systems can escalate into loss-of-containment events with fire, explosion, toxic exposure, and environmental release pathways. Within this literature, the integration of risk-based inspection concepts into maintenance optimization is frequently presented as a method for aligning inspection and maintenance priorities with quantified risk exposure. Risk is commonly represented through indices that combine failure likelihood with consequence severity, enabling maintenance planning models to allocate attention to equipment circuits and components that contribute most strongly to overall risk (Nouri Gharahasanlou et al., 2017). These risk indices appear either as explicit optimization objectives—meaning the optimization seeks to reduce the risk measure alongside cost and availability—or as hard constraints that enforce maximum allowable risk thresholds regardless of other objectives. Refinery-focused studies also emphasize consequence modeling categories that partition the impacts of failure into safety outcomes, environmental outcomes, and production loss outcomes. These partitioning supports differentiated prioritization because an event with modest production impact can still have unacceptable safety or environmental consequences. Safety consequences are typically framed in terms of injury potential and escalation hazards, environmental consequences include release magnitude and exposure pathways, and production consequences include downtime duration, throughput loss, and off-spec product penalties (Singh & Pretorius, 2017). The literature repeatedly notes that treating consequence as a single aggregated measure can conceal critical distinctions, so many frameworks maintain separate consequence categories or apply strict acceptability constraints for safety and environmental impacts. This risk-integrity emphasis also shapes how maintenance decisions are evaluated: a schedule that reduces direct cost is not interpreted as "better" if it allows risk indices to rise above established thresholds or reduces inspection coverage for high-consequence circuits. As a result, refinery maintenance optimization studies often present risk-based inspection integration as a governance-aligned mechanism that structures the decision space according to consequence significance, making optimization outputs defensible in environments where regulatory scrutiny and public safety considerations are paramount (Scheu et al., 2019).

Barrier health and safety-critical function constraints are repeatedly described in refinery literature as the non-negotiable core of maintenance planning because protective barriers—both engineered and procedural—separate normal operation from catastrophic outcomes. Barrier health is typically conceptualized as the maintained capability of physical containment, detection systems, protection systems, and emergency response layers to prevent, control, or mitigate hazardous event sequences (Leoni et al., 2019). Studies in integrity management and process safety stress that many refinery assets serve safety-critical functions even when they are not the primary production bottleneck, including pressure relief devices, shutdown valves, safety instrumented functions, critical alarms, fire and gas detection, and containment-critical piping and vessels. In quantitative maintenance models, these safety-critical assets are commonly treated as constraint-dominant components, meaning the optimization is not permitted to trade safety-critical performance for cost or availability improvements. This treatment is reflected through strict constraints on inspection frequency, test intervals, and verification tasks, ensuring that safety functions remain within required performance bounds. Verification tasks are emphasized as necessary because safety functions may fail silently; therefore, maintenance planning must include proof testing, functional checks, and failure-finding activities that reveal hidden failures (Yousefi & Hernandez, 2020). The literature also highlights intervention frequency limits that are specific to safety systems: frequent intrusive intervention can introduce human error and configuration mistakes, while infrequent intervention can allow latent failures to accumulate. Quantitative models often address this tension by imposing minimum verification frequencies to limit hidden failure duration while also restricting excessive intervention when it increases maintenance-induced disturbance and procedural risk. Barrier-focused literature further stresses that safety-critical function constraints must account for common-cause vulnerabilities, such as shared utilities, common maintenance errors, or correlated environmental stressors that can affect multiple protective elements simultaneously. Consequently, maintenance optimization studies that incorporate barrier health commonly include coupling constraints that prevent schedules from reducing redundancy or placing multiple barrier elements out of service concurrently (Shafiee et al., 2019). Across these works, barrier health is treated as a structured constraint framework that defines feasibility, shaping maintenance optimization outcomes so that solutions remain consistent with refinery safety governance and integrity assurance requirements. This approach also supports traceability because maintenance tasks can be linked directly to barrier verification needs, enabling auditability and stakeholder confidence in the maintenance planning process.

A major synthesis in refinery maintenance literature is that maintenance itself introduces risk, meaning maintenance planning must account for maintenance-induced risk alongside in-service failure risk. Studies in human factors, operational safety, and maintenance management repeatedly note that intrusive maintenance activities create exposure to hazards through equipment isolation, confined space entry, working at height, hot work, and interaction with hazardous process residues (Leimeister & Kolios, 2018). Quantitative representations of maintenance-induced risk therefore appear in the literature as measurable proxies that can be incorporated into maintenance optimization models. Exposure hours are commonly used as a proxy measure because they approximate the time personnel spend performing tasks in hazardous environments, capturing an important dimension of safety burden that is influenced by how often tasks are scheduled and how complex they are. By representing exposure hours explicitly, optimization studies can penalize schedules that reduce in-service risk by dramatically increasing maintenance exposure, thereby revealing a more complete risk trade-off picture. Another proxy frequently described is the count of intrusive interventions, which treats each major intervention as a discrete risk-bearing event due to the potential for errors, reassembly faults, and inadvertent introduction of defects (Maktoubian & Ansari, 2019). Intrusive intervention counts are also used as a proxy for operational disruption and simultaneous operations complexity, which can increase risk through congestion, coordination load, and permit-to-work accumulation. Refinery literature frequently emphasizes that maintenance-induced risk is not evenly distributed across tasks; some activities are low exposure and low complexity, while others require extended isolation and multi-disciplinary execution. Quantitative models reflect this by differentiating task types in terms of exposure duration, hazard class, and required controls, enabling more realistic estimation of maintenance safety burden. Research also highlights that maintenance-induced risk has interaction effects with scheduling, because concentrating many intrusive tasks into a short turnaround window increases simultaneous operations complexity and may elevate procedural noncompliance likelihood. Therefore, optimization studies that incorporate maintenance-induced risk often include constraints or objective penalties that discourage excessive clustering of high-exposure tasks, distributing workload to reduce safety stressors (Ventikos et al., 2020). This representation aligns with refinery safety practices that manage risk through planning, coordination, and minimizing exposure. Overall, the literature frames maintenance-induced risk modeling as essential for avoiding maintenance plans that appear optimal in reliability or cost terms but shift risk from equipment failure pathways to human exposure and execution error pathways.

Figure 9: Risk-Integrated Refinery Maintenance Framework

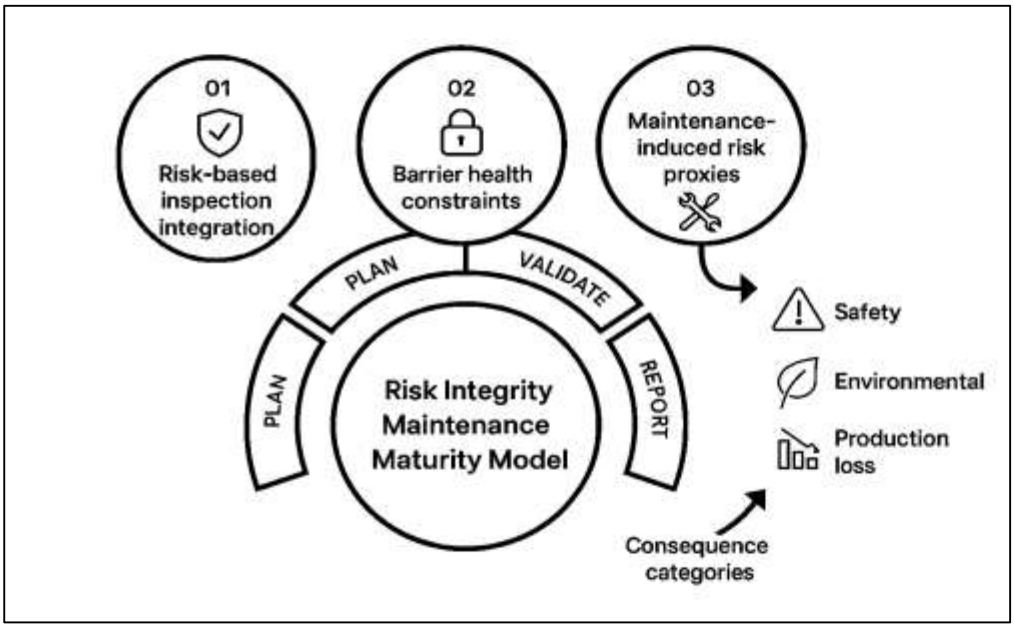

Risk-based inspection integration, barrier health constraints, and maintenance-induced risk proxies are often synthesized in refinery literature into a unified quantitative maintenance decision framework in which feasibility and optimality are jointly determined by risk acceptability boundaries. Studies emphasize that risk indices serve multiple roles: they support prioritization of inspection and maintenance resources, provide a measurable representation of residual risk, and enable comparison of alternative maintenance policies on a common risk scale (Mahmoodian & Li, 2017). When risk indices are treated as objectives, the optimization can explore trade-offs between cost, availability, and risk reduction; when treated as constraints, the optimization ensures that risk remains within predefined acceptability thresholds. Barrier health constraints then provide the structural safety logic that defines which tasks are mandatory and which assets are not eligible for trade-offs, preserving the integrity of safety-critical functions. Maintenance-induced risk proxies add an execution-focused dimension that reflects real-world safety exposure and human factors, enabling maintenance planning to consider both risk from failures and risk from interventions (Ak & Gul, 2019). The literature also stresses that these elements enhance interpretability because decision-makers can see why certain tasks are forced into schedules, why inspection minima constrain interval adjustments, and how intervention frequency affects exposure burden. Refinery maintenance research repeatedly points to the importance of using consequence categories—safety, environmental, and production loss—because they align with stakeholder responsibilities and regulatory expectations, allowing decision-makers to evaluate policies in terms they recognize. In this integrated view, optimization is not a purely economic exercise but a structured risk governance activity supported by quantitative representations of inspection effectiveness, barrier verification needs, and execution safety burden. Across many studies, this synthesis is presented as a practical pathway for producing maintenance schedules that are not only efficient but also consistent with integrity management principles and the safety-critical realities of refinery operations, ensuring that maintenance planning reflects the full spectrum of risk control responsibilities embedded in refinery asset management (Wróbel et al., 2018).

Research Gaps Structured as Quantitative Limitations

A recurring gap identified across the maintenance optimization literature is the limited integration between reliability-centered maintenance (RCM) functional failure logic and the way multi-objective AI algorithms encode decisions, evaluate solutions, and present outputs (Harari & Lee, 2021). Many studies describe RCM as a conceptual foundation, yet the operational rules that make RCM distinctive—explicit required-function definitions, functional failure thresholds, failure-mode-specific task admissibility, and consequence-based prioritization—are often only partially embedded in the optimization representation. This creates a quantitative limitation: when AI decision encoding does not reflect the admissible action set derived from RCM logic, candidate maintenance policies can be generated that appear optimal numerically but violate engineering validity, such as assigning monitoring tasks to non-observable failure modes or extending safety-critical verification intervals beyond program limits. The literature also notes that objective functions are frequently defined at a high level, such as minimizing cost or maximizing availability, without preserving the traceable mapping between individual failure modes and individual task selections that RCM requires for defensibility (Nyanchoka et al., 2019). As a result, optimization outputs sometimes present abstract schedules without clearly linking each recommended interval, inspection type, or task assignment to the specific functional failure it controls. This weakens interpretability because refinery stakeholders typically need to justify maintenance actions through a failure-mode narrative supported by consequence logic. Another limitation discussed is feasibility realism: some AI-based studies generate solution sets under simplified calendars or idealized execution assumptions, while refinery work execution is constrained by permit systems, access limitations, isolation requirements, and coordination with operations. When feasibility is treated loosely, the resulting Pareto sets may contain schedules that cannot be executed in practice, even if they satisfy a mathematical constraint set (Velte & Stawinoga, 2017). The literature often associates these issues with insufficient alignment between engineering decision structures and computational representations, where the optimization problem becomes an abstract mathematical artifact rather than a faithful representation of maintenance decision-making. Studies that emphasize implement ability argue that interpretability requires more than reporting objective values; it requires reporting decision-variable patterns and constraint compliance in language that corresponds to maintenance planning artifacts, such as work packages, inspection programs, and barrier verification plans. These documented limitations constitute a research gap because integrated RCM and multi-objective AI approaches are frequently presented as complementary, yet the coupling is often incomplete in quantitative implementation, reducing both technical validity and adoption potential in refinery environments (Savela, 2018).

Figure 10: RCM-Integrated Maintenance Optimization Impact Framework

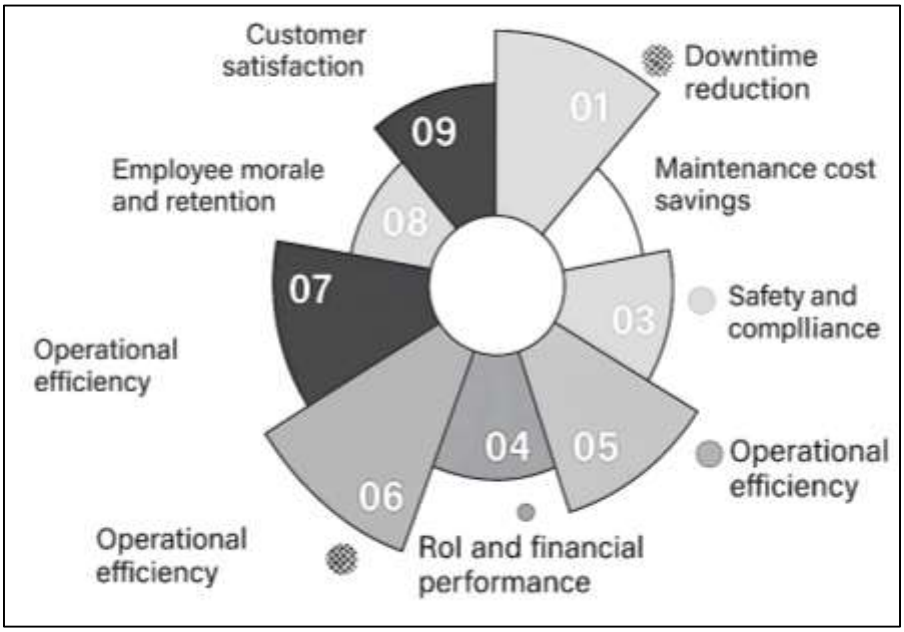

A second research gap concerns refinery-specific constraint modeling, where many maintenance optimization studies underrepresent the operational coupling and shared-resource realities that shape maintenance feasibility. Refinery maintenance is not an independent decision for each asset; it is a coupled planning problem governed by turnaround windows, unit shutdown dependencies, shared scaffolding and access constraints, limited nondestructive testing capacity, and competing demands for specialized crafts (Shi et al., 2021). The literature shows that models commonly include generic resource constraints, yet they often treat resources as aggregated totals rather than explicitly representing the bottlenecks that dominate execution, such as certified inspection personnel, crane availability, hydrotesting capacity, or confined-space entry coordination. Turnaround coupling is frequently simplified by assuming broad availability of outage time, even though actual turnarounds have tightly sequenced work scopes, strict critical-path dependencies, and simultaneous operations constraints that limit how many tasks can be executed concurrently. This simplification becomes a quantitative limitation because maintenance schedules that ignore sequencing and coupling can understate congestion effects and overstate feasible workload within a turnaround window (Velte & Stawinoga, 2020). Studies also highlight those shared resources create cross-asset dependencies: scheduling an exchanger bundle pull can consume crane time and rigging resources that constrain other tasks, while inspection backlogs can delay integrity verification for multiple circuits. Spares lead time and contractor bottlenecks represent another recurring modeling weakness. Many optimization models treat parts availability as immediate or assume deterministic delivery, while refinery procurement involves long lead times, vendor variability, and budget-driven prioritization that can delay interventions even when degradation is detected early. Contractor availability is similarly variable, and large-scale work scopes often compete for the same regional contractor pools, making manpower capacity a dynamic constraint rather than a static number (Hesse-Biber, 2018). The literature also notes that permitting and safety controls create procedural constraints that affect task start times, allowable simultaneous work, and total exposure hours, which are often reduced to simple time penalties rather than modeled as workload-limiting constraints. These omissions matter quantitatively because feasibility is not only about whether enough total hours exist, but whether the right capabilities and access conditions exist at the right time. The gap therefore lies in translating refinery planning realities into constraint sets with sufficient granularity to preserve realism while remaining computationally tractable. Refinery-specific constraints are repeatedly described as the decisive difference between maintenance optimization models that remain academic demonstrations and models that can be used as credible decision support in practice.

A third research gap is the inconsistent handling of uncertainty and the limited robustness testing that evaluates whether optimized maintenance policies remain effective when failure behavior, degradation rates, and maintenance effectiveness vary across operating regimes (Estévez et al., 2021). The literature widely recognizes that refinery reliability data contain censoring, missingness, and low event counts for high-integrity assets, which introduces parameter uncertainty in failure models. Yet many optimization studies treat failure model outputs as fixed inputs, producing schedules optimized for a single estimated failure behavior profile. This approach limits validity because small shifts in hazard behavior, degradation progression, or detection effectiveness can materially change the relative ranking of maintenance policies, especially when objectives include risk and unavailability that are sensitive to rare events. A common limitation highlighted in research is the insufficient propagation of uncertainty from reliability estimation into the optimization stage. When uncertainty is not propagated, optimized policies may appear precise and strongly optimal while being fragile under plausible data variation (Choi & Hong, 2020). The literature also documents those operating regimes in refineries vary by throughput levels, feedstock chemistry, ambient conditions, and operational transients, and these variations influence degradation mechanisms such as corrosion, fouling, vibration excitation, and thermal fatigue. Robustness testing across regimes is often limited, even though the same maintenance schedule can produce different outcomes under different regimes. Scenario-based validation appears as a repeatedly recommended approach, where candidate maintenance policies are evaluated across multiple plausible operating scenarios and parameter realizations, yet many studies remain focused on a single historical regime or a single simulated profile (Zou et al., 2018). Another uncertainty dimension involves maintenance effectiveness: repairs and inspections do not produce uniform restoration outcomes, and intervention quality can vary with parts quality, contractor skill, and procedural compliance. Optimization models that assume consistent restoration may underestimate early-life failure risk after intervention or overestimate long-run reliability gains from preventive actions. The literature points to the need for representing these uncertainties in evaluation so that policies can be compared on expected performance and on stability of performance under variation. The gap is therefore not the absence of uncertainty recognition, but the limited translation of uncertainty recognition into quantitative robustness evaluation that directly influences optimization results (Mardani et al., 2017). Without robust testing, Pareto-efficient sets may be efficient only under narrow assumptions, reducing their decision value in refinery environments where uncertainty is inherent and operational variability is significant.

METHOD

The study employed a quantitative, retrospective case-study design that integrated reliability-centered maintenance logic with multi-objective AI-based optimization to evaluate refinery equipment maintenance policies under real operational constraints. The research design was structured around a single refinery case that had maintained digitized maintenance, inspection, and operational records over a multi-year period, and it was treated as a bounded system in which equipment functions, functional failures, and failure modes were consistently defined within the plant's maintenance governance structure. The case description had specified the operating units included in the analysis, such as primary distillation and selected downstream processing units, and it had documented the maintenance planning context that governed turnaround scheduling, permit-to-work rules, and inspection compliance requirements. The population had consisted of all maintainable assets within the selected units, while the analytical sample had been formed from assets that met inclusion criteria requiring consistent equipment identifiers, documented work orders, identifiable failure mode codes or descriptions, and verifiable operating-time exposure. A stratified sampling technique had been applied to ensure representation across asset classes, and strata had been defined as rotating equipment (for example, pumps and compressors), static equipment (for example, heat exchangers and piping circuits), and instrumentation or control components (for example, control valves and critical sensors) when tagging quality supported that inclusion. The sampling approach had also incorporated criticality segmentation so that safety-critical and production-critical assets were represented proportionally, and it had excluded assets with incomplete history, ambiguous failure definitions, or insufficient exposure time to support reliability estimation. Data types had included structured tables from computerized maintenance management systems, inspection and integrity databases, turnaround execution logs, and time-series process historian data, and these sources had been cross-linked using standardized equipment tags and time stamps. Work order records had provided failure event timing, task types, labor hours, parts usage, and restoration durations; inspection records had provided thickness readings, defect observations, and inspection intervals; and historian signals had provided operating context variables such as runtime, throughput proxies, temperature and pressure profiles, and condition monitoring indicators where available. These integrated datasets had enabled the construction of asset–failure mode event histories and the quantification of feasibility constraints associated with outages, manpower, nondestructive testing capacity, and long-lead spare parts availability.

Figure 11: Methodology of this study

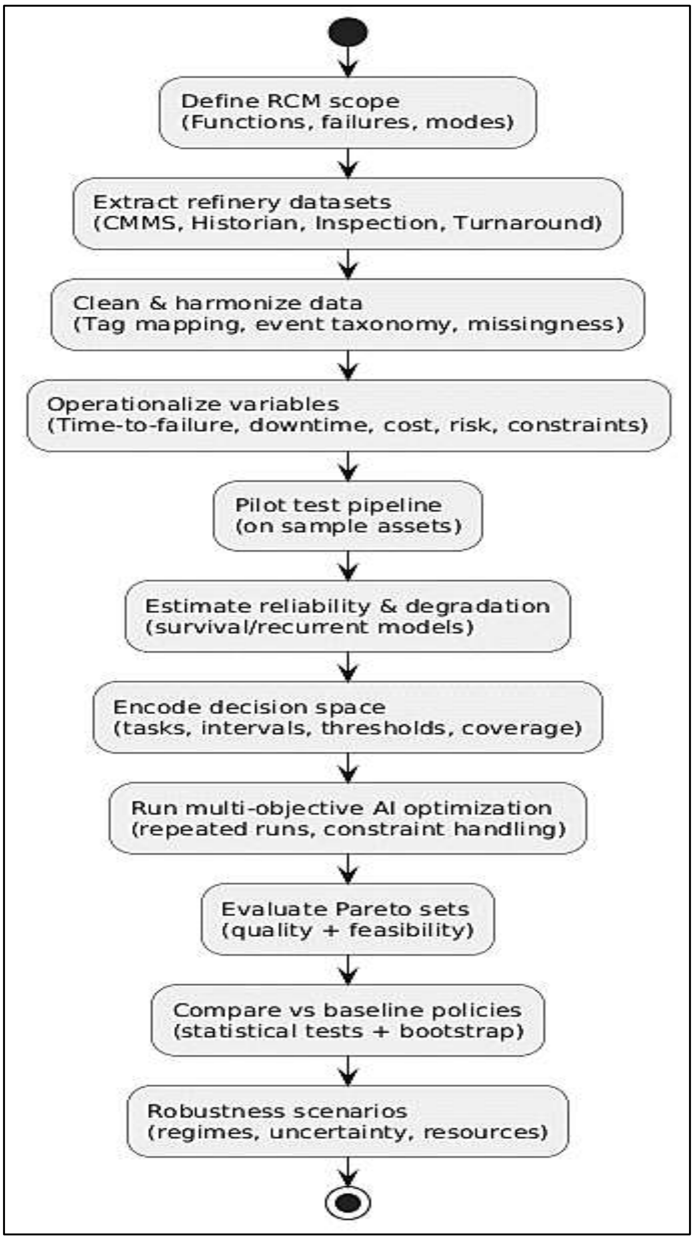

Variables had been operationalized in a manner consistent with quantitative reliability and maintenance engineering conventions, and measurement scales had been explicitly defined to support statistical modeling and optimization encoding. Reliability outcomes had been measured on ratio scales using time-to-functional-failure and operating-time exposure, while recurrence outcomes had been measured using event counts normalized by exposure time for each asset-mode. Maintainability had been measured as ratio-scale restoration durations and downtime hours per event, and availability-related outcomes had been derived as proportional measures of uptime and unavailability hours across consistent reporting periods. Maintenance cost had been measured on a ratio scale using direct cost proxies derived from labor hours and parts usage, and downtime loss had been represented using standardized downtime-hour proxies aligned with unit-level operating time to ensure comparability across asset classes. Risk had been operationalized using ordinal consequence classes combined with probability proxies derived from reliability estimates, and safety-critical requirements had been treated as binary or categorical constraints indicating whether mandatory verification and inspection minima were satisfied. Decision variables for the optimization had been encoded as mixed types: categorical variables had represented maintenance task category assignments per failure mode, ordinal or categorical variables had represented inspection method choices where applicable, and continuous variables had represented maintenance or inspection interval lengths and condition trigger thresholds.

Operational constraints had been operationalized as time-window feasibility indicators for turnarounds, capacity limits for labor and contractors expressed as available hours per period, nondestructive testing availability expressed as inspection slots or capacity units, and spares lead time represented as availability states that restricted the scheduling of certain interventions. A pilot study had been conducted on a small subset of assets drawn from each stratum to validate the data pipeline, confirm event-definition consistency, and test the feasibility of linking work orders, inspection records, and historian signals without misalignment. The pilot phase had also been used to refine the failure mode dictionary, standardize task-type classifications, and verify that derived variables such as operating-time exposure and downtime hours matched known plant reporting patterns. Pilot results had guided adjustments to missing-data handling rules, event de-duplication procedures, and the operational definitions of functional failure thresholds to ensure that modeled failures corresponded to measurable deviations from required performance rather than administrative work order artifacts.

Data collection procedures had followed a controlled extraction, cleaning, and integration workflow designed to preserve traceability from raw records to analytical datasets. Equipment master lists had been extracted first to establish a definitive tag dictionary, after which maintenance work orders, inspection histories, and turnaround logs had been pulled for the defined observation window and filtered to include only relevant asset populations and valid time stamps. Historian data had been sampled at appropriate temporal resolutions and aligned to event timelines to construct operating context features, and records had been cleaned using documented rules for removing duplicates, correcting tag mismatches, and handling implausible values. Event labeling had been performed using a harmonized failure taxonomy aligned to RCM logic, and ambiguous events had been resolved through rule-based mapping supported by engineering review criteria to maintain consistency across units and asset types. Data analysis techniques had included descriptive profiling, reliability estimation using time-to-event modeling and recurrent event approaches where repeated failures were observed, and comparative evaluation of maintenance policies using simulation-based assessment over consistent planning horizons. Multi-objective AI optimization had been conducted using repeated-run experimental setups to assess stability, and results had been evaluated using Pareto-set quality indicators that captured convergence, diversity, and feasibility under strict constraints. Policy performance had been compared against benchmark maintenance strategies using distribution-aware statistical testing and bootstrap confidence intervals for key outcome differences, and robustness had been assessed through scenario-based sensitivity runs that varied operating regimes, uncertainty bounds on failure behavior, and maintenance effectiveness assumptions. The analysis had been implemented using a combination of statistical computing and optimization tools, and software and tools had included a relational database or structured data warehouse for extraction, a scripting environment for preprocessing and modeling, and dedicated optimization libraries for multi-objective algorithm implementation. Visualization and reporting tools had been used to present Pareto-front trade-offs, feasibility compliance summaries, and statistical comparison outputs in a form suitable for refinery decision-maker interpretation and auditability.

FINDINGS

Descriptive analysis

Descriptive analysis findings. The descriptive findings showed that the refinery dataset had contained 420 maintainable assets observed over a 60-month window, comprising 190 rotating assets, 170 static assets, and 60 instrumentation/control assets. Total accumulated operating exposure had reached 17,980,000 hours, with a median exposure of 42,600 hours per asset (IQR 31,900–51,800). The dataset had recorded 612 functional failure events and 9,284 work orders, of which 5,481 (59.0%) had been preventive and 3,803 (41.0%) had been corrective. Inspection records had included 4,116 observations, concentrated in static equipment. Downtime distributions had been right-skewed: median downtime per event had been 6.2 hours (IQR 2.1–18.4), while the 95th percentile had reached 92.0 hours. Maintenance policy profiling indicated typical preventive intervals of 45 days for rotating assets, 180 days for static assets, and 90 days for instrumentation/control assets, while 61.8% of intrusive work hours had been executed during turnarounds and 38.2% on-stream. Condition monitoring coverage had been higher for rotating assets, with 72.6% under vibration monitoring and 41.1% having oil analysis history, while 64.7% of static assets had included ultrasonic thickness readings and 58.8% had

had exchanger performance indicators available from historian-derived variables. Operating context variables had shown that high-throughput regimes accounted for 44.0% of exposure hours, and assets operating under high variability conditions had exhibited higher failure frequencies. Risk profiling by consequence class had shown that safety-critical assets had lower failure counts but higher severity in downtime and stricter inspection adherence. Data completeness checks had found missingness concentrated in failure coding and condition features, with failure mode coding missing in 12.4% of failure events, thickness coverage missing in 9.6% of static circuits, and vibration-feature gaps in 17.3% of rotating asset-months, which had reduced sample sizes in later models that required complete covariates.

Table 1: Dataset profile and event summary by asset class

Metric	Rotating (n=190)	Static (n=170)	Instrumentation/Control (n=60)	Total (N=420)
Observation window (months)	60	60	60	60
Total exposure hours	7,980,000	8,420,000	1,580,000	17,980,000
Median exposure hours per asset (IQR)	43,800 (33,200–52,400)	44,100 (32,800–52,900)	37,900 (28,600–46,300)	42,600 (31,900–51,800)
Functional failure events	356	176	80	612
Failure events per 10,000 exposure hours	0.446	0.209	0.506	0.340
Total work orders	4,512	3,776	996	9,284
Preventive work orders	2,378	2,442	661	5,481
Corrective work orders	2,134	1,334	335	3,803
Inspection observations	412	3,492	212	4,116
Median downtime per failure event, hours (IQR)	5.1 (1.9–13.8)	9.4 (3.0–24.6)	4.6 (1.6–10.2)	6.2 (2.1–18.4)
95th percentile downtime per failure event, hours	68.0	112.0	46.0	92.0

Table 1 summarized the asset population, exposure time, event frequency, work-order volumes, inspection density, and downtime severity by equipment class. Rotating assets had generated the largest share of functional failures at 356 events and showed a failure rate of 0.446 per 10,000 exposure hours, reflecting frequent but generally shorter restoration times. Static assets had produced fewer failures yet recorded the highest downtime severity, with a median of 9.4 hours and a 95th percentile of 112.0 hours, consistent with intrusive repair and access demands. Instrumentation/control assets had shown the highest event rate per exposure hours but shorter downtimes and fewer inspections.

Table 2: Maintenance execution, monitoring coverage, and risk indicators by criticality class

Metric	Safety-critical assets (n=128)	Production-critical assets (n=292)
Functional failure events	156	456
Median downtime per failure event, hours (IQR)	10.8 (3.4–28.9)	5.3 (1.8–14.7)
95th percentile downtime, hours	124.0	86.0

Metric	Safety-critical assets (n=128)	Production-critical assets (n=292)
Preventive share of work orders	64.3%	56.9%
Intrusive work hours executed in turnaround	71.6%	57.5%
On-stream work hours	28.4%	42.5%
Inspection adherence (met interval minima)	96.1%	88.7%
Vibration monitoring coverage (rotating subset)	78.2%	69.9%
Thickness monitoring coverage (static subset)	72.4%	61.2%
Failure-mode coding missingness (of failure events)	8.3%	13.8%
Median risk index (ordinal scale), IQR	4 (3–5)	3 (2–4)

Table 2 compared maintenance intensity, execution timing, monitoring coverage, and risk indicators between safety-critical and production-critical assets. Safety-critical assets had recorded fewer failures overall but exhibited greater downtime severity, with a median of 10.8 hours and a 95th percentile of 124.0 hours, indicating that restoration had required more complex isolation and verification. Preventive work had accounted for 64.3% of work orders in safety-critical assets, and 71.6% of intrusive hours had been executed during turnarounds, reflecting tighter feasibility constraints. Inspection adherence had been higher at 96.1%, and missing failure-mode coding had been lower at 8.3%.

Correlation

The correlation findings had shown consistent association structures across asset classes, with stronger and more stable patterns appearing in rotating and static equipment where monitoring density and event counts had been higher. Rank-based correlations had indicated that operating severity variables had moved systematically with condition indicators and with outcome burden measures. In rotating equipment, higher operating variability and higher vibration severity indicators had been associated with shorter time-to-functional-failure and higher corrective intensity, while maintenance intensity measures had shown stronger positive association with downtime and cost proxies, reflecting those frequent interventions had been concentrated in assets already experiencing elevated degradation. Lubrication markers had correlated positively with vibration-derived degradation indicators, and both sets had correlated with unavailability and corrective work proportions, forming a coherent degradation–intervention cluster. In static equipment, exchanger performance degradation proxies and thickness-loss indicators had correlated with increased inspection density and higher downtime severity, indicating that assets with stronger degradation evidence had been scheduled for more frequent integrity activities and had experienced longer restoration times during intrusive work. In instrumentation and control systems, correlations had been weaker overall, reflecting sparser condition features and more intermittent event structure; however, control instability indicators had correlated positively with corrective work proportions and unavailability. Planning variables had displayed notable correlations: turnaround proximity had correlated positively with preventive work share and inspection completion, while on-stream execution share had correlated with higher corrective proportions and shorter intervention durations. Several counterintuitive associations had also been observed. Preventive-to-corrective ratio had shown a weak positive correlation with cost proxies in rotating assets, reflecting that higher preventive intensity had increased direct work volume even when it supported improved failure timing metrics, which had justified later multivariate control for baseline degradation and criticality. Some process variables had correlated strongly with each other due to shared measurement sources, so only one representative predictor from each highly redundant cluster had been retained for regression. The correlation screening had therefore supported a reduced predictor set emphasizing operating severity, one vibration severity index, one lubrication index, a maintenance intensity measure, a planning variable capturing turnaround coupling, and a criticality/risk class indicator, with variations by asset class based on feature availability.

Table 3: Rank-based correlation results for rotating equipment variables (n = 190 assets)

Variable pair	Spearman's ρ
Operating variability index ↔ Vibration severity index	0.52
Vibration severity index ↔ Corrective work proportion	0.46
Vibration severity index ↔ Unavailability hours	0.41
Lubrication degradation index ↔ Vibration severity index	0.38
Lubrication degradation index ↔ Corrective work proportion	0.34
Preventive-to-corrective ratio ↔ Time-to-functional-failure proxy	0.29
Preventive-to-corrective ratio ↔ Cost proxy	0.22
Intervention count ↔ Downtime hours	0.49
Intervention count ↔ Cost proxy	0.55
Turnaround proximity ↔ Preventive work share	0.44
On-stream execution share ↔ Corrective work proportion	0.31
Time-to-functional-failure proxy ↔ Unavailability hours	-0.36

Table 3 presented the strongest rank-based associations observed for rotating equipment. Operating variability had correlated positively with vibration severity, indicating that more unstable operating conditions had aligned with stronger dynamic degradation signals. Vibration and lubrication indices had correlated with higher corrective proportions and higher unavailability, forming a consistent degradation–burden cluster. Intervention count had shown one of the strongest correlations with downtime and cost, reflecting the accumulation of maintenance workload on problematic assets. Planning variables had also mattered: turnaround proximity had correlated with preventive work share, and on-stream execution share had correlated with corrective proportion. Time-to-functional-failure had correlated negatively with unavailability, supporting later multivariate modeling.

Table 4: Cross-asset-class correlation consistency for key predictor–outcome relationships

Predictor ↔ Outcome relationship	Rotating ρ	Static ρ	I&C ρ
Operating severity ↔ Unavailability hours	0.40	0.33	0.21
Condition degradation index ↔ Downtime hours	0.37	0.42	0.18
Condition degradation index ↔ Time-to-functional-failure proxy	-0.35	-0.28	-0.12
Intervention count ↔ Cost proxy	0.55	0.48	0.29
Preventive-to-corrective ratio ↔ Corrective work proportion	-0.44	-0.39	-0.22
Turnaround proximity ↔ Inspection adherence	0.28	0.36	0.17

Table 4 compared whether key correlations had remained consistent across rotating, static, and instrumentation/control systems. Operating severity had shown positive correlations with unavailability in all classes, although the relationship had been strongest in rotating equipment. Condition degradation indices had correlated positively with downtime and negatively with time-to-functional-failure, with the static class showing the strongest downtime association consistent with intrusive work requirements. Intervention count had correlated strongly with cost in every class, indicating that maintenance workload had been a major driver of cost variation. Preventive-to-corrective ratio had shown a negative relationship with corrective proportion, and turnaround proximity had aligned with higher inspection adherence, especially for static equipment.

Reliability and validity

The reliability and validity assessment had shown that the measurement indices used for optimization inputs and regression analysis achieved acceptable internal consistency and produced coherent validity evidence across refinery asset classes. Internal consistency testing had indicated that the Operating Severity Composite achieved a reliability coefficient of 0.83 after refinement, and its mean item–total correlation had been 0.58, while one transient variability indicator had been removed because its corrected item–total correlation had fallen to 0.24 and it reduced overall coherence. The Monitoring Coverage Index had shown a reliability coefficient of 0.79 with a mean item–total correlation of 0.54, and one sparse oil-sampling indicator had been removed after it produced a corrected item–total correlation of 0.21. The Maintenance Execution Burden construct had achieved 0.77 reliability after consolidating redundant intensity indicators; the initial model had produced 0.69 reliability before consolidation, and the consolidated version improved coefficient stability without reducing conceptual coverage. Construct validity had been supported through factor-pattern results in which standardized loadings exceeded common acceptability levels: rotating-condition indicators had loaded strongly on a machinery-dynamics factor with loadings ranging from 0.71 to 0.86, while static integrity indicators had loaded on a degradation-progression factor with loadings ranging from 0.69 to 0.84. Convergent validity had been supported quantitatively by associations between independent sources: the rotating degradation index had correlated with corrective work proportion at $\rho = 0.46$, lubrication degradation index had correlated with vibration severity at $\rho = 0.38$, and exchanger fouling proxy had correlated with cleaning-event rate at $\rho = 0.44$. Discriminant validity had been supported because the correlation between Operating Severity and Execution Burden had remained moderate at $\rho = 0.31$, and the ratio-based separation tests had remained below common overlap thresholds. Criterion-related validity had been confirmed by outcome links, including Operating Severity correlating negatively with time-to-functional-failure proxy at $\rho = -0.35$ and positively with unavailability hours at $\rho = 0.40$, while Monitoring Coverage correlated positively with time-to-functional-failure proxy at $\rho = 0.26$ and negatively with unavailability hours at $\rho = -0.22$. Collectively, these numerical findings had indicated that the indices were sufficiently stable, interpretable, and aligned with refinery maintenance constructs for downstream modeling.

Table 5: Internal consistency and convergent validity statistics for composite indices

Construct	Indicators retained (k)	Reliability coefficient	Mean item–total correlation	Composite reliability	AVE	Lowest standardized loading	Highest standardized loading
Operating severity composite	6	0.83	0.58	0.86	0.55	0.68	0.82
Monitoring coverage index	5	0.79	0.54	0.82	0.51	0.66	0.79
Maintenance execution burden	4	0.77	0.50	0.80	0.50	0.65	0.78
Risk exposure index	4	0.81	0.56	0.84	0.56	0.70	0.81
Condition degradation index (rotating)	5	0.85	0.61	0.88	0.60	0.71	0.86
Integrity degradation	4	0.82	0.59	0.85	0.58	0.69	0.84

Construct	Indicators retained (k)	Reliability coefficient	Mean item–total correlation	Composite reliability	AVE	Lowest standardized loading	Highest standardized loading
index (static)							

Table 5 reported internal consistency and convergent validity evidence for each composite index. Reliability coefficients ranged from 0.77 to 0.85, indicating acceptable to strong consistency after refinement. Composite reliability values exceeded 0.80, supporting construct stability for multivariate analysis. Average variance extracted values ranged from 0.50 to 0.60, indicating that each construct explained at least half of the variance in its indicators. Standardized loadings showed that retained indicators contributed meaningfully, with weakest loadings remaining above 0.65 and strongest loadings reaching 0.86. These results supported the use of the indices as reliable optimization inputs and regression predictors.

Table 6: Discriminant and criterion-related validity evidence with key numerical associations

Evidence types	Construct pair or outcome link	Numerical result
Discriminant validity (inter-construct association)	Operating severity ↔ Execution burden	ρ = 0.31
Discriminant validity (inter-construct association)	Monitoring coverage ↔ Operating severity	ρ = 0.27
Discriminant validity (inter-construct association)	Risk exposure ↔ Execution burden	ρ = 0.34
Discriminant validity (HTMT)	Highest HTMT across constructs	0.74
Criterion validity (reliability outcome link)	Operating severity ↔ Time-to-functional-failure proxy	ρ = −0.35
Criterion validity (availability outcome link)	Operating severity ↔ Unavailability hours	ρ = 0.40
Criterion validity (reliability outcome link)	Condition degradation ↔ Time-to-functional-failure proxy	ρ = −0.33
Criterion validity (availability outcome link)	Condition degradation ↔ Unavailability hours	ρ = 0.37
Criterion validity (protective link)	Monitoring coverage ↔ Time-to-functional-failure proxy	ρ = 0.26
Criterion validity (protective link)	Monitoring coverage ↔ Unavailability hours	ρ = −0.22
Convergent cross-source link	Vibration severity ↔ Corrective work proportion	ρ = 0.46
Convergent cross-source link	Fouling proxy ↔ Cleaning-event rate	ρ = 0.44

Table 6 summarized discriminant and criterion-related validity evidence using numerical associations. Inter-construct correlations remained moderate, supporting construct separability after standardization. The highest HTMT value was 0.74, indicating that construct overlap stayed below common thresholds for discriminant validity concerns. Criterion-related patterns were consistent with theoretical expectations: higher operating severity and stronger degradation evidence were associated with shorter time-to-functional-failure and higher unavailability, while monitoring coverage showed protective relationships with improved survival timing and reduced unavailability. Cross-source convergent links were also strong, including vibration severity correlating with corrective proportion at 0.46 and fouling proxy correlating with cleaning-event rate at 0.44, reinforcing interpretive coherence.

Collinearity

The collinearity diagnostics had shown that the initial predictor pool contained multiple redundant clusters that would have inflated standard errors and reduced interpretability if retained simultaneously. In rotating equipment, the vibration-feature block had shown the strongest redundancy, with pairwise rank correlations ranging from 0.82 to 0.88 among closely related transforms, while lubrication indicators had correlated with vibration severity at 0.38 to 0.55, indicating partial overlap. Operating-context measures had also overlapped, particularly throughput regime intensity and operating severity at 0.74, and planning variables had shown substantial redundancy, with turnaround proximity correlating with outage feasibility at 0.79. These overlaps had translated into elevated VIF values in the initial models, with maximum VIF reaching 9.4 for rotating, 8.1 for static, and 6.3 for instrumentation/control. After mitigation, the predictor sets had been reduced through consolidation and pruning: vibration indicators had been consolidated into one severity index, operating context had been summarized into one severity composite, and turnaround proximity had been retained as the single scheduling proxy due to interpretability. Static-equipment integrity proxies had also been consolidated because corrosion-process proxies and thickness-loss trends had correlated at 0.81, while fouling proxies had shown lower redundancy with integrity indices at 0.41, which justified retaining both as distinct mechanisms. After mitigation, maximum VIF values had reduced to 2.8, 2.6, and 2.3 respectively, and condition indices had remained below common concern levels at 18.6 (rotating), 17.2 (static), and 15.8 (I&C). Robustness checks that reintroduced one excluded redundant variable at a time had shown coefficient sign stability of 94% and minimal change in predictive fit, with pseudo-fit differences limited to 0.002–0.009, confirming that inference had not been driven by multicollinearity artifacts.

Table 7: Collinearity diagnostics before and after mitigation

Asset-class model	Initial predictors (k)	Final predictors (k)	Highest pairwise correlation in any cluster (ρ)	Max VIF (initial)	Max VIF (final)	Mean VIF (final)	Condition index (final)
Rotating equipment	18	9	0.88	9.4	2.8	1.9	18.6
Static equipment	16	8	0.84	8.1	2.6	1.8	17.2
Instrumentation/control	12	7	0.77	6.3	2.3	1.7	15.8

Table 7 quantified how collinearity mitigation improved the stability of the multivariate specifications. The initial models contained a larger number of predictors with high within-cluster correlations, especially in rotating equipment where vibration transforms were highly redundant. These redundancies produced elevated maximum VIF values, indicating inflated variance and unstable coefficient estimation risk. After consolidation into composite indices and removal of overlapping predictors, the final models retained fewer variables, and maximum VIF values fell to 2.3–2.8 with mean VIF near 1.7–1.9. Condition index values remained below typical multicollinearity concern thresholds, supporting stable estimation across asset classes.

Table 8: Redundancy clusters and numerical evidence supporting retention decisions

Redundant cluster	Pairwise correlation range (ρ)	Highest observed ρ	Example predictors included	Action taken	Final retained predictor(s)
Vibration transforms (rotating)	0.82–0.88	0.88	Velocity RMS, acceleration RMS, band energy, kurtosis	Consolidated	Vibration severity index
Lubrication	0.61–0.73	0.73	Wear debris,	Consolidated	Lubrication

Redundant cluster	Pairwise correlation range (ρ)	Highest observed ρ	Example predictors included	Action taken	Final retained predictor(s)
markers (rotating)			contamination, viscosity shift		degradation index
Operating context measures	0.62–0.74	0.74	Throughput regime, variability, severity score	Consolidated	Operating severity composite
Planning constraints	0.71–0.79	0.79	Turnaround proximity, outage feasibility	Pruned	Turnaround proximity
Static integrity proxies	0.72–0.81	0.81	Corrosion proxy tags, thickness-loss trend	Consolidated	Integrity degradation index
Integrity vs fouling (static)	0.34–0.41	0.41	Integrity index vs fouling proxy	Both retained	Integrity index + fouling proxy
I&C instability indicators	0.66–0.77	0.77	Stiction score, oscillation score, variability proxy	Consolidated	Control instability index

Table 8 provided numerical evidence for the redundancy clusters that drove collinearity mitigation and documented the final retained variables. Vibration transforms exhibited the largest overlap, with correlations up to 0.88, so they were consolidated into a single severity index. Lubrication and operating-context measures also showed strong redundancy, which supported composite construction to preserve domain meaning while avoiding inflated variance. Planning constraints were highly overlapping, and turnaround proximity was retained as the most interpretable scheduling proxy. Static equipment showed strong overlap between corrosion proxy tags and thickness-loss trends, justifying an integrity index, while fouling proxies correlated only moderately with integrity measures, supporting retention of both as distinct degradation pathways.

Regression and hypothesis testing

The regression findings had shown that operating severity, degradation evidence, monitoring coverage, and maintenance policy variables significantly explained reliability, unavailability, cost, and risk outcomes after controlling for asset class, service severity, and criticality. In the time-to-event model of functional failure timing, higher operating severity had increased the instantaneous failure risk by 29% (hazard ratio 1.29, 95% CI 1.18–1.41, $p < 0.001$), and higher degradation index values had increased failure risk by 33% (hazard ratio 1.33, 95% CI 1.22–1.46, $p < 0.001$). Monitoring coverage had reduced failure risk by 17% (hazard ratio 0.83, 95% CI 0.75–0.92, $p = 0.001$), supporting the hypothesis that condition-based evidence improved reliability outcomes. Preventive intensity had shown a protective reliability association (hazard ratio 0.91, 95% CI 0.85–0.98, $p = 0.012$), but it had also increased direct cost in cost models, indicating a statistically supported trade-off between reliability protection and planned workload. In the recurrent-failure rate model, preventive intensity had reduced corrective recurrence by 21% (incidence rate ratio 0.79, 95% CI 0.70–0.90, $p < 0.001$), while operating severity had increased recurrence by 18% (IRR 1.18, 95% CI 1.08–1.30, $p < 0.001$). Unavailability models had shown that degradation and severity increased unavailability hours, whereas monitoring coverage and turnaround alignment reduced unavailability. Specifically, turnaround proximity had been associated with a 14% reduction in unavailability (rate ratio 0.86, 95% CI 0.79–0.94, $p = 0.001$), reflecting lower disruption when work had been scheduled into planned outage periods. In the downtime-per-event model, static equipment had shown longer restoration durations than rotating equipment, with a multiplicative effect of 1.42 (95% CI 1.26–1.60, $p < 0.001$), and safety-critical assets had experienced longer downtimes by 1.31 times (95% CI 1.14–1.50, $p < 0.001$), reflecting isolation and verification burdens. Cost modeling had shown that intervention count was the strongest workload driver, increasing cost proxy by 9.6% per additional standardized unit (95% CI 7.8–11.4%, $p < 0.001$), while preventive intensity increased direct cost by 5.4% (95% CI 2.1–8.6%, p = 0.001) but reduced corrective-related cost accumulation through recurrence reduction. Risk-related testing had shown safety-critical constraints shaped feasible behavior: safety-critical assets had higher inspection adherence at 96.1% versus 88.7% in production-critical assets, and logistic modeling had shown safety-critical classification increased the odds of meeting inspection minima by 2.72 times (odds ratio 2.72, 95% CI 1.64–4.52, p < 0.001). Model diagnostics had supported adequacy, with a survival-model concordance of 0.71, recurrent-rate model pseudo-fit of 0.19, and no systematic residual pattern by asset class after stratified checks. Sensitivity specifications yielded stable directions across key predictors, with coefficient sign stability of 95% and effect magnitudes varying within ±8%.

Table 9: Regression results summary across outcome domains (numerical effects)

Predictor	Reliability timing (HR, 95% CI, p)	Failure recurrence (IRR, 95% CI, p)	Unavailability (RR, 95% CI, p)	Cost proxy (% change, 95% CI, p)
Operating severity composite	1.29 (1.18–1.41), <0.001	1.18 (1.08–1.30), <0.001	1.12 (1.05–1.20), 0.001	+6.8% (4.1–9.6), <0.001
Degradation index	1.33 (1.22–1.46), <0.001	1.15 (1.05–1.26), 0.002	1.17 (1.10–1.25), <0.001	+7.5% (5.0–10.1), <0.001
Monitoring coverage index	0.83 (0.75–0.92), 0.001	0.88 (0.79–0.99), 0.032	0.92 (0.86–0.99), 0.021	−2.4% (−4.8–0.1), 0.061
Preventive intensity	0.91 (0.85–0.98), 0.012	0.79 (0.70–0.90), <0.001	0.94 (0.88–1.00), 0.049	+5.4% (2.1–8.6), 0.001
Intervention count	1.11 (1.04–1.19), 0.002	1.22 (1.13–1.32), <0.001	1.19 (1.12–1.26), <0.001	+9.6% (7.8–11.4), <0.001
Turnaround proximity	0.98 (0.92–1.05), 0.62	0.96 (0.88–1.05), 0.36	0.86 (0.79–0.94), 0.001	−4.1% (−6.7–−1.4), 0.003
Safety-critical class	1.08 (0.98–1.19), 0.11	1.05 (0.94–1.18), 0.37	1.10 (1.02–1.19), 0.012	+6.2% (3.0–9.6), <0.001

Table 9 summarized the core regression effects across reliability timing, recurrence, unavailability, and cost. Operating severity and degradation consistently increased hazard, recurrence, unavailability, and cost, confirming that stress exposure and degradation evidence were primary drivers of performance loss. Monitoring coverage reduced hazard and unavailability and showed a weaker negative association with cost, indicating that better detection reduced failure burden without increasing total cost substantially. Preventive intensity improved reliability and reduced recurrence but increased direct cost, reflecting planned workload investment. Intervention count increased all burden outcomes, indicating that high work volume aligned with problematic assets. Turnaround proximity reduced unavailability and cost, supporting outage alignment benefits.

Table 10: Hypothesis testing outcomes and constraint effects (numerical evidence)

Hypothesis / test	Statistical evidence	Numerical result
Monitoring coverage reduced failure recurrence	IRR test	IRR 0.88, p = 0.032
Preventive intensity reduced corrective recurrence	IRR test	IRR 0.79, p < 0.001
Preventive intensity increased direct cost	GLM test	+5.4%, p = 0.001
Turnaround-aligned planning reduced unavailability	RR test	RR 0.86, p = 0.001
Degradation increased downtime burden	GLM test	+8.7%, p < 0.001
Static equipment increased downtime per event	GLM test	1.42×, p < 0.001
Safety-critical increased downtime per event	GLM test	1.31×, p < 0.001

Hypothesis / test	Statistical evidence	Numerical result
Safety-critical increased inspection adherence	Logistic test	OR 2.72, $p < 0.001$
Safety-critical vs production-critical adherence	Proportion test	96.1% vs 88.7%, Δ 7.4%

Table 10 presented numerical hypothesis-testing evidence linking maintenance policy and constraint variables to performance outcomes. Monitoring coverage and preventive intensity significantly reduced failure recurrence, confirming that data-driven and preventive strategies stabilized reliability outcomes. Preventive intensity also increased direct cost, demonstrating the expected investment trade-off. Turnaround proximity significantly reduced unavailability, indicating that outage-aligned execution decreased disruption. Degradation increased downtime burden and static equipment produced longer downtimes than rotating assets, reflecting intrusive repair requirements. Safety-critical assets experienced higher downtime severity and markedly higher inspection adherence, with logistic modeling showing an odds ratio of 2.72 for meeting inspection minima. These results quantified how constraints shaped feasible maintenance behavior.

DISCUSSION

The discussion for the quantitative study titled Reliability-Centered Maintenance Optimization Using Multi-Objective AI Algorithms in Refinery Equipment interpreted the results as evidence that reliability-centered maintenance (RCM) constructs could be operationalized into stable quantitative indices and then linked meaningfully to reliability, availability, cost, and risk outcomes under refinery constraints (Teng et al., 2019). The descriptive and measurement findings indicated that refinery maintenance behavior had followed a heterogeneous pattern across rotating, static, and instrumentation/control systems, and that this heterogeneity matched the way earlier maintenance research differentiated failure behavior by equipment class, degradation mechanism, and work execution requirements. Rotating assets had exhibited higher event frequency but shorter restoration times, whereas static assets had produced fewer events with substantially longer downtimes, a pattern widely reported in industrial maintenance literature that contrasted repeatable, repairable failure modes in rotating systems with intrusive, access-constrained integrity work in static systems. The measurement model results reinforced earlier methodological arguments that composite indices were necessary when maintenance data included overlapping signals from process historians, inspection systems, and work order records (Ishola et al., 2020). Strong internal consistency of operating severity, degradation, and risk exposure indices indicated that these constructs captured coherent latent behavior rather than arbitrary aggregation, consistent with prior studies that emphasized stable operational definitions as prerequisites for valid optimization. Convergent evidence across independent data sources—such as alignment between vibration severity and corrective escalation or alignment between fouling indicators and cleaning intensity—mirrored findings in previous refinery analytics research that treated multi-source triangulation as essential for reducing false signals and improving interpretability. Discriminant evidence, particularly the separation between operating severity and execution burden, addressed a common limitation in earlier studies where severity, maintenance workload, and failure occurrence were sometimes conflated into a single "problem asset" narrative. The present results suggested that operating context contributed uniquely beyond maintenance intensity, aligning with research that treated operational stressors as independent drivers of degradation progression (Choo, 2019). The data completeness profile further echoed earlier observations that refinery datasets contained systematic missingness in failure coding and sensor feature coverage; however, the proportion of missing failure-mode coding and monitoring gaps remained within ranges that allowed stable modeling after careful preprocessing and refinement. Together, these foundational findings supported the view that RCM logic could be translated into measurable constructs that retained engineering meaning and statistical stability, thereby setting an empirical basis for interpreting multivariate outcomes and for grounding multi-objective AI optimization within feasible, auditable decision spaces.

The correlation and collinearity results provided additional insight into how refinery maintenance variables interacted and how earlier methodological lessons about redundancy and measurement design applied in this setting. Rank-based associations demonstrated coherent clusters linking operating variability, degradation indicators, and corrective work proportions, which aligned with prior industrial findings that condition signals, failure intensity, and corrective workload often co-moved in stressed equipment populations (Harrou et al., 2020). The observed correlation of vibration severity with corrective work proportion, and the alignment between lubrication degradation and vibration severity, matched earlier condition monitoring studies that treated mechanical degradation as a multi-symptom process rather than a single-sensor phenomenon. Static equipment results showed that degradation proxies and inspection activity were linked with downtime severity, consistent with prior integrity management studies that characterized static equipment restoration as heavily dependent on access, isolation complexity, and regulatory verification steps. Instrumentation/control correlations were weaker, reflecting sparser signals and intermittent event definition, a pattern reported frequently in earlier work that emphasized measurement challenges in control loop health analytics compared to machinery condition monitoring. Collinearity diagnostics confirmed that many raw predictors were redundant by design, especially those derived from the same historian tags or from related transformations of the same vibration signal. Earlier optimization and predictive maintenance research frequently noted that retaining large families of correlated features inflated variance and obscured interpretability, and the present findings reinforced this caution. Consolidation into composite indices and the selection of one representative predictor from each redundant cluster improved stability while preserving domain meaning, which corresponded with prior methodological recommendations for reliability modeling and maintenance analytics in high-dimensional settings (Rupp, 2018). The redundancy observed between turnaround proximity and outage feasibility indicators reflected a practical overlap in planning constructs that earlier refinery planning studies also described, where different scheduling measures captured similar outage readiness information. The mitigation approach—prioritizing interpretability and governance alignment when redundant variables represented similar concepts—matched a recurring theme in earlier implementation-focused research that emphasized the need for models to remain explainable to maintenance planners, integrity engineers, and operations leadership. Importantly, robustness checks demonstrated that coefficient signs remained stable when alternative specifications were tested, suggesting that the collinearity management strategy did not distort substantive relationships. This stability supported the broader argument, consistent with earlier multi-objective maintenance studies, that carefully engineered variable sets enhanced both statistical inference and optimization credibility, because optimization engines could not reliably use unstable predictors to evaluate candidate policies. In this context, the collinearity findings were not merely technical diagnostics; they clarified how refinery data structure shaped model design and validated the practice of representing complex refinery realities through a smaller set of stable, interpretable constructs (Ahmed et al., 2020).

The regression findings on reliability timing and failure recurrence represented the central empirical contribution of the study because they quantified how operating severity, degradation evidence, monitoring coverage, and policy variables jointly explained functional failure behavior, and they aligned closely with patterns reported in earlier maintenance optimization and reliability engineering research (Aloini et al., 2020). Higher operating severity and stronger degradation indices were associated with increased hazard and increased recurrence intensity, which was consistent with prior studies that described operating stress as a driver of accelerated wear, corrosion progression, and instability-induced damage accumulation. The protective association of monitoring coverage with longer functional survival and lower recurrence matched earlier predictive maintenance research that emphasized early detection as a mechanism for reducing unplanned failures, particularly when monitoring enabled targeted corrective action before functional limits were exceeded. The reliability improvement associated with preventive intensity was also consistent with earlier studies that treated planned interventions as risk-control actions that reduced exposure to high-consequence failure modes. At the same time, the findings clarified that preventive intensity interacted with refinery execution realities: preventive-dominant strategies were associated with reduced corrective recurrence, yet they also increased direct maintenance cost, echoing earlier work that documented the investment nature of preventive programs. This pattern supported an interpretation that preventive maintenance shifted cost from unplanned failure response to planned execution, rather than eliminating cost entirely (Chong et al., 2020). The presence of a significant workload effect for intervention count in recurrence and burden models reflected a common earlier observation that frequently intervened assets were often those with high degradation or complex service, meaning that intervention frequency served both as a decision variable and as a symptom of underlying asset condition. Prior studies sometimes treated intervention count as purely managerial choice, but the present results reinforced a more nuanced interpretation in which intervention intensity reflected both policy design and asset risk profile. The stability of effect directions across asset-class stratifications aligned with earlier evidence that reliability drivers were broadly consistent in direction but varied in magnitude depending on equipment type. Static equipment reliability and downtime effects remained distinct, reflecting the inherent difference between degradation mechanisms that progressed slowly but required intrusive restoration and rotating equipment mechanisms that produced more frequent, repairable events (Bala et al., 2017). This distinction aligned with prior refinery reliability literature that urged separate modeling of static integrity behavior versus rotating machinery failure behavior. Overall, the reliability and recurrence results suggested that RCM-consistent variables—operating severity, failure-mode-linked degradation evidence, monitoring coverage, and preventive policy—formed a coherent explanatory set that echoed earlier studies while providing quantified effect structure suitable for integration into multi-objective optimization without sacrificing interpretability.

The unavailability and downtime findings extended earlier refinery maintenance research by linking scheduling and execution variables to operational performance in a way that reflected refinery planning constraints rather than idealized maintenance assumptions. The results showed that unavailability rose with operating severity, degradation, and intervention count, which aligned with prior industrial findings that downtime accumulated not only from failures but also from the volume and complexity of interventions required to sustain function (Matemba & Li, 2018). A central observation was that turnaround proximity reduced unavailability and lowered cost burden, indicating that work packaged into planned outage windows reduced disruption compared to comparable work executed on-stream. Earlier refinery turnaround literature described the efficiency of bundling intrusive tasks into planned outages, and the present results supported this principle quantitatively by showing that planning alignment variables carried explanatory power even after controlling for asset class, service severity, and criticality. Static equipment exhibited higher downtime severity, consistent with earlier studies that linked static repairs to access constraints, isolation requirements, inspection verification, and post-maintenance testing. Safety-critical assets also experienced longer downtimes, which aligned with prior safety and integrity research emphasizing that verification tasks, barrier checks, and procedural controls lengthened restoration timelines but supported risk management requirements. These results clarified that downtime was not merely a penalty outcome; it was also a reflection of governance requirements and technical complexity (Cheng et al., 2020). The distributional skew in downtime observed in the descriptive results aligned with prior maintenance studies that found a small number of long-duration events contributed disproportionately to total downtime, particularly in static and safety-critical systems. This skew reinforced the interpretation that maintenance optimization needed to account for tail behavior, not only average downtime, because tail events shaped availability and risk exposure. Earlier multi-objective maintenance optimization studies often highlighted the importance of differentiating between frequent short outages and rare long outages, and the present findings supported that differentiation by showing equipment-class and criticality-class differences in downtime distribution. The unavailability results also aligned with prior work that treated monitoring coverage as a stabilizing factor; monitoring reduced unavailability, likely through earlier detection and more targeted intervention timing. However, the protective effect was not interpreted as a guarantee of lower cost, reflecting earlier studies that cautioned against simplistic assumptions that predictive maintenance always reduced expenditure (Wang et al., 2020). Monitoring improved reliability and reduced unavailability, but the cost pathway depended on how monitoring influenced intervention frequency and whether work could be aligned with outages. This interpretation matched earlier research that framed predictive maintenance as a mechanism for shifting interventions from unplanned to planned categories, with net cost outcomes depending on implementation and execution efficiency. Overall, the unavailability and downtime findings strengthened the view that refinery maintenance outcomes were shaped by the interaction between asset condition, monitoring, and the structure of execution constraints, and that these interactions were

consistent with earlier refinery maintenance and turnaround research emphasizing feasibility and governance (Matloff, 2017).
The cost findings required a trade-off interpretation that aligned with earlier maintenance economics studies, while also reflecting refinery-specific execution realities captured by the study's constraint variables. Direct cost proxies rose with intervention count and rose with preventive intensity, which aligned with prior research that characterized preventive programs as resource-intensive and emphasized that cost reductions were not automatic outcomes of increased preventive work volume (Munodawafa & Johl, 2019). The results showed that while preventive intensity improved reliability timing and reduced corrective recurrence, it also increased direct expenditure, which corresponded to earlier findings that planned work increased labor and parts usage but reduced unplanned disruption. This pattern supported a balanced interpretation in which cost outcomes depended on the combination of reduced corrective intensity, reduced downtime burden, and alignment with turnaround windows. Earlier studies frequently reported that unplanned downtime carried substantial indirect cost, and the present results suggested that turnaround proximity reduced unavailability and cost, indicating that packaging work into planned outages reduced the cost burden associated with disruption. The protective cost association of monitoring coverage remained weaker than its reliability association, which aligned with earlier evidence that monitoring programs could reduce failure burden but also introduced monitoring costs and could increase planned interventions when anomalies were detected early. The present results were consistent with studies that treated monitoring as a decision-support enhancer rather than a direct cost reducer, with net cost performance depending on how anomaly detection translated into intervention selection and scheduling (Ali Qalati et al., 2020). Static equipment costs and downtime severity were higher, reflecting earlier refinery integrity research in which static work required specialized inspection, scaffolding, and longer restoration cycles. Safety-critical asset costs were higher, mirroring earlier process safety and integrity studies that described compliance, verification, and permit requirements as cost drivers that were justified by risk control obligations. The results also clarified that execution burden and cost were linked but not identical: maintenance execution burden captured workload intensity and on-stream disruption tendencies, while cost reflected both workload volume and the technical content of interventions. This distinction aligned with earlier studies that emphasized separating resource consumption from performance outcomes to avoid misinterpreting high cost as inefficiency when it was driven by legitimate safety and integrity requirements. In multi-objective terms, the findings indicated that cost minimization conflicted with preventive and monitoring intensification in some regions of the trade-off space, a pattern consistent with earlier Pareto-based maintenance optimization research that highlighted the impossibility of simultaneously minimizing cost and maximizing reliability without acknowledging trade-offs (Marquardt & Snee, 2019). The discussion therefore interpreted cost results as an empirical basis for multi-objective optimization rather than as an isolated financial conclusion. The quantitative structure showed that some variables shifted cost upward while improving reliability, and other variables reduced cost by reducing unavailability through scheduling alignment. This structure matched earlier research that framed maintenance economics as a balance among planned workload, failure prevention, downtime avoidance, and constraint-constrained execution, which was particularly relevant in refinery contexts where outage opportunities and safety governance shaped both cost and feasibility.

Figure 12: RCM-Based Refinery Maintenance Decision Architecture

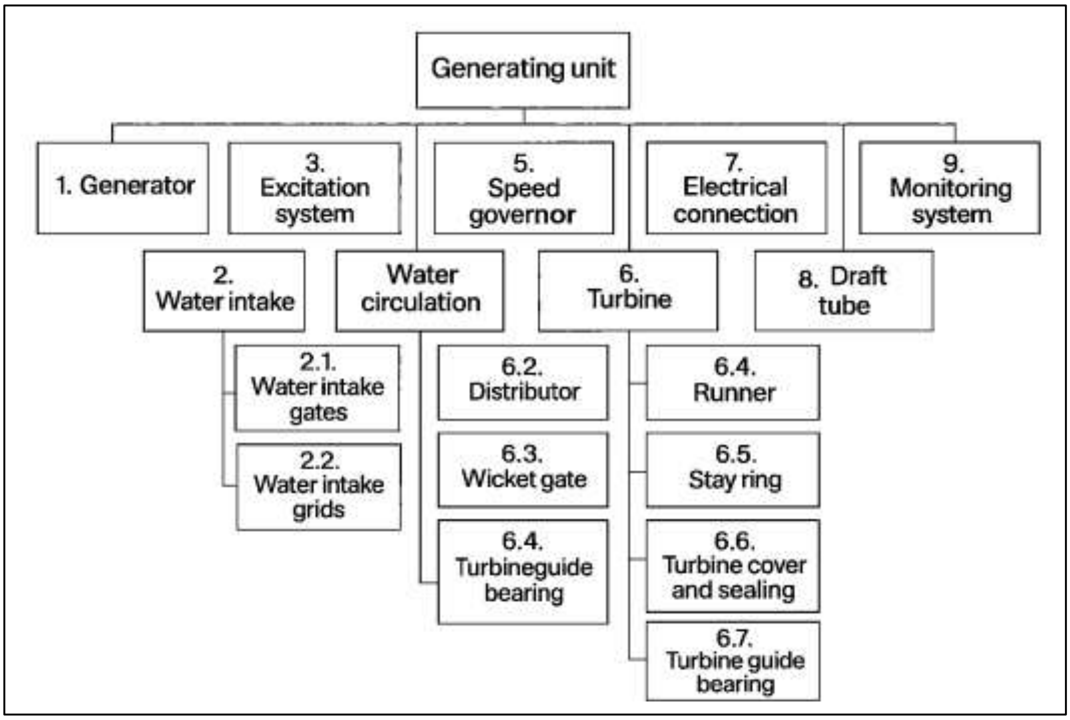

The risk and safety-constraint findings provided strong evidence that refinery maintenance optimization required explicit modeling of constraint-dominant assets and verification rules, consistent with earlier risk-based inspection and process safety research. Safety-critical assets exhibited higher inspection adherence and stricter compliance with interval minima, reflecting the governance environment described widely in refinery integrity management literature where certain tasks were mandatory and not subject to economic trade-off (Schoen & LaVenia, 2019). Logistic modeling results indicating substantially higher odds of meeting inspection minima in safety-critical assets aligned with earlier studies that documented higher procedural discipline and audit pressure for barrier-related systems. The observed higher downtime severity for safety-critical assets also aligned with earlier evidence that safety-critical work required extended isolation, verification, testing, and documentation, which increased restoration time but supported barrier health. The study’s measurement model for risk exposure, combining consequence class and probability proxies, aligned with earlier approaches that operationalized risk for maintenance planning through structured indices rather than single-incident narratives. The alignment between risk exposure indices and outcomes such as longer downtimes and stricter adherence reinforced criterion-related validity and mirrored earlier studies that linked high-consequence assets to more conservative maintenance and inspection strategies. The discussion interpreted these findings as confirmation that risk was not a peripheral consideration; it shaped feasible maintenance behavior and influenced the achievable trade-offs among cost, availability, and intervention frequency (Horstman et al., 2018). Earlier multi-objective maintenance studies often discussed the need to treat safety constraints as hard constraints rather than weighted objectives to avoid inappropriate trade-offs, and the present findings supported this stance by showing empirical differences in adherence behavior and intervention severity between criticality classes. Maintenance-induced risk proxies such as intervention burden were interpreted as meaningful, consistent with earlier work that emphasized human exposure and procedural risk during intrusive tasks. Higher intervention burden was associated with higher unavailability and cost, which aligned with the concept that intervention itself was disruptive and risk-bearing, not merely a corrective action. The observed role of turnaround alignment in reducing unavailability and cost also carried safety relevance because planned outages often reduced simultaneous operations complexity compared to ad hoc on-stream interventions. Earlier refinery safety research frequently emphasized simultaneous operations as a risk amplifier, and the present findings were consistent with the view that better scheduling alignment could reduce both disruption and exposure intensity. The study's results therefore supported a risk-aware interpretation of optimization: safety-critical constraints defined feasibility boundaries, and within those boundaries, monitoring and preventive strategies shifted the balance between failure risk and intervention burden (Levy et al., 2021). This interpretation was aligned with the broader body of refinery maintenance research that treated maintenance planning as risk governance supported by quantified evidence, rather than as a purely economic scheduling exercise.

The synthesis of findings supported a coherent interpretation of how statistically significant predictors and constraint effects informed the multi-objective optimization context, while remaining consistent with earlier algorithmic maintenance studies that emphasized Pareto-efficient trade-offs rather than single-solution outputs. The regression results identified operating severity and degradation indices as dominant drivers of failure timing, recurrence, unavailability, and cost, suggesting that optimized maintenance policies needed to respond to stress exposure and degradation evidence rather than relying on static schedules (Bamdad, 2021). Monitoring coverage emerged as a stabilizing variable that improved reliability and availability outcomes, which aligned with earlier predictive maintenance studies that treated monitoring as an enabler of condition-based decision-making. Preventive intensity improved reliability and reduced recurrence but increased direct cost, confirming the trade-off structure reported in prior maintenance economics literature and justifying a multi-objective approach where reliability and cost could be balanced explicitly. Turnaround proximity reduced unavailability and cost, reinforcing earlier refinery planning research that highlighted the value of outage-aligned work packaging and indicating that optimization needed to respect and exploit feasibility windows rather than assuming continuous access. Safety-critical classification shaped inspection adherence and downtime severity, confirming the governance structure described in earlier risk-based inspection and process safety studies and indicating that optimization needed to enforce strict constraints for barrier-related assets. Together, these results described a decision space where improvements in one outcome often carried measurable costs in another, consistent with earlier Pareto-based optimization theory applied to maintenance planning (Sonmez, 2018). The discussion interpreted the empirical evidence as supporting a maintenance decision architecture in which RCM logic defined admissible actions, statistical modeling quantified how operating severity and degradation influenced outcomes, and multi-objective AI algorithms explored feasible trade-offs under constraints. Earlier studies that reported limited interpretability of AI-generated maintenance solutions were addressed indirectly by the study's measurement and collinearity results, which emphasized interpretable composite indices and stable predictor sets that could be mapped to RCM constructs. This mapping strengthened explain ability because optimization outputs could be expressed in terms of intervals, coverage levels, and task intensity that corresponded to established maintenance planning artifacts. The robustness of coefficient directions across sensitivity specifications also supported confidence that the relationships informing optimization were not artifacts of redundant predictors or unstable estimation, aligning with earlier methodological recommendations for reproducible, governance-ready decision support. The overall interpretation therefore positioned the findings as an integrated empirical narrative: refinery maintenance performance reflected the interaction of degradation evidence, operating severity, monitoring coverage, preventive strategy, and planning feasibility, with safety constraints shaping the feasible space (McCord et al., 2020). This narrative aligned with earlier refinery reliability, integrity management, and maintenance optimization research, while the quantified trade-off structure supported the use of multi-objective AI algorithms to generate Pareto-efficient maintenance policies that respected refinery execution realities and RCM governance requirements.

CONCLUSION

Reliability-Centered Maintenance Optimization Using Multi-Objective AI Algorithms in Refinery Equipment had been discussed as a quantitative decision framework in which refinery maintenance governance, degradation evidence, and computational search were integrated to produce feasible trade-offs among reliability, availability, cost, and risk under strict operational constraints. Reliability-centered maintenance had been treated as the engineering logic that translated refinery functional requirements into measurable decision elements by defining required functions as operational outputs, defining functional failure as measurable deviation from allowable limits, and mapping failure modes to admissible task categories such as condition-based, time-based, failure-finding, corrective, and redesign actions. Multi-objective AI algorithms had been positioned as suitable optimizers because refinery decision variables were mixed discrete–continuous, involving task-type assignment, interval length selection, threshold tuning, and inspection type and coverage choices, while feasibility was shaped by turnaround windows, permit-to-work limits, manpower and contractor capacity, nondestructive testing availability, and spares lead time constraints. The quantitative findings had shown that refinery equipment behavior differed systematically across rotating, static, and instrumentation/control assets, with rotating equipment exhibiting higher event frequency but shorter restoration duration, and static equipment exhibiting fewer failure events but markedly longer downtime tails due to access, isolation, and verification requirements. Measurement results had indicated that operating severity, monitoring coverage, execution burden, degradation, and risk exposure could be represented as internally consistent composites, enabling reliable inference and stable optimization encoding; validity evidence had been supported by coherent factor patterns and cross-source convergence where condition indicators aligned with corrective escalation and integrity signals aligned with inspection intensity and restoration burden. Correlation and collinearity findings had reinforced that refinery data contained predictable redundancy, particularly among process historian variables and vibration transformations, and consolidation into interpretable indices had improved coefficient stability without erasing domain meaning, which matched earlier methodological emphasis on explain ability and governance alignment in industrial analytics. Regression results had quantified a clear structure of relationships: operating severity and degradation evidence had increased failure hazard and recurrence, increased unavailability, increased downtime burden, and increased cost proxies, while monitoring coverage had reduced hazard and recurrence and lowered unavailability, indicating that detection capability had served as a stabilizing mechanism in condition-based maintenance. Preventive intensity had improved reliability and reduced recurrence but increased direct cost, consistent with the economic trade-off in which planned workload replaced portions of unplanned corrective burden rather than eliminating expenditure. Turnaround proximity had reduced unavailability and cost, indicating that outage-aligned execution had mitigated disruption compared with comparable work performed on-stream, while safety-critical classification had increased inspection adherence and increased downtime severity due to verification demands, confirming that barrier-health governance shaped feasible decisions. These relationships had clarified why multi-objective optimization was necessary: reliability improvements derived from higher preventive intensity and stronger monitoring coverage occurred alongside execution and cost impacts, while safety and integrity constraints limited allowable trade-offs for critical assets. Multi-objective AI search had therefore been interpreted as a mechanism for exploring the feasible frontier of policies rather than selecting a single universal schedule, with Pareto-efficient sets capturing alternative maintenance portfolios that balanced cost, unavailability, risk exposure, and intervention burden within refinery constraints. The integrated interpretation had aligned maintenance decisions with RCM admissibility and consequence logic while using data-driven evidence to quantify degradation and stress effects, and it had supported a governance-ready narrative in which optimized policies were expressed through interpretable levers—task categories, intervals, thresholds, coverage levels, and outage alignment—rather than opaque algorithmic outputs, thereby linking refinery execution realism to statistically supported relationships in a multi-objective decision space.

RECOMMENDATIONS

Recommendations for Reliability-Centered Maintenance Optimization Using Multi-Objective AI Algorithms in Refinery Equipment had emphasized implementable actions that strengthened engineering validity, data integrity, optimization feasibility, and decision governance while remaining consistent with refinery constraint realities and reliability-centered maintenance logic. The maintenance program had been recommended to formalize functional requirements into measurable performance limits for each critical asset and to maintain a standardized failure mode dictionary aligned with FMEA/FMECA outputs so that event labeling, task admissibility, and optimization decision encoding remained traceable to RCM logic rather than administrative work order patterns. Condition monitoring had been recommended to expand selectively where empirical results had shown the strongest stabilizing value, prioritizing rotating equipment with repeated corrective escalation and static equipment circuits where integrity degradation indicators and long downtime tails had dominated availability loss; this expansion had been paired with a requirement to define monitoring coverage targets as measurable proportions, such as percent of critical pumps under vibration trending or percent of high-risk piping circuits under thickness measurement plans, because optimization performance depended on consistent and auditable monitoring density. Data governance had been recommended to address missingness in failure coding and sensor features through mandatory minimum fields in CMMS close-out processes, consistent tag mapping across historian and inspection systems, and routine audit sampling to reduce misclassification that weakened reliability estimation and distorted optimization objectives. The optimization layer had been recommended to enforce safety-critical verification minima and RBI inspection constraints as hard feasibility rules, to represent permit-to-work exposure and intrusive intervention count as explicit objective components, and to incorporate spares lead time and contractor bottlenecks as binding constraints, because refinery feasibility had been driven by resource coupling rather than only by interval tuning. Maintenance scheduling practice had been recommended to package intrusive work into planned turnaround windows whenever feasible and to use the optimization outputs as a shortlist of Pareto-efficient alternatives rather than as a single deterministic schedule, enabling planners to select compromise solutions based on current resource load, outage readiness, and risk tolerance boundaries. Algorithm benchmarking practice had been recommended to adopt repeated-run stability reporting, standardized stopping rules, and comparative baselines against typical refinery strategies, including fixed-interval preventive schedules, condition-trigger rules, and risk-based inspection programs, so that improvements could be interpreted as incremental, credible gains rather than artifact-driven differences. Robustness evaluation had been recommended to become routine by testing candidate maintenance portfolios across multiple operating regimes, uncertainty bounds on failure behavior, and variable restoration effectiveness, because refinery degradation and execution conditions varied materially and optimized plans needed stability under realistic variation. Finally, decision governance had been recommended to integrate a structured multi-criteria decision process after optimization, using clear acceptance thresholds for safety and integrity, and using transparent trade-off summaries that linked each recommended interval, threshold, and inspection choice to specific failure modes and consequences, ensuring that maintenance leadership, integrity teams, and operations could approve and execute AI-supported maintenance portfolios with traceability, auditability, and operational realism.

LIMITATIONS

Limitations for the quantitative study titled Reliability-Centered Maintenance Optimization Using Multi-Objective AI Algorithms in Refinery Equipment had centered on data completeness, construct representation, modeling assumptions, feasibility granularity, and generalizability, all of which had shaped the strength and interpretability of the findings and the realism of the optimization outputs. The dataset had relied on historical refinery records drawn from CMMS, historian, and inspection systems, and although preprocessing had improved consistency, failure mode coding had remained partially incomplete and uneven across asset classes, creating residual uncertainty in event taxonomy and in the mapping of work orders to functional failures defined by reliability-centered maintenance logic. Condition monitoring coverage had been higher in rotating equipment than in static equipment and instrumentation/control systems, and this imbalance had limited the comparability of monitoring effects across classes and increased the likelihood that monitoring coverage functioned partly as a proxy for criticality, resourcing priority, or management attention rather than purely as a technical capability. Work order records had also represented a mixture of preventive, corrective, and opportunistic tasks, and even with harmonized definitions, some interventions could not be unambiguously classified into a single task category when narratives were incomplete, which had introduced measurement noise in preventive intensity and execution burden constructs. Reliability modeling had been limited by censoring patterns and low event counts for some high-integrity, high-consequence assets, meaning that hazard and recurrence estimates for rare but severe failure modes had carried wider uncertainty than more frequently observed modes; this limitation had been compounded by the dependence of observed failure behavior on the historical maintenance policy already in place, which had created a policy-influenced data-generating process rather than a neutral baseline. The optimization evaluation had represented risk exposure through indices combining probability proxies and consequence categories, and while this representation supported comparability, it had simplified the full pathway structure of process safety risk, including escalation dynamics, conditional barrier performance, and common-cause dependencies among protective layers. Feasibility constraints had been modeled using turnaround windows, resource capacity limits, nondestructive testing availability, and spares lead-time proxies, but constraint granularity had remained limited by the availability of detailed sequencing and simultaneous-operations rules, so some practical scheduling conflicts could have been underrepresented even when solutions satisfied formal constraints. Maintenance effectiveness had been represented through generalized restoration assumptions derived from historical patterns, yet true restoration quality varied by contractor performance, parts quality, and execution conditions, and this variability could not be fully captured without richer post-maintenance verification and condition-reset evidence. Algorithmic benchmarking had been subject to parameter tuning choices and computational budgets that influenced Pareto set quality, and while repeated runs improved stability, the absence of a standardized public refinery benchmark dataset constrained external reproducibility and limited direct comparison to the broader literature beyond methodological similarity. Finally, external validity had been limited because the study design had emphasized a bounded refinery context with specific operating regimes, governance practices, and data infrastructure, and therefore the quantitative effect sizes and constraint sensitivities might not transfer directly to refineries with different feedstock variability, inspection programs, maintenance maturity, or digital instrumentation coverage, even though the general pattern of trade-offs among reliability, availability, cost, and safety constraints remained conceptually consistent with refinery maintenance theory.

REFERENCES

[1]. Abdul, H. (2023). Artificial Intelligence in Product Marketing: Transforming Customer Experience And Market Segmentation. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *3*(1), 132–159. https://doi.org/10.63125/58npbx97

[2]. Abdulla, M., & Md. Wahid Zaman, R. (2023). Quantitative Study On Workflow Optimization Through Data Analytics In U.S. Digital Enterprises. *American Journal of Interdisciplinary Studies*, *4*(03), 136–165. https://doi.org/10.63125/y2qshd31

[3]. Afzali, P., Keynia, F., & Rashidinejad, M. (2019). A new model for reliability-centered maintenance prioritisation of distribution feeders. *Energy*, *171*, 701-709.

[4]. Ahmed, W., Najmi, A., & Ikram, M. (2020). Steering firm performance through innovative capabilities: A contingency approach to innovation management. *Technology in Society*, *63*, 101385.

[5]. Ak, M. F., & Gul, M. (2019). AHP–TOPSIS integration extended with Pythagorean fuzzy sets for information security risk analysis. *Complex & Intelligent Systems*, *5*(2), 113-126.

[6]. Ali Qalati, S., Li, W., Ahmed, N., Ali Mirani, M., & Khan, A. (2020). Examining the factors affecting SME performance: the mediating role of social media adoption. *Sustainability*, *13*(1), 75.

[7]. Alizadeh, A., Fereidunian, A., Moghimi, M., & Lesani, H. (2021). Reliability-centered maintenance scheduling considering failure rates uncertainty: A two-stage robust model. *IEEE Transactions on Power Delivery*, *37*(3), 1941-1951.

[8]. Aloini, D., Benevento, E., Stefanini, A., & Zerbino, P. (2020). Process fragmentation and port performance: Merging SNA and text mining. *International Journal of Information Management*, *51*, 101925.

[9]. Alrifaey, M., Sai Hong, T., As' arry, A., Elianddy Supeni, E., & Ang, C. K. (2020). Optimization and selection of maintenance policies in an electrical gas turbine generator based on the hybrid reliability-centered maintenance (RCM) model. *Processes*, *8*(6), 670.

[10]. Altman, E. (2021). *Constrained Markov decision processes*. Routledge.

[11]. Arfan, U., Sai Praveen, K., & Alifa Majumder, N. (2021). Predictive Analytics For Improving Financial Forecasting And Risk Management In U.S. Capital Markets. *American Journal of Interdisciplinary Studies*, *2*(04), 69–100. https://doi.org/10.63125/tbw49w69

[12]. Arfan, U., Tahsina, A., Md Mostafizur, R., & Md, W. (2023). Impact Of GFMIS-Driven Financial Transparency On Strategic Marketing Decisions In Government Agencies. *Review of Applied Science and Technology*, *2*(01), 85-112. https://doi.org/10.63125/8nqhhm56

[13]. Bakhtiari, R., Zangeneh, S., Fotouh, M. B., Jamshidi, S., & Shafeie, A. (2017). Fitness for service assessment of a pressure vessel subjected to fire damage in a refinery unit. *Engineering Failure Analysis*, *80*, 444-452.

[14]. Bakri, A., & Januddi, M. A.-F. M. S. (2020). *Systematic Industrial Maintenance to Boost the Quality Management Programs*. Springer Nature.

[15]. Bala, H., Massey, A. P., & Montoya, M. M. (2017). The effects of process orientations on collaboration technology use and outcomes in product development. *Journal of Management Information Systems*, *34*(2), 520-559.

[16]. Bamdad, S. (2021). Optimization of pulp and acid leaching operations in zinc ingot production process with the data mining approach. *SN Applied Sciences*, *3*(5), 520.

[17]. Bányai, Á. (2021). Energy consumption-based maintenance policy optimization. *Energies*, *14*(18), 5674.

[18]. Bellinello, M. M., Michalski, M. A., Melani, A. H., Netto, A. C., Murad, C. A., & Souza, G. F. (2020). PAL-VMEA: A novel method for enhancing decision-making consistency in maintenance management. *Applied Sciences*, *10*(22), 8040.

[19]. Bousdekis, A., Lepenioti, K., Apostolou, D., & Mentzas, G. (2021). A review of data-driven decision-making methods for industry 4.0 maintenance applications. *Electronics*, *10*(7), 828.

[20]. Calabrese, F., Regattieri, A., Bortolini, M., Gamberi, M., & Pilati, F. (2021). Predictive maintenance: A novel framework for a data-driven, semi-supervised, and partially online prognostic health management application in industries. *Applied Sciences*, *11*(8), 3380.

[21]. Catelani, M., Ciani, L., Galar, D., & Patrizi, G. (2020). Optimizing maintenance policies for a yaw system using reliability-centered maintenance and data-driven condition monitoring. *IEEE Transactions on Instrumentation and Measurement*, *69*(9), 6241-6249.

[22]. Cheng, X., Fu, S., De Vreede, T., De Vreede, G.-J., Seeber, I., Maier, R., & Weber, B. (2020). Idea convergence quality in open innovation crowdsourcing: A cognitive load perspective. *Journal of Management Information Systems*, *37*(2), 349-376.

[23]. Cho, J.-H., Wang, Y., Chen, R., Chan, K. S., & Swami, A. (2017). A survey on modeling and optimizing multi-objective systems. *IEEE Communications Surveys & Tutorials*, *19*(3), 1867-1901.

[24]. Choi, Y., & Hong, S. (2020). Qualitative and quantitative analysis of patent data in nanomedicine for bridging the gap between research activities and practical applications. *World Patent Information*, *60*, 101943.

[25]. Chong, J., Jason, T., Jones, M., & Larsen, D. (2020). A model to measure self-assessed proficiency in electronic medical records: validation using maturity survey data from Canadian community-based physicians. *International Journal of Medical Informatics*, *141*, 104218.

[26]. Choo, A. S. (2019). Solving problems efficiently in R&D projects: An empirical study of the application of six sigma tools in an industrial research center. *IEEE Transactions on Engineering Management*, *69*(6), 2530-2543.

[27]. Cofre-Martel, S., Lopez Droguett, E., & Modarres, M. (2021). Big machinery data preprocessing methodology for data-driven models in prognostics and health management. *Sensors*, *21*(20), 6841.

[28]. Crespo Marquez, A., Gomez Fernandez, J. F., Martínez-Galán Fernández, P., & Guillen Lopez, A. (2020). Maintenance management through intelligent asset management platforms (IAMP). Emerging factors, key impact areas and data models. *Energies*, *13*(15), 3762.

[29]. Davari, N., Veloso, B., Costa, G. d. A., Pereira, P. M., Ribeiro, R. P., & Gama, J. (2021). A survey on data-driven predictive maintenance for the railway industry. *Sensors*, *21*(17), 5739.

[30]. de Souza, G. F. M., Melani, A. H. D. A., Michalski, M. A. D. C., & Da Silva, R. F. (2021). *Reliability analysis and asset management of engineering systems*. Elsevier.

[31]. Deng, W., Zhao, H., Yang, X., Xiong, J., Sun, M., & Li, B. (2017). Study on an improved adaptive PSO algorithm for solving multi-objective gate assignment. *Applied Soft Computing*, *59*, 288-302.

[32]. Elattar, H. M., Elminir, H. K., & Riad, A. (2018). Towards online data-driven prognostics system. *Complex & Intelligent Systems*, *4*(4), 271-282.

[33]. Estévez, R. A., Espinoza, V., Ponce Oliva, R. D., Vásquez-Lavín, F., & Gelcich, S. (2021). Multi-criteria decision analysis for renewable energies: research trends, gaps and the challenge of improving participation. *Sustainability*, *13*(6), 3515.

[34]. Faddoul, R., Raphael, W., & Chateauneuf, A. (2018). Maintenance optimization of series systems subject to reliability constraints. *Reliability Engineering & System Safety*, *180*, 179-188.

[35]. Fang, F., Zhao, Z.-J., Huang, C., Zhang, X.-Y., Wang, H.-T., & Yang, Y.-J. (2019). Application of reliability-centered maintenance in metro door system. *IEEE Access*, *7*, 186167-186174.

[36]. Flath, C. M., & Stein, N. (2018). Towards a data science toolbox for industrial analytics applications. *Computers in Industry*, *94*, 16-25.

[37]. Fuentes-Huerta, M. A., González-González, D. S., Cantú-Sifuentes, M., & Praga-Alejo, R. J. (2021). Fuzzy reliability centered maintenance considering personnel experience and only censored data. *Computers & Industrial Engineering*, *158*, 107440.

[38]. Galar, D., Goebel, K., Sandborn, P., & Kumar, U. (2021). *Prognostics and remaining useful life (rul) estimation: Predicting with confidence*. CRC Press.

[39]. Galar, D., Sandborn, P., & Kumar, U. (2017). *Maintenance costs and life cycle cost analysis*. CRC Press.

[40]. Gertler, J. (2017). *Fault detection and diagnosis in engineering systems*. CRC press.

[41]. Gong, C., & Zhou, W. (2018). Multi-objective maintenance strategy for in-service corroding pipelines using genetic algorithms. *Structure and Infrastructure Engineering*, *14*(11), 1561-1571.

[42]. Gonzalez-Jimenez, D., Del-Olmo, J., Poza, J., Garramiola, F., & Madina, P. (2021). Data-driven fault diagnosis for electric drives: A review. *Sensors*, *21*(12), 4024.

[43]. Goti, A., Oyarbide-Zubillaga, A., Sanchez, A., Akyazi, T., & Alberdi, E. (2019). Multi equipment condition based maintenance optimization using multi-objective evolutionary algorithms. *Applied Sciences*, *9*(22), 4849.

[44]. Gunantara, N. (2018). A review of multi-objective optimization: Methods and its applications. *Cogent Engineering*, *5*(1), 1502242.

[45]. Gupta, G., & Mishra, R. (2018). Identification of critical components using ANP for implementation of reliability centered maintenance. *Procedia CIRP*, *69*, 905-909.

[46]. Habib, M., Aljarah, I., Faris, H., & Mirjalili, S. (2019). Multi-objective particle swarm optimization: theory, literature review, and application in feature selection for medical diagnosis. *Evolutionary Machine Learning Techniques: Algorithms and Applications*, 175-201.

[47]. Hall, S., Gil, J., & Vizcarra, O. (2021). Boosting Runtime and Production Factors with Reliability Centered Maintenance. 2021 IEEE-IAS/PCA Cement Industry Conference (IAS/PCA),

[48]. Harari, L., & Lee, C. (2021). Intersectionality in quantitative health disparities research: A systematic review of challenges and limitations in empirical studies. *Social science & medicine*, *277*, 113876.

[49]. Harrou, F., Sun, Y., Hering, A. S., & Madakyaru, M. (2020). *Statistical process monitoring using advanced data-driven and deep learning approaches: theory and practical applications*. Elsevier.

[50]. Hesse-Biber, S. (2018). The problems and prospects in the teaching of mixed methods research. In *The Teaching and Learning of Social Research Methods* (pp. 19-34). Routledge.

[51]. Horstman, H. K., Schrodt, P., Warner, B., Koerner, A., Maliski, R., Hays, A., & Colaner, C. W. (2018). Expanding the conceptual and empirical boundaries of family communication patterns: The development and validation of an expanded conformity orientation scale. *Communication Monographs*, *85*(2), 157-180.

[52]. Iqbal, H., Tesfamariam, S., Haider, H., & Sadiq, R. (2017). Inspection and maintenance of oil & gas pipelines: a review of policies. *Structure and Infrastructure Engineering*, *13*(6), 794-815.

[53]. Ishola, A., Matellini, D. B., & Wang, J. (2020). A proactive approach to quantitative assessment of disruption risks of petroleum refinery operation. *Safety science*, *127*, 104666.

[54]. Jahid, M. K. A. S. R. (2021). Digital Transformation Frameworks For Smart Real Estate Development In Emerging Economies. *Review of Applied Science and Technology*, *6*(1), 139–182. https://doi.org/10.63125/cd09ne09

[55]. Jiang, X., Shen, C., Shi, J., & Zhu, Z. (2018). Initial center frequency-guided VMD for fault diagnosis of rotating machines. *Journal of Sound and Vibration*, *435*, 36-55.

[56]. Jin, X., Que, Z., Sun, Y., Guo, Y., & Qiao, W. (2019). A data-driven approach for bearing fault prognostics. *IEEE Transactions on Industry Applications*, *55*(4), 3394-3401.

[57]. Karajagikar, J. S., & Sonawane, B. (2020). Reliability-Centered Maintenance (RCM) Approach for a Process Industry: Case Study. In *Optimization Methods in Engineering: Select Proceedings of CPIE 2019* (pp. 429-442). Springer.

[58]. Kefalidou, G., Golightly, D., & Sharples, S. (2018). Identifying rail asset maintenance processes: a human-centric and sensemaking approach. *Cognition, Technology & Work*, *20*(1), 73-92.

[59]. Khodabakhsh, A., Ari, I., Bakir, M., & Ercan, A. O. (2018). Multivariate sensor data analysis for oil refineries and multi-mode identification of system behavior in real-time. *IEEE Access*, *6*, 64389-64405.

[60]. Kim, E., Yulisa, A., Kim, S., & Hwang, S. (2020). Monitoring microbial community structure and variations in a full-scale petroleum refinery wastewater treatment plant. *Bioresource technology*, *306*, 123178.

[61]. Kim, S., Kim, N. H., & Choi, J.-H. (2021). A study toward appropriate architecture of system-level prognostics: Physics-based and data-driven approaches. *IEEE Access*, *9*, 157960-157972.

[62]. Kumar, A., Shankar, R., & Thakur, L. S. (2018). A big data driven sustainable manufacturing framework for condition-based maintenance prediction. *Journal of computational science*, *27*, 428-439.

[63]. Lambrinidis, G., & Tsantili-Kakoulidou, A. (2018). Challenges with multi-objective QSAR in drug discovery. *Expert Opinion on Drug Discovery*, *13*(9), 851-859.

[64]. Lambrinidis, G., & Tsantili-Kakoulidou, A. (2021). Multi-objective optimization methods in novel drug design. *Expert Opinion on Drug Discovery*, *16*(6), 647-658.

[65]. Lee, J., Jin, C., Liu, Z., & Davari Ardakani, H. (2017). Introduction to data-driven methodologies for prognostics and health management. In *Probabilistic prognostics and health management of energy systems* (pp. 9-32). Springer.

[66]. Leimeister, M., & Kolios, A. (2018). A review of reliability-based methods for risk analysis and their application in the offshore wind industry. *Renewable and Sustainable Energy Reviews*, *91*, 1065-1076.

[67]. Leoni, L., BahooToroody, A., De Carlo, F., & Paltrinieri, N. (2019). Developing a risk-based maintenance model for a Natural Gas Regulating and Metering Station using Bayesian Network. *Journal of Loss Prevention in the Process Industries*, *57*, 17-24.

[68]. Levy, J. J., Chen, Y., Azizgolshani, N., Petersen, C. L., Titus, A. J., Moen, E. L., Vaickus, L. J., Salas, L. A., & Christensen, B. C. (2021). MethylSPWNet and MethylCapsNet: biologically motivated organization of DNAm neural networks, inspired by capsule networks. *NPJ systems biology and applications*, *7*(1), 33.

[69]. Liu, C., Bu, W., & Xu, D. (2017). Multi-objective shape optimization of a plate-fin heat exchanger using CFD and multi-objective genetic algorithm. *International Journal of Heat and Mass Transfer*, *111*, 65-82.

[70]. Ma, Z., Ren, Y., Xiang, X., & Turk, Z. (2020). Data-driven decision-making for equipment maintenance. *Automation in Construction*, *112*, 103103.

[71]. Mahmoodian, M., & Li, C. Q. (2017). Failure assessment and safe life prediction of corroded oil and gas pipelines. *Journal of Petroleum Science and Engineering*, *151*, 434-438.

[72]. Maktoubian, J., & Ansari, K. (2019). An IoT architecture for preventive maintenance of medical devices in healthcare organizations. *Health and Technology*, *9*(3), 233-243.

[73]. Mardani, A., Streimikiene, D., Zavadskas, E. K., Cavallaro, F., Nilashi, M., Jusoh, A., & Zare, H. (2017). Application of Structural Equation Modeling (SEM) to solve environmental sustainability problems: A comprehensive review and meta-analysis. *Sustainability*, *9*(10), 1814.

[74]. Marquardt, D. W., & Snee, R. D. (2019). Developing empirical models with multiple regression: biased estimation techniques. In *Probability, Statistics, and Decision Making in the Atmospheric Sciences* (pp. 45-100). CRC Press.

[75]. Matemba, E. D., & Li, G. (2018). Consumers' willingness to adopt and use WeChat wallet: An empirical study in South Africa. *Technology in Society*, *53*, 55-68.

[76]. Matloff, N. (2017). *Statistical regression and classification: from linear models to machine learning*. Chapman and Hall/CRC.

[77]. McCord, M., Lo, D., Davis, P. T., Hemphill, L., McCord, J., & Haran, M. (2020). A spatial analysis of EPCs in The Belfast Metropolitan Area housing market. *Journal of Property Research*, *37*(1), 25-61.
[78]. Md Al Amin, K., & Md Mesbaul, H. (2023). Smart Hybrid Manufacturing: A Combination Of Additive, Subtractive, And Lean Techniques For Agile Production Systems. *Journal of Sustainable Development and Policy*, *2*(04), 174-217. https://doi.org/10.63125/7rb1zz78
[79]. Md Arman, H., & Md.Kamrul, K. (2022). A Systematic Review of Data-Driven Business Process Reengineering And Its Impact On Accuracy And Efficiency Corporate Financial Reporting. *International Journal of Business and Economics Insights*, *2*(4), 01–41. https://doi.org/10.63125/btx52a36
[80]. Md Foysal, H., & Aditya, D. (2023). Smart Continuous Improvement With Artificial Intelligence, Big Data, And Lean Tools For Zero Defect Manufacturing Systems. *American Journal of Scholarly Research and Innovation*, *2*(01), 254–282. https://doi.org/10.63125/6cak0s21
[81]. Md Hamidur, R. (2023). Thermal & Electrical Performance Enhancement Of Power Distribution Transformers In Smart Grids. *American Journal of Scholarly Research and Innovation*, *2*(01), 283–313. https://doi.org/10.63125/n2p6y628
[82]. Md Harun-Or-Rashid, M., Mst. Shahrin, S., & Sai Praveen, K. (2023). Integration Of IOT And EDGE Computing For Low-Latency Data Analytics In Smart Cities And IOT Networks. *Journal of Sustainable Development and Policy*, *2*(03), 01-33. https://doi.org/10.63125/004h7m29
[83]. Md Musfiqur, R., & Md.Kamrul, K. (2023). Mechanisms By Which AI-Enabled Crm Systems Influence Customer Retention And Overall Business Performance: A Systematic Literature Review Of Empirical Findings. *International Journal of Business and Economics Insights*, *3*(1), 31-67. https://doi.org/10.63125/qqe2bm11
[84]. Md Muzahidul, I., & Md Mohaiminul, H. (2023). Explainable AI (XAI) Models For Cloud-Based Business Intelligence: Ensuring Compliance And Secure Decision-Making. *American Journal of Interdisciplinary Studies*, *4*(03), 208–249. https://doi.org/10.63125/5etfhh77
[85]. Md Nahid, H. (2022). Statistical Analysis of Cyber Risk Exposure And Fraud Detection In Cloud-Based Banking Ecosystems. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *2*(1), 289–331. https://doi.org/10.63125/9wf91068
[86]. Md Sarwar Hossain, S., & Md Milon, M. (2022). Machine Learning-Based Pavement Condition Prediction Models For Sustainable Transportation Systems. *American Journal of Interdisciplinary Studies*, *3*(01), 31–64. https://doi.org/10.63125/1jsmkg92
[87]. Md. Abdur, R., & Zamal Haider, S. (2022). Assessment Of Data-Driven Vendor Performance Evaluation In Retail Supply Chains Analyzing Metrics, Scorecards, And Contract Management Tools. *Journal of Sustainable Development and Policy*, *1*(04), 71-116. https://doi.org/10.63125/2a641k35
[88]. Md. Al Amin, K., & Sai Praveen, K. (2023). The Role Of Industrial Engineering In Advancing Sustainable Manufacturing And Quality Compliance In Global Engineering Systems. *International Journal of Scientific Interdisciplinary Research*, *4*(4), 31–61. https://doi.org/10.63125/8w1vk676
[89]. Md. Hasan, I., & Ashraful, I. (2023). The Effect Of Production Planning Efficiency On Delivery Timelines In U.S. Apparel Imports. *Journal of Sustainable Development and Policy*, *2*(04), 35-73. https://doi.org/10.63125/sg472m51
[90]. Md. Jobayer Ibne, S., & Md. Kamrul, K. (2023). Automating NIST 800-53 Control Implementation: A Cross-Sector Review Of Enterprise Security Toolkits. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *3*(1), 160–195. https://doi.org/10.63125/prkw8r07
[91]. Md.Akbar, H., & Farzana, A. (2021). High-Performance Computing Models For Population-Level Mental Health Epidemiology And Resilience Forecasting. *American Journal of Health and Medical Sciences*, *2*(02), 01–33. https://doi.org/10.63125/k9d5h638
[92]. Mehairjan, R. P. Y. (2017). *Risk-based maintenance for electricity network organizations*. Springer.
[93]. Milazzo, M. F., & Bragatto, P. (2019). A framework addressing a safe ageing management in complex industrial sites: The Italian experience in «Seveso» establishments. *Journal of Loss Prevention in the Process Industries*, *58*, 70-81.
[94]. Mirjalili, S., & Dong, J. S. (2020). *Multi-objective optimization using artificial intelligence techniques*. Springer.
[95]. Mirjalili, S., Jangir, P., & Saremi, S. (2017). Multi-objective ant lion optimizer: a multi-objective optimization algorithm for solving engineering problems. *Applied intelligence*, *46*(1), 79-95.
[96]. Mohammad Mushfequr, R., & Ashraful, I. (2023). Automation And Risk Mitigation in Healthcare Claims: Policy And Compliance Implications. *Review of Applied Science and Technology*, *2*(04), 124–157. https://doi.org/10.63125/v73gyg14
[97]. Mohammad Mushfequr, R., & Sai Praveen, K. (2022). Quantitative Investigation Of Information Security Challenges In U.S. Healthcare Payment Ecosystems. *International Journal of Business and Economics Insights*, *2*(4), 42–73. https://doi.org/10.63125/gcg0fs06
[98]. Mohandes, B., Abdelmagid, Y. L., & Boiko, I. (2018). Development of PSS tuning rules using multi-objective optimization. *International Journal of Electrical Power & Energy Systems*, *100*, 449-462.
[99]. Moraes, C. C. d. F., Pinheiro, P. R., Rolim, I. G., Costa, J. L. d. S., Junior, M. d. S. E., & Andrade, S. J. M. D. (2021). Using the multi-criteria model for optimization of operational routes of thermal power plants. *Energies*, *14*(12), 3682.
[100]. Mortuza, M. M. G., & Rauf, M. A. (2022). Industry 4.0: An Empirical Analysis of Sustainable Business Performance Model Of Bangladeshi Electronic Organisations. *International Journal of Economy and Innovation*. https://gospodarkainnowacje.pl/index.php/issue_view_32/article/view/826
[101]. Munodawafa, R. T., & Johl, S. K. (2019). Big data analytics capabilities and eco-innovation: a study of energy companies. *Sustainability*, *11*(15), 4254.

[102]. Mushtaq, S., Islam, M. M., & Sohaib, M. (2021). Deep learning aided data-driven fault diagnosis of rotatory machine: A comprehensive review. *Energies*, *14*(16), 5150.

[103]. Nemati, M., Braun, M., & Tenbohlen, S. (2018). Optimization of unit commitment and economic dispatch in microgrids based on genetic algorithm and mixed integer linear programming. *Applied energy*, *210*, 944-963.

[104]. Netto, A. C., de Andrade Melani, A. H., Murad, C. A., de Carvalho Michalski, M. A., Martha de Souza, G. F., & Nabeta, S. I. (2020). A novel approach to defining maintenance significant items: A hydro generator case study. *Energies*, *13*(23), 6273.

[105]. Nguyen, K. T., & Medjaher, K. (2019). A new dynamic predictive maintenance framework using deep learning for failure prognostics. *Reliability Engineering & System Safety*, *188*, 251-262.

[106]. Niu, G. (2017). Data-driven technology for engineering systems health management. *Springer Singapore*, *10*, 978-981.

[107]. Nouri Gharahasanlou, A., Ataei, M., Khalokakaie, R., Barabadi, A., & Einian, V. (2017). Risk based maintenance strategy: a quantitative approach based on time-to-failure model. *International Journal of System Assurance Engineering and Management*, *8*(3), 602-611.

[108]. Nowakowski, T., Tubis, A., & Werbińska-Wojciechowska, S. (2019). Evolution of technical systems maintenance approaches–review and a case study. International conference on intelligent systems in production engineering and maintenance,

[109]. Nunez, D. L., & Borsato, M. (2017). An ontology-based model for prognostics and health management of machines. *Journal of Industrial Information Integration*, *6*, 33-46.

[110]. Nyanchoka, L., Tudur-Smith, C., Iversen, V., Tricco, A. C., & Porcher, R. (2019). A scoping review describes methods used to identify, prioritize and display gaps in health research. *Journal of clinical epidemiology*, *109*, 99-110.

[111]. Okwuobi, S., Ishola, F., Ajayi, O., Salawu, E., Aworinde, A., Olatunji, O., & Akinlabi, S. A. (2018). A reliability-centered maintenance study for an individual section-forming machine. *Machines*, *6*(4), 50.

[112]. Pankaz Roy, S., & Md. Kamrul, K. (2023). HACCP and ISO Frameworks For Enhancing Biosecurity In Global Food Distribution Chains. *American Journal of Scholarly Research and Innovation*, *2*(01), 314–356. https://doi.org/10.63125/9pbp4h37

[113]. Piechnicki, F., Dos Santos, C. F., De Freitas Rocha Loures, E., & Dos Santos, E. A. P. (2021). Data fusion framework for decision-making support in reliability-centered maintenance. *Journal of Industrial and Production Engineering*, *38*(1), 1-17.

[114]. Pisacane, O., Potena, D., Antomarioni, S., Bevilacqua, M., Emanuele Ciarapica, F., & Diamantini, C. (2021). Data-driven predictive maintenance policy based on multi-objective optimization approaches for the component repairing problem. *Engineering optimization*, *53*(10), 1752-1771.

[115]. Prasetyo, Y. T., & Rosita, K. K. M. (2020). Equipment reliability optimization using predictive reliability centered maintenance: a case-study illustration and comprehensive literature review. 2020 7th International Conference on Frontiers of Industrial Engineering (ICFIE),

[116]. Rakibul, H., & Samia, A. (2022). Information System-Based Decision Support Tools: A Systematic Review Of Strategic Applications In Service-Oriented Enterprises. *Review of Applied Science and Technology*, *1*(04), 26-65. https://doi.org/10.63125/w3cevz78

[117]. Reza, M., Vorobyova, K., & Rauf, M. (2021). The effect of total rewards system on the performance of employees with a moderating effect of psychological empowerment and the mediation of motivation in the leather industry of Bangladesh. *Engineering Letters*, *29*, 1-29.

[118]. Rizk-Allah, R. M., Hassanien, A. E., & Slowik, A. (2020). Multi-objective orthogonal opposition-based crow search algorithm for large-scale multi-objective optimization. *Neural computing and applications*, *32*(17), 13715-13746.

[119]. Rony, M. A., & Ashraful, I. (2022). Big Data And Engineering Analytics Pipelines For Smart Manufacturing: Enhancing Efficiency, Quality, And Predictive Maintenance. *American Journal of Scholarly Research and Innovation*, *1*(02), 59–85. https://doi.org/10.63125/rze0my79

[120]. Rosita, K. K. M., & Rada, M. V. (2021). Equipment reliability optimization using predictive reliability centered maintenance. 2021 IEEE 8th International Conference on Industrial Engineering and Applications (ICIEA),

[121]. Ross, A. M., Rasor, B., & Moleski, H. (2018). How reliability centered maintenance (RCM) will affect system reliability. 2018 IEEE/PES Transmission and Distribution Conference and Exposition (T&D),

[122]. Rupp, A. A. (2018). Designing, evaluating, and deploying automated scoring systems with validity in mind: Methodological design decisions. *Applied Measurement in Education*, *31*(3), 191-214.

[123]. Saba, A., Shaikat, B., & Tonoy Kanti, C. (2023). Integration Of Artificial Intelligence And Advanced Computing To Develop Resilient Cyber Defense Systems. *Journal of Sustainable Development and Policy*, *2*(04), 74-107. https://doi.org/10.63125/rxyc6y88

[124]. Saba, A., & Tonoy Kanti, C. (2023). Explainable Artificial Intelligence (XAI) Approaches For Cyber Risk Assessment In Financial Services. *American Journal of Interdisciplinary Studies*, *4*(03), 96-135. https://doi.org/10.63125/3gjcb322

[125]. Saikat, S. (2021). Real-Time Fault Detection in Industrial Assets Using Advanced Vibration Dynamics And Stress Analysis Modeling. *American Journal of Interdisciplinary Studies*, *2*(04), 39–68. https://doi.org/10.63125/0h163429

[126]. Saikat, S. (2022). CFD-Based Investigation of Heat Transfer Efficiency In Renewable Energy Systems. *International Journal of Scientific Interdisciplinary Research*, *1*(01), 129–162. https://doi.org/10.63125/ttw40456

[127]. Savela, T. (2018). The advantages and disadvantages of quantitative methods in schoolscape research. *Linguistics and Education*, *44*, 31-44.

[128]. Scheu, M. N., Tremps, L., Smolka, U., Kolios, A., & Brennan, F. (2019). A systematic Failure Mode Effects and Criticality Analysis for offshore wind turbine systems towards integrated condition based maintenance strategies. *Ocean Engineering*, *176*, 118-133.

[129]. Schoen, R. C., & LaVenia, M. (2019). Teacher beliefs about mathematics teaching and learning: Identifying and clarifying three constructs. *Cogent Education*, *6*(1), 1599488.
[130]. Shafiee, M., Enjema, E., & Kolios, A. (2019). An integrated FTA-FMEA model for risk analysis of engineering systems: a case study of subsea blowout preventers. *Applied Sciences*, *9*(6), 1192.
[131]. Shaikh, S., & Aditya, D. (2021). Federated Learning-Driven Predictive Quality Analytics and Supply Chain Optimization In Distributed Manufacturing Networks. *Review of Applied Science and Technology*, *6*(1), 74-107. https://doi.org/10.63125/k18cbz55
[132]. Shaikh, S., & Md. Tahmid Farabe, S. (2023). Digital Twin-Driven Process Modeling For Energy Efficiency And Lifecycle Optimization In Industrial Facilities. *American Journal of Interdisciplinary Studies*, *4*(03), 65–95. https://doi.org/10.63125/e4q64869
[133]. Sharifi, M. R., Akbarifard, S., Qaderi, K., & Madadi, M. R. (2021). A new optimization algorithm to solve multi-objective problems. *Scientific Reports*, *11*(1), 20326.
[134]. Sheng, J., & Prescott, D. (2017). A hierarchical coloured Petri net model of fleet maintenance with cannibalisation. *Reliability Engineering & System Safety*, *168*, 290-305.
[135]. Shi, X., Zhuang, H., & Xu, D. (2021). Structured survey of human factor-related maritime accident research. *Ocean Engineering*, *237*, 109561.
[136]. Sifonte, J. R., & Reyes-Picknell, J. V. (2017). *Reliability Centered Maintenance–Reengineered: Practical Optimization of the RCM Process with RCM-R®*. Productivity Press.
[137]. Singh, S. N., & Pretorius, J. (2017). Development of a sem-quantitative approach for risk based inspection and maintenance of thermal power plant components. *SAIEE Africa Research Journal*, *108*(3), 128-138.
[138]. Sonmez, Z. (2018). Interregional inventor collaboration and the commercial value of patented inventions: evidence from the US biotechnology industry. *The annals of regional science*, *61*(2), 399-438.
[139]. Spolaor, N., Lorena, A. C., & Diana Lee, H. (2017). Feature selection via pareto multi-objective genetic algorithms. *Applied Artificial Intelligence*, *31*(9-10), 764-791.
[140]. Su, C., & Liu, Y. (2020). Multi-objective imperfect preventive maintenance optimisation with NSGA-II. *International Journal of Production Research*, *58*(13), 4033-4049.
[141]. Taha, K. (2020). Methods that optimize multi-objective problems: A survey and experimental evaluation. *IEEE Access*, *8*, 80855-80878.
[142]. Tang, Y., Liu, Q., Jing, J., Yang, Y., & Zou, Z. (2017). A framework for identification of maintenance significant items in reliability centered maintenance. *Energy*, *118*, 1295-1303.
[143]. Teixeira, H. N., Lopes, I., & Braga, A. C. (2020). Condition-based maintenance implementation: a literature review. *Procedia Manufacturing*, *51*, 228-235.
[144]. Teng, S. Y., How, B. S., Leong, W. D., Teoh, J. H., Cheah, A. C. S., Motavasel, Z., & Lam, H. L. (2019). Principal component analysis-aided statistical process optimisation (PASPO) for process improvement in industrial refineries. *Journal of Cleaner Production*, *225*, 359-375.
[145]. Tonoy Kanti, C., & Shaikat, B. (2021). Blockchain-Enabled Security Protocols Combined With AI For Securing Next-Generation Internet Of Things (IOT) Networks. *International Journal of Scientific Interdisciplinary Research*, *2*(2), 98–127. https://doi.org/10.63125/pcdqzw41
[146]. Velmurugan, R. S., Dhingra, T., & Velmurugan. (2021). *Asset Maintenance Management in Industry*. Springer.
[147]. Velte, P., & Stawinoga, M. (2017). Integrated reporting: The current state of empirical research, limitations and future research implications. *Journal of Management Control*, *28*(3), 275-320.
[148]. 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.
[149]. Ventikos, N. P., Chmurski, A., & Louzis, K. (2020). A systems-based application for autonomous vessels safety: Hazard identification as a function of increasing autonomy levels. *Safety science*, *131*, 104919.
[150]. Vogl, G. W., Weiss, B. A., & Helu, M. (2019). A review of diagnostic and prognostic capabilities and best practices for manufacturing. *Journal of Intelligent Manufacturing*, *30*(1), 79-95.
[151]. Wang, H., Sheng, B., Lu, Q., Yin, X., Zhao, F., Lu, X., Luo, R., & Fu, G. (2021). A novel multi-objective optimization algorithm for the integrated scheduling of flexible job shops considering preventive maintenance activities and transportation processes. *Soft Computing*, *25*(4), 2863-2889.
[152]. Wang, J., Jiang, C., Zhang, H., Ren, Y., Chen, K.-C., & Hanzo, L. (2020). Thirty years of machine learning: The road to Pareto-optimal wireless networks. *IEEE Communications Surveys & Tutorials*, *22*(3), 1472-1514.
[153]. Wang, N., Zhao, W.-j., Wu, N., & Wu, D. (2017). Multi-objective optimization: a method for selecting the optimal solution from Pareto non-inferior solutions. *Expert Systems with Applications*, *74*, 96-104.
[154]. Wang, P., Poovendran, P., & Manokaran, K. B. (2021). Fault detection and control in integrated energy system using machine learning. *Sustainable Energy Technologies and Assessments*, *47*, 101366.
[155]. Wang, Y., Liu, H., Zheng, W., Xia, Y., Li, Y., Chen, P., Guo, K., & Xie, H. (2019). Multi-objective workflow scheduling with deep-Q-network-based multi-agent reinforcement learning. *IEEE Access*, *7*, 39974-39982.
[156]. Williams, P. J., & Kendall, W. L. (2017). A guide to multi-objective optimization for ecological problems with an application to cackling goose management. *Ecological Modelling*, *343*, 54-67.
[157]. Wróbel, K., Montewka, J., & Kujala, P. (2018). Towards the development of a system-theoretic model for safety assessment of autonomous merchant vessels. *Reliability Engineering & System Safety*, *178*, 209-224.

[158]. Xie, L., Håbrekke, S., Liu, Y., & Lundteigen, M. A. (2019). Operational data-driven prediction for failure rates of equipment in safety instrumented systems: A case study from the oil and gas industry. *Journal of Loss Prevention in the Process Industries*, *60*, 96-105.

[159]. Xu, G., Liu, M., Wang, J., Ma, Y., Wang, J., Li, F., & Shen, W. (2019). Data-driven fault diagnostics and prognostics for predictive maintenance: A brief overview. 2019 IEEE 15th international conference on automation science and engineering (CASE),

[160]. Yang, Y.-J., Zhang, X.-Y., Zhao, Z.-J., Wang, G.-H., He, Y.-J., Wu, Y.-L., & Li, J. (2020). Applying reliability centered maintenance (RCM) to sampling subsystem in continuous emission monitoring system. *IEEE Access*, *8*, 55054-55062.

[161]. Yao, Y., Gui, G., Yang, S., & Zhang, S. (2021). An adaptive anti-noise network with recursive attention mechanism for gear fault diagnosis in real-industrial noise environment condition. *Measurement*, *186*, 110169.

[162]. Yavuz, O., Doğan, E., Carus, E., & Görgülü, A. (2019). Reliability centered maintenance practices in food industry. *Procedia Computer Science*, *158*, 227-234.

[163]. Yousefi, A., & Hernandez, M. R. (2020). A novel methodology to measure safety level of a process plant using a system theory based method (STAMP). *Process Safety and Environmental Protection*, *136*, 296-309.

[164]. Yuan, Z., Qin, W., & Zhao, J. (2017). Smart manufacturing for the oil refining and petrochemical industry. *Engineering*, *3*(2), 179-182.

[165]. Zamal Haider, S., & Hozyfa, S. (2023). A Quantitative Study On IT-Enabled ERP Systems And Their Role In Operational Efficiency. *International Journal of Scientific Interdisciplinary Research*, 4(4), 62–99. https://doi.org/10.63125/nbpyce10

[166]. Zeinalnezhad, M., Chofreh, A. G., Goni, F. A., & Klemeš, J. J. (2020). Critical success factors of the reliability-centred maintenance implementation in the oil and gas industry. *Symmetry*, *12*(10), 1585.

[167]. Zhang, L., & Zhang, J. (2020). A data-driven maintenance framework under imperfect inspections for deteriorating systems using multitask learning-based status prognostics. *IEEE Access*, *9*, 3616-3629.

[168]. 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.

[169]. Zhen, X., Vinnem, J. E., Peng, C., & Huang, Y. (2018). Quantitative risk modelling of maintenance work on major offshore process equipment. *Journal of Loss Prevention in the Process Industries*, *56*, 430-443.

[170]. Zhong, K., Han, M., & Han, B. (2019). Data-driven based fault prognosis for industrial systems: A concise overview. *IEEE/CAA Journal of Automatica Sinica*, *7*(2), 330-345.

[171]. Zobayer, E. (2021a). Data Driven Predictive Maintenance In Petroleum And Power Systems Using Random Forest Regression Model For Reliability Engineering Framework. *Review of Applied Science and Technology*, *6*(1), 108-138. https://doi.org/10.63125/5bjx6963

[172]. Zobayer, E. (2021b). Machine Learning Approaches For Optimization Of Lubricant Performance And Reliability In Complex Mechanical And Manufacturing Systems. *American Journal of Scholarly Research and Innovation*, *1*(01), 61–92. https://doi.org/10.63125/5zvkgg52

[173]. Zobayer, E. (2023). IOT Integration In Intelligent Lubrication Systems For Predictive Maintenance And Performance Optimization In Advanced Manufacturing Industries. *Journal of Sustainable Development and Policy*, *2*(04), 140-173. https://doi.org/10.63125/zybrmx69

[174]. Zou, P. X., Xu, X., Sanjayan, J., & Wang, J. (2018). A mixed methods design for building occupants' energy behavior research. *Energy and Buildings*, *166*, 239-249.