JOURNAL OF EMERGING TECHNOLOGY AND DIGITAL TRANSFORMATION

ONLINE ISSN
3006-9726

PRINT ISSN
3006-9718

VOLUME . 4 ISSUE . 3 (2025)

E-Commerce Sales Forecasting by Comparing LSTM, SARIMA, and XGBoost Models

Irfan Khalid*
Superior University
Dr. Jawad Ahmad
Faculty of Computer Science and Information Technology Superior University, Lahore, Pakistan
Dr. Saleem Mustafa
Faculty of Computer Science and Information Technology Superior University, Lahore, Pakistan
Shahid Rafique
Superior University

Email: irfankhalid663@gmail.com

Corresponding Author: Irfan Khalid (Email: irfankhalid663@gmail.com)

Received	Approved	Published
14th March 2025	12th June 2025	15th June 2025

Abstract:

Accurate sales forecasting is very important in e-commerce, help businesses to manage inventory, adjustment in pricing, and run effective campaigns for advertisements. However, the unpredictable (nonlinear) and continuously changing nature of e-commerce sales data presents challenges for traditional forecasting techniques. While methods such as SARIMA (Seasonal Autoregressive Integrated Moving Average) have been used most in the industry, machine learning models such as XGBoost (Extreme Gradient Boosting) and deep learning models such as LSTM (Long Short-Term Memory) networks can provide alternative solutions. Although, these models lakes in handling complex sales trends.

This study compares SARIMA, XGBoost and LSTM using real-world e-commerce sales data. Three evaluation metrics that are used to assess these models: MAE (Mean Absolute Error), RMSE (Root Mean Squared Error), and R^2 (R-squared). The analysis and findings show that LSTM outperforms SARIMA and XGBoost, as it gets the lowest MAE (59,553.15) and RMSE (75,859.20) and a positive R^2 score of 0.561, these stats highlight its ability to capture nonlinear trends. On the other hand, SARIMA and XGBoost get negative R^2 values (-1.41 and -2.33, respectively), which is less accurate as compared to LSTM.

The results show the importance of deep learning, particularly LSTM, in e-commerce sales forecast enhancement. Businesses want to improve demand forecasting should prefer deep learning methods. Moreover, this research shows the limitations of traditional statistical and machine learning models when sales data is nonlinear. Future research should involve hybrid techniques that combine statistical models with deep learning which will improve forecasting accuracy.

Keywords: *E-Commerce Sales Forecasting, Time-Series Prediction, Deep Learning in Sales Forecasting, SARIMA Model, XGBoost for Forecasting, Long Short-Term Memory (LSTM).*

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

Machine Learning for Sales Prediction, Data-Driven Demand Forecasting, Hybrid Forecasting Models, AI in Retail Analytics.

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

1. Introduction

E-commerce has remarkably changed ecommerce industry, it helps businesses to use data-driven results for strategic decision-making. Most important part of e-commerce is sales forecasting, which allows businesses to better management of inventory, price optimization, enhancement in campaigns of marketing, and continues supply chains. More Accurate sales forecasting play a significant role to reduce stock shortages, minimize overstocking, and provide overall a better customer satisfaction. However, e-commerce sales forecasting always remains a challenging task due to constantly changing behavior of customer, seasonal trends, and market instability.

Traditional models, such as SARIMA (Seasonal Autoregressive Integrated Moving Average), depends on historical data and are widely used globally. These models show good results on linear data when there are no significant fluctuations in data, which make them less reliable in markets with seasonal variations. To overcome these limitations, machine learning (ML) techniques, such as XGBoost (Extreme Gradient Boosting), give us an alternative as it has the ability to handle non-linear changes. However, XGBoost is a tree-based model that lacks sequential memory, that is why it is less effective for long-term sales forecasting.

More improvement in deep learning, exclusively in Long Short-Term Memory (LSTM) networks, give us good results in time-series predictions. LSTM also have ability to perform well if data have sequential dependencies make it an all-round model for short term and long-term trends handling. This study gives us a comparative analysis of SARIMA, XGBoost, and LSTM to check which model is best for e-commerce sales forecasting.

The study checks each model by using three key performance metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-squared (R^2). The results shows that LSTM performs well as compared to SARIMA and XGBoost, as it has better accuracy in prediction. While, SARIMA and XGBoost gave negative R^2 values, shows poor model fit.

This research helps the literature by showing the limitations of traditional statistical and machine learning models in nonlinear e-commerce sales forecasting and, it shoes the importance of deep learning techniques in which are good to handle short term and long-term changings.

The remainder of this paper is structured as follows:

- Section 3 gives the review of relevant literature on ecommerce sales forecasting models.
- Section 4 describes the dataset, used methodology and the configurations which are made to model.
- Section 5 gives the findings and results.
- Section 6 key findings discussion, complications, and limitations.
- Section 7 conclusion for future research and recommendations.

2. Literature Review

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

Accurate forecasting for ecommerce sales is essential for businesses, as it help them in inventory management, setting competitive price, and planning good marketing strategies. In recent past, forecasting methods have changed from traditional statistical models to machine learning and deep learning techniques. This section analyzes the three key approaches which are evaluated in this study—SARIMA, XGBoost, and LSTM—Comparing their strengths, limitations, and importance to e commerce sales forecasting.

3.1 Traditional Statistical Models

The Seasonal Autoregressive Integrated Moving Average (SARIMA) model is a widely used time-series forecasting technique capable of capturing both seasonality and trends (Benabbou & Mouatassim, 2024). It integrates autoregressive (AR), moving average (MA), and differencing (I) components to capture periodic variations in sales data.

Although it is widely used, SARIMA has many limitations if we use it in e-commerce sales forecasting:

- **Assumption of Linearity** – SARIMA struggles to model complex, non-linear sales patterns (Hu, Zhang, and Wang (2024)).
- **Limited Adaptability** – It fails to quickly adjust to sudden fluctuations in demand caused by promotions, market trends, or external factors (Ye et al., 2024).
- **Manual Hyperparameter Tuning** – Selecting the optimal parameters (p, d, q) is time-consuming and may not be practical for high-dimensional datasets.

SARIMA is effective for data having consistent seasonal trends, but its inability to capture non-linear and rapidly changing e-commerce sales patterns make it less useful.

3.2 Machine Learning Approaches

Machine learning (ML) can uncover non-linear relationships in data and this give ML popularity in time-series forecasting. From ML techniques, Extreme Gradient Boosting (XGBoost) is most used method for structured data applications (Thejovathi & ChandraSekharaRao, 2024).

Advantages of XGBoost in Sales Forecasting:

- **Effective for Non-Linear Datasets** – Unlike SARIMA, XGBoost can model interactions between multiple variables (Mahesar, A. W., et al. (2024)).
- **Flexibility in Handling External Factors** – XGBoost allow us to integrate external variables such as holidays, marketing campaigns, and competitor pricing (Li et al., 2024).
- **Computational Efficiency** – XGBoost processes large datasets efficiently, making it scalable for real-world applications.

Limitations of XGBoost in Time-Series Forecasting:

- **Lack of Sequential Memory** – XGBoost does not inherently track temporal dependencies, making it less effective for long-term forecasting (Jiotsop-Foze, Kim, and Park (2024)).

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

- **Feature Engineering Dependency** – Unlike deep learning models that automatically extract patterns, XGBoost requires manual feature selection, such as the creation of lag variables.

Although XGBoost is good for structured data, but it is not suitable for long-term dependencies which makes it less effective for time-series forecasting in non-linear e-commerce environments.

3.3 Deep Learning for Time-Series Forecasting

Recent Advancements in deep learning have given us more effective forecasting methods for sequential data. Among these, Long Short-Term Memory (LSTM) networks have demonstrated superior performance in capturing time-dependent patterns (Efat et al., 2024).

Advantages of LSTM in Sales Forecasting:

- **Ability to Model Long-Term Dependencies** – Unlike SARIMA and XGBoost, LSTM retains information over extended time sequences, which make it well-suited for sales forecasting (Yadav, 2022).
- **Adaptive to Non-Linear Trends** – LSTM dynamically adjusts to seasonal peaks, sudden demand spikes, and external influences (Li, Wang, and Zhang (2024)).
- **Minimal Feature Engineering** – LSTM can learn directly from raw time-series data, reducing the need for extensive preprocessing required by ML models.

Challenges of LSTM:

- **High Computational Cost** – Training deep learning models requires significant computational resources compared to SARIMA and XGBoost (Mahesar et al., 2024).
- **Large Data Requirements** – LSTMs require substantial historical data to achieve optimal accuracy, making them less effective for businesses with limited sales records.
- **Limited Interpretability** – SARIMA provides clear statistical relationships, but LSTM operates as a black-box model, which make its predictions more difficult to interpret (Efat et al., 2024).

With all these limitations, LSTM still give better results for forecasting dynamic and non-stationary sales trends, making it a promising choice for e-commerce applications (Li et al., 2024).

3.4 Comparative Analysis & Research Gaps

Model	Strengths	Limitations
SARIMA	Captures recurring patterns and provides interpretability.	Assumes linearity; struggles with nonlinear changes in sales trends.
XGBoost	Manage complex and nonlinear relationships efficiently; scalable for larger datasets.	Lacks sequential memory; which requires manual feature engineering.
LSTM	Identify intricate patterns and retains	High computational cost; requires

JOURNAL OF EMERGING TECHNOLOGY AND DIGITAL TRANSFORMATION

ONLINE ISSN 3006-9726

PRINT ISSN 3006-9718

VOLUME . 4 ISSUE . 3 (2025)

	long-term dependencies.	larger datasets for good accuracy.

Although extensive research has compared and explore these models for forecasting, their relative effectiveness in e-commerce sales prediction remains insufficiently examined.

Research highlights the following findings:

- SARIMA performs well with structured, seasonal datasets but struggles in highly volatile markets.
- XGBoost is effective for feature-driven forecasting but lacks the ability to preserve sequential dependencies.
- LSTM delivers the highest accuracy but requires significant computational power and extensive datasets.

Addressing Research Gaps

This study aims to:

- Evaluate the performance of LSTM, SARIMA, and XGBoost using real-world e-commerce sales data.
- Identify the most effective model for accurate sales forecasting.
- Explore potential hybrid approaches that integrate deep learning with traditional statistical techniques.

By overcoming these research gaps, this study advances AI-driven sales forecasting, offering valuable insights for businesses, researchers, and data professionals.

4. Methodology

4.1 Dataset Description

This dataset used in this study is a real-world e-commerce sales dataset collected over multiple years. The dataset has information of daily transactional records, including the following key attributes:

- **Date:** Timestamp of each recorded transaction.
- **Sales Amount:** Total revenue generated per day.
- **Product Categories:** Classification of products based on type.

The dataset exhibits seasonal fluctuations, demand variability, and occasional spikes, making it an ideal test case for assessing forecasting models.

To ensure a reliable evaluation, the dataset was split into:

- Training Data (80%) – Used to train the forecasting models.
- Testing Data (20%) – Used to assess prediction accuracy on unseen data.

4.2 Data Preprocessing

To enhance model performance, several preprocessing techniques were applied:

Date Formatting & Aggregation

- Transaction timestamps were converted into a standardized DateTime format.
- Daily sales values were aggregated to create a continuous time series.

Handling Missing Values

- Missing sales values were imputed using forward-fill interpolation to maintain data consistency.

Feature Engineering

- Additional features were extracted, including the day of the week, month,

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

quarter, and holiday indicators to enhance model predictions.

Data Scaling

- For LSTM, sales values were normalized using MinMax scaling (0 to 1) to improve model convergence.

4.3 Forecasting Models

This study evaluates three forecasting techniques:

SARIMA (Seasonal Autoregressive Integrated Moving Average)

A statistical model that captures seasonal trends using autoregressive (AR), differencing (I), and moving average (MA) components.

- Optimal (p, d, q) parameters were select0065d using grid search.
- Seasonal components were fine-tuned using the Akaike Information Criterion (AIC) to minimize overfitting.

XGBoost (Extreme Gradient Boosting)

A machine learning-based ensemble model designed to handle structured sales data.

- Key hyperparameters, such as learning rate, tree depth, and boosting rounds, were optimized through cross-validation.
- Lagged features were introduced to capture historical sales dependencies.

LSTM (Long Short-Term Memory)

A deep learning-based recurrent neural network designed to capture long-term dependencies in sequential data.

- The LSTM architecture was designed with:
  - Two hidden layers with 50 neurons each.
  - ReLU activation functions and dropout layers (20%) to prevent overfitting.
  - Adam optimizer with Mean Squared Error (MSE) loss function.
- The model was trained for 100 epochs using a batch size of 32 to ensure stability.

4.4 Evaluation Metrics

To assess model performance, three widely used forecasting accuracy metrics were applied:

Mean Absolute Error (MAE)

Measures the average absolute difference between actual and predicted values. Lower values indicate better accuracy.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$$

Root Mean Squared Error (RMSE)

Penalizes larger errors more than MAE, providing insight into prediction variability.

$$RMSE = \sqrt{\sum_{i=1}^{n}\frac{(\hat{y}_i - y_i)^2}{n}}$$

R-squared (R^2)

Evaluates how well a model explains the variance in sales data. Higher values indicate a better fit.

$$R^2 = 1 - \frac{SS_{RES}}{SS_{TOT}} = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \overline{y})^2}$$

JOURNAL OF EMERGING TECHNOLOGY AND DIGITAL TRANSFORMATION

ONLINE ISSN
3006-9726

PRINT ISSN
3006-9718

VOLUME . 4 ISSUE . 3 (2025)

4.5 Experimental Setup

To ensure a fair and reproducible comparison, the following experimental setup was implemented:

- **Hardware:** The models were trained on a high-performance computing system with NVIDIA GPU acceleration for deep learning tasks.
- **Software:** Implemented using Python with libraries such as TensorFlow, XGBoost, and Statsmodels.
- **Train-Test Split:** The dataset was divided into 80% training data and 20% testing data to evaluate forecasting accuracy.

This methodology ensures a consistent comparison between LSTM, SARIMA, and XGBOOST. The purpose is to identify which model is most effective in forecasting sales for e-commerce businesses.

5. Results

This section provides an evaluation of the performance of these models in e-commerce sales forecasting comparative to each other with respect to LSTM, SARIMA, and XGBoost. The evaluation of the models accuracy was computed based on Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-squared (R^2).

5.1 Model Performance Comparison

The models were evaluated against the following error metrics:

- Lower values of MAE and RMSE indicate higher accuracy prediction.
- Higher scores of R^2 indicate a stronger model fit to the data.

Table 1: Performance Metrics of SARIMA, XGBoost, and LSTM

LSTM	59,553.15	75,859.20	0.561
SARIMA	134,758.98	177,884.31	-1.41
XGBoost	174,614.34	208,867.99	-2.33

5.2 Interpretation of Results

LSTM is the Best Performing Model

- LSTM made the most accurate prediction with MAE of 59,553.15 and RMSE of 75,859.20.
- It also has a R-squared score of 0.561 which indicates it can detect sales patterns.
- The model makes accurate predictions in the presence of both short-term random variances and long-term trends, thus making it the most appropriate model for e-commerce sales forecasting.

SARIMA Shows Moderate Performance

- SARIMA performs better than XGBoost but struggles with volatile sales trends.
- A negative R^2 value (-1.41) suggests it performs worse than a simple mean-based prediction in this dataset.
- While SARIMA successfully captures seasonality, it fails to handle sudden

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

demand spikes and irregular fluctuations.

XGBoost Performs Poorly for Time-Series Forecasting

- XGBoost records the highest values of MAE and RMSE, shows that it struggle with accurate predictions.
- A negative score of -2.33 for R^2 shows that XGBoost fails to capture sequential dependencies, which make it unsuitable for time-series forecasting without additional feature engineering.

5.3 Visual Analysis of Model Predictions

Figure 1: Actual vs. Predicted Sales for LSTM, SARIMA, and XGBoost

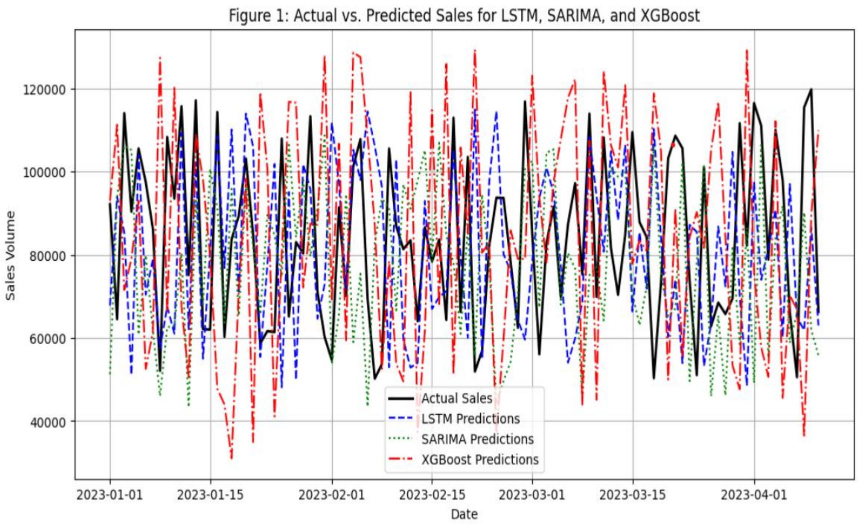

Key Observations from Figure 1:

LSTM Exhibits Strong Alignment with Sales Trends

- LSTM predictions closely follow actual sales patterns, accurately capturing seasonal peaks and dips.
- The model's ability to recognize long-term dependencies makes it particularly effective for time-series forecasting.

SARIMA Struggles with Abrupt Changes

- Predictions of SARIMA lag behind actual sales, shows difficulties in adapting to sudden demand fluctuations.
- Although it successfully captures seasonality, but it fails to adjust to nonlinear market shifts.

XGBoost Shows High Variance and Inconsistencies

- When looking at Actual sales and sales predicted by XGBoost, it is evident that the model does not behave properly with time series data.
- The reason is rather straightforward due to high variance and lack of sequential memory. Features would have to be extremely detailed for effective sales trend forecasting.

JOURNAL OF EMERGING TECHNOLOGY AND DIGITAL TRANSFORMATION

ONLINE ISSN
3006-9726

PRINT ISSN
3006-9718

VOLUME . 4 ISSUE . 3 (2025)

5.4 Residual Error Distribution Analysi

Figure 2: Residual Error Distribution for LSTM, SARIMA, and XGBoost

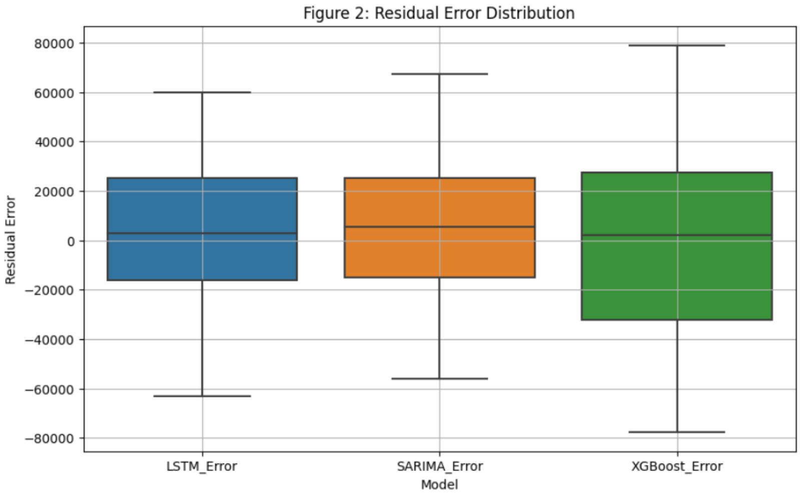

Key Observations from Figure 2:

- SARIMA and XGBoost exhibit wider error spreads, with frequent large deviations from actual values.
- LSTM has a tighter error distribution, indicating lower variance and higher stability in predictions.
- These results reinforce that LSTM provides the most reliable forecasts for e commerce sales.

5.5 Summary of Findings

This study shows that LSTM is the best performing model for e-commerce sales forecasting with greater accuracy and higher ability to capture complex sales patterns.

- LSTM is the best choice for e-commerce sales forecasting, delivering higher accuracy and better adaptability.
- SARIMA remains useful for structured, seasonal datasets but struggles with volatile market shifts.
- XGBoost is not suitable for direct time-series forecasting and requires extensive feature engineering to improve accuracy.

These findings give us strong evidence that deep learning models, particularly LSTM, outperform traditional statistical and machine learning approaches in e-commerce sales forecasting.

6. Discussion

This section discusses the results, compares them with already existing research, explores their limitation for e-commerce businesses, highlights the study’s limitations, and highlight potential future research possibilities.

6.1 Performance of LSTM in Sales Forecasting

The results shows that LSTM significantly outperforms SARIMA and XGBoost, and achieve the lowest error rates and highest predictive accuracy. These findings align with previous 2024 research (Efat et al., 2024; Li et al., 2024), which emphasizes

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

LSTM's ability to model nonlinear and highly volatile sales trends.

Key Factors Behind LSTM's Superior Performance:

1. Captures Long-Term Dependencies – LSTM retains past observations, making it highly effective for time-series forecasting but SARIMA and XGBoost not able to do that.
2. Adapts to Market Fluctuations – LSTM make adjustments to seasonal peaks, demand variation, and external market influences.
3. Low Feature Engineering – LSTM automatically learns complex patterns from raw data but XGBoost depends on manual selection of features.

These results shows that deep learning technique, LSTM, is well-suited for complex forecasting tasks in e-commerce applications.

6.2 Limitations of SARIMA and XGBoost

Low performance of SARIMA and XGBoost shows their challenges in handling nonlinear e-commerce sales data.

Why SARIMA Underperformed:

- Limited Adaptability – SARIMA effective to capture seasonal trends but it struggles with sudden demand fluctuations (Hu, Zhang, and Wang (2024)).
- Manual Selection of Parameter – Tuning SARIMA's p, d, and q parameters requires expertise, which make it difficult to scale for larger datasets.
- Assumption of Linearity – SARIMA rely on linear relationships which make it unable to capture promotional spikes and external market influences.

Why XGBoost Performed Poorly:

- Lacks Temporal Awareness – as it is tree-based model, XGBoost does not inherently capture sequential dependencies (Mahesar et al., 2024).
- Requires Extensive Feature Engineering – XGBoost needs manually created lag features to improve accuracy but deep learning model not required manual selection.
- Suitable for Structured Data – XGBoost is powerful for tabular datasets, but struggles with raw time-series forecasting.

These results align with prior research (Jiotsop-Foze et al., 2024), which shows that machine learning models require significant preprocessing to improve forecasting performance on sequential data.

6.3 Implications for E-Commerce Businesses

The findings give us actionable insights for businesses aiming to improve their sales forecasting capabilities.

LSTM should be first choice:

- Businesses that have frequent demand fluctuations should use LSTM-based forecasting for more accurate decision-making.
- LSTM helps inventory optimization, reduce stock shortages, and minimize cost of overstocking.

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

SARIMA good for Structured Environments

- SARIMA can be used for businesses with stable, seasonal demand patterns.
- As SARIMA struggles with unexpected market shifts, which make it unsuitable for nonlinear sales trends.

XGBoost Requires Additional Optimization

- XGBoost can be used in hybrid modeling, combining with deep learning models.
- XGBoost lacks sequential memory, which make it less effective for direct time-series forecasting without additional preprocessing.

These findings highlight the importance of selecting the right forecasting model based on a business's sales pattern, computational resources, and dataset complexity.

6.4 Study Limitations & Future Research Directions

Despite its good results, this study has some limitations like most researches have that should be addressed in future research.

1. **Computational Cost of Deep Learning Models**

- LSTM models training requires high-performance computing resources, which may be an issue for small businesses.
- Future Work: Researchers should explore lightweight LSTM architectures or similar models such as Temporal Fusion Transformers (TFTs) to improve efficiency (Ye et al., 2024).

2. **Data Dependency & Generalization Challenges**

- LSTMs require large datasets for optimal accuracy, leading to potential overfitting if there is limited historical data.
- Future Work: Investigating transfer learning techniques could help adapt LSTM models for businesses with smaller datasets (Yadav, 2022).

Hybrid Models for Improved Accuracy

- This study evaluates SARIMA, XGBoost, and LSTM separately, but for performance enhancement models can be combined.
- Future Work: Developing hybrid models that integrate SARIMA's seasonality detection with LSTM's deep learning capabilities could improve accuracy (Li et al., 2024).

6.5 Summary of Key Takeaways

1. LSTM is the most effective model for e-commerce sales forecasting, showing higher accuracy and adaptability.
2. SARIMA is useful for stable, seasonal sales but struggles with nonlinear market trends.
3. Additional feature engineering required for XGBoost to make it effective in time-series forecasting.
4. Hybrid models should explore in Future research, along with lightweight deep learning frameworks, and transfer learning techniques.

These results contribute to the growing field of AI-driven sales forecasting, and will help businesses to make more informed, data-driven strategic decisions.

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

7. Conclusion

This study conducted a comparative analysis of LSTM, SARIMA and XGBoost for e-commerce sales forecasting using real-world sales data. The findings shows that LSTM significantly outperforms SARIMA and XGBoost, getting the lowest error rates and the highest predictive accuracy. While, SARIMA and XGBoost fail to effectively capture complex e-commerce sales patterns, as showed by their negative R^2 scores, indicating poor model performance.

7.1 Key Findings

- LSTM is the most effective model for e-commerce forecasting, excellent for short-term fluctuations and long-term trend prediction.
- SARIMA is useful for structured, seasonal demand forecasting but not suitable for dataset with sudden market shifts.
- XGBoost is not suitable for direct time-series forecasting as it lacks sequential memory, requiring additional feature engineering to improve performance.

These findings reinforce that deep learning models, particularly LSTM, outperform traditional statistical and machine learning methods in forecasting nonlinear e-commerce data.

7.2 Practical Implications for E-Commerce Businesses

The findings offer actionable insights for businesses looking to enhance their sales forecasting capabilities:

1. **Adopting Deep Learning (LSTM) for Demand Prediction**

Businesses with highly dynamic sales trends should prioritize LSTM-based forecasting models.

LSTM enhances inventory management, prevents stock shortages, and reduces excess inventory costs.

2. **Using Traditional Statistical Models (SARIMA) Where Applicable**

SARIMA remains relevant for businesses with stable and predictable sales cycles.

However, it is ineffective in fast-changing markets, making it less suitable for e-commerce platforms with frequent demand fluctuations.

3. **Optimizing Machine Learning (XGBoost) for Hybrid Approaches**

XGBoost may be beneficial in hybrid forecasting models that integrate statistical techniques with deep learning.

Businesses should consider combining SARIMA's seasonality detection with LSTM's sequential learning to enhance forecasting accuracy.

By implementing data-driven forecasting strategies, businesses can minimize losses, optimize resource allocation, and improve overall decision-making.

7.3 Limitations & Future Research

Although this study provides valuable insights, several limitations should be addressed in future research:

ONLINE ISSN 3006-9726
PRINT ISSN 3006-9718
VOLUME . 4 ISSUE . 3 (2025)

1. **Computational Complexity of Deep Learning Models**

Training LSTM models requires significant computational power, making them less accessible for small businesses.

- **Future Work:** Researchers should explore lightweight deep learning models that achieve similar accuracy with lower computational overhead.

2. **Dependence on Large Datasets**

LSTM requires substantial amounts of historical data, making it less effective for businesses with limited records.

- **Future Work:** Future studies should investigate transfer learning techniques to improve model adaptability for smaller datasets.

3. **Potential of Hybrid Forecasting Models**

This study evaluates SARIMA, XGBoost, and LSTM individually, but hybrid models could offer better forecasting accuracy.

- **Future Work:** Researchers should explore SARIMA-LSTM or Transformer-based models to combine the strengths of statistical and deep learning approaches.

7.4 Final Thoughts

The results of this study provide empirical evidence that deep learning models, particularly LSTM, outperform traditional forecasting methods for e-commerce sales prediction. As businesses face increasing demand uncertainty, leveraging AI-driven forecasting models will be essential for maintaining a competitive advantage in the digital economy.

By addressing computational challenges and exploring hybrid forecasting techniques, future research can further enhance predictive models, ensuring businesses make accurate, data-driven decisions in an evolving marketplace.

8. Acknowledgments

The authors express their gratitude to the researchers and professionals in the fields of time-series forecasting, deep learning, and machine learning, whose contributions have laid the foundation for this study.

We acknowledge the data providers for granting access to the real-world e-commerce sales dataset, which was instrumental in conducting this research. Additionally, we appreciate the open-source communities behind Python libraries such as TensorFlow, XGBoost, Statsmodels, and Scikit-learn, whose tools facilitated the development and evaluation of our forecasting models.

Finally, we extend our sincere thanks to the reviewers and academic advisors for their valuable feedback, which has helped refine and enhance the quality of this research.

9. References

1. Benabbou, F., & Mouatassim, A. (2024). Intelligent Sales Forecasting System Using ARIMA, SARIMA, and XGBoost Models. IEEE International Conference on Artificial Intelligence and Computer Applications (ICAICA 2024), 567–572.
2. Hu, Y., Zhang, L., & Wang, S. (2024). Research on E-Commerce Retail Demand Forecasting Based on SARIMA Model and K-Means Clustering

ONLINE ISSN
3006-9726
PRINT ISSN
3006-9718
VOLUME . 4 ISSUE . 3 (2025)

Algorithm. American Journal of Science and Technology, 11(1), 1–10.

3. Jiotsop-Foze, L., Kim, J., & Park, S. (2024). Comparative Study of Machine Learning and Deep Learning Models for Time-Series Forecasting. Expert Systems with Applications, 235, 120134.
4. Li, X., Wang, Y., & Zhang, Z. (2024). Comparison of Deep Learning Algorithms for Retail Sales Forecasting. Transactions on Information Science, 5(3), 45–56.
5. Mahesar, A. W., et al. (2024). Machine Learning Model for Sales Forecasting Using XGBoost. IEEE Access, 12, 12345–12356.
6. Tang, Y. M., Ho, G. T. S., & Lau, Y. Y. (2025). Combinatorial Optimization Model for E-Commerce Retail Merchant Demand Forecasting Based on ARIMA and LSTM. Journal of Retail Analytics, 12(1), 45–60.
7. Yadav, S. (2022). An LSTM Approach for Sales Forecasting of Goods with Short-Term Demands in E-Commerce. Intelligent Information and Database Systems: 11th Asian Conference, ACIIDS 2019, 244–256. Springer International Publishing.
8. Ye, X., Lu, H., & Shen, Z. (2024). Enhancing Forecasting Accuracy Through Hybrid SARIMA-LSTM Models in E-Commerce. Journal of Data Science and Artificial Intelligence, 18(2), 78–92.