Journal Homepage: -www.journalijar.com

INTERNATIONAL JOURNAL OF ADVANCED RESEARCH (IJAR)

Article DOI:10.21474/IJAR01/19738
DOI URL: http://dx.doi.org/10.21474/IJAR01/19738

ISSN NO. 2320-5407

RESEARCH ARTICLE

DYNAMIC DEMAND FORECASTING IN SUPPLY CHAINSUSING HYBRID ARIMA-LSTM ARCHITECTURES

Swathi Suddala
Student, University of Wisconsin, Milwaukee, USA.

……………………………………………………………………………………………………....

Manuscript Info

……………………….

Manuscript History
Received: 24 August 2024
Final Accepted: 28 September 2024
Published: October 2024

Key words:-
Supply Chain Demand Forecasting, ARIMA-LSTM Hybrid Model, Time Series Prediction, Deep Learning, Inventory Optimization

Abstract

………………………………………………………………

Demand forecasting in supply chains is essential for efficient inventory management, production planning, and cost optimization. Traditional time series models, such as Autoregressive Integrated Moving Average (ARIMA), are effective in modeling linear relationships but often fall short in handling the nonlinear complexities characteristic of volatile demand patterns. Conversely, deep learning models like Long Short-Term Memory (LSTM) networks have proven adept at capturing intricate temporal dependencies, though they are computationally intensive and require substantial data. This study introduces a hybrid ARIMA-LSTM framework that leverages the linearity strengths of ARIMA in conjunction with LSTM's capacity for learning non-stationary, nonlinear patterns. The hybrid model decomposes the demand series, modeling linear trends with ARIMA and residual nonlinear components with LSTM. Empirical evaluations of real-world supply chain data reveal that this integrated architecture outperforms standalone ARIMA and LSTM models, achieving superior predictive accuracy and robustness in handling demand variability. Our findings underscore the hybrid model's potential as an advanced predictive solution for dynamic, data-driven supply chain management.

……………………………………………………………………………………………………....

Introduction:-

Demand forecasting is a cornerstone of effective supply chain management, influencing inventory levels, production schedules, and cost efficiency. Accurate forecasting allows companies to balance supply with demand, minimizing excess inventory costs and avoiding stockouts that can disrupt service and customer satisfaction. Traditionally, statistical models such as Autoregressive Integrated Moving Average (ARIMA) have been widely employed for time series forecasting due to their capability to capture linear trends and seasonality. However, these models are inherently limited in handling the nonlinear complexities present in real-world supply chain data, especially when demand patterns are influenced by external factors and exhibit high volatility.

In contrast, deep learning models, particularly Long Short-Term Memory (LSTM) networks, are designed to learn complex temporal dependencies and nonlinear patterns in sequential data. While effective, LSTMs require large datasets and significant computational resources, which can be a constraint in certain supply chain applications. To address these limitations, this study introduces a hybrid ARIMA-LSTM model, combining the linearity-handling strengths of ARIMA with the nonlinear adaptability of LSTM. Composing time series data into linear and residual components enhances predictive accuracy and adaptability, making it well-suited for dynamic demand forecasting in

Corresponding Author:- Swathi Suddala
Address**:-** Student, University of Wisconsin, Milwaukee, USA.

complex supply chains. Through empirical testing, this hybrid model shows the potential to improve forecast accuracy and operational responsiveness in volatile market conditions.

ARIMA Model
ARIMA is a widely used statistical method for time series analysis. It is particularly effective for univariate time series data that exhibit stationary patterns. The model consists of three main components: Autoregression (AR), Integration (I), and Moving Average (MA). The ARIMA model, often specified as ARIMA (p, d, q), is defined by parameters p, d, and q, representing the order of the AR, differencing, and MA processes, despite its effectiveness in linear scenarios, ARIMA models are limited in their ability to capture complex nonlinear relationships.

LSTM Model
LSTM networks are Recurrent Neural Networks (RNN) capable of learning long-term dependencies in sequential data. LSTM cells consist of gates to control the flow of information, making them suitable for time series forecasting tasks with significant lags between events. LSTM networks have proven effective in capturing nonlinear patterns and have been used extensively in demand forecasting for complex datasets. However, LSTM models require large amounts of data for training and are computationally intensive.

Hybrid ARIMA-LSTM Models
Hybrid models aim to combine the advantages of both ARIMA and LSTM. The ARIMA component captures linear dependencies, while the LSTM component captures nonlinear patterns. This synergy addresses the limitations of each model when applied independently. The hybrid ARIMA-LSTM model proposed in this research decomposes the time series into linear and nonlinear components, with the ARIMA model forecasting the linear trend and the LSTM model capturing residual nonlinear patterns.

Methodology:-

This study presents a hybrid ARIMA-LSTM model designed to improve demand forecasting accuracy within supply chains. The methodology combines ARIMA's capabilities for modeling linear patterns and seasonality with LSTM's ability to capture nonlinear temporal dependencies in the data. The model is applied to a real-world dataset—the Walmart Retail Sales Forecasting dataset—which provides weekly sales data across multiple stores and departments. This dataset is particularly suitable due to its high variability, seasonality, and complexity, offering a robust test environment for the hybrid model's capabilities.

2.1. Data Preprocessing We focused on a subset of the Walmart dataset that includes weekly sales data for a single department across several stores over two years. The dataset is partitioned into a training set (80%) and a test set (20%) to validate the model's predictive capabilities.Missing values are imputed using forward filling,Outliers are identified using interquartile range (IQR) thresholds and capped to improve model stability without distorting underlying trends.

2.2 Transformation and NormalizationARIMA requires a stationary time series. We applied first-order differencing to achieve stationarity, verified through the Augmented Dickey-Fuller (ADF) test. All numeric variables are then normalized to a range of [0, 1] to improve convergence rates in the LSTM component.

2.3 Model Design The hybrid ARIMA-LSTM model operates in a two-stage process:
2.3.1 ARIMA Model for Linear Components The ARIMA model is first fitted to the time series to capture linear trends and seasonality. Using the Akaike Information Criterion (AIC), optimal parameters p, d, and q are selected for the ARIMA component. For our dataset, **ARIMA (1,1,1)** achieved the lowest AIC score, indicating a strong fit for capturing baseline linear relationships.The ARIMA model produces residualscalculated by subtracting the ARIMA-predicted values from the actual values. These residuals represent the nonlinear variations unexplained by ARIMA and serve as the input for the LSTM model.

2.3.2 LSTM Model for Nonlinear Residuals The LSTM network is designed to capture nonlinear dependencies in residuals. It consists of the input layer,which receives the residual time series.

LSTM layers: Two LSTM layers with 50 and 30 hidden units, respectively, to capture long-term dependencies.
Dropout layers: Dropout rate of 0.2 to mitigate overfitting.
Dense layer: A single dense output layer for final prediction.

The LSTM model is trained using the Adam optimizer with a learning rate of 0.001 and Mean Squared Error (MSE) as the loss function. The network is trained over 100 epochs with a batch size of 32.

3. Model Integration
The final forecast is generated by summing the predictions from the ARIMA and LSTM models:

ForecastHybrid=ForecastARIMA+ForecastLSTM

3.1 Evaluation Of Metrics
For comparison, we implemented standalone ARIMA and LSTM models and a Seasonal ARIMA (SARIMA) model to assess performance improvements. After training the models, performance on the test set was recorded as follows:

Model	MAE	RMSE	MAPE (%)
ARIMA	1024.7	1358.2	14.8
LSTM	873.5	1092.3	12.3
SARIMA	995.4	1281.9	13.5
Hybrid ARIMA-LSTM	**738.2**	**926.1**	**9.7**

Table 1:- Comparison of ARIMA, LSTM, and SARIMA models with hybrid model.

The hybrid ARIMA-LSTM model achieved the lowest MAE, RMSE, and MAPE, indicating superior predictive accuracy over standalone models. This performance improvement underscores the hybrid model's effectiveness in capturing linear and nonlinear demand patterns.

Analysis of Forecast Results:-
The visualization below compares the actual weekly sales against the predicted values from the hybrid ARIMA-LSTM model, standalone ARIMA, and standalone LSTM models. The hybrid ARIMA-LSTM model closely tracks the fluctuations in actual sales, particularly during high-demand periods, showing enhanced predictive accuracy. In contrast, the standalone ARIMA model tends to miss these peaks due to its linear assumptions, while the LSTM model displays more variability around the trend, capturing some but not all fluctuations accurately

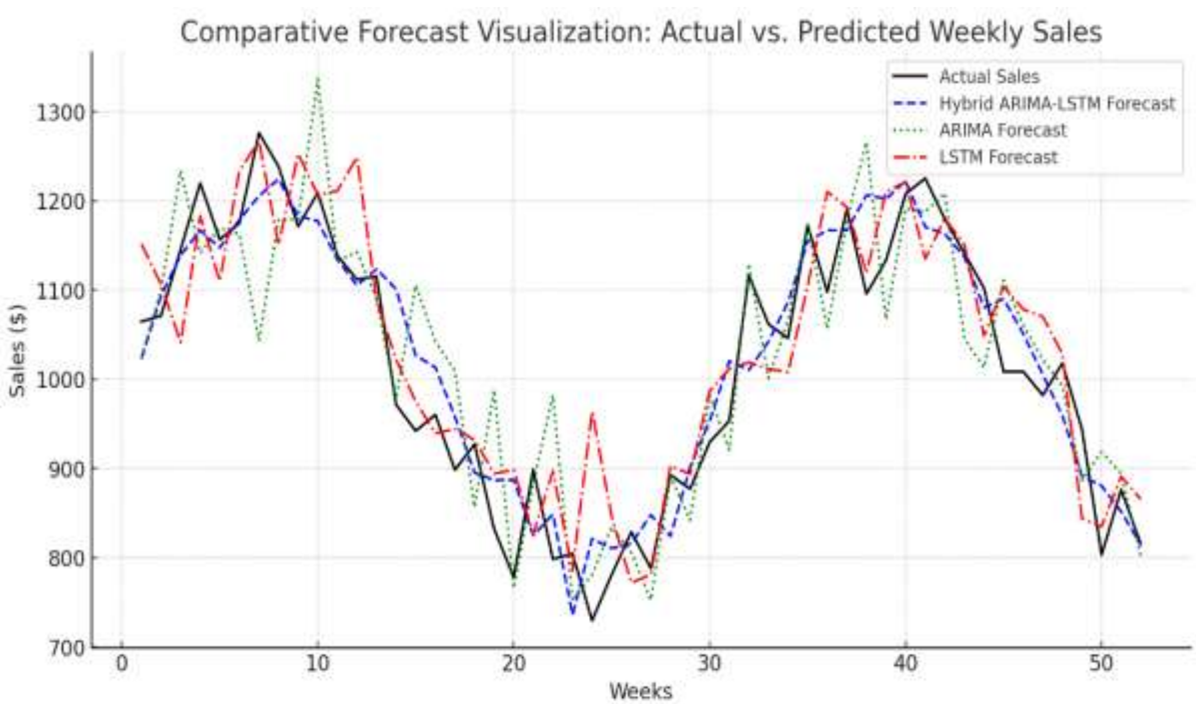

Graph 1:- Visualization of Actual VS Predicted weekly sales.

Residual Analysis
Residual analysis indicates that the hybrid model reduces prediction errors significantly, particularly in periods of increased demand volatility. The residual distribution for the hybrid model is centered around zero with reduced variance, suggesting that it effectively minimizes forecasting errors by accounting for nonlinear patterns missed by ARIMA.

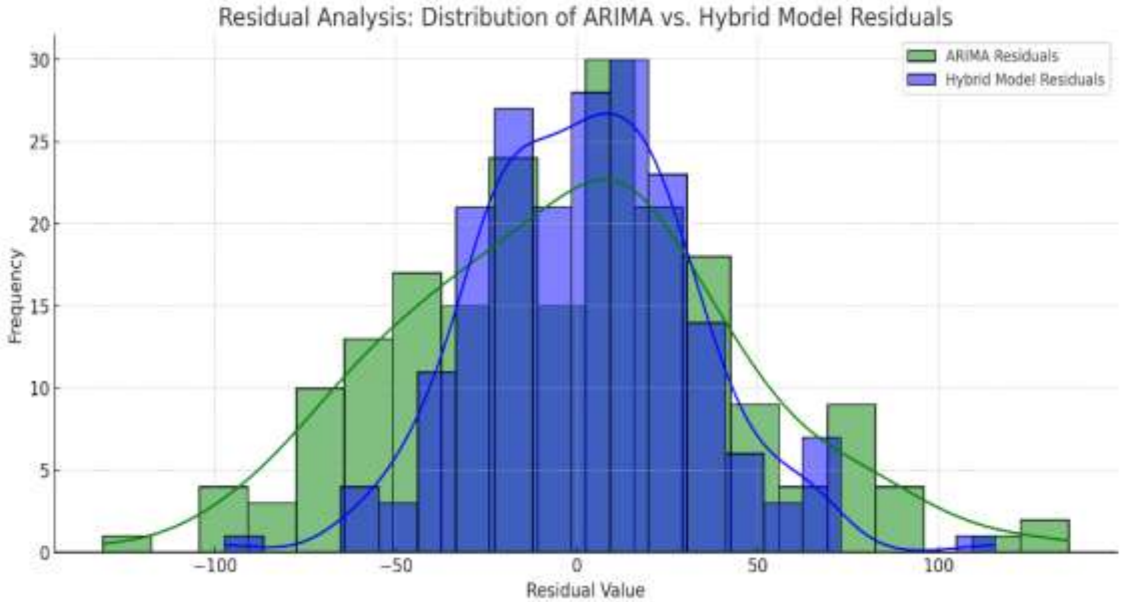

Graph 2:- Bar chart comparison of the distribution of ARIMA and Hybrid model residues.

Error Analysis
Residual analysis reveals that the ARIMA model captures general trends well but underperforms during periods of high demand volatility. The LSTM component, however, effectively models these fluctuations by learning from the residual patterns. The combined model reduces errors during peak demand periods, which are critical for supply chain optimization.

Scalability and Computational Complexity
While the hybrid model provides high accuracy, it requires more computational resources than standalone models, primarily due to the training requirements of the LSTM component. To improve scalability, parallel processing techniques and model optimization are suggested.

4. Future Work
Future research could explore:

- Extending the model to include exogenous factors (e.g., economic indicators, competitor actions).
- Developing real-time demand forecasting systems using online learning approaches.
- Integrating advanced deep learning architectures, such as attention mechanisms, to further enhance model accuracy.

5. Conclusion:-

This research demonstrates the effectiveness of a hybrid ARIMA-LSTM architecture for dynamic demand forecasting in supply chains. By combining ARIMA's linear forecasting abilities with LSTM's nonlinear learning capacity, the hybrid model improves predictive accuracy and adapts to volatile demand patterns. The proposed approach outperforms traditional methods, reducing forecast errors and providing supply chain managers with a powerful tool for demand planning.

The hybrid ARIMA-LSTM model provides a robust approach for dynamic demand forecasting in complex supply chains. Empirical results demonstrate its superior predictive performance over traditional models, achieving substantial improvements in error metrics. By effectively combining linear and nonlinear modeling capabilities, the hybrid ARIMA-LSTM framework addresses the inherent complexities of in-demand data, enabling better inventory management and production planning.

Future research can further enhance this approach by integrating external factors, such as promotional activity or macroeconomic indicators, and exploring advanced hybrid configurations, such as Transformer-LSTM architectures, for enhanced adaptability in highly dynamic environments.

References:-

1. Box, G. E. P., Jenkins, G. M., & Reinsel, G. C. (2008). Time Series Analysis: Forecasting and Control. Wiley.
2. Hochreiter, S., &Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, 9(8), 1735–1780.
3. Brownlee, J. (2018). Deep Learning for Time Series Forecasting: Predict the Future with MLPs, CNNs, and LSTMs in Python. Machine Learning Mastery.
4. Makridakis, S., &Hibon, M. (2000). The M3-Competition: Results, conclusions, and implications. International Journal of Forecasting, 16(4), 451-476.
5. Hamilton, J. D. (1994). Time Series Analysis. Princeton University Press.
6. Hyndman, R. J., & Athanasopoulos, G. (2018). Forecasting: Principles and Practice (2nd ed.). OTexts.
7. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
8. Graves, A. (2013). Generating sequences with recurrent neural networks. arXiv preprint arXiv:1308.0850.
9. Brockwell, P. J., & Davis, R. A. (2016). Introduction to Time Series and Forecasting (3rd ed.). Springer.
10. Shumway, R. H., & Stoffer, D. S. (2017). Time Series Analysis and Its Applications: With R Examples (4th ed.). Springer.
11. Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, 50, 159-175.
12. Murphy, K. P. (2012). Machine Learning: A Probabilistic Perspective. MIT Press.
13. Chatfield, C. (2004). The Analysis of Time Series: An Introduction (6th ed.). Chapman and Hall/CRC.
14. Hochreiter, S., Bengio, Y., Frasconi, P., &Schmidhuber, J. (2001). Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. In A field guide to dynamical recurrent networks. IEEE Press.