International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2362-2372
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0318

Real-Time Intelligent Supply Chain Demand Forecasting and Dynamic Repricing System Using Machine Learning

Ann Mariya Ajay Lazar [1], Dr. Rahul B[2]
[1]PG Scholar, Dept. of CSE, Royal College of Engineering and Technology, Thrissur, Kerala, India
[2]Associate Professor, Dept. of CSE, Royal College of Engineering and Technology, Thrissur, Kerala, India
Email ID: *annmariya3al@gmail.com[1], rahul.rahul.b9@gmail.com[2]*

Abstract

In modern retail environments, effective supply chain management requires accurate demand forecasting and adaptive pricing strategies to handle demand uncertainty and inventory fluctuations. This paper presents a real-time intelligent supply chain demand forecasting and dynamic repricing system based on machine learning techniques. The proposed system leverages the Walmart Sales Forecasting dataset, integrating historical sales data, weather-related attributes, and holiday indicators to predict product demand across multiple store locations and departments. Feature engineering is performed using time-series lag variables and rolling statistics to capture seasonality and temporal dependencies. A Light Gradient Boosting Machine (LightGBM) regression model is employed for demand forecasting due to its efficiency and suitability for structured retail data. The forecasted demand is further utilized in a dynamic repricing module that adjusts prices based on demand–inventory imbalance, demand elasticity, and simulated competitor pricing. To emulate real-world deployment scenarios, a real-time sales event simulation module is implemented, representing a Kafka-compatible streaming architecture. The trained model is exposed through a RESTful API using FastAPI, enabling on-demand predictions. An interactive dashboard provides visual insights into demand trends, pricing recommendations, and inventory risk alerts to support data-driven decision-making. Experimental results demonstrate that the proposed system effectively captures demand patterns and offers practical pricing recommendations, making it a scalable and deployable solution for intelligent retail supply chain optimization.

Keywords: *Demand Forecasting, Dynamic Repricing, Machine Learning, LightGBM, Time-Series Analysis, Decision Support Systems*

1. Introduction

1.1. Background and Motivation

Modern retail supply chains operate in highly dynamic and uncertain environments, where customer demand is influenced by multiple factors such as seasonality, holidays, weather conditions, and economic indicators. Inaccurate demand estimation often leads to overstocking or stockouts, increasing operational costs, revenue loss, and reducing customer satisfaction. Traditional demand forecasting approaches based on static rules or statistical averages are limited in their ability to capture complex and non-linear demand patterns present in large-scale retail data. With the increasing availability of historical sales data and external contextual information, machine learning techniques have emerged as effective tools for demand forecasting in retail supply chains. Recent studies have highlighted the increasing adoption of machine learning and deep learning techniques for demand forecasting in supply chain management due to their ability to model complex temporal patterns and nonlinear relationships (Scientific Reports, 2024). Data-driven models are capable of learning temporal dependencies and interactions among multiple influencing factors, enabling more accurate and adaptive demand predictions across different store locations and product categories. Accurate demand forecasting further enables intelligent pricing decisions, allowing retailers to dynamically adjust prices based on demand variations and inventory

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2362-2372
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0318

levels. Predictive analytics powered by artificial intelligence has become a key enabler for improving forecasting accuracy and operational efficiency in modern supply chains (Radhakrishnan & Shankar, 2024). In addition to forecasting accuracy, real-time decision-making has become a critical requirement in modern retail systems. The ability to continuously ingest sales data, generate updated predictions, and visualize insights through interactive dashboards supports proactive supply chain management. Motivated by these challenges, this work focuses on developing an integrated and deployable system that combines machine learning–based demand forecasting, dynamic repricing strategies, real-time data simulation, and decision-support visualization to enhance retail supply chain optimization [1].

1.2.Problem Statement

Retail organizations face significant challenges in accurately forecasting demand and making timely pricing decisions amid demand variability and uncertain market conditions. Demand patterns vary across different stores and product categories and are influenced by temporal factors such as seasonality, holidays, and external environmental conditions. In many practical scenarios, retailers rely on static forecasting models and fixed pricing strategies that fail to adapt to changing demand and inventory situations, leading to inefficient supply chain operations. Another major challenge is the lack of integrated systems that combine demand forecasting, pricing optimization, and real-time decision support within a unified framework. Existing approaches often address these components independently, limiting their effectiveness in practical deployment. Furthermore, real-time data ingestion and scalable deployment mechanisms are frequently absent in academic implementations, reducing their applicability to real-world retail environments. Therefore, there is a need for an intelligent, data-driven supply chain system that can accurately forecast demand across multiple store locations and product categories, dynamically adjust pricing based on demand and inventory conditions, and provide real-time visibility through interactive visualization tools. Addressing these challenges forms the core problem addressed in this research.

1.3.Contributions of the Proposed System

Recent research emphasizes the need for end-to-end intelligent supply chain systems that integrate data engineering, forecasting models, and decision-support mechanisms (Rabbani et al., 2025). The primary contributions of this research are summarized as follows End-to-End Intelligent Supply Chain Framework: This work proposes an integrated end-to-end system that combines demand forecasting, dynamic repricing, real-time data simulation, API deployment, and visualization within a unified retail supply chain analytics framework.

- Machine Learning–Based Demand Forecasting: A time-series–aware demand forecasting model based on Light Gradient Boosting Machine (LightGBM) is developed using real-world retail data. The model incorporates historical sales, weather-related features, holiday indicators, and engineered temporal features to predict demand across multiple store locations and product departments.
- Feature Engineering for Retail Time-Series Data: Effective feature engineering techniques, including lag-based features and rolling statistics, are employed to capture seasonality and temporal dependencies inherent in retail sales data, improving forecasting accuracy and robustness.
- Dynamic Repricing Strategy Based on Demand and Inventory: A practical and explainable dynamic pricing mechanism is introduced, where price recommendations are adjusted based on demand–inventory imbalance, demand elasticity rules, and simulated competitor pricing, enabling adaptive pricing decisions.
- Real-Time Data Ingestion Simulation and Deployment: A real-time sales event simulation module is implemented to emulate Kafka-style streaming behaviour, demonstrating the system's readiness for real-world deployment. The trained forecasting model is exposed via a RESTful API using

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2362-2372
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0318

FastAPI for on-demand prediction services.

- Interactive Decision-Support Dashboard: An interactive dashboard is developed to visualize demand trends, pricing recommendations, and inventory risk alerts, providing actionable insights for retail decision-makers.

1.4.SYSTEM DESIGN AND METHODOLOGY

To support accurate demand forecasting and adaptive pricing decisions, the proposed system is designed as a modular, end-to-end pipeline integrating data acquisition, preprocessing, machine learning–based prediction, and real-time deployment components. Figure 1 illustrates the overall system architecture, highlighting the sequential flow from raw data ingestion to decision-oriented outputs. The architecture consists of major modules such as dataset ingestion, preprocessing and feature engineering, demand forecasting, dynamic repricing, API-driven deployment, and visualization dashboard. Each module is designed to operate independently while ensuring seamless data flow across the pipeline. This structured design enables the system to handle large-scale retail datasets, incorporate temporal and external influencing factors, and support real-time prediction and pricing functionalities essential for practical retail supply-chain environments.

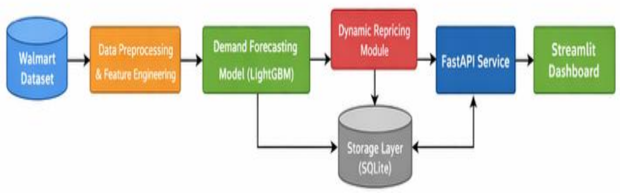

Figure 1. Overall System Architecture Diagram

1.5.Dataset Description

The proposed system is evaluated using the Walmart Recruiting – Store Sales Forecasting Dataset, a real-world retail dataset that captures historical weekly sales across multiple store locations and product departments. The experiments in this study are conducted using the publicly available Walmart Store Sales Forecasting dataset released through the Kaggle data science platform. The dataset is widely used for research in retail analytics and demand forecasting due to its large scale, temporal structure, and inclusion of external influencing factors. The dataset consists of four primary files as shown in Figure 2: test.csv, train.csv, stores.csv, and features.csv. The training data contains weekly sales values recorded for each store–department combination along with corresponding dates. Store-level information includes store identifiers and store type, while the features dataset provides additional contextual attributes such as temperature, fuel price, consumer price index (CPI), unemployment rate, and a binary indicator representing holiday weeks. Each record in the dataset is identified by a combination of store number, department number, and date, making it suitable for multi-store and multi-product demand forecasting. The inclusion of weather-related and macroeconomic variables enables the modelling of external factors that influence consumer purchasing behaviour. Holiday indicators capture seasonal demand variations typically observed in retail environments. The temporal granularity of weekly sales data allows the application of time-series forecasting techniques while maintaining sufficient historical depth for model training. Owing to these characteristics, the Walmart dataset serves as an appropriate benchmark for developing and evaluating machine learning–based demand forecasting and dynamic pricing systems in retail supply chain scenarios.

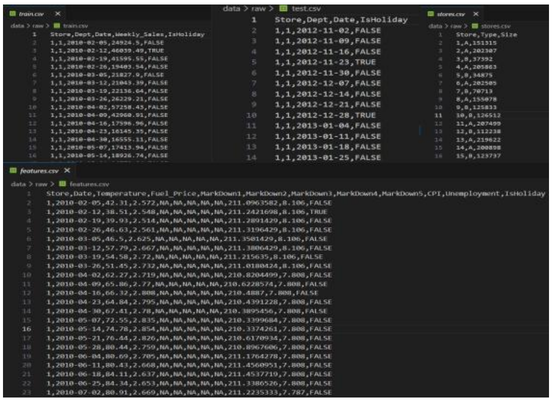

Figure 2 Walmart Dataset

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2362-2372
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0318

1.6.Data Preprocessing and Feature Engineering

Figures should be provided separately from the main text. Data preprocessing is performed to ensure consistency, completeness, and suitability of the dataset for time-series demand forecasting. The sales, store, and feature datasets are merged using store identifiers and date attributes to create a unified dataset containing both sales values and external influencing factors. The date field is converted into a standardized datetime format, and records are chronologically sorted for each store–department combination to preserve temporal order. To capture temporal patterns and seasonality in retail demand, several time-based features are engineered from the date attribute. These include calendar features such as week number, month, and year, which help the model learn recurring seasonal demand trends. Additionally, lag-based features are created using historical sales values, including one-week and four-week lag variables, to incorporate recent demand behaviour into the forecasting process. Rolling statistical features are further introduced to smooth short-term fluctuations and represent local demand trends. In particular, rolling mean values over a four-week window are computed for each store–department pair. These rolling features enhance the model's ability to capture underlying demand dynamics while reducing noise in weekly sales observations. Rows containing missing values introduced by lag and rolling operations are removed to ensure data integrity. The resulting feature-engineered dataset is then used as input for training the machine learning–based demand forecasting model. This preprocessing pipeline is designed in a modular manner, making it scalable and compatible with distributed data processing frameworks for large-scale retail deployments. To illustrate the structure of the processed datasets used in the forecasting pipeline, Figure 3 presents samples of the merged data and the feature-engineered dataset. Only the initial rows are shown for brevity[2].

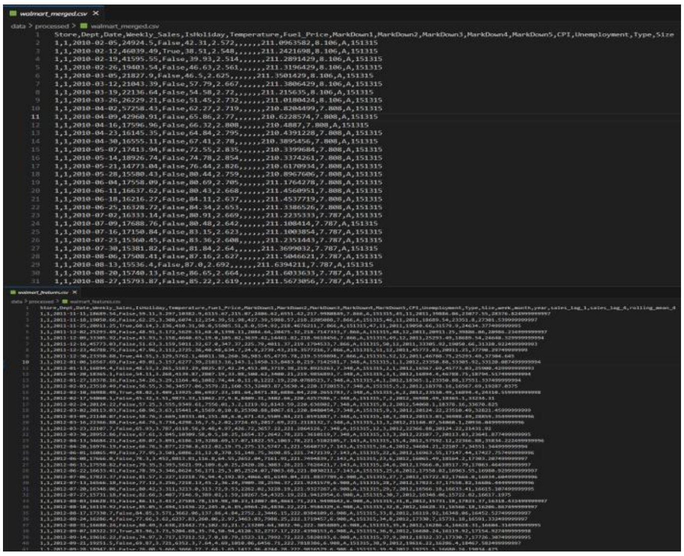

Figure 3. Feature-engineered Dataset

1.7. Demand Forecasting Model

Machine learning–based demand forecasting engines have demonstrated strong performance in retail analytics by leveraging historical sales data and engineered temporal features (Anjaneyulu & Reddy, 2024). To predict product demand across multiple store locations and departments, a supervised machine learning–based regression approach is adopted. Specifically, a Light Gradient Boosting Machine (LightGBM) regressor is employed due to its efficiency, scalability, and strong performance on structured tabular data commonly found in retail applications. Gradient boosting models are well-suited for capturing non-linear relationships and complex feature interactions without extensive feature normalization. The demand forecasting model is trained using the feature-engineered dataset described in the previous subsection. Input features include historical sales lag variables, rolling statistical features, calendar-based features, holiday indicators, and external factors such as temperature, fuel price, consumer price index, and unemployment rate. The target variable is the weekly sales value for each store–department combination. To ensure temporal consistency and avoid data leakage, a time-aware train–test split strategy is applied. Historical records prior to a specified cutoff date are used for training, while more recent observations are reserved for testing. This approach closely reflects real-world forecasting scenarios,

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2362-2372
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0318

where future demand must be predicted using only past information. Model performance is evaluated using standard regression metrics, including Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). These metrics provide quantitative insight into the average prediction error and the variability of forecasting performance. The trained LightGBM model demonstrates the ability to capture demand trends and seasonal patterns across different products and store locations, making it suitable for downstream pricing optimization and decision-support tasks. Figure 4 displays the line plot of Demand Forecast vs Actual Sales [3].

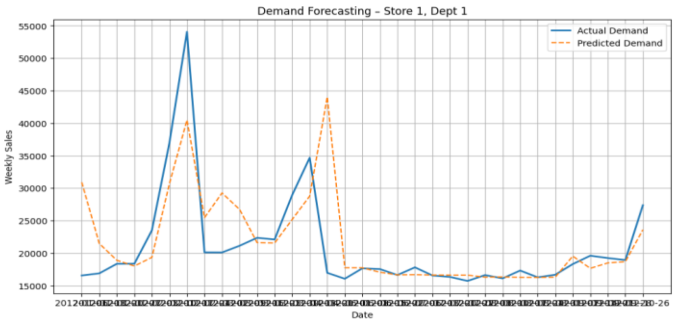

Figure 4 Demand Forecast vs Actual Sales

1.8. Dynamic Repricing Strategy

Dynamic pricing plays a critical role in modern retail supply chains by enabling organizations to adapt prices in response to changing demand, inventory levels, and market conditions. Dynamic pricing strategies that leverage real-time demand signals and market conditions have been shown to improve competitiveness and revenue optimization in retail environments (Dynamic Pricing Strategies in Retail, 2025). Artificial intelligence–driven pricing mechanisms are increasingly used to adapt prices dynamically in response to demand fluctuations and sustainability objectives (Abdullah et al., 2025). Integrating dynamic pricing decisions with demand-driven supply planning enables coordinated optimization of inventory utilization and pricing strategies (Integrating Dynamic Pricing Strategies…, 2024). In the proposed system, demand forecasts generated by the machine learning model are utilized as inputs to a rule-based dynamic repricing module designed to simulate real-world pricing decisions. The repricing strategy considers three primary factors: predicted demand, available inventory, and competitor pricing signals. Predicted demand values obtained from the forecasting model are compared against simulated inventory levels to estimate demand–supply imbalance. Inventory values are generated to reflect realistic stock constraints typically encountered in retail environments, while competitor pricing is simulated to represent market-driven price fluctuations. A demand–inventory gap metric is computed to quantify the relative difference between forecasted demand and current inventory. Based on this gap, pricing adjustments are applied using predefined thresholds. When predicted demand significantly exceeds available inventory, prices are increased to manage demand pressure and maximize revenue. Conversely, when inventory exceeds predicted demand, prices are reduced to stimulate sales and minimize overstock losses. If the demand–inventory gap remains within acceptable bounds, prices are maintained at a base level. This rule-based pricing mechanism provides interpretability and controllability, making it suitable for integration into decision-support systems. While the current implementation employs deterministic rules for pricing adjustments, the framework is extensible and can be enhanced in future work using reinforcement learning or elasticity-based optimization models. The dynamic repricing module operates in conjunction with the demand forecasting pipeline, enabling end-to-end intelligent pricing recommendations under varying supply chain conditions. Figure 5 shows an example of the price recommended based on the predicted demand and inventory stock availability

Figure 5. Example of price recommendation

1.9. Real-Time Data Ingestion Simulation

Modern retail systems require continuous ingestion of sales transactions to enable timely demand forecasting and dynamic decision-making. In real-world deployments, such data streams are typically handled using distributed event-streaming platforms such as Apache Kafka. However, due to infrastructural and experimental constraints, this work implements a simulated real-time data ingestion pipeline that closely mirrors Kafka-based streaming behaviour as shown in Figure 6. The simulation is designed to emulate live sales updates by sequentially feeding transaction records from the Walmart dataset into the forecasting pipeline at fixed time intervals. Each incoming record represents a new weekly sales observation for a specific store and department, thereby mimicking the arrival of real-time sales events. This incremental ingestion approach allows the system to continuously generate updated demand predictions and pricing recommendations without retraining the model from scratch. The simulated streaming mechanism is integrated with the prediction and storage layers of the system. As new data events are ingested, they are processed by the trained forecasting model, and the resulting demand estimates and pricing decisions are stored for downstream consumption. This design demonstrates how the proposed architecture can seamlessly transition to a fully distributed streaming setup using Kafka or similar platforms in future deployments. By incorporating real-time data ingestion simulation, the system effectively bridges the gap between batch-trained machine learning models and live retail environments, ensuring responsiveness, scalability, and adaptability to changing demand patterns [4].

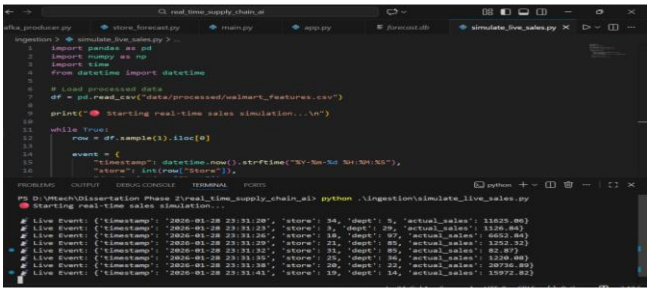

Figure 6. Simulated real-time data ingestion

1.10. API and Deployment Layer

To enable real-time access to demand forecasts and dynamic pricing recommendations, the proposed system is deployed as a lightweight RESTful microservice using FastAPI as it can be seen in Figure 7. FastAPI is selected due to its high performance, asynchronous request handling, and built-in support for API documentation through OpenAPI specifications. This design allows seamless integration of the forecasting system with external applications, dashboards, or enterprise decision-support systems. The API layer acts as an interface between the machine learning pipeline and end users. Upon receiving a request, the service retrieves the most recent demand predictions and pricing outputs from the storage layer and returns them in a structured JSON format. This on-demand inference capability enables near real-time decision-making without exposing the internal model complexity. The service is deployed locally using an ASGI server (Uvicorn), ensuring rapid response times and ease of scalability. The modular design of the API allows future extensions such as authenticated endpoints, batch prediction requests, and cloud-native deployment using containerization technologies. Additionally, automatic API documentation facilitates testing and validation during development and experimentation. Overall, the API and deployment layer ensures that the forecasting and pricing intelligence generated by the system is accessible, scalable, and production-ready, thereby completing the end-to-end supply chain optimization pipeline [5].

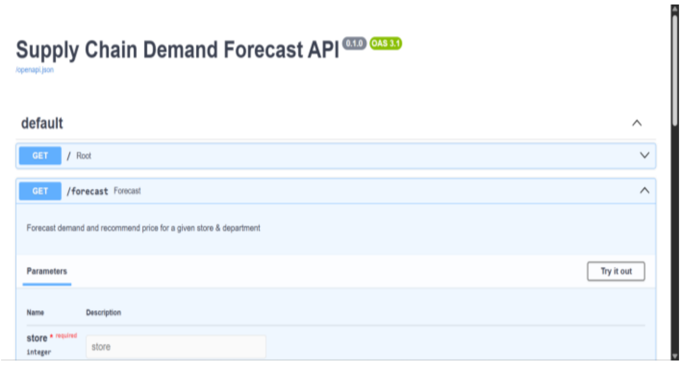

Figure 7 Supply Chain Demand Forecast API

2. RESULTS AND DISCUSSION

2.1. Experimental Results

The proposed demand forecasting and dynamic repricing system is evaluated using the Walmart Store Sales Forecasting dataset to assess its predictive accuracy and practical applicability in retail supply chain scenarios. The dataset is divided into training and testing sets using a time-aware split to ensure that future sales are predicted only from historical information. The performance of the demand forecasting model is quantified using Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE), which are widely adopted metrics for evaluating regression-based forecasting models (Table 1). The trained LightGBM model achieved an MAE of 2126.81 and an RMSE of 6303.62 on the test dataset. These results indicate that the model effectively captures overall demand trends and seasonal patterns across multiple store–department combinations, while maintaining robustness against outliers. Visual inspection of the demand prediction plots further confirms the model's ability to follow actual sales behaviour over time. The predicted demand curves closely align with observed weekly sales, particularly during non-holiday periods, while slight deviations are observed during extreme seasonal peaks. This behaviour is expected in retail environments where promotional events and holidays introduce sudden demand fluctuations. The dynamic repricing module responds to the forecasted demand by adjusting recommended prices based on simulated inventory levels and competitor pricing. During periods of projected excess demand relative to available inventory, the system recommends price increases, whereas price reductions are suggested when forecasted demand falls below inventory thresholds (Table 2). This adaptive pricing behaviour demonstrates the practical usefulness of integrating demand forecasting with pricing optimization. Overall, the experimental results validate the effectiveness of the proposed system in generating accurate demand forecasts and meaningful pricing recommendations, supporting its applicability for intelligent retail supply chain management [10].

Table 1 Model Performance Metrics

Metric	Value
Mean Absolute Error (MAE)	2126.81
Root Mean Squared Error (RMSE)	6306.62

Table 2 Repricing Module Behaviour

Condition	Pricing Action
Forecasted demand > inventory	Price increase
Forecasted demand < inventory	Price decrease
Stable demand	Minimal or no price adjustment

2.2. Discussion

The experimental results demonstrate that the proposed machine learning–driven demand forecasting and dynamic repricing system can effectively model sales behaviour in a complex retail environment. Data sparsity and demand volatility remain persistent challenges in retail forecasting, motivating the use of advanced data augmentation and synthetic data techniques in supply chain analytics (Leveraging Synthetic Data…, 2025). Data-driven machine learning approaches have shown consistent improvements over traditional statistical models in demand prediction and supply chain risk management (Zhang et al., 2024). The obtained MAE and RMSE values indicate that the LightGBM-based forecasting model is capable of learning meaningful temporal patterns from historical sales data while incorporating external factors such as holidays, economic indicators, and weather conditions. Although prediction errors increase during peak seasonal periods, such as major holidays, this behaviour reflects the inherent volatility of retail demand rather than model deficiency. Sudden demand surges caused by promotions and consumer behaviour shifts are challenging to forecast precisely using historical data alone. Nonetheless, the model maintains stable performance across regular sales periods, which constitute the majority of retail

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2362-2372
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0318

operations. The integration of demand forecasting with dynamic pricing further enhances the practical value of the system. By adjusting prices based on predicted demand and inventory availability, the system supports proactive decision-making aimed at reducing stock-outs and minimizing excess inventory. Even though competitor pricing and inventory levels are simulated in this study, the pricing logic demonstrates realistic and interpretable behaviour aligned with retail pricing strategies. From a system design perspective, the modular architecture enables seamless interaction between data ingestion, forecasting, pricing, storage, and visualization components. This separation of concerns improves scalability and allows individual components to be upgraded independently. The simulated real-time ingestion and API-driven deployment further highlight the system's readiness for real-world integration. Overall, the discussion highlights that the proposed approach strikes a balance between predictive accuracy, interpretability, and deployability, making it suitable for intelligent retail supply chain optimization while leaving scope for further enhancements [6].

3. Visualization And Decision Support

An interactive visualization layer is developed to translate the outputs of the demand forecasting and dynamic pricing modules into actionable insights for retail decision-makers as shown in Figure 8. While predictive models generate numerical outputs, effective visualization is essential for interpreting trends, identifying risks, and supporting operational decisions in real time [11]. The visualization component is implemented as a web-based dashboard that interacts with the forecasting and pricing outputs stored in the system. The dashboard enables users to explore demand patterns, pricing recommendations, and inventory risks across different store and department combinations. This layer serves as the primary interface between the intelligent backend system and business stakeholders. By integrating forecasting intelligence with intuitive visual representations, the dashboard enhances transparency and usability, allowing users to understand not only what decisions are recommended but also why those decisions are made. This decision-support capability is particularly important in retail environments where timely and informed interventions can significantly impact profitability and customer satisfaction.

Figure 8. Visualization Dashboard

3.1. Interactive Dashboard Design

The interactive dashboard is designed to provide an intuitive and user-friendly interface for exploring demand forecasts and dynamic pricing recommendations across multiple retail dimensions as shown in Figure 9. The dashboard allows users to select specific store and department combinations through interactive filters, enabling focused analysis of localized demand behaviour. Demand forecasts are visualized using time-series plots that compare historical sales with predicted demand values. This side-by-side representation allows users to assess forecasting accuracy and identify trends, seasonality, and anomalies over time. Such visual feedback is essential for building trust in predictive systems and supporting informed decision-making. In addition to demand visualization, the dashboard presents dynamic pricing trends derived from the repricing module. Recommended prices are displayed as temporal curves, illustrating how pricing strategies evolve in response to changing demand and inventory conditions. This visualization helps stakeholders understand the relationship between demand forecasts and pricing actions. The dashboard is designed with scalability in mind, allowing future integration of real-time streaming data and additional analytical views. By combining interactivity with predictive insights, the dashboard serves as an effective decision-support tool for retail supply chain management [7].

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2362-2372
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0318

Figure 9. Interactive Dashboard Design

4. Inventory Risk Monitoring

Effective inventory management is a critical component of retail supply chain operations, as both stock-outs and overstocking can lead to significant financial losses [8] . To address this challenge, the proposed system incorporates an inventory risk monitoring mechanism that leverages demand forecasts to proactively identify potential inventory imbalances. Forecast-driven inventory monitoring plays a critical role in identifying potential stock-out risks and reducing operational losses in retail supply chains (Machine Learning Approach to Inventory Stockout Prediction, 2024). Within the dashboard, inventory levels are compared against predicted demand to detect scenarios where forecasted sales exceed available stock. Such conditions are flagged as potential stock-out risks, enabling early intervention through replenishment or pricing adjustments. These alerts provide decision-makers with timely insights into high-risk products and store locations. The inventory risk monitoring component complements the dynamic pricing strategy by offering an additional layer of operational intelligence. While pricing adjustments help regulate demand, inventory alerts support logistical planning and supply chain coordination. Together, these capabilities enable a holistic approach to demand-driven retail management. By visualizing inventory risks alongside demand and pricing information, the system enhances situational awareness and supports data-driven decisions aimed at maintaining service levels and optimizing inventory utilization. Figure 10 showcases alerts in case of potential stock out situations while Figure 11 shows scenarios where there are no inventory stockout risks [9].

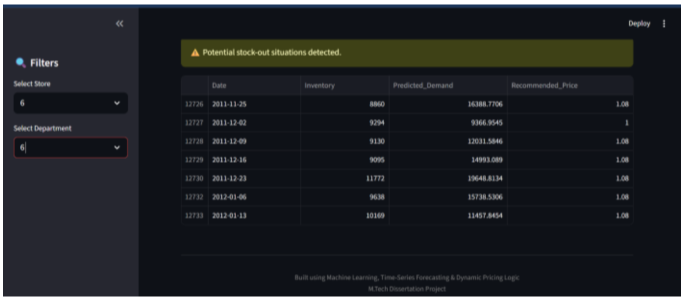

Figure 10. Potential Stockout Risk Detection

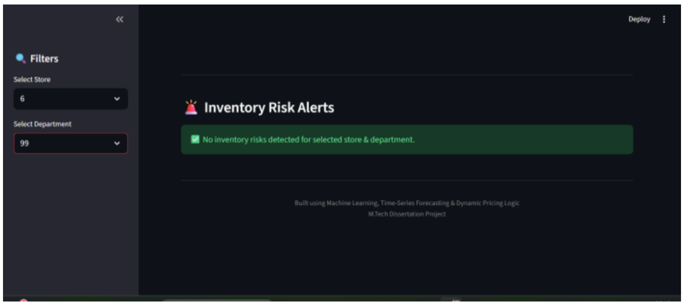

FIGURE 11. No Stockout Risk Scenarios

Conclusion

This paper presented a real-time intelligent supply chain demand forecasting and dynamic repricing system designed to support data-driven decision-making in retail environments. Using the Walmart Store Sales Forecasting dataset, the proposed system integrates historical sales data, external influencing factors, and machine learning techniques to generate accurate demand predictions across multiple store and department combinations. A gradient boosting–based forecasting model was developed and evaluated using standard regression metrics, demonstrating its ability to capture temporal demand patterns and seasonal variations. The forecasted demand was further utilized to drive a dynamic pricing strategy that adapts prices based on demand intensity and inventory availability. In addition, a simulated real-time data ingestion pipeline and a RESTful API layer were implemented to reflect real-world deployment scenarios. An interactive visualization dashboard was developed to present demand trends, pricing recommendations, and inventory risk alerts intuitively. Together, these components form a complete end-to-end intelligent retail supply chain system that bridges predictive analytics and operational decision support. The results indicate that the proposed architecture is effective, scalable, and adaptable, making it suitable for modern retail applications where responsiveness and data-driven strategies are essential. While the proposed system demonstrates the effectiveness of machine learning–driven demand forecasting and dynamic repricing, several extensions can further enhance its accuracy, scalability, and real-world applicability. Future extensions of this work may explore multi-agent deep reinforcement learning frameworks to jointly optimize demand forecasting and inventory management decisions (Yang et al., 2025). Future work can explore the integration of deep learning–based time series models such as Long Short-Term Memory (LSTM) networks and Transformer architectures to better capture long-term temporal dependencies and complex demand patterns. Deep reinforcement learning–based pricing models can be investigated to enable adaptive pricing strategies under supply chain uncertainty and disruption risks (Zhu et al., 2024). The simulated real-time ingestion framework can be replaced with a fully distributed streaming infrastructure using Apache Kafka or similar platforms to enable continuous learning and low-latency updates in production environments. Additionally, incorporating NoSQL databases such as MongoDB or Cassandra can improve scalability and flexibility for handling large volumes of real-time transactional data. Cloud-native deployment using containerization and orchestration technologies can further enhance system reliability and scalability. Integration of real competitor pricing feeds and real-time inventory management systems would also improve the realism and effectiveness of the dynamic repricing strategy. Overall, these enhancements provide a clear pathway for transforming the proposed system into a fully production-ready intelligent retail supply chain platform.

Acknowledgements

The author sincerely thanks the project guide and project coordinator of the M.Tech program at Royal College of Engineering and Technology for their valuable guidance and continuous support throughout this work. The author also extends gratitude to fellow M.Tech classmates for their encouragement, discussions, and collaborative support during the project.

References

[1] Machine Learning and Deep Learning Models for Demand Forecasting in Supply Chain Management: A Critical Review, Scientific Reports, Nature Publishing Group, 2024.

[2] G. V. Radhakrishnan and U. Shankar, “Predictive Analytics in Supply Chain Management: The Role of AI and Machine Learning in Demand Forecasting,” Journal of Informatics Education and Research, vol. 4, no. 3, 2024.

[3] A. Rabbani, N. Sohu, N. Zaki, R. Muhammad, M. U. Hakeem, F. E. A. Longa, and A. Fatima, “Towards Next-Generation Smart Supply Chain Management: AI-Driven Big Data Analytics for Robust Demand Forecasting,” Center for Management Science Research, vol. 3, no. 5, pp. 391–431, 2025.

[4] N. Anjaneyulu and C. K. Reddy, “Modernising Supply Chain Analytics with Demand Forecasting Engine,” International Journal for Multidisciplinary Research (IJFMR), 2024.

[5] Dynamic Pricing Strategies in Retail: Leveraging Real-Time Data Analytics for Competitive Advantage, ResearchGate Preprint, 2025.

[6] N. Abdullah, A. K. Shaikh, and A. Almusharraf, “Towards a Sustainable Retail Food Chain: Artificial Intelligence Driven Dynamic Pricing and Promotions for Food Waste Reduction,” Journal of Posthumanism, vol. 5, no. 6, pp. 571–588, 2025.

[7] Integrating Dynamic Pricing Strategies and Demand-Driven Supply Planning in Wood Pellet Supply Chains: A Stochastic

Optimization Approach, Applied Sciences, vol. 14, no. 18, p. 8110, MDPI, 2024.
[8] Leveraging Synthetic Data to Tackle Machine Learning Challenges in Supply Chains: Methods and Applications, Computers & Industrial Engineering, Elsevier, 2025.
[9] J. Zhang, Y. Wang, and Z. Wang, "Enhancing Supply Chain Forecasting with Machine Learning: A Data-Driven Approach to Demand Prediction, Risk Management, and Demand Supply Optimization," Journal of Fintech and Business Analysis, vol. 2, no. 1, 2024.
[10] A Machine Learning Approach to Inventory Stockout Prediction, Journal of Manufacturing Systems, ScienceDirect, 2024.
[11] Y. Yang, M. Wang, J. Wang, P. Li, and M. Zhou, "Multi-Agent Deep Reinforcement Learning for Integrated Demand Forecasting and Inventory Optimization in Sensor-Enabled Retail Supply Chains," Sensors, vol. 25, no. 8, p. 2428, MDPI, 2025.
[12] C. Zhu, J. Xin, and D. Zhang, "A Deep Reinforcement Learning Approach to Dynamic E-Commerce Pricing Under Supply Chain Disruption Risk," Annals of Applied Sciences, vol. 5, no. 1, 2024