Academy for Global Science and Innovation (AGSI)

ARTICLE

Machine Learning–Enabled Predictive Maintenance in Reducing Parts Shortages Across International Supply Networks

Nguyen Thanh Khoa,[1] Zhao Meilin,[2] and Pham Thi Hong Lan[3]

[1]Saigon Pacific University, Department of Computer Science, Le Duan Boulevard, Ho Chi Minh City 700000, Vietnam

[2]Shandong Bay Institute of Finance, Department of Business Administration, Haier Road, Qingdao 266000, China

[3]Mekong Frontier University, Department of Information Systems and E-Business, Hung Vuong Avenue, Can Tho 900000, Vietnam

Abstract

Globalized manufacturing and service systems depend on complex international supply networks that must provide the right spare parts at the right time to geographically dispersed maintenance sites. Traditional maintenance strategies that rely on fixed-interval schedules or corrective interventions after failure can generate lumpy, poorly forecasted demand for spare parts, increasing the risk of local stockouts and reactive emergency logistics. At the same time, advances in sensing, connectivity, and data storage have created large multivariate time series describing condition, usage, and environment for critical equipment. Machine learning techniques can exploit these data to estimate the probability of failure or remaining useful life, offering the possibility of transforming maintenance demand for spare parts from a largely stochastic phenomenon into a partially controllable process. This paper examines how machine learning enabled predictive maintenance can be mathematically integrated with multi echelon international spare parts networks to reduce parts shortages. The study develops a linear algebraic representation linking failure risk predictions, maintenance scheduling decisions, and inventory state across locations, and combines this representation with probabilistic machine learning models for component degradation. A simulation based evaluation with stylized yet realistic network structures illustrates how the coupling between predictive models and inventory control policies influences stockout frequency, fill rate, and inventory investment under varying lead times and forecast error profiles. The discussion highlights conditions under which predictive maintenance most strongly affects spare parts availability, emphasizes the role of uncertainty quantification in the predictions, and outlines practical considerations for embedding such models into existing planning processes.

1. Introduction

International supply networks for spare parts support fleets of high value assets such as aircraft, power generation units, mining equipment, and complex manufacturing tools deployed across multiple regions with heterogeneous operating conditions [1]. The reliability of these networks is often measured by the ability to supply the required part when a maintenance task is initiated, a property closely related to fill rate and stockout frequency. Traditional planning methods treat maintenance induced spare parts demand as an exogenous stochastic process estimated from historical failure rates and preventive schedules. This exogenous view can be a limitation when operating environments are highly variable, when fleets are heterogeneous, and when risk tolerances for unplanned downtime differ across operators, since all

of these factors influence the true demand for spare parts in ways that may not be visible in aggregated historical averages.

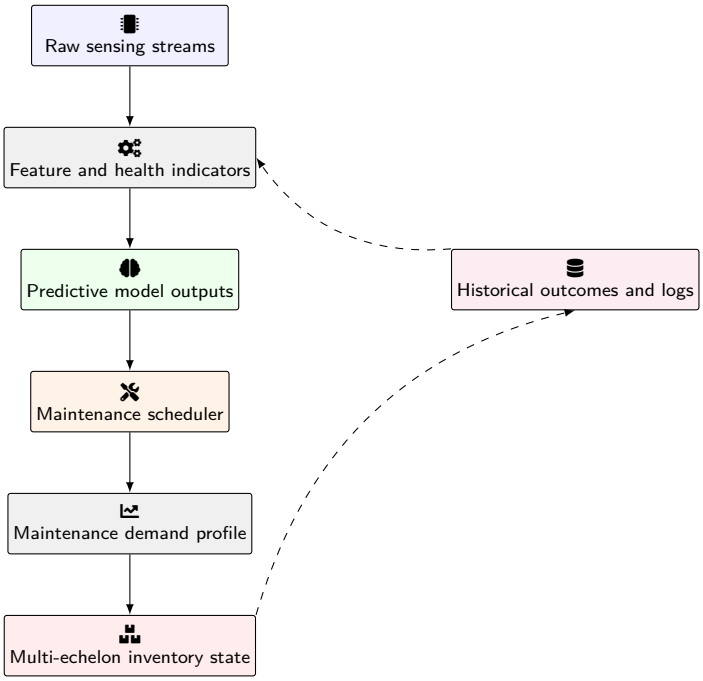


Figure 1. End-to-end flow from raw sensing streams to multi-echelon inventory state. Sensor data are transformed into health indicators, fed into predictive models, transformed into maintenance schedules, and then aggregated into a demand profile that drives inventory evolution. Historical outcomes and transactional logs feed back into the feature construction and model training stages to refine predictions over time.

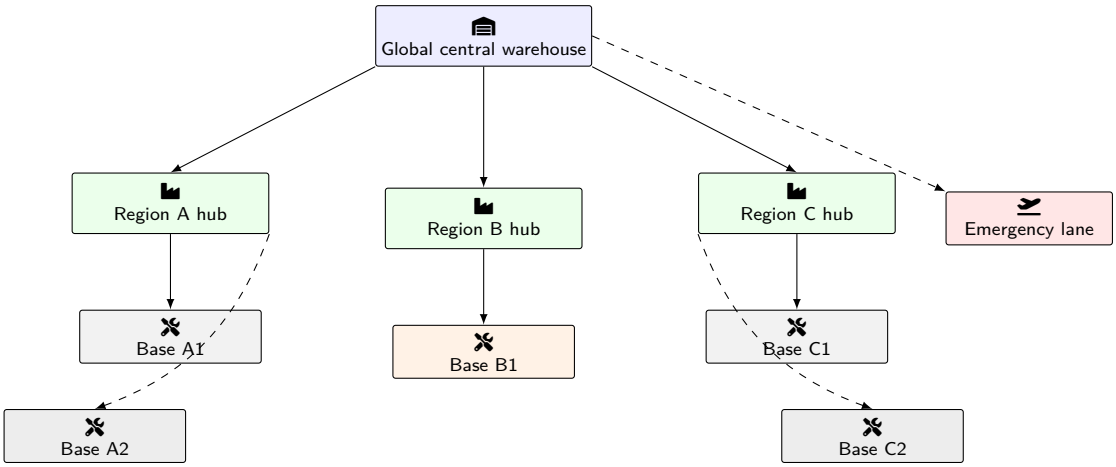


Figure 2. Stylized international spare parts network with a global warehouse, regional hubs, and local maintenance bases. Vertical links illustrate regular replenishment flows, while dashed arcs capture secondary flows such as alternative base allocations and expedited shipments. The configuration highlights how shortages or surges at specific bases propagate only through a limited number of key links rather than across all locations simultaneously.

Predictive maintenance leverages condition monitoring, on board sensors, and operational logs to estimate the evolving health state of components and to forecast their probability of

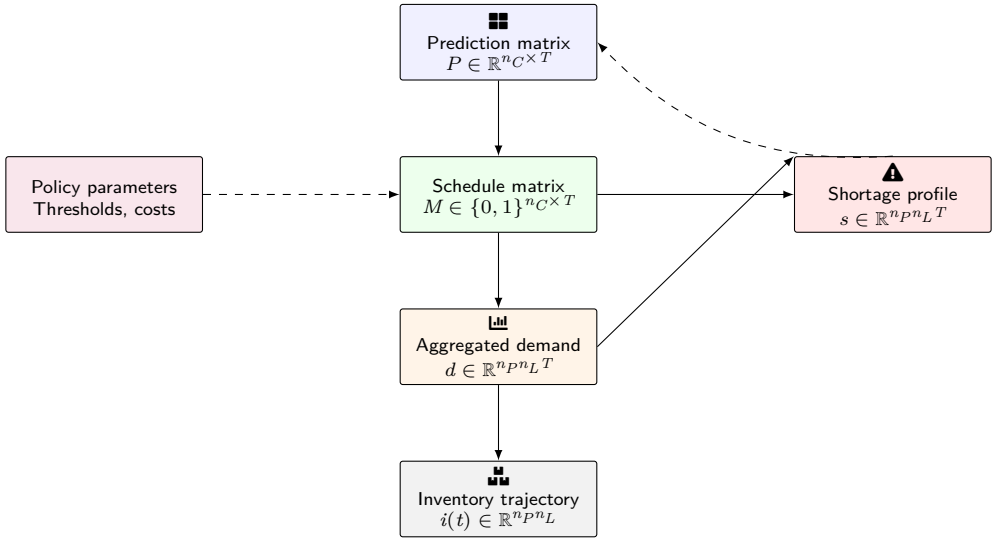


Figure 3. Matrix-based view of the coupling between predictive maintenance and inventory. The prediction matrix encodes component-level risk over the horizon, the schedule matrix represents maintenance decisions, and the aggregated demand vector drives inventory trajectories at partlocation level. Shortage profiles arise as a function of schedules and demand, while policy parameters such as thresholds and cost weights influence schedules. Feedback from observed shortages can be used to refine predictive models and decision rules.

failure over a future horizon. Machine learning methods, including tree based ensembles and deep learning architectures, have expanded the range of degradations and operating regimes that can be modeled using high dimensional data. These methods can produce both point predictions for metrics such as remaining useful life and probabilistic forecasts of failure events. If such predictive information is integrated into planning processes, maintenance tasks can be scheduled proactively in time windows that respect both operational requirements and supply constraints. In principle, this partial control over the timing and distribution of maintenance tasks could attenuate demand peaks for critical parts and reduce the likelihood of shortages across the network [2].

However, the benefits of predictive maintenance for spare parts availability are not automatic. Predictions are subject to model error, distribution shifts, and missing data, and maintenance decisions must respect constraints such as limited workshop capacity, regulatory requirements, and contract terms. International supply networks further introduce long and uncertain transportation lead times, customs delays, currency fluctuations, and regional safety stock policies that may differ across business units. A predictive maintenance policy that schedules large numbers of tasks in a narrow window in response to elevated predicted risk could unintentionally generate localized demand spikes, exhausting available stock and pushing the network into reactive expedites. The relationship between predictive maintenance and parts shortages is therefore mediated by both the quality of the predictions and the structure of the supply network.

This work examines this relationship through a modeling framework that connects machine learning based failure predictions with multi echelon inventory states using linear algebra. The framework represents equipment health features, predicted failure probabilities, maintenance scheduling decisions, and part flows between locations as vectors and matrices. This representation supports analysis of how perturbations in predictions or policy parameters propagate through the network and alter the expected shortage profile [3]. Within this framework,

machine learning models are treated as mappings from high dimensional feature spaces into parameters governing demand distributions, while inventory and replenishment dynamics are captured through linear and piecewise linear transformations.

The remainder of the paper develops this idea in several steps. After outlining the structure of international supply networks and the role of maintenance induced demand, the paper describes data flows and feature engineering strategies for constructing learning ready representations from heterogeneous sensor and event streams. It then surveys machine learning architectures appropriate for predicting failure risk at the component level and discusses calibration and uncertainty estimation. Building on these elements, the paper formulates a linear algebraic model that links predicted risks, maintenance decisions, and spare part inventory across locations and over time. The model is used to define optimization problems and simulation experiments that examine parts shortages under different policy and prediction scenarios. Implementation aspects and limitations, including data quality, organizational integration, and scaling to large fleets, are discussed before concluding.

2. International Supply Networks and Maintenance Induced Demand

Table 1. Representative data sources feeding machine learning enabled predictive maintenance.

Data source	*Typical content*	*Temporal resolution*
On-board sensors	Vibration, temperature, pressure, speed traces	Milliseconds to seconds
Maintenance logs	Work orders, failure codes, replaced parts	Events with timestamps
Inventory transactions	Receipts, issues, stock adjustments	Hours to days
Operational records	Duty cycles, load profiles, mission plans	Minutes to hours
Environmental feeds	Ambient temperature, humidity, dust indices	Hours to days

Table 2. Illustrative feature groups extracted for degradation and failure modelling.

Feature group	*Construction*	*Example indicators*	*Typical size*
Time-domain statistics	Window-based aggregation	Moving mean, variance, peak-to-peak	10–40 per channel
Spectral descriptors	Transform-domain analysis	Band power, spectral centroid, kurtosis	20–80 per channel
Usage and load	Aggregated utilization	Cumulative cycles, load spectrum bins	10–30 per asset
Environment encoded	Normalized external drivers	Temperature z-scores, humidity indices	5–15 per region
Event history	Encoded maintenance trace	Time since last overhaul, intervention flags	10–25 per component

International spare parts networks are typically organized as multi echelon structures comprising central distribution centers, regional hubs, and local stocking points co located with maintenance facilities or customer sites. Each echelon interacts through transportation links with distinct lead time distributions. Let there be a set of locations indexed by the finite set $L = \{1, \ldots, n_L\}$ and a set of parts indexed by $P = \{1, \ldots, n_P\}$. For each part $p \in P$ and location $l \in L$, the network maintains stock $I_{l,p}(t)$ at discrete time t, and there are replenishment flows $Q_{l,p}(t)$ that result from previously placed orders. Transportation between locations is constrained by lead times and capacity, and cross border shipments are subject to additional variability related to customs and regulatory inspections.

Table 3. Example model classes used in predictive maintenance pipelines.

Model family	*Target quantity*	*Strengths*	*Typical deployment role*
Regularized logistic model	Per-horizon failure probability	Simple, interpretable coefficients	Baseline risk scoring and calibration anchor
Gradient boosted trees	Failure classification or RUL quantiles	Flexible non-linear interactions	Fleet-scale deployment with tabular features
Temporal convolutional network	Sequence-to-risk mapping	Handles long effective context	End-to-end modelling of sensor streams
Discrete-time survival model	Hazard over time bins	Direct survival-time representation	Remaining useful life prediction with censoring
Domain-adapted encoder	Latent health representation	Shared structure across regions	Transfer of knowledge between fleets and climates

Table 4. Hierarchical levels in the international maintenance and inventory model.

System level	*Index set*	*State vector example*
Component level	$c = 1, \ldots, n_C$	Health indicators, predicted risk, installation age
Asset level	$a = 1, \ldots, n_A$	Active components, operational status, mission queue
Location level	$l = 1, \ldots, n_L$	Local inventory, capacity, scheduled work orders
Part type level	$p = 1, \ldots, n_P$	Global demand rate estimates, lead time statistics
Network level	Graph over L	Flow constraints, transshipment options, pooling structure

Demand for spare parts arrives from maintenance activities triggered by different policies. Corrective maintenance occurs after a failure event and generates immediate demand for the failed part and often for associated consumables. Time based preventive maintenance schedules tasks at fixed intervals, generating more predictable demand patterns but sometimes replacing components that have substantial remaining useful life. Condition based maintenance monitors specific measurements such as vibration amplitude, temperature, or fluid properties and triggers interventions when thresholds are exceeded. Predictive maintenance moves beyond threshold based rules by estimating the probability that a component will fail over a future horizon or the remaining time to failure, potentially allowing maintenance timing to be adjusted flexibly within operational constraints.

In many real networks, maintenance induced demand is partially planned and partially unplanned. Planned demand arises from scheduled preventive tasks and known overhaul events, which can be encoded as deterministic time series. Unplanned demand arises when a failure occurs earlier or later than expected or when inspections detect unexpected degradation. Let $D_{l,p}(t)$ denote the random demand for part p at location l and time t, decomposed into $D_{l,p}(t) = D_{l,p}^{\mathrm{plan}}(t) + D_{l,p}^{\mathrm{unplan}}(t)$. Traditional inventory control methods estimate the distribution of $D_{l,p}^{\mathrm{unplan}}(t)$ using historical averages or parametric failure models assuming stationary conditions and independence across assets. In international contexts, these assumptions are often weak, since operating conditions vary by region, fleets differ in age and configuration, and external factors such as ambient temperature or duty cycle patterns shift over time.

Predictive maintenance has the potential to reduce the variance of $D_{l,p}^{\mathrm{unplan}}(t)$ by replacing some of the uncertainty with forecastable structure. By observing sensor data and usage histories, machine learning models can estimate for each component instance its current degradation state and risk profile. These estimates can be aggregated over all assets served by

Table 5. Main cost elements in the joint maintenanceinventory objective.

Cost component	*Symbolic representation*	*Driver*	*Planning effect*
Maintenance effort	J_{maint}	Scheduled interventions and labor	Encourages temporal smoothing of tasks
Shortage and backorder	J_{short}	Unfilled demand and delay	Drives safety stock and buffer allocation
Unplanned failure risk	J_{risk}	Residual failure probabilities	Promotes early replacement for high-risk items
Holding and logistics	Embedded in inventory terms	Inventory depth and transport usage	Balances central versus regional stocking
Policy change overhead	Implicit penalty on adjustments	Reconfiguration of plans and shifts	Moderates aggressive rescheduling moves

Table 6. Comparison of maintenance policies and their qualitative impact on parts demand.

Policy type	*Trigger*	*Demand pattern*	*Inventory implication*
Reactive corrective	Observed failure	Highly spiky and uncertain	High expedite usage, sensitive to lead times
Time-based preventive	Calendar or usage threshold	Regular but conservative	Larger average stock, lower emergency moves
Simple condition-based	Threshold on single indicator	Local smoothing around threshold	Depends on indicator stability and noise
Predictive risk-based	Multi-feature risk estimate	Tunable smoothing via thresholds	Joint optimisation with base-stock levels
Hybrid region-specific	Combined local rules	Mixed regular and reactive demand	Region-tailored safety stock and flows

a given location to form predictive distributions of near term demand. Furthermore, maintenance planners can decide to advance or defer specific tasks to smooth anticipated peaks. However, predictive information is imperfect, and false positives or negatives translate into deviations between forecasted and realized demand. When many sites are coupled through shared central stocks, correlation in prediction errors across sites can influence the joint shortage profile [4].

The international dimension adds several layers to this interaction. First, long replenishment lead times from central warehouses to distant regions make local safety stock policies more sensitive to forecast error. Second, differences in regulatory or contractual requirements for maximum allowable risk can lead to heterogeneous maintenance policies across regions, even when using similar predictive models. Third, currency and cost differences influence tradeoffs between holding inventory locally versus relying on cross shipments or emergency logistics. A model that aims to capture the effect of machine learning enabled predictive maintenance on parts shortages must therefore represent not only local demand modulation but also the propagation of that demand through a network with diverse constraints [5].

In the modeling approach developed later in the paper, the international network is represented as a directed graph with nodes corresponding to locations and edges to transportation links. Lead times, shipping capacities, and cost parameters are encoded in matrices and vectors, while the maintenance induced demand process is represented as a function of predicted failure risks at the level of individual components. This representation allows the dependencies between predictive maintenance decisions at one site and stock availability at other sites to be expressed as linear transformations, facilitating analysis and simulation of shortage patterns under different scenarios.

Table 7. Selected parameters for sensitivity analysis in simulation experiments.

Scenario parameter	*Variation range*	*Primary metric affected*	*Secondary impact*
Lead time mean	Short to long	Fill rate and backorder depth	Emergency transport frequency
Lead time variability	Low to high	Shortage volatility	Required safety stock margins
Prediction discrimination	Low to high	Avoided failures	Peaks in scheduled demand
Workshop capacity	Tight to flexible	Fraction of tasks deferred	Temporal clustering of orders
Pooling strength	Local to highly pooled	Aggregate service level	Cross-region dependency pattern

Table 8. Key uncertainty dimensions and their role in planning formulations.

Uncertainty aspect	*Summary measure*	*Treatment in modelling*	*Planning usage*
Prediction spread	Ensemble variance	Scenario sampling and bounds	Robust thresholds for risk-based scheduling
Demand uncertainty	Lead-time demand variance	Safety factor computation	Base-stock level selection by part and site
Model shift across regions	Calibration residuals	Domain adaptation layers	Region-specific correction of risk scores
Data incompleteness	Missingness pattern	Imputation and feature masking	Conservative treatment of weakly observed assets

3. Data Infrastructure and Feature Construction Across Fleets and Regions

Machine learning enabled predictive maintenance depends critically on the availability and quality of data that describe component condition, usage, environment, and maintenance history [6]. In international supply networks, these data originate from multiple fleets, manufacturers, and regions, often using heterogeneous formats and varying levels of granularity. Building a coherent data infrastructure involves defining consistent identifiers for assets and components, aligning time bases across different logging systems, and integrating sensor streams with maintenance and inventory records.

Consider a set of assets indexed by $A = \{1, \dots, n_A\}$. Each asset $a \in A$ contains instances of parts $p \in P$ installed and replaced over time. Let $H_{a,p}(t)$ denote the health related measurements for part p on asset a at time t, structured as a vector in $\mathbb{R}^{d_H}$ that may include vibration statistics, temperatures, pressures, and derived features such as spectral energy in specific bands. Usage and environment variables, such as load, speed, ambient temperature, and duty cycle indicators, are represented by $U_a(t) \in \mathbb{R}^{d_U}$. Maintenance events are recorded as event logs with timestamps indicating inspections, replacements, and repairs, and inventory systems record the times at which parts are issued for maintenance tasks and received into stock.

To structure the prediction problem, one can discretize time into intervals indexed by $t = 1, \dots, T$. For each installed component instance c, with associated asset $a(c)$ and part type $p(c)$, feature vectors $x_c(t)$ are constructed by aggregating sensor, usage, and environment data over time windows preceding t. For example, the feature vector may include moving averages, moving standard deviations, exponentially weighted statistics, and lagged values of key measurements. Formally, $x_c(t) \in \mathbb{R}^d$ can be written as a function $x_c(t) = \Phi(H_{a(c),p(c)}(s), U_{a(c)}(s); s \le t)$, where Φ denotes a deterministic feature extraction map. In practice, Φ may incorporate domain knowledge, such as known frequency bands associated with specific failure modes, or may be learned jointly with the predictive model

Table 9. Evaluation metrics for predictive maintenance and inventory strategies.

Metric	*Definition*	*Primary scope*	*Interpretation*
Fill rate	Fraction of demand served from stock	Part–location–horizon	Direct indicator of service availability
Backorder depth	Time-integrated unsatisfied demand	Part–location	Measures cumulative delay burden
Emergency shipment share	Expedites over total shipments	Network	Proxy for logistical stress and cost
Planned-to-unplanned ratio	Scheduled over corrective tasks	Fleet or region	Indicates effectiveness of predictive policy
Aggregate cost index	Weighted sum of cost components	Scenario level	Compact comparison of competing policies

using deep architectures.

Label construction requires associating each feature vector with a target variable representing either a failure event within a future horizon or a remaining useful life quantity [7]. For binary failure prediction over a horizon of length τ, one can define $y_c(t) = 1$ if component c fails in the interval $(t, t+\tau]$ and $y_c(t) = 0$ otherwise, with censoring handled appropriately when the observation period ends or the component is replaced for reasons other than failure. For remaining useful life, one can define $r_c(t)$ as the time until failure or replacement, subject to censoring. In a multi fleet international context, it is often necessary to align labeling across regions with different maintenance practices, since some fleets may replace components proactively, while others run to failure. This introduces complexities in interpreting what constitutes a failure and how to treat planned replacements in the data.

Inventory and maintenance planning data must also be integrated. For each location l and part p, time series $I_{l,p}(t)$, $Q_{l,p}(t)$, and $D_{l,p}(t)$ provide the stock levels, receipts, and issues, respectively. Linking these series to asset level events requires mapping each issued component instance to its destination asset and maintenance task. This mapping allows the derivation of aggregate demand that originates from predictive maintenance decisions, as opposed to corrective actions, and enables attribution of shortages or backorders to specific patterns in predictive model outputs or scheduling policies.

Feature construction must address missing data, sensor dropouts, and differing sampling rates. Techniques such as interpolation, last observation carried forward, or model based imputation can be applied, but they introduce uncertainty that should ideally propagate into the predictive model. Normalization across fleets and regions is particularly important, since sensor calibrations, typical operating ranges, and environmental conditions can vary [8]. A common approach is to standardize features within homogeneous subsets defined by asset class and region, but this can reduce comparability across groups. Alternatively, features can be transformed into dimensionless quantities, such as ratios relative to design limits or z scores computed using global statistics.

To support linear algebraic modeling later in the paper, it is useful to gather the feature vectors for all component instances and time points into large matrices. Let N denote the total number of training examples, potentially on the order of millions when multiple fleets and long histories are involved. The feature matrix $X \in \mathbb{R}^{N \times d}$ has rows $x_i^\top$ corresponding to examples and columns corresponding to features. Label vectors $y \in \mathbb{R}^N$ or $r \in \mathbb{R}^N$ collect the target variables. Additional matrices can encode asset identifiers, part types, and locations, enabling the construction of block structures or group specific parameters in the machine learning models. This matrix representation will facilitate the integration of predictive models

with inventory and maintenance decisions in later sections.

4. Machine Learning Models for Failure Risk and Remaining Useful Life

Once features and labels are constructed, machine learning models can be trained to estimate failure probabilities or remaining useful life for individual components. Let $X \in \mathbb{R}^{N \times d}$ denote the feature matrix and $y \in \{0,1\}^N$ the binary failure labels over a prediction horizon. A generic supervised learning model can be written as a function $f_\theta : \mathbb{R}^d \to [0,1]$ parameterized by θ, with predicted probabilities $\hat{p}_i = f_\theta(x_i)$ for example i. In matrix form, the vector of predictions is $\hat{p} = f_\theta(X)$, understood row wise. Training typically minimizes a loss function such as the logistic loss, expressed as [9]

$$\mathcal{L}(\theta) = -\sum_{i=1}^{N} \left[y_i \log \hat{p}_i + (1 - y_i) \log(1 - \hat{p}_i) \right] + \Omega(\theta),$$

where $\Omega(\theta)$ denotes a regularization term that may be an ℓ_2 or ℓ_1 norm or a more structured penalty.

For linear models, $f_\theta(x) = \sigma(w^\top x + b)$, where $w \in \mathbb{R}^d$ is a weight vector, $b \in \mathbb{R}$ is a bias, and σ is the logistic function. In matrix notation, this becomes $\hat{p} = \sigma(Xw + b\mathbf{1})$, where $\mathbf{1}$ is an N dimensional vector of ones. Linear models provide interpretability, but may be insufficient to capture complex degradation dynamics that depend on interactions among features and non linear effects of operating conditions.

Tree based ensemble methods such as gradient boosted decision trees can capture non linear relationships and interactions without explicit feature engineering. While such methods are typically described algorithmically, they can still be seen as mapping the feature matrix X into predicted probabilities via a function composed of basis functions corresponding to decision tree leaves. Although there is no simple closed form linear algebraic expression for the model, the predictions can be treated as a vector $\hat{p}$ whose elements are differentiable with respect to continuous hyperparameters in some implementations.

Deep learning architectures, particularly those designed for sequence data, are often used when raw time series are fed directly into the model with minimal hand crafted feature extraction. Recurrent neural networks, long short term memory networks, temporal convolutional networks, and transformer based encoders can process multivariate sequences and output failure risk estimates at each time step. When the architecture is viewed at a high level, it defines a mapping from a three dimensional tensor of shape $N \times T \times d_0$ into a vector of probabilities. For integration with inventory models, it is convenient to think of the output at a given current time as a vector $\hat{p}(t)$ across all component instances, which will drive anticipated maintenance induced demand.

For remaining useful life prediction, models estimate a continuous variable r or a survival distribution. Survival models based on proportional hazards or accelerated failure time formulations can incorporate covariates linearly in the log hazard or log lifetime. If $h_0(t)$ is a baseline hazard and $x \in \mathbb{R}^d$ a feature vector, a proportional hazards model has hazard $h(t|x) = h_0(t) \exp(w^\top x)$. Discretizing time into intervals and interpreting the conditional failure probabilities as outputs of a neural network leads to discrete time survival models, which can be trained with cross entropy on a multi label representation of failure or censoring at each interval [10]. In matrix form, the hazard parameters for all time intervals and samples can be arranged into a tensor, and training involves operations that resemble matrix multiplications and elementwise non linearities.

Uncertainty quantification is important for downstream planning decisions, since maintenance scheduling and inventory control are sensitive not only to expected failure rates but also to the spread of possible realizations. Techniques such as model ensembles, Monte Carlo dropout, or Bayesian neural networks can provide approximate predictive distributions. Let $\hat{p}^{(k)}$ denote the prediction vector from ensemble member $k = 1, \ldots, K$. The ensemble mean $\bar{p} = \frac{1}{K} \sum_{k=1}^{K} \hat{p}^{(k)}$ and empirical variance $v = \frac{1}{K-1} \sum_{k=1}^{K} (\hat{p}^{(k)} - \bar{p})^2$ provide a simple measure of predictive uncertainty at each component instance. These quantities can be collected into vectors $\bar{p}, v \in \mathbb{R}^N$ and used in downstream optimization as parameters influencing safety stock or maintenance thresholds.

Domain adaptation and transfer learning play roles when training data from one region or fleet must support predictions in another. Suppose there are two domains, source and target, with feature matrices $X^S \in \mathbb{R}^{N_S \times d}$ and $X^T \in \mathbb{R}^{N_T \times d}$, and labels available primarily in the source domain. Methods such as domain adversarial training or feature alignment seek a transformation $T : \mathbb{R}^d \to \mathbb{R}^{d'}$ such that the transformed features $Z^S = X^S W$ and $Z^T = X^T W$ share similar distributions for some matrix $W \in \mathbb{R}^{d \times d'}$. Training then proceeds on Z^S with an auxiliary objective encouraging closeness of distributions between Z^S and Z^T. This linear algebraic view emphasizes the role of matrix transformations in mapping heterogeneous data sources into a common representation suitable for predictive modeling.

Calibration of predicted probabilities is also important, particularly when predictions enter inventory and maintenance models as quantitative inputs. Miscalibrated probabilities can lead to systematic over or under stocking. Calibration methods such as temperature scaling or isotonic regression operate on the prediction vector $\hat{p}$ and true labels y to fit a monotonic transformation g such that $g(\hat{p})$ is better aligned with observed frequencies. When g is parameterized by a small number of parameters, such as a scalar temperature in the softmax setting, the calibration step can be viewed as a small optimization on the prediction vector [11].

In the following sections, predicted failure probabilities and remaining useful life estimates are treated not only as outputs of a learning system but as inputs to a linear algebraic representation of inventory and maintenance dynamics. The structure of the predictive model influences this representation primarily through the statistical properties of $\hat{p}$ and related quantities, including their distribution across locations, parts, and time, as well as the degree of correlation in prediction errors.

5. Modeling of Maintenance and Inventory Coupling

To analyze how machine learning enabled predictive maintenance affects parts shortages, it is useful to construct a linear algebraic model linking component level failure probabilities, maintenance scheduling decisions, and inventory states across a multi echelon international network. The key idea is to represent the relevant quantities as vectors and matrices, so that the impact of changes in predictions or policy parameters can be expressed as matrix vector operations and their compositions.

Consider a finite planning horizon divided into T discrete periods. Let C denote the set of component instances currently in service at the beginning of the horizon, with $|C| = n_C$. Each component $c \in C$ is associated with a part type $p(c) \in P$, an asset $a(c) \in A$, and a location $l(c) \in L$ where the asset resides. For each period $t = 1, \ldots, T$, the predictive maintenance system provides failure probability estimates $\hat{p}_{c,t}$ or remaining useful life estimates that can be converted into per period failure risk. Collect these into a matrix $P \in \mathbb{R}^{n_C \times T}$ with entries $P_{c,t} = \hat{p}_{c,t}$. The matrix P summarizes the machine learning model output over the planning

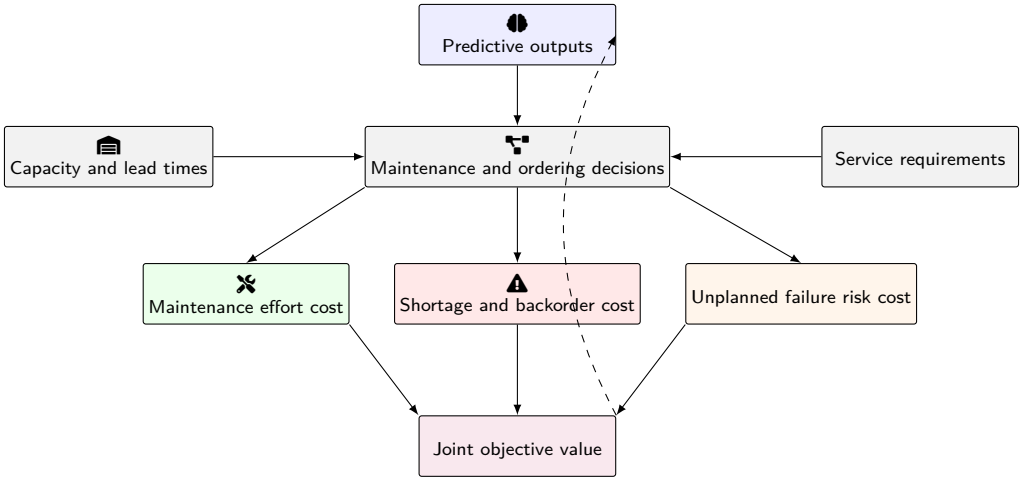


Figure 4. Conceptual decomposition of the joint cost function used for planning. Predictive outputs inform a decision layer that accounts for capacity constraints and service requirements. The resulting decisions contribute to maintenance effort cost, shortage and backorder cost, and residual failure risk cost, which aggregate into a single objective value. Feedback from observed objective realizations can be used to adjust predictive models and decision rules.

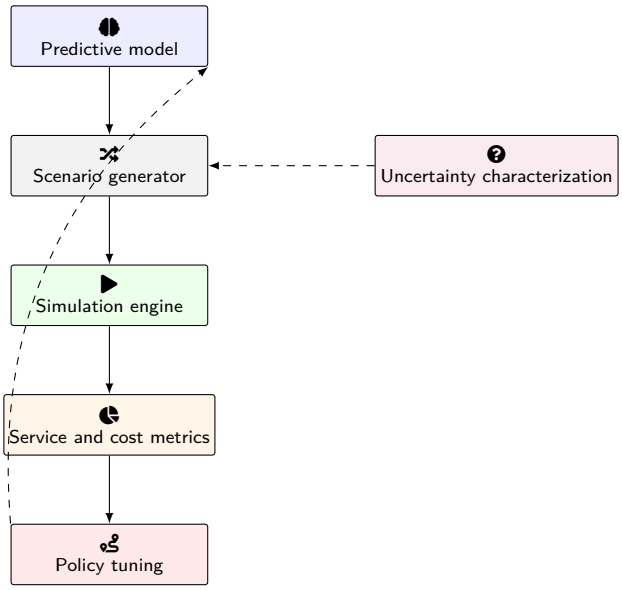


Figure 5. Scenario-driven evaluation loop for predictive maintenance policies. A predictive model feeds a scenario generator that samples possible degradation and demand paths, which are evaluated in a simulation engine producing service and cost metrics. A policy tuning layer adjusts decision rules based on these metrics, while separate uncertainty characterization supports the definition and sampling of scenarios.

horizon.

Maintenance scheduling decisions can be represented by a binary matrix $M \in \{0,1\}^{n_C \times T}$, where $M_{c,t} = 1$ indicates that component c is scheduled for replacement in period t. A feasible policy must satisfy logical and resource constraints. For example, each component

can be replaced at most once within the horizon, which can be expressed as

$$\sum_{t=1}^{T} M_{c,t} \leq 1 \quad \text{for all } c \in C.$$

Workshop capacity constraints at each location and period can be encoded by inequalities of the form [12]

$$\sum_{c:l(c)=l} M_{c,t} \leq K_{l,t},$$

where $K_{l,t}$ denotes the maximum number of replacements that can be processed at location l in period t.

The maintenance schedule induces demand for spare parts. Let $d_{p,l,t}$ denote the demand for part type p at location l in period t resulting from scheduled replacements, excluding unforeseen failures. The schedule induced demand can be written as

$$d_{p,l,t} = \sum_{c \in C: p(c)=p, l(c)=l} M_{c,t}.$$

This can be expressed compactly using assignment matrices. Define $A^P \in \{0,1\}^{n_C \times n_P}$ by $A^P_{c,p} = 1$ if part type of component c is p, and zero otherwise. Similarly, define $A^L \in \{0,1\}^{n_C \times n_L}$ by $A^L_{c,l} = 1$ if component c resides at location l. The Kronecker product can be used to construct a mapping from the component by time schedule matrix M into a part by location by time demand tensor, but an alternative is to flatten indices.

Flatten the triple index (p, l, t) into a single index j of length $n_P n_L T$. Define a mapping matrix $S \in \{0,1\}^{n_P n_L T \times n_C T}$ that captures how each component time decision contributes to each part location time demand. Let $m \in \mathbb{R}^{n_C T}$ denote the vectorization of M by stacking its columns, and let $d \in \mathbb{R}^{n_P n_L T}$ denote the vector of demands $d_{p,l,t}$ similarly stacked. Then the relationship can be written as

$$d = Sm.$$

The matrix S encodes the mapping from component level maintenance decisions to part location time demands through its block structure determined by A^P, A^L, and the period index.

Unexpected failures that are not prevented by predictive maintenance also generate demand. Let $F_{c,t}$ be a Bernoulli variable indicating an unplanned failure of component c in period t that was not preempted by scheduled maintenance. Conditional on the schedule M and predictive probabilities P, one can define failure probabilities $\tilde{p}_{c,t}$ that may differ from the original $\hat{p}_{c,t}$ due to policy adjustments. In a simple model, $\tilde{p}_{c,t} = \hat{p}_{c,t}$ for all $t < t^*(c)$ where $t^*(c)$ is the scheduled replacement time if any, and $\tilde{p}_{c,t} = 0$ thereafter. Let $f_{p,l,t}$ denote the expected unplanned failure demand for part type p at location l and time t, obtained by aggregating $\tilde{p}_{c,t}$ over relevant components. Collect these expectations into a vector $f \in \mathbb{R}^{n_P n_L T}$. This vector is a non linear function of the schedule m and predictions P, but for analytical purposes it can be treated as a given parameter once a schedule and predictive model are fixed [13].

Total expected demand for each part at each location and time is then $d^{\text{tot}} = d + f$. Inventory dynamics relate this demand to stock levels. Let $i(t) \in \mathbb{R}^{n_P n_L}$ denote the inventory vector across parts and locations at time t, and $q(t) \in \mathbb{R}^{n_P n_L}$ the vector of receipts resulting from previous orders. The issue vector $u(t) \in \mathbb{R}^{n_P n_L}$ corresponds to the realized demand, which may be truncated by stock availability. Under a simple first come first served rule, the inventory recursion is

$$i(t+1) = i(t) + q(t) - u(t),$$

with the constraint $0 \leq u(t) \leq i(t) + q(t)$ elementwise. Stockouts occur when desired demand exceeds available inventory. Let $d^{\text{tot}}(t)$ denote the planned total demand at time t corresponding to d^{tot}. The shortage vector $s(t)$ can be defined as

$$s(t) = \max\{d^{\text{tot}}(t) - i(t) - q(t), 0\},$$

where the maximum is taken elementwise. Over the horizon, shortages accumulate and can be quantified by aggregating across parts, locations, and time with appropriate weights reflecting the impact on service levels [14].

To connect predictive maintenance explicitly with shortages, consider the expected shortage as a function of the schedule vector m and prediction matrix P. In deterministic approximations, stochastic demand from unplanned failures is replaced by its expectation, leading to a deterministic demand vector $d^{\text{tot}} = Sm + f(P, m)$. The resulting inventory trajectory can be written in closed form by unrolling the recursion, yielding a linear function of initial inventory and order decisions and a piecewise linear function of m due to the max operator in the shortages. Let C^S be a diagonal matrix of shortage cost weights for each part location time element, and stack shortage vectors into a single vector s. An aggregate shortage cost can be expressed as

$$J_{\text{short}}(m; P) = s^\top C^S \mathbf{1},$$

where s depends on m and P through the demand and inventory recursions.

Maintenance costs and operational risk costs can be included similarly. Let $c^M \in \mathbb{R}^{n_C T}$ denote a vector of costs associated with performing maintenance actions at specific times on specific components. These costs may incorporate labor, downtime, and opportunity costs related to taking assets out of service. The total maintenance cost is $J_{\text{maint}}(m) = c^{M\top} m$, which is linear in m. Risk costs related to unplanned failures can be approximated as a function of residual failure probabilities $\tilde{p}_{c,t}$ that depend on the schedule. For example, an expected penalty for unplanned failures over the horizon can be written as

$$J_{\text{risk}}(m; P) = \sum_{c=1}^{n_C} \sum_{t=1}^{T} c^F_{c,t} \tilde{p}_{c,t},$$

where $c^F_{c,t}$ quantifies the cost impact of a failure of component c in period t. This term is non linear in m via $\tilde{p}_{c,t}$ but additive over components and periods.

Combining these elements, one can pose an optimization problem that chooses a maintenance schedule m to balance shortage, maintenance, and risk costs: [15]

$$\min_{m \in \{0,1\}^{n_C T}} J_{\text{short}}(m; P) + J_{\text{maint}}(m) + J_{\text{risk}}(m; P)$$

subject to capacity and logical constraints on m. Although this is a mixed integer non linear optimization problem, the linear algebraic representation clarifies the structure of interactions between predictive maintenance decisions and inventory outcomes. Approximations can be obtained by relaxing binary constraints to continuous variables in $[0, 1]$, linearizing the shortage terms, or decomposing across parts or locations.

From a linear algebra perspective, an important quantity is the sensitivity of expected shortages to changes in the prediction matrix P. At a high level, the mapping from P to d^{tot} to s can be viewed as a composition of functions, and its Jacobian captures how small perturbations in predicted failure probabilities propagate to shortage levels. While computing this Jacobian exactly may be challenging due to integer decisions and non linearities, approximate sensitivities can be estimated by differentiating through relaxed versions of the model or by finite differences. This provides insight into how prediction accuracy for specific components or regions influences overall parts availability.

6. Integration of Predictive Models into Planning Algorithms

The linear algebraic representation of maintenance and inventory coupling provides a foundation for designing planning algorithms that utilize machine learning predictions. In practice, planning processes operate under time and computational constraints, and must often accommodate additional considerations such as budget limits, regulatory requirements, and stakeholder preferences [16]. Integrating predictive models into such processes involves several layers, including the translation of predictions into actionable risk indicators, the formulation of heuristic or approximate optimization rules, and the interaction between tactical planning and operational execution.

One common integration approach is a two stage structure. In the first stage, the machine learning model produces, for each component instance, failure probability estimates over a fixed horizon or a remaining useful life estimate. These outputs are transformed into risk scores or priority indices that combine predicted risk with criticality measures and cost parameters. In the second stage, a scheduling algorithm selects a subset of components for maintenance within each period, subject to capacity and resource constraints, and the induced demand is communicated to inventory planners. In linear algebra terms, the first stage constructs the prediction matrix P and possibly additional parameter vectors such as failure cost weights, while the second stage solves an optimization problem over the schedule vector m that depends on these parameters.

Heuristic policies can often be expressed conveniently using matrix operations. For example, consider a class of threshold based policies that schedule component c for maintenance in period t if its predicted failure probability over a horizon of length τ exceeds a component specific threshold γ_c, and capacity is available. Let $r_{c,t} = \sum_{s=t}^{t+\tau} \hat{p}_{c,s}$ denote the cumulative predicted risk. Define a matrix R with entries $R_{c,t} = r_{c,t}$, which can be computed from P by multiplying with a banded Toeplitz matrix representing the horizon summation. Formally, if B is a $T \times T$ matrix where $B_{t,s} = 1$ if $s \in [t, t+\tau]$ and zero otherwise, then $R = PB$. A naive threshold policy would set $M_{c,t} = 1$ when $R_{c,t} \geq \gamma_c$, but this must be adjusted to avoid scheduling the same component multiple times and to respect capacity. Capacity constraints can be handled by ranking candidate maintenance actions within each period by $R_{c,t} - \gamma_c$ or by a cost adjusted priority, and selecting the top actions up to capacity.

Inventory planners can use predictive information to adjust safety stocks and replenishment orders [17]. In a periodic review system with lead time L, target inventory positions for each part and location can be expressed as base stock levels that depend on forecasted mean demand and safety factors. Let $\mu_{l,p}$ and $\sigma_{l,p}$ denote the predicted mean and standard deviation of demand over the lead time for part p at location l, derived from the schedule induced and expected unplanned demand. A classical formula for the base stock level is $S_{l,p} = \mu_{l,p} + z\sigma_{l,p}$, where z is a safety factor chosen according to the desired service level. In vector form, let $\mu \in \mathbb{R}^{n_P n_L}$ and $\sigma \in \mathbb{R}^{n_P n_L}$ collect these quantities, and define $S = \mu + z \odot \sigma$, where $\odot$ denotes elementwise multiplication. The mean and variance of demand can be computed from the prediction matrix P and schedule vector m using linear operations and variance aggregation formulas that assume certain independence structures.

When inventory is centralized or partially pooled, planning must consider not only local demand but also the option of reallocating stock between locations. This can be represented using a transfer matrix $T \in \mathbb{R}^{n_P n_L \times n_P n_L}$ that maps the vector of inventory at central and regional nodes into effective stock availability at local nodes, accounting for transshipment policies and lead times. Predictive maintenance can influence the design of T by altering expectations of which locations will experience elevated risk in future periods. For instance,

an algorithm may recommend repositioning stock proactively toward locations where the aggregated predictive risk is higher over the next several periods.

Optimization based integration approaches formulate joint problems over maintenance schedules and replenishment decisions. Let $o(t) \in \mathbb{R}^{n_P n_L}$ denote order quantities placed at time t for future periods, and stack these into a vector o. The joint decision vector is then $z = [m^\top, o^\top]^\top$. The linear algebraic model described earlier provides a mapping from z and P into expected shortages, inventory levels, and maintenance outcomes. A joint cost function can be expressed as

$$J(z; P) = z^\top C z + c^\top z + \phi(P, z),$$

where C is a symmetric positive semidefinite matrix capturing quadratic penalties on orders and maintenance variations, c is a cost vector, and $\phi(P, z)$ encapsulates non linear terms related to shortages and risk [18]. Approximate solutions can be obtained by linearizing ϕ around a baseline schedule and order plan, leading to a quadratic programming problem, or by using decomposition methods that alternate between updating maintenance and inventory decisions.

Scenario based planning is another integration mechanism. Instead of relying solely on expected failure probabilities, one can generate multiple scenarios of failure realizations and maintenance demands by sampling from the predictive distributions produced by the machine learning model. For each scenario k, a realization of the demand vector $d^{(k)}$ is constructed, and corresponding inventory trajectories and shortages are simulated. The scenario set can be represented as a matrix $D^{\text{scen}} \in \mathbb{R}^{n_P n_L T \times K}$ whose columns correspond to scenarios. Planning decisions can then be chosen to minimize an objective function defined on this matrix, such as expected shortage or a risk measure like conditional value at risk. While this increases computational requirements, the use of matrix operations allows efficient implementation, particularly when exploiting linear structure and parallelization.

In all these approaches, the quality and structure of the prediction matrix P are central. Correlated errors across components or locations can cause the planning model to over or underestimate demand in particular regions systematically. Sensitivity analysis using the linear algebraic model can help identify which elements of P have the greatest impact on shortage outcomes, guiding efforts to refine models or collect additional data for those components or regions [19].

7. Numerical Experiments and Sensitivity Analysis

To illustrate how machine learning enabled predictive maintenance interacts with international spare parts networks, it is informative to consider numerical experiments based on stylized but representative network configurations. While specific parameter values depend on the application, the goal is to explore qualitative relationships between prediction quality, maintenance scheduling policies, and parts shortages. The linear algebraic model described earlier provides a convenient framework for such experiments, since it allows the generation of synthetic instances and the evaluation of shortage metrics through matrix operations.

Consider a network with three echelons: a global central warehouse, several regional hubs, and multiple local maintenance bases in each region. Let the number of parts be on the order of tens to hundreds, and the number of component instances in service be on the order of several thousands. Lead times from the central warehouse to regional hubs are assumed to range from several weeks to months, with variability representing customs and transportation uncertainty, while lead times from hubs to local bases are shorter but still non negligible.

Initial inventory levels and base stock policies are set to achieve target fill rates under a baseline maintenance policy that does not use predictive information and relies on historical average failure rates.

Synthetic failure processes are generated using stochastic degradation models [20]. For each component instance, a latent health state evolves according to a discrete time stochastic process, such as a random walk with drift or a linear state space model

$$x_c(t+1) = Ax_c(t) + Bu_c(t) + w_c(t),$$

where $x_c(t) \in \mathbb{R}^{d_x}$ is the health state, $u_c(t)$ represents operational stress, and $w_c(t)$ is process noise. Failure occurs when a linear functional $h^\top x_c(t)$ crosses a threshold, defining a stopping time. Sensor measurements and derived features form the observable process used to train a machine learning model that approximates the mapping from observed history to failure risk. Time series are simulated across all component instances and locations, producing a feature matrix X and subsequent labels for training the predictive model.

Different levels of prediction quality can be modeled by altering the noise structure, feature availability, and model capacity. For example, a high quality prediction scenario may assume that the model class is well matched to the true degradation process, that the training dataset is large, and that features are informative, leading to high discrimination and well calibrated probabilities. A lower quality scenario may impose mismatched model structures, reduced feature sets, or distribution shifts between training and deployment domains. These differing conditions can be reflected in the prediction matrix P through changes in discrimination metrics and calibration properties.

Maintenance scheduling policies are applied using the prediction outputs [21]. A purely reactive policy schedules maintenance only upon failure, resulting in unplanned demands that directly reflect the failure process. A basic preventive policy schedules replacements at fixed intervals independent of predictions, yielding more regular but potentially inefficient demand. Predictive policies use thresholds on predicted risk or remaining useful life to schedule maintenance within allowable windows. These policies are implemented as algorithms that construct the schedule matrix M given the prediction matrix P and capacity constraints.

Inventory control is modeled using base stock policies adjusted or not adjusted by predictive information. In some scenarios, inventory planners ignore predictive information and maintain base stocks based on historical averages. In other scenarios, planners recalculate base stock levels periodically using updated demand forecasts derived from predictive maintenance schedules and expected residual failure risk. In each case, the inventory recursion is simulated over the horizon, and shortage events are recorded [22].

Key performance metrics include average fill rate by part and location, frequency and magnitude of stockouts, total inventory investment, and frequency of emergency expedites or cross shipments. These metrics can be computed from the shortage vectors and inventory trajectories. For each scenario, multiple simulation replications can be generated by resampling degradation paths and observation noise, leading to empirical distributions of metrics. The results can be summarized using means, variances, and quantiles, providing insight into variability and robustness.

Sensitivity analysis is conducted by systematically varying parameters that characterize prediction quality, capacity constraints, lead time distributions, and safety stock policies. For instance, one can vary the receiver operating characteristic curve of the predictive model by adjusting noise levels and assess how the relationship between prediction discrimination

and fill rate behaves under different lead time regimes. Another set of experiments can vary workshop capacities to explore how the ability to flexibly schedule maintenance influences the benefits of predictive information. Similarly, the impact of correlated prediction errors across locations can be studied by introducing covariance structures into the generation of prediction matrices and observing their effect on simultaneous shortages [23].

The linear algebraic representation simplifies these experiments by enabling parameters to be encoded in matrices and vectors that can be perturbed systematically. For example, lead time distributions can be represented by convolution matrices that transform order vectors into receipt vectors, and adjustments to lead time means or variances correspond to modifying the entries of these matrices. Prediction errors can be modeled by adding structured noise to the prediction matrix P, such as additive Gaussian noise with a specified covariance matrix capturing cross component correlations.

Preliminary numerical experiments constructed in this way generally show that predictive maintenance can reduce the frequency of parts shortages under certain conditions, particularly when lead times are moderate, prediction quality is sufficient to distinguish high and low risk components, and maintenance capacities allow for smoothing of demand. At the same time, the experiments highlight that predictive information can increase peak demand variability in some configurations, especially when many components are driven over threshold by similar operating conditions and policies respond synchronously. Under such circumstances, careful coordination between maintenance scheduling and inventory planning is required to avoid shifting variability rather than reducing it.

8. Implementation Considerations, Organizational Integration, and Limitations

Translating the modeling concepts and numerical insights into practice within international supply networks involves several implementation considerations. Data acquisition and governance are major challenges, since predictive maintenance relies on consistent and high quality data streams from assets located in different countries, operated by different organizational entities, and maintained by different service providers. Establishing data pipelines that collect sensor data, event logs, and inventory transactions in near real time requires harmonization of protocols, adherence to privacy and regulatory-requirements, and agreements on data sharing between stakeholders [24].

Model training and deployment must address heterogeneity across fleets and regions. In some cases, transferring models trained on one fleet to another may be appropriate, while in other cases regional adaptations are necessary due to differing operating conditions. This suggests a modeling strategy where a global base model captures common degradation patterns, and regional adaptation layers or calibration factors adjust predictions to local conditions. In linear algebra terms, the feature matrix may be multiplied by a global weight matrix followed by region specific transformation matrices, effectively factorizing model parameters into shared and region specific components. Monitoring model performance in each region and adjusting transformations as new data become available are continuous activities.

Integration with existing planning systems is another major aspect. Many organizations use enterprise resource planning systems and dedicated maintenance and inventory planning tools that were not originally designed to ingest high dimensional predictive risk information. Bridging this gap often requires intermediate services that aggregate predictions into simpler indicators, such as component risk scores or expected demand adjustments for specific parts and locations over specified horizons [25]. These indicators can be represented as low dimen-

sional vectors derived from the prediction matrix by linear transformations or aggregation operations, making them easier to integrate with established planning workflows.

Organizational processes and decision rights influence how predictive information is used. Maintenance planners may have discretion to override model recommendations based on local knowledge, regulatory constraints, or customer preferences. Inventory planners may prioritize service levels for certain customer segments or asset types. These qualitative factors complicate the relationship between model outputs and actual maintenance and replenishment decisions, introducing variability that is not captured in purely quantitative models. Capturing such behavior in a modeling framework would require additional structures that represent decision rules as functions of both predictions and contextual variables.

Limitations of the modeling approach should be acknowledged. The linear algebraic representation simplifies many aspects of the system, such as assuming that expected unplanned failure demand can be aggregated linearly from predictive probabilities and that inventory dynamics can be represented with relatively simple recursions [26]. In reality, maintenance and inventory processes may involve non linear interactions, state dependent lead times, and complex priority rules that are difficult to capture exactly. Mixed integer representations may grow very large when modeling thousands of components and multiple echelons, making exact optimization computationally challenging. Approximate methods such as heuristics, decomposition, or simulation based optimization are necessary in such cases, and the gap between approximate and optimal solutions may be difficult to quantify.

Another limitation relates to uncertainty quantification. While the modeling framework treats predicted probabilities as inputs, uncertainties in these predictions can be substantial, particularly for rare failure modes or newly introduced component designs. Incorporating prediction uncertainty into planning decisions requires stochastic optimization or robust optimization formulations, which can increase complexity. Scenario based approaches can mitigate this but rely on the availability of realistic sampling schemes for predictive distributions, which in turn depend on modeling assumptions.

Environmental and regulatory differences across countries can also constrain the adoption of predictive maintenance for parts planning [27]. For example, some regulatory regimes may prescribe fixed inspection intervals or replacement schedules that limit the flexibility to shift maintenance tasks based on predicted risk. Customs and trade regulations can affect the feasibility of repositioning inventory across borders in response to predicted demand shifts. These factors must be considered when designing policies, and may require region specific adaptations of the general modeling framework.

Finally, organizational learning and change management play a role in realizing potential benefits. Introducing predictive maintenance alters the division of responsibility between maintenance, operations, and supply chain functions, and may change performance metrics and incentives. Establishing governance structures that monitor both predictive model performance and system level outcomes, such as parts shortages and service levels, is important for iterative improvement. From a modeling standpoint, this implies the need for feedback loops in which observed shortages and deviations between predicted and realized demand are used not only to refine machine learning models but also to adjust planning heuristics and policy parameters.

9. Conclusion

Empirical validation is another important area for future work. Stylized experiments can illustrate structural patterns, but operational data from actual international supply networks are required to determine how predictive maintenance signals interact with real maintenance practices, heterogeneous asset conditions, and region specific supply constraints. Such data would support the examination of how prediction accuracy varies across operating environments, how maintenance decisions are adjusted in response to predicted risks, and how these adjustments influence observed demand trajectories and stock availability.

Another direction involves studying how prediction errors propagate through multi echelon inventory systems. Prediction uncertainty affects not only the timing of maintenance actions but also the distribution of demand across both part types and locations. Quantifying how different forms of uncertainty influence demand variability, reorder frequencies, and backorder depth could help determine which components of the predictive pipeline require further refinement. For example, systematic biases in predicting degradation for certain asset classes may lead to recurring misalignment between forecasted and actual demand, while high variance in predictions could result in inconsistent scheduling decisions across regions.

The interaction between capacity constraints and predictive maintenance outputs also warrants deeper investigation [28]. Maintenance depots, workshops, and regional hubs typically operate under personnel, tooling, and time limitations. If predictive outputs indicate elevated failure risk across multiple assets simultaneously, but available capacity cannot accommodate all recommended maintenance actions, the resulting scheduling backlog may generate compressed demand patterns. Such compression affects inventory consumption and can increase the likelihood of shortages. Modeling these interactions requires an integrated approach that combines predictive risk estimates with queueing effects, resource constraints, and temporal coupling in maintenance operations.

The structural properties of supply networks further influence how predictive information should be used. Different network topologies exhibit different sensitivities to demand fluctuations. Networks with centralized inventory pools may absorb moderate shifts in demand without significant shortages, while networks with localized storage and segmented supply channels may exhibit more pronounced reactions to predictive deviations. Examining how predictive maintenance interacts with network design choicessuch as centralization versus decentralization, flexibility in transshipment, or regional autonomymay help clarify under which conditions predictive information leads to smoother demand patterns and under which conditions it introduces additional variability [29].

Another aspect involves evaluating the alignment between predictive maintenance horizons and supply chain lead times. Predictive horizons that are short relative to procurement times may have limited influence on inventory decisions because alerts arrive too late to adjust replenishment strategies. Conversely, very long predictive horizons may introduce noise if long term predictions are highly uncertain. Matching predictive horizons with logistical realities requires a careful assessment of model limitations, degradation characteristics, and procurement cycles. Studying this alignment may clarify situations where predictive maintenance influences inventory decisions effectively and situations where its impact is minimal.

Differences in environmental conditions, such as temperature, humidity, and operational load, affect both degradation processes and prediction performance. Predictive models may produce different risk profiles for assets operating in cold climates, desert regions, or maritime environments. These contextual differences influence maintenance timing, spare parts demand

patterns, and replenishment needs across the network. Investigating how environmental heterogeneity affects the coupling between predictive maintenance and inventory planning may help determine how feature representations, calibration strategies, and decision thresholds should be adjusted by region [30].

Integration with existing enterprise systems presents another set of considerations. Maintenance management systems, inventory planning tools, and enterprise resource planning systems differ across regions and often rely on legacy data formats or inconsistent coding practices. Ensuring that predictive outputs are incorporated consistently across these systems requires processes for data synchronization, version control, and model monitoring. Studying the implications of integration constraints could help identify where inconsistencies in data flow or system behavior may create divergences between predicted and actual maintenance driven demand.

Human decision making also plays a role. Even when predictive estimates are available, planners and technicians may adjust or override recommendations based on local experience, operational priorities, or risk assessments. These adjustments influence realized maintenance timing and, therefore, actual demand patterns. Investigating how human judgment interacts with predictive outputs may provide insight into the degree to which predictive maintenance can be expected to influence inventory consumption in practice [31]. It would also support an understanding of how decision variability affects the accuracy of inventory forecasts.

Further work could examine how different policy formulations respond to the same set of predictive inputs. Threshold based policies, cost shaped optimization approaches, rolling horizon planners, and heuristic decision rules may yield different demand trajectories even when applied to identical predictive matrices. Comparing these policy types under varying levels of prediction quality, network segmentation, and capacity constraints may help clarify how different decision mechanisms amplify or dampen demand fluctuations.

Another area for investigation concerns the computational tractability of joint maintenance and inventory optimization models. As the dimensionality of the system increases, exact optimization may become impractical, requiring approximations or decomposition techniques. Analysis of which constraints, decision variables, and coupling terms contribute most strongly to complexity could help guide the development of tractable surrogate models or reduced order representations. This could include examining whether certain interactionssuch as cross regional coupling or multi part compatibility relationshipscan be simplified without materially altering the behavior of the system [32].

The distributional characteristics of demand induced by predictive maintenance also merit closer study. Demand shaped by predictive scheduling may exhibit different autocorrelation structures, burstiness levels, or spectral properties compared to demand driven by reactive maintenance. These properties influence safety stock calculation methods, reorder point strategies, and transshipment policies. Understanding how predictive maintenance reshapes demand distributions may support more accurate inventory control rules.

Another dimension involves studying the long term effects of predictive maintenance on part consumption. Predictive maintenance may shift the balance between early replacement and extended use of components. This could alter failure modes, part lifecycles, and overall consumption rates. Evaluating how these shifts influence aggregate demand over multi year horizons could help assess the implications for procurement planning, inventory strategy, and supplier relationships [33].

It is also relevant to examine scenarios where predictive maintenance does not materially

reduce shortages or where shortages arise due to factors unrelated to predictive accuracy. Examples include supply chain disruptions, material shortages, geopolitical constraints, and abrupt changes in operational mission profiles. In such cases, predictive maintenance may have limited influence compared to external constraints. Understanding the relative contribution of predictive driven demand variation versus structural supply limitations may help contextualize the conditions under which predictive maintenance meaningfully affects parts availability.

Further research could explore the development of models that incorporate endogenous feedback between predictive maintenance and asset behavior. Maintenance actions triggered by predictive estimates may alter degradation trajectories, introduce new failure modes, or change the statistical properties of operational data. Over time, these feedback loops may influence model training and the behavior of the predictive system. Studying these dynamics may help clarify the stability and evolution of predictive maintenance strategies [34].

In addition, exploring the interactions between maintenance induced downtime, operational readiness, and inventory planning may offer further insight. Maintenance timing affects asset availability, which in turn influences mission scheduling, workload distribution, and the loading applied to other assets. These operational adjustments feed back into degradation rates and spare parts consumption patterns. Modeling these interactions may support a more complete understanding of how predictive maintenance integrates with broader operational planning systems.

Examining cross fleet interactions may be relevant in multinational or multi platform contexts. Spare parts may be shared across platforms with different degradation behaviors, and predictive signals may indicate different maintenance priorities across fleets. Analyzing how predictive information should be aggregated, prioritized, or segregated across platforms could clarify how cross fleet coupling affects inventory outcomes. The relationship between machine learning enabled predictive maintenance and parts shortages in international supply networks involves multiple interacting dimensions, including data quality, predictive modeling, capacity constraints, network structure, human decision making, and logistical considerations. Extensions of this work could focus on empirical validation, uncertainty propagation, policy comparison, computational tractability, environmental heterogeneity, and long term consumption effects to further clarify the conditions under which predictive maintenance influences parts availability [35].

References

[1] Ejaz Ghani, Arti Grover, and William R. Kerr. “Highway to Success in India: The Impact of the Golden Quadrilateral Project for the Location and Performance of Manufacturing - Highway to success in India : the impact of the golden quadrilateral project for the location and performance of manufacturing”. In: (Jan. 28, 2013), pp. 1–39. DOI: 10.1596/1813-9450-6320.

[2] Panagiotis Barlas and Cathal Heavey. “Automation of input data to discrete event simulation for manufacturing: A review”. In: *International Journal of Modeling, Simulation, and Scientific Computing* 07.01 (Mar. 28, 2016), pp. 1630001–. DOI: 10.1142/s1793962316300016.

[3] Daniel K. Reef et al. “Equal Access to Hematopoietic Stem Cell Transplant in a Minority-Rich, Ethnically Diverse, Socioeconomically Disadvantaged Population with Peripheral T-Cell Lymphomas at the Montefiore Medical Center/Albert Einstein College of Medicine”. In: *Blood* 138.Supplement 1 (Nov. 5, 2021), pp. 3035–3035. DOI: 10.1182/blood-2021-149271.

[4] James Hayes. "Storage shapes up - [IT data management]". In: *Engineering & Technology* 4.17 (Oct. 10, 2009), pp. 50–52. DOI: 10.1049/et.2009.1711.

[5] Shreesha Hegde Kukkuhalli. "Global Parts Management through Data and AI Leveraging Structured and Unstructured Data". In: *ESP International Journal of Advancements in Computational Technology* 1.1 (2023).

[6] Xiaoling Sun, Zihao Xiong, and Haocheng Yang. "The Transformation of Enterprise Data and Business Under the Data Center Structure-Taking Kuaishou Group as an Example". In: *BCP Business & Management* 36 (Jan. 13, 2023), pp. 100–105. DOI: 10.54691/bcpbm.v36i.3393.

[7] Yun Ma and Hao Du. "Enterprise Data at Huawei - Data Service Development Targeting Self-service Consumption". In: Springer Singapore, Nov. 23, 2021, pp. 111–161. DOI: 10.1007/978-981-16-6823-4_6.

[8] Nguyen Thi Huu Ai, Nguyen Thi Thia, and Le Thu Trang. "The Influence of the Three-Component Commitment Model on the Intention to Leave of Employees Working in Insurance Enterprises". In: *INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH AND ANALYSIS* 6.9 (Sept. 29, 2023). DOI: 10.47191/ijmra/v6-i9-67.

[9] Akhlaq Ahmad. "Examining Green HR Practices and Innovative Work Behaviors: A Moderated Mediation Model in Hungarian Automotive Enterprises". In: *Journal of Digitovation and Information System* 3.2 (Dec. 27, 2023), pp. 250–265. DOI: 10.54433/jdiis.2023100033.

[10] Xi Li, Anthony Ventresque, John Murphy, and James Thorburn. "SOC: Satisfaction-Oriented Virtual Machine Consolidation in Enterprise Data Centers". In: *International Journal of Parallel Programming* 44.1 (Oct. 19, 2014), pp. 130–150. DOI: 10.1007/s10766-014-0333-6.

[11] Kannan Subramanian R and Dr. Sudheesh Kumar Kattumannil. "Enterprise Data Management". In: Apress, Jan. 3, 2022, pp. 457–577. DOI: 10.1007/978-1-4842-7440-8_6.

[12] Liuxu Wang. "A study on the impact of environmental information disclosure on corporate financial performance". In: *SHS Web of Conferences* 169 (May 29, 2023), pp. 1037–01037. DOI: 10.1051/shsconf/202316901037.

[13] Julia Whitlow Yarahuan, Amy Billet, and Jonathan D. Hron. "A Quality Improvement Initiative to Decrease Platelet Ordering Errors and a Proposed Model for Evaluating Clinical Decision Support Effectiveness." In: *Applied clinical informatics* 10.03 (July 10, 2019), pp. 505–512. DOI: 10.1055/s-0039-1693123.

[14] null Dr. Ravi Malhotra. "Deep Learning Models for Cybersecurity in Indian Enterprises: A Comparative Study". In: *Universal Research Reports* 8.4 (Dec. 29, 2021). DOI: 10.36676/urr.v8.i4.1410.

[15] S. Berski and Martyna Bilau. "Safety Mechanisms in Relational Database as Part of the IT System of the Enterprise". In: *New Trends in Production Engineering* 2.2 (Dec. 1, 2019), pp. 12–23. DOI: 10.2478/ntpe-2019-0068.

[16] Robert N. Stewart, Devin A. White, Marie L. Urban, April Morton, Clayton G. Webster, Miroslav Stoyanov, Eddie A Bright, and Budhendra L. Bhaduri. "Uncertainty quantification techniques for population density estimates derived from sparse open source data". In: *SPIE Proceedings* 8747 (May 23, 2013), pp. 874705–. DOI: 10.1117/12.2015795.

[17] Josh Juneau and Tarun Telang. "Jakarta Security". In: Apress, July 19, 2022, pp. 551–585. DOI: 10.1007/978-1-4842-8079-9_15.

[18] ... "DEVELOPMENT OF THE REFERENCE MODEL FOR MODERN MINING AND METALLURGICAL ENTERPRISE". In: *Organizer of Production* 4 (Feb. 9, 2022), pp. 25–34. DOI: 10.36622/vstu.2021.99.44.003.

[19] John Scarisbrick. "DSP: Driving Convergence in Digital Communications". In: *NIP & Digital Fabrication Conference* 16.1 (Jan. 1, 2000), pp. 1–5. DOI: 10.2352/issn.2169-4451.2000.16.1.art00001_1.

[20] "A Novel Cloud Data Center Resource Scheduling Optimization Model based on KPCA and Markov Chain". In: *Revista Tecnica De La Facultad De Ingenieria Universidad Del Zulia* (Jan. 15, 2017). DOI: 10.21311/001.39.12.011.

[21] Xiaofei Wang, Yuchun Sun, and Guoqiang Li. "The Relationship Between Collaborative Innovation Risk and Performance of Industrial Parks". In: *Journal of Global Information Management* 30.9 (Mar. 1, 2022), pp. 1–19. DOI: 10.4018/jgim.298673.

[22] Chun-wei Wang. "Data Loss Prevention System Based on Big Data". In: *DEStech Transactions on Computer Science and Engineering* aita (Nov. 28, 2017). DOI: 10.12783/dtcse/aita2017/16032.

[23] Changyou Ye, Xiaowei Song, G. N. Vivekananda, and V. Savitha. "Intelligent physical systems for strategic planning and management of enterprise information". In: *Peer-to-Peer Networking and Applications* 14.4 (Sept. 19, 2020), pp. 2501–2510. DOI: 10.1007/s12083-020-00966-7.

[24] Ruiyang Liu. "Data Security Threats of Log Aggregation". In: *SHS Web of Conferences* 170 (June 14, 2023), pp. 1007–01007. DOI: 10.1051/shsconf/202317001007.

[25] Aleksandra Stoiljkovi and Slobodan Mari. "Financing and Success/Performance of Small Businesses in Serbia". In: *Management:Journal of Sustainable Business and Management Solutions in Emerging Economies* (Dec. 7, 2020). DOI: 10.7595/management.fon.2020.0022.

[26] Yongkang Li and Boyan Duan. "Research on Enterprise Data Management Strategy Analysis System Based on Knowledge Mining Model". In: *Mobile Information Systems* 2022 (Sept. 29, 2022), pp. 1–9. DOI: 10.1155/2022/6359293.

[27] Kelly Briner and John Rome. "The Need for and Value of Data Governance". In: Springer International Publishing, Feb. 2, 2018, pp. 67–80. DOI: 10.1007/978-3-319-71162-1_5.

[28] Rutvij Khanolkar. "Impact of COVID19-related non-pharmacologic interventions on healthcare utilization for other virally-triggered respiratory illnesses". In: *Canadian Journal of Emergency Nursing* 45.2 (July 12, 2022), pp. 8–9. DOI: 10.29173/cjen174.

[29] Ian A. Willson. "The Evolution of the Massively Parallel Processing Database in Support of Visual Analytics". In: *Information Resources Management Journal* 24.4 (Oct. 1, 2011), pp. 1–26. DOI: 10.4018/irmj.2011100101.

[30] . .. " ". In: 27 (Nov. 17, 2020), pp. 58–61. DOI: 10.36477/2522-1256-2020-27-10.

[31] Manishankar S and S. Sathayanarayana. "Enhanced Big Data Platform for Visualization of Employee Data." In: *JOIV : International Journal on Informatics Visualization* 2.3 (May 23, 2018), pp. 169–174. DOI: 10.30630/joiv.2.3.132.

[32] Mehdi Kargar. "Enhancing knowledge accessibility in digital economy with search over enterprise data". In: *Personal and Ubiquitous Computing* 27.4 (Feb. 24, 2022), pp. 1645–1652. DOI: 10.1007/s00779-022-01668-y.

[33] Xiaoyi Yuan and Yu Chen. "Design and Implementation of Enterprise Refined Marketing System Under the Background of Big Data". In: Atlantis Press International BV, Dec. 20, 2022, pp. 763–772. DOI: 10.2991/978-94-6463-030-5_76.

[34] Haihua Mo, Yingchang You, Liping Wu, Fenghua Yan, Ming Chang, Weiwen Wang, Peng Wang, and Xuemei Wang. "Potential Impact of Industrial Transfer on Pm2.5 and Economic Development Under Scenarios with Multiple Objectives in Guangdong, China". In: *SSRN Electronic Journal* (Jan. 1, 2022). DOI: 10.2139/ssrn.4208636.

[35] M.I. Lali, R.U. Mustafa, F. Ahsan, M.S. Nawaz, and W. Aslam. “Performance Evaluation of Software Defined Networking vs. Traditional Networks”. In: *The Nucleus* 54.1 (Mar. 30, 2017), pp. 16–22. DOI: 10.71330/thenucleus.2017.95.