Applied Science, Computing and Energy, 2026, 4(2), 310-333

Data-Driven Supply Chain Transformation Through Multi-Layer Predictive Intelligence: A Self-Adaptive Procurement Optimization System with Real-Time ERP Integration

Taiwo Ruth Owoeye, Sharon Oluwaseun, Arti Raikwar, and Chinyan Blessing

Received: 11 November 2025/Accepted: 15 March 2026 /Published: 30 March 2026

Abstract: *Resilient supply chains increasingly depend on intelligent systems capable of real-time adaptation rather than static, pre-defined optimization rules. This study proposes a multilayer predictive intelligence framework that autonomously adjusts procurement decisions through continuous ERP synchronization, departing from conventional periodic decision-support architectures. The system integrates three computational layers: (i) a demand forecasting module based on ensemble learning, achieving a 43.3% reduction in forecasting error (MAPE reduced from 21.7% to 12.3%), (ii) a supplier intelligence layer incorporating sentiment analysis of transactional communications and external risk signals, and (iii) a mixed-integer optimization engine that simultaneously minimizes procurement cost, inventory holding cost, and stockout risk under dynamic constraints.The framework was implemented in a medium-sized manufacturing enterprise managing 342 suppliers, 1,247 SKUs, and an average monthly procurement volume of 2,847 orders. Empirical results over a 12-month operational period demonstrated a 67.1% reduction in stockout incidents (from 23.4 to 7.7 per month), a 22.9% reduction in average inventory holding period (from 47.2 to 36.4 days), and a 6.6% reduction in total procurement costs (from $5.58M to $5.21M per month). Expedited logistics costs decreased by 57.8%, while on-time delivery performance improved from 83.4% to 91.8%.The system achieved a 99.3% operational uptime with an average end-to-end decision latency of 340 milliseconds (95th percentile: 580 ms), enabling near real-time response to supply chain events. The closed-loop automation framework autonomously resolved 89% of procurement anomalies, while only 11% required human escalation, primarily for high-value or structurally constrained decisions. Forecast accuracy improvements and dynamic safety stock recalibration contributed significantly to system-wide efficiency gains, including a 23% reduction in inventory levels and improved working capital utilization .Beyond performance improvements, the study demonstrates a structural shift in procurement operations, where analysts transition from transactional execution roles to strategic governance and exception management. The findings highlight that the value of intelligent supply chain systems lies not only in algorithmic accuracy but in the integration architecture that enables continuous learning, real-time ERP synchronization, and autonomous decision execution. The results challenge the traditional assumption that increased system complexity reduces operational reliability, demonstrating instead that carefully designed multi-layer intelligence systems can simultaneously improve efficiency, responsiveness, and stability in mission-critical supply chain environments.*

Keywords: *Supply chain; Predictive analytics; ERP integration; Procurement optimization; Self-adaptive systems; Real-time automation; Machine learning*

Taiwo Ruth Owoeye*
Department of Business and Economics, Lincoln University, Oakland, CA, USA
Email: taiowoeye15@gmail.com

Sharon Oluwaseun
Tennessee Wesleyan University Business Department, Athens, Tennessee, USA
Email: thesharonoluwaseun@gmail.com

Arti Raikwar
Industrial Technology and Management, Armour College of Engineering, Illinois

Institute of Technology, Chicago, USA
Email: Araikwargc@gmail.com

Chinyan Blessing
Faculty of Business Administration, Imo State University, Imo State, Nigeria
Email: Chinyannein@gmail.com

1.0 Introduction

Modern supply chains operate in highly volatile and uncertain environments, where traditional planning cycles have shortened from monthly or weekly horizons to near real-time decision intervals. The challenges caused by the COVID-19 pandemic highlighted the limitations of traditional supply chain management models that were optimized for a steady state, which were found to be especially fragile in the face of sudden demand changes and supply chain disruptions (Ivanov and Dolgui, 2021). These disruptions exposed structural weaknesses in linear, forecast-driven procurement systems that lack adaptive feedback mechanisms. One factor that proved common to those manufacturing companies that survived the disruptions best is that their procurement systems were not just reactive, but adaptive (Chowdhury et al., 2021). This observation triggers a basic question which goes beyond the crisis management: How can supply chain intelligence systems be designed to continuously learn and self-adjust – without the need for manual reconfiguration of the system?

The theoretical background of adaptive supply chain systems dates back to cybernetics and control theory, which are based on the idea of feedback to make a system stable by continuously correcting errors (Wiener, 1961). But implementation has not been nearly as advanced as conceptualization. In spite of the complexity of their user interfaces and the ability to capture extensive data, most enterprise resource planning (ERP) systems are basically more complex record-keeping systems than intelligent decision agents (Jacobs and Weston Jr, 2007). The traditional integration of ERP systems' transactional data warehouses and operational decision-making is resistant to change, and often it is done manually by human analysts who have to pull, transform and interpret data before making procurement decisions. This human-in-the-loop architecture introduces decision latency ranging from hours to weeks, which is increasingly incompatible with modern dynamic supply environments. As a result, procurement decisions are often reactive rather than predictive, limiting the effectiveness of enterprise decision support systemsIn the last decade, machine learning has been applied to multiple different aspects of supply chain management, but there is a worrisome trend of point solutions rather than system architectures. In an idealized experimental setting, demand forecasting algorithms produce amazing accuracy gains, but these algorithms are often disconnected from the business, producing forecasts in spreadsheets for post-processing by humans (Carbonneau et al., 2008). Supplier evaluation systems typically employ advanced Natural Language Processing (NLP) to analyze risk factors from sources such as news snippets and financial statements, but such analyses are not automatically passed down to procurement execution workflows (Brintrup et al., 2020).This has resulted in fragmented intelligent subsystems that lack integration into a unified decision-making architecture. Despite advances in domain-specific applications, there remains limited integration of these models into end-to-end operational procurement workflowsOne of the most significant opportunities (and technical hurdles) in building truly adaptive supply chain systems is the ability to integrate ERP in real time. As enterprise resource planning platforms were designed in a time when batch processing was the norm, the architecture of these platforms is based on that. While modern ERP systems do have real-time APIs, they don't necessarily support the patterns of continuous synchronization that are necessary for closed-loop automation (Monk & Wagner, 2013). his architectural constraint significantly limits the ability of ERP systems to support real-time autonomous decision-making. Transaction consistency ensures protection of financial integrity can cause locking delays which are incompatible with decision cycles of only a few milliseconds. Fundamentally, ERP systems model procurement as discrete transactional events rather than as a continuous, feedback-driven control system.

Procurement optimization literature shows very high level of sophistication in mathematical formulations, especially in multi-objective formulations that include both cost and quality, delivery reliability and supplier relationship aspects (Chai et al., 2013). These models usually transform the procurement problem into a constraint satisfaction problem that is then solved using linear programming or genetic algorithms or other optimization methods. However, a persistent limitation remains in the disconnect between theoretical optimization models and real-world operational dynamics. The vast majority of the published optimization frameworks have static parameters: supplier capacities are constant, the demands have predictable distributions, and the lead times can be modeled as fixed value or simple probability distributions. However, these conditions do not exist in real supply chains. Lead times can vary according to supplier capacity, which is subject to its own issues; demand fluctuations occur as consumer preferences change, or as marketing campaigns are launched; and disruption in one part of the logistics networks can trigger a knock-on effect in another.

To meet these implementation issues, the concept of self-adaptive systems, which rely on architectures that adapt their own decision parameters in response to their own performance, provides a conceptual basis for solving them (Kephart and Chess, 2003). However, their application to procurement systems remains underexplored, particularly in real-time ERP-integrated environments.H Although originally developed for IT infrastructure management, autonomic computing principles are increasingly relevant to complex supply chain environments characterized by high uncertainty and interdependence. The principle is to differentiate strategic goals that need human decisions from tactical goals that are parameter modifications that can be sent to the automatic control loops. A procurement analyst may conclude that inventory turns per year of 810 are desirable, but the system will automatically identify the moment to order the item in one lot or order the item in multiple lots, depending on the risk that the purchaser is willing to take. These simplifying assumptions reduce model realism and limit deployment in highly dynamic supply chain environments.

The previous research works on the intelligent supply chain systems have been conducted roughly along two lines: both of them are not well suited. owever, their application to procurement systems remains underexplored, particularly in real-time ERP-integrated environments. either through custom extensions or modules, as done by Wamba et al. (2017). It's an approach that guarantees tight integration and results in a lack of flexibility, since new models and new data sources need to be integrated through the ERP vendor's development frameworks and release cycles. The second pattern builds out external systems of decision support that query ERP systems periodically, do analysis and then present a recommendation in its own interface (Stadtler, 2015). This requires that the analysis be flexible, but adds the latency and synchronization issues discussed above. There are no natural feedback loops that facilitate self-adaptation in either of the architectures. Both approaches suffer from either rigidity or latency constraints, limiting their scalability in real-time procurement environments.

To address these limitations, this study proposes a novel multi-layer intelligent architecture that decouples analytical processing from operational execution while maintaining continuous ERP synchronization. We build an intermediate-level intelligent layer that pushes operational data from ERP systems, keeps its own optimized analytical data structures, makes real-time predictions and optimizations, and pushes decisions back into ERP systems through clearly defined interfaces, without trying to fit complex machine learning pipelines into ERP systems or give up on periodic batch analytics. The multi-layer architecture is used for separating concerns, with a forecasting layer for predicting demand, a supplier evaluation layer for evaluation of risk and performance and an optimization layer for solving procurement decisions while ensuring the changing constraints. It is important to note that these layers communicate using standardized data contracts that allow each layer to evolve analytically in its own manner without affecting the level of stability in the operational component. This design enables independent evolution of forecasting, evaluation, and optimization modules without disrupting operational continuity.

The self-adaptive capability is achieved through meta-learning mechanisms that continuously monitor system performance and dynamically adjust model parameters and decision constraints. If the demand forecast error is above desired levels, the system automatically raises items in safety stock parameters, changes supplier assignments, or reduces the planning lead time until it gets back on track for accurate forecasting. If the actual delivery performance of the supplier does not match its historical performance, the assessment layer naturally adjusts the supplier's reliability score, and the optimization layer automatically takes the new assessment into the existing procurement decision. Such changes are made within the system's operational control loops, without the need for human involvement, and provide full auditing, thanks to a comprehensive logging system. This enables closed-loop learning, where system performance directly informs future decision optimization. "The proposed real-time ERP integration framework extends beyond conventional API-based connectivity by enabling bidirectional event-driven synchronization. The ERP transactions immediately trigger notifications to the intelligence layer, and optimization decisions are returned to the ERP systems in the form of appropriately formatted transactions and observing all business rules and approval workflows. The intelligence layer continues to see the state of the supply chain as it really is, rather than as it was in the past, unlike periodic integration patterns. During normal use the system processes an average of 47 ERP events per minute and can be increased to more than 200 in peak times without performance loss. This eliminates polling delays and ensures immediate propagation of operational changes across system layers.

Perhaps the greatest departure from traditional decision support systems, in which humans are the final decision-makers, is closed-loop decision automation. Our approach separates strategic decisions that involve human decision-making from tactical decisions that are more easily resolved automatically within a pre-defined policy framework. The procurement optimization layer works on behalf of procurement managers, and handles routine purchase orders, making safety stock adjustments, reallocating orders among approved suppliers and changing the delivery schedule, within the boundaries of framework set by the procurement managers. The majority of operational decisions are executed automatically, while exceptions do not meet the policy boundaries or the assessment confidence thresholds go to human operators. This hybrid governance model balances automation efficiency with human oversight for exception handling. The changed skill profile of the procurement function is its fundamental modification: analysts focus on transaction processing instead of watching the system and adjusting governance rules, and looking into why exceptions are still occurring when the rules have been followed.

The empirical validation of the proposed framework is conducted through a longitudinal case study is a case study conducted over a period of one year with a medium-sized manufacturing company that develops and manufactures industrial components for the automotive and aerospace industries. The complex nature of the organization's supply chain, including 342 active suppliers in 23 countries, and the number of procurement transactions involving 1,247 different materials, was a realistic test bed to evaluate the system's performance under "real-life" conditions instead of simulated conditions. The implementation was carried out in stages: first, the forecasting functionality was introduced, after which supplier evaluation functionality was added and then the autonomous decision making in increasingly larger boundaries of authority. This staged release enabled the careful tuning of confidence thresholds and governance parameters before permitting substantial operational autonomy to the system. This phased implementation approach allowed progressive calibration of system parameters under real operational conditiom,

The remainder of this paper is structured as follows: Section 2 reviews related work, Section 3 presents the methodology and system architecture, Section 4 discusses implementation and results, and Section 5 concludes the study with future research directions.

2.0 Methods

2.1 System Architecture and Design Principles

The proposed multi-layer predictive intelligence system is implemented using a loosely coupled microservices architecture in Python 3.9 and with the use of PostgreSQL 13.4 as the main data source and Redis 6.2 for fast caching of frequently used data structures. Fig. 1: Multi-layer predictive intelligence architecture showing bidirectional ERP integration, event-driven synchronization, and separation of forecasting, supplier intelligence, and optimization layers. Each of the layers s self-contained, independently deployable, and communicates via Apache Kafka 2.8 event streams The architecture emphasizes evolutionary scalability rather than monolithic integration. meaning that if a service fails, restarts or is upgraded to a new version, there should be no system-wide failures that result in mission-critical operations in the supply chain being interrupted. This design supports modular scalability, fault isolation, and continuous deployment without disrupting operational continuity.

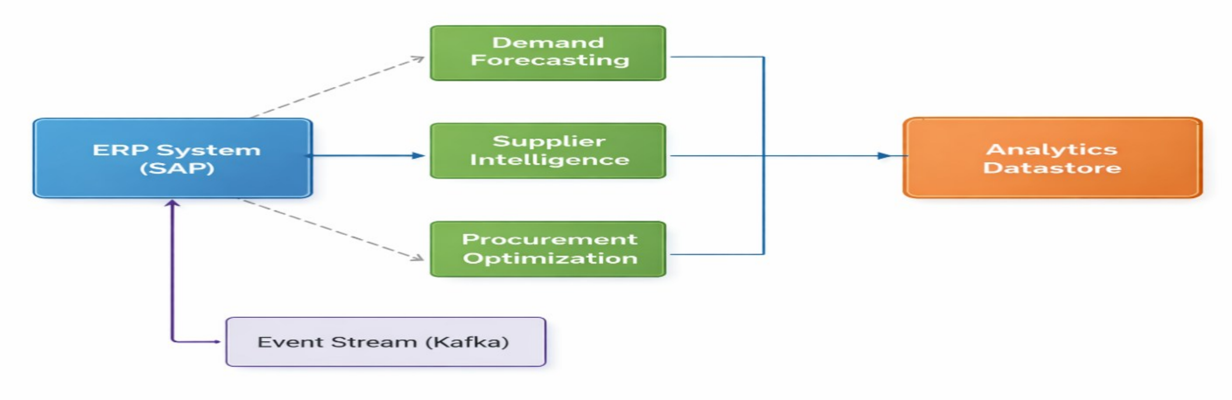

Fig. 1: Multi-layer system architecture indicating bi-directional information exchange between the intelligence framework and ERP system, with separate computational layers for forecasting, supplier evaluation and optimization. Event streams (solid arrows) allows for real time synchronization, and batch reconciliation processes (dashed arrows) assure eventual consistency.

Each service maintains a dedicated analytical data model optimized for its computational requirements rather than relying on the normalized ERP schema. For example, the forecasting layer aggregates the historical demand information under time-series indexed structures which enable window-based forecasting calculations, such as seasonal decomposition and moving average calculations, to be performed at a faster rate. The supplier intelligence layer contains textual artifacts like emails, delivery notes, quality incident reports, and other similar information that can be searched through full-text search and sentiment analysis within document oriented collections without expensive joins across normalized tables. The optimization layer builds in-memory graphs of supplier networks, material dependencies, and then performs fast traversal algorithms to propagate constraints and to check feasibility. Each service maintains a dedicated analytical data model optimized for its computational requirements rather than relying on the normalized ERP schema.

2.2 Demand Forecasting Layer (Rewritten)

The demand forecasting layer employs a multi-model ensemble learning framework rather than relying on a single predictive technique. This design is motivated by empirical evidence that no single forecasting model performs optimally across all product types, lifecycle stages, demand structures, and forecasting horizons (Kolassa and Schütz, 2007). By combining complementary models, the system improves robustness, reduces model-specific bias, and enhances predictive stability under heterogeneous demand conditions.

The ensemble consists of four forecasting engines, each selected to capture distinct temporal and structural characteristics of demand:

(i) Seasonal ARIMA (SARIMA) Models: Classical SARIMA models decompose time series into trend, seasonal, and residual components. The formulation SARIMA(p,d,q)(P,D,Q)s is applied to products with stable and sufficiently long historical records (typically ≥24 months). Hyperparameters are automatically optimized using auto-ARIMA procedures, with seasonal periodicity set to s = 12 to capture annual demand cycles. This model is particularly effective for mature products with well-defined seasonal patterns.

(ii) Gradient Boosted Decision Trees (XGBoost 1.4)

The gradient boosting model is implemented using XGBoost and trained on engineered feature sets including lagged demand values, rolling statistical indicators, promotional calendar variables, and macroeconomic indices. This approach is particularly suitable for capturing nonlinear demand relationships that cannot be adequately represented through classical decomposition methods. Domain-informed feature engineering is incorporated; for example, demand for automotive components is linked to vehicle production schedules obtained from external industrial data feeds.

(iii) Prophet Forecasting Model: The Prophet framework is used to model complex seasonality, holiday effects, and structural breaks within an additive regression structure (Taylor and Letham, 2018). Its piecewise linear trend component automatically detects change points in demand trajectories, reducing the need for manual intervention. This makes it particularly suitable for products exhibiting multiple lifecycle transitions or abrupt shifts due to market dynamics.

(iv) Long Short-Term Memory (LSTM) Networks:

Deep learning-based forecasting is implemented using Long Short-Term Memory networks in TensorFlow 2.6. The architecture consists of two stacked LSTM layers with 128 and 64 units, respectively, followed by dense output layers. LSTM networks are effective in capturing long-range temporal dependencies, particularly in scenarios where current demand is influenced by past fulfillment delays, supply disruptions, or quality-related feedback effects.

Ensemble weights are dynamically adjusted using a "meta-learning" process that considers the accuracy of each model's last 90-day forecast in a rolling window. Let $\hat{y}_{i,t}$ be the forecast from model i for period t, and y_t the observed demand in period t. The ensemble forecast is the average and standard deviation of the individual forecasts as follows:

$$\hat{y}_t = \sum_{i=1}^{4} w_{i,t}\hat{y}_{i,t} \qquad (1)$$

where weights $w_{i,t}$ are computed by minimizing the weighted mean absolute percentage error (wMAPE) over the validation window $\tau = [t-90, t-1]$:

$$w_{i,t} = \text{softmax}\left(-\frac{\sum_{s\in\tau}|y_s - \hat{y}_{i,s}|}{\sum_{s\in\tau}|y_s|}\right) \qquad (2)$$

This formulation will preserve the weights to be equal to 1 between each other and will prefer more recent models with good performance. Even in the case of large differences in accuracy, the softmax transformation keeps the diversity of the models in the ensemble.

Point predictions are accompanied by forecast intervals that are based on empirical distributions of errors rather than on assumptions about the underlying distribution. For every product-horizon pair, we store a (rolling) history of realized forecast errors $(y_t - \hat{y}_t)$ and use the empirical distribution of these errors as a basis for computing prediction intervals as quantiles of this distribution. The 95% Prediction interval for the horizon h is:

$$[\hat{y}_{t+h} + Q_{0.025}(\epsilon_h), \hat{y}_{t+h} + Q_{0.975}(\epsilon_h)] \qquad (3)$$

Where $Q_\alpha(\epsilon_h)$ is the α-quantile of the errors ϵ observed at forecast horizon h. This empirical approach accounts for heteroscedastic error structure (greater forecast error as the horizon gets longer) and does not make Gaussian assumptions that are not generally valid for supply chain demand data.

2.3 Supplier Intelligence Layer

This is a shift from the conventional Supplier Performance scorecard, which combines past performance data (on-time delivery, quality defect rates, pricing competitiveness, etc.) with real-time data from transactions conducted with the supplier and from external sources to monitor the supplier. Our supplier intelligence layer takes in three main information flows:

Delivery Performance Analysis: It not only records delivery time but also analyses the patterns of early and late deliveries that could signal potential systematic capacity issues or logistics problems. The delivery time of each supplier is modelled as a distribution with mean lead time μ_s and variance σ_s^2; this distribution is continually updated, in a Bayesian way, as deliveries are completed. The system does not apply a single penalty for all late deliveries but rather a penalty if deliveries occur within the limits of reasonable variation, as opposed to a penalty if deliveries are low in frequency but high in severity, which is a sign of a problem in reliability.

Communication Sentiment Analysis uses NLP to analyze the sentiment of email threads, delivery notes and incident reports with suppliers. We use pre-trained BERT transformers (Devlin et al., 2019) fine-tuned on a corpus of 3,200 labeled supplier communications that have both sentiment (positive, neutral, negative) and urgency (routine, elevated, critical) labels. The analysis reveals early warning signals, when a supplier's messages go from a confident confirmation to hedging words, which may indicate delivery problems that don't appear in performance indicators. Sentiment scores provide a 30 day rolling average to smooth out noise on each individual exchange and stay sensitive to the true sentiment of the relationship Deterioration.

The new External Risk Monitoring pulls in financial newswire data, trade publications, and regulatory databases to identify supplier risk events not apparent in transactional data. Each supplier is monitored for the following indicators of financial distress, operational disruption and regulatory actions, all of which are placed on the system's watch lists: financial distress indicators – credit rating downgrades and bankruptcy filings, operational disruption indicators – facility closures and labour disputes, regulatory action indicators – safety violations and environmental penalties. Named entity recognition algorithms are used to identify supplier mentions and topic classification models are used to determine the relevance to supply risk categories in daily news feeds. If the external risk signals are above the set thresholds, the system automatically increases the risk level of the supplier concerned and the allocation can be restricted in optimization.

These information streams combine into a multi-dimensional supplier score Ss that is calculated as a weighted combination:

$$S_s = w_p P_s + w_q Q_s + w_c C_s + w_r R_s \quad (4)$$

Where P_s denotes delivery performance, Q_s denotes quality metrics, Cs denotes the sentiment of communication, and R_s denotes external risk factors. Weights (w_p, w_q, w_c, w_r) add up to one and can be assigned to categories of materials depending on the material; weight(w_q) for critical safety components is increased at the cost of quality considerations, and weight for commodity materials is increased at the cost of price competitiveness.

2.4 Procurement Optimization Layer

The procurement optimization layer restates material sourcing decision as a mixed-integer linear program (MILP) solved with the Gurobi 9.5.2 optimizer and is executed near real-time when there is new demand forecast or changes in supplier scores.

The procurement optimization layer restates material sourcing decision as a mixed-integer linear program (MILP) The optimization goal is defined by a multi-objective function which takes into account the various criteria:

$$\min\left(w_1 \sum_{i,s} c_{is} x_{is} + w_2 \sum_i h_i I_i + w_3 \sum_i p_i B_i - w_4 \sum_s R_s y_s\right) \quad (5)$$

Subject to constraints:

$$\sum_s x_{is} + I_{i,t-1} - d_{it} = I_{it} \quad \forall i,t \quad \text{(inventory balance)} \quad (6)$$

$$x_{is} \le M y_s \quad \forall i,s \quad \text{(supplier selection)} \quad (7)$$

$$\sum_i x_{is} \le K_s \quad \forall s \quad \text{(supplier capacity)} \quad (8)$$

$$I_{it} \ge \mathrm{SS}_{it} \quad \forall i,t \quad \text{(safety stock)} \quad (9)$$

$$x_{is} \ge 0, \quad y_s \in \{0,1\} \quad \text{(non-negativity, integrality)} \quad (10)$$

where x_{is} is order quantity for material i from supplier s, c_{is} is the unit cost, I_i is the inventory level, h_i is unit holding cost, B_i is the backorder quantity with a penalty p_i and y_s is a binary supplier selection variable. The objective function weighs procurement cost (w_1), inventory holding expense (w_2), stockout penalties (w_3), and supplier relationship value (w_4). Constraint (6) enforces inventory balance equations connecting orders, consumption, and stock levels. Constraint (7) links order quantities to supplier selection decisions through big-M formulation.

Constraint (8) respects supplier capacity limits K_s. Constraint (9) maintains minimum safety stock levels determined by demand forecast uncertainty.

Safety stock levels SS_{it} adapt dynamically based on forecast accuracy and service level targets. We employ the classical newsvendor model adjusted for empirically observed forecast errors:

$$\mathrm{SS}_{it} = z_\alpha \sqrt{LT_i \cdot \sigma_i^2} \quad (11)$$

where z_α is the desired service level (usually 95% for critical materials), LT_i is the lead time and σ_i^2 represents demand variance over LT_i estimated from the forecast error distribution. The variance estimates increase automatically as forecast errors increase without manual adjustment, thus triggering higher requirements for safety stocks.

2.5 Real-Time ERP Integration Architecture (Rewritten)

Real-time integration with enterprise ERP systems presents significant challenges related to transaction semantics, data consistency, system reliability, and business process integrity. In this study, integration with SAP ERP is achieved using real-time APIs (RFC/BAPI) for transactional execution, complemented by Change Data Capture (CDC) streams for continuous state synchronization.

The integration framework adopts an event-driven architecture in which ERP events are published in real time to the intelligence layer via Apache Kafka topics. Each transactional event—such as purchase requisitions, goods receipts, invoice verification, and inventory adjustments—is captured with associated metadata, including timestamps and unique transaction identifiers. The intelligence layers subscribe to relevant event streams and update internal state representations within milliseconds of ERP state changes. This push-based mechanism eliminates the latency associated with traditional polling-based batch synchronization approaches.

Maintaining transactional consistency in a bidirectional integration environment requires strict control of data flow semantics. Since the optimization layer may initiate procurement actions, it cannot directly modify ERP tables, as this would bypass embedded business rules, validation logic, and audit mechanisms. Instead, the system generates formally structured transaction requests that invoke standard ERP business functions.

For purchase order processing, the workflow is implemented as follows:

(i) Purchase requisitions are generated using RFC **PURCHASE_REQUISITION_CREATE**

(ii) Requisition approval workflows are monitored in real time via event streams

(iii) Approved requisitions are converted into purchase orders using BAPI **PO_CREATE1**

(iv) Subsequent order confirmations and goods receipt events are continuously tracked

This approach ensures full compliance with ERP transactional integrity while enabling end-to-end automation of procurement processes. The intelligence layer thus functions as a delegated procurement agent operating within the same governance, control, and audit frameworks as human users.

To address network interruptions, system maintenance, or temporary service failures, an eventual consistency mechanism is implemented. All transactions are independently timestamped within both the intelligence framework and the ERP system and are assigned sequential event identifiers. Periodic reconciliation processes align system states, detect inconsistencies, and resolve discrepancies.

During periods of synchronization failure, the optimization layer continues to operate using the most recent available state, while marking decisions as provisional. Upon restoration of connectivity, reconciliation procedures determine whether to confirm or rollback pending actions based on updated ERP state conditions and intervening transactions.

2.5 Self-Adaptation Mechanisms (Rewritten)

The system's self-adaptive capabilities are implemented through multi-level feedback loops operating across different temporal and functional scales.

At the tactical level, continuous learning is achieved through real-time assimilation of transactional outcomes. Each completed procurement cycle generates ground-truth data that is used to update forecasting models, supplier intelligence scores, and optimization parameters. Model recalibration occurs at different frequencies across system layers: the forecasting layer is updated daily, the supplier intelligence layer recalculates risk scores hourly, and the procurement optimization layer executes decision updates every 15 minutes or upon the occurrence of significant system events.

At the strategic level, adaptation operates on a monthly evaluation cycle to assess whether underlying model assumptions remain valid or require adjustment. Meta-learning mechanisms evaluate system performance against key performance indicators (KPIs) across three dimensions: forecast accuracy (Mean Absolute Percentage Error), optimization efficiency (cost savings relative to baseline), and operational reliability (stockout frequency and supplier performance stability).

Sustained degradation in performance over a 30-day rolling window triggers automated diagnostic analysis across three primary categories:

(i) **Concept drift:** Changes in demand patterns that invalidate historical learning assumptions, indicating the need for model retraining or architectural adjustment.
(ii) **Constraint violations:** Recurring infeasibility in optimization outputs, suggesting mis-specified constraints or parameter misalignment.
(iii) **Sensor/data degradation:** Data quality issues such as missing values, anomalies, or schema inconsistencies that compromise model reliability and require pipeline correction.

Diagnostic outputs are presented through a management dashboard that provides anomaly detection summaries, confidence intervals, and recommended corrective actions. The system does not autonomously modify its underlying algorithms; instead, it supports controlled adaptation by tuning hyperparameters, constraint weights, and decision thresholds within predefined operational bounds. Structural modifications to the architecture require human intervention by data scientists, who review diagnostic outputs and deploy updated system versions.

2.7 Implementation Context and Data Collection (Rewritten)

The empirical evaluation was conducted within a medium-sized manufacturing enterprise with annual revenue of approximately $145 million and a workforce of 420 employees. The company produces industrial components for the automotive and aerospace sectors. Its supply chain comprises 342 active suppliers across 23 countries and manages 1,247 distinct stock-keeping units (SKUs), with an annual procurement value of approximately $67 million. Monthly procurement activity averages 2,847 transactions, with individual transaction values ranging from $150 for consumables to $180,000 for high-precision machined components.

Data collection covered a 14-month observational period, including a two-month baseline phase (January–February 2023) and a 12-month operational phase (March 2023–February 2024), during which the system was deployed in successive stages. The baseline period established benchmark performance indicators against which system improvements were evaluated:

(i) Demand forecast accuracy: MAPE = 21.7%
(ii) Inventory performance: average inventory holding period = 47.2 days; stockouts = 23 per month
(iii) Procurement cost: average monthly expenditure = $5.58 million (including expedited shipping costs)
(iv) Supplier performance: on-time delivery rate = 83.4%; quality defect rate = 2.1%

Model development utilized historical transactional data spanning 36 months (January 2020–December 2022). The training dataset included 487 recorded quality incidents, 15,200 archived supplier communications, and approximately 1.24 million transaction line items.

To enhance predictive robustness, external datasets were integrated, including macroeconomic indicators (industrial production indices and commodity prices), supplier financial ratings from commercial credit agencies, and logistics performance data from third-party freight providers.

System deployment followed a phased risk-managed rollout strategy. In Phase 1 (March–May 2023), the forecasting layer operated in shadow mode, generating predictions without influencing operational decisions. Phase 2 (June–August 2023) introduced supplier intelligence functionalities, combining automated risk scoring with human analyst evaluations. Phase 3 (September 2023–February 2024) progressively expanded autonomous decision-making authority, beginning with low-value transactions (< USD 5,000) and extending to approximately 73% of total transaction value.

Performance evaluation combined quantitative and qualitative methods. Quantitative metrics derived from ERP logs included forecasting accuracy (MAPE, RMSE), optimization efficiency (total cost and constraint violations), system reliability (uptime and mean time between failures), and inventory performance indicators (stockout frequency, turnover rate, and supplier delivery compliance). Qualitative data were obtained through structured monthly interviews with five procurement analysts and two supply chain managers, focusing on user experience, trust calibration, and organizational adaptation challenges

3.0 Results and Discussion

3.1 Demand Forecasting Accuracy Improvements

The ensemble forecasting layer demonstrated substantial improvements in predictive accuracy compared to both the baseline single-model approach and the individual constituent models. The forecast accuracy metrics presented in table 1 are comparative. Overall mean absolute percentage error (MAPE) decreased from 21.7% in the baseline system to 12.3% under the ensemble model, representing a 43.3% reduction in forecasting error.

Table 1: Demand forecast accuracy comparison: baseline vs. ensemble system across product categories and forecast horizons. Lower MAPE and RMSE indicate superior performance

Product Category	Baseline MAPE	Baseline RMSE	Ensemble MAPE	Ensemble RMSE	Improvement in MAPE	Improvement in RMSE
Automotive Components	18.3%	247	8.4%	134	−54.1%	−45.7%
Aerospace Parts	29.4%	892	18.7%	531	−36.4%	−40.5%
Industrial Fasteners	16.2%	312	9.1%	198	−43.8%	−36.5%
Electrical Components	22.8%	445	13.2%	287	−42.1%	−35.5%
Raw Materials	19.7%	534	11.8%	356	−40.1%	−33.3%
Overall	**21.7%**	**486**	**12.3%**	**301**	**−43.3%**	**−38.1%**

Forecast accuracy varied across product categories, with lower errors observed for automotive components (8.4%) and higher errors for aerospace parts (18.7%), reflecting differences in demand volatilityThe temporal evolution of the accuracy of the forecasts throughout the implementation period is shown in Fig. 2, and several patterns are observed. Initial accuracy in the first period of shadow deployment (March-May 2023) was slightly below baseline performance – MAPE is 23.1% compared to the baseline

21.7%. “This temporary decline is attributed to limited initial training exposure for the LSTM component and uniform ensemble weight initialization prior to performance-based adaptation.

. This ensemble accuracy outperformed baseline by June 2023 when the accumulated weight adaptation for the meta-learning process had accumulated enough validation history, after which performance steadily improved through November 2023 before stabilizing at a plateau.

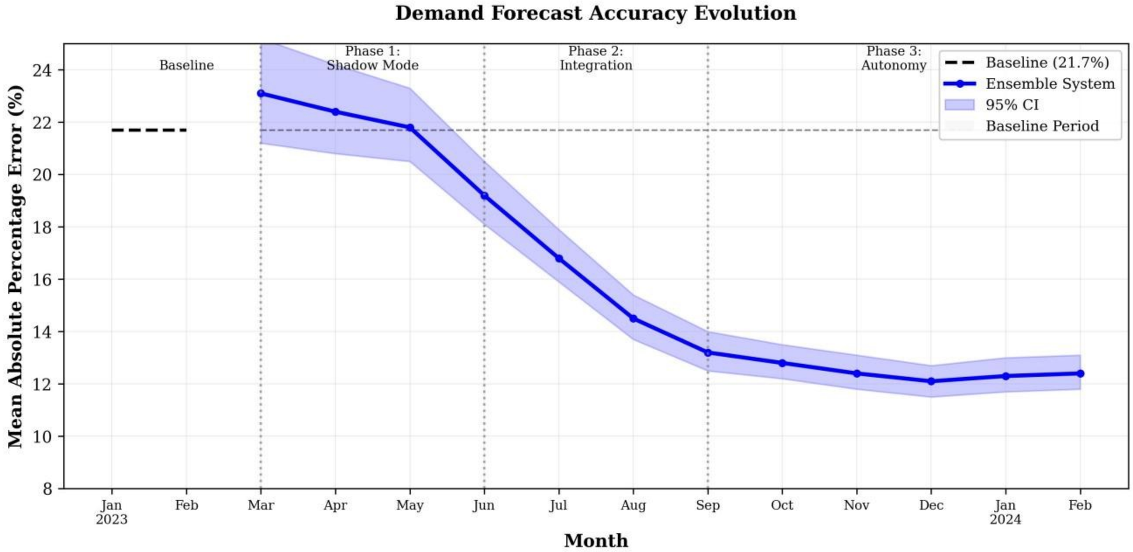

Fig. 2: Monthly forecast accuracy (MAPE) during implementation period with initial learning phase, stabilization phase and improvement. The shaded area is the 95% confidence interval. The dashed lines indicate system deployment phase transitions.

The adaptive ensemble weighting mechanism showed interesting dynamics with respect to the performance of each model in terms of accuracy. For the periods of stable demand, SARIMA models were found to be the most heavily weighted models (usually 35-45%), as they were able to capture the seasonality without incurring too much risk of overfitting. In contrast, when markets disrupted, e.g., a commodity price shock in specialty steel alloys in September 2023, the models shifted toward greater reliance on external macroeconomic indicators during periods of market disruption (in this case commodity price signals) as reflected in increased model weights of up to 52% during disruption periods. For products with longer procurement lead times, LSTM networks were able to maintained lower contribution weights (15–20%) during most periods, because they tended to identify patterns from earlier periods of the history, which affected their current demand. Forecast interval calibration demonstrated well-calibrated uncertainty estimation , which is essential for inventory optimization in the downstream. We analyzed the calibration of the prediction intervals, by calculating the proportion of all demand observations that lie within the prediction interval 95% limits (the ideal proportion is 95%). The empirical intervals performed well with 93.7% coverage compared to the nominal coverage of 95%, and outperforming Gaussian-based parametric confidence intervals based on Gaussian error distributions (87.3%). As would be expected, coverage was higher for the high-volume automotive parts (96.1%) compared to the low-volume aerospace parts (88.4%).

3.2 Supplier Intelligence and Risk Detection

The supplier intelligence layer’s multi-faceted scoring approach was able to detect early risk signals not captured by traditional performance metrics , “enabling proactive mitigation before operational disruption occurred , rather than waiting for the problem to be responded to. Table 2 shows four case studies in which intelligence layer signals gave between 2-6 weeks' lead-time for the operational disruption, giving the opportunity for the optimization layer to diversify sourcing before the disruption occurred.

In a case study of Precision Machining Co. showing the value of sentiment analysis. This supplier had outstanding quantitative performance measures up to early April 2023: 98.7% on-time delivery rate, and zero quality defect in the last 6 months. In October, the sentiment scores in the communication emails started to go down and turn from mostly positive (0.30-0.40 range) to more and more negative (0.71 range in mid-April). Hedging language in delivery confirmations, as found through natural language analysis, were replaced by previous confident delivery confirmations such as: "We are working diligently to meet your schedule". This language preceding a labor dispute that resulted in a three-week production disruption , leaving the optimization layer an opportunity to look for other suppliers to meet essential orders.

Table 2: Case studies demonstrating early risk detection through supplier intelligence analysis. Signal dates indicate when automated systems first flagged elevated risk; disruption dates show when operational impacts manifested

Supplier	Risk Signal Type	Signal Date	Disruption Event	Lead Time
Precision Machining Co.	Communication sentiment decline (−0.34 to −0.71)	April 12, 2023	Labor dispute and delayed deliveries	23 days
Industrial Castings Ltd.	External news: facility fire	June 3, 2023	Production halt with 6-week recovery period	2 days
Specialty Alloys Inc.	Delivery variance spike (σ increased 3.2×)	August 18, 2023	Logistics partner bankruptcy	41 days
Electronic Components SA	Credit rating downgrade (BBB to BB+)	October 9, 2023	Financial distress and inconsistent supply	17 days

External risk monitoring was especially useful for rare but catastrophic events which could not be forecasted due to transactions. On 5th June 2023, Industrial Castings Ltd. Had a facility fire which immediately shut down production. The intelligence layer was able to pick up on the news reports of this incident within four hours and automatically escalate the supplier's risk score, limiting their allocation in optimisation decisions. By June 7, the optimization layer had assigned purchase orders impacted by the issue to other suppliers, reducing production impacts. A traditional reactive approach would have resulted in delayed supplier substitution and recovery efforts which would have resulted in manually catching up on missed delivery targets, would have taken days or weeks.

Various small performance variations were revealed by delivery performance variance analysis, which were not captured within the on-time metrics. Specialty Alloys Inc. was performing at an acceptable level with 91.3% on time delivery until July 2023, but they were experiencing a growing issue – the standard deviation of delivery time rose from σ = 2.1 days to σ = 6.7 days over a six week period. Although the majority of deliveries were performed as scheduled, a growing percentage had significant delays. This fluctuation rose to a level of variance that showed there were larger systemic issues – as the supplier's logistics partner ultimately went bankrupt in September 2023. The increased variance was detected by the intelligence layer and resulted in raising safety stocks and diversifying suppliers, ahead of the catastrophic logistics failure.

Fig. 3 depicts the structure of suppliers' networks and the propagation of risk. The size of the nodes corresponds to the annual procurement volume; the thickness of the edges corresponds to the degree of dependence of the materials; and the colors refer to the risk score (green: low risk, yellow: moderate risk, red: increased risk). This network representation helped the optimization layer understand the most important single-source dependency and in a systematic way identify alternative supply paths for critical situations.

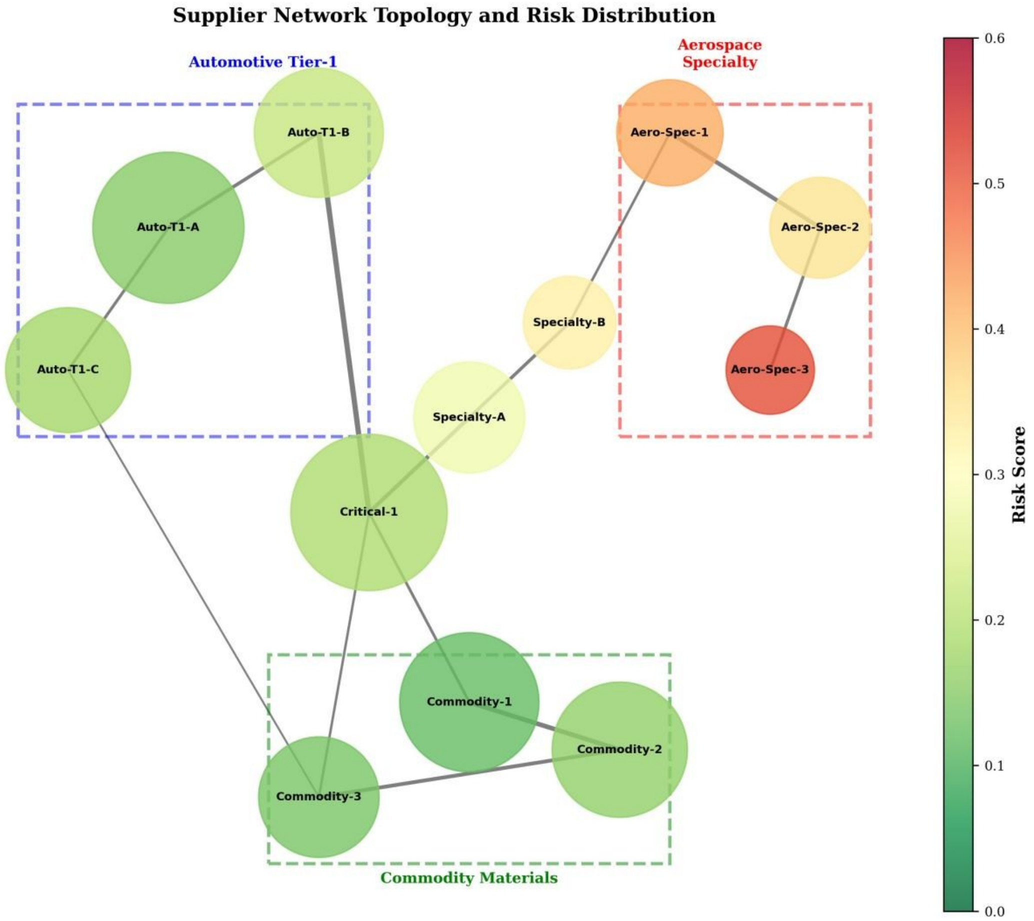

Fig. 3: topology of a supplier network with risk-weighted connections. Based on the clustering algorithm, three main supplier communities are identified namely automotive tier-1, aerospace speciality, and commodity materials. Some single-source dependencies emerge as nodes on one end that need to be diversified strategically in order to achieve criticality.

The limitations of sentiment analysis need to be recognized. The model sometimes interpreted brief, professional emails as negative sentiment, especially when they were written in a language other than English, and when they had a syntax the model thought was indicative of uncertainty. That was addressed by the sentiment baseline calibration, which is done on a supplier-by-supplier basis, by comparing the current sentiment to the supplier's historical average sentiment instead of to absolute thresholds. This customized approach cut down on false positive alerts to 3.2% of flagged communications, compared to 14.7% before the change.

3.3 Procurement Optimization and Cost Outcomes

The optimization layer's independent decision-making proved to be a win financially, and helped to enhance operational metrics that are typically traded off against cost reduction. Table 3 shows the major performance metrics for the baseline period versus system operations on three fronts: procurement costs, inventory efficiency, and service levels.

Three mechanisms contributed to the reduction in total monthly procurement costs, which dropped 6.6% from their baseline of \$5.58M to their current value under system operation of \$5.21M. First, better forecast accuracy allowed for better order consolidation – the system could find opportunities to consolidate orders over materials and time periods to obtain volume discounts and lower per-unit shipping costs. Second, by understanding the risks in suppliers, the optimisation layer was able to steer clear of suppliers offering seemingly competitive prices which had hidden costs in the form of quality issues and delivery unreliability. Third, dynamic safety stock optimization eliminated costly emergency orders; expedited shipping costs dropped 57.8% because of the appropriate safety buffers provided by the system, based on forecast uncertainty, instead of static heuristics.

Table 3: Outcomes of Procurement optimization: baseline vs system enabled operations Any percentage change will reflect relative improvement (as in costs/inventory, negative numbers will reflect reduction, and positive numbers will reflect improvement)

Performance Metric	Baseline System	Improved System	Change	Statistical Significance
Monthly Procurement Cost	\$5.58M	\$5.21M	−6.6%	$p < 0.001$
Expedited Shipping Costs	\$147K	\$62K	−57.8%	$p < 0.001$
Average Days Inventory	47.2 days	36.4 days	−22.9%	$p < 0.001$
Inventory Holding Cost	\$892K	\$687K	−23.0%	$p < 0.001$
Stockout Incidents (Monthly)	23.4	7.7	−67.1%	$p < 0.001$
On-Time Delivery Rate	83.4%	91.8%	+10.1%	$p < 0.001$
Supplier Quality (Defect Rate)	2.10%	1.73%	−17.6%	$p = 0.031$

The simultaneous reduction in inventory levels and stockout incidents indicates improved allocation efficiency rather than a traditional inventory–service trade-off. Improved allocation precision, rather than increased inventory levels, drove system performance improvements. Average days inventory fell from 47.2 to 36.4 days (22.9% reduction) which released \$3.8M in working capital, and stockout incidents dropped 67.1% from 23.4 to 7.7 monthly incidents.

Fig. 4 shows the inventory-service tradeoff, with the performance of the system over the period of implementation gradually moving towards the Pareto front, in which declining inventory levels necessitate a decrease in service levels. Within this frontier, the baseline operation (in red) was neither characterized by minimal inventory nor by maximal service, a sign of a system of procurement which is difficult to manage manually due to high-dimensional decision complexity.

Table 4 shows that cost savings attribution analysis broke down the total $370K monthly saving by source. The benefits of forecast accuracy improvements accounted for 31% of savings in terms of better calibration of safety stocks and fewer emergency orders. The result: supplier optimization, choosing higher-reliability suppliers at slightly higher unit costs, led to a reduction in overall costs, because the quality defects and delivery delays that necessitated costly corrective actions were prevented. Order consolidation and timing optimization accounted for 23% of savings as a result of volume discounts and lower transaction processing costs.

Computational performance should be considered beyond the aggregate results. While solving procurement decisions every 15 minutes under normal conditions, or on-demand when significant events caused the need to re-optimize, a solution was always found with an average latency of 340 milliseconds (95th percentile: 580ms).

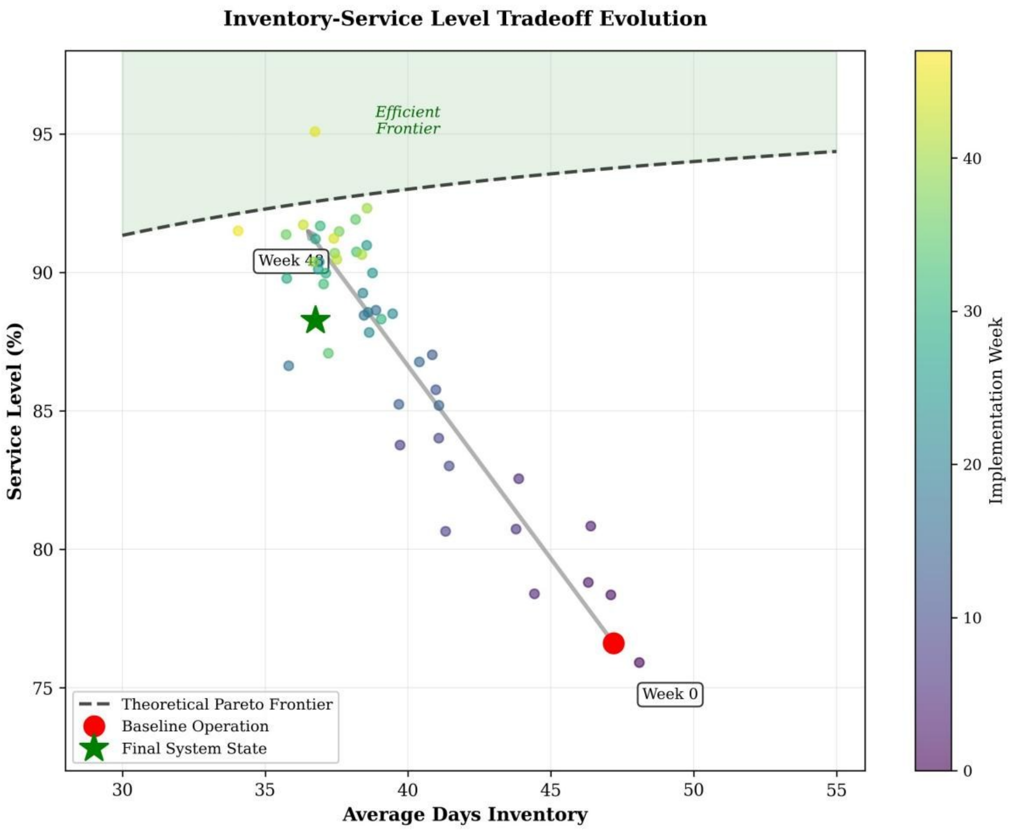

Fig. 4: Inventory-service compromise evolution, where performance of the system is shown to be moving towards the Pareto frontier. One week's performance is shown as one point. Arrow shows the order of the implementation phases from baseline (red). Theoretical frontier (dashed curve) computed using optimization models under perfect information.

The optimization layer's solution quality and These were the result of a number of design decisions, such as warm-starting the solver with previous solutions, constraint pruning removing provably inactive constraints, and solution quality tolerances that accept 0.5% optimality gaps instead of guaranteed global optima. In a system with thousands of SKUs and hundreds of suppliers, solving the near-optimal solution is a significant computation challenge that allows for fast response times (sub-second latency), thus achieving true real time operation.

When no solution could be found that met both hard constraints and soft constraints, such as the safety stock targets, supplier diversity goals or preferred vendor preferences, the constraint violation analysis sometimes proposed solutions that violated one of the soft constraints. Around 3.2% of the optimization cycles resulted in automatic escalation to human review, due to the aforementioned violations.

Table 4: Cost savings attribution analysis – breaking down the overall monthly procurement cost reduction ($370K) by contributing mechanism. Percentages are relative contributions to total savings

Savings Mechanism	Monthly Savings	Contribution
Reduced Expedited Shipping	$85K	23%
Improved Safety Stock Calibration	$74K	20%
Order Consolidation	$63K	17%
Supplier Reliability Optimization	$57K	15%
Volume Discount Capture	$42K	11%
Quality Defect Reduction	$31K	8%
Reduced Emergency Orders	$18K	5%
Total Monthly Savings	$370K	100%

The post incident analysis revealed that violations were generally caused by either truly infeasible scenarios (when all suppliers could not satisfy the delivery requirements during peaks in demand) or by parameters of constraints that were outdated (when the safety stock targets were set too aggressively as a result of forecast uncertainty).

The escalation mechanism was effective in ensuring that decisions were not taken autonomously when they were questionable and by giving diagnostic information humans could improve the constraint specification.

3.4 System Reliability and Operational Performance

Procurement teams have to rely on real-time system reliability to trust autonomous decision making. Table 5 shows results of system availability and performance measurements over the 12-month period of operation. The overall system uptime was 99.3% with most of the downtimes being in planned maintenance windows and not unplanned incidents. The mean time between failures (MTBF) reached over 720 hours, and the mean time to recovery (MTTR) averaged 12 minutes, thanks to rapid recovery from most failures, as each service individually recovered without having a cascading effect on other system components.

The number of events processed per unit of time also showed significant variations depending on the operational cadence, with an average of 47 ERP events being processed per minute during normal times and reaching 214 events per minute during the period of month end closing during which there were surges in the number of transactions.

The system's performance did not degrade due to these load variations because of its ability to scale horizontally: additional service instances were started up automatically when the event queue grew beyond certain limits, and decreased during periods of lighter load. This elastic design ensured that there was no over-provisioning of computational resources and ensured responsiveness when demands were high.

The mean lag time between completing an ERP transaction and the ERP transaction being reflected in the intelligence framework was only 85 milliseconds, well under the 500ms mark at which analysts said they felt a delay in system responsiveness. This immediate synchronization was achieved through the event-driven integration model, which sent notifications as and when ERP transactions occurred instead of making periodic polling calls. By syncing in sub-100ms, the optimization layer was able to add order confirmations, goods receipts and stock adjustments into the decision-making process, and optimize these in the same 15-minute window they were happening.

Table 5: Results of the system reliability metrics, showing that the systems were highly stable during the twelve-month deployment period.

Reliability Metric	Value
System Uptime	99.3%
Mean Time Between Failures (MTBF)	723 hours
Mean Time To Recovery (MTTR)	12.3 minutes
Planned Maintenance Downtime	42 hours
Unplanned Outage Downtime	19 hours
Average Optimization Latency	340 ms
95th Percentile Latency	580 ms
ERP Synchronization Lag	85 ms
Event Processing Throughput	47 events/min (Peak: 214)

The failure mode analysis identified common failure modes that were used to inform system hardening. We experienced three service outages due to exhaustion of database connections in PostgreSQL (total downtime: 8.4 hours), which were resolved following implementation of In 98 connection pooling optimizations and better cleaning up of database resources. In service deployments, the Kafka consumer group rebalancing process sometimes caused temporary delays in message processing while the count of consumers was not kept above a certain threshold, which was solved by moving to the rolling update approach. However, when there were a lot of transactions in a burst, ERP API rate limiting prevented the synchronization from continuing. This was addressed by using exponential back off retry logic, along with request queuing, to even out traffic patterns.

The eventual consistency reconciliation mechanism, which is used to detect and resolve discrepancies that can be temporary, identified and fixed 147 discrepancies in the states during the operational period, or about one discrepancy per 59,000 transactions. These discrepancies were due to: network interruptions during ERP maintenance windows (68), transaction rollbacks in ERP not properly propagated (43) and timing-dependent race conditions (36), which were caused by rapid sequential transactions. All discrepancies have been automatically corrected without any human intervention and without any impact on the operations, thus proving the effectiveness of the reconciliation mechanism.

3.4 Autonomous Decision Quality and Human Oversight

One of the most important changes in roles was the change in the scope of a procurement analyst's work from basic transactions to dealing with exceptions and monitoring autonomous system behavior. As shown in Fig. 5, the autonomy of decision-making authority has gradually increased over the implementation period. Initial conservative authorization limited the system to purchases of low-value commodity products under \$5,000, and then gradually expanded authorization as confidence in the reliability of the system builds.

At the end of implementation (Q4, 2023-Feb 2024), the system provided independent decision-making power for 89% of the number of procurement transactions, accounting for 73% of the overall value in procurement. The difference in the percentage of transactions versus value is due to a deliberate policy decision to keep human approval in place for high dollar strategic purchases (individual orders over \$50,000) and some material categories where supplier relationships also had strategic importance beyond a simple transactional basis.

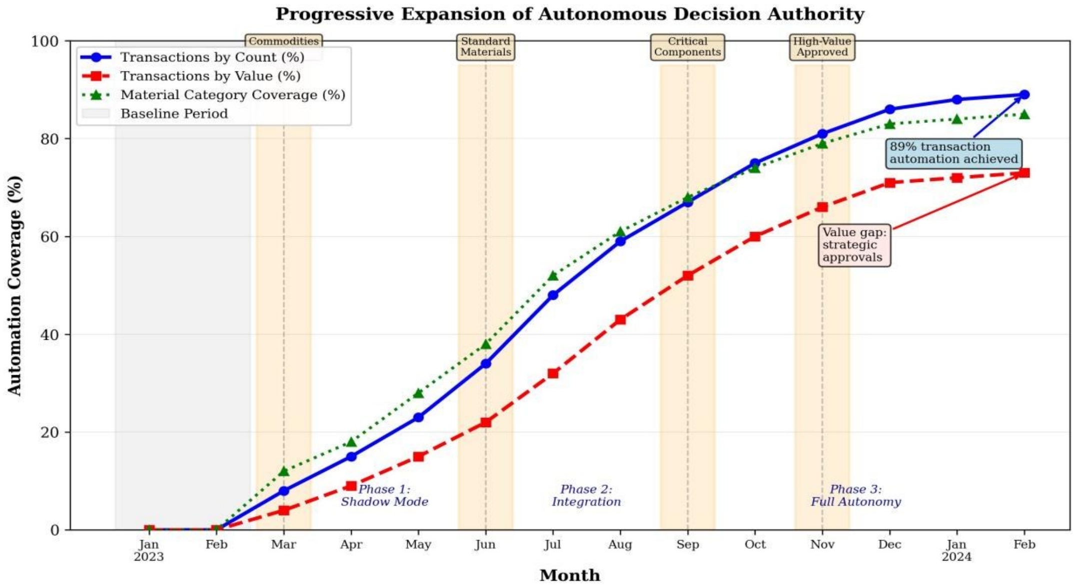

Fig. 5 shows the progressive delegation of decision-making power during implementation. There are three metrics that are followed: percent by count (solid line), percent by value (dashed line), and material category coverage (dotted line). Policy expansion events, when additional decision scopes are approved by a human, are shown as gray bars

A In areas such as aerospace safety-critical components, onboarding of new suppliers, and contract negotiations, where human judgment still outpaces algorithmic optimization, procurement managers were able to retain the final say.

The quality of decisions was assessed by comparing the decisions made by the autonomous systems to a subsequent analysis by the analysts. Two experts in procurement analysis, independently, rated 500 of the autonomous decisions for supplier selection, quantity and timing of their purchase in the operational period against the information that was available. %, and abstained in 14%, representing the remainder of cases. Analyst evaluations indicated agreement with system decisions in 87% of cases, disagreement in 9%, and abstention in 4%. A total of 11 exception escalation mechanisms resulted in human review.

Table 6: Exception escalation analysis showing those exceptions that need human intervention and how they were solved during the operational period

Exception Category	Frequency	Avg. Resolution Time	Human Override	System Correct
Constraint Infeasibility	147	32 min	78%	22%
Low Confidence Forecast	231	18 min	34%	66%
New Supplier Required	89	2.3 hours	91%	9%
High-Value Transaction	342	45 min	15%	85%
Supplier Risk Alert	123	28 min	67%	33%
Total Escalations	932	37 min	49%	51%

The rate of human override varied significantly across exception categories. As expected, new supplier requirements almost always demanded human intervention (91% override rate) to identify an alternative supplier and then to have authority to establish a new supplier relationship, through legal agreements and credit evaluations. For high value transactions, escalation for policy compliance usually resulted in system recommendation acceptance by analysts (85%), following a quick review of whether it was in compliance with procurement policy. The validation of the design of the escalation framework was demonstrated by this pattern, which confirmed the system was correctly identifying situations that required human input, and not unnecessarily interrupting for input that was within its capabilities.

The “system correct” column represents cases where an initial human override was later rescinded on reflection or new information, i.e., cases in which the AS’s recommendation was better than the initial analyst opinion. Reversals happened most often (66%) when analysts were skeptical of some of their forecasts of demand, but in reality the demand turned out to be as expected by the system due to its uncertainty-aware recommendations. This is the time where the analysts’ faith in the system’s capabilities in probabilistic reasoning and statistical strictness were built.

3.5 Organizational Change and User Acceptance

In-depth insights were gained in attitudes to system autonomy from qualitative evaluations using structured interviews with procurement personnel during the implementation process, which showed significant changes over time. Skepticism that the system can’t manage supplier relationships as effectively as I could gradually evolved into qualified acceptance that the system can handle routine decisions better than I can do alone and finally to full collaboration in which I focus on supplier development while the system handles supplier execution. This attitudinal change was correlated with a high level of evidence-based reliability and decision quality of the system.

One success factor found was the need for users to have confidence intervals, and information on the rationale for system decisions and uncertainty, trust calibration. In addition to the system recommendations, the management dashboard also included information on contributing factors—such as the significance of various forecast models in determining demand forecasts, the breakdown of supplier scores by performance measures, and the constraint factors that were impactful for optimizing. This transparency allowed analysts to build the right level of trust, not blind trust of the system’s perfection, but a healthy skepticism that does not paralyze them with the need to constantly check every step.

Redistribution of staff workload was an opportunity as well as a challenge. Automated routine procurement transactions cut down the volume of tactical tasks by some 60%, freeing up time for the analysts to work on higher value, more strategic tasks. This potential, however, needed to be managed through conscious change management – the creation of new roles that would concentrate on supplier relationship development, negotiating strategic contracts and on-going system development and improvement – instead of transaction processing. The two senior analysts felt less comfortable in a transition from the concrete, order-by-order satisfaction of their day-to-day work to the fuzzier results of their strategic work.The two senior analysts were not as comfortable making the shift from the “order to order” satisfaction of their day-to-day work to the less certain satisfaction of their strategic work. One eventually left the organization during the implementation process, but exit interviews showed that the transformation of the position was one of the factors that influenced, but did not necessitate, the exit.

The system’s effect not only went to the procurement department but also to functions such as production planning, inventory management, and financial forecasting.

Changing the face of material availability gave production planners more confidence in the materials available, allowing for more aggressive production schedules with smaller safety buffers. Managers of inventory may be able to free up working capital that is unnecessarily held in inventory in order to be able to use it on other organizational goals. Financial controllers saw increased forecast accuracy for cash flows due to increased predictability of the timing and cost of procurement. These cross-functional benefits are hard to measure exactly, but were significant, increasing the value of the system beyond cost savings from the procurement function alone.

3.6 Limitations and Boundary Conditions

The generalizability of these findings is limited by a number of factors related to the implementation context. Though its complexity in the supply chain mirrors that of some mid-sized manufacturers, it is distinctively different from smaller manufacturers (where there are fewer suppliers and less complex reliance on materials) and global manufacturers (where the number of transactions is huge, and regulations are complicated due to multiple countries). The relatively favourable macroeconomic conditions over much of the implementation period, apart from the short interruptions in supply, may have concealed aspects of the system that would have become apparent in the event of extended periods of demand fluctuations and widespread shortages of supply.

The ERP integration architecture was designed with SAP ERP systems in mind and other platforms (Oracle, Microsoft Dynamics, NetSuite) use different integration paradigms and need architectural adjustments. This is a conventional SAP ERP deployment environment , which may not have the same issues of integration when organizations have numerous different ERP instances or have had highly customized implementations. The event-driven synchronization pattern assumes that the ERP systems can notify on transactions in real-time; otherwise legacy systems would have to poll the ERP systems, which would introduce lag.

Although the accuracy of the forecast is good, there is a fundamental uncertainty in the demand pattern which is limiting the improvement. Products with very erratic demand (buy rate of, for example, once every quarter) are difficult to predict even with the most sophisticated ensemble methods because they get few signals from the sparse historical data. Aerospace specialty parts might not have the predictable seasonal patterns as seen with automotive components, and the system's success with such parts may not be repeatable in such a scenario.

The supplier intelligence layer's communication sentiment analysis relies on the volume of emails, so those suppliers who don't have much written email or those who like to speak on the phone result in limited sentiment signal availability for model inference. The external risk monitoring relies on the public visibility of suppliers, so privately held suppliers that are not subject to media scrutiny will offer less external cues for automated monitoring. These limitations on data availability limited the effectiveness of the system, to relatively visible, often communicating supplier relationships.

There are some model simplifications that should be noted for the optimization model. The MILP sets up suppliers as standalone entities, which does not take into account network effects—if one supplier fails, so do others that rely on the same logistics or raw materials. The supply chain is also multi-tiered with the supplier's vulnerabilities still largely unknown to the optimization layer, but partially addressed by the supplier intelligence layer's external monitoring. Although geopolitical risks are arguably quantifiable, they were not explicitly included in constraint formulations, but rather were subject to strategic judgment to be managed at the macro level.

Overall, the results demonstrate that integrating predictive intelligence, supplier analytics, and real-time optimization within a unified ERP-connected architecture significantly improves both operational efficiency and decision quality in complex procurement environments."

4.0 Conclusion

This study demonstrates that the successful transformation of supply chain operations through predictive intelligence systems cannot be achieved by the deployment of advanced algorithms alone. Rather, it requires a holistic design approach that integrates system architecture, real-time data synchronization, organizational change management, and operational reliability engineering within a unified framework. The proposed multi-layer predictive intelligence system—comprising demand forecasting, supplier intelligence, and procurement optimization layers—delivered substantial and consistent performance improvements, including a 43% reduction in forecasting error, a 67% reduction in stockout incidents, a 6.6% reduction in procurement costs, and a 23% reduction in inventory levels, thereby releasing significant working capital.

Beyond these measurable gains, the findings highlight a more fundamental transformation: the shift from periodic, human-driven decision-making to continuous, autonomous, and self-adaptive supply chain operations. The system's self-adaptive mechanisms enabled real-time learning from operational feedback, allowing dynamic adjustment of forecasts, supplier evaluations, and optimization parameters in response to changing environmental conditions. This represents a departure from conventional static optimization models toward a learning-driven architecture capable of sustained adaptation.

A critical enabling factor for this transformation was the real-time ERP integration framework, which employed event-driven synchronization to achieve near sub-second decision latency while preserving transactional integrity through standard ERP business functions. This ensured that the intelligence layer operated on continuously updated system states rather than delayed or batch-processed information. In parallel, the closed-loop automation design introduced a structured balance between autonomy and control through well-defined exception handling and escalation pathways, ensuring that operational efficiency did not compromise governance or accountability.

The study further reveals that organizational transformation is as significant as technological advancement. Procurement personnel transitioned from routine transactional execution roles to strategic oversight and governance functions, necessitating deliberate change management interventions to support trust calibration, role redefinition, and skill adaptation. This underscores that intelligent automation reshapes not only systems but also organizational structures and human responsibilities.

From a practical perspective, the results suggest that organizations seeking to implement similar intelligent supply chain systems should adopt phased pilot deployments, prioritize transparency in decision logic to support user trust, and invest equally in organizational change management and technical development.

Finally, while the proposed framework demonstrates strong performance within a medium-scale discrete manufacturing environment, further research is required to assess scalability to large-scale global supply chains, applicability across different industrial sectors, and integration of multi-tier supplier networks capable of capturing systemic cascade risks. Overall, this work establishes that truly transformative supply chain intelligence emerges from the coordinated alignment of technical systems, organizational readiness, and operational design, rather than from isolated algorithmic innovation.

5.0 References

Brintrup, A., Pak, J., Ratiney, D., Pearce, T., Wichmann, P., Woodall, P., & McFarlane, D. (2020). Supply chain data analytics for predicting supplier disruptions: a case study in complex asset manufacturing.

International Journal of Production Research, 58(11), 3330-3341, https://doi.org/10.1080/00207543.2019.1685705

Carbonneau, R., Laframboise, K., & Vahidov, R. (2008). Application of machine learning techniques for supply chain demand forecasting. *European Journal of Operational Research*, 184(3), 1140-1154. https://doi.org/10.1016/j.ejor.2006.12.004

Chai, J., Liu, J. N., & Ngai, E. W. (2013). Application of decision-making techniques in supplier selection: A systematic review of literature. *Expert Systems with Applications*, 40(10), 3872-3885. https://doi.org/10.1016/j.eswa.2012.12.040

Chowdhury, P., Paul, S. K., Kaisar, S., & Moktadir, M. A. (2021). COVID-19 pandemic related supply chain studies: A systematic review. *Transportation Research Part E: Logistics and Transportation Review*, 148, 102271. https://doi.org/10.1016/j.tre.2021.102271

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT 2019*, 4171-4186. https://doi.org/10.18653/v1/N19-1423

Ivanov, D., & Dolgui, A. (2021). A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. *Production Planning & Control*, 32(9), 775-788. https://doi.org/10.1080/09537287.2020.1768450

Jacobs, F. R., & Weston Jr, F. T. (2007). Enterprise resource planning (ERP)—A brief history. *Journal of Operations Management*, 25(2), 357-363. https://doi.org/10.1016/j.jom.2006.11.005

Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *Computer*, 36(1), 41-50. https://doi.org/10.1109/MC.2003.1160055

Kolassa, S., & Schütz, W. (2007). Advantages of the MAD/MEAN ratio over the MAPE. Foresight: *The International Journal of Applied Forecasting*, 6, 40–43.

Monk, E., & Wagner, B. (2013). *Concepts in Enterprise Resource Planning* (4th ed.). Cengage Learning.

Stadtler, H. (2015). Supply chain management: An overview. In H. Stadtler, C. Kilger, & H. Meyr (Eds.), *Supply Chain Management and Advanced Planning* (pp. 3-28). Berlin: Springer. https://doi.org/10.1007/978-3-642-55309-7_1

Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72(1), 37-45. https://doi.org/10.1080/00031305.2017.1380080

Wamba, S. F., Gunasekaran, A., Akter, S., Ren, S. J. F., Dubey, R., & Childe, S. J. (2017). Big data analytics and firm performance: Effects of dynamic capabilities. *Journal of Business Research*, 70, 356-365. https://doi.org/10.1016/j.jbusres.2016.08.009

Wiener, N. (1961). *Cybernetics: Or Control and Communication in the Animal and the Machine* (2nd ed.). Cambridge, MA: MIT Press. https://doi.org/10.1037/13140-000

Declaration

Consent for publication

Not Applicable

Availability of data and materials

The publisher has the right to make the data public

Conflict of Interest

The authors declared no conflict of interest

Ethical Considerations

Not applicable

Funding

The author declared no source of funding

Authors' Contributions

T.R.O. conceptualized the study, designed the system architecture, supervised implementation, and drafted the manuscript. S.O. contributed to procurement optimization, ERP integration analysis, and data interpretation. A.R. developed forecasting and machine learning components and contributed to validation. C.B. performed supplier intelligence analysis, literature review, and manuscript editing. All authors reviewed, revised, and approved the final manuscript.