International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 11-No.4 (2025) pp. 7260-7269
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

Machine Learning Framework for Retail Sales Forecasting

Padmanabhan Venkiteela*

Senior Enterprise Integration Architect, USA
* **Corresponding Author Email:** padmanabham.research@gmail.com - **ORCID:** 0009-0002-2562-5624

Article Info:

DOI: 10.22399/ijcesen.3993
Received : 28 July 2025
Accepted : 27 September 2025

Keywords

Sales Forecasting,
Machine Learning,
Retail Analytics,
Predictive Modeling,
Low-Code Deployment,
Feature Engineering

Abstract:

Sales forecasting plays a critical role in retail operations, directly impacting inventory management, supply chain planning, and revenue optimization. This paper presents a comprehensive machine learning–based forecasting framework developed for SuperKart, a multi-city retail chain. The proposed approach integrates structured data preprocessing, exploratory data analysis (EDA), feature engineering, and comparative evaluation of regression models including Decision Tree, Random Forest, Gradient Boosting, and XGBoost. Among the tested algorithms, XGBoost outperformed others, achieving an R^2 score of 0.87 and demonstrating strong generalization on test data. A low-code deployment strategy was implemented using Streamlit and Hugging Face Spaces, enabling business users to generate real-time sales forecasts with minimal technical intervention. This research contributes a scalable, accessible, and data-driven forecasting solution for retail environments, bridging the gap between advanced machine learning techniques and practical business adoption.

1. Introduction

Accurate sales forecasting is a cornerstone of retail operations, directly influencing procurement, inventory management, marketing, and financial planning. For large retail chains, even small inaccuracies in forecasts can lead to significant inefficiencies, including overstocking of slow-moving products, stockouts of high-demand items, missed opportunities for targeted promotions, and ultimately reduced profitability. As competition intensifies and customer preferences evolve rapidly, the need for robust, data-driven forecasting solutions has become critical.Traditional forecasting approaches such as spreadsheet-based methods and classical time-series models often fail to capture the complexities of retail environments where outcomes depend on multiple interacting factors, including product attributes (e.g., price, category, display allocation) and store characteristics (e.g., size, location, establishment year). These limitations typically result in forecasts that are overly simplistic and insufficient for guiding optimal business strategies.Recent advancements in machine learning (ML) have opened new opportunities for improving forecasting accuracy by modeling non-linear relationships between diverse variables. Ensemble methods, particularly Gradient Boosting and XGBoost, have demonstrated strong performance on structured datasets due to their ability to handle categorical and numerical features, mitigate overfitting, and provide interpretable results. However, despite their predictive power, adoption of ML solutions is often hindered by operational challenges, as many models are complex to use and require specialized technical expertise.To address the dual challenge of accuracy and usability, this study focuses on SuperKart, a retail chain with operations across Tier 1, Tier 2, and Tier 3 cities. SuperKart faced persistent difficulties in generating reliable sales forecasts, limiting its ability to optimize inventory allocation and tailor regional sales strategies. In response, this research proposes a comprehensive ML-based forecasting framework that integrates data preprocessing, feature engineering, model evaluation, and low-code deployment. The objective is to deliver forecasts that are both accurate and accessible to non-technical business stakeholders.The main contributions of this study are threefold: (1) the development of an end-to-end ML pipeline for retail sales forecasting using SuperKart's dataset, integrating preprocessing, feature engineering, and model evaluation; (2) a comparative analysis of multiple forecasting algorithms, with XGBoost identified as the most effective model (R^2 = 0.852); and (3) the implementation of a low-code deployment framework using Streamlit and Hugging Face Spaces to democratize access to forecasts across the retail network

2. Related Work

Sales forecasting has been extensively studied in both academic and industrial contexts, with approaches evolving from traditional forecasting [1] models to modern machine learning and deep learning techniques.Early methods relied on statistical time-series models such as Autoregressive Integrated Moving Average (ARIMA) and Holt-Winters exponential smoothing [2]. These techniques are effective for univariate series but often fail to capture complex relationships in high-dimensional retail datasets that include diverse products and store-level attributes.With the growth of structured retail datasets, machine learning (ML) models significantly improved forecasting accuracy in retail environments. Tree-based methods such as Random Forest and Gradient Boosting have been applied to predict sales performance by learning non-linear relationships between features [3]. Among these, Extreme Gradient Boosting (XGBoost) [4] has become a widely adopted algorithm due to its scalability, ability to handle mixed data types, and robustness against overfitting [5].More recently, deep learning approaches [6] such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and Temporal Convolutional Networks (TCNs) have been explored for time-series forecasting in retail. While these models excel in capturing sequential dependencies, they typically require large amounts of training data and significant computational resources, making them less feasible for smaller organizations.Despite these advancements, a gap remains in translating ML forecasting solutions into business-friendly applications. Many research implementations stop model evaluation, without addressing challenges in deployment, integration, and usability. Studies focusing on low-code or no-code deployment for forecasting remain scarce, even though such approaches are critical for enabling adoption among non-technical business stakeholders.This research builds upon existing ML forecasting literature by not only comparing model performance across several regression techniques but also proposing a low-code deployment pipeline that bridges the gap between data science research and real-world retail operations.

3. Business Problem Definition

Retail organizations operate in highly dynamic and competitive environments where customer preferences, regional demand variations, and product characteristics continuously shape sales outcomes. For businesses managing diverse product portfolios across multiple geographic markets, generating accurate forecasts becomes particularly challenging. Traditional manual or spreadsheet-based methods often fall short, as they fail to capture the complex interactions between product-level and store-level attributes. SuperKart, a multi-city retail chain operating across Tier 1, Tier 2, and Tier 3 cities, encountered similar challenges in its forecasting processes [7]. The absence of a robust, data-driven mechanism led to operational inefficiencies such as inventory imbalances, with frequent stockouts in fast-moving categories like Snacks and Dairy and overstocking in low-performing categories such as Frozen Foods and Seafood. Regional demand variations also went unaddressed, as manual methods could not capture the higher purchasing power consistently observed in Tier 1 cities compared to Tier 2 and Tier 3 locations. Furthermore, reliance on manual, spreadsheet-based forecasts was both time-consuming and error-prone, while the lack of integration with decision-support systems prevented real-time application of insights in supply chain and marketing strategies.To address these limitations, SuperKart required a data-driven forecasting solution with clearly defined objectives. The system needed to accurately forecast quarterly sales revenue at the store–product level, incorporate product and store attributes such as MRP, Product Type, Store Size, and City Tier into predictive modeling [8] [9], and support regional strategies by enabling demand-aware marketing and targeted promotions. Additionally, the solution had to be deployable through low-code frameworks to make forecasts accessible to non-technical stakeholders like store managers and supply chain planners, while also integrating seamlessly with existing business dashboards to enhance decision-making across procurement, inventory allocation, and expansion planning. By framing the problem in this way, the study highlights the dual challenge facing modern retailers: achieving technical forecasting accuracy while ensuring that solutions remain practical, usable, and scalable in real-world business contexts.

4. Dataset and Exploratory Data Analysis (EDA)

4.1. Dataset Description

The data set [10] used in this study integrates both store-level and product-level information, enabling multi-dimensional analysis of sales performance. Table 1 summarizes the key attributes.This dataset structure provides a comprehensive view of both product demand and store characteristics, supporting the development of robust predictive models.

4.2. Exploratory Data Analysis

A comprehensive exploratory data analysis (EDA) was performed to uncover key relationships and sales drivers, providing foundational business intelligence for subsequent modeling steps.1 The analysis began with visual exploration of the data through plots and histograms, as recommended for retail analytics. [9].**Impact of Price on Sales:** A significant correlation was observed between product price (Product_MRP) and sales revenue. As shown in Figure 1, products priced in the low-to-medium range achieved higher sales volumes, while very high MRP products exhibited diminishing returns. This finding highlights the role of price sensitivity in customer purchasing behavior.**Product Category Contributions:** Product categories such as Snack Foods, Dairy, and Household Items were identified as the strongest contributors to revenue. Categories like Seafood and Starchy Foods exhibited higher variability in performance, influenced by store type and regional preferences (Figure 2).Store-related characteristics revealed further insights. Larger stores consistently achieved higher sales, reflecting greater product variety and customer traffic. Tier 1 locations exhibited stronger purchasing power, followed by Tier 2 and Tier 3 cities. Interestingly, store age was not a significant determinant of revenue, as newer outlets performed comparably to older ones, suggesting that demand stabilizes quickly after opening. Additionally, shelf space allocation was found to be critical, with products receiving greater allocated display area achieving significantly higher sales, emphasizing the importance of planogram optimization.

4.3. Key Insights from EDA

The analysis highlighted several important drivers that shaped feature engineering [11][12] and model development. On the product side, MRP, Product Type, and Allocated Area emerged as strong predictors of sales. On the store side, Store Size and City Tier were shown to be critical determinants of revenue. Finally, category-level dynamics revealed that high-performing product groups should be prioritized in inventory planning, while underperforming categories may require promotional strategies to improve turnover. Collectively, these insights ensured that the forecasting framework was grounded in the most relevant business drivers, improving both predictive performance and practical applicability.

5. Methodology

The forecasting framework was developed following a structured data science lifecycle, ensuring reproducibility, scalability, and integration with business workflows. The methodology consisted of four key phases: data preprocessing, feature engineering, model building, and model evaluation.The overall architecture of the proposed sales forecasting framework [13] is illustrated in Figure 3, which depicts the sequential flow from raw data ingestion to business decision-making

5.1. Data Preprocessing

To ensure data quality and prepare the dataset for modeling, several preprocessing steps were applied. Missing values were handled by imputing Product_Weight using the median grouped by product category, while categorical attributes such as Store_Type and Product_Sugar_Content were filled using mode or labeled as “Unknown.” Outlier detection and treatment were conducted using boxplots and histograms to identify extreme values in Product_MRP and Product_Allocated_Area; these outliers were capped at the 95th percentile to reduce skewness without discarding important observations. For encoding categorical variables, ordinal features such as Store_Size were processed using Label Encoding, while nominal features including Product_Type and Store_Type were transformed using One-Hot Encoding. Continuous variables (Product_MRP, Product_Allocated_Area, and Product_Weight) were standardized using StandardScaler to normalize ranges and stabilize model training [14]. The complete preprocessing workflow is illustrated in Figure 4, which shows how numerical and categorical transformations were integrated.

5.2. Feature Engineering

To enhance model performance and capture additional business-relevant patterns, new features were engineered. Store_Age was calculated as the difference between the current year and the store’s establishment year to represent store maturity. Is_Tier1_City was introduced as a binary variable to distinguish Tier 1 locations from others, while Normalized_Product_Area applied Min-Max scaling to standardize display space allocation. MRP_Bins grouped product prices into Low, Medium, and High to reflect non-linear pricing effects, and Product_Category_Code provided an encoded representation of product types for efficient model input. These engineered features are summarized in Table 2, which highlights their purpose and relevance in forecasting.

5.3. Model Building

The forecasting problem was formulated as a supervised regression task to predict quarterly sales revenue at the store–product level. Multiple algorithms were evaluated to ensure robustness and interpretability, including Decision Tree Regressor, Random Forest Regressor, AdaBoost Regressor, Gradient Boosting Regressor [15], and XGBoost Regressor [16]. The dataset was split into training and test sets using an 80/20 ratio, with stratified sampling applied to preserve balanced representation across store types and product categories. Model generalization was further assessed using five-fold cross-validation. For XGBoost, hyperparameter tuning was conducted using GridSearchCV, focusing on key parameters such as n_estimators (number of boosting rounds), max_depth (tree complexity), learning_rate (step size), and subsample and colsample_bytree (sampling ratios for robustness).

5.4. Model Evaluation

Model performance was assessed using standard regression metrics. The R^2 score measured the proportion of variance explained by the predictors, with higher values indicating better fit, though it can be misleading for time-series applications. Root Mean Squared Error (RMSE) quantified the average magnitude of squared errors in the same units as the target variable, offering interpretability but also sensitivity to outliers due to error squaring. Mean Absolute Error (MAE) provided a more robust alternative by measuring the average absolute deviation, making it less sensitive to extreme values. Both RMSE and MAE are scale-dependent, limiting their comparability across datasets. While these measures form a strong foundation for evaluation, more advanced, scale-independent metrics such as Mean Absolute Percentage Error (MAPE) and Weighted Mean Absolute Percentage Error (WMAPE) are often better suited for retail forecasting because they provide consistency across products with different sales volumes. In this study, RMSE and MAE were chosen for their interpretability and practicality, though future work may explore MAPE and WMAPE to further enhance evaluation robustness.

6. Results and Evaluation

The forecasting framework was evaluated using multiple regression algorithms, including Decision Tree, Random Forest, Gradient Boosting, AdaBoost, and XGBoost. Model performance on the test dataset was assessed using R^2, RMSE, and MAE as evaluation metrics [17].

6.1 Comparative Model Performance

The comparative results are summarized in Table 3. Among the models, XGBoost achieved the highest performance, explaining 85.2% of the variance in sales revenue while maintaining the lowest RMSE and MAE. Gradient Boosting also performed well but slightly underperformed compared to XGBoost, while Random Forest achieved solid accuracy though with marginally higher error. In contrast, Decision Tree and AdaBoost lagged in both accuracy and stability, confirming their limited suitability for structured retail datasets.

6.2. Impact of Hyperparameter Tuning

The performance of XGBoost improved notably after hyperparameter optimization, particularly tuning the number of estimators, maximum depth, and column sampling ratio. These adjustments increased R^2 while reducing both RMSE and MAE, demonstrating the value of careful model tuning in enhancing predictive power for retail forecasting.

6.3. Feature Importance Analysis

The tuned XGBoost model was further examined to identify the most influential predictors of sales. As shown in Figure 5, Product MRP emerged as the strongest predictor, confirming customer sensitivity to price. Store Size and City Tier were also highly influential, emphasizing the role of store-level characteristics in shaping sales outcomes. Product Type and Allocated Area contributed significantly, underscoring the impact of product mix and planogram optimization. Store Age, though less influential, still had a measurable effect. These findings validated earlier insights from exploratory data analysis (EDA) and offered actionable guidance for inventory allocation, pricing, and marketing strategies.

6.4. Model Generalization

The difference between training and test results was modest, indicating that the tuned XGBoost model generalized effectively to unseen data. This level of stability is particularly valuable in retail environments, where sudden demand shifts and product mix variations are common.

7. Low-Code Deployment Framework

While model development is critical for predictive accuracy, the real value of a forecasting system is realized only when it can be operationalized at scale and made accessible to business stakeholders. To this end, the sales forecasting framework for SuperKart was deployed using a low-code approach, enabling seamless interaction between machine learning outputs and decision-making processes.

7.1. Deployment Objectives

The deployment architecture was designed with four primary objectives: accessibility, scalability, integration, and maintainability. Accessibility was achieved by providing business managers with an intuitive interface for generating forecasts, while scalability ensured that predictions could be extended across multiple stores and product categories without major re-engineering. Integration was a key goal to enable seamless compatibility with existing dashboards and supply chain systems, and maintainability was addressed by adopting modular, low-code components that minimize IT overhead.

7.2. Frontend Implementation

The frontend interface was implemented using Streamlit and Gradio to create a lightweight web application accessible to business users. Through this interface, managers could input product attributes such as Product MRP, Product Type, and Allocated Area, along with store characteristics including Store Size and City Tier. The system generated real-time forecasts of sales revenue accompanied by confidence scores, with results displayed in both tabular and graphical formats for ease of interpretation. This low-code design eliminated the need for technical expertise, empowering store managers to independently utilize the forecasting system.

7.3. Backend Integration

At the backend, the trained XGBoost model and preprocessing pipeline were hosted on Hugging Face Spaces. This provided a containerized deployment environment with REST APIs for communication with external applications and dashboards, as well as built-in version control to ensure reproducibility and streamlined model updates. The backend accepted structured JSON inputs, processed them through the pipeline, and returned predictions in real time, ensuring both reliability and efficiency in operational use.

7.4. Deployment Architecture

The complete deployment workflow is illustrated in Figure 6.

The architecture consisted of three layers:

- **Frontend**: A Streamlit/Gradio user interface for forecast generation.
- **Backend**: HuggingFace Spaces hosting the XGBoost model and preprocessing pipeline.
- **Decision Systems**: Integration with SuperKart’s business dashboards and supply chain systems to support inventory planning and marketing strategies.

7.5. The MLOps Imperative

The deployment strategy for this project was guided by the principles of Machine Learning Operations (MLOps). MLOps aims to streamline the process of taking a model from research and development to a production environment, ensuring it is scalable, maintainable, and seamlessly integrated into business workflows. The goal is to move the model from a static state in a notebook to a dynamic, production-ready application that can deliver real-time insights at scale.

7.6. Benefits of Low-Code Deployment

The low-code and no-code paradigm is a growing trend that addresses the operational challenge of MLOps by providing tools that enable non-technical users to build and deploy applications with minimal to no coding. These platforms accelerate development cycles and empower business users, such as store managers and supply chain planners, to directly utilize AI-driven forecasts without relying on a specialized data science team. Examples of such platforms, which aim to democratize AI and reduce the total cost of ownership, include Akkio, Hubler, Dataiku, and DataRobot. These platforms provide a user-friendly interface for designing, developing, and deploying AI solutions, fostering collaboration between technical and non-technical staff and accelerating the speed-to-value of data initiatives.

7.6 Comparison to Enterprise Solutions

While the Streamlit/Hugging Face [18] approach provides a powerful and practical solution, it is important to situate it within the broader landscape of enterprise MLOps platforms. Platforms [19] such as Amazon SageMaker, Microsoft Azure, Dataiku, and DataRobot offer fully managed, end-to-end MLOps capabilities, including automated machine learning (AutoML), governance, and robust monitoring. These solutions are designed for large-scale, enterprise-level deployments and provide comprehensive features that a bespoke open-source solution might lack. However, the Streamlit/Hugging Face approach provides a flexible, low-cost alternative that is ideal for pilot projects, smaller organizations, or scenarios where a quick proof-of-concept is required. The decision to use this framework demonstrates a practical understanding of the trade-offs between comprehensive, high-cost enterprise solutions and a more agile, accessible open-source approach.

8. Discussion

The sales forecasting framework for SuperKart demonstrates the dual value of machine learning models: enhancing predictive accuracy while supporting strategic retail decision-making. The findings show implications for both managerial practice and academic research

8.1 Managerial Implications

At the operational level, strong predictive performance at the store–product level supports dynamic inventory optimization, enabling retailers to prioritize high-demand categories such as Snacks, Dairy, and Household while reducing overstock in weaker categories like Frozen Foods and Seafood. City Tier emerged as a key driver, confirming the stronger purchasing power of Tier 1 markets and justifying targeted promotions and hyperlocal marketing campaigns. Product Type and Allocated Area were also shown to influence sales outcomes, offering actionable insights for planogram design and product mix rationalization. In addition, the consistent superiority of larger stores suggests that expansion strategies should favor large-format outlets in Tier 1 and select Tier 2 cities.

8.2. Academic Contributions

This study contributes in three ways. First, it bridges the gap between accuracy and usability by demonstrating how low-code deployment can operationalize machine learning, making forecasting tools accessible to business stakeholders. Second, it validates the effectiveness of ensemble methods for structured retail datasets, with XGBoost showing strong robustness, interpretability, and predictive performance. Finally, comparative evaluation confirms XGBoost’s superiority over traditional and alternative ensemble approaches, reinforcing its role as a preferred model for retail forecasting.

8.3 Limitations

The study also acknowledges several limitations. Quarterly data aggregation restricted the ability to capture short-term seasonality such as holiday effects. Benchmarking against advanced deep learning models like LSTMs or Transformers [20] was not conducted, which may further improve sequential forecasting performance. Moreover, the framework did not explicitly account for sudden external disruptions, such as supply chain shocks or macroeconomic downturns, which may reduce forecasting accuracy in volatile environments.

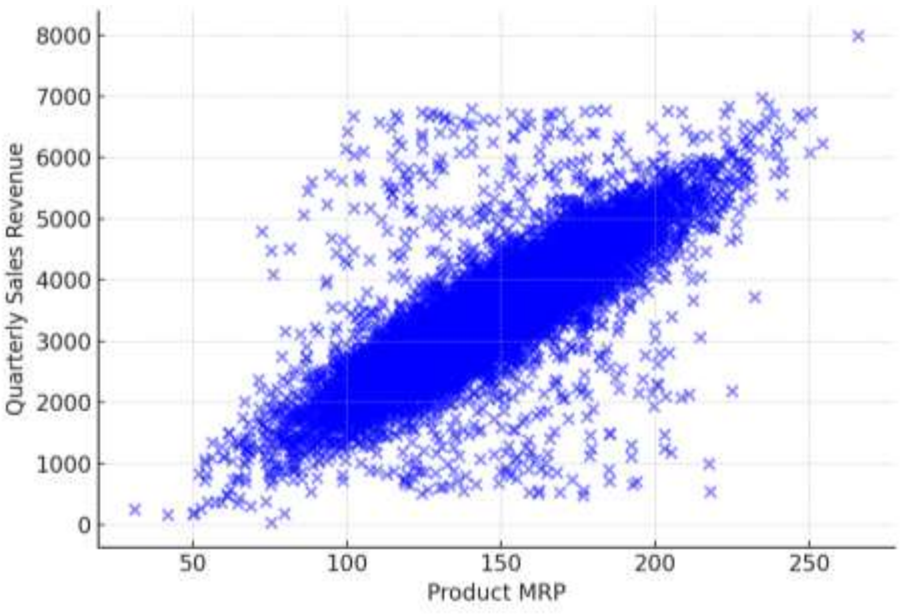

Figure 1. *Relationship between Product MRP and average sales revenue.*

Table 1: *Dataset attributes and descriptions*

Attribute	Description
Store_ID	Unique store identifier
Store_Size	Physical size of store (Small/Medium/High)
Store_Location_City_Type	City tier (Tier 1, Tier 2, Tier 3)
Store_Establishment_Year	Year the store was launched
Store_Type	Format (Supermarket, Departmental Store, Food Mart)
Product_ID	Unique product identifier
Product_Weight	Weight of the product
Product_Sugar_Content	Sugar level (Low, Regular, No Sugar)
Product_Allocated_Area	Shelf space allocated to the product
Product_Type	Product category (Snacks, Dairy, Household, etc.)
Product_MRP	Maximum Retail Price of the product
Product_Store_Sales_Total	Quarterly sales revenue (target variable)

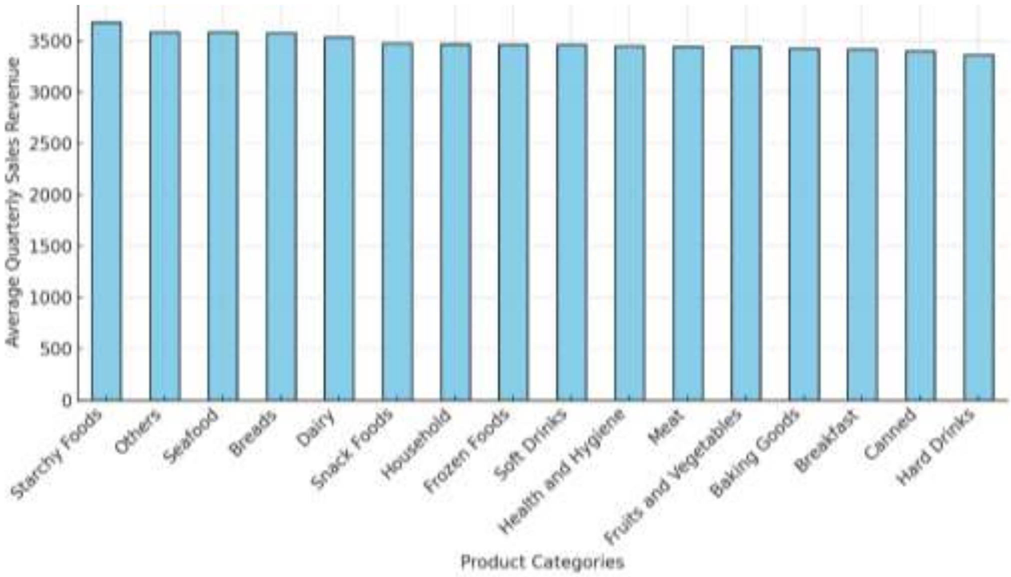

Figure 2. *Average sales across major product categories.*

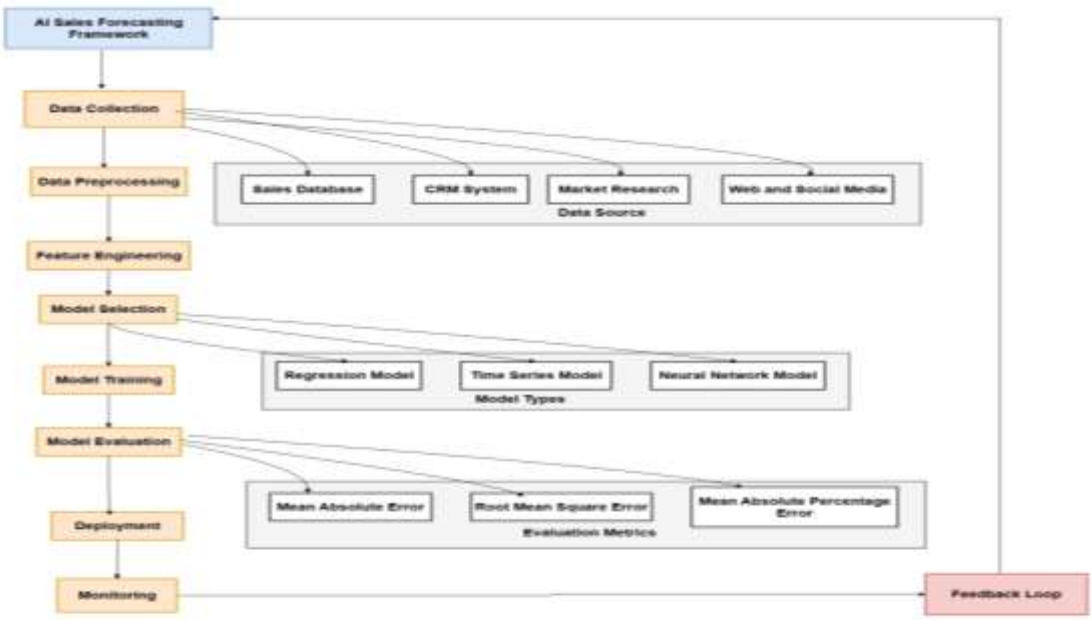

Figure 3 – *Sales Forecasting Framework*

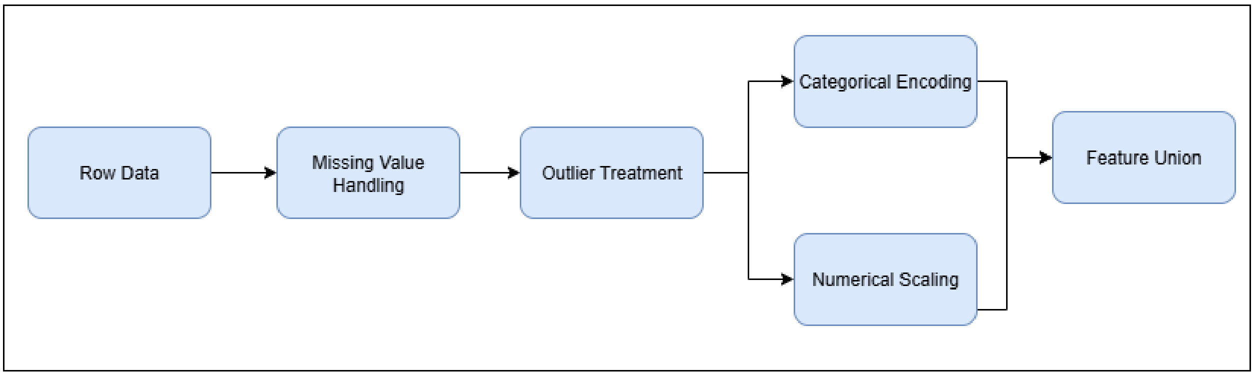

Figure 4. *Preprocessing pipeline combining numerical and categorical transformations.*

Table 2. *Derived features used in the forecasting framework.*

Feature Name	**Description**
Store_Age	Current Year – Establishment Year (maturity measure)
Is_Tier1_City	Binary flag (Tier 1 = 1, otherwise 0)
Normalized_Area	Normalized display space allocation
MRP_Bins	Categorization of MRP into Low/Medium/High groups
Product_Category_Code	Encoded representation of product categories

Table 3. *Model performance comparison on test data.*

Model	**R²**	**RMSE**	**MAE**
Decision Tree	0.605	1214.6	842.7
Random Forest	0.811	845.9	589.2
Gradient Boosting	0.829	810.4	573.1
AdaBoost	0.742	1002.5	704.6
XGBoost	0.852	755.8	542.3

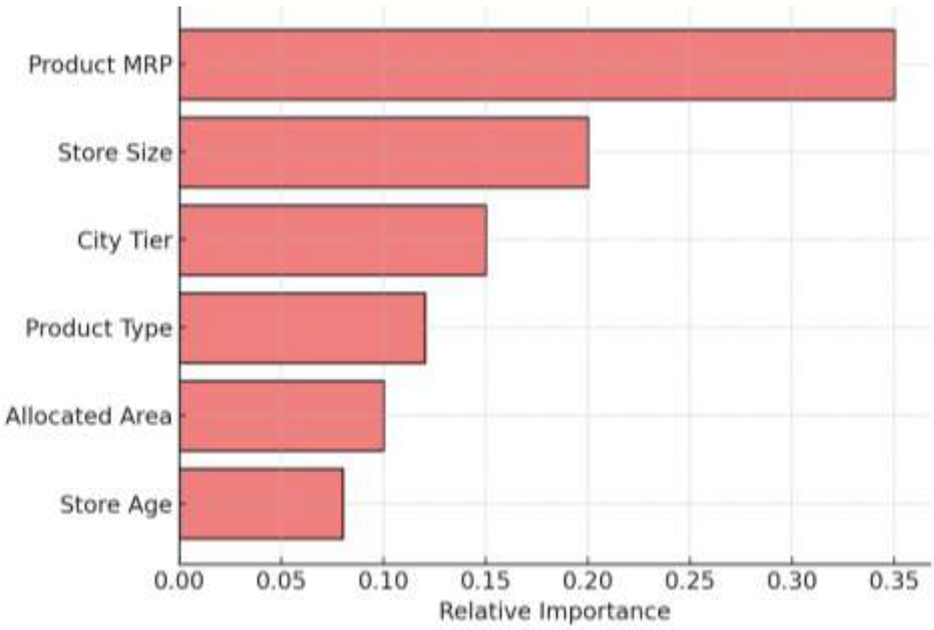

Figure 5. *Feature importance ranking from the tuned XGBoost model.*

4. Conclusions and Future Work

This study developed a comprehensive machine learning–based sales forecasting framework for SuperKart, integrating preprocessing, feature engineering, model evaluation, and low-code deployment. Among the models evaluated, XGBoost achieved the highest accuracy (R^2 =

0.852), outperforming Decision Tree, Random Forest, Gradient Boosting, and AdaBoost baselines. Feature importance analysis confirmed that Product MRP, Store Size, and City Tier were the most influential predictors, with Product Type and Allocated Area also playing key roles. These insights demonstrate the alignment of data-driven approaches with established retail management practices.

A notable contribution to this work was the operationalization of forecasting through a low-code deployment framework using Streamlit [21] and Hugging Face Spaces. This approach ensures accessibility for non-technical users such as store managers and planners, thereby bridging the gap between research and practice.

The research makes three primary contributions: (1) development of an end-to-end ML pipeline validated on real-world retail data; (2) empirical demonstration of tuned XGBoost as the most effective model for structured datasets; and (3) implementation of a scalable, low-code deployment framework that accelerates business adoption.

Future research could explore time-series deep learning architectures (e.g., LSTMs, Transformers) for improved seasonality capture, AutoML for automated optimization, and multi-objective forecasting to balance revenue with efficiency and resilience. Additional directions include incorporating external macroeconomic and competitive factors, enhancing robustness against disruptions, and advancing real-time streaming forecasts for adaptive decision-making.

In summary, by combining advanced machine learning with low-code deployment, this study offers both academic and practical value, enabling retailers to transform forecasting into a scalable, data-driven capability that strengthens inventory management, marketing, and growth strategies.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper

- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.

- **Author contributions:** The authors declare that they have equal right on this paper.

- **Funding information:** The authors declare that there is no funding to be acknowledged.

- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

[1] Forecasting Retail Sales using Machine Learning Models - ResearchGate, accessed September 13,2025:[Online].Available: https://www.researchgate.net/publication/390949992_Forecasting_Retail_Sales_using_Machine_Learning_Models

[2] J. Brownlee, Deep Learning for Time Series Forecasting: Predict the Future with MLPs, CNNs and LSTMs in Python, San Juan, PR, USA: Machine Learning Mastery, 2018. [Online]. Available: https://machinelearningmastery.com/deep-learning-for-time-series-forecasting

[3] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001. [Online]. Available: https://link.springer.com/article/10.1023/A:1010933404324

[4] Implementation of Extreme Gradient Boosting (XGBoost) for Retail Sales Prediction, accessed September 13, 2025. [Online]. Available : https://www.researchgate.net/publication/394485351_Implementation_of_Extreme_Gradient_Boosting_XGBoost_for_Retail_Sales_Prediction

[5] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD), San Francisco, CA, USA, 2016, pp. 785–794. [Online]. Available: https://doi.org/10.1145/2939672.2939785

[6] Deep learning for time series forecasting: a survey, accessed September 13, 2025. Online]. Available: https://www.researchgate.net/publication/388826200_Deep_learning_for_time_series_forecasting_a_survey

[7] R. J. Hyndman and G. Athanasopoulos, Forecasting: Principles and Practice, 3rd ed. Melbourne, Australia: OTexts, 2021. [Online]. Available: https://otexts.com/fpp3

[8] Data Driven Retail Price Optimization Using XGBoost and Predictive Modeling | Request PDF - ResearchGate, accessed September 13, 2025, https://www.researchgate.net/publication/389356193_Data_Driven_Retail_Price_Optimization_Using_XGBoost_and_Predictive_Modeling

[9] Unlocking Retail Insights: Predictive Modeling and Customer - MDPI, accessed September 13, 2025. [Online]. Available: https://www.mdpi.com/0718-1876/20/2/59

[10] Super kart data set file. [Online] Available: https://github.com/padmanabhamgit/SuperKart/blob/main/SuperKart.csv

[11] Advanced Feature Engineering for Classification - Kaggle, accessed September 13, 2025. [Online].

Available: https://www.kaggle.com/code/azminetoushikwasi/advanced-feature-engineering-for-classification
[12] What is Feature Engineering? - Machine Learning - GeeksforGeeks, accessed September 13, 2025. [Online]. Available: https://www.geeksforgeeks.org/machine-learning/what-is-feature-engineering/
[13] Sales Forecasting Based on Transformer-LSTM Model, accessed September 13, 2025. [Online]. Available: https://www.researchgate.net/publication/379520643_Sales_Forecasting_Based_on_Transformer-LSTM_Model
[14] Best Practices for Time Series Forecasting (Tutorial Paper), accessed September 13, 2025, [Online] Available: https://www.researchgate.net/publication/332079043_Best_Practices_for_Time_Series_Forecasting_Tutorial_Paper
[15] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," Annals of Statistics, vol. 29, no. 5, pp. 1189–1232, 2001. [Online]. Available: https://projecteuclid.org/journals/annals-of-statistics/volume-29/issue-5/Greedy-function-approximation-A-gradient-boosting-machine/10.1214/aos/1013203451.full
[16] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD), San Francisco, CA, USA, 2016, pp. 785–794. [Online]. Available: https://doi.org/10.1145/2939672.2939785
[17] S. Raschka, V. Mirjalili, and J. Hearty, Python Machine Learning, 3rd ed. Birmingham, UK: Packt Publishing, 2019 — includes evaluation metrics (R^2, RMSE, MAE) and hyperparameter tuning practices. [Online]. Available: https://www.packtpub.com/product/python-machine-learning-third-edition/9781789955750
[18] Samalba/hf-model-ops: MLOps with Hugging Face Spaces and Dagger - GitHub, accessed September 13, 2025. [Online] Available: https://github.com/samalba/hf-model-ops
[19] Best Data Science and Machine Learning Platforms Reviews 2025 ..., accessed September 13, 2025. [Online] Available: https://www.gartner.com/reviews/market/data-science-and-machine-learning-platforms
[20] F. Chollet, Deep Learning with Python, 2nd ed. Shelter Island, NY, USA: Manning, 2021. [Online]. Available: https://www.manning.com/books/deep-learning-with-python-second-edition
[21] Streamlit Spaces - Hugging Face, accessed September 13, 2025. [Online] Available: https://huggingface.co/docs/hub/spaces-sdks-streamlit