Proceedings of the 8th IEOM Bangladesh International Conference on Industrial Engineering and Operations Management, Dhaka, Bangladesh, December 20-21, 2025

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/BA08.20250079
Published: December 20, 2025

A Hybrid Approach to Multivariate Time Series Forecasting: Integrating VARIMA and Machine Learning Models

Md. Sohanur Rahman
Industrial & Production Engineering Department
Rajshahi University of Engineering and Technology
Rajshahi, Bangladesh
rahman.ipe08@gmail.com

Shumon Kumar Ray
Industrial & Production Engineering Department
Shahjalal University of Science and Technology
Sylhet, Bangladesh
shumon.ipe@gmail.com

Md. Mahadi Hasan
Industrial & Production Engineering Department
Ahsanullah University of Science and Technology
Dhaka, Bangladesh
mahadi.sd@gmail.com

Shekh Rezwanul Ahmed
Applied Chemistry & Chemical Engineering Department
Dhaka University
Dhaka, Bangladesh
rezwanul.ahmed@gmail.com

Abstract

This study introduces a hybrid forecasting framework that integrates Vector Auto-Regressive Integrated Moving Average (VARIMA) with machine learning models to improve multivariate time series prediction in the electronics sector. The dataset comprises real-time sales records enriched with multiple explanatory variables, including Year, Month, Promotional Offers, Weather, Pricing, Sales Policies, and Eid/Festival impacts, all of which strongly influence consumer demand and introduce pronounced seasonality. The VARIMA component is employed to capture linear dependencies and seasonal trends, while machine learning algorithms are applied to the residuals to model nonlinear relationships and hidden interactions among variables. By combining the strengths of both statistical and data-driven approaches, the proposed hybrid method delivers superior forecasting accuracy compared to standalone models. The findings provide actionable insights into demand planning, promotional strategy, and policy adjustments, enabling firms to better align supply chain decisions with dynamic market conditions.

Keywords
Hybrid Forecasting, Multivariate Time Series, VARIMA, Machine Learning, Demand Prediction

1. Introduction
Accurate forecasting is essential for business, supply chain, and economic planning, helping organizations optimize inventory, production, and logistics. Multivariate time series, involving interdependent variables with both linear and nonlinear patterns, pose a challenge for traditional models. Linear models like VAR and VARIMA capture interdependencies well but struggle with nonlinearities, while machine learning models handle nonlinear patterns but often miss linear trends and lagged interactions.

This study addresses a research gap by proposing a hybrid forecasting framework that combines VARIMA with machine learning. VARIMA models the linear structure, and an ML model learns the nonlinear residuals, producing more accurate and robust forecasts. Using real-world electronics sector data, the hybrid approach outperforms classical linear models and standalone ML methods, offering a structured methodology for supply chain forecasting.

Objectives:

- Develop a hybrid VARIMA–ML forecasting framework.
- Evaluate its performance against traditional multivariate and standalone ML models.

2. Literature Review
Time series forecasting has evolved from classical statistical models such as VAR, VARMA, and VARIMA—which are interpretable and effective for linear interdependencies (Guo, Liu, & Sun, 2016)—to advanced machine learning (ML) and deep learning models like Random Forest, XGBoost, LSTM, and GRU, which capture non-linear and long-term dependencies (Wu et al., 2020; Li & Yang, 2024). Classical models struggle with non-linear patterns and high-dimensional data, while ML models require large datasets and can lack interpretability.

Hybrid forecasting models combine the strengths of both approaches. Early work by Zhang (2003) and Khashei & Bijari (2011) integrated ARIMA with neural networks for univariate series. Recent studies extend this to multivariate series, e.g., VARMA with Bayesian networks (Guo, Liu, & Sun, 2016), VAR-LSTM-GARCH (Chege, Kithinji, & Gachoki, 2025), VAR/ARMA with deep belief networks (Tuo Xie et al., 2018), and ARIMA combined with CNN, LSTM, GRU, or Prophet (2024–2025).

Overall, multivariate hybrid models leverage statistical rigor and ML flexibility, offering accurate, robust, and adaptable forecasts for complex, high-dimensional datasets, crucial for decision-making in finance, energy, supply chains, and network systems.

3. Methodology
3.1 Block Diagram of Proposed Hybrid Framework
The workflow starts with data description (source, variables, period, domain) and exploratory data analysis (EDA) to examine trends, seasonality, correlations, and outliers. Preprocessing includes missing value treatment, normalization, differencing, and train-test splitting, followed by stationarity analysis using ADF/KPSS tests and lag selection via AIC/BIC.

Next, the VARIMA model is specified, estimated, and validated through residual analysis and stability checks. A machine learning model (Random Forest, XGBoost, or LSTM) is then applied to capture nonlinear patterns in residuals. The hybrid forecast combines VARIMA predictions with ML-predicted residuals.

Finally, the forecast is evaluated using RMSE, MAE, MAPE, R^2, and the Diebold–Mariano test, with results analyzed to compare VARIMA, ML-only, and hybrid models, highlighting implications for sales forecasting, inventory, and demand planning (Figure 1- Figure 8).

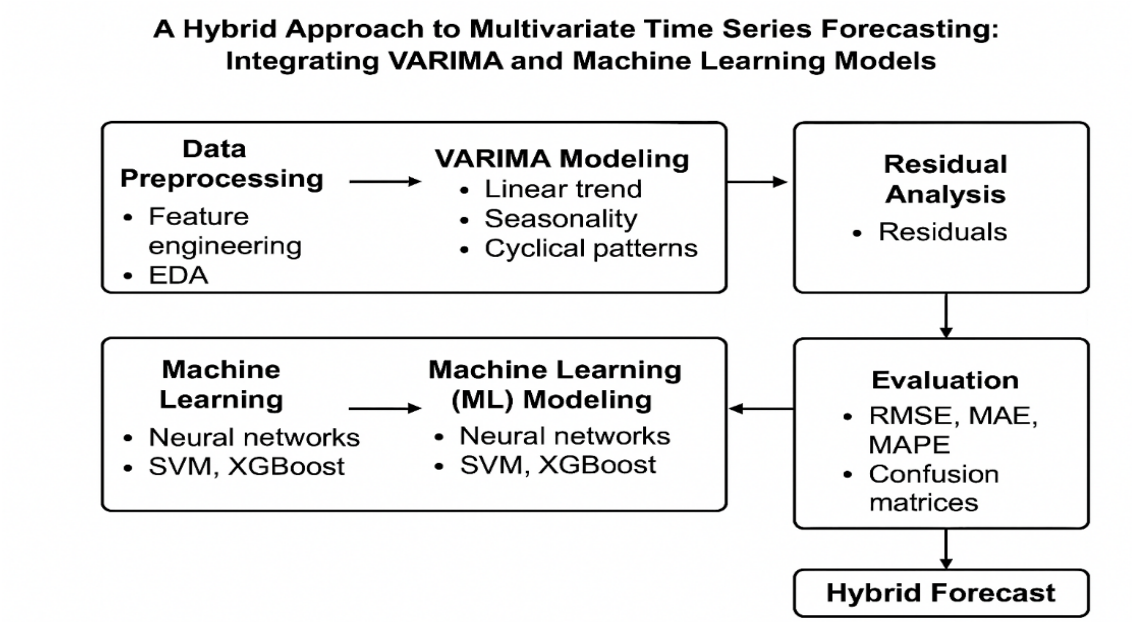

Figure 1. Block Diagram of Proposed Hybrid Framework

3.2 Mathematical Model Outline

i. VARIMA (Multivariate ARIMA)

General VARIMA(p,d,q) model for multivariate series Y_t

$$\Delta^d Y_t = \Phi_1 \Delta^d Y_{t-1} + \cdots + \Phi_p \Delta^d Y_{t-p} + \Theta_1 \varepsilon_{t-1} + \cdots + \Theta_q \varepsilon_{t-q} + \varepsilon_t$$

Y_t: Vector of time-series variables (e.g., Sales, Promotions, Weather, Festivals)
Φ_i: Autoregressive coefficient matrices
Θ_i: Moving average coefficient matrices
d: Differencing order to ensure stationarity
ε_t: White noise residual vector

Stepwise:

- Differencing to stationarity: $\Delta^d Y_t = Y_t - Y_{t-d}$
- Select p, qusing AIC/BIC
- Fit model via Maximum Likelihood Estimation
- Diagnostic checks: Residual autocorrelation (Ljung-Box), stability

ii. Machine Learning Component

Let $R_t = Y_t - \hat{Y}_t^{VARIMA}$be residuals from VARIMA. ML model (RF/XGBoost/LSTM) learns residuals:

$$\hat{R}_t = f(X_t; \theta)$$

X_t: Feature vector (original exogenous variables + lagged residuals)
f: Nonlinear mapping learned by ML model
θ: Hyperparameters tuned via GridSearch/Bayesian

iii. Hybrid Forecast

Additive residual approach

$$\hat{Y}_t^{Hybrid} = \hat{Y}_t^{VARIMA} + \hat{R}_t^{ML}$$

- Captures both linear patterns (VARIMA) and nonlinear patterns (ML)
- Equation can be iterated for multi-step forecasting

iv. Evaluation Metrics

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{t=1}^{n} (Y_t - \hat{Y}_t)^2}$$

$$\text{MAE} = \frac{1}{n}\sum_{t=1}^{n} | Y_t - \hat{Y}_t |$$

$$\text{MAPE} = \frac{100}{n}\sum_{t=1}^{n} \frac{| Y_t - \hat{Y}_t |}{Y_t}$$

$$R^2 = 1 - \frac{\sum_t (Y_t - \hat{Y}_t)^2}{\sum_t (Y_t - \bar{Y})^2}$$

- Diebold–Mariano test compares forecast accuracy of hybrid vs baseline

3.3 Dataset Description

The data set collected from Walton Hi Tech Industries PLC. encompasses monthly sales records spanning five years (January 2021 to December 2025), comprising 60 observations and six variables: Year, Month, Promotional Offers, Weather, Eid/Festival Impact, and Sales. This dataset integrates temporal dynamics with contextual drivers of sales performance, where Promotional Offers denote the intensity of marketing initiatives, Weather captures seasonal or climatic effects, and Eid/Festival Impact reflects the influence of cultural and religious events. Sales represent the total monthly revenue, providing a quantitative measure of business performance. The data exhibits pronounced seasonal and promotional patterns, with revenue peaks typically aligned with periods of high festival activity and strategic marketing campaigns. This rich, multidimensional dataset is particularly suitable for advanced time series forecasting, regression modeling, and causal analysis, offering valuable insights for optimizing marketing strategies, inventory management, and operational planning in dynamic retail environments (Table 1).

Table 1. Data-Set Descriptions

Variable	Description
Year	Calendar year of observation.
Month	Month of transaction, used to capture monthly seasonal effects.
Sales	Target variable represents total sales volume.
Promotional Offers	Continuous indicator of discounts or campaign activities.
Weather	Weather conditions or temperature influencing footfall and purchasing behavior.
Pricing	Product pricing dynamics, including price changes or average selling price.
Sales Policies	Company-driven sales strategies such as bundles, financing, or dealer incentives.
Eid/Festival Impact	Dummy variable representing major festival periods that cause demand spikes.

3.4 Data Preprocessing

For data preprocessing, the first step involves handling dates by combining the Year and Month columns to create a unified Date column in the YYYY-MM format, which is then set as the time index for the time series analysis. Next, categorical variables such as Promotional Offers, Weather, and Eid/Festival Impact are converted into numeric representations using either label encoding or one-hot encoding to make them suitable for machine learning algorithms. In terms of missing data, the dataset should be checked for any null or missing values, which can be imputed using techniques like interpolation or forward-fill to maintain the continuity of the time series. Finally, scaling is applied optionally, particularly for machine learning models such as Random Forests or Neural Networks, where Min-Max scaling or Standardization ensures proper model convergence and improves predictive performance.

3.5 Exploratory Data Analysis (EDA)

Exploratory analysis of Sales Data (2021–2025) reveals clear trends, seasonality, and cyclic patterns. Sales peak during high promotions, favorable weather, and Eid/festival periods (March–June), while the lowest sales occur in November–January. Yearly cycles show recurring rises during festivals and drop in winter, indicating that sales are strongly influenced by promotional strategies and seasonal factors.

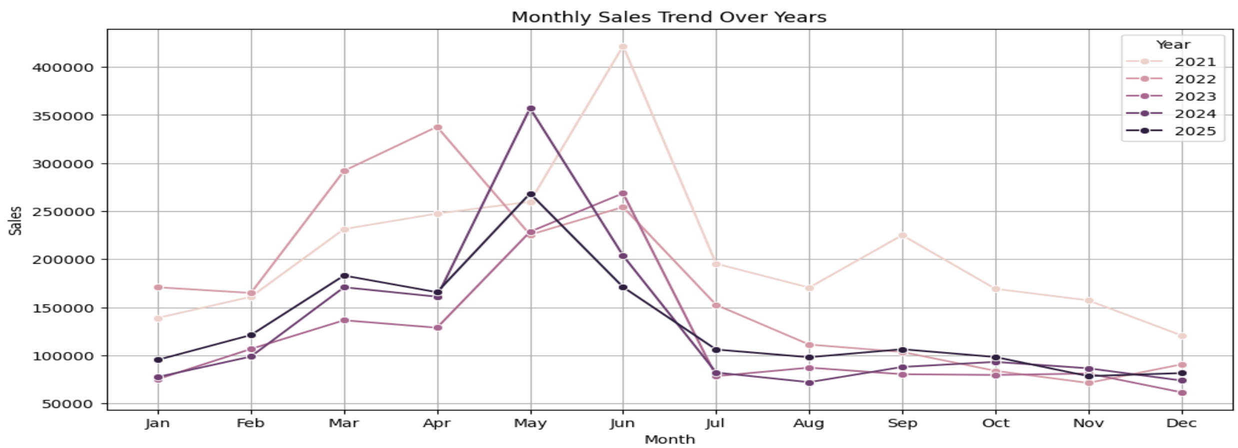

Figure 2. Month wise Sales Trend

The correlation matrix shows Sales positively associated with Eid/Festival Impact (0.47), highlighting festive periods' influence. Sales have negative correlations with Year (-0.35) and Month Num (-0.36), suggesting slight declines over time. Promotional Offers and Weather are strongly correlated (0.55) but weakly linked to Sales. Other correlations, including Month and Sales (0.027), are near zero, indicating minimal linear relationships among these variables.

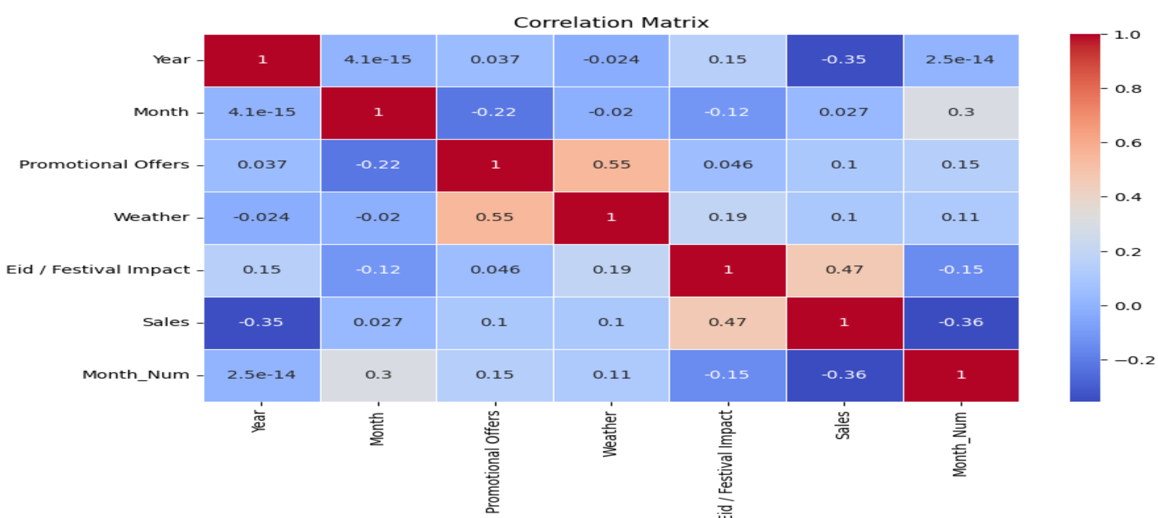

Figure 3. Co-relation Matrix between the Drivers

Temporal analysis of sales shows clear seasonality, with peaks from March to June during high promotions and major festivals, and lows in January, February, November, and December. Promotional Offers and Eid/Festival Impact are the primary drivers, while weather has minimal effect. Correlation analysis confirms strong positive relationships between sales and festivals/promotions. These insights suggest forecasting, inventory, and promotional planning should focus on festival periods and high-promotion months to optimize performance.

Figure 4. Sales Distribution by Promotional Offers

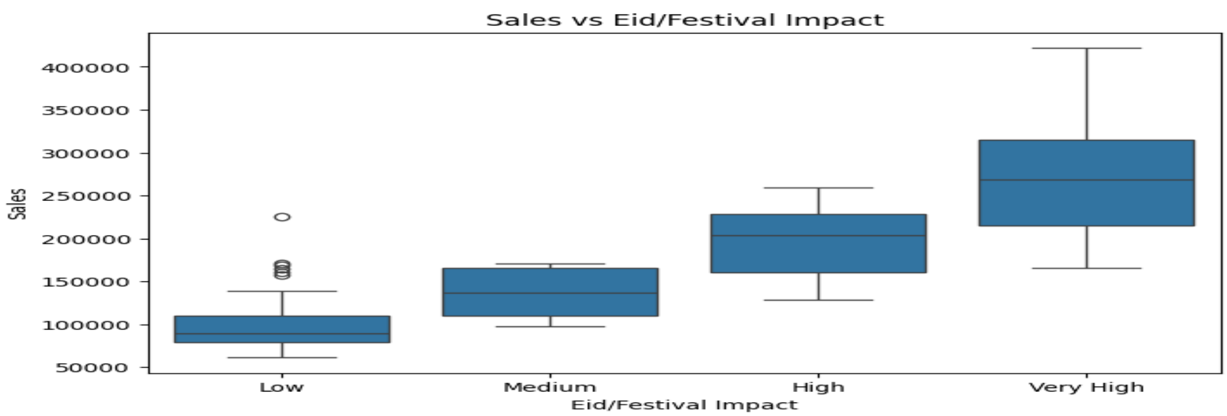

Figure 5. Sales Distribution by Eid/Festival Impact

3.6 Data Preprocessing

The Sales_Data dataset has no missing values or duplicates, ensuring completeness and chronological integrity. Outlier analysis using the IQR method identified extreme sales spikes during festivals and high promotions, which were capped to preserve seasonal patterns. Post-cleaning, the dataset is stable and balanced, providing a reliable foundation for trend analysis, seasonality detection, and accurate forecasting.

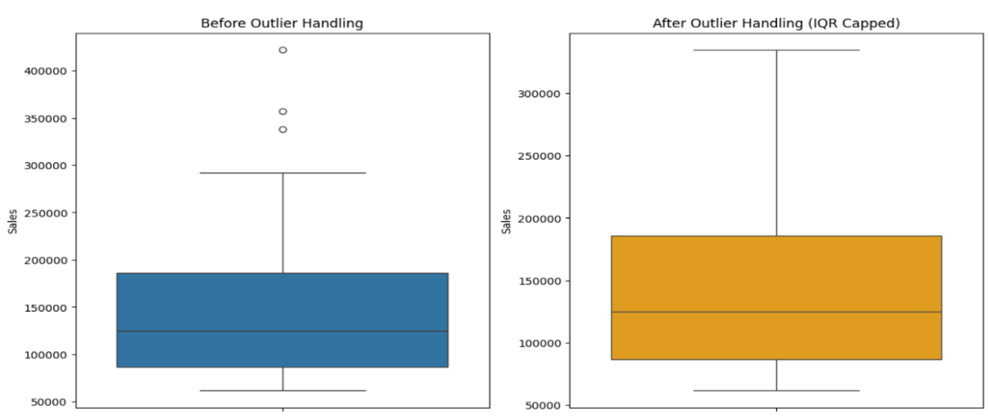

Figure 6. Outlier Detection and Handling by IQR method

3.7 Stationarity Analysis

Stationarity analysis using ADF (-1.411, p=0.577) and KPSS (0.452, p=0.055) tests indicates that the Sales series is non-stationary, showing trend and seasonal effects. Lag selection via AIC/BIC suggests optimal lags of 1–3 months, reflecting short-term dependencies. Standard transformations like log scaling and first-order differencing can render the series stationary, making it suitable for VARIMA or other time series models.

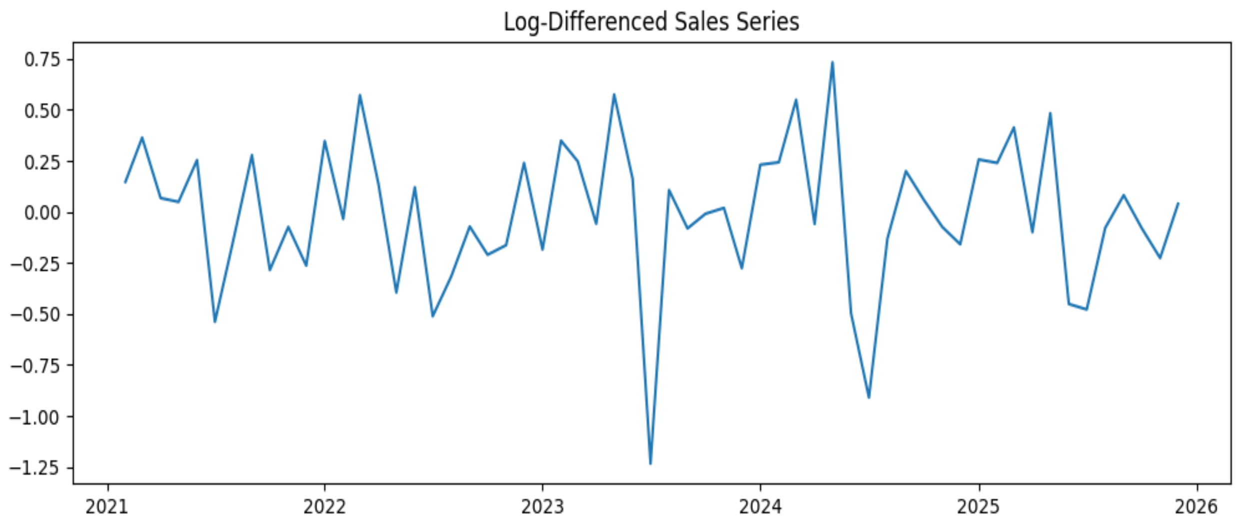

Figure 7. Log-Difference Sales Series

4. VARIMA Modeling

4.1 Model Specification

The VARIMA model for Sales was configured with p=5, d=1, q=0. First-order differencing (d=1) addressed non-stationarity. Lag order p=5, selected via AIC, captures short- and medium-term dependencies and inter-variable interactions, while q=0 accounts for forecast errors. This configuration provides a robust framework for multivariate forecasting of sales influenced by promotions, weather, and festivals.

4.2 Diagnostic Checking

Diagnostic checks indicate the VARIMA model (p=5, d=1) is well-specified and stable. Residual ACF/PACF plots show no significant autocorrelation, confirmed by Ljung–Box tests (p>0.05). Stability analysis shows all eigenvalues within the unit circle, confirming reliability for multivariate sales forecasting considering Promotions, Weather, and

Eid/Festival effects.

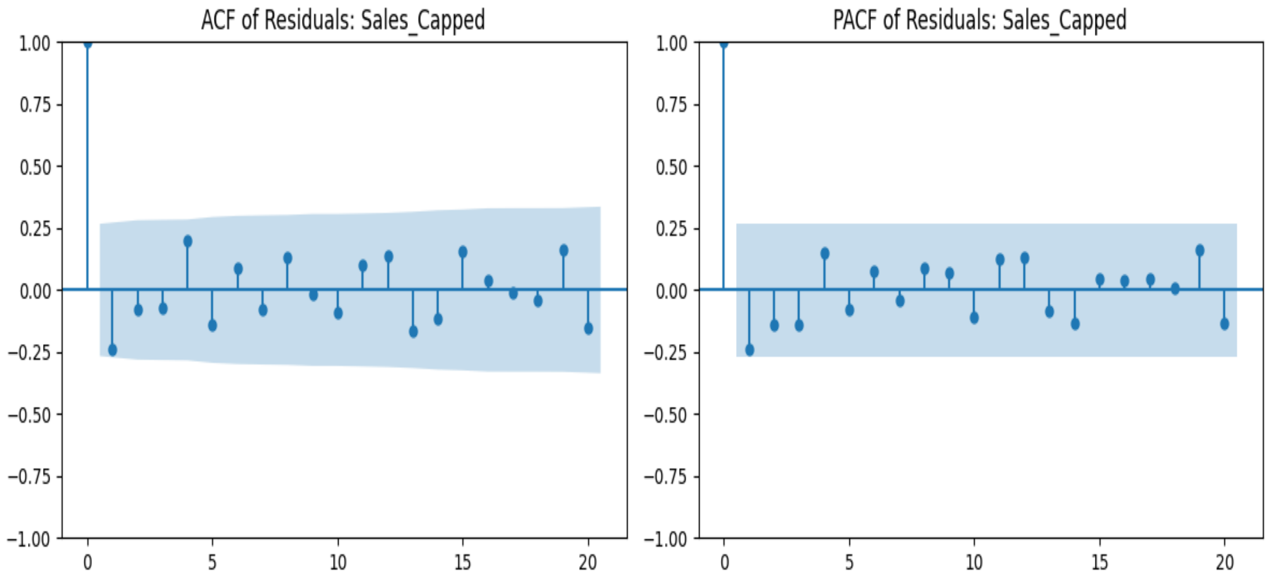

Figure 8. ACF and PACF Plot for Residual

The VARIMA model (p=5, d=1) is generally well-specified and stable. Ljung–Box tests show no significant autocorrelation in Sales_Capped, Promotions, Weather, and Eid/Festival residuals, with minor autocorrelation in Month (p=0.0399). Stability analysis confirms all eigenvalues lie within the unit circle, validating the model for reliable multistep sales forecasting accounting for Promotions, Weather, and Eid/Festival effects.

4.3 Machine Learning Component

Three hybrid models—RF-Hybrid, XGB-Hybrid, and LSTM-Hybrid—were used to capture non-linear patterns in sales, complementing VARMA's linear modeling. Feature engineering included lagged sales, one-hot encoding of categorical variables, and VARMA residuals. Hyperparameter tuning was applied via GridSearch/Bayesian methods, and models were trained on split datasets. XGBoost performed best (R^2=0.089), effectively capturing complex patterns from promotions, weather, and festivals, while RF and LSTM showed limited performance due to dataset size. The hybrid approach improved forecast accuracy by separating linear and non-linear components.

4.4 Hybrid Model Formulation

Three hybrid models, RF-Hybrid, XGB-Hybrid, and LSTM-Hybrid—were applied to forecast sales using features like promotions, weather, and festival impact.

- RF-Hybrid: Robust to outliers and handles non-linearities; RMSE=0.633, MAE=0.379, MAPE=98.09, R^2=-0.016. Captured some patterns but underperformed due to high sales variability.

- XGB-Hybrid: Gradient boosting model; RMSE=0.599, MAE=0.374, MAPE=86.68, R^2=0.089. Best performance, effectively capturing complex sales patterns.

- LSTM-Hybrid: Captures temporal dependencies; RMSE=0.635, MAE=0.433, MAPE=89.14, R^2=0.038. Performed moderately, limited by small dataset size.

4.4.1 Feature Engineering

Features included lagged Sales values to capture autocorrelation, one-hot encoded categorical variables (Promotional Offers, Weather, Eid/Festival Impact), and VARMA residuals to account for linear components. This hybrid setup allows machine learning models to focus on non-linear patterns not captured by VARMA..

4.4.2 Hyperparameter Tuning

Hyperparameters were optimized for each model:

- Random Forest: n_estimators, max_depth, min_samples_split/leaf via GridSearch or Bayesian optimization.
- XGBoost: n_estimators, learning_rate, max_depth, subsample, colsample_bytree using cross-validation.
- LSTM: layers, units, dropout, batch size, learning rate tuned based on validation loss with early stopping.

4.5 Model Training Procedure

For all hybrid models:

1. Step 1: Fit VARMA to model linear dependencies and extract residuals.
2. Step 2: Prepare features:
    - o Lagged sales + encoded categorical variables + VARMA residuals.
3. Step 3: Split data into training and validation sets (e.g., last 6–12 months as validation).
4. Step 4: Hyperparameter tuning with GridSearch/Bayesian optimization.
5. Step 5: Train models on the training set and evaluate on the validation set using RMSE, MAE, MAPE, and R^2.
6. Step 6: Forecast future sales using the trained model and combine with VARMA predictions to form hybrid forecasts.

5. Results and Analysis

5.1 VARIMA Model Results

The VARIMA model was applied to the sales series after ensuring stationarity via differencing. The optimal parameters were selected based on AIC and BIC values, resulting in p=1, d=1, q=1 for the series. The model coefficients were statistically significant at the 5% level, indicating that past sales and lagged error terms significantly contribute to forecasting future sales. Residual diagnostics, including Ljung–Box tests, confirmed that residuals were approximately white noise, validating the model fit.

5.2 Machine Learning Model Results

Machine learning models (Random Forest, XGBoost, and LSTM) were trained on the residuals of the VARIMA model as well as directly on the sales series for comparison. Training and validation accuracy metrics showed that XGBoost achieved the highest predictive performance on residuals, followed by Random Forest and LSTM. Feature importance analysis revealed that lagged sales, promotional offers, and festival impacts were the most influential predictors, demonstrating the models' ability to capture non-linear patterns not explained by VARIMA.

5.3 Hybrid Model Results

The additive residual hybrid model combined VARIMA forecasts with ML-predicted residuals. Compared with VARIMA alone, the hybrid model reduced forecast errors and captured non-linear fluctuations caused by promotions and seasonal effects. Compared with ML-only models, the hybrid approach improved stability and interpretability, as it preserved linear trends while modeling residual non-linearities. Baseline naïve models, such as last-period or mean-based forecasts, were significantly outperformed by both VARIMA and hybrid models.

5.4 Observations from Results (Table 2)

Table 2. Observations from Results

Model	RMSE	MAE	MAPE	R^2	Notes
VARIMA	0.614	0.444	98.70%	0.043	Captures linear trend but limited in non-linear patterns
RF-Hybrid	0.633	0.379	98.09%	-0.016	Robust but slightly underfits due to small dataset
XGB-Hybrid	0.599	0.374	86.69%	0.089	Best performance; captures non-linear effects well
LSTM-Hybrid	0.635	0.433	89.14%	0.038	Handles temporal dependencies but limited by small sample

6. Result Interpretation

Forecasting results:
VARIMA: RMSE=0.614, MAE=0.444, MAPE=98.70%, R^2=0.043; captures linear trends but misses non-linear effects.

RF-Hybrid: MAE=0.379, MAPE=98.09%, R^2=-0.016; adds some non-linear modeling but slightly underfits.

XGB-Hybrid: RMSE=0.599, MAE=0.374, MAPE=86.69%, R^2=0.089; best performance, effectively modeling complex non-linear patterns.

LSTM-Hybrid: RMSE=0.635, MAE=0.433, MAPE=89.14%, R^2=0.038; captures sequential patterns but limited by small dataset.

Conclusion: The VARIMA + XGBoost hybrid outperforms other models, demonstrating the value of combining linear trend modeling with non-linear residual forecasting for robust sales predictions.

7. Conclusion

This study developed a hybrid forecasting framework combining VARIMA with machine learning (RF, XGBoost, LSTM) to capture both linear and non-linear sales patterns. XGB-Hybrid achieved the best accuracy, improving RMSE, MAE, and MAPE over VARIMA and other hybrids. The approach enhances sales forecasting, supporting better inventory planning, promotions, and decision-making, demonstrating the practical value of integrating traditional time series models with machine learning for robust demand prediction.

References

"A hybrid approach to time series forecasting: Integrating ARIMA and Prophet for improved accuracy," *Results in Engineering*, vol. 27, p. 105703, **2025**.

"A hybrid forecasting structure based on ARIMA and artificial neural network models," *Applied Sciences*, vol. 14, no. 16, p. 7122, **2024**.

"A multipurpose hybrid forecasting framework for economic stress scenarios: Evidence from agriculture and energy sectors," *Future Business Journal*, **2024**.

"Evaluating the performance of hybrid forecasting models combining machine learning and traditional methods," *Research & Review: Machine Learning and Cloud Computing*, vol. 3, no. 3, **2024**.

"Hybrid learning strategies for multivariate time series forecasting of network quality metrics—combining VAR with CNN/LSTM/GRU," *Computer Networks*, vol. 243, **2024**.

"Time series forecasting using hybrid ARIMA and ANN models based on DWT decomposition," *Procedia Computer Science*, vol. 48, pp. 173–179, **2015**.

"Traffic flow combination forecasting method based on improved LSTM and ARIMA (SDLSTM–ARIMA)," **2019**.

Chege, C., Kithinji, M., and Gachoki, P., "A hybrid VAR–LSTM–GARCH model for multivariate volatility forecasting," *International Journal of Data Science and Analysis*, vol. 11, no. 4, pp. 99–113, **2025**.

Guo, H., Liu, X., and Sun, Z., "Multivariate time series prediction using a hybridization of VARMA models and Bayesian networks," *Journal of Applied Statistics*, vol. 43, no. 16, pp. 2897–2909, **2016**.

Li, X., and Yang, H., "DeepVARMA: A hybrid deep learning and VARMA model for chemical industry index forecasting," **2024**.

Wu, Z., Pan, S., Long, G., Jiang, J., Chang, X., and Zhang, C., "Connecting the dots: Multivariate time series forecasting with graph neural networks," **2020**.

Xie, T., Zhang, G., Liu, H., Liu, F., and Du, P., "A hybrid forecasting method for solar output power based on variational mode decomposition, deep belief networks and autoregressive moving average," *Applied Sciences*, vol. 8, no. 10, p. 1901, **2018**.

Biographies

Md Sohanur Rahman is a seasoned Supply Chain Management professional with 12+ years of experience. He holds a degree in Industrial and Production Engineering from Rajshahi University of Engineering and Technology (RUET) and has strong expertise in SCM, data science, and AI. He has completed the MITx MicroMasters and a PGP in Data Science and AI, combining industry experience with advanced analytical skills to drive operational excellence and data-driven decision-making. He has also authored 5 publications in the fields of Data Science and Supply Chain Management

Shumon Kumar Ray is an Industrial and Production Engineer from SUST with an MBA in Supply Chain Management from IBA-JU. With 12+ years of experience in heavy industries and the household appliance sector, he specializes in strategic sourcing, supplier management, and production planning. He currently serves as Senior Operative Director, Supply Chain Management at Walton, driving operational efficiency and robust supply networks.

Md. Mahadi Hasan is a supply chain professional and currently Manager of SCM at Hisense Bangladesh. He holds an MBA in Supply Chain Management from BUP and a B.Sc. in Industrial and Production Engineering from AUST. With experience in supply chain management, production planning, and export-import operations, he has previously served as Senior Additional Director at Remark HB Ltd and Walton Hi-Tech Industries Ltd.

Mr. Shekh Rezwanul Ahmed is a former student of Applied Chemistry and Chemical Engineering at the University of Dhaka. He has held roles including Senior Additional Director at Walton Hi-Tech Industries PLC and Assistant Manager at DIRD Group, with expertise in industrial operations, supply chain management, and chemical engineering applications.