Proceedings of the International Conference on Inventive Computation Technologies (ICICT-2025)
IEEE Xplore Part Number: CFP25F70-ART; ISBN: 979-8-3315-1224-8

Demand Forecasting with XGBoost and LSTM: An Exploratory Analysis

R. Kanagaraj
School of Computer Science and Engineering
Vellore Institute of Technology
Vellore, Tamil Nadu, India
kanagaraj.r@vit.ac.in

Nuvvula Revanth
School of Computer Science and Engineering
Vellore Institute of Technology
Vellore, Tamil Nadu, India
revanth.nuvvula2021@vitstudent.ac.in

K. Vishnu
School of Computer Science and Engineering
Vellore Institute of Technology
Vellore, Tamil Nadu, India
karumanchi.vishnu2021@vitstudent.ac.in

Rahul Saxena
School of Computer Science and Engineering
Vellore Institute of Technology
Vellore, Tamil Nadu, India
rahul.saxena2021@vitstudent.ac.in

2025 International Conference on Inventive Computation Technologies (ICICT) | 979-8-3315-1224-8/25/$31.00 ©2025 IEEE | DOI: 10.1109/ICICT64420.2025.11005215

Abstract—**Predictive models are of great importance to any organization seeking to know and predict demand behavior. With an accurate demand forecast, any business can make plans to prepare, manage inventory, produce goods, and develop logistics systems that will be in place to meet the expected demand. This results in long-term re-investment and growth in operational efficiency and resource management. Effective demand forecasting makes it possible for firms to develop specific and affirmative strategies that reduce costs while improving customer satisfaction. There are two well-known models used in time series analysis: Long Short-Term Memory (LSTM) networks and eXtreme Gradient Boosting (XGBoost). The LSTM is a deep learning architecture that excels in learning and modeling long-range dependencies in a sequence, which is one of the main reasons why it is suitable for this work which deals with predicting demand. XGBoost is a modern and practical implementation of the gradient boosting algorithm that achieves great speed and performance in almost any predictive modeling task. By implementing both LSTM and XGBoost models and analyzing the demand forecasting performance of the two models from key assessment metrics. The two models are evaluated independently against vast quantities of Test Data, which gives a better understanding of the advantages and disadvantages of both approaches. This analysis then provides suggestions regarding the choice of models for forecasting demand within a business where there are several factors to consider, such as the need for explainability, the performance of the model, and the time taken to run the process through looking at results from the individual models.**

Keywords – LSTM, XGBoost, Neural Networks, ARIMA, Forecasting

I. INTRODUCTION

The process of predicting future consumption trends based on an evaluation of historical data is typically known as demand forecasting. Such an analysis is important for businesses as it enables them to meet customer demand while reducing inventory and other related costs by making their supply chain operations more effective. Demand forecasting is an independent method that relies on time series data, which contains trends, seasonality, variations, and other phenomena, which are all essential for effective, accurate, and valid forecasting. Numerous techniques, like a well-analyzed statistical approach to advanced Machine Learning Algorithms, have been employed in enhancing the precision of demand prediction [16]. Through Significant Post-Analysis Testing, the success and the failure of the forecasting models can be improved by fine-tuning the data, which provides ways of improving the prediction methods Without significantly changing the approach. A correct model allows one to make better decisions and increases the efficiency of work.

Demand forecasting has been a cornerstone in accurately predicting future consumer needs, helping businesses in making strategic decisions [17-18]. However, while traditional methods have been successful in the past, the evolution in consumer practices and the scale of innovation has led to a huge increase in the complexity of the marketplace, due to which post-forecast evaluations of traditional models often lead to gaps in the accuracy and scope of the predictions. As a solution to that, this paper introduces two different demand forecasting methods using both Long Short-Term Memory (LSTM) networks and eXtreme Gradient Boosting (XGBoost). Through various post-forecasting results and modifications, we aim to improve the accuracy of these two models over other conventional forecasting methods. The dataset will be split into training and testing sets, while pre-processing steps like data cleaning, transformation, and normalization ensure

Authorized licensed use limited to: University of Pittsburgh. Downloaded on July 07,2026 at 12:25:02 UTC from IEEE Xplore. Restrictions apply.

compatibility with the models, and post-forecasting modifications will lead to an increase in accuracy and scope.

The main objective of this project is to implement and evaluate two advanced machine learning models, namely XGBoost and LSTM [1]. The Main training and testing dataset will be of a time series format. By implementing both models, this project tries to showcase the improvement of advanced ML models over traditional models and forecasting practices. This project also conducts a detailed comparison between the two models (LSTM and XGboost) in order to provide deeper insights into their differences in accuracy and scope of the results, offering deeper insights into their applicability and ultimately guiding businesses in selecting the most suitable model for their needs.

II. Existing System

For decades, companies have used the same methods to estimate demand, which typically included a combination of various predictive modeling techniques that examined demand patterns using large amounts of historical data [2-3]. Accurate demand forecasting has always been critical for businesses to manage inventory and prepare for changes in customer demand while efficiently managing their supply chain operations, resulting in minimum downtime [14]. In the past, for example, despite advances in software applications, it was not uncommon for organizations to use the moving average method and exponential smoothing techniques to eliminate cyclic variations and prepare for short-term demand. Some of the more basic Machine Learning algorithms, such as regression analysis, were used, allowing companies to develop a demand equation that could account for seasonality and price but frequently failed to explain demand due to its nonlinear nature[5-7]. Among the various time series models, ARIMA rose to prominence due to the convenience of combining patterns and seasonality into data. However, these methods have limitations when dealing with complex patterns, relationships, or structures in data, and they typically necessitate extensive manual parameter tuning [8]. While they have their own pros and cons, older systems are generally less precise and accommodating than the newer machine-learning techniques available today[9-13]. With changes in demand forecasting, the demand for data-driven insights has increased, indicating the need for the implementation of advanced and flexible systems to produce accurate and consistent forecasts that reflect the ever-changing course of demand [15].

III. Dataset

In the following Study a CSV file consisting of time-series-based data is being used as the foundation for working on implementing and deploying both LSTM (Long Short-Term Memory) &XGBoost models. The Features and Targets in the data are Datetime, demand and isActive. The Datetime column is very important to create time-based features for temporal sequence modeling as it creates the timestamps hourly at which dependency changes, we will now be able to analyze when the demand occurs over period. The demand column contains the numeric demand values, showing demand volume or magnitude of observed demands at each hour. We will predict how much of the target variable should be produced/sold, to use our resources. Finally, the isActive column is binary (0, 1) and could allow to identify demand patterns or exceptions when looking at specific hours that are charged as actives vs non-active.

IV. Proposed System

The methodology of this study depends on the implementation of two machine learning models (i.e., XGBoost and Long Short-Term Memory, LSTM) for demand forecasting. Both models use the same time series dataset, and hence a comparison of their prediction abilities in similar data conditions can be well-supervised.

The LSTM model was implemented first to leverage the temporal dependencies in sequential data using the architecture of LSTM. The demand values were pre-processed to normalize the data and changed into a supervised learning format by creating lagged variables. This led to the creationof a model ofLSTM that could learn the demand pattern over a period. The model was trained and validated on a historical data split, with hyper-parameters tuned to minimize performance metrics like Mean Squared Error (MSE).

Similarly, the XGBoost Model was set up to predict demand. XgBoost model has the powerful ability to from complex interactions between features. The model was Fine-Tuned by using methods like grid–searching over parameters. Finally, we used cross-validation to verify the performance and prevent overfitting.

These two models were then compared in terms of how well each model predicted external data, computational efficiency, and applicability to seasonal variation.

4.1.Data Collection:

The first stage of this study consisted of compiling the cleaned dataset that pools hourly sales data of a particular product. TheResultant dataset was compiled by creating a seamless timeseries of sales activity which measured the volume and demand dynamics of the particular product by the hour. This dataset takes an approach more focused on temporal granularity, providing detailed and precise mappings to fluctuations in consumer demand on the hour. Such structuring of the training data ensures better training of models like XGBoost and LSTM, which can capture more well-defined temporal features for accurate predictions in real-world scenario.

Authorized licensed use limited to: University of Pittsburgh. Downloaded on July 07,2026 at 12:25:02 UTC from IEEE Xplore. Restrictions apply.

4.2Data Preprocessing:

Pre-processing is an important step in prepping the data for demand forecasting with both LSTM and XGBoost models. We will be formatting the demand data, which was fluctuating hourly into a structure so that the model could effectively capture temporal aspects and other important time series patterns.

4.2.1. Data Scaling:

In the LSTM model implementation,we have scaled demand and activity features to a range of 0–1 using Min-MaxScaler so that all the features are on the same scale. This scaling is important for the effective learning of neural networks and avoids problems that large numerical different numbers. Cause due to over-representation.

4.2.2. Feature Engineering:

TheXGBoostModel requires more features thanthe simple dataset to better understand the time-based patterns in the training dataset. We also created a number of time-series-related features like hour, day-of-week, month, and year, rolling means/standard deviations over windows of 24 hours (daily) along with those for the last week. The model can learn how sales change across the different months of a year and use this data to make more accurate future forecasts.

4.2.3. Activity Indicator:

The isActive metric in each model is used to denote that the hour of the day falls within active business hours (5 am — 1 pm) or inactive downtime. This binary variable provides the model with some context regarding demand, as it will allow differentiating between standard hours and non-standard hours that has a lower probability.

4.2.4. Train-Test Split and Reshaping:

The LSTM model used Historical data sequences as input because each hour sequence was mapped to the demand target of next hour. The XGBoost model used a similar split, focusing on reshaping without requiring sequence reshaping.

4.3. Implementation of LSTM and XGBoost Models:

Each stage of the implementation is structured in a way such that it will always contribute to Successful improvements in accuracy over multiple iterations in order to capture temporal patterns and forecast demand accurately.

4.3.1. Data Preprocessing and Feature Engineering:

For, LSTM the model will be utilizing Min-Max-Scalar to normalize both demand and is-Active between 0 and 1.

Feature Engineering (XGBoost): We first add time-series features including hour, day of week, and month as well as rolling statistics in the 24-hour and past 168 periods to the XGBoost model. Passing rolling averages and standard deviations of this data into the model helps it adjust to changing demand patterns over both seasonal and weekly time horizons.

4.3.2. Model Training

In LSTM model training, we utilized layers of dropout and batch normalization in order to prevent overfitting. It is trained on sequential data and uses a Timeseries Generator, which prepares sequences based on time. An early stopping mechanism is used to end the training early when there is no improvement in the validation performance and provide the best model weights and LSTM system architecture is illustrated in the Figure 2.

Training theXGBoost Model: Our Training Mechanism trains a baseline model using GridSearchCV in order to find the most optimal Boosts. They train the final model with the most optimal parameters. An early stopping Mechanism is added which is responsible for preventing over-fitting,which will be triggered when the improvement in accuracy is low.

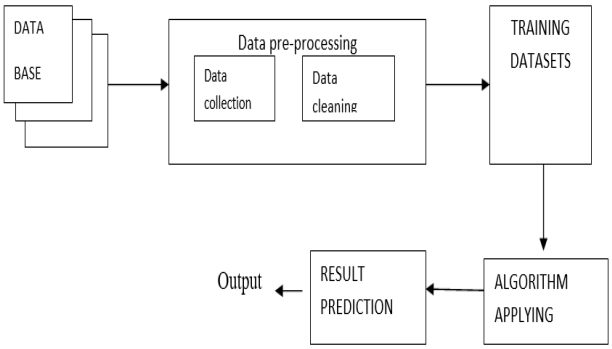


Fig.1.Flowchart of the Proposed System

4.3.3. Making Predictions and Evaluating Results:

LSTM predictions --- The predictions made by the model are now inverse-transformed to the original scale for evaluation. XGBoost predictions --- This model directly outputs the data required which will be stored for evaluation is illustrated in Figure 3.

4.3.4. Evaluation Metrics and Visualization:

Root Mean Squared Error (RMSE): For both models, the RMSE metric is calculated to evaluate the accuracy of predictions. This error metric reflects the average deviation of predictions from actual values.Visual Comparison: Finally, both models' results are visualized.

Authorized licensed use limited to: University of Pittsburgh. Downloaded on July 07,2026 at 12:25:02 UTC from IEEE Xplore. Restrictions apply.

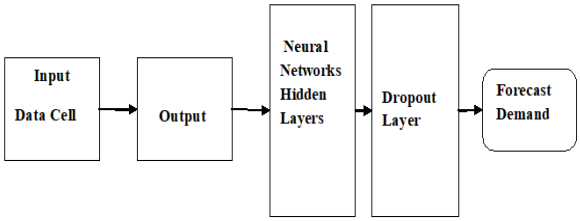


Fig.2. Basic LSTM System Architecture

V. Experimental Results

The performance of the models was evaluated through two primary metrics: Root Mean Squared Error (RMSE) and R^2 score. These metrics provide insights into the accuracy and explanatory power of each model on the test dataset, with lower RMSE and higher R^2 indicating better model performance.

5.1. Error Metrics and Model Evaluation

The RMSE and R^2 scores for the LSTM and XGBoost models were calculated as shown below. Figure 5 and Figure 6.LSTM Model: The LSTM model achieved an RMSE of 1568.7 and an R^2 score of 0.9422 on the test set, indicating its strength in capturing temporal dependencies over sequences. XGBoost Model: The XGBoost model reported an RMSE of 1419.21 and an R^2 score of 0.9516, benefiting from engineered features that enhance its prediction accuracy on high-dimensional data.The result is displayed in Figure

```
Root Mean Squared Error (RMSE): 1568.7459794662655
R² Score: 0.9422780305900494
```

Fig.5. LSTM Score

```
RMSE Score on Test set: 1419.21
R^2 Score on Test set: 0.9516
```

Fig.6.. XGBoost Score

5.2. Visualization of Actual vs. Predicted Demand

To further examine the models' performance, the predicted demand values were compared against the actual demand values over time. Figure 5.3 shows the line plot for the LSTM model, where actual demand values are plotted alongside predicted values, demonstrating the model's ability to follow demand trends over time. Figure 4 provides the corresponding visualization for the XGBoost model, illustrating the impact of feature engineering on demand prediction.

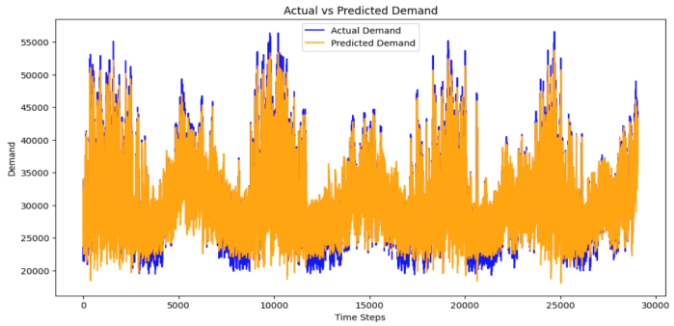


Fig.3. Actual vs. Predicted Demand for LSTM Model

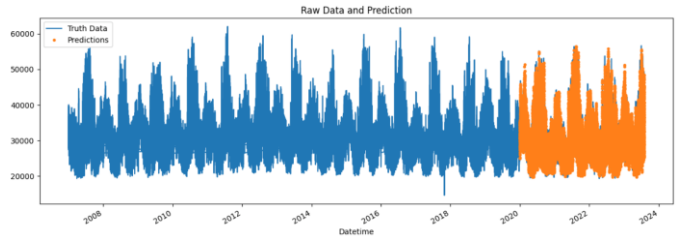


Fig.4.Actual vs. Predicted Demand for XGBoost Model

TABLE I.

COMPARISON OF ALGORITHMS

S.No	*Algorithm*	*Testing Accuracy*
1.	LSTM	0.942
2.	XGBoost	0.9516

VI. Conclusion

The rise of machine learning models like LSTM and XGBoost which has taken demand forecasts from good to already great, as it provided a step change performance over conventional methods for organizations that relied on manual algorithms suited to stationary data streams but struggled with larger datasets or those where intermittent seasonality was present. As opposed to the antiquated methodologies that operate under linear assumptions, this approach enables modelling based on intricate patterns inherent in data and hence, is more accurate and can be used universally. For example, one type of neural network; LSTM is excellent to analyze long time dependencies which are critical for forecasting scenarios following a seasonality trend. However, they require a lot of data and computing power, which can be difficult.XGBoost, by contrast, is built for handling large, feature-rich datasets and excels in situations where short-term accuracy is key. It works well with engineered features and trains quickly, which makes it practical and efficient, though it's less effective for sequential data unless features are specifically designed for

Authorized licensed use limited to: University of Pittsburgh. Downloaded on July 07,2026 at 12:25:02 UTC from IEEE Xplore. Restrictions apply.

that purpose.In short, choosing between LSTM and XGBoost depends on your needs: LSTM is better for long-term trend analysis, while XGBoost is great for accurate short-term forecasts. Both models outperform traditional methods, helping businesses make smarter, data-driven decisions.

REFERENCES

[1] Gers, F. A., Schraudolph, N. N., & Schmidhuber, J. "Learning precise timing with LSTM recurrent networks." *Journal of Machine Learning Research*, 2002.

[2] Sundermeyer, M., Schlüter, R., & Ney, H. "LSTM neural networks for language modeling." *Proceedings of Interspeech*, 2012.

[3] Pulver, A., & Lyu, S. "LSTM with working memory." *2017 International Joint Conference on Neural Networks (IJCNN)*, IEEE, 2017.

[4] Luo, Y., Ren, J., Wang, Z., Sun, W., & Pan, J. "LSTM Pose Machines." *Proceedings of CVPR*, 2018.

[5] El-Khamy, M., & Lee, J. "Residual LSTM: Design of a deep recurrent architecture for distant speech recognition." *arXiv preprint*, 2017.

[6] Smagulova, K., & James, A. P. "A survey on LSTM memristive neural network architectures and applications." *The European Physical Journal Special Topics*, 2019.

[7] Landi, F., Baraldi, L., Cornia, M., & Cucchiara, R. "Working memory connections for LSTM." *Neural Networks*, 2021.

[8] Sherstinsky, A. "Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network." *Physica D: Nonlinear Phenomena*, 2020.

[9] Xu, J., He, J., Gu, J., Wu, H., Wang, L., & Zhu, Y. "Financial time series prediction based on XGBoost and generative adversarial networks." *Journal of Circuits, Systems, and Signal Processing*, 2022.

[10] Liwei, T., Li, F., Yu, S., & Yuankai, G. "Forecast of LSTM-XGBoost in stock price based on Bayesian optimization." *Intelligent Automation & Soft Computing*, 2021.

[11] . Kanagaraj, R., Elakiya, E., Kumaresan, A., ... Suganya, J., Shruthi, K., Predictive Classification Model of Crop Yield Data Using Artificial Neural Network, Proceedings of the 5th International Conference on Inventive Research in Computing Applications, ICIRCA 2023, 2023, pp. 747–751

[12] Tan, B., Gan, Z., & Wu, Y. "The measurement and early warning of daily financial stability index based on XGBoost and SHAP: Evidence from China." *Expert Systems with Applications*, 2023.

[13] Petropoulos, A., & Siakoulis, V. "Can central bank speeches predict financial market turbulence? Evidence from an adaptive NLP sentiment index analysis using XGBoost machine learning." *Central Bank Review*, 2021.

[14] Qin, R. "The construction of corporate financial management risk model based on XGBoost algorithm." *Journal of Mathematics*, 2022.

[15] Liu, Y. "Design of XGBoost prediction model for financial operation fraud of listed companies." *International Journal of System Assurance Engineering and Management*, 2023.

[16] E. Elakiya, R. Kanagaraj, T. Paturu, S. D. Nivethika and R. S.Kanmani, "Text Feedback Classification using Machine LearningTechniques," 2023 2nd International Conference on Edge ComputingandApplications(ICECAA),Namakkal,India,2023,pp.933-939,doi:10.1109/ICECAA58104.2023.10212398

[17] Vuong, P. H., Dat, T. T., & Mai, T. K. "Stock-price forecasting based on XGBoost and LSTM." *Systems Science & Engineering*, 2022.

[18] Neves, R. F. "Combining principal component analysis, discrete wavelet transform and XGBoost to trade in the financial markets." *Expert Systems with Applications*, 2019.

Authorized licensed use limited to: University of Pittsburgh. Downloaded on July 07,2026 at 12:25:02 UTC from IEEE Xplore. Restrictions apply.