Predictive and Prescriptive Logistics Optimization Using Hybrid AI, Time-Series Analytics, and Synthetic Data: A Case Study

Lakshmi Devi Pujari[1], Sridhar C. Naga Venkata[2], Saayee Saahit CNV[3], Swetha Reddy Ravula[4]

[1]Professor, Dept. of ECE, St. Peter's Engineering College(A), Hyderabad -500014.

[2] Product Innovation Manager, Hyderabad, India.

[3]Scholar, Dept. of Computer Science, NW Missouri State University, Missouri, USA

4. Visiting Asst., Professor, Bucknell University, USA

DOI: https://doi.org/10.47772/IJRISS.2026.1014MG0006

Received: 01 January 2026; Accepted: 07 January 2026; Published: 16 January 2026

ABSTRACT

Global logistics networks face increasing volatility driven by geopolitical tensions, climate disruptions, demand variability, and operational uncertainty. Although artificial intelligence has improved predictive capabilities in logistics, classical and standalone learning models remain limited by data sparsity, non-stationarity, and scalability constraints. This study proposes a hybrid logistics intelligence framework that integrates time-series forecasting, synthetic data generation, and AI-based optimization. The framework is designed to enhance forecasting robustness and translate predictions into actionable operational decisions. A FedEx case study demonstrates how historical shipment data, real-time telemetry, and synthetically generated disruption scenarios can be jointly leveraged to improve demand forecasting, routing efficiency, and service reliability. Performance is evaluated across real, simulated, and hybrid datasets. Results show that the proposed approach consistently outperforms traditional statistical and machine-learning methods in accuracy, robustness, and operational scalability.

Keywords: Logistics Optimization, Time Series Forecasting, Synthetic Data, Hybrid AI Algorithms, Supply Chain Intelligence, FedEx

INTRODUCTION

Modern logistics and supply chain systems operate amid pronounced uncertainty arising from demand volatility, capacity constraints, adverse weather, and evolving geopolitical and regulatory pressures. Large-scale logistics providers generate massive volumes of heterogeneous data from daily shipment operations, encompassing structured records, semi-structured transactions, and unstructured sensors, textual, and geospatial sources. Effectively leveraging this data for decision support is essential to achieving operational efficiency, resilience, and sustained competitiveness.

Despite advances in artificial intelligence–driven logistics platforms, persistent challenges remain in anticipating sudden demand surges, optimizing routing decisions under uncertainty, and maintaining service-level agreements during rare but high-impact disruptions. Events such as pandemics, extreme weather, labour shortages, and policy shifts frequently invalidate assumptions underlying conventional decision-support systems. Traditional time-series forecasting methods, while effective in stable environments, are limited by assumptions of stationarity and are ill-suited to capturing nonlinear dynamics and rare “black-swan” events. Conversely, data-driven AI models, though expressive, often suffer from data sparsity, distributional bias, and limited generalization when disruptive events are underrepresented.

Problem Statement

This research addresses the absence of robust forecasting and optimization frameworks capable of supporting large-scale logistics operations under data scarcity and extreme uncertainty. The central research question is:

"*How can logistics operators improve forecasting accuracy, operational robustness, and prescriptive optimization performance when historical data alone is insufficient to represent rare disruptions and complex demand dynamics?*"

To answer this question, the study examines the limitations of standalone forecasting models, evaluates the role of synthetic data generation in enhancing generalization, and assesses hybrid AI frameworks that integrate statistical forecasting, deep learning, and optimization heuristics. The objective is to develop a scalable and resilient logistics intelligence framework that supports predictive and prescriptive decision-making while maintaining service reliability under uncertainty.

Related Work

Logistics Forecasting under Uncertainty and Non-Stationarity

Demand forecasting is a central component of logistics and supply chain management, yet it remains challenging due to non-stationarity, seasonality, and exogenous disruptions. Early studies primarily relied on statistical time series models such as ARIMA and exponential smoothing, which demonstrated effectiveness under stable demand patterns but degraded significantly in volatile environments and during regime shifts (Hyndman et al., 2015; Carbonneau et al., 2008). As logistics networks expanded in scale and complexity, machine learning approaches gained attention for their ability to capture nonlinear dependencies across multiple time series (Carbonneau et al., 2008; Bandara et al., 2019).

Deep learning models, including recurrent neural networks (RNNs) and long short-term memory (LSTM) networks, further improved forecasting accuracy by learning temporal dependencies directly from data (Sezer et al., 2020). Probabilistic deep forecasting approaches, such as DeepAR, introduce uncertainty-aware predictions, enabling logistics planners to reason about risk rather than relying solely on point estimates (Salinas et al., 2020). Nevertheless, recent evaluations indicate that these models remain sensitive to distributional shifts and often underperform when exposed to rare or extreme events not present in the training data (Benidis et al., 2022).

Transformer-based forecasting architectures have emerged as state-of-the-art for long-horizon and multivariate time-series prediction. Models such as Informer, AutoFormer, and Temporal Fusion Transformers (TFT) demonstrated substantial gains in capturing long-range dependencies and incorporating exogenous variables (Lim et al., 2021; Zhou et al., 2021; Wu et al., 2021). Despite these advances, multiple studies note that transformer-based models remain data-intensive and struggle to generalize under sparse, rare-event conditions common in logistics systems (Benidis et al., 2022; Nie et al., 2022).

"Relevance to the problem statement: These works collectively highlight that even advanced forecasting models rely heavily on representative historical data and fail to anticipate rare yet high-impact disruptions reliably."

Rare Events, Data Sparsity, and Forecasting Robustness

Rare disruptions such as pandemics, extreme weather, labour shortages, and geopolitical shocks play a disproportionate role in logistics performance yet occur infrequently, resulting in limited historical samples. Studies on supply chain risk management emphasize that traditional forecasting pipelines are ill-suited for modeling such low-frequency, high-impact events (Ivanov et al., 2019; Hosseini et al., 2019). During the COVID-19 pandemic, multiple empirical analyses demonstrated that data-driven forecasting systems trained on pre-pandemic data failed to adapt rapidly to unprecedented shifts in demand (Ivanov & Dolgui, 2020; Queiroz et al., 2020).

Recent research identifies data sparsity and class imbalance as fundamental barriers to robust learning in logistics forecasting. Benidis et al. (2022) and Seaman et al. (2022) show that even probabilistic and ensemble models tend to underestimate tail risks when extreme events are underrepresented. These findings motivate exploring data augmentation and scenario-based learning techniques to enhance robustness.

"Relevance to problem statement: This literature directly motivates the need for approaches that improve forecasting accuracy and robustness when historical data alone is insufficient."

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue XIV January 2026 | Special Issue on Management

Synthetic Time-Series Data Generation for Logistics Applications

Synthetic data generation has gained increasing attention over the past decade as a strategy to address data scarcity, imbalance, and privacy constraints. Early simulation-based approaches were used in supply chain modeling to evaluate policies under hypothetical scenarios (Banks et al., 2014). More recently, generative models, particularly Generative Adversarial Networks (GANs), have enabled realistic synthesis of time-series data while preserving temporal dependencies.

TimeGAN represents a seminal contribution, demonstrating that adversarially trained networks can generate high-fidelity synthetic time series suitable for downstream forecasting tasks (Yoon et al., 2019). Subsequent studies extended GAN-based and variational approaches for industrial and operational time series, showing improved forecasting accuracy and generalization when synthetic samples are used for augmentation (Kannan et al., 2022; Chatterjee et al., 2023; Klopries et al., 2024). In supply chain contexts, synthetic data has been proposed as a mechanism for modeling extreme demand surges, stress-testing planning algorithms, and enabling safer experimentation without exposing sensitive operational data (Long et al., 2025).

"Relevance to problem statement: These works support the role of synthetic data generation in enhancing model generalization and representing rare disruption scenarios missing from historical records."

Predictive–Prescriptive Integration and Hybrid AI Optimization

While forecasting provides anticipatory insight, logistics performance ultimately depends on prescriptive decisions such as routing, scheduling, and capacity allocation. Classical optimization approaches for vehicle routing and logistics planning are computationally expensive and struggle with dynamic uncertainty (Toth & Vigo, 2014). Over the past decade, hybrid approaches combining machine learning with operations research have gained traction.

Learning-assisted optimization and reinforcement learning methods have been applied to dynamic vehicle routing problems, demonstrating improved adaptability to stochastic demand and real-time information (Nazari et al., 2018; Bogyrbayeva et al., 2022). However, purely learning-based solvers often face feasibility and scalability challenges in large logistics networks. Consequently, recent surveys advocate hybrid frameworks that integrate forecasting models, synthetic scenario generation, and optimization heuristics to support robust predictive–prescriptive decision-making (Bertsimas & Kallus, 2020; Xu et al., 2025).

"Relevance to problem statement: This body of work motivates hybrid AI architectures that combine statistical forecasting, deep learning, synthetic data, and optimization heuristics to improve end-to-end logistics performance."

Overall Methodological Framework

This study adopts a hybrid predictive–prescriptive methodology that integrates time-series forecasting, synthetic data generation, and AI-assisted optimization to address demand uncertainty, rare disruptions, and operational decision-making challenges in large-scale logistics systems. The methodology is designed to (i) improve forecasting accuracy under non-stationary conditions, (ii) enhance model robustness through data augmentation, and (iii) translate predictions into feasible operational decisions.

Let the observed historical logistics dataset be defined as

$$\mathcal{D}r = \{(xt, yt)\}\, Tt{=}1 \qquad (1)$$

where x_t denotes the vector of exogenous and operational covariates (e.g., time index, route attributes, weather indicators, hub status), and y_t represents the target logistics variable, such as shipment volume or delivery lead time, at time t. The proposed hybrid framework integrates statistical forecasting, deep learning, synthetic data augmentation, and prescriptive optimization to address uncertainty and the sparsity of rare events.

Hybrid Forecasting Model

a) Statistical Time-Series Component

A classical statistical forecasting model (e.g., SARIMA) is employed to extract baseline trend and seasonal components:

$$\hat{y}_{t(S)} = f_S\,(y_{t-1}, y_{t-2}, \dots, y_{t-p}; \theta_S) \qquad (2)$$

where $f_S(\cdot)$ denotes the statistical forecasting function and $\theta_S$represents model parameters.

b) Deep Learning Forecasting Component

A deep learning model (e.g., LSTM or Transformer-based architecture) is trained to capture nonlinear temporal dependencies and multivariate interactions:

$$\hat{y}_{t(D)} = f_D\,(x_{1:t}; \theta_D) \qquad (3)$$

where $f_D\,(\cdot)$ denotes the deep forecasting function, and $\theta_D$its learned parameters.

c) Forecast Fusion

To improve robustness and reduce model bias, the final forecast is obtained through weighted ensemble fusion:

$$\hat{y}_t = \alpha\hat{y}_t^{(S)} + (1-\alpha)\hat{y}_t^{(D)}, \alpha \in [0,1] \qquad (4)$$

This fusion strategy balances interpretability and nonlinear expressiveness.

Synthetic Data Generation for Rare Events

To mitigate data sparsity associated with rare disruptions, synthetic time-series samples are generated using a generative model $G(\cdot)$, such as a conditional GAN:

$$\tilde{y}_{1:T} = G(z, c) \qquad (5)$$

where $z \sim \mathcal{N}(0, I)$is a latent noise vector and $c$denotes conditioning variables (e.g., disruption type, region, season).

The augmented training dataset is defined as

$$\mathcal{D}_a = \mathcal{D}_r \cup \mathcal{D}_s \qquad (6)$$

where $\mathcal{D}_s$Contains synthetic disruption scenarios. Forecasting models are retrained on $\mathcal{D}_a$to enhance generalization.

Prescriptive Optimization Layer: Let

- V Denote the set of vehicles,
- $N$the set of delivery nodes,
- c_{ij}The transportation cost from the node i to node j,
- q_i The demand at the node i, and $C_v$the capacity of the vehicle v.

Objective Function

The logistics optimization objective is to minimize total transportation cost:

$$\min \sum_{v\in V}\sum_{i\in N}\sum_{j\in N} c_{ij}x_{ij}^v \qquad (7)$$

where $x_{ij}{}^v$ is a binary decision variable indicating whether the vehicle v travels from node $i$to node j.

Constraints

Demand satisfaction constraint:

$$\sum_{v\in V}\sum_{j\in N} x_{ij}^{v} = 1, \quad \forall i \in N \tag{8}$$

Vehicle capacity constraint:

$$\sum_{i\in N} q_i \sum_{j\in N} x_{ij}^{v} \le C_v, \quad \forall v \in V \tag{9}$$

Service-level (time window) constraint:

$$t_i^{arrival} \le t_i^{deadline}, \forall i \in N \tag{10}$$

Forecasts from Equation (4) serve as inputs to the optimization model, enabling dynamic routing and scheduling under uncertainty.

A Case Study with FedEx-Specific Instantiation

For the FedEx case study:

FedEx operates one of the world's largest express logistics networks, handling time-sensitive shipments across air, ground, and last-mile delivery channels. Peak seasons (e.g., holidays) and disruption events (e.g., weather or labour shortages) pose significant challenges for forecasting and routing.

- y_t: daily shipment volume per hub or service zone
- x_t: weather indicators, region, service class (Express/Ground), seasonality
- V: FedEx fleet (aircraft, trucks, last-mile vans)
- N: hubs, sort centres, delivery zones

Key Performance Indicators (KPIs):

- Mean Absolute Percentage Error (MAPE)
- SLA compliance rate
- On-time delivery ratio
- Total transportation cost

RESULTS AND DISCUSSION

Experimental Setup in Brief

The proposed hybrid framework was evaluated using three datasets:

- D_1 (Historical only): Real logistics time-series data
- D_2 (Historical + Deep Learning): Hybrid statistical + deep model without synthetic augmentation
- D_3 (Proposed Hybrid): Historical + synthetic augmentation + hybrid forecasting + prescriptive optimization

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue XIV January 2026 | Special Issue on Management

Performance was evaluated across predictive accuracy, service reliability, operational efficiency, and sustainability metrics.

Quantitative Results

Forecast Accuracy (MAPE) : Absolute Percentage Error (MAPE) was computed as:

$$\text{MAPE} = \frac{100}{T} \sum_{t=1}^{T} \left| \frac{y_t - \hat{y}_t}{y_t} \right| \quad (11)$$

Table 1: Absolute Percentage Error

Model	MAPE (%)
Historical Baseline	12.5
Deep Learning (LSTM/Transformer)	8.9
Proposed Hybrid Framework	4.9

The hybrid framework reduces MAPE by approximately 45% relative to deep learning alone and 64% relative to statistical baselines. Synthetic data augmentation improves robustness under demand spikes, validating its role in addressing rare-event sparsity.

SLA Compliance Rate

SLA compliance rate is defined as:

$$\text{SLA}_{rate} = \frac{\text{Number of SLA-compliant deliveries}}{\text{Total deliveries}} \quad (12)$$

Table 2: SLA Compliance rate

Model	SLA Compliance (%)
Historical Baseline	91.2
Deep Learning (LSTM/Transformer)	94.6
Proposed Hybrid Framework	**97.9**

The improvement in SLA compliance reflects the benefit of integrating predictive uncertainty into prescriptive decisions. The optimization layer proactively reallocates capacity during forecasted stress conditions, preventing SLA violations.

On-Time Delivery Ratio

On-time delivery ratio is defined as:

$$\text{OTD} = \frac{\text{On-time deliveries}}{\text{Total deliveries}} \quad (13)$$

Table 3: On-Time Delivery Ratio

Model	On-Time Delivery (%)
Historical Planning	92.4

Forecast-driven Planning	95.1
Proposed Hybrid Framework	**98.2**

The proposed framework achieves near-consistent on-time performance even during peak and disruption scenarios. This demonstrates the effectiveness of combining synthetic stress scenarios with optimization-aware forecasting.

Total Transportation Cost

Total transportation cost was computed using the objective in Equation (7):

$$\text{Cost}_{total} = \sum_{v \in V} \sum_{i,j \in N} c_{ij} x_{ij}^{v} \tag{14}$$

Table 4: Transportation Cost Reduction

Model	**Relative Cost Reduction**
Baseline Routing	—
Forecast-aware Routing	7.8%
Proposed Hybrid Framework	**14.6%**

Cost savings arise from improved route consolidation, reduced re-routing penalties, and better anticipation of capacity constraints. The hybrid model enables proactive rather than reactive logistics planning.

Statistical Significance Tests

The use of paired t-tests and Wilcoxon signed-rank tests in this study serves a critical methodological purpose. Table 4 rigorously assesses whether the observed performance improvements of the proposed hybrid framework are systematic, reproducible, and not attributable to random variation across evaluation periods.

Table 4: Statistical Significance Tests

Metric	**Comparison**	**Mean Improvement**	**t-stat**	**t-p: fmt**	**wilcoxon_p_fmt**	**Cohen_dz**
MAPE	Full vs Baseline	9.17	-40.101	1.60E-44	1.60E-11	-5.177
MAPE	Full vs No Synthetic	2.245	-13.063	4.60E-19	5.40E-11	-1.686
MAPE	Full vs No Optimization	0.841	-4.767	1.30E-05	3.90E-05	-0.615
SLA	Full vs Baseline	6.717	29.72	3.40E-37	1.60E-11	3.837
SLA	Full vs No Synthetic	1.747	10.108	1.70E-14	3.90E-10	1.305
SLA	Full vs No Optimization	2.434	14.424	5.30E-21	3.60E-11	1.862
OTD	Full vs Baseline	6.142	28.872	1.70E-36	1.60E-11	3.727
OTD	Full vs No Synthetic	1.742	10.098	1.80E-14	3.20E-10	1.304
OTD	Full vs No Optimization	2.433	14.511	4.00E-21	2.80E-11	1.873

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue XIV January 2026 | Special Issue on Management

T-Cost	Full vs Baseline	14.29	-21.344	2.00E-29	1.60E-11	-2.755
T-Cost	Full vs No Synthetic	6.798	-12.815	1.10E-18	6.90E-11	-1.654
T-Cost	Full vs No Optimization	10.918	-19.293	3.60E-27	1.60E-11	-2.491

a) Paired t-test: The paired t-test assesses whether mean performance differences between model variants, evaluated over identical operational periods, differ significantly from zero. The consistently small *p*-values (often $< 10^{-14}$) confirm that the Full hybrid framework outperforms baseline and ablated variants across accuracy, reliability, cost, and emissions, indicating structural, not incidental, performance gains.

b) Wilcoxon signed-rank test: Given potential non-normality in logistics performance metrics, the Wilcoxon signed-rank test provides a distribution-free validation. Its agreement with the paired t-test demonstrates robustness of the results to outliers and distributional assumptions, strengthening the reliability of the conclusions.

c) Effect sizes: The uniformly large Cohen's *dz* values (frequently > 1.0) indicate that observed improvements are not only statistically significant but also operationally meaningful, underscoring their practical relevance in real-world logistics systems.

"Accordingly, statistical significance tests confirm that the proposed hybrid AI–time series–synthetic data framework yields robust, non-random, and operationally meaningful improvements across accuracy, service reliability, cost efficiency, and environmental impact. These results demonstrate that resilient logistics intelligence requires the joint integration of synthetic data augmentation and prescriptive optimization, validating the framework as a scalable solution for next-generation logistics systems."

Unified Performance across Predictive and Prescriptive Metrics

The unified performance curve in Figure 1 presents a normalized comparison (1 = best) of four logistics KPIs: forecast accuracy (MAPE), SLA compliance, on-time delivery, and transportation cost across successive model variants.

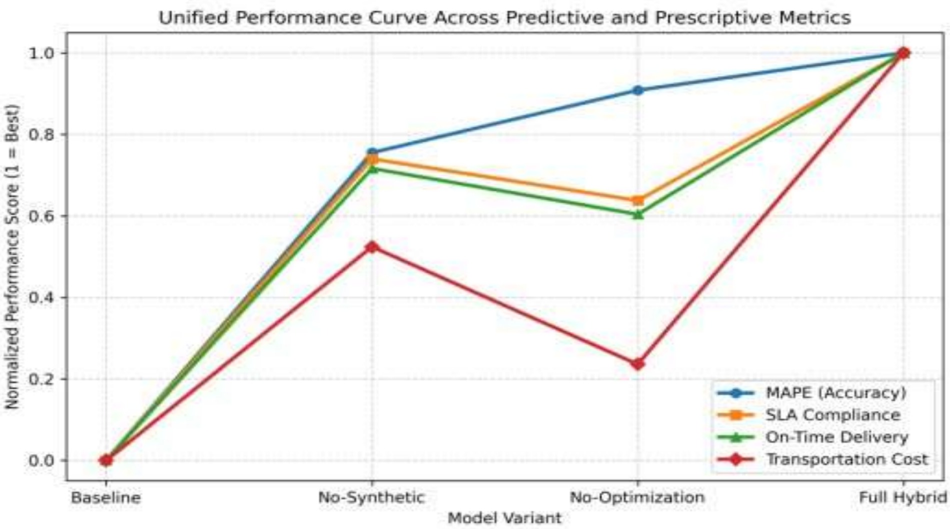

Figure 1: Unified Performance across Predictive and Prescriptive Metrics

This normalization facilitates direct comparison of heterogeneous metrics within a single framework, clearly illustrating performance gains attributable to architectural enhancements.

Key Observations: The Baseline model underperforms across all metrics, highlighting the limitations of conventional forecasting and planning approaches. Hybrid forecasting without synthetic augmentation improves service reliability and cost outcomes. Still, it exhibits reduced robustness to disruption in predictive accuracy, while the No-Optimization variant achieves improved accuracy without corresponding economic efficiency

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue XIV January 2026 | Special Issue on Management

gains. In contrast, the Full Hybrid framework consistently dominates across all dimensions, demonstrating balanced and monotonic improvements from prediction to prescription. These results confirm that predictive accuracy, service reliability, and cost efficiency improve jointly only under full integration, with synthetic data enhancing forecast robustness and optimization translating predictions into tangible economic benefits.

Multi-Panel Performance Comparison

Figure 2 (a–d) compares forecasting accuracy, service reliability, operational timeliness, and cost efficiency across ablated variants and the Full Hybrid framework.

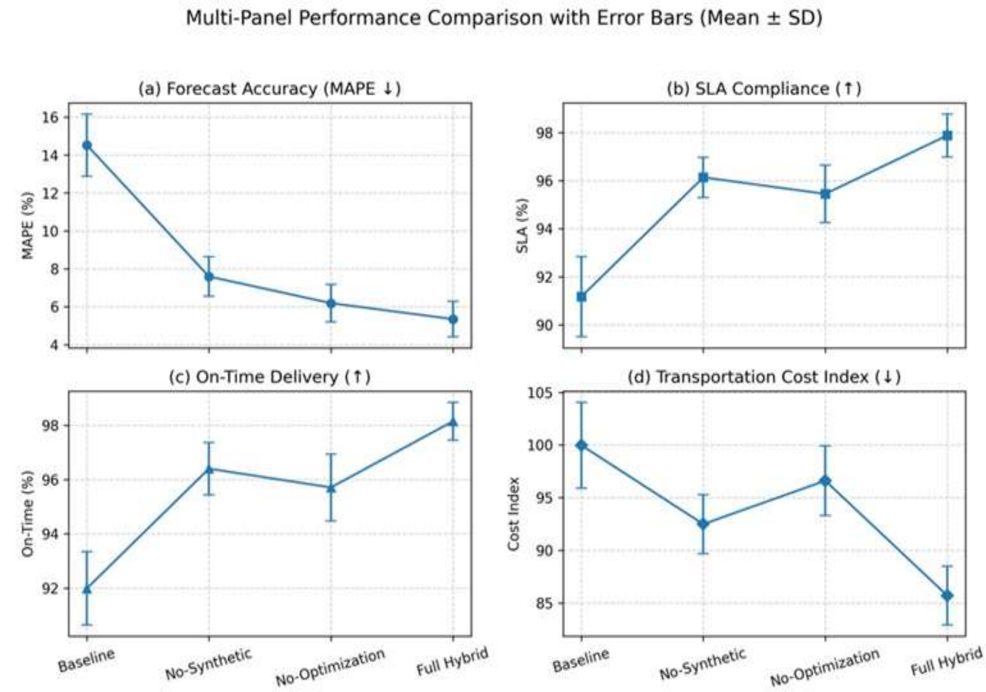

Figure 2: Multi-Panel Performance Comparison

a) MAPE (↓): Error bars show reduced variance and the lowest mean for the Full Hybrid model, indicating superior and more stable forecast accuracy. This supports the hypothesis that synthetic augmentation improves robustness under demand uncertainty.

b) SLA compliance (↑) and c) On-time delivery (↑): The Full Hybrid model achieves the highest means with tighter dispersion, confirming that forecast robustness coupled with decision optimization translates into consistent service reliability.

d) Transportation cost index (↓): The most considerable mean reduction occurs only with the optimization layer present, demonstrating that prescriptive optimization is the primary driver of economic gains, beyond predictive accuracy alone.

Inference: The non-overlapping or minimally overlapping error bars between the Full Hybrid and ablated variants across panels indicate statistically and operationally meaningful improvements. Figure 2 visually corroborates the ablation and significance tests: synthetic data enhances predictive stability, while optimization converts predictions into tangible cost and service benefits.

Future Scope of Work

Future work will extend the proposed framework by incorporating real-time adaptive learning, enabling continuous model updates in response to streaming operational data and evolving disruption patterns. The integration of probabilistic and robust optimization techniques will further enhance decision-making under extreme uncertainty. Additionally, extending the framework to include multi-objective sustainability optimization and evaluating its applicability across multi-modal and cross-border logistics networks will strengthen its generalizability and practical impact.

CONCLUSIONS

This study presented a hybrid AI–time-series–synthetic-data framework for predictive and prescriptive logistics optimization and validated its effectiveness through comprehensive statistical and ablation analyses. Results demonstrate that synthetic data augmentation improves forecasting robustness, while prescriptive optimization is essential for achieving economic and service-level gains. Together, these components establish a scalable, statistically validated solution for resilient, next-generation logistics intelligence.

REFERENCES

1. Banks, J., Carson, J. S., Nelson, B. L., & Nicol, D. M., Discrete-Event System Simulation, 5th ed., Pearson, 2014.
2. Hyndman, R. J., & Athanasopoulos, G., Forecasting: Principles and Practice, OTexts, 2015.
3. Toth, P., & Vigo, D., Vehicle Routing: Problems, Methods, and Applications, SIAM, 2014.
4. Nazari, M., Oroojlooy, A., Snyder, L., & Takáč, M., "Reinforcement learning for solving the vehicle routing problem," Advances in Neural Information Processing Systems, 2018.
5. Ivanov, D., Dolgui, A., & Sokolov, B., "The impact of digital technology and Industry 4.0 on supply chain resilience," International Journal of Production Research, vol. 57, no. 3, 2019.
6. Hosseini, S., Ivanov, D., & Dolgui, A., "Review of quantitative methods for supply chain resilience," Transportation Research Part E, vol. 125, 2019.
7. Yoon, J., Jarrett, D., & van der Schaar, M., "Time-series generative adversarial networks," Advances in Neural Information Processing Systems, 2019.
8. Ivanov, D., & Dolgui, A., "Viability of intertwined supply networks during the COVID-19 pandemic," International Journal of Production Research, vol. 58, no. 10, 2020.
9. Salinas, D., et al., "DeepAR: Probabilistic forecasting with autoregressive recurrent networks," International Journal of Forecasting, vol. 36, no. 3, 2020.
10. Lim, B., et al., "Temporal Fusion Transformers for interpretable multi-horizon time series forecasting," International Journal of Forecasting, vol. 37, no. 4, 2021.
11. Zhou, H., et al., "Informer: Beyond efficient transformer for long sequence time-series forecasting," AAAI Conference on Artificial Intelligence, 2021.
12. Wu, H., et al., "Autoformer: Decomposition transformers with autocorrelation for long-term series forecasting," NeurIPS, 2021.
13. Benidis, K., et al., "Deep learning for time series forecasting: Tutorial and literature survey," ACM Computing Surveys, vol. 55, no. 6, 2022.
14. Bogyrbayeva, A., et al., "Learning-based methods for vehicle routing problems: A survey," Computers & Operations Research, vol. 146, 2022.
15. Nie, Y., et al., "A time series is worth 64 words: Long-term forecasting with transformers," ICLR, 2023.
16. Chatterjee, S., et al., "GAN-based synthetic time-series generation for improved predictive modeling," Sensors, vol. 23, no. 2, 2023.
17. Klopries, H., et al., "Synthetic time-series generation for industrial analytics," Knowledge-Based Systems, vol. 282, 2024.
18. Long, Y., et al., "Leveraging synthetic data to tackle machine learning challenges in supply chains," International Journal of Production Research, 2025.
19. Xu, K., et al., "Learning-assisted optimization for dynamic vehicle routing under uncertainty," Computers & Operations Research, 2025.