

## AI-Driven Spare Parts Optimisation for Smart Manufacturing: A Predictive Decision-Support Framework for Saudi Arabian Industries

Syed Nadeemuddin<sup>1\*</sup> 

<sup>1</sup>Independent Researcher

**\*Corresponding Author**  
**Syed Nadeemuddin**  
Independent Researcher

### Article History

Received: 14.08.2026

Accepted: 06.10.2026

Published: 12.10.2026

**Abstract:** In asset-intensive manufacturing, spare-parts demand is irregular, equipment condition changes over time, and conventional reorder rules may not exploit enterprise and machine data. This article examines how artificial intelligence can connect demand forecasting, asset-health information and inventory decisions within a predictive decision-support framework for manufacturing in Saudi Arabia. A structured narrative review synthesises 29 journal publications from 2020 to 2025 covering spare-parts forecasting, predictive maintenance, digital twins, inventory optimisation, reinforcement learning and explainable artificial intelligence. The evidence is organised by decision problem, data foundation, analytical method, operational output and boundary of application; Saudi policy material provides context, while Project Management Institute guidance informs a separate lifecycle-governance adaptation. The review is not presented as an exhaustive systematic search or meta-analysis. The synthesis supports evaluating forecasts through the inventory decisions they inform and shows how demand estimates and equipment condition information can be connected to replenishment and maintenance decisions. Building on this evidence, the paper proposes a five-layer framework linking industrial data, context and feature engineering, predictive services, constrained inventory optimisation, and governed execution and learning. Proposed governance controls include review checkpoints, assigned decision rights, bounded automation and auditable feedback. The framework is conceptual and has not been validated in Saudi manufacturing plants. It provides a staged pilot pathway from data readiness and criticality segmentation to condition-informed forecasting and closed-loop optimisation, while requiring plant specific assessment of technical compatibility, organisational readiness and economic value.

**Keywords:** Artificial intelligence, decision governance, predictive maintenance, smart manufacturing, spare-parts optimization.

**Copyright © 2026 The Author(s):** This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

## I. INTRODUCTION

Smart manufacturing has increased the quantity and variety of information available for operational decisions, but spare parts management remains a difficult issue since the demand process involved is essentially irregular. A large number of maintenance items have periods when there is no

demand whatsoever, followed by sudden and at times grouped requirements resulting from breakdowns, planned interventions or engineering changes. The economic consequences of a poor decision are asymmetric: having too much stock ties up working capital and may cause the stock to become obsolete, while a lack of an important component can result in longer periods of downtime, delay production and

**Citation:** Nadeemuddin, S. (2026). AI-Driven Spare Parts Optimisation for Smart Manufacturing: A Predictive Decision-Support Framework for Saudi Arabian Industries. *Glob Acad J Econ Buss*, 8(5), 1682-1696.  
<https://doi.org/10.36348/gajeb.2026.v08i05.055>

increase the cost of emergency procurement. That is the reason why studies on forecasting for intermittent demand emphasise that the problem cannot be solely regarded as being similar to that of conventional high-volume demand prediction [1]. Recent research investigates both univariate machine-learning models and methods that transfer information across demand series. Their usefulness depends on data quality, demand characteristics and the operational decision in which the forecast is applied [2], [3].

The application of artificial intelligence (AI) offers the chance to improve this decision-making process. Instead of basing its actions solely on previous withdrawals, a smart manufacturing system could incorporate data from enterprise resource planning (ERP) records, maintenance work orders, bills of material, asset criticality, supplier lead times, the characteristics of the installed base, and machine condition data. As has been shown by research into predictive maintenance, sensor data and machine-learning models are able to estimate the risk of deterioration or failure before any maintenance action is carried out [4], [5]. Moreover, the research into digital twins extends this method by linking physical assets with continuously updated digital replicas that can be used for condition assessment, simulation, and maintenance planning [6], [7]. In the case of spare parts, these capabilities are significant because an anticipated maintenance intervention can serve as a leading indicator of future demand for materials rather than having the demand only become apparent after a stock request has been made.

The literature addresses forecasting, maintenance, inventory control and intelligent automation through different, sometimes overlapping, research streams. Certain studies focus on intermittent-demand prediction using support-vector, ensemble and deep-learning methods [2], [8], [9]. Perez *et al.*, [10] compare inventory optimisation approaches in a supply-chain simulation, while Rodrigues and Yoneyama [11] explicitly connect prognostics with spare-parts control. Wesendrup and Hellingrath [12] integrate spare-parts sourcing, production and maintenance planning for a single-machine system, and Burtea and Tsay [13] address constrained reinforcement learning for inventory management. Integration is therefore not absent from the literature. The remaining design question addressed here is how to connect these capabilities with item segmentation, enterprise data and accountable human decisions in a modular implementation architecture.

### **The challenge is relevant to manufacturing in Saudi Arabia**

In the context of the Kingdom's industrial transformation. An official account of the Future Factories programme describes an ambition to upgrade 4,000 factories using technologies including artificial intelligence, automation and 3D printing [14]. This establishes a policy context, not evidence that any particular spare-parts architecture is effective. The present framework considers plants with existing ERP systems, variable supplier lead times and different consequences of part unavailability as intended application conditions. Their prevalence and importance must be established for each participating Saudi plant. A decision-support tool for such settings should address managerial integration, traceability and governance alongside predictive performance.

The review meets this need by summarising selected evidence published between 2020 and 2025 concerning AI-enabled spare parts forecasting and other smart manufacturing technologies, and by proposing a predictive decision-support framework for potential use in industries that depend heavily on assets. The main point put forward is that operational value should be evaluated when forecasting, knowledge of the condition of assets and inventory optimisation are regarded as part of a single, integrated decision system. As a result, the paper studies not just the algorithms but also the data architecture, the decision variables, the question of explainability, the level of maturity of implementation, and the institutional controls necessary for reliable use.

To make these institutional controls explicit, the paper adapts the Cognitive Project Management for AI (CPMAI) methodology associated with Project Management Institute (PMI) Certified Professional in Managing AI (PMI-CPMAI). PMI describes an iterative life-cycle that connects business objectives, data work, model development, evaluation and operational use [15]. Here, decision governance means assigning authority, specifying the evidence required before an action is permitted, and retaining accountability for its consequences. The proposed adaptation connects lifecycle review with individual spare-parts decisions; it remains a design contribution requiring empirical evaluation.

## **II. CONCEPTUAL BASIS, AIM AND REVIEW QUESTIONS**

Spare-parts optimisation differs from many finished-goods forecasting settings in three respects. Firstly, demand is often intermittent or irregular, so the choice of error measure can affect model selection. For example, absolute-error loss targets a median: if more than half of the observed periods

have zero demand, a zero forecast minimises in-sample absolute error among constant point forecasts, even though it need not support an acceptable stocking policy. This is not a property of all error measures, and forecast evaluation should be connected to the inventory objective [1]. Secondly, demand is influenced by maintenance policy: inspection frequency, preventive-maintenance intervals and replacement decisions affect when parts are required. Thirdly, a service failure has a technical consequence, since the lack of a low-value component can stop an expensive asset. For this reason, the decision objective should consider availability, inventory cost, downtime exposure, obsolescence and procurement responsiveness rather than forecast accuracy alone.

The benefits of tailoring the model to the demand structure are shown by recent forecasting studies. Abbasimehr *et al.*, [16] investigate an optimised long short-term memory (LSTM) network for demand forecasting in a furniture-company setting; this provides adjacent methodological evidence rather than direct validation for intermittent spare-parts demand. Intermittent spare-parts settings require explicit treatment of sparse observations and item-specific features. Various methods have been suggested to handle these features, including support-vector methods, feature-adaptive models and hybrid resampling approaches [2], [3], [17]. Ensemble learning can likewise make use of a variety of predictors and decrease sensitivity to any one modelling assumption, as has been demonstrated in aftermarket applications [8], [18]. However, the operational importance of any predicted quantity is still not complete unless it is linked to an inventory action.

The second conceptual foundation is provided by condition based information. Studies on predictive maintenance and digital twins have shown that failure precursors, the level of usage and the patterns of degradation could be turned into maintenance warnings or information about the remaining life [19], [20]. Rodrigues and Yoneyama [11] demonstrate this link in a periodic-review inventory model for non-repairable items, using remaining-useful-life estimates to update stock-control parameters under a fill-rate constraint. The significance of this lies in the fact that when the probability of failure changes, the expected timing and urgency of parts requirements may also change. A predictive inventory system should therefore adjust its stocking decisions as the evidence regarding asset health changes rather than waiting for historical consumption to indicate a new pattern.

This review critically examines selected evidence published between 2020 and 2025 and

develops an AI-driven decision support framework which connects intermittent-demand forecasting, asset-condition intelligence and inventory policy for potential use in Saudi manufacturing. The analysis is guided by five objectives: to compare the assumptions and reported applications of AI methods for sparse spare-parts demand; to look at how maintenance and prognostic signals can function as demand predictors; to evaluate optimisation and reinforcement-learning approaches that turn the forecasts into replenishment decisions; to identify the explainability, integration and governance requirements needed for enterprise deployment; and to convert the combined body of evidence into a staged implementation architecture.

The following three review questions are posed. The first, RQ1, investigates which data and AI methods are most defensible under various spare-parts demand scenarios. RQ2 looks at how information relating to condition, maintenance, and digital twin technology should affect inventory decisions. RQ3 examines how a predictive decision architecture and CPMAI-based governance controls could support accountable deployment within Saudi industrial organisations, subject to local validation. These questions intentionally go beyond the task of picking out a single “best” algorithm and instead concentrate on the conditions under which analytical methods generate operational value.

### III. REVIEW METHODOLOGY

A structured narrative review is used because the research problem spans forecasting, inventory control, maintenance engineering and smart manufacturing. The bounded corpus comprises 29 journal publications dated 2020–2025. It covers spare parts, intermittent demand, forecasting, inventory optimisation, predictive maintenance, prognostics, digital twins, artificial intelligence, machine learning, reinforcement learning, explainable AI and smart manufacturing. These concepts define the thematic scope; they are not presented as a reproducible database search strategy. The review does not report an auditable record of database searches, initial retrieval counts or screening decisions and therefore makes no claim of exhaustive or systematic coverage.

Within this bounded corpus, the analytical rationale for retention is purposive and tied to the three review questions rather than to statistical representativeness. Each retained publication contributes to at least one defined decision problem or evidence stream, and Appendix A states the role and boundary of inference used for it. Direct spare-parts studies provide the core evidence; adjacent forecasting, maintenance, digital-twin, reinforcement-learning and broader manufacturing

studies are retained only where they contribute a method, capability, constraint or deployment consideration needed to interpret the proposed architecture. This explains the analytical composition of the 29-publication corpus without reconstructing an undocumented database-search history, retrieval count or screening log.

The synthesis includes publications relevant to industrial forecasting, maintenance, inventory decisions or manufacturing integration, with identifiable titles, authors, publication details and DOIs. General demand-forecasting and supply chain studies are retained as adjacent evidence and distinguished from direct spare-parts applications. Journal issue years are used where these differ from online-first dates. Bibliographic identities were checked against publisher records and publisher-deposited DOI metadata, with verification updated on 3 October 2026. DOI traceability establishes the identity of a publication, not the validity of its findings. The Saudi policy reference is used only to establish application context and is excluded from the 29-publication research corpus. Two official PMI explanatory sources, consulted on 4 October 2026, provide the methodology basis for the governance extension [15], [21]. These professional sources are also excluded from the research corpus and the study-level evidence map. The six-phase lifecycle is mapped to the proposed architecture, after which spare-parts-specific review evidence, approval responsibilities and escalation rules are specified as design propositions. This mapping is interpretive; PMI guidance is not treated as evidence of inventory-performance improvement.

The evidence is organised around the decision problem, data foundation, modelling approach, operational output and boundary of application. Table I summarises the evidence streams, while Appendix A provides a study-level map of the retained publications and the roles they play in the proposed framework. Reviews supply background and identify research directions; primary studies supply examples of models or applications. They are not counted as independent replications of the same result. No pooled effect estimate, numerical quality score or formal risk-of-bias assessment is claimed. Where verification was limited to original abstracts or publisher previews, interpretation is restricted to the stated scope and reported contribution rather than unexamined experimental details.

The synthesis follows an evidence-to-decision approach, distinguishing prediction from replenishment, maintenance and allocation decisions. The framework in Fig. 1 is an interpretive synthesis of requirements discussed across the

corpus: contextual data, condition information, uncertainty, constrained optimisation, explainability and feedback. Its proposed routing rules, information reconciliation and governance controls are design recommendations, not findings that every cited study has tested. Existing integrated approaches are discussed explicitly in Section IV to delimit the contribution. No new plant-level performance data are claimed.

The evidence is compared in terms of the decisions it supports rather than ranked by reported forecasting performance. Metrics are not directly comparable across different item populations, horizons, sparsity levels, cost assumptions or train-test designs. Results from simulated inventory systems and single-site applications are therefore treated as conditional evidence. Transfer to Saudi plants requires examination of equipment fleets, supplier behaviour, data latency, maintenance practice and digital maturity. The synthesis asks which method may be useful, for which decision and under which data conditions; the proposed deployment controls remain subject to empirical assessment.

#### IV. SYNTHESIS OF THE EVIDENCE AND KEY FINDINGS

Forecasting under intermittency requires attention to both the occurrence of demand and the magnitude of non-zero events. Pınçe *et al.*, [1] review the diversity of approaches and the importance of demand characteristics. Jiang *et al.*, [2] develop an adaptive univariate support-vector model using heavy-duty vehicle spare-parts data; their work should not be interpreted as evidence that external covariates are always necessary for machine learning. Yang *et al.*, [23] combine scenario analysis with an improved stacking model, while Fan *et al.*, [3] transfer information across demand series through intermittent feature adaptation. Zhuang *et al.*, [24] combine classification of demand occurrence with regression of non-zero demand using automotive aftermarket data. Taken together, these studies motivate testing different model routes. The segmentation proposed here also incorporates maintenance criticality and repairability, which are operational design choices rather than a shared classification validated by all these studies.

Deeper and ensemble models expand the available methods, but their value remains dependent on the setting. Abbasimehr *et al.*, [16] provide an LSTM demand-forecasting example outside the spare-parts domain. Chien *et al.*, [8] develop an aftermarket forecasting framework with classification and model monitoring, while Dodin *et al.*, [18] report a deployed forecasting pipeline at Bombardier that combines machine learning and

time-series models. Ifraz *et al.*, [25] examine regression and machine-learning methods using bus-fleet demand and maintenance-related variables. Zhang *et al.*, [9] test transformers on 925 intermittent-demand items from an airline spare-parts provider. These settings support method comparison, not an assumption of equivalent performance in Saudi manufacturing. Caserta and D'Angelo [22] specifically address little or no historical demand information. Accordingly, the proposed architecture treats low history, poor data quality and model complexity as separate problems: additional model capacity cannot by itself repair missing asset links or inconsistent transaction records.

The second point concerns the relationship between maintenance intelligence and material demand. Reviews of predictive maintenance point to machine learning, sensing and connected industrial systems as the key enablers of condition-based intervention [4], [5]. Silvestri *et al.*, [19] explain how Industry 4.0 technologies alter maintenance processes by boosting connectivity and analytical competencies. These developments are important for inventory since maintenance plans act as demand-generating mechanisms; if the degradation trend of an asset shows an increasing chance of having to replace a component, the spare-parts model can modify the expected demand before a replacement requirement is issued. Rodrigues and Yoneyama [11] make this link explicit in a single-site, single-component inventory model with backorders and a fillrate constraint. Extension to repairable items, multiple sites or shared components requires additional modelling.

Digital twins can strengthen the connection between asset state and maintenance decisions. van Dinter *et al.*, [6] review digital-twin applications in predictive maintenance and identify data and model-complexity challenges; Chen *et al.*, [7] examine the role of machine learning in this relationship. Feng *et al.*, [20] propose a multi-level maintenance framework that accounts for component dependencies and maintenance resources. These contributions support investigating condition informed decisions, but do not establish the performance of the present inventory architecture. As a starting design, this paper proposes a linked asset-state record containing equipment hierarchy, installed configuration, usage, health indicators, maintenance history and bills of material. Such a record could improve attribution of requirements to assets; a static database alone should not be described as a digital twin. The degree of updating, connection to the physical asset and modelling required must be specified for each implementation.

The transition from prediction to inventory policy requires explicit objectives and constraints. Perez *et al.*, [10] compare deterministic optimisation, stochastic optimisation and reinforcement learning for a single-product supply-chain inventory problem in a stochastic simulation environment. Their study illustrates the need to evaluate decision outcomes as well as forecasts, but is not a direct test of condition-informed spare-parts control. In the proposed framework, inventory cost and downtime exposure would be balanced subject to service targets and feasible procurement, storage, repair and substitution rules. A reduction in forecast error need not change the selected inventory action or create an economic benefit. Conversely, an improvement that changes a decision at a critical threshold could matter operationally. These effects must be measured rather than inferred from forecast accuracy alone.

Reinforcement learning provides a different method since it learns a policy based on sequences of state, action, and reward rather than merely giving a single forecast. Papers on smart manufacturing emphasise its potential for use in adaptive control, scheduling and the allocation of resources [26], [27]. Wesendrup and Hellingrath [12] combine prognostic information with demand, production, spare parts and maintenance planning by means of reinforcement learning, in a single-machine setting, illustrating coordination within the assumptions of that model. Burtea and Tsay [13] apply constrained continuous-action reinforcement learning to inventory management in the supply chain, stressing the need to include operational constraints in the learning process. These studies motivate further testing in changing supply environments. They do not establish that a learned policy will remain effective when the simulated and operational environments differ; constraints, reward design and validation must therefore be explicit. Alginahi *et al.*, [28] make a similar point that the adoption of reinforcement learning in industry is restricted by issues of interpretability, validation and capability for deployment.

The fourth point concerns explainability and human monitoring. Inventory planners and maintenance engineers should be able to question recommendations when engineering knowledge, supplier disruption or a planned shutdown is not represented in the model. Albayrak Ünal *et al.*, [30] review the growing use of AI in inventory management. Dereci and Tuzkaya [31] connect explainable predictive maintenance with an early concept of spare-parts management; this should not be treated as validation of a complete inventory optimisation system. Building on this connection, the proposed framework would present the drivers of a recommendation, the associated asset or

intervention, and the estimated effects of alternative decisions. Predictive explanations alone do not establish causality or quantify an inventory trade-off: service and cost comparisons require an explicit decision model and stated assumptions.

The fifth point is integration. Reviews of smart manufacturing place AI within a broader technological and organisational setting [29], [32]. For the proposed spare-parts architecture, processing could be distributed between equipment-level systems and enterprise applications according

to latency, connectivity and computational requirements. Item master data, purchase orders and stock balances should retain an identified authoritative system. The decision layer must reconcile different update frequencies: a frequent condition alert should update an asset-risk estimate, not automatically create a new purchase order. Replenishment recommendations should also account for open orders, reservations, lead time, criticality and planned maintenance. This is a proposed integration principle, not a demonstrated deployment result.

**Table I: Evidence Streams Supporting AI-Driven Spare-Parts Decision Support**

| Evidence stream                           | Representative publications                                                                                                                                                    | Decision contribution                                            | Main boundary                                            |
|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|----------------------------------------------------------|
| Intermittent and low-history forecasting  | Pinçe <i>et al.</i> , [1]; Jiang <i>et al.</i> , [2]; Fan <i>et al.</i> , [3]; Zhang <i>et al.</i> , [9]; Affonso <i>et al.</i> , [17]; Caserta and D'Angelo [22]              | History-based, transfer, deep-learning and resampling options    | Sparsity, cold start, item heterogeneity                 |
| Machine-learning and ensemble forecasting | Abbasimehr <i>et al.</i> , [16]; Yang <i>et al.</i> , [23]; Zhuang <i>et al.</i> , [24]; Chien <i>et al.</i> , [8]; Dodin <i>et al.</i> , [18]; Ifraz <i>et al.</i> , [25]     | Forecasting using learned patterns and, where available, context | Different industrial settings and data availability      |
| Predictive maintenance and prognostics    | Dalzochio <i>et al.</i> , [4]; Zonta <i>et al.</i> , [5]; Silvestri <i>et al.</i> , [19]; Rodrigues and Yoneyama [11]                                                          | Condition signals and examples of inventory linkage              | Failure information, asset mapping and model assumptions |
| Digital twins                             | van Dinter <i>et al.</i> , [6]; Chen <i>et al.</i> , [7]; Feng <i>et al.</i> , [20]                                                                                            | Asset-state representation and maintenance decisions             | Model complexity and integration requirements            |
| Optimisation and reinforcement learning   | Perez <i>et al.</i> , [10]; del Real Torres <i>et al.</i> , [26]; Li <i>et al.</i> , [27]; Wesendrup and Hellingrath [12]; Burtea and Tsay [13]; Alginahi <i>et al.</i> , [28] | Constrained inventory and adaptive planning policies             | Environment fidelity, reward design and transfer         |
| Explainability and enterprise integration | Sahoo and Lo [29]; Albayrak Ünal <i>et al.</i> , [30]; Dereci and Tuzkaya [31]; Gao <i>et al.</i> , [32]                                                                       | Context for understandable, operational recommendations          | Explanation is not proof of decision quality             |

Table I groups the literature into six overlapping evidence streams. Statistical and intermittent-demand methods provide benchmarks; machine-learning and ensemble approaches provide additional forecasting options; predictive maintenance and digital twins contribute asset-state information; and optimisation or reinforcement learning can connect information to actions. Existing integration is evident in Chien *et al.*'s [8] forecasting and monitoring framework, Rodrigues and Yoneyama's [11] prognostics-based stock control, Wesendrup and Hellingrath's [12] joint planning model, and Dereci and Tuzkaya's [31] explainability work. The present contribution is to bring model routing, reconciliation of demand signals, inventory constraints and human authority into one modular design with an implementation pathway. It is not a claim to be the first integrated architecture, and the relative value of this combination remains to be tested.

For potential Saudi applications, the framework is designed to accommodate mixed equipment generations, separate planning systems and domestic or international suppliers where these conditions are present. These are application assumptions, not measured characteristics of all Saudi factories. A staged implementation would allow each plant to establish its own data readiness and integration needs. Criticality should remain explicit: a seldom-used component serving a production bottleneck should not be downgraded solely because its forecast demand is low. The proposed division of responsibilities is to use prediction to support anticipation, optimisation to expose trade-offs, and accountable human judgement for exceptional or high-consequence decisions.


## V. PROPOSED PREDICTIVE DECISION-SUPPORT FRAMEWORK

The proposed framework has five layers, as illustrated in Fig. 1. CPMAl-based governance spans these layers: its phases organise project decisions over time, while the layers describe the technical functions of the operating system. The first layer consists of the industrial data foundation and includes ERP item and transaction records, warehouse balances, purchase orders, supplier performance, maintenance-management data, asset hierarchies, bills of material and, where available, condition or sensor information. The main design principle is entity linkage; a forecast must be traceable from an item to the assets which use it, to the maintenance activities that cause replacement, to the suppliers that restock it and to the sites where inventory is held. Without these relationships, statistical associations may be difficult to translate into justified inventory actions. Here, ERP denotes enterprise resource planning, CMMS/EAM computerised maintenance management or enterprise asset management, and WMS warehouse management systems.

Layer 2 involves carrying out context and feature engineering. The spare parts are divided according to how frequently they are in demand, their level of criticality, ease of repair, exposure to lead time, the asset family to which they belong, and their stage in the lifecycle. Maintenance signals are transformed into decision features such as the

predicted intervention window, the failure-risk band, the remaining-life range or the need for a planned overhaul. Instead of sending each item through a single model, a model-routing rule is used to select the right method; items with low demand may be handled using intermittent-demand or low-history methods, while those with abundant data can make use of gradient-boosting, ensemble, recurrent or transformer models, and items that are linked to condition can include prognostic features. This routing rule is a proposed design informed by the importance of demand characteristics and information availability [1], [22].

Layer 3 delivers predictive services through a demand distribution over the relevant planning horizon and an asset-driven requirement signal. The latter indicates how condition evidence or planned maintenance may change the timing and quantity of parts requirements. These outputs must be reconciled before optimisation: the asset signal is not an additional demand quantity to be added automatically to a forecast that already includes maintenance-related consumption. One implementation would condition a single demand model on health and maintenance features. Another would separate identifiable planned requirements from residual demand, map each requirement to an asset, component and intervention, and remove overlaps. The chosen rule, information cut-off and treatment of repeated alerts must be documented. These are proposed safeguards whose effectiveness requires testing.



Uncertainty should be retained through the decision chain. Demand should be evaluated over the replenishment lead time, or the lead time plus review

interval where appropriate, rather than an unrelated forecast horizon. Service targets and the relative consequences of shortages and excess stock should

determine the relevant forecast quantile or decision criterion. Criticality alone does not specify a universal stock level. Forecast calibration and the assumptions used to represent demand and lead-time uncertainty should be reported to the optimiser and planner.

Layer 4 involves inventory optimisation with constraints. It takes into account the reconciled demand distribution together with the stock currently on hand, open orders, minimum order quantities, repair returns, substitution rules, storage constraints, service targets and the degree of criticality of downtime. The result is not an automatic purchase order by default, but rather a set of actions: to expedite, to replenish, to transfer between sites, to repair, to defer, to substitute, to increase safety stock, to reduce excess stock or to take no action. Reinforcement learning can be incorporated at this stage only after a sufficiently representative simulation or digital environment has been evaluated and operational constraints can be enforced. Traditional optimisation and existing planning rules should remain available as benchmarks and fallbacks. Service targets should distinguish, for example, a unit fill rate from the probability of avoiding a stockout during a replenishment cycle; these measures are not interchangeable.

The fifth layer deals with decision support and learning; Fig. 2 shows a conceptual cycle of sensing, prediction, prioritisation, optimisation and learning, with human approval governing execution. Each recommendation should include its main drivers, the estimated effects on service and inventory cost, and the uncertainty and assumptions behind those estimates. The planners have the option of accepting, altering or rejecting the action, and must provide reasons for their decision. Executed actions and observed outcomes should be recorded separately from planner overrides: acceptance is not

proof that a recommendation was correct, and an override is not automatically an error. Feedback must be assessed before it is used to retrain a model. Table II lists the main decisions, the AI outputs, the operational indicators and the controls in this system.

The architecture is deliberately modular. An initial implementation could use ERP and maintenance records before adding advanced sensing. Application programming interfaces (APIs), scheduled exchanges or analytical interfaces offer possible integration routes, but compatibility depends on data definitions, interface access, latency, cybersecurity and local testing. Modularity is intended to support staged evaluation; it does not guarantee compatibility, measurable savings or successful adoption. Prediction estimates future requirements, optimisation evaluates feasible actions, and governance assigns authority for approval and automation. Section VI specifies how lifecycle approval and operational decision authority interact.

## VI. PMI-CPMAI-BASED DECISION GOVERNANCE

### A. Lifecycle Review and Evidence Requirements

PMI's CPMAI lifecycle comprises six iterative phases, shown in Table III. Phases can be revisited when evidence changes [15]. PMI also positions governance, ethical oversight and business alignment across the lifecycle, including operational monitoring and accountability [21]. The following review checkpoints, role assignments and execution rules are this paper's proposed application to spare-parts management, rather than a PMI-prescribed inventory-control scheme. They supplement plant procedures and do not constitute certification or a claim of compliance.

**Table II: Operational Design of the Proposed Predictive Decision-Support Framework**

| Layer                           | Core decision                           | Inputs                                               | Output                                          | Evaluation                                     | Control                                           |
|---------------------------------|-----------------------------------------|------------------------------------------------------|-------------------------------------------------|------------------------------------------------|---------------------------------------------------|
| Data foundation                 | Can records be linked and trusted?      | ERP, CMMS/EAM, WMS, suppliers, sensors               | Linked data and quality flags                   | Completeness and linkage errors                | Ownership, access, lineage                        |
| Context and features            | Which model route fits?                 | Intermittency, criticality, repairability, lifecycle | Part class and candidate route                  | Route performance and exceptions               | Engineering-approved rules                        |
| Predictive services             | What is needed, and when?               | History, context, health signals                     | Reconciled demand distribution and asset signal | Calibration and decision-relevant error        | Time cut-offs, deduplication, drift               |
| Inventory optimisation          | Which feasible action meets objectives? | Stock, orders, lead time, costs, constraints         | Replenish, transfer, repair, expedite or defer  | Defined service target, excess, emergency buys | Constraint and scenario checks                    |
| Governed execution and learning | Approve or automate?                    | Recommendation, uncertainty, explanation             | Authorised action and outcome record            | Cost, service, downtime, overrides             | CPMAI-aligned release, authority, audit, fallback |

The proposed checkpoints are evidence-based decisions to proceed, revise, hold or stop an iteration. They are not a one-way sequence: a supplier disruption may reopen data understanding and evaluation, while a change in service objectives may reopen business understanding. The business case applies across all five layers; data understanding and preparation primarily govern the first two layers; development and evaluation cover predictive services and inventory policy; operationalization connects the approved system to governed execution. This many-to-many mapping prevents governance from being confined to the final approval screen.

For this adaptation, the sponsor approves the business case and residual operational risk; the operations owner is accountable for live use. Maintenance engineering determines technical criticality and substitution acceptability, procurement retains purchasing authority, and the data owner accepts data quality exceptions. The analytics team develops the models, while an evaluator outside the development role reviews the release evidence where staffing permits. An AI project manager coordinates the evidence, decisions and unresolved issues without replacing these domain authorities. Each plant should name the people holding these roles and record delegated limits before pilot operation.

## B. Decision Rights and Execution Limits

Lifecycle approval authorises a defined use case and model– policy version; it does not authorise every recommendation that version produces. Each operational action must also pass checks on current data quality, applicability, uncertainty, engineering constraints and delegated authority. The proposed policy distinguishes three modes. In advisory mode, recommendations are reviewed without automatic transaction release. In bounded automation, only explicitly approved routine actions may execute within item, quantity, expenditure, service-risk and validity limits. In escalation mode, high-consequence actions or breached limits require the designated engineering, procurement or sponsor decision before execution.

The pilot should begin in shadow mode, followed by planner-approved operation. Bounded automation is considered only after evaluation demonstrates that the intended limits are workable. Safety-critical substitutions, reductions in insurance-spare coverage and actions outside validated conditions remain subject to authorised human procedures. An inexpensive part is not necessarily low risk: consequence of unavailability and reversibility matter alongside purchase value. Plants must set quantitative thresholds from their risk assessment and validation evidence; this paper does

not prescribe universal values. Missing authorisation, stale inputs or a failed control should block automatic release and route the case to the responsible person and an approved fallback procedure.

## C. Decision Records, Monitoring and Reapproval

Each recommendation should generate a decision record linking the item, asset and site to the data snapshot, forecast horizon, model and policy versions, uncertainty, feasible alternatives and expected service and cost effects. The record should preserve the triggered authority rule, approver or approved automation policy, timestamp, override rationale, executed transaction and subsequently observed outcome. Permissions and change history should protect this evidence from untraceable alteration. Expected benefits and realised outcomes must remain distinguishable, and emergency overrides must retain their author, justification and review status.

The operations owner should monitor service failures, emergency procurement, excess inventory, calibration, data freshness, constraint breaches and override patterns at an agreed cadence. A serious control breach should suspend the affected automatic actions immediately. Persistent deterioration, a new asset population, a supplier-regime change or a revised optimisation objective should trigger review of the relevant CPMAL phases. Retraining alone does not restore execution authority: changed models and policies must repeat the applicable evaluation and release decisions. Rollback or manual planning should remain available, and feedback should be checked for bias introduced by prior approvals, stockouts or maintenance interventions before reuse. These controls are hypotheses about accountable operation to be evaluated in the proposed pilot.

## VII. IMPLEMENTATION IN SAUDI ARABIAN INDUSTRIES

Implementation in a Saudi manufacturing plant should commence by identifying a clearly defined operational problem rather than by launching a company-wide AI programme. The four capability stages below sit within repeated CPMAL lifecycle iterations; they do not replace the six phases. Before data preparation, the sponsor should approve the pilot charter, baseline, decision owners and acceptance criteria in Table III. A suitable pilot project would be a set of critical spares that phases; lifecycle approval defines the permitted scope of execution. are associated with production bottlenecks, the high cost of emergency purchases, or frequent stockouts. The first stage establishes data readiness; this means getting rid of duplicate item codes, standardising the units of measure, linking the spares to the assets, correlating the maintenance and

warehouse records, and establishing the lead-time distributions based on evidence from the purchase orders. This stage may identify obsolete stock, missing bill-of-material links, and inconsistent planning parameters. Any resulting benefit should be

recorded against an agreed baseline rather than assumed. The data owner should approve readiness or document a remediation decision before model development proceeds.

**Table III: proposed spare-parts governance checkpoints mapped to the CPMAl lifecycle**

| CPMAI phase                 | Proposed governance question                                                   | Evidence required for the spare-parts pilot                                                                                                                            | Proposed accountable approver                                          |
|-----------------------------|--------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| 1. Business understanding   | Is the decision worth supporting with AI, and what authority may be delegated? | Pilot item and asset scope; current-policy baseline; service and cost objectives; criticality rules; budget; prohibited actions; agreed acceptance criteria.           | Plant sponsor, with engineering and finance review.                    |
| 2. Data understanding       | Can available records support the intended decision?                           | Source and owner register; access permissions; item-asset linkage; reporting latency; lead-time coverage; missingness and stockout-censoring assessment.               | Data owner, with maintenance and procurement review.                   |
| 3. Data preparation         | Are decision-time inputs traceable and fit for evaluation?                     | Versioned dataset; reconciled units and item codes; maintenance-demand deduplication; lineage; temporal splits; leakage checks; accepted quality exceptions.           | Data owner, supported by the data engineering lead.                    |
| 4. Model development        | Is the candidate reproducible and consistent with approved constraints?        | Forecast and optimiser versions; baseline comparisons; feature rationale; uncertainty representation; constraint tests; stated limits for sparse or unseen items.      | Analytics lead for submission to evaluation; no live-action authority. |
| 5. Model evaluation         | Does the whole decision policy satisfy agreed release criteria?                | Holdout and stress-test results; service, cost and critical-stockout measures; explanation review; shadow-mode observations; residual-risk assessment.                 | Plant sponsor, following validation and engineering review.            |
| 6. Model operationalization | Can the authorised policy operate within a monitored and reversible scope?     | Approved execution limits; named approvers; access and audit checks; monitoring triggers; incident procedure; tested fallback; review and retirement responsibilities. | Operations owner within sponsor-approved limits.                       |

#### Closed-Loop Decision Cycle for Smart Spare-Parts Management




**Fig. 2: Operational decision cycle subject to CPMAl-based lifecycle governance and human accountability. The five operational steps are not the six CPMAl**

In the second stage, demand segmentation and benchmark forecasting are carried out. The traditional intermittent-demand methods must be kept as comparisons, while the machinelearning models are tested on items which have sufficient contextual data. Assessment should include operational measures as well as statistical accuracy, such as critical-stockout frequency, service level, excess inventory, emergency procurement, planner overrides and downtime caused by part unavailability. A complex model should not be adopted solely because it reduces forecast error if it also increases unacceptable service risk or implementation cost. Evaluation should use chronological holdouts or rolling forecast origins, with only information available at the decision date. Recorded withdrawals should be checked for stockout censoring, since fulfilled demand can understate actual requirements. Benchmark and candidate policies should be compared under the same lead-time, cost and service assumptions. The evaluation checkpoint should document which criteria passed, which remain unresolved and whether shadow operation or a limited live pilot is authorised.

In the third stage, maintenance intelligence is incorporated. Plants that have the capability for condition monitoring can slowly introduce health indicators, beginning with the assets whose failure modes and replacement parts are well understood. When sensors are in short supply, information from work orders, the amount of time the equipment has been in operation, preventive maintenance schedules and inspection results can serve as useful stand-ins. A linked asset-state model can be introduced before considering a more detailed digital twin, with its update mechanisms and intended decisions stated explicitly. This staged design is informed by the broader discussion of manufacturing integration [29], [32]. Adding condition signals changes the evidence base and should reopen the relevant data, development and evaluation reviews before it changes live recommendations.

At the fourth stage, cross-site optimisation and controlled automation may be considered after the pilot has met agreed service, cost and reliability criteria; in the beginning, recommendations will need the approval of the planner, while the thresholds for automatic execution will be limited to low-risk decisions. Actions of greater consequence—such as reducing the stock of insurance spares or replacing a safety-critical component—should remain subject to the plant's authorised engineering and procurement procedures. Cybersecurity, access control, model versioning and audit trails must be treated as operational requirements from the outset, not be incorporated afterwards. The operations owner

should confirm the execution limits, incident response and fallback test before enabling bounded automation, and each additional site should undergo its own applicability review.

The pathway is consistent with the technology-adoption direction described for the Future Factories programme [14]. Its potential value lies in improved asset availability, lower inventory exposure and more disciplined sourcing decisions, but these are outcomes to evaluate rather than benefits already established. For a Saudi pilot, the implementation assessment should document the sector and equipment population, local and imported procurement routes, observed lead-time distributions, system interfaces and decision authority. These characteristics would determine whether the generic architecture needs adaptation and provide a basis for evaluating transfer between plants.

## VIII. RESEARCH AGENDA AND PRAGMATIC IMPLICATIONS

Future research should focus on integration rather than on separate model competitions. To begin with, empirical studies should compare condition-informed demand forecasting with strong baselines for intermittent demand throughout a number of plants and asset types. Secondly, probabilistic forecasts should be assessed by looking at the results of inventory management, not just by using error measures. Thirdly, reinforcement-learning policies need digital environments that are realistic and which reflect lead-time disruptions, repair loops, substitutions, and maintenance dependencies. Fourthly, research into explainability should assess whether the explanations lead to better decisions by planners or simply whether they meet the technical standards of interpretability. Fifthly, transfer learning and low-data methods warrant investigation where critical spares have too few events for reliable itemlevel supervised learning. Tests should also examine whether maintenance signals provide information beyond demand history and whether the reconciliation rules prevent duplicate requirements. The modular framework enables such ablation comparisons, but no results from these comparisons are claimed here.

For practitioners the main conclusion is that spare-parts optimisation involves a socio-technical system in which data ownership, maintenance engineering, procurement, warehouse operations and IT all have to agree on the meanings of criticality, service risk and acceptable automation. A successful programme should therefore have joint governance, include the clear tracking of benefits and incorporate a model-monitoring procedure which is able to pick

up any drift caused by changes to the equipment, the suppliers or the maintenance policies.

The governance extension creates an additional evaluation question: does explicit decision authority and evidence-based release improve decision quality relative to existing planning governance? A prospective pilot could compare the current process, AI recommendations under existing approvals, and AI recommendations under the proposed CPMAI-based controls, using comparable item groups and phased introduction. All arms should retain mandatory plant safety procedures. Alongside service and inventory outcomes, evaluation should measure approval latency, traceability completeness, unauthorised action attempts, override reasons, incident-response time and the effort required to maintain the controls. A simple before– after gain would not isolate the governance contribution from data cleansing, model changes or demand conditions.

## IX. LIMITATIONS

This structured narrative review is bounded by a selected corpus of 29 journal publications from 2020 to 2025 and is not an exhaustive account of the field. The purposive retention rationale in Section III clarifies how the 29 publications contribute to the synthesis, but it does not reconstruct an auditable search and screening record; assessment of selection bias and completeness therefore remains limited. The journal and date boundaries may omit relevant conference, industrial and earlier foundational work; DOI availability should not be confused with peer review or research quality. Reviews and primary studies overlap, and their findings are not independent observations suitable for pooling. Access to some original sources was limited to abstracts or publisher previews, restricting detailed assessment of experimental design and limitations. The study-level map is therefore an interpretive evidence map, not a substitute for a formal full-text quality appraisal.

The framework is conceptual and has not been validated using Saudi plant data, stakeholder assessment or a controlled pilot. The Saudi policy reference establishes relevance but not local technical suitability. The PMI sources provide professional methodology guidance, not independent evidence that the proposed controls improve spare-parts outcomes. The adaptation has not been assessed for CPMAI conformity or endorsed by PMI. Its role assignments, review burden, approval delays and automation limits need stakeholder testing; their transferability will depend on plant size and available expertise. Proposed model-routing rules, demand reconciliation, explanations and automation controls require operational testing. Future research should compare the framework with current planning

policies and simpler alternatives under consistent cost and service assumptions, accounting for maintenance interventions, censored demand and differences in digital maturity.

## X. CONCLUSION

AI can support spare-parts management when forecasting is connected to operational decisions that affect availability and inventory cost. The selected literature from 2020 to 2025 covers intermittent-demand forecasting, machine learning, predictive maintenance, digital twins, reinforcement learning and explainability. It includes integrated approaches as well as studies focused on individual tasks. This paper addresses the practical design problem of connecting those capabilities with enterprise data and accountable execution; it does not establish that integration is absent from prior research.

The review proposes a five-layer predictive decision-support framework connecting industrial data, context and feature engineering, predictive services, constrained inventory optimisation, and governed decision support and learning. By retaining uncertainty and reconciling maintenance signals with forecast demand, the design aims to support a staged progression from ERP-based analysis to condition-informed planning. The CPMAI-based governance extension connects lifecycle evidence to named decision owners, bounded execution authority, traceable actions and reapproval when conditions change. These are proposed operating controls whose practical value must be tested alongside model performance. For Saudi manufacturers, it offers a framework for pilot design and evaluation within the industrial modernisation context. Its feasibility, compatibility and economic value must be established in the plants where it is applied.

The main conclusion is that the most valuable AI system is not the one which has the most complex model, but rather the one that turns reliable evidence into timely, explainable and economically defensible decisions regarding spare parts. Research and implementation should therefore be evaluated in terms of service performance, exposure to downtime, inventory efficiency and the quality of decisions rather than on the basis of forecast precision alone.

## APPENDIX A STUDY-LEVEL EVIDENCE MAP

This map identifies the contribution used from each of the 29 retained research publications. The final column states the boundary applied in this synthesis; it is not a claim that every limitation was reported by the original authors. The entries do not constitute a formal quality score or an independent replication. Full bibliographic details appear in the reference list.

**Table IV: Study-Level Evidence Map**

| <b>Publication</b>                   | <b>Evidence and setting</b>                                                                                                 | <b>Contribution and boundary of inference</b>                                                                                 |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Abbasimehr <i>et al.</i> , [16]      | Demand forecasting; furniture-company case; optimised multilayer LSTM.                                                      | Adjacent forecasting evidence; does not validate intermittent spare-parts decisions or Saudi transfer.                        |
| Affonso <i>et al.</i> , [17]         | Spare-parts forecasting; heuristics combined with bootstrapping; iron-ore company case.                                     | Supports a hybrid statistical route; do not classify bootstrapping alone as machine learning.                                 |
| Albayrak Ünal <i>et al.</i> , [30]   | Systematic review of AI applications in inventory management.                                                               | Background on method families; not an independent test of this architecture.                                                  |
| Alginahi <i>et al.</i> , [28]        | Bibliometric review of reinforcement learning in industrial automation.                                                     | Background on deployment and research directions; no review percentages are used as evidence of this framework's performance. |
| Burtea and Tsay [13]                 | Constrained continuous-action reinforcement learning for supplychain inventory decisions.                                   | Supports explicit feasibility constraints; applicability remains tied to the modelled environment.                            |
| Caserta and D'Angelo [22]            | Forecasting framework for spare parts with little or no demand history.                                                     | Informs low-history routing; not proof that any critical spare can be forecast accurately.                                    |
| Chen <i>et al.</i> , [7]             | Review of machine learning in digital twins for predictive maintenance.                                                     | Supports asset-state and maintenance integration; inventory-policy benefits require separate testing.                         |
| Chien <i>et al.</i> , [8]            | Aftermarket forecasting framework incorporating product classification, ensemble methods and monitoring.                    | An existing integrated forecasting approach; the present contribution must be distinguished from it.                          |
| Dalzochio <i>et al.</i> , [4]        | Review of machine learning, reasoning and integration in predictive maintenance.                                            | Provides maintenance and integration context; failure prediction does not alone specify a parts order.                        |
| del Real Torres <i>et al.</i> , [26] | Review of deep reinforcement learning for smart manufacturing.                                                              | Motivates adaptive decision policies and environment models; not spareparts deployment validation.                            |
| Dereci and Tuzkaya [31]              | Explainable predictive maintenance with an early concept of spareparts management.                                          | Supports intelligible recommendations; does not establish a fully validated inventory optimiser.                              |
| Dodin <i>et al.</i> , [18]           | Deployed Bombardier aftermarket forecasting pipeline combining machine learning, time-series models and ensemble forecasts. | Demonstrates an industrial forecasting application; aviation results are not assumed to transfer unchanged.                   |
| Fan <i>et al.</i> , [3]              | Intermittent feature adaptation and transfer across demand series; two equipment-manufacturer datasets.                     | Informs cross-series learning; transfer should be evaluated on the target item population.                                    |
| Feng <i>et al.</i> , [20]            | Digital-twin-driven multi-level maintenance decisions using optimisation and matheuristics.                                 | Supports coordinated maintenance planning; distinct from validation of this replenishment architecture.                       |
| Gao <i>et al.</i> , [32]             | Review of AI methods, manufacturing applications and research directions.                                                   | Broader manufacturing context; no plant-specific ERP compatibility is implied.                                                |
| I'fraz <i>et al.</i> , [25]          | Bus-fleet demand forecasting using regression, machine learning and maintenance-related variables.                          | Supports considering operational predictors; distinguish the fleet application from factory inventory control.                |
| Jiang <i>et al.</i> , [2]            | Adaptive univariate support-vector machine for intermittent demand; heavy-duty vehicle spare-parts case.                    | Demonstrates a history-based machine-learning route; not evidence that external covariates are required.                      |
| Li <i>et al.</i> , [27]              | Review of deep reinforcement learning across the smartmanufacturing lifecycle.                                              | Supports mapping adaptive-policy applications and challenges; not a test of the proposed framework.                           |
| Perez <i>et al.</i> , [10]           | Single-product supply-chain simulation comparing deterministic and stochastic optimisation with reinforcement learning.     | Supports policy-level comparison; demand and network assumptions bound interpretation.                                        |

| Publication                    | Evidence and setting                                                                                                                                           | Contribution and boundary of inference                                                                             |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Pinçe <i>et al.</i> , [1]      | Critical review of intermittent-demand forecasting for spare parts.                                                                                            | Supports matching methods and evaluation to demand characteristics; does not establish one universally best model. |
| Rodrigues and Yoneyama [11]    | Periodic-review, single-site and single-component inventory model for non-repairable items; prognostics and health-monitoring inputs and fill-rate constraint. | Existing prognostics–inventory integration; repair loops and cross-site decisions require extensions.              |
| Sahoo and Lo [29]              | Review of technological developments supporting smart manufacturing.                                                                                           | Provides system-level context; the modular implementation sequence here is a design proposal.                      |
| Silvestri <i>et al.</i> , [19] | Systematic review of Industry 4.0 technologies in maintenance.                                                                                                 | Supports the maintenance-transformation context; does not measure benefits of this inventory framework.            |
| van Dinter <i>et al.</i> , [6] | Systematic review of predictive maintenance using digital twins.                                                                                               | Informs representation and integration choices; data and model complexity remain constraints.                      |
| Wesendrup and Hellingrath [12] | Reinforcement-learning model integrating spare-parts sourcing, production and maintenance for a single-machine system.                                         | Existing joint decision approach; multi-asset and cross-site generalisation needs testing.                         |
| Yang <i>et al.</i> , [23]      | Scenario-based equipment spare-parts forecasting using improved feature selection and stacking.                                                                | Supports contextual feature construction; reported model comparisons do not validate local inventory outcomes.     |
| Zhang <i>et al.</i> , [9]      | Transformer forecasting tested on 925 intermittent-demand items from an airline spare-parts provider.                                                          | Supports a deep-learning candidate under sparsity; target-plant performance remains an empirical question.         |
| Zhuang <i>et al.</i> , [24]    | Automotive aftermarket demand occurrence classification combined with non-zero demand regression.                                                              | Supports separating occurrence and size; no universal transfer to manufacturing is assumed.                        |
| Zonta <i>et al.</i> , [5]      | Systematic review and taxonomy of predictive maintenance in Industry 4.0.                                                                                      | Provides predictive-maintenance context and limitations; not an inventory policy evaluation.                       |

## REFERENCES

1. Ç. Pinçe, L. Turrini, and J. Meissner, "Intermittent demand forecasting for spare parts: a critical review," *Omega*, vol. 105, Art. no. 102513, 2021, doi: 10.1016/j.omega.2021.102513.
2. P. Jiang, Y. Huang, and X. Liu, "Intermittent demand forecasting for spare parts in the heavy-duty vehicle industry: a support vector machine model," *International Journal of Production Research*, vol. 59, no. 24, pp. 7423–7440, 2021, doi: 10.1080/00207543.2020.1842936.
3. L. Fan, X. Liu, W. Mao, K. Yang, and Z. Song, "Spare parts demand forecasting method based on intermittent feature adaptation," *Entropy*, vol. 25, no. 5, Art. no. 764, 2023, doi: 10.3390/e25050764.
4. J. Dalzochio et al., "Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges," *Computers in Industry*, vol. 123, Art. no. 103298, 2020, doi:10.1016/j.compind.2020.103298.
5. T. Zonta et al., "Predictive maintenance in the Industry 4.0: a systematic literature review," *Computers & Industrial Engineering*, vol. 150, Art. no. 106889, 2020, doi: 10.1016/j.cie.2020.106889.
6. R. van Dinter, B. Tekinerdogan, and C. Catal, "Predictive maintenance using digital twins: a systematic literature review," *Information and Software Technology*, vol. 151, Art. no. 107008, 2022, doi: 10.1016/j.infsof.2022.107008.
7. C. Chen, H. Fu, Y. Zheng, F. Tao, and Y. Liu, "The advance of digital twin for predictive maintenance: the role and function of machine learning," *Journal of Manufacturing Systems*, vol. 71, pp. 581–594, 2023, doi: 10.1016/j.jmsy.2023.10.010.
8. C.-F. Chien, C.-C. Ku, and Y.-Y. Lu, "Ensemble learning for demand forecast of after-market spare parts to empower data-driven value chain and an empirical study," *Computers & Industrial Engineering*, vol. 185, Art. no. 109670, 2023, doi: 10.1016/j.cie.2023.109670.
9. G. P. Zhang, Y. Xia, and M. Xie, "Intermittent demand forecasting with transformer neural networks," *Annals of Operations Research*, vol. 339, no. 1–2, pp. 1051–1072, 2024, doi: 10.1007/s10479-023-05447-7.
10. H. D. Perez, C. D. Hubbs, C. Li, and I. E. Grossmann, "Algorithmic approaches to inventory management optimization," *Processes*, vol. 9, no. 1, Art. no. 102, 2021, doi: 10.3390/pr9010102.
11. L. R. Rodrigues and T. Yoneyama, "A spare parts inventory control model based on prognostics and health monitoring data under a fill rate constraint," *Computers & Industrial Engineering*, vol. 148, Art. no. 106724, 2020, doi: 10.1016/j.cie.2020.106724.

12. K. Wesendrup and B. Hellingrath, "Post-prognostics demand management, production, spare parts and maintenance planning for a single-machine system using reinforcement learning," *Computers & Industrial Engineering*, vol. 179, Art. no. 109216, 2023, doi: 10.1016/j.cie.2023.109216.
13. R. Burtea and C. Tsay, "Constrained continuous-action reinforcement learning for supply chain inventory management," *Computers & Chemical Engineering*, vol. 181, Art. no. 108518, 2024, doi: 10.1016/j.compchemeng.2023.108518.
14. Saudi Press Agency, "Industry minister showcases Saudi industrial vision at WEF's 'Next-Gen Industrial Infrastructure' session," Jan. 25, 2025. Accessed: Oct. 3, 2026. [Online]. Available: <https://www.spa.gov.sa/en/N2249583>.
15. K. Walch and R. Schmelzer, "6 phases to run a successful AI project," Project Management Institute, Jun. 11, 2025. Accessed: Oct. 4, 2026. [Online]. Available: <https://www.pmi.org/blog/6-stages-to-run-a-successful-ai-project>.
16. H. Abbasimehr, M. Shabani, and M. Yousefi, "An optimized model using LSTM network for demand forecasting," *Computers & Industrial Engineering*, vol. 143, Art. no. 106435, 2020, doi: 10.1016/j.cie.2020.106435.
17. T. B. Affonso, S. V. Conceição, L. R. Muniz, J. F. de Freitas Almeida, and J. C. de Lima, "A new hybrid forecasting method for spare part inventory management using heuristics and bootstrapping," *Decision Analytics Journal*, vol. 10, Art. no. 100415, 2024, doi:10.1016/j.dajour.2024.100415.
18. P. Dodin et al., "Bombardier aftermarket demand forecast with machine learning," *INFORMS Journal on Applied Analytics*, vol. 53, no. 6, pp. 425–445, 2023, doi: 10.1287/inte.2023.1164.
19. L. Silvestri, A. Forcina, V. Introna, A. Santolamazza, and V. Cesarotti, "Maintenance transformation through Industry 4.0 technologies: a systematic literature review," *Computers in Industry*, vol. 123, Art. no.103335, 2020, doi: 10.1016/j.compind.2020.103335.
20. Q. Feng et al., "Multi-level predictive maintenance of smart manufacturing systems driven by digital twin: a matheuristics approach," *Journal of Manufacturing Systems*, vol. 68, pp. 443–454, 2023, doi: 10.1016/j.jmsy.2023.05.004.
21. K. Walch and R. Schmelzer, "PMI-CPMAI: The best certification to lead AI projects," Project Management Institute, Sep. 29, 2025. Accessed: Oct. 4, 2026. [Online]. Available: <https://www.pmi.org/blog/the-best-ai-certification-to-lead-ai-projects>.
22. M. Caserta and L. D'Angelo, "Intermittent demand forecasting for spare parts with little historical information," *Journal of the Operational Research Society*, vol. 76, no. 2, pp. 294–309, 2025, doi: 10.1080/01605682.2024.2349734.
23. Y. Yang et al., "An improved stacking model for equipment spare parts demand forecasting based on scenario analysis," *Scientific Programming*, vol. 2022, Art. no. 5415702, 2022, doi: 10.1155/2022/5415702.
24. X. Zhuang, Y. Yu, and A. Chen, "A combined forecasting method for intermittent demand using the automotive aftermarket data," *Data Science and Management*, vol. 5, no. 2, pp. 43–56, 2022, doi: 10.1016/j.dsm.2022.04.001.
25. M. Ifraz, A. Aktepe, S. Ersöz, and T. Çetinyokus,, "Demand forecasting of spare parts with regression and machine learning methods: application in a bus fleet," *Journal of Engineering Research*, vol. 11, no. 2, Art. no. 100057, 2023, doi: 10.1016/j.jer.2023.100057.
26. del Real Torres, D. S. Andreiana, Á. Ojeda Roldán, A. Hernández Bustos, and L. E. Acevedo Galicia, "A review of deep reinforcement learning approaches for smart manufacturing in Industry 4.0 and 5.0 framework," *Applied Sciences*, vol. 12, no. 23, Art. no. 12377, 2022, doi: 10.3390/app122312377.
27. Li, P. Zheng, Y. Yin, B. Wang, and L. Wang, "Deep reinforcement learning in smart manufacturing: a review and prospects," *CIRP Journal of Manufacturing Science and Technology*, vol. 40, pp. 75–101, 2023, doi: 10.1016/j.cirpj.2022.11.003.
28. Y. M. Alginahi, O. Sabri, and W. Said, "Reinforcement learning for industrial automation: a comprehensive review of adaptive control and decision-making in smart factories," *Machines*, vol. 13, no. 12, Art. no. 1140, 2025, doi: 10.3390/machines13121140.
29. S. Sahoo and C.-Y. Lo, "Smart manufacturing powered by recent technological advancements: a review," *Journal of Manufacturing Systems*, vol. 64, pp. 236–250, 2022, doi: 10.1016/j.jmsy.2022.06.008.
30. Ö. Albayrak Ünal, B. Erkeyman, and B. Usanmaz, "Applications of artificial intelligence in inventory management: a systematic review of the literature," *Archives of Computational Methods in Engineering*, vol. 30, no. 4, pp. 2605–2625, 2023, doi: 10.1007/s11831-022-09879-5.
31. U. Dereci and G. Tuzkaya, "An explainable artificial intelligence model for predictive maintenance and spare parts optimization," *Supply Chain Analytics*, vol. 8, Art. no. 100078, 2024, doi: 10.1016/j.sca.2024.100078.
32. R. X. Gao, J. Krüger, M. Merklein, H.-C. Möhring, and J. Váncza, "Artificial intelligence in manufacturing: state of the art, perspectives, and future directions," *CIRP Annals*, vol. 73, no. 2, pp. 723–749, 2024, doi: 10.1016/j.cirp.2024.04.101.