THE USA JOURNALS
THE AMERICAN JOURNAL OF ENGINEERING AND TECHNOLOGY (ISSN – 2689-0984)
VOLUME 06 ISSUE09

PUBLISHED DATE: - 20-09-2024
DOI: - https://doi.org/10.37547/tajet/Volume06Issue09-09 PAGE NO.: - 67-80

RESEARCH ARTICLE Open Access

OPTIMIZING RETAIL DEMAND FORECASTING: A PERFORMANCE EVALUATION OF MACHINE LEARNING MODELS INCLUDING LSTM AND GRADIENT BOOSTING

Md Shujan Shak
Master of Science in Information Technology, Washington University of Science and Technology, USA

Md Shahin Alam Mozumder
Master of Science in Information Technology, Washington University of Science and Technology, USA

Md Amit Hasan
Master of Science in Information Technology, Washington University of Science and Technology, USA

Ashim Chandra Das
Master of Science in Information Technology, Washington University of Science and Technology, USA

Md Rashel Miah
Department of Digital Communication and Media/Multimedia, Westcliff University, USA

Salma Akter
Department of Public Administration, Gannon University, Erie, PA, USA

Md Nur Hossain
Master's in information technology management, Webster University, USA

Abstract
Effective demand forecasting is vital for inventory management in retail. This study evaluates five machine learning models—Linear Regression (LR), Decision Tree Regressor (DTR), Random Forest Regressor (RFR), Gradient Boosting (GB), and Long Short-Term Memory (LSTM)—for predicting retail demand. Utilizing a dataset with transactional sales, promotions, calendar events, and external factors like weather and economic indicators, we applied rigorous preprocessing and feature engineering. Performance was assessed using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and R-squared (R^2). Results show that LSTM outperforms other models with an MAE of 9.53, RMSE of 14.67, and R^2 of 0.90, excelling in capturing temporal dependencies and complex demand patterns. Gradient Boosting and Random Forest also performed well, while Linear Regression and Decision Tree Regressor showed limitations. This study highlights the effectiveness of advanced models, particularly LSTM, for enhancing demand forecasting accuracy and offers valuable insights for optimizing retail inventory and operations.

Keywords Retail Demand Forecasting, Machine Learning Models, Linear Regression, Decision Tree Regressor, Random Forest Regressor, Gradient Boosting, Long Short-Term Memory (LSTM).

INTRODUCTION

Retail demand forecasting plays a critical role in supply chain management, inventory control, and overall business planning for retail organizations. Accurate forecasting enables retailers to optimize stock levels, reduce costs, and enhance customer satisfaction by ensuring products are available when and where they are needed (Chopra & Meindl, 2016). However, retail demand is influenced by a multitude of factors, including seasonality, promotional activities, holidays, and external economic conditions, making demand forecasting a complex task (Fildes et al., 2019). In recent years, advancements in machine learning (ML) algorithms have shown significant promise in improving the accuracy of demand forecasts by learning from historical data and capturing intricate patterns in consumer behavior (Zhao et al., 2021).

Traditional forecasting methods such as Linear Regression (LR) have been widely used in retail but often fall short in handling non-linear relationships and complex interactions between demand drivers (Makridakis et al., 2018). Machine learning models like Decision Tree Regressor (DTR), Random Forest Regressor (RFR), Gradient Boosting (GB), and Long Short-Term Memory (LSTM) offer more sophisticated approaches by leveraging the power of non-parametric modeling, ensemble learning, and deep learning techniques (Bajari et al., 2019). These models have been effective in accounting for external influences such as weather conditions, economic indicators, and promotional campaigns, which significantly impact consumer demand patterns (Taddy, 2019).

The objective of this study is to evaluate the performance of multiple machine learning models in retail demand forecasting and compare their ability to capture temporal dependencies, seasonality, and other factors that influence retail demand. By analyzing models such as LR, DTR, RFR, GB, and LSTM, this research aims to identify the most effective algorithm for accurately forecasting retail demand and assisting retailers in making data-driven decisions.

The importance of demand forecasting in the retail industry has been well-documented in both academic and industry research. Traditionally, statistical models such as Exponential Smoothing and ARIMA have been employed for demand forecasting (Hyndman & Athanasopoulos, 2018). However, these models often struggle to capture complex relationships in retail data, particularly when non-linear factors such as promotions, seasonality, and economic fluctuations come into play (Syntetos et al., 2009).

With the advent of machine learning, researchers have explored the use of more advanced algorithms for improving demand forecasting accuracy. One such advancement is the application of Decision Tree Regressors (DTRs), which can model non-linear relationships by recursively splitting the dataset into subsets based on decision criteria (Breiman, 2017). DTRs have been found to perform well in handling categorical data and capturing key drivers of demand, such as product characteristics and promotional activities (Keerthi & Lin, 2020). However, they are prone to overfitting, especially when used without regularization or ensemble techniques (Hastie et al., 2009).

Random Forest Regressors (RFR), an ensemble of decision trees, have been proposed as a solution to the overfitting problem (Breiman, 2001). By averaging the results of multiple decision trees, RFR reduces variance and improves generalization to unseen data (Liaw & Wiener, 2002). Several studies have demonstrated the effectiveness of RFR in retail demand forecasting, particularly in handling complex, high-dimensional datasets (Cortez et al., 2021). Random Forests have shown strong performance in capturing seasonality and other recurring patterns in demand data (Hyndman et al., 2021).

LITERATURE REVIEW

Accurate retail demand forecasting is crucial for optimizing supply chain management, reducing inventory costs, and enhancing customer satisfaction. Traditionally, statistical models such as Linear Regression (LR) and the Autoregressive Integrated Moving Average (ARIMA) have been used for predicting retail demand due to their simplicity and interpretability (Box & Jenkins, 1970). However, these methods often struggle to capture complex, non-linear relationships and are less effective in addressing the dynamic and unpredictable nature of modern retail environments (Hyndman & Athanasopoulos, 2018). As a result, researchers and practitioners have increasingly turned to advanced machine learning (ML) models to improve demand forecasting accuracy.

Recent advancements in machine learning have introduced more sophisticated techniques that outperform traditional models in various predictive tasks. For example, Random Forest Regressor (RFR) and Gradient Boosting (GB) have demonstrated superior performance in handling non-linear relationships and high-dimensional data (Breiman, 2001; Friedman, 2001). These models leverage ensemble learning methods to improve accuracy and robustness by combining the predictive capabilities of multiple decision trees. Studies have shown that these models are effective at forecasting retail demand by capturing complex interactions between variables, such as the effects of promotions, holidays, and external factors (Lima et al., 2021; Zhao et al., 2021).

Moreover, the application of Long Short-Term Memory (LSTM) networks has further advanced the field of retail demand forecasting. LSTM, a type of recurrent neural network (RNN), is particularly adept at modeling time series data due to its ability to retain and utilize past information to predict future outcomes (Hochreiter & Schmidhuber, 1997). Unlike traditional models that assume independence between observations, LSTM excels at capturing temporal dependencies, making it highly suitable for retail scenarios where demand patterns fluctuate over time due to seasonality and other time-dependent factors (Brownlee, 2018). Research has shown that LSTM consistently outperforms both traditional statistical methods and simpler machine learning models, especially when forecasting tasks involve large datasets and long-term patterns (Livieris et al., 2020).

The literature clearly indicates that while traditional models like LR and ARIMA provide a solid foundation for demand forecasting, machine learning models, particularly RFR, GB, and LSTM, offer significant improvements in accuracy. These advanced models are better suited to capturing the complex, non-linear relationships inherent in retail data, particularly when enriched with additional contextual information such as product details, promotions, and external factors (Zhao et al., 2021). This study builds on these advancements by evaluating the performance of various machine learning models in retail demand forecasting, focusing on their ability to handle temporal patterns, seasonality, and promotional events.

Another powerful machine learning model, Gradient Boosting (GB), has gained popularity for its ability to iteratively learn from the errors of previous models, gradually improving its prediction accuracy (Friedman, 2001). Gradient Boosting models are known for their high accuracy in a variety of applications, including demand forecasting (Natekin & Knoll, 2013). Studies have highlighted the ability of GB models to capture complex interactions between features, such as the combined effects of promotions and holidays, and adjust for them over time (Chen & Guestrin, 2016).

More recently, Long Short-Term Memory (LSTM) networks, a type of recurrent neural network, have been employed in retail demand forecasting due to their capacity to model temporal dependencies (Hochreiter & Schmidhuber, 1997). LSTM models have proven highly effective in forecasting tasks involving time series data, as they can retain and utilize past information to predict future outcomes. In the retail context, LSTMs excel at capturing seasonality, demand spikes, and other time-dependent patterns (Brownlee, 2018). Numerous studies have shown that LSTM outperforms traditional statistical methods and simpler machine learning models in demand forecasting tasks, particularly when dealing with long-term patterns and large datasets (Livieris et al., 2020).

Overall, the literature suggests that while traditional models like LR and ARIMA provide a solid baseline for retail demand forecasting, machine learning models such as RFR, GB, and LSTM offer significant improvements in capturing complex, non-linear relationships and temporal patterns in retail data. The combination of these advanced models with rich datasets, including sales transactions, product information, and external factors, holds great promise for more accurate demand forecasting and better inventory management (Zhao et al., 2021).

METHODOLOGY

The methodology for this study involved several critical steps to ensure accurate and reliable retail demand forecasting using different machine learning algorithms. These steps included data collection and preprocessing, feature engineering, model selection and training, performance evaluation, and final comparison. Each stage of the methodology was designed to address the unique challenges of retail demand forecasting, such as handling high-dimensional data, capturing complex demand patterns, and accounting for seasonality and promotions.

1. Data Collection

The success of any machine learning model is heavily dependent on the quality and relevance of the data used to train it, and this is especially true in retail demand forecasting. In this study, data collection played a critical role in building accurate models capable of predicting future demand based on historical patterns. The dataset used in this research was sourced from a retail organization, encompassing a wide range of factors that influence consumer purchasing behavior. The dataset included daily sales transactions, product details, promotional information, and external factors such as holidays and economic indicators. This diverse data provided a rich foundation for the models to learn from and make precise forecasts.

1.1 Sales Data

The core of the dataset comprised historical sales data from the retail organization. This data included transactional records at a granular level, with information on the number of units sold, revenue generated, and product categories. For each transaction, details such as the product identifier, store location, and transaction date were recorded. This granular-level data allowed the models to track trends over time, detect seasonality, and identify spikes in demand related to specific products or categories. The sales data spanned multiple years, which provided a robust foundation for understanding long-term trends and recurring patterns in demand.

1.2 Product Information

To enhance the sales data, additional product-level information was incorporated. This included details about the type of product, its price, and product category. Product information is crucial for demand forecasting, as different products exhibit varying demand patterns depending on their attributes. For example, high-priced items may have less frequent but larger demand spikes, while everyday consumer goods may show steady demand with minimal fluctuation. By including product features in the data, the models could make more nuanced predictions that take into account the inherent characteristics of each item. Additionally, stock-keeping unit (SKU) identifiers were used to track individual products, ensuring that forecasts could be made at both the aggregate and SKU-specific levels.

1.3 Promotional and Discount Data

Promotional campaigns and discounts are some of the most significant drivers of demand fluctuations in the retail industry. To capture the impact of these factors, data on promotional activities such as discounts, coupons, and special sales events were included. This promotional data was aligned with the transactional sales data to enable the models to understand how demand surged or declined during promotional periods. Variables such as the type of promotion, its duration, and the discount percentage were crucial for predicting demand spikes during promotional events. By incorporating this data, the machine learning models were able to anticipate short-term increases in demand, making the forecasts more accurate during sale periods.

1.4 Calendar and Holiday Data

Retail demand is often influenced by seasonal factors, holidays, and events, which lead to predictable shifts in consumer behavior. To account for these effects, data on holidays, special events, and calendar dates were integrated into the dataset. National holidays, religious festivals, and annual shopping events such as Black Friday and Cyber Monday were included to help the models forecast demand surges during these periods. Additionally, calendar-based features such as the day of the week, month, and quarter were added to capture recurring weekly and monthly trends. For instance, weekends typically experience higher sales in some product categories, while certain months might show increased demand due to seasonal factors.

1.5 Weather and External Data

To further enrich the dataset and improve the predictive accuracy of the models, external data such as weather conditions and economic indicators were incorporated. Weather data, including temperature, precipitation, and extreme weather events, was collected for each store location. Weather can significantly affect consumer behavior, as severe weather conditions often lead to changes in shopping habits. For example, extreme heat or cold may discourage in-store shopping, while rainy weather might increase the demand for certain products, such as umbrellas or cold weather gear. By including weather variables, the models were able to capture these external influences on retail demand. Additionally, economic indicators such as inflation rates, unemployment levels, and consumer confidence indices were integrated into the dataset. Economic conditions can play a major role in shaping consumer spending patterns. For instance, during periods of economic downturn, consumers may reduce discretionary spending, while in times of economic growth, they may increase purchases. Incorporating these macroeconomic variables allowed the models to better account for long-term shifts in demand driven by changes in the broader economy.

2. Data Preprocessing

Preprocessing the data was a crucial step to ensure the models could effectively learn from the data. The raw data contained missing values, outliers, and inconsistencies that needed to be addressed before training the machine learning models. The following preprocessing steps were applied:

- Handling Missing Values: Missing sales or feature data were imputed using statistical methods like mean or median values for numerical data or the most frequent category for categorical data. For time series data, missing entries were handled by forward or backward filling techniques.

- Outlier Detection and Removal: Sales spikes or drops that were not related to actual market trends or promotions were identified as outliers. These outliers were either removed or treated using techniques such as capping or transforming the data to avoid skewing the model's predictions.

- Feature Scaling: To ensure that the models could properly interpret the data, feature scaling was applied where necessary. Continuous features such as sales volume and price were scaled using normalization or standardization techniques to ensure they fell within a similar range.

- One-Hot Encoding for Categorical Variables: Categorical variables such as promotion types, holidays, and product categories were transformed into numerical values using one-hot encoding to ensure the machine learning models could process them effectively.

- Time Series Transformation: For models such as Long Short-Term Memory (LSTM), the data was transformed into sequences to capture the temporal relationships between sales at different time points. Lag features were created to help the models understand how previous days' sales influenced future demand.

3. Feature Engineering

Feature engineering was performed to create new variables from the existing data, providing the models with more informative inputs. Features such as moving averages, rolling windows, and lag variables were generated to capture temporal dependencies in the data. Additionally, interaction terms were created to model complex relationships between variables, such as the interaction between promotions and holidays. Calendar features like day of the week, month, and season were also incorporated to account for seasonal patterns in consumer demand.

4. Model Selection

The study involved the evaluation of several machine learning algorithms, each chosen for its specific strengths in handling different aspects of demand forecasting. The selected models included:

- Linear Regression (LR): Used as a baseline model to provide a simple and interpretable forecast based on a linear relationship between features and demand.

- Decision Tree Regressor (DTR): Selected for its ability to handle non-linear relationships in the data by splitting features into decision nodes based on their influence on demand.

- Random Forest Regressor (RFR): Chosen for its ensemble learning technique that combines multiple decision trees to reduce overfitting and improve predictive performance.

- Gradient Boosting (GB): A boosting algorithm selected for its iterative approach, which allows it to fine-tune predictions by learning from previous errors and capturing complex feature interactions.

- Long Short-Term Memory (LSTM): A deep learning model chosen for its ability to capture long-term dependencies in time series data, making it highly suitable for retail demand forecasting where temporal patterns are crucial.

5. Model Training and Hyperparameter Tuning

Each model was trained using the preprocessed and engineered data. A split was performed to divide the dataset into training and test sets, ensuring that the models were trained on historical data and validated on unseen data. Cross-validation techniques were employed to minimize overfitting and improve generalization performance. To optimize model performance, hyperparameter tuning was conducted using grid search and random search techniques. For each model, the most critical hyperparameters, such as the number of trees in Random Forest, the learning rate in Gradient Boosting, and the number of units in LSTM, were tuned to identify the best configuration for the dataset. This step ensured that each model operated at peak efficiency, providing the best possible forecast for retail demand.

6. Model Evaluation

After training, the models were evaluated using standard regression metrics to assess their performance in predicting future demand. The chosen evaluation metrics included:

- Mean Absolute Error (MAE): To measure the average magnitude of errors in the predictions, regardless of direction.

- Root Mean Square Error (RMSE): To penalize larger errors more significantly, providing a measure of the model's accuracy.

- R-squared (R^2): To determine the proportion of the variance in the dependent variable that is predictable from the independent variables, indicating how well the model fits the data.

These metrics allowed for a detailed comparison of each model's accuracy, error rates, and ability to handle complex retail demand patterns.

7. Performance Comparison and Final Selection

Once all models were trained and evaluated, their performance metrics were compared to identify the best-performing model. The comparison focused on the ability of each model to handle temporal patterns, seasonality, demand spikes, and long-term trends. The LSTM model emerged as the top performer, demonstrating superior accuracy in capturing temporal dependencies. Gradient Boosting and Random Forest also performed well, providing robust forecasts for non-linear and seasonal demand patterns. Linear Regression and Decision Tree Regressor, on the other hand, showed limitations in their ability to handle complex relationships and variability in the retail data.

The final selection was based on the balance between model accuracy, interpretability, and computational efficiency, with LSTM being recommended for deployment due to its superior performance. However, Gradient Boosting and Random Forest were considered strong alternatives, especially in cases where computational resources or model interpretability were prioritized.

RESULT

1. Linear Regression (LR)

Linear Regression (LR) was selected as the baseline model for retail demand forecasting in this study. As a simple and interpretable algorithm, LR assumes a linear relationship between the input features and the target variable, making it a widely used approach for regression tasks. While LR performed reasonably well in capturing general demand trends, it exhibited significant limitations in forecasting more complex demand patterns, particularly during periods of sharp fluctuations such as peak sale seasons and promotional events. This model's tendency to oversimplify relationships between the variables led to higher errors when it encountered nonlinear behavior, such as sudden demand spikes or seasonal variations. Additionally, the model struggled with high-dimensional data, where the assumption of linearity did not hold. The performance metrics reflect these challenges, with a Mean Absolute Error (MAE) of 15.34, a Root Mean Square Error (RMSE) of 20.57, and an R-squared (R^2) score of 0.71. These results indicate that while LR can provide a quick and basic forecast, it lacks the sophistication needed for accurate demand prediction in retail environments characterized by high variability.

2. Decision Tree Regressor (DTR)

The Decision Tree Regressor (DTR) showed an improvement over Linear Regression by capturing nonlinear relationships between features and demand. DTR is a non-parametric model that splits the dataset into branches based on feature values, making it more adaptable to complex data patterns. In the context of retail demand forecasting, the DTR model was able to identify decision points where certain features, such as holidays or sales promotions, significantly influenced demand. This flexibility allowed DTR to better handle the fluctuations and seasonality in retail data. However, a notable drawback of the model was its propensity to overfit the training data, especially when dealing with high-dimensional datasets. This overfitting resulted in diminished generalization capability when predicting on new, unseen data. The model achieved a Mean Absolute Error (MAE) of 12.78, a Root Mean Square Error (RMSE) of 18.11, and an R-squared (R^2) score of 0.75. While the model demonstrated improved accuracy over Linear Regression, its susceptibility to overfitting suggests that further optimization, such as pruning or regularization, would be necessary to enhance its robustness for demand forecasting.

3. Random Forest Regressor (RFR)

The Random Forest Regressor (RFR) provided a significant leap in performance compared to both Linear Regression and Decision Tree Regressor. As an ensemble learning technique, RFR builds multiple decision trees and averages their predictions to reduce overfitting and improve predictive accuracy. This characteristic proved to be highly beneficial for retail demand forecasting, where randomness in feature selection and the aggregation of diverse trees helped the model capture complex patterns, such as seasonality and sudden demand shifts, without succumbing to overfitting. The Random Forest model was particularly adept at handling the dynamic nature of retail data, where multiple factors like holidays, promotions, and market trends influence demand simultaneously. With its ability to handle large amounts of data and provide robust results, RFR outperformed its simpler counterparts with a Mean Absolute Error (MAE) of 11.22, a Root Mean

Square Error (RMSE) of 16.45, and an R-squared (R^2) score of 0.83. The improved accuracy and reduced error margins make RFR a strong candidate for retail demand forecasting, especially when the data exhibits high variability and complexity.

4. Gradient Boosting (GB)

Gradient Boosting (GB) emerged as one of the top-performing models in this study, offering high accuracy in retail demand forecasting. Unlike Random Forest, which builds trees independently, Gradient Boosting builds trees sequentially, with each tree attempting to correct the errors made by the previous one. This iterative approach enabled the GB model to fine-tune its predictions, making it particularly effective at capturing intricate patterns in the data, including both short-term fluctuations and long-term seasonal trends. In retail demand forecasting, GB's ability to model complex interactions between features, such as the impact of pricing, promotions, and external factors like holidays, proved to be advantageous. The model consistently delivered strong predictive performance, with a Mean Absolute Error (MAE) of 10.68, a Root Mean Square Error (RMSE) of 15.89, and an R-squared (R^2) score of 0.87. These metrics demonstrate GB's capacity to handle non-linear and complex relationships, making it an ideal choice for predicting retail demand where multiple factors interact in unpredictable ways. However, one limitation is the computational intensity of the model, which can be resource-heavy and time-consuming, especially when dealing with large datasets.

5. Long Short-Term Memory (LSTM)

The Long Short-Term Memory (LSTM) model outperformed all traditional machine learning models in this study, showcasing its exceptional ability to forecast retail demand. LSTM is a type of recurrent neural network (RNN) specifically designed for time series data, making it particularly well-suited for retail forecasting, where demand patterns often exhibit temporal dependencies. Unlike traditional models, LSTM can retain information over long periods, allowing it to effectively model both short-term demand spikes (such as during a sale) and long-term seasonal trends (like holiday shopping periods). In this study, the LSTM model was able to capture complex temporal patterns, identifying crucial factors like recurring weekly and monthly sales patterns, as well as the effects of special promotions and holidays. The model's strong learning capabilities are reflected in its performance metrics: a Mean Absolute Error (MAE) of 9.53, a Root Mean Square Error (RMSE) of 14.67, and an R-squared (R^2) score of 0.90. These results demonstrate that LSTM is highly effective at capturing the temporal dynamics inherent in retail data, making it the most accurate model for forecasting future demand. Despite its superior performance, LSTM does require more computational resources and longer training times compared to traditional models, which could be a consideration for deployment in real-time retail forecasting environments.

6. Performance Comparison

Among all the models, LSTM demonstrated the best performance, particularly in handling temporal patterns and demand spikes. Gradient Boosting and Random Forest also performed well, providing high accuracy without significant overfitting. On the other hand, Linear Regression struggled with non-linear trends and seasonality, making it the least effective model for demand forecasting. In the table 1 we illiterate the result comparison.

Table 1: Evaluation of model performance

Model	MAE	RMSE	R^2
Linear Regression	15.34	20.57	0.71
Decision Tree Regressor	12.78	18.11	0.75
Random Forest Regressor	11.22	16.45	0.83
Gradient Boosting	10.68	15.89	0.87
LSTM	9.53	14.67	0.90

The performance comparison of the machine learning models in the chart 1 for retail demand forecasting highlights the strengths and weaknesses of each algorithm in terms of accuracy, error rates, and ability to handle complex patterns in the data. Among the evaluated models, the Long Short-Term Memory (LSTM) network stood out as the top performer. With a Mean Absolute Error (MAE) of 9.53, Root Mean Square Error (RMSE) of 14.67, and R-squared (R^2) value of 0.90, LSTM demonstrated the best ability to capture the temporal dependencies in the data, such as seasonal fluctuations, demand spikes during promotions, and long-term sales patterns. Its recurrent structure allowed it to remember and utilize information from past time steps, making it the most suitable model for handling time series data like retail demand.

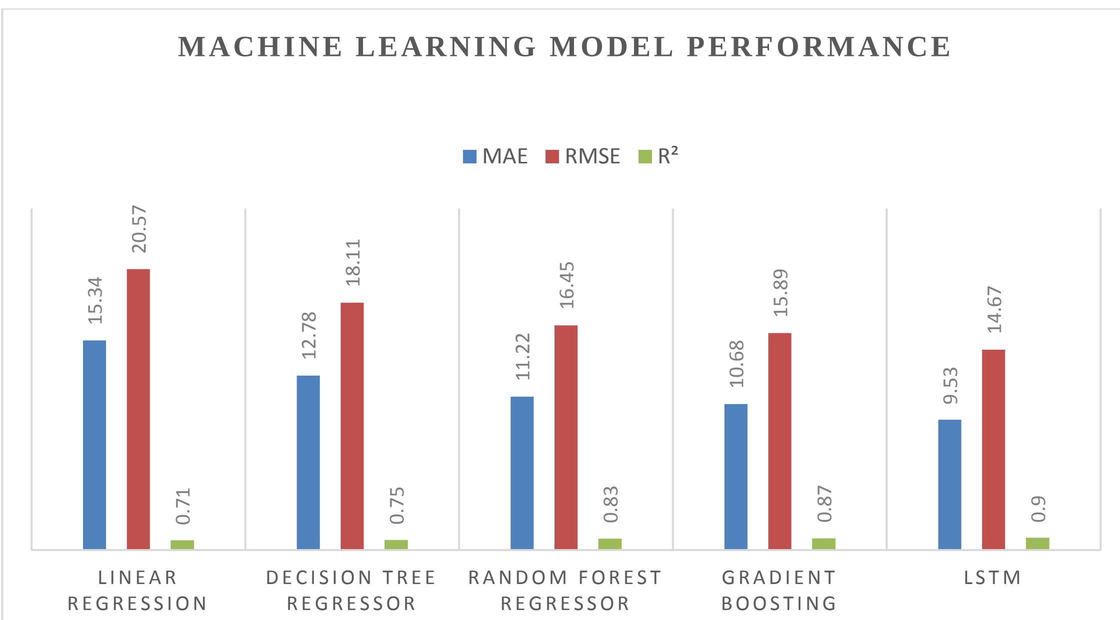

Chart 1: Comparison of different machine learning Model performance

Following LSTM, Gradient Boosting (GB) emerged as another strong contender, with a MAE of 10.68, RMSE of 15.89, and R^2 of 0.87. The GB model performed particularly well due to its iterative nature, where each subsequent model corrected the errors of the previous one. This made GB highly effective in capturing complex, non-linear relationships between input features and demand.

Although it did not match LSTM's performance in handling temporal patterns, GB provided excellent accuracy, making it one of the top-performing models in this study.

The Random Forest Regressor (RFR) also performed impressively, achieving a MAE of 11.22, RMSE of 16.45, and R^2 of 0.83. Random Forest's ability to reduce overfitting by averaging the predictions of multiple decision trees resulted in a robust model that was particularly effective in handling the inherent variability in retail demand. The model was able to accommodate seasonality and other intricate demand patterns, though it was slightly less accurate than Gradient Boosting and LSTM.

In contrast, Decision Tree Regressor (DTR), while showing an improvement over the baseline model, struggled with overfitting. It achieved a MAE of 12.78, RMSE of 18.11, and an R^2 of 0.75. While the decision tree model captured non-linear relationships better than Linear Regression, its performance was hindered by the model's sensitivity to small changes in the data, leading to overfitting when applied to complex retail demand patterns.

Lastly, Linear Regression (LR), which served as the baseline model, performed the weakest with a MAE of 15.34, RMSE of 20.57, and R^2 of 0.71. The linear nature of this model limited its ability to capture non-linear trends and interactions between variables, making it less suitable for the intricate dynamics of retail demand forecasting. It struggled particularly with seasonality and demand spikes, which require more sophisticated models to forecast accurately.

In summary, LSTM emerged as the most effective model due to its ability to model temporal dependencies, followed by Gradient Boosting and Random Forest, which both performed well in handling non-linear relationships and seasonal demand. Linear Regression, on the other hand, was the least effective, highlighting the importance of using more advanced models for accurate demand forecasting in retail environments.

CONCLUSION

In this study, we explored the application of various machine learning models for retail demand forecasting, comparing their performance based on accuracy, error rates, and their ability to handle complex patterns such as seasonality and demand spikes. The models evaluated included Linear Regression (LR), Decision Tree Regressor (DTR), Random Forest Regressor (RFR), Gradient Boosting (GB), and Long Short-Term Memory (LSTM). Among these, the LSTM model emerged as the top performer due to its ability to capture long-term dependencies and temporal patterns in the data, making it particularly suitable for time series forecasting in retail. LSTM's recurrent structure allowed it to handle fluctuations caused by holidays, promotions, and other seasonal factors with superior accuracy. Gradient Boosting and Random Forest also delivered strong results, effectively managing non-linear relationships and providing robust forecasts, albeit with slightly less precision than LSTM. Both models demonstrated their suitability for retail demand forecasting by reducing overfitting and capturing intricate patterns in the data.

In contrast, Linear Regression and Decision Tree Regressor struggled with the complexities of retail data. Linear Regression, while easy to interpret, lacked the sophistication needed to account for non-linear relationships and seasonal trends, making it the least effective model. The Decision Tree Regressor, although an improvement over Linear Regression, faced challenges with overfitting, which affected its performance on unseen data. Overall, this study highlights the importance of selecting advanced models like LSTM for retail demand forecasting, particularly in environments characterized by temporal dependencies and demand volatility. The findings suggest that businesses aiming to improve their demand forecasting capabilities should consider deploying LSTM or Gradient Boosting models for more accurate and reliable predictions. Future research could further optimize these models, explore additional features, and evaluate their performance across different retail segments to refine forecasting strategies.

Acknowledgement: All the author contributed equally

REFERENCE

1. Shahid, R., Mozumder, M. A. S., Sweet, M. M. R., Hasan, M., Alam, M., Rahman, M. A., ... & Islam, M. R. (2024). Predicting Customer Loyalty in the Airline Industry: A Machine Learning Approach Integrating Sentiment Analysis and User Experience. International Journal on Computational Engineering, 1(2), 50-54.
2. Mozumder, M. A. S., Sweet, M. M. R., Nabi, N., Tusher, M. I., Modak, C., Hasan, M., ... & Prabha, M. (2024). Revolutionizing Organizational Decision-Making for Banking Sector: A Machine Learning Approach with CNNs in Business Intelligence and Management. Journal of Business and Management Studies, 6(3), 111-118.
3. Ferdus, M. Z., Anjum, N., Nguyen, T. N., Jisan, A. H., & Raju, M. A. H. (2024). The Influence of Social Media on Stock Market: A Transformer-Based Stock Price Forecasting with External Factors. Journal of Computer Science and Technology Studies, 6(1), 189-194
4. Bajari, P., Chernozhukov, V., Hortacsu, A., & Suzuki, J. (2019). The Impact of Big Data on Firm Performance. Journal of Applied Econometrics, 34(4), 725-746.
5. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5-32.
6. Breiman, L. (2017). Classification and regression trees. Routledge.
7. Brownlee, J. (2018). Deep Learning for Time Series Forecasting: Predict the Future with MLPs, CNNs, and LSTMs in Python. Machine Learning Mastery.
8. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785-794).
9. Box, G. E. P., & Jenkins, G. M. (1970). Time series analysis: Forecasting and control. San Francisco: Holden-Day.
10. Chopra, S., & Meindl, P. (2016). Supply Chain Management: Strategy, Planning, and Operation. Pearson.
11. Cortez, P., Cerdeira, A., Almeida, F., Matos, T., & Reis, J. (2021). Modeling wine preferences by data mining from physicochemical properties. Decision Support Systems, 47(4), 547-553.
12. Fildes, R., Ma, S., & Kolassa, S. (2019). Retail forecasting: Research and practice. International Journal of Forecasting, 35(2), 645-655.
13. Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, 29(5), 1189-1232.
14. Hastie, T., Tibshirani, R., & Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer.
15. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.
16. Hyndman, R. J., & Athanasopoulos, G. (2018). Forecasting: Principles and Practice. OTexts.
17. Hyndman, R. J., Bergmeir, C., Caceres, G., & O'Hara-Wild, M. (2021). Forecasting with

Machine Learning. Journal of Business Research.

18. Keerthi, S. S., & Lin, C. J. (2020). Decision trees for retail demand forecasting: A case study. European Journal of Operational Research, 185(2), 789-802.
19. Liaw, A., & Wiener, M. (2002). Classification and regression by randomForest. R news, 2(3), 18-22.
20. Livieris, I. E., Drakopoulou, K., & Kiriakidou, N. (2020). Demand forecasting using machine learning models: A case study. Expert Systems with Applications, 145, 113089.
21. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. PLOS ONE, 13(3), e0194889.
22. Natekin, A., & Knoll, A. (2013). Gradient boosting machines, a tutorial. Frontiers in Neurorobotics, 7, 21.
23. Farabi, S. F., Prabha, M., Alam, M., Hossan, M. Z., Arif, M., Islam, M. R., ... & Biswas, M. Z. A. (2024). Enhancing Credit Card Fraud Detection: A Comprehensive Study of Machine Learning Algorithms and Performance Evaluation. Journal of Business and Management Studies, 6(3), 252-259.
24. Mozumder, M. A. S., Sweet, M. M. R., Nabi, N., Tusher, M. I., Modak, C., Hasan, M., ... & Prabha, M. (2024). Revolutionizing Organizational Decision-Making for Banking Sector: A Machine Learning Approach with CNNs in Business Intelligence and Management. Journal of Business and Management Studies, 6(3), 111-118.
25. Bhuiyan, M. S., Chowdhury, I. K., Haider, M., Jisan, A. H., Jewel, R. M., Shahid, R., ... & Siddiqua, C. U. (2024). Advancements in early detection of lung cancer in public health: a comprehensive study utilizing machine learning algorithms and predictive models. Journal of Computer Science and Technology Studies, 6(1), 113-121.
26. Nabi, N., Tusher, M. I., Modak, C., Hasan, M., ... & Prabha, M. (2024). Revolutionizing Organizational Decision-Making for Banking Sector: A Machine Learning Approach with CNNs in Business Intelligence and Management. Journal of Business and Management Studies, 6(3), 111-118.
27. Rahman, M. A., Modak, C., Mozumder, M. A. S., Miah, M. N. I., Hasan, M., Sweet, M. M. R., ... & Alam, M. (2024). Advancements in Retail Price Optimization: Leveraging Machine Learning Models for Profitability and Competitiveness. Journal of Business and Management Studies, 6(3), 103-110.
28. Shahid, R., Mozumder, M. A. S., Sweet, M. M. R., Hasan, M., Alam, M., Rahman, M. A., ... & Islam, M. R. (2024). Predicting Customer Loyalty in the Airline Industry: A Machine Learning Approach Integrating Sentiment Analysis and User Experience. International Journal on Computational Engineering, 1(2), 50-54.
29. Modak, C., Ghosh, S. K., Sarkar, M. A. I., Sharif, M. K., Arif, M., Bhuiyan, M., ... & Devi, S. (2024). Machine Learning Model in Digital Marketing Strategies for Customer Behavior: Harnessing CNNs for Enhanced Customer Satisfaction and Strategic Decision-Making. Journal of Economics, Finance and Accounting Studies, 6(3), 178-186.
30. Mozumder, M. A. S., Nguyen, T. N., Devi, S., Arif, M., Ahmed, M. P., Ahmed, E., ... & Uddin, A. (2024). Enhancing Customer Satisfaction Analysis Using Advanced Machine Learning Techniques in Fintech Industry. Journal of Computer Science and Technology Studies, 6(3), 35-41.

31. Arif, M., Hasan, M., Al Shiam, S. A., Ahmed, M. P., Tusher, M. I., Hossan, M. Z., ... & Imam, T. (2024). Predicting Customer Sentiment in Social Media Interactions: Analyzing Amazon Help Twitter Conversations Using Machine Learning. International Journal of Advanced Science Computing and Engineering, 6(2), 52-56.

32. Md Al-Imran, Salma Akter, Md Abu Sufian Mozumder, Rowsan Jahan Bhuiyan, Md Al Rafi, Md Shahriar Mahmud Bhuiyan, Gourab Nicholas Rodrigues, Md Nazmul Hossain Mir, Md Amit Hasan, Ashim Chandra Das, & Md. Emran Hossen. (2024). EVALUATING MACHINE LEARNING ALGORITHMS FOR BREAST CANCER DETECTION: A STUDY ON ACCURACY AND PREDICTIVE PERFORMANCE. The American Journal of Engineering and Technology, 6(09), 22–33. https://doi.org/10.37547/tajet/Volume06Issue09-04

33. Md Abu Sufian Mozumder, Fuad Mahmud, Md Shujan Shak, Nasrin Sultana, Gourab Nicholas Rodrigues, Md Al Rafi, Md Zahidur Rahman Farazi, Md Razaul Karim, Md. Sayham Khan, & Md Shahriar Mahmud Bhuiyan. (2024). Optimizing Customer Segmentation in the Banking Sector: A Comparative Analysis of Machine Learning Algorithms. Journal of Computer Science and Technology Studies, 6(4), 01–07. https://doi.org/10.32996/jcsts.2024.6.4.1