iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VI **Month of publication:** June 2026

DOI: https://doi.org/10.22214/ijraset.2026.83297

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

Time-Series and Machine Learning Integration for Walmart Sales Forecasting

Lokesh Reddy Ujjelli[1], Dr. K. Santhi Sree[2]
[1]*Post Graduate Student, M. Tech (Data Science), Department of Computer Science and Engineering, Jawaharlal Nehru Technological University, Hyderabad, India*
[2]*Professor, Department of Computer Science and Engineering, Jawaharlal Nehru Technological University, Hyderabad, India*

Abstract: ***Accurate sales forecasting is essential for retail planning, inventory optimization, and supply chain management. Traditional statistical models often struggle with nonlinear patterns and complex seasonality, while standalone machine learning approaches may ignore inherent temporal structures. This study proposes a hybrid weekly retail sales forecasting framework that integrates Facebook Prophet and XGBoost. Prophet is first employed to model trend, seasonality, and holiday effects, producing baseline forecasts and residual errors. An XGBoost regression model is then used to capture nonlinear residual components through engineered temporal features, including lag variables, rolling statistics, growth rates, and holiday indicators. The final prediction is obtained by combining Prophet forecasts with XGBoost-corrected residuals. Model performance is evaluated using MAE, RMSE, and MAPE metrics. Experimental results demonstrate that the proposed hybrid Prophet–XGBoost model significantly improves forecasting accuracy compared to the standalone Prophet approach, providing a scalable, interpretable, and data-driven solution for retail sales prediction.***

Keywords: Sales Forecasting, Time Series Analysis, Hybrid Model, Prophet, XGBoost, Residual Learning, Retail Analytics, Inventory Optimization, Feature Engineering, Performance Metrics.

I. INTRODUCTION

Accurate sales forecasting plays a vital role in retail business strategy, directly impacting inventory optimization, supply chain efficiency, workforce planning, and financial performance. In large retail environments, weekly sales data are influenced by multiple dynamic factors, including seasonality, promotional campaigns, holiday effects, and market fluctuations. Inaccurate forecasts may lead to overstocking, stock shortages, increased operational costs, and reduced profitability. Retail sales time series typically exhibit nonlinear trends, recurring seasonal patterns, and sudden demand spikes. Traditional statistical approaches often struggle to model such complexities effectively, while standalone machine learning models may overlook inherent temporal structures such as long-term trends and seasonal components. Therefore, a hybrid modelling approach that integrates statistical time-series analysis with machine learning techniques is essential.

Recent advancements in artificial intelligence have enabled the development of hybrid forecasting frameworks. Time-series models such as Prophet are capable of modelling trend, seasonality, and holiday effects, while machine learning algorithms like XGBoost effectively capture nonlinear residual patterns through feature engineering. This research proposes a multi-stage hybrid forecasting framework that integrates time-series decomposition, residual modelling, engineered temporal features, and machine learning refinement into a unified predictive system. By combining baseline forecasts generated by Prophet with residual correction using XGBoost, the proposed approach enhances forecasting accuracy, robustness, and interpretability in weekly retail sales prediction.

A. Motivation of the study

Retail organizations operate in highly dynamic and competitive environments where accurate sales forecasting is essential for efficient inventory management, workforce planning, and strategic decision-making. Inaccurate forecasts may result in overstocking, stock shortages, increased operational costs, and revenue loss. Traditional time-series models such as Prophet effectively capture trend, seasonality, and holiday effects but may not adequately model nonlinear fluctuations in sales data. Conversely, machine learning models such as XGBoost can capture complex nonlinear relationships but often lack explicit temporal interpretability. Motivated by these limitations, this study proposes a hybrid forecasting framework that integrates Prophet and XGBoost to improve weekly retail sales prediction accuracy and forecasting reliability.

©IJRASET: All Rights are Reserved

B. Objectives

The primary objective of this study is to develop an accurate, scalable, and interpretable hybrid retail sales forecasting system. Specifically, the study aims to preprocess and aggregate weekly retail sales data, model temporal patterns using Prophet, predict residual errors using XGBoost with engineered temporal features, combine both models to generate hybrid forecasts, and evaluate forecasting performance using MAE, RMSE, and MAPE metrics. The proposed framework also generates future sales forecasts with confidence intervals to support data-driven retail decision-making.

C. Problem Statement

Despite the wide range of forecasting techniques available, achieving high predictive accuracy in retail sales forecasting remains challenging. Traditional statistical models often struggle to capture nonlinear fluctuations and irregular residual patterns in sales data. In contrast, standalone machine learning models may overlook essential temporal components such as seasonality and long-term trends.

Another key limitation is the lack of integration between baseline time-series forecasting and residual correction mechanisms. Many existing approaches rely on a single modelling paradigm, resulting in reduced performance when applied to real-world retail datasets characterized by seasonal peaks, holiday effects, and demand volatility.

Therefore, there is a need to develop a comprehensive forecasting framework that:

- Accurately models temporal patterns such as trend and seasonality
- Captures nonlinear residual behavior using advanced machine learning techniques
- Reduces overall prediction error through hybrid modelling
- Provides interpretable and reliable outputs for retail decision-making

The central problem addressed in this study is the development of an integrated hybrid forecasting system that combines Prophet-based time-series modelling with XGBoost-based residual learning to improve accuracy and robustness in weekly retail sales prediction.

II. LITERATURE REVIEW

Machine learning and hybrid forecasting techniques have been widely applied in retail sales prediction to address complex demand patterns and seasonal variations. Xu [1] proposed a machine-learning-based framework for Walmart sales forecasting using Decision Trees, Random Forests, and Neural Networks. The study achieved better accuracy than traditional statistical methods; however, it relied only on standalone machine learning models without incorporating time-series decomposition or residual correction. Hybrid forecasting approaches have shown improved predictive capability by combining statistical and deep learning models. Wang and Jiang [2] developed an ARIMA-BIGRU hybrid framework in which ARIMA captured linear trends while BIGRU learned nonlinear residual patterns. Although the approach improved forecasting accuracy, it was mainly designed for stock forecasting and required high computational resources.

Ram et al. [3] enhanced SARIMAX forecasting through automated seasonal interval optimization, improving seasonal trend modelling. However, the method depended entirely on statistical assumptions and lacked nonlinear residual learning. Similarly, Brykin [4] compared ARIMA, LSTM, and Prophet models and concluded that forecasting performance depends on data characteristics, but no hybrid integration was explored.

Mahin et al. [5] applied ensemble regression methods for supply-chain forecasting and demonstrated improved predictive stability. Nevertheless, the study did not explicitly model time-series components or incorporate residual learning strategies. Overall, existing research mainly focuses on either standalone statistical models or independent machine learning approaches. Limited work integrates structured time-series decomposition with gradient-boosting-based residual correction in a unified retail forecasting framework. The proposed work addresses this gap by combining Prophet-based forecasting with XGBoost residual learning to improve accuracy and interpretability in weekly retail sales prediction.

III. METHODOLOGY

The proposed sales forecasting framework is developed as a hybrid multi-stage pipeline that combines statistical time-series modelling with machine-learning-based residual correction to improve forecasting accuracy. Initially, the Walmart weekly sales dataset undergoes preprocessing, including handling missing values, formatting date features, and performing exploratory data analysis to identify trends and seasonal patterns. Additional features such as lag values and rolling averages are generated to strengthen the forecasting process.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

In the first stage, the Prophet model is used to capture trend and seasonality components of weekly sales data. The residual errors produced by Prophet are then extracted and modelled using XGBoost to learn nonlinear demand variations that remain unaddressed by the baseline model. Finally, the Prophet forecast and XGBoost residual predictions are combined to generate the final hybrid forecast. The model performance is evaluated using metrics such as MAE, RMSE, and MAPE to compare forecasting accuracy and robustness.

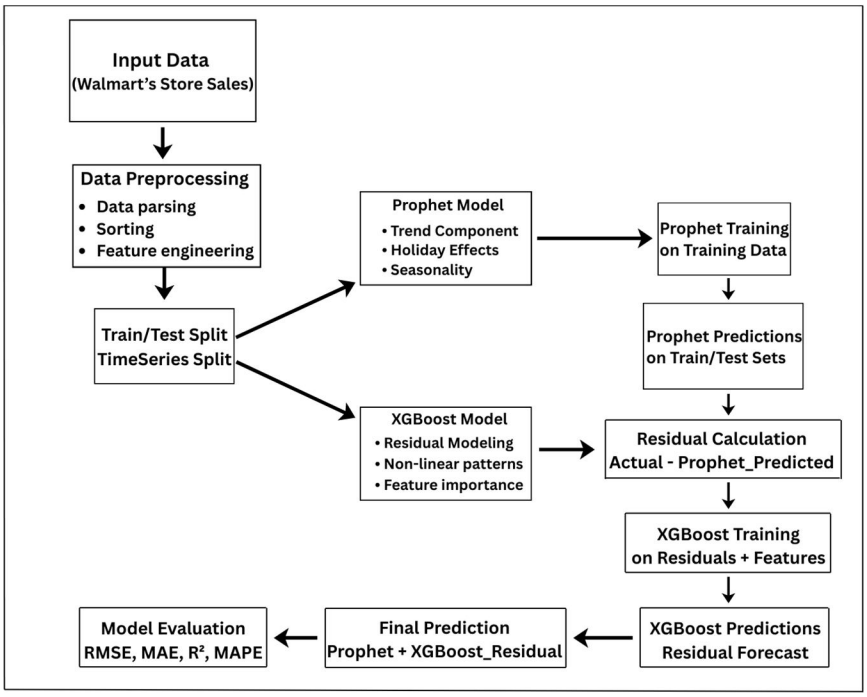


Fig. 1 System Architecture

A. Dataset Description and Preprocessing

The experimental evaluation was conducted using the Walmart weekly retail sales dataset containing 6,435 records collected from 45 stores between 05 February 2010 and 26 October 2012. The dataset includes weekly sales information along with economic indicators such as fuel price, consumer price index (CPI), unemployment rate, and holiday flags, which influence retail demand behavior.

Prior to modelling, the dataset was preprocessed to ensure consistency and compatibility with time-series forecasting techniques. The preprocessing steps included:

- Conversion of date fields into date-time format
- Chronological sorting of records
- Aggregation of sales data into weekly intervals
- Handling missing values using forward filling or interpolation
- Removal of extreme outliers where necessary

After preprocessing and aggregation, the final dataset characteristics are summarized below:

- Total weekly observations: 143
- Date range: 2010-02-05 to 2012-10-26
- Training samples: 122 weeks (85.31%)
- Testing samples: 21 weeks (14.69%)
- Training period: 2010-02-05 to 2012-06-01
- Testing period: 2012-06-08 to 2012-10-26

A chronological train-test split was adopted to avoid data leakage and simulate realistic forecasting conditions. The final processed dataset was indexed as a univariate time-series for forecasting model implementation.

B. Train-Test Splitting Strategy

Since this is time-series data, a chronological split is performed instead of random sampling.

- The first 80–85% of observations are used for training.
- The remaining 15–20% are reserved for testing.

This approach ensures that future values are never used to predict past values, thereby maintaining realistic forecasting conditions.

C. Prophet-Based Time-Series Modelling

Prophet is employed in the first stage to model structured components of the sales time series.

The additive model formulation is given by:

$$y(t) = g(t) + s(t) + h(t) + \varepsilon_t \quad (1)$$

Where:

- g(t) represents the trend
- s(t) represents seasonality
- h(t) represents holiday effects
- ε_t represents residual error

Prophet automatically detects change-points in trend and models weekly and yearly seasonality. It generates baseline forecasts along with confidence intervals.

This stage captures long-term growth and recurring seasonal fluctuations effectively. However, nonlinear irregular variations may remain ε_t, which motivates the use of residual correction.

D. Residual Error Computation

After generating predictions from Prophet, residual errors are calculated as:

$$Residual_t = Actual_t - Prophet_t \quad (2)$$

These residuals represent unexplained nonlinear variations caused by sudden demand shifts, promotional campaigns, or external market factors.

Modelling residual errors allows the system to refine baseline forecasts and improve overall accuracy.

E. Residual Learning Using XGBoost

To improve forecasting accuracy, residual errors obtained from the Prophet model were further modelled using XGBoost, a gradient boosting framework optimized for regression tasks. Additional time-dependent features were engineered to capture short-term dependencies and nonlinear sales patterns not explicitly learned by Prophet.

The engineered features include:

- Lag features (Sales at t_1, t_2, t_3, t_4)
- Rolling mean (moving average of previous weeks)
- Rolling standard deviation
- Sales growth rate
- Holiday indicator variables
- Week number and month index

Before model training, all features were normalized using StandardScaler to ensure balanced feature influence. XGBoost iteratively constructs an ensemble of decision trees to minimize prediction error while reducing overfitting through regularization.

The model optimizes an objective function consisting of:

- Loss function (Mean Squared Error)
- Regularization term (to reduce overfitting)

The residual learning function can be expressed as:

$$\hat{r_t} = f(X_t) \quad (3)$$

Where:

- $\hat{r_t}$ is predicted residual
- X_t represents engineered feature vector

F. Hybrid Forecast Generation

The final hybrid forecast is obtained by combining baseline Prophet prediction and predicted residual correction:

$$Final\ Forecast_t = Prophet_t + \hat{r_t} \quad (4)$$

This hybrid formulation integrates structured time-series decomposition and nonlinear correction, resulting in improved predictive performance.

G. Performance Evaluation Metrics

The performance of the model is evaluated using the following regression metrics:

TABLE I
PERFORMANCE METRICS

1	Mean Absolute Error (MAE)	$\frac{1}{n}\sum\lvert y_i - \hat{y}_i\rvert$
2	Root Mean Squared Error (RMSE)	$\sqrt{\frac{1}{n}\sum(y_i - \hat{y}_i)^2}$
3	Mean Absolute Percentage Error (MAPE)	$\frac{100}{n}\sum\left\lvert\frac{y_i - \hat{y}_i}{y_i}\right\rvert$

These metrics provide comprehensive evaluation of forecast accuracy and error magnitude.

IV. RESULTS AND DISCUSSION

A. Prophet Model Performance

The Prophet model was first trained to capture:

- Long-term trend
- Weekly seasonality
- Yearly seasonality
- Holiday effects

The additive forecasting model is:

$$y(t) = g(t) + s(t) + h(t) + \epsilon_t \quad (5)$$

Prophet Performance (Test Set)

- MAE : 1,066,281
- RMSE : 1,456,249
- MAPE : 2.25%

Although Prophet effectively modelled seasonal patterns, residual analysis showed nonlinear fluctuations remaining during peak demand periods.

B. Hybrid Prophet–XGBoost Model Performance

To improve predictive accuracy, Residual errors were modelled using the XGBoost algorithm with engineered features.
The final hybrid prediction was computed as:

$$Final\ Prediction = Prophet\ Prediction + XGBoost\ Residual \quad (6)$$

Hybrid Model Results (Test Set)

- MAE : 664,917
- RMSE : 822,145
- MAPE : 1.41%

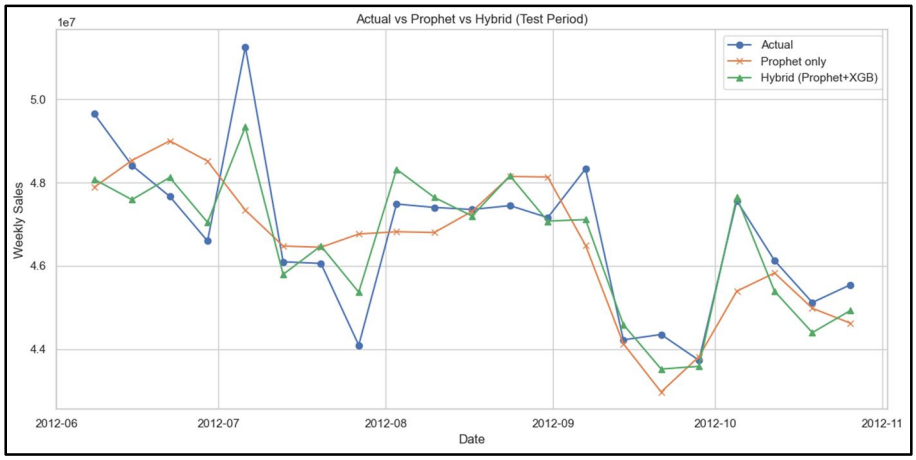


Fig. 2 Actual vs Prophet vs Hybrid (Test Period)

Figure illustrates the effectiveness of the proposed hybrid forecasting framework. While the Prophet model captures overall trend and seasonal patterns, minor deviations remain during high-demand periods. By modelling residual errors, the Prophet–XGBoost hybrid model more closely follows actual sales behavior and better adapts to nonlinear fluctuations, resulting in improved forecasting accuracy.

C. Performance Comparison

TABLE II
PERFORMANCE COMPARISON

Model	MAE	RMSE	MAPE
Prophet	10,66,281	14,56,249	2.25%
Hybrid (Prophet + XGBoost)	6,64,917	8,22,145	1.41%

Improvement Achieved

- MAE Improvement: 37.64%
- Significant reduction in RMSE
- Improved percentage forecasting accuracy

The hybrid model substantially outperformed the standalone Prophet model, confirming the effectiveness of residual learning.

D. Feature Importance Analysis

XGBoost feature importance analysis revealed:

- Sales growth rate had highest contribution
- Prophet baseline prediction was highly influential
- Rolling mean captured short-term smoothing
- Lag features modelled temporal dependencies
- Holiday indicators improved peak-week adjustments

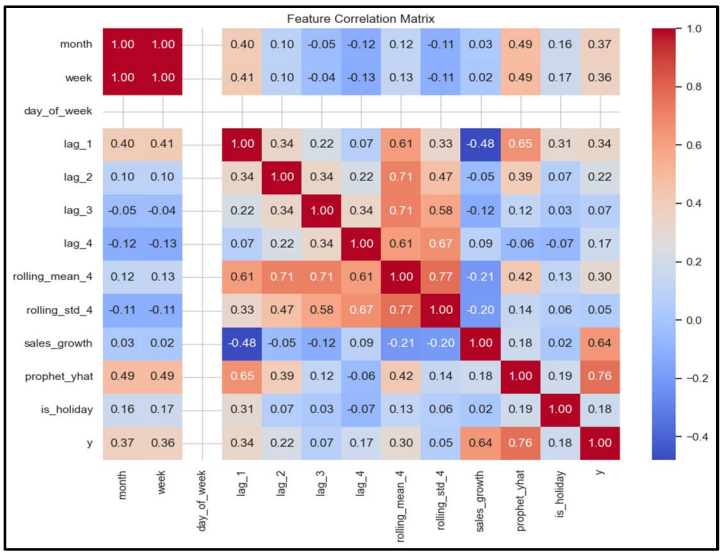


Fig. 3 Feature Correlation Matrix

The correlation analysis illustrated in Figure 3 provides insight into the relationships among engineered temporal features used for residual modelling. Lag variables and rolling statistical features exhibit strong positive correlations with weekly sales, indicating that historical sales behavior strongly influences future demand. Additionally, the Prophet baseline prediction shows high correlation with the target variable, validating its effectiveness in capturing structured time-series components. These observations justify the use of lag-based temporal features and rolling statistics within the XGBoost residual learning stage.

C. Residual Distribution Analysis

The residual distribution of the hybrid model showed:

- Reduced variance compared to Prophet
- Near-normal error distribution
- Absence of extreme systematic bias

This indicates improved model stability and robustness.

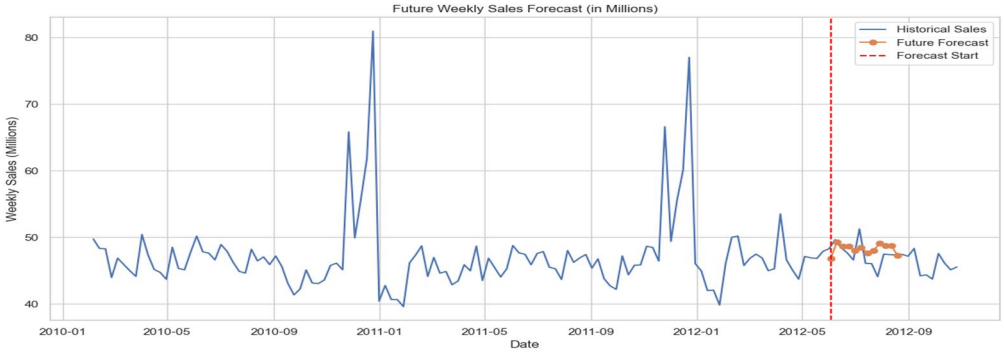


Fig. 4 Residual Distribution (Hybrid Model)

The residual distribution presented in Figure 4 provides further evidence of the robustness of the hybrid forecasting model. The error values are centered around zero and display an approximately symmetric distribution, indicating the absence of systematic prediction bias. Compared to the baseline Prophet model, the hybrid framework produces lower variance in residual errors, suggesting improved predictive stability. This behavior confirms that the integration of gradient boosting residual modelling effectively refines the baseline time-series forecasts.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

D. Future Sales Forecast

To evaluate the practical applicability of the proposed forecasting framework, the trained hybrid model was used to generate predictions for future weekly sales beyond the available historical data. Forecasting future demand is essential for retail planning, inventory optimization, and supply chain management. The predicted results provide insights into expected sales trends for upcoming weeks and help assess the model's ability to generalize beyond the training period.

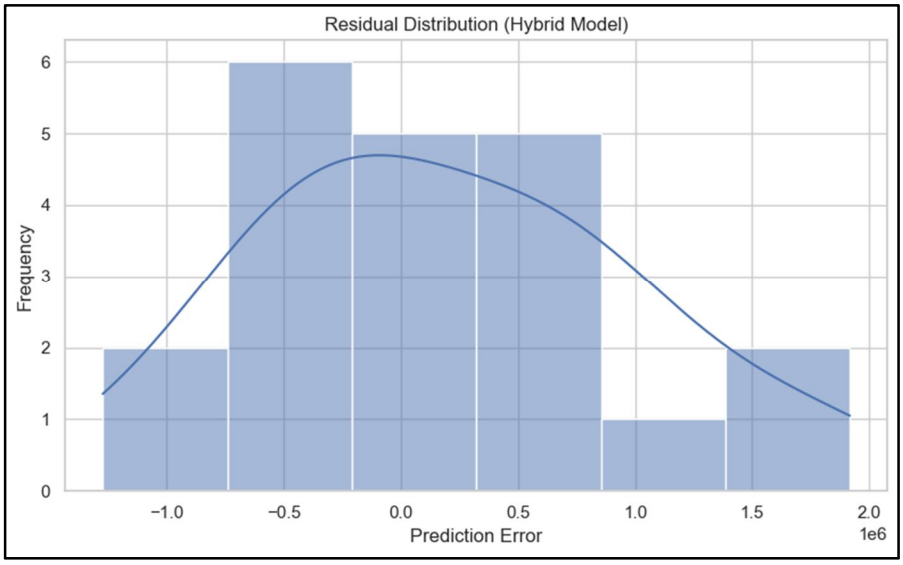


Fig. 5 Future Weekly Sales Forecast Generated by the Hybrid Model

The trained hybrid forecasting model was further used to generate predictions for future weekly sales beyond the available historical data. Figure 5 illustrates the predicted sales trajectory produced by the hybrid Prophet–XGBoost framework. The dashed vertical line indicates the beginning of the forecasting period, separating historical observations from predicted values.

The forecasting results show that the predicted sales follow the general pattern observed in the historical data while maintaining realistic variability. The hybrid model captures underlying temporal patterns and produces stable forecasts for upcoming weeks. Such predictions can assist retail managers in proactive decision-making related to inventory planning, supply chain coordination, and workforce allocation.

E. Practical Implications

The proposed hybrid forecasting framework has practical applications in retail analytics and operational planning. Accurate weekly sales forecasts enable retailers to optimize inventory allocation, reduce stock shortages, and minimize excess storage costs. In addition, improved demand prediction supports better supply chain coordination, workforce scheduling, and promotional planning. By combining interpretable time-series modelling with machine learning-based residual correction, the proposed system provides a scalable forecasting solution that can be applied to large retail environments and data-driven business decision systems.

V. CONCLUSION

The experimental evaluation shows that the hybrid Prophet–XGBoost framework significantly improves forecasting accuracy compared with the standalone Prophet model. The hybrid approach reduces MAE, RMSE, and MAPE errors by effectively combining structured time-series decomposition with nonlinear residual learning. This integrated forecasting strategy provides reliable predictions for weekly retail sales and offers valuable insights for inventory planning, supply chain optimization, and data-driven retail decision-making.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

VI. ACKNOWLEDGMENT

I express my sincere gratitude to my guide, Dr. K. Santhi Sree, Professor and Head, Department of Computer Science and Engineering, Jawaharlal Nehru Technological University Hyderabad, for her valuable guidance, continuous support, and constructive suggestions throughout the development of this research work. Her encouragement and technical insights greatly contributed to the successful completion of this project.

I also extend my heartfelt thanks to the faculty members of the Department of Computer Science and Engineering for their support and motivation during the course of this study.

REFERENCES

[1] Xu, R., A Method for Wal-Mart Sales Forecasting Based on Machine Learning, 2024 International Conference on Cloud Computing and Big Data (ICCBD 2024), 6 pages, 2024, ACM ISBN 979-8-4007-1022-3, https://dl.acm.org/doi/10.1145/3695080.3695141

[2] Tianyu Wang, Xiantao Jiang, ARIMA-BIGRU Stock Forecast Model Based on Bayesian Optimization, Proceedings of the 2025 3rd International Conference on Communication Networks and Machine Learning (CNML 2025), 314–319, 2025, ACM ISBN 979-8-4007-1323-1, https://dl.acm.org/doi/10.1145/3728199.3728219

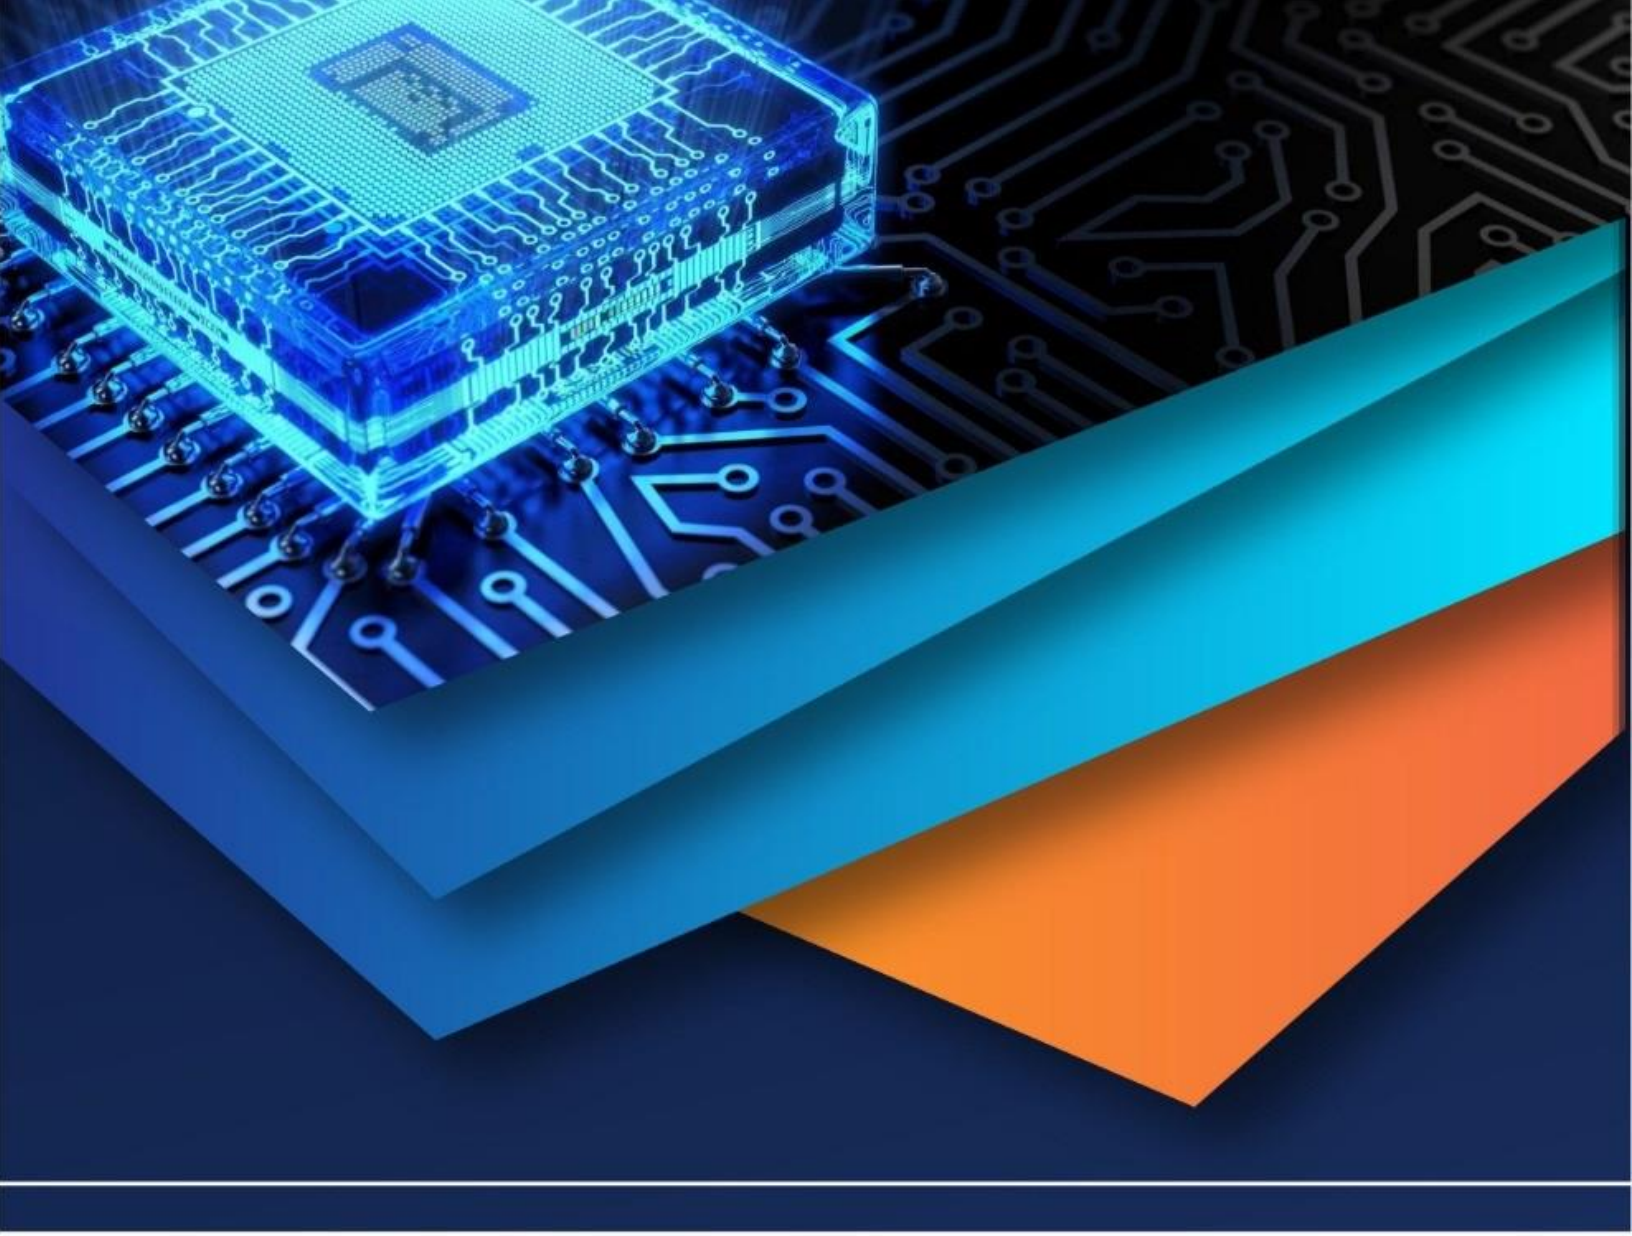
[3] Chandra Shekhar Ram, Manish Raja, Rajnish Kumar Chaturvedi, Boosting Time-Series Forecasting Accuracy with SARIMAX Seasonal Interval Automation, Procedia Computer Science, 814–821, Vol. 260, 2025, ISSN: 1877-0509, 10.1016/j.procs.2025.03.262/ https://dl.acm.org/doi/10.1016/j.procs.2025.03.262

[4] Dmitry Brykin, Sales Forecasting Models: Comparison between ARIMA, LSTM and Prophet, Journal of Computer Science, 1222-1230, Vol. 20, 2024, DOI: 10.3844/jcssp.2024.1222.1230/ https://doi.org/10.3844/jcssp.2024.1222.1230.

[5] Mahin, M.P.R., Shahriar, M., Das, R.R., Roy, A., Reza, A.W., 2025. Enhancing sustainable supply chain forecasting using machine learning for sales prediction. Procedia Computer Science 252, 470–479. https://doi.org/10.1016/j.procs.2025.01.006

[6] Kong, X., Chen, Z., Liu, W., Ning, K., Zhang, L., Marier, S. M., Liu, Y., Chen, Y., & Xia, F. (2025). Deep learning for time series forecasting: a survey. International Journal of Machine Learning and Cybernetics, 5079–5112, 16, 2025, https://doi.org/10.1007/s13042-025-02560-w

[7] Md Kamal Ahmed, Md Ekrim Hossin, Mohammad Muzahidur Rahman Bhuiyan, Sazzat Hossain, Fahmida Binte Khair, Shafaete Hossain, Mia Md Tofayel Gonee Manik, Forecasting Sales Trends Using Time Series Analysis: A Comparative Study Of Traditional And Machine Learning Models, Membrane Technology, 668-682, Vol: 2025, 2025, ISSN (online): 1873-4049, https://www.researchgate.net/publication/389411074



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)