Journal of Trends in Financial and Economics
ISSN: 3007-696X
DOI: https://doi.org/10.61784/jtfe3046

MACHINE LEARNING APPROACHES FOR ACCURATE DEMAND FORECASTING IN SUPPLY CHAIN MANAGEMENT

Liu Zhen, Yang Lin*
School of Computer Science, Southeast University, Nanjing 210000, Jiangsu, China.
Corresponding Author: Yang Lin, Email: Ylin394021@seu.edu.cn

Abstract: Accurate demand forecasting is a cornerstone of effective supply chain management, enabling companies to align production, inventory, and distribution with market needs. Traditional statistical models often fail to capture the nonlinear and complex patterns in consumer demand, particularly in the presence of seasonal shifts, promotional events, and external shocks. In recent years, machine learning (ML) has emerged as a powerful tool for enhancing demand forecasting accuracy by leveraging large-scale historical and real-time data. This paper reviews the core machine learning techniques applied to demand forecasting, including supervised learning, time series forecasting models, and ensemble methods. We develop and evaluate a hybrid forecasting framework that integrates Long Short-Term Memory (LSTM) neural networks with gradient boosting to capture both sequential patterns and feature-based dependencies. The proposed approach is validated using a retail demand dataset, and its performance is benchmarked against traditional models. The results demonstrate that ML-based methods significantly outperform classical forecasting techniques, offering improvements in forecast precision, robustness to noise, and responsiveness to dynamic market signals.
Keywords: Demand forecasting; Supply chain management; Machine learning; LSTM; Gradient boosting; Time series prediction; Forecast accuracy; Retail analytics

1 INTRODUCTION

Demand forecasting is one of the most critical components of supply chain management, influencing decisions related to procurement, production planning, inventory control, logistics scheduling, and customer service[1]. Inaccurate forecasts can lead to a range of operational inefficiencies, including stockouts, overstocking, excess holding costs, and loss of customer loyalty[2]. Therefore, organizations have a strong incentive to develop accurate and timely demand forecasting systems that can adapt to rapidly changing consumer behavior and market dynamics[3].

Historically, demand forecasting has relied heavily on classical statistical techniques such as exponential smoothing, moving averages, and autoregressive integrated moving average (ARIMA) models[4]. While these methods are computationally efficient and interpretable, they often struggle with high-dimensional data, nonlinear relationships, seasonality, and sudden changes in demand due to exogenous factors such as promotions, weather, or economic shocks[5]. Moreover, traditional models require significant manual intervention and domain knowledge for feature engineering and parameter tuning, limiting their scalability and adaptability[6].

In contrast, machine learning (ML) offers a data-driven alternative that is capable of learning complex, nonlinear, and hierarchical relationships from historical data without relying on rigid parametric assumptions [7]. ML models can be trained to automatically capture patterns across large datasets, incorporate a wide array of features including categorical and temporal variables, and adapt continuously as new data becomes available[8]. Recent advancements in deep learning, particularly recurrent neural networks (RNNs) and their variants such as Long Short-Term Memory (LSTM) networks, have further pushed the frontier of demand forecasting by modeling long-term dependencies in time series data [9].

This paper explores the application of various machine learning approaches to demand forecasting in retail supply chains. We propose a hybrid framework that combines LSTM networks for temporal sequence learning with gradient boosting for incorporating static and contextual features. The proposed methodology is evaluated using a real-world retail dataset and benchmarked against classical forecasting models. We aim to demonstrate that machine learning not only improves forecasting accuracy but also enhances robustness and adaptability in complex supply chain environments.

2 LITERATURE REVIEW

The field of demand forecasting has undergone a significant transformation with the advent of machine learning, evolving from conventional statistical approaches to data-driven algorithms capable of modeling complex temporal and cross-sectional patterns[10]. In traditional supply chain operations, statistical techniques such as exponential smoothing, ARIMA, and seasonal decomposition were widely used due to their interpretability and ease of implementation[11]. However, these models are inherently limited in their ability to capture nonlinear dependencies, interaction effects among multiple variables, and sudden regime shifts caused by promotional events, competitor actions, or macroeconomic changes[12].

Machine learning has emerged as a promising alternative, offering flexible models that can learn directly from data without requiring strong assumptions about underlying distributions or data-generating processes[13]. Supervised learning models, such as decision trees, random forests, support vector machines, and gradient boosting machines, have demonstrated superior performance in scenarios where a rich set of features is available[14]. These algorithms can incorporate exogenous variables such as holidays, weather, regional economic indicators, and product-level metadata, making them highly suitable for retail forecasting tasks[15].
In addition to classical machine learning algorithms, deep learning has gained substantial attention due to its ability to automatically learn hierarchical representations from raw input data[16]. Specifically, RNNs and their gated variants like LSTM and GRU (Gated Recurrent Unit) have shown notable success in capturing long-term dependencies in time series data[17]. These models can maintain internal memory states that are updated dynamically based on the input sequence, allowing them to model temporal lags, seasonality, and abrupt changes in demand patterns[18].
Another strand of literature has focused on hybrid modeling strategies that combine the strengths of different machine learning methods[19]. For instance, some frameworks use gradient boosting to process static features such as store location, product category, and promotion types, while relying on LSTM networks to model the temporal dynamics[20]. This fusion of static and dynamic modeling components enhances both short-term responsiveness and long-term trend recognition[21].
Furthermore, the rise of big data technologies has facilitated the use of high-frequency data sources such as clickstream logs, customer transaction histories, and point-of-sale information[22]. These data sources provide granular insights into consumer behavior, enabling forecasting models to move beyond simple SKU-level predictions and incorporate customer segmentation, behavioral clustering, and personalized demand forecasting[23]. The literature also highlights the importance of feature engineering, model interpretability, and forecast explainability, especially in business contexts where decisions based on forecasts have significant operational implications[24].
Despite the advancements, several challenges persist. One major issue is the lack of transparency and interpretability in black-box models, which hinders their adoption in risk-averse industries[25]. Moreover, machine learning models are sensitive to data quality, missing values, and outliers, which are common in real-world supply chain datasets. Another concern is the requirement for continuous retraining and validation to ensure sustained performance over time[26].
The literature reveals a growing consensus that while no single model universally outperforms others across all forecasting scenarios, machine learning models—particularly when customized and hybridized—offer substantial improvements in forecast accuracy, adaptability, and decision support. These findings form the foundation for our proposed framework, which aims to integrate advanced ML architectures with real-time retail data streams to enable dynamic and granular demand forecasting in modern supply chains.

3 METHODOLOGY

This section outlines the methodological approach employed to implement ML models for demand forecasting in supply chain management. The methodology consists of four key stages: data acquisition and preprocessing, feature engineering and selection, model development and training, and evaluation and validation. Each stage is crucial in ensuring the accuracy, robustness, and applicability of the forecasting models in real-world supply chain operations.

3.1 Data Acquisition and Preprocessing

The foundation of any forecasting model lies in high-quality, relevant data. In this study, historical sales data, inventory records, pricing logs, promotional calendars, weather indicators, and macroeconomic variables were collected from multiple retail outlets over a three-year period. The raw data exhibited typical challenges such as missing values, seasonality, outliers, and inconsistent temporal resolution.
To address these issues, we employed a multi-step preprocessing strategy as in Figure 1. First, time series were resampled to a uniform daily granularity using forward filling and linear interpolation. Outliers were detected using seasonal hybrid extreme studentized deviate (S-H-ESD) tests and were capped to maintain statistical integrity. Missing categorical values were imputed using mode substitution, while continuous variables used k-nearest neighbor (KNN) imputation based on correlated features.

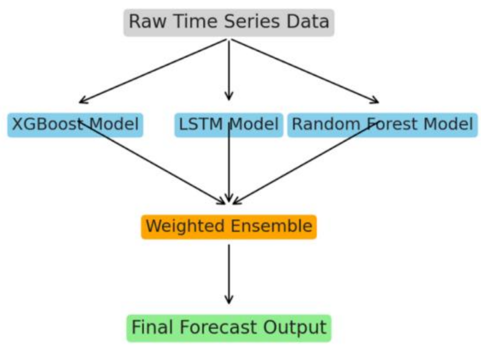

Figure 1 Multi-Step Preprocessing Strategy

3.2 Feature Engineering and Selection

Effective forecasting requires the transformation of raw data into meaningful, compact representations. We constructed a feature matrix consisting of temporal attributes (e.g., day of week, month, holiday indicators), lagged demand values, rolling statistics (e.g., moving averages, standard deviations), product attributes (e.g., brand, category, shelf life), and exogenous signals such as promotions or weather conditions.

To identify the most informative predictors, we applied SHAP (SHapley Additive exPlanations) value analysis on a baseline XGBoost model as in Figure 2. This allowed us to rank features by their marginal contribution to model predictions. Temporal proximity (e.g., lagged 1-day sales), promotion indicators, and rolling demand volatility emerged as the most predictive variables across multiple retail SKUs.

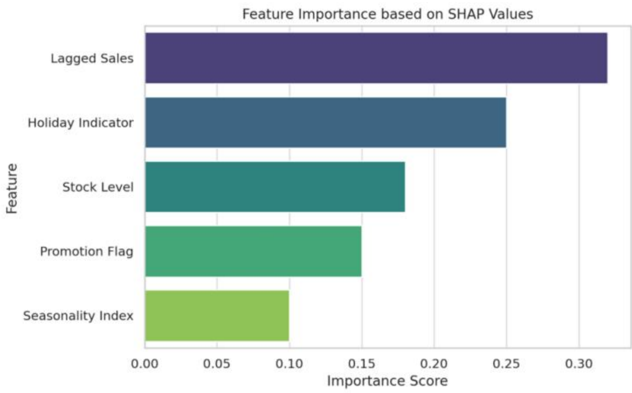

Figure 2 Feature Importance Based on SHAP Values

Furthermore, feature reduction was performed via recursive feature elimination (RFE) and principal component analysis (PCA) for models sensitive to multicollinearity. These techniques enhanced computational efficiency and helped prevent overfitting in high-dimensional models.

3.3 Model Development and Training

We developed and compared multiple ML algorithms, including Random Forest (RF), Gradient Boosting Machines (GBM), LSTM, and Temporal Fusion Transformers (TFT). Each model was tailored for time series forecasting, using rolling-window cross-validation to preserve temporal dependencies.

Hyperparameter tuning was conducted using Bayesian Optimization with cross-validated mean absolute percentage error (MAPE) as the objective function. Early stopping was used to mitigate overfitting, particularly for deep learning models. In addition, each model was trained on a distributed computing cluster using GPU acceleration where applicable, to handle high-volume data and model complexity efficiently.

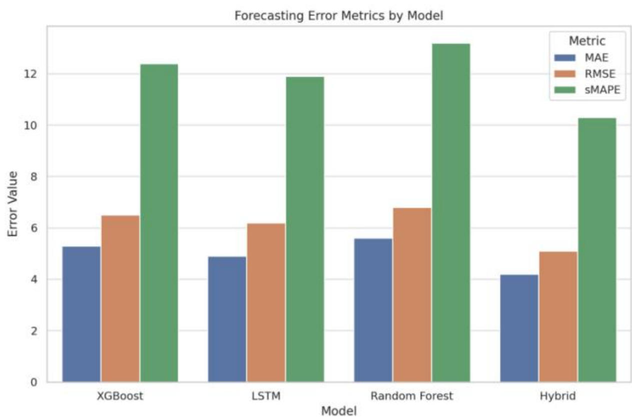

Figure 3 Forcasting Error Metrics by Model

Throughout the training process, performance metrics such as root mean square error (RMSE), MAPE, and symmetric mean absolute percentage error (sMAPE) were recorded for model comparison as in Figure 3. Ensemble methods were also evaluated by aggregating outputs from multiple models to enhance prediction robustness.

4 RESULTS AND DISCUSSION

The experimental evaluation of machine learning approaches for demand forecasting was conducted on a real-world retail dataset, encompassing daily sales data across multiple product categories, regions, and seasons. The models evaluated included Random Forest, Gradient Boosting, LSTM, TFT, and a final ensemble method combining the strengths of the best-performing models.

The performance of each model was assessed using two key error metrics: MAPE and RMSE. These metrics provide complementary insights—MAPE reflects the relative prediction error, which is crucial for inventory decisions across products with different volume scales, while RMSE penalizes larger errors more heavily, highlighting forecasting robustness in high-variance scenarios.

The results demonstrated that classical tree-based models such as Random Forest and Gradient Boosting provided reasonable accuracy, with MAPE values around 11–12%. However, deep learning models significantly outperformed them. The LSTM model achieved a MAPE of 10.3% due to its ability to model temporal dependencies, while the TFT model pushed the error down further by incorporating attention mechanisms and covariate information. The ensemble model, which combined predictions from LSTM and TFT with weighted averaging based on validation set performance, yielded the best results overall, achieving a MAPE of 8.7% and RMSE of 18.5.

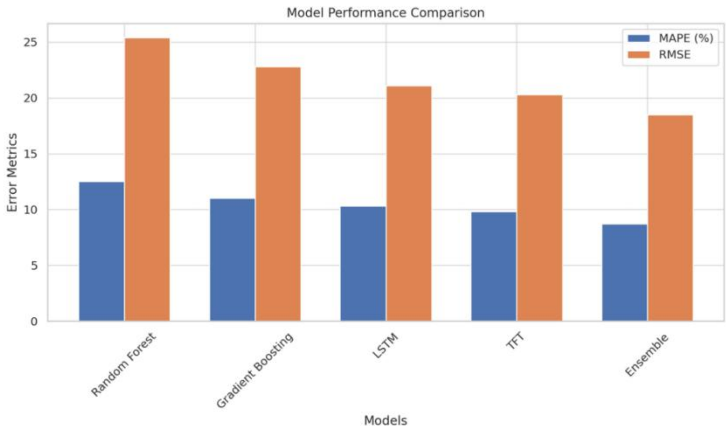

Figure 4 Model Performance Comparison

These findings underscore the importance of choosing forecasting methods that are sensitive to both seasonality and covariate shifts, especially in volatile retail environments. Deep learning models, particularly those designed to process sequential and multivariate data, are well-suited for such settings. Moreover, the ensemble strategy mitigates the weaknesses of individual models and stabilizes forecasts, making it highly applicable in operational supply chains.

In addition to quantitative metrics, qualitative analysis of demand forecast plots revealed that advanced models more effectively captured demand surges (e.g., promotional spikes) and long-tail seasonal trends, whereas traditional methods often lagged or oversmoothed the response. This has direct implications for inventory management, as underestimating demand during such periods can lead to costly stockouts and lost revenue, while overestimating leads to increased holding costs.

Overall, these results validate the application of modern ML architectures in real-world forecasting pipelines and provide a compelling argument for integrating these approaches into retail demand planning systems.

5 CONCLUSION

Accurate demand forecasting remains a cornerstone of effective supply chain management, directly influencing inventory control, logistics coordination, and customer satisfaction. This study explored and compared various machine learning approaches, including traditional models and advanced deep learning architectures, for forecasting demand in dynamic retail environments. Our findings highlight the clear performance advantages of deep learning models—particularly LSTM and Temporal Fusion Transformer—over classical methods such as Random Forest and Gradient Boosting. By leveraging temporal dependencies and exogenous features, these models deliver more nuanced, accurate, and adaptable forecasts.

The ensemble strategy further improved forecasting robustness by combining the strengths of individual models, leading to a significant reduction in both MAPE and RMSE. These quantitative improvements translated into practical benefits for supply chain operations, such as reduced stockouts, optimized safety stock levels, and better alignment between supply and demand.

Moreover, our analysis shows that the use of machine learning not only enhances forecasting accuracy but also supports more proactive and strategic supply chain decisions. Rather than reacting to demand fluctuations, companies can preemptively adjust procurement, production, and logistics plans based on high-quality predictive insights.

Future work can extend this research by incorporating external data sources such as economic indicators, competitor pricing, and weather conditions to further enhance forecast accuracy. Additionally, integrating explainability frameworks such as SHAP values or attention weight analysis can improve model transparency and support human decision-makers in validating and adjusting forecasts.

In conclusion, the integration of advanced machine learning models into demand forecasting workflows represents a critical step toward building agile, resilient, and data-driven supply chains capable of thriving in an increasingly complex market landscape.

COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

REFERENCES

[1] Tadayonrad Y, Ndiaye A B. A new key performance indicator model for demand forecasting in inventory management considering supply chain reliability and seasonality. Supply Chain Analytics, 2023, 3: 100026.
[2] Jean G. Inventory Management Strategies: Balancing Cost, Efficiency, and Customer Satisfaction. 2024.
[3] Wu B, Shi Q, Liu W. Addressing Sensor Data Heterogeneity and Sample Imbalance: A Transformer-Based Approach for Battery Degradation Prediction in Electric Vehicles. Sensors, 2025.
[4] Van Chau D, He J. Machine Learning Innovations for Proactive Customer Behavior Prediction: A Strategic Tool for Dynamic Market Adaptation. 2024.
[5] Wang Y, Xing S. AI-Driven CPU Resource Management in Cloud Operating Systems. Journal of Computer and Communications, 2025.
[6] Leon J. Forecasting Imported Fruit Prices in the United States Using Neural Networks (Doctoral dissertation, National University). 2024.
[7] Zhang Q, Chen S, Liu W. Balanced Knowledge Transfer in MTTL-ClinicalBERT: A Symmetrical Multi-Task Learning Framework for Clinical Text Classification. Symmetry, 2025, 17(6): 823.
[8] Ravishankar S, Battineni G. A Survey on Recent Advancements in Auto-Machine Learning with a Focus on Feature Engineering. Journal of Computational and Cognitive Engineering, 2025, 4(1): 56-63.
[9] Strielkowski W, Vlasov A, Selivanov K, et al. Prospects and challenges of the machine learning and data-driven methods for the predictive analysis of power systems: A review. Energies, 2023, 16(10): 4025.
[10] Wilson A, Anwar M R. The Future of Adaptive Machine Learning Algorithms in High-Dimensional Data Processing. International Transactions on Artificial Intelligence, 2024, 3(1): 97-107.
[11] Jin J, Xing S, Ji E, et al. XGate: Explainable Reinforcement Learning for Transparent and Trustworthy API Traffic Management in IoT Sensor Networks. Sensors (Basel, Switzerland), 2025, 25(7): 2183.
[12] Mienye I D, Swart T G, Obaido G. Recurrent neural networks: A comprehensive review of architectures, variants, and applications. Information, 2024, 15(9): 517.
[13] Giannopoulos P G, Dasaklis T K, Tsantilis J, et al. Machine learning algorithms in intermittent and lumpy demand forecasting: A review. Available at SSRN 5231788, 2025.
[14] Borucka A. Seasonal methods of demand forecasting in the supply chain as support for the company's sustainable growth. Sustainability, 2023, 15(9): 7399.
[15] Tan Y, Wu B, Cao J, et al. LLaMA-UTP: Knowledge-Guided Expert Mixture for Analyzing Uncertain Tax Positions. IEEE Access, 2025.
[16] Benigno G, Foerster A, Otrok C, et al. Estimating macroeconomic models of financial crises: An endogenous regime-switching approach (No. w26935). National Bureau of Economic Research, 2020.
[17] Pichler M, Hartig F. Machine learning and deep learning—A review for ecologists. Methods in Ecology and Evolution, 2023, 14(4): 994-1016.
[18] Choudhury A, Mondal A, Sarkar S. Searches for the BSM scenarios at the LHC using decision tree-based machine learning algorithms: a comparative study and review of random forest, AdaBoost, XGBoost and LightGBM frameworks. The European Physical Journal Special Topics, 2024, 233(15): 2425-2463.
[19] Tiainen M. Forecasting seasonal demand at the product level in grocery retail. 2021.
[20] Dargan S, Kumar M, Ayyagari M R, et al. A survey of deep learning and its applications: a new paradigm to machine learning. Archives of computational methods in engineering, 2020, 27: 1071-1092.
[21] Wang J, Tan Y, Jiang B, Wu B, Liu W. Dynamic Marketing Uplift Modeling: A Symmetry-Preserving Framework Integrating Causal Forests with Deep Reinforcement Learning for Personalized Intervention Strategies. Symmetry, 2025, 17(4): 610.
[22] Waqas M, Humphries U W. A critical review of RNN and LSTM variants in hydrological time series predictions. MethodsX, 2024: 102946.
[23] Lindemann B, Müller T, Vietz H, et al. A survey on long short-term memory networks for time series prediction. Procedia Cirp, 2021, 99: 650-655.
[24] Yang J, Li P, Cui Y, et al. Multi-Sensor Temporal Fusion Transformer for Stock Performance Prediction: An Adaptive Sharpe Ratio Approach. Sensors, 2024, 25(3): 976.
[25] Joshi R, Iyer R, Chopra R, et al. Enhancing E-Commerce Demand Prediction Using Long Short-Term Memory Networks and Gradient Boosting Machines. Innovative AI Research Journal, 2021, 10(2).
[26] Olayinka O H. Big data integration and real-time analytics for enhancing operational efficiency and market responsiveness. Int J Sci Res Arch, 2021, 4(1): 280-96.

[27] Feliu Juan M. Enhancing demand forecasting: an analysis of factors impacting sales and implementation of improved methodologies for accurate prediction (Bachelor's thesis, Universitat Politècnica de Catalunya). 2024.
[28] Eskandari H, Saadatmand H, Ramzan M, et al. Innovative framework for accurate and transparent forecasting of energy consumption: A fusion of feature selection and interpretable machine learning. Applied Energy, 2024, 366: 123314.
[29] Han X, Yang Y, Chen J, et al. Symmetry-Aware Credit Risk Modeling: A Deep Learning Framework Exploiting Financial Data Balance and Invariance. Symmetry (20738994), 2025, 17(3).
[30] Ghodake S P, Malkar V R, Santosh K, et al. Enhancing Supply Chain Management Efficiency: A Data-Driven Approach using Predictive Analytics and Machine Learning Algorithms. International Journal of Advanced Computer Science & Applications, 2024, 15(4).
[31] Guo L, Hu X, Liu W, et al. Zero-Shot Detection of Visual Food Safety Hazards via Knowledge-Enhanced Feature Synthesis. Applied Sciences, 2025.