Machine Learning Framework for Food Demand Forecasting and Inventory Optimization

Diya Shah[1*], Harsh Majithia[2], Parmi Kenia[3], Tirth Shah[4], Avinash Tandle[5]
Department of Electronics and Telecommunications Engineering, Mukesh Patel School of Technology Management Engineering, SVKM's Narsee Monjee Institute of Management Studies (NMIMS) University, Mumbai-400056, India

diyashah2300603@gmail.com [1*], harshmajithia99@gmail.com [2], keniaparmi9@gmail.com[3], tirthshah1202@gmail.com
Avinash.tandle@nmims.edu[5]

Abstract. In an evolving supply chain management landscape, accurate demand forecasting and inventory management are critical for cost reduction and meeting customer demands. This paper proposes a machine learning model to forecast food demand and manage inventory using real-world data such as fulfilment centre operations, meal data, and past order quantities. We use various machine learning algorithms such as Gradient Boosting to develop and predict demand with high accuracy. The developed solution combines demand forecasting and simple inventory logic to determine stock needs and prevent over- and understocking. The result is that the Gradient Boosting model had a lower prediction error than other algorithms. The proposed solution demonstrates how data-driven forecast models could significantly improve operational efficiency and decision-making processes in food logistics.

Keywords: Mean absolute error (MAE), Mean squared error (MSE), Root mean square error (RMSE), R^2 Score, Machine Learning, Demand Forecasting, Inventory Optimization, Gradient Boosting, XGBoost, Ensemble Learning, Time-Series Analysis, Supply Chain Analytics, Feature Engineering

1 Introduction

Accurate demand forecasting is an essential component of supply chain and inventory management in today's fast-paced, data-driven economy. For food delivery, where seasonality, holidays, customer preferences, and location can drastically affect demand, traditional methods are often insufficient [1]. Overstocking results in higher storage and waste costs, while understocking can lead to lost sales and unhappy customers. Therefore, having a solid and reliable forecasting model can help optimize resource usage, reduce operating costs, and make customers happier [2].

Prediction modeling has seen many improvements with the introduction of machine learning and data analytics. These technologies are capable of extracting complex relationships from historical data and are capable of providing more.

The main scope of this project is the application of machine learning algorithms to forecast food demand and the implementation of a simple inventory management system to support decision-making in supply chain management [5]. The data set considered in this article includes details of meals, fulfilment centres, and historical demand records. The study applies and compares multiple ML models, including Gradient Boosting. The study investigates the most effective approach to forecasting demand. Integration of these models with a simple rule-based inventory management system provides insight to maintain the most appropriate stock levels and balance supply and demand. Research on this not only improves the accuracy of demand forecasting, but also shows the feasibility of integrating machine learning with operational decision-making in the context of food logistics. [1][5]. Although various statistical-based traditional approaches are available for statistical-based forecasts, there is still a failure to recognize the intricate non-linear relationships as well as the changing demand trends with regard to the food supply chain. Secondly, none of the existing research has successfully incorporated the forecasting mechanisms with the existing

S. Sharma et al. (eds.), *Proceedings of the International Conference on Recent Advances in Intelligent and Sustainable Technologies (RAIST 2026)*, Atlantis Highlights in Intelligent Systems 17,
https://doi.org/10.2991/978-94-6239-707-1_12

operational mechanisms of an inventory system. This piece of research attempts to solve these concerns by developing an efficient machine learning-based model.

This paper is structured as follows: Section II provides a review of the literature on demand forecasting and inventory management and elucidates the high performance of various machine learning models. Section III explains the method employed in this research, Summarizing the integration of the dataset, data pre-processing, and model building. Section IV provides the results obtained from the demand forecasting models, their performance comparison, and the evaluation of the application interface. Section V presents results and Discussion, including exploratory data analysis, model evaluation, and insights into application interface design. Finally, the paper is summarized in Section VI, which discusses the findings and suggests avenues for future research in inventory optimization and food logistics.

2 Literature Review

Demand forecasting, inventory management and other supply chain capabilities are critical to efficient operations, especially within those industries such as food delivery where demand volatility drives operating inefficiencies. In recent years there has been a substantial surge in the use of machine learning (ML) methods for such problems, offering new opportunities to enhance forecast accuracy and enable data-driven decision making. One of the popular machine learning model that performs time series forecasting and regression is XGBoost (Extreme Gradient Boosting). XGBoost [1] was used to model complex non-linear relationships and interactions between features, with a recorded RMSE (0.7015), for inventory demand prediction. The findings in this study stressed that the regularization techniques and tree-pruning algorithms of XGBoost are useful to avoid overfitting in noisy, real-world data. To enhance predicting power, the paper also advocated to include additional context of seasonality and external factors. Another ensemble learning approach GBR was studied in [2] where it modeled variance on retail demand. By achieving superior forecast accuracy over linear regression and ARIMA, the model proved particularly good at capturing seasonality, outliers, and long-tail nature of product demand. For both inventory and sales forecast, Gradient Boosting is a potential candidate for practitioners due to its potential flexibility to fit different kind of data. In retail and e-commerce datasets, a more generic comparison study has evaluated the performance of several other ML models such as Random Forest, Linear Regression, SVR and Gradient Boosting [3]. That brings us the conclusion repeatedly where ensemble methods outperform machines, when comparing MAE, MSE and R2 scores. The authors finally concluded that GradientBoost and its derivatives better manage multicollinearity and offer more immunity to feature correlation. Hybrid- techniques, which combine the strengths of different methods to achieve better generalization and forecast stability, are a new research area of interest. Mixing XGBoost with Gradient Boosting was proposed in [4] by taking advantage of XGBoost's excellent performance on big datasets and GBR's ability to capture complex relationships over time. This ensemble-fusing method exhibited great antioverfitting effect and surpassed single model over a range of evaluation metrics [5] [6]. The authors noted that hybrid models offer a viable means of improving short- and mid-term demand forecasting in volatile markets. Inventory-related ML systems of stock levels have been discussed in [7] where the prediction is counted using regression models empowered with time-series decomposition. The research pointed out that combining forecast models with inventory logic rules could facilitate automatic generation of reorder decision and shall increase profitability by reducing holding cost. Indeed, as further highlighted as in [8] there is also the need of predictive models that decide on supply planning and restocking policies and cost optimization across an entire operational context.

3 Methodology

This paper presents an end-to-end inventory management and demand forecasting solution based on the hybrid model architecture. Its approach is based on 5 main components which are described as : (i) dataset integration, data preparation and feature engineering; (ii) model training and hybridization; (iii) inventory logic implementation. including an application-level integration using a web interface, as well as blockchain communication. Fig.1 Flowchart of the proposed solution The overall architecture of the proposed solution is illustrated in Fig 1. It begins with the integration of datasets and pre-processing, moves on to model training using Gradient Boosting, XGBoost and a hybrid ensemble then dives in inventory logic developments and deployment over web application and communicating with blockchain for secure transferring data. Each block represents a stage on the pipeline from raw data to insights.

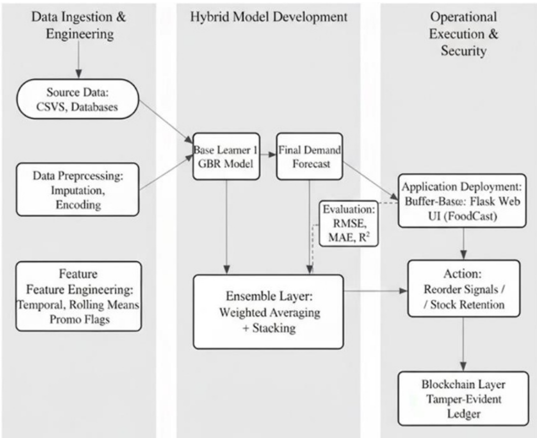

Fig. 1. **End-to-end workflow of the proposed food demand forecasting and inventory optimization framework.**

3.1 Dataset Integration

It is a structured set of data-sets relating to cooking proceedings or food delivery demand at the fulfilment centre. The primary dataset comprises transaction logs of meal orders placed weekly along with attributes like price, promotions, meals identifier and distribution centre identifier. A few additional datasets just provide context information such as the meal category, cuisine type, and geographic / operational characteristics of a fulfilment centre. These two data sets were integrated into enriched common dataset suitable for machine learning. The dataset was very neatly formatted, requiring minimal modification and presented sufficient diversity for effective model training and testing.

3.2 Data Preprocessing and Feature Engineering

After the data integration, several pre-processing procedures were performed to convert the integrated dataset into a format suitable for building machine learning models As both numerical and categorical variables are present in the merged datasets, categorical variable encoding was ap- plied to transform these features—such as meal type, cuisine and fulfilment centre attributes —into machine-readable representations[3][2]. We implemented feature engineering to craft features that enhance the predictive power of models. These parameters included price differences to account for promotional sales, binary variables that indicated if promotion was offered on the home- page or through email, and time features (i.e., week and month) to adjust for seasonality [4][5]. Rolled features for history of demand per centre, including the rolling mean, were introduced to help models learn short- and long-term trend patterns [6]. These techniques, widely promoted in retailing and logistics literature [7][8][9], improved model generalization, provided more accurate demand forecast of different meal types across regions.

3.3 Model Development and Hybrid Ensemble Learning

For the purpose of accurate predicting food demand, three machine learning approach like Gradient Boosting Regressor, XGBoost Regressor and a hybrid ensemble model combining both are employed in this research. These models were chosen, as they are capable of learning complex non-linear relations, dealing with high dimensional data efficiently and performing well in real world forecasting application [10]. The Gradient Boosting Regressor builds strong learners by boosting the weak ones sequentially and its hyperparameters was tuned such as learning rate, maximum depth and number of estimation [1] [3] [8]. It is assessed that XGBoost optimizes gradient boosting, regularization, parallel computation and tree pruning that result in better speed and accuracy especially when they are dealing with big and noisy datasets [4][11][12]. To further improve prediction accuracy, a mixed ensemble model with weighted averaging and stacking between the predictions of the above two models is proposed where predictions are fed to a meta-learner [6][13][14]. All models were evaluated with cross-validation and common regression measures-RMSE, MAE, R^2. The ensemble model was superior to the base learners due to better generalization in test data [3][9].

Let the dataset be represented as **D = {(xi, yi)}**,

where xi denotes the feature vector

(price, promotions, temporal and location attributes)

yi denotes weekly food demand.

The Gradient Boosting model minimizes squared error loss:

$$\mathbf{L = \Sigma (yi - \hat{y}i)^2}$$

by iteratively adding weak learners:

$$\mathbf{Fm(x) = Fm{-}1(x) + \eta\ hm(x)}$$

where η is the learning rate.

The hybrid ensemble prediction is computed as:

$$\mathbf{\hat{Y}\ ensemble = \alpha\ \hat{y}\ GB + (1 - \alpha)\ \hat{y}\ XGB}$$

where α is empirically selected based on validation performance.

3.4 Inventory Logic Implementation

To close the loop between demand planning and operational execution, rudimentary rule based stock logic was integrated with the project. The software translates forecast demand into actionable supply chain signals that can be acted upon in real-time to prevent stock outs and overstocks. It compares forecasted meal orders to inventory levels and, when demand exceeds available stock plus an assigned buffer, creates a reorder signal. Such a buffer-based system takes into account demand uncertainty, delivery lead time and surge phenomenon as techniques used to automated stock replenishment systems [15][16][17]. Simple but scalable and modular architecture Focus has been on keeping the logic layer free of any entanglement, so that it can later be integrated with ERP systems or reporting tools in a way which suits your needs. It can be further generalized to support non-deterministic buffer levels, variable lead time and cost-based optimization [5][12]. In overall, integration of predictive models and inventory decision logic enhances supply chain agility and operational efficiency, particularly in the dynamic environment such as food logistics [4][6].

Inventory optimization can be defined as the process of setting optimal stock levels that result in a balance involving holding costs and costs, shortage risks, and service-level requirements. Although metaheuristic methods such as GA: Genetic Algorithms and PSO: Particle Swarm Optimization are commonly used for complex inventory systems, the current study employs a forecast-driven rule-based approach appropriate in real-time systems. operational environments. The proposed approach emphasizes interpretability, scalability, and ease of deployment while maintaining effective inventory control.

3.5 Application Deployment and Blockchain Communication

To aid in usability and open access, a user-friendly web interface was developed with Flask. Users are presented with input fields to input parameters such as meal type, cuisine, price, promotions and distribution center parameters. This is then fed into a serialized machine learning model (.pkl) to process payfiles7, after which they produce real-time demand forecast and display them on the frontend for decision making. Firebase was used to enable real-time message-based communication between applications of different stakeholders including restaurants, warehouses and suppliers for open and coordinated supply chain operations [13][10][14]. Further, the platform includes an in-house developed blockchain module which generates transaction (test orders and messages) into a tamper-evident decentralized record maintained with proof-of-work [18][19]. This ensures integrity, traceability of data and trust between distributed stakeholders. Combining machine learning, real-time messaging and blockchain, the system offers security as well as intelligence and demonstrates how a prediction engine can be re-invented into an intelligent distributed app for a dynamic logistics world.

4 Results and Discussion

4.1 Exploratory Data Analysis (EDA)

Preliminary and non-parametric EDA was conducted to examine the data frame, distribution and relationship between aggregated variables. The authors found significant bursts of time patterns in meal demand, characterized by the appearance of two-weekly repeating cycles and peaks of demand around holidays. Order volumes in urban fulfilment centers were consistently higher when compared with rural, indicating that location influences customer behaviour. Correlation analysis showed that the main variables, such as base price, whether with promotion or not and cuisine type had a significant effect on demand. Few missing data were easy to address with simple imputa- tion. Outliers (in particular at peak demand hours) were retained to replicate real-world variability. Notions such as this led to the development at both the feature-engineering and model-architecture levels — particularly incorporating temporal and promotional features to better capture demand variation.

a. Weekly Order Trend: As shown in Figure 2, there exists extreme fluctuations with weekly food orders. The x-axis represents weeks and the y-axis shows number of orders, with many sharp rises and declines along the plot. The fluctuation like this needs the strong demand forecasting model to handle seasonality and non-stationarity.

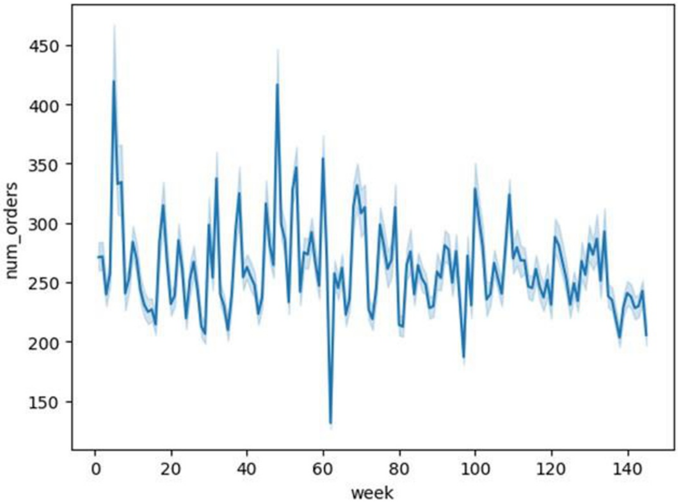

Fig. 2. **Weekly variation in the number of food order**

b. Category-wise Cuisine Distribution: In Figure 3, a multi-dimensional perspective is shown by intersecting category and cuisine types. Italian cuisine is prevalent in categories such as 'Salad' and 'Sandwich', whereas Indian cuisine is prevalent in 'Rice Bowl' and 'Desert'. These types of relationships are useful for cross-category demand analysis.

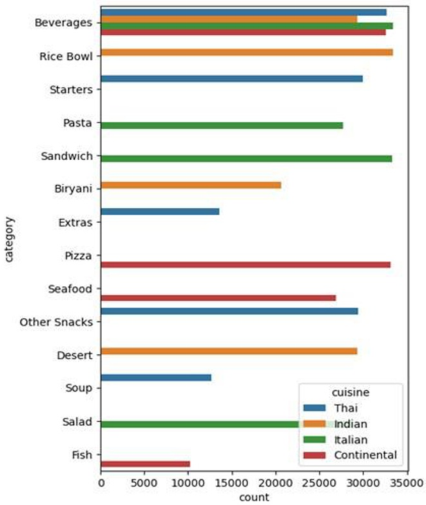

Fig. 3. **Distribution of food categories by cuisine.**

c. Correlation Matrix of Training Features The heatmap of correlation in Figure 4 indicates that the highest positive correlation (in dark red) is between base price and checkout price (0.95), reflecting a strong price relationship. The lowest correlation (in dark blue) is between meal id and Num orders (–0.28), reflecting little impact of meal ID on demand. Also, moderate positive correla- tion exists between Num orders and promotion features such as email for promotion and homepage featured, whereas pricing features have a minimal negative correlation with demand, reflecting price sensitivity.

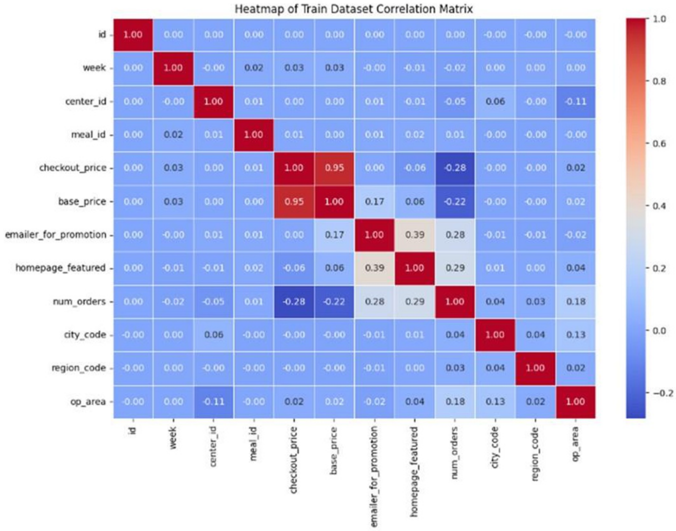

Fig. 4. **Heatmap of correlation matrix**

4.2 Model Evaluation and Comparison

The table 1 represents predictive performance of the Gradient Boosting Regressor, XGBoost Regressor, and the proposed hybrid ensemble model was evaluated using three standard regression metrics: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2).

Table 1: Performance Comparison of Regression Models

Metric	Gradient Boosting	XGBoost	XGB+GB
RMSE (Test)	159.26	155.57	**153.78**
RMSE (Train)	105.64	82.84	**74.02**
MAE (Test)	79.98	**78.97**	80.00
R^2 Score (Test)	0.8356	0.8432	**0.8467**
R^2 Score (Train)	0.9293	0.9565	**0.9653**

The Hybrid Model repeatedly outperformed the individual models in virtually all of the primary metrics, specifically in RMSE, R^2, and variance scores. Though XGBoost displayed marginally improved MAE, the Hybrid Model had better overall generalization and stability.

The hybrid ensemble approach clearly has better performance in terms of generalization, clearly illustrated by the RMSE values and the R^2 values obtained in the hybrid approach, which are better than the others. This is significant for food supply chains, where the demand volatility has a real effect on the cost of the inventory, food waste, and service levels.

Figure 5 illustrates the predicted versus actual values for the XGBoost, Gradient Boosting Regressor, and the proposed Hybrid model(XGB+ Gradient Boosting). In each of the three scatter plots, the best possible situation—where prediction matches exactly with the actual value—is shown as the diagonal blue line. The red data points represent the model outputs, and the distance of the red data points from the blue line shows prediction accuracy. Although both Gradient Regressor and XGBoost are equally well-performing with fairly compact grouping around the line, the Hybrid model depicts the tightest prediction point distribution with less variance and better alignment. This pictorial proof again proves the greater predictive consistency and precision of the Hybrid model over stand-alone models.

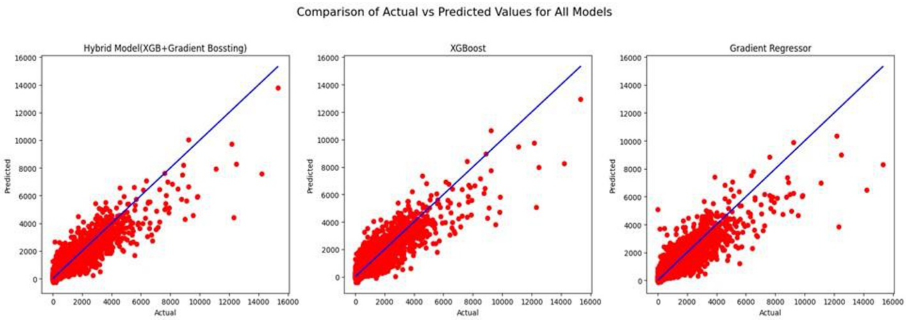

Fig. 5. **Actual vs Predicted Comparison Across Regression Models**

4.3 Application Interface

The collection of images demonstrates an intuitive interface of the FoodCast application, a genius platform aimed at forecasting food demand.

- **Login Screen:** The Figure 6 is a clean and contemporary login screen prompting users to provide a username and password to log into the system.
- **Prediction Input Panel:** The Figure 6 indicates the submission in the core demand forecasting module in which users can enter different parameters like food type (e.g., Bever- ages), cuisine type (e.g., Thai), prediction weeks, base price, and checkout price, promotional information, location codes, and operational area. The system uses these inputs to predict the quantity of food orders.
- **Prediction Output:** The Figure 6 also shows the output produced by the system , a rough estimate of 403 orders in 3 weeks as output from the user-supplied inputs.

This kind of system will most likely employ machine learning-based models to make reliable and decision-relevant forecasting predictions for stock planning and food industry supply.

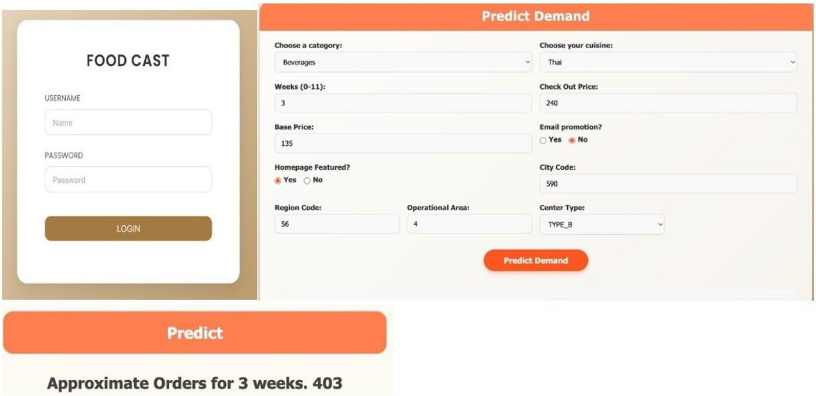

Fig. 6. **FoodCast: Food Demand Forecasting Interface**

5 Conclusion

The present study enables a holistic and practical approach to food demand forecasting and inventory management within machine learning which in the case includes Gradient Boosting, XG-Boost as well as merged version of both. The comparative model showed that the improved predictive performance (RMSE, MAE and R2) were better than each of the individual models. And further, by incorporating rule-based inventory logic system the forecast output was converted to executeable inventory control signals reaching an optimal balance of demand and supply. The development of the web-based application FoodCast with an easy-to-use online interface and a blockchain-supported communication module shows that it is possible to introduce intelligent forecasting tools on real logistics background. This modular and scalable approach improves not only the accuracy of prediction, but also transparency, traceability, and operational response to stakeholders. Overall, this study further validates the essential role of machine learning in smart supply chain management especially for those volatile and sensitive sectors such as food delivery.

6 Future Work

While the present system proves to be promising, many possibilities for future work lie ahead. One of the most interesting directions is extending the blockchain module. Smart contracts can be further integrated in future work for automatic ordering, payment and compliance checking among supply chain partners and thus reducing human intervention and increasing trust. Further, transitioning to the a consortium blockchain model with established members (e.g., suppliers, warehouses, restaurant) can improve data management and collaborative decision making. The forecasting and inventory side of the operation could benefit from more advanced optimization methods, including dynamic buffer calculations, programming lead time estimation and cost sensitive inventory policies which would make them faster and cheaper. Including reinforcement learning or simulation-based approaches may also increase a warehouse's resilience in the face of fluctuating demand trends. Another potential avenue is to generalize the model for multi-objective forecasting so that waste minimized, profit optimized and service level guaranteed, in addition to demand accuracy.

References

[1] S. Sharma and T. Patel, "Inventory Management using Machine Learning," International Journal of Scientific Research in Computer Science, vol. 8, no. 2, pp. 45–50, 2020.

[2] M. Gupta and P. Rathi, "Predictive Analysis for Big Mart Sales using Machine Learning Algorithms," IJERT Conference Proceedings, vol. 10, 2018.

[3] A. Roy and K. Banerjee, "Comparative Analysis of Supervised Machine Learning Models for Sales Forecasting," International Journal of Computer Applications, vol. 182, no. 21, pp. 1–6, 2021.

[4] S. Kumar and R. Chauhan, "Demand Forecasting in Retail using XGBoost and ARIMA," International Journal of Advanced Research in Computer Science, vol. 12, no. 5, pp. 27–32, 2021.

[5] Tekinerdogan, Bedir and Verdouw, Cor and van Haaren, Bram, A Framework for Mapping Inventory Decisions to Machine Learning Methods in Perishable Food Supply Chains 3 Dec 2025.

[6] Wang, L. Zhang, and T. Huang, "Improved Demand Forecasting Model Based on Gradient Boosting and Feature Fusion," IEEE Access, vol. 8, pp. 98841–98849, 2020.

[7] K. Singh and M. Jain, "Machine Learning Techniques for Demand Forecasting: A Review," Journal of Retail and Distribution Management, vol. 47, no. 5, pp. 456–464, 2021.

[8] Y. Chen, G. He, and S. Lin, "Ensemble Learning Approaches for Retail Forecasting," Expert Systems with Applications, vol. 147, p. 113249, 2020.

[9] D. Joshi, R. More, and N. Kadam, "A Comparative Study of Forecasting Models for Inventory Management," International Journal of Engineering Research Technology, vol. 9, no. 3, pp. 143–147, 2020.

[10] H. Zhang and Q. Zhao, "Hybrid Ensemble Models for Retail Sales Forecasting," Elsevier Procedia Computer Science, vol. 186, pp. 556–563, 2021.

[11] Miguel Rodrigues, Vera Miguéis, Susana Freitas, Telmo Machado,Machine learning models for short-term demand forecasting in food catering services: A solution to reduce food waste,Journal of Cleaner Production,Volume 435,2024,140265

[12] ISSN 0959-6526,M. Lee and D. Park, "Smart Inventory Systems using Predictive Analytics," IEEE Transactions on Industrial Informatics, vol. 15, no. 4, pp. 2450–2458, 2019.

[13] N. Sharma, S. Roy, and A. Kumar, "Deployment of ML Models in Web-Based Applications," International Journal of Scientific Technology Research, vol. 9, no. 4, pp. 478–483, 2020.

[14] T. Fang and Y. Yang, "Stacking-Based Ensemble Model for Improved Forecasting Accuracy," Expert Systems with Applications, vol. 182, p. 115196, 2021.

[15] J. Kim, "Automated Stock Replenishment in Retail Using ML," Procedia Computer Science, vol. 160, pp. 620–626, 2019.

[16] A. Srivastava and R. Mehta, "ML-Powered Forecasting in Smart Supply Chains," IEEE Internet of Things Journal, vol. 8, no. 2, pp. 980–987, 2021.

[17] B. Ramesh, P. Kulkarni, and V. Patil, "Real-Time Decision Systems in Logistics Using ML and Firebase," International Journal of Computer Applications, vol. 177, no. 17, pp. 15–20, 2019.

[18] F. Caro and M. S. Martinez-de-Albeniz, "Blockchain in the Supply Chain: Applications and Challenges," Operations Research, vol. 67, no. 6, pp. 1449–1464, 2019.

[19] Pasupuleti, V.; Thuraka, B.; Kodete, C.S.; Malisetty, S. Enhancing Supply Chain Agility and Sustainability through Machine Learning: Optimization Techniques for Logistics and Inventory Management. Logistics 2024, 8, 73. https://doi.org/10.3390/logistics8030073

[20] R. Saberi, M. Kouhizadeh, J. Sarkis, and L. Shen, "Blockchain Technology and Its Relationships to Sustainable Supply Chain Management," International Journal of Production Research, vol. 57, no. 7, pp. 2117–2135, 2019.