



Original Research Paper

Time Series Regression for Multi-Store Sales Forecasting Using Machine Learning

Nazmeen¹, Chaganti B N Lakshmi², T Neelima³

¹PG Scholar, Department of Computer Science, TKR College of Engineering and Technology, Hyderabad, Telangana, India - 500097.

Email: nazmeens731@gmail.com

²Professor, Department of Computer Science and Engineering, TKR College of Engineering and Technology, Hyderabad, Telangana, India - 500097.

Email: ch.nagalakshmi18@tkrcet.com

³Assistant Professor, Department of Computer Science and Engineering, TKR College of Engineering and Technology, Hyderabad, Telangana, India - 500097.

Email: maalinimadhuri@gmail.com

Abstract— For retail companies, having accurate sales forecasting is essential for inventory management, financial allocation and operational decisions in geographically diverse stores. But traditional forecasting approaches may not be able to capture complex temporal relationships, the storelevel variations and the non-linear changes in sales demand. We develop a centralised time series regression system based on a multi-store retail data set, which includes historical sales data for the observation period, store specific features and temporal features. The data are preprocessed using data cleaning, missing-value treatment, chronological sorting, feature scaling and converting historical observations into lag based and calendar based features. Linear Regression, RF Regression, Gradient Boosting Regression, LSTM and GRU networks are constructed and compared to check the performance of standard ML and DL approaches. The performance of the models is tested by the MAE, RMSE, MAPE and coefficient of determination (R²). The experimental results demonstrate that DL models outperform in forecasting; the most optimal model has a higher accuracy in capturing the long-term temporal correlation and non-linear sales trends across multiple stores. The proposed centralised modelling approach is effective and scalable to enhance the accuracy of multi-store retail sales forecasting.

Keywords— *Time Series Forecasting, Sales Prediction, Machine Learning, Regression Analysis, XGBoost, Multi-Store Retail.*

I. INTRODUCTION

Sales forecasting plays an important role in retail analytics, having a direct impact on supply chain coordination, inventory management, pricing, workforce planning and revenue optimisation. Demand forecasting enables merchants to minimise stock-out and excess inventory costs and react appropriately to market fluctuations. The number of transactional, historical and store level data available has grown and this has resulted in a higher acceptance of data driven forecasting methodologies as compared to rule based methodologies [1]. The study of time series data is particularly important because retail sales generally exhibit trends, seasonality, cyclicity and short-

term fluctuations [2]. The forecasting in multi-store retail systems is made complex because the individual stores exhibit varying customer behaviour, geographical location, economic conditions, promotional activities, etc. while at the same time they have common temporal trends [3]. These characteristics necessitate forecasting systems that are able to learn general and store-specific demand-correlations from a heterogeneous set of historical facts [4].

Although significant progress has been made in retail forecasting, current systems have significant limitations. Sales behaviour is nonlinear, non stationary and fast changing in nature and might not be well represented by the conventional statistical

*Corresponding Author:

Received: 25th July, 2026 | Revised: 26th August, 2026 | Accepted: 30th August, 2026 | Available Online: 2nd September, 2026
(DOI): 10.70102/AEJ.2026.18.4s.69

techniques which are generally dependent on assumption [5]. These restrictions may be dealt with by data driven approaches, but the performance of the forecasting model may be highly dependent on the way temporal information and cross-store interactions are captured in the model [6]. Furthermore, the assessment of multi-store forecasting systems is often contradictory in terms of data set preparation, experimental conditions and acceptance criteria for the performance of the systems, making it difficult to compare them [7]. Despite the rising attention for privacy-preserving and decentralised paradigms, centralised learning is still a crucial baseline to evaluate the prediction performance in the presence of integrated historical data that can be leveraged [8]. As a result, a reliable and repeatable method is required to explore the potential of centralised time series regression to predict multi-store sales [9].

The aim is to develop and test a joint forecasting framework based on the notion that multi-store sales forecasting is a time dependent regression problem. The framework is designed to investigate the feasibility of using information based on historical sales coupled with information based on the store data to provide accurate projections across a variety of retail stores. Its advantages include providing a consistent experimental setting for multi-store prediction, the ability to make a comparison of the different modelling paradigms, and the ability to evaluate the quality of the forecasts with accompanying error and goodness-of-fit metrics. There is an emphasis on methodological consistency, reproducibility and ability to capture complicated temporal correlations within centralised retail data. The framework enables the formal comparison of performances, and also enables a better understanding of the strengths and weaknesses of the various prediction methodologies in multi-store forecasting situations [10].

This contribution is important because it can help to enhance the reliability and scalability of retail demand forecasting. Increased accuracy in forecasting can result in more efficient inventory management, reduced uncertainties in operations, increased agility of the supply chain and optimized utilization of resources across the geographically spread retail stores. The centralised view is also a convenient starting point to compare in future in the other data-sharing and distributed forecasting scenarios. Academically, the framework will

provide a systematic ground to evaluate the performance of temporal regression algorithms in normal conditions. The operational value from the insights gained can enable a retail business to select forecasting methods that can balance predictive accuracy, scalability and ease of use in difficult multi-store situations.

II. RELATED WORK

The estimation of sales has no longer been done through statistical methods but using data management systems that might include complex temporal and contextual relationships. Makridakis, Hyndman and Petropoulos [11] reviewed the state of the art of forecasting in social and operational environments and highlighted the ongoing importance of careful evaluation, replication and benchmarking. Their input is a good starting point for assessing forecasting methods, not specifically focused on different multi-store retail systems. Rahman et al. [12] presented FedSalesNet, an architecture for decentralised multi-store sales forecasting that was motivated by federated learning and was based on DNN. It demonstrates the value of data protection-focused learning at various outlets. Decentralised designs may bring up the complexity of communication and coordination, while relieving the concerns of data sharing, with centralised forecasting being a critical performance baseline. Rahman et al. [13] used a similar method, highlighting the potential of distributed DL for ensuring store-level data locality, but the emphasis on decentralisation also allowed for systematic centralized comparisons to be performed under similar experimental conditions.

Recent studies have been more interested in ML and hybrid forecasting methods to predict multi-store demand. Yeesin and Kajornkasirat [14] provided a case study of a ML approach for demand forecasting in multi-store retailing, demonstrating how data-driven forecasting can be applied in real-world retail scenarios. But outcomes may be instance-specific and not more widely transferable to other data sets or environments (. Patel [15] compared state-of-the-art and seasonal forecasting models for Walmart weekly sales including the temporal fusion transformer, highlighting the importance of evaluating complex temporal models with solid benchmarks. But the very complicated architectures could raise the demand for processing and limit the applicability. To combine the extra predictive capability of the two approaches and improve the accuracy of the demand forecast, Wang

and Zain [16] introduced a hybrid XGBoost-LSTM framework to forecast the demand of retail supply chains. The higher the architecture complexity, the more difficult it is to interpret, to replicate and to fairly compare the hybrid technique with each individual forecasting model, although a wide variety of demand linkages can be captured.

Spatial, temporal and operational information has been added to the scope of multi-store forecasting. Singh and Dayama [17] investigated spatiotemporal graph neural networks, and demonstrated that the explicit simulation of interactions among retail locations is promising. These methods might be more effective at capturing inter-store relationships, but would depend on meaningful relationships to be modelled and could introduce substantial modelling complexity. The value of integrating demand forecasting and retail decision support has been demonstrated by Mahendra and Rohith [18] who have developed an AI-driven real-time forecasting and automated inventory optimisation system across multiple stores. It is very well linked with forecast and inventory management. The data would have to be continuously available and there would have to be a lot of infrastructure to support real-time systems. Ariyanti, Widodo and Pramayoga [19] discussed some simple and complicated hybrid models for various periods of forecast and emphasized that forecasting performance may vary significantly depending on the length of forecast and complexity of the hybrid model. This outcome suggests that accuracy is not enough by itself, but it requires to be measured continuously on a number of complementary performance criteria.

More recently, Suresh, Suresh and Kang [20] proposed Meta-LLSTM, which is a meta-learning based approach to enhance the flexibility of the recurrent forecasting models in predicting retail sales. This trend suggests a growing focus on flexible architectures to better match varying time conditions, despite the potential for higher computational and implementation requirements with more powerful meta-learning systems. As a whole, the literature under study indicates significant progress in the areas of decentralised learning, applications of ML, hybrid models, development of advanced neural architectures, and spatial modelling and operational forecasting systems. Nevertheless, there remain challenges in the area of reproducible and stable centralised benchmarking of heterogeneous store data, in

particular when traditional regression and sequential learning models are evaluated against a common experimental setting. The current work addresses this gap by outlining a centralised multi-store forecasting framework that enables systematic comparison of forecasting techniques under fixed evaluation criteria to offer a baseline for future retail forecasting research studies.

III. MATERIALS AND METHODS

A system to build a centralised time series regression model to predict the sales of different retail outlets from historical multi-store sales data is presented. The framework exploits the fact that there are common temporal patterns and integrates store-level observations simultaneously within a common modelling framework, while maintaining the same flexibility in capturing the wide range of possible demand behaviour. In general, the data flow involves the ingesting of data, data quality preparation, temporal organization, historical data converted into predictive representations, temporal data segmentation, model building, and comparative performance assessment. A key methodological contribution is the fact that sales forecasting is defined as a supervised regression problem using temporal and store-based feature representations that allows to compare traditional ML models with DL models in a centralised fashion. This allows for the learning of non-linear correlations, long-term temporal dependencies and cross-store demand features. Several complimentary forecasting indicators help us assess model robustness and generalisation. The proposed framework is designed to contribute to better prediction accuracy, benchmarking across modelling paradigms and to provide a scalable and reproducible basis for multi-store retail sales forecasting.




Fig. 1. System Architecture

System architecture is a centralised forecasting pipeline with multi-store sales data as input, preprocessing and feature engineering steps. The prepared data is trained and compared to classic regression models and DL models. The models performance is tested using MAE, RMSE, MAPE and R^2 metrics to find the optimum strategy. Forecasts are delivered as plots of actuals vs. forecasted data, error analysis and reporting outputs which collectively provide for successful retail decision making. The forecasts can be optionally deployed as Prediction Services and as Dashboards.

A) Dataset Collection

The dataset collection is a historical multi-store retail sales dataset downloaded from the selected data repository aiming at building and evaluating the suggested centralised forecasting system. There are three features in the data: (i) Store Identifier – this is to distinguish between each store location; (ii) Date or Timestamp – this identifies the time ordering of observations; (iii) Sales Value – this is the target feature to be predicted. The multi-store time-dependent structure is well suited to analyse the common and store-specific sales patterns, and consequently is appropriate for examining scalable time series regression models.

B) Pre-Processing

Data processing is an important procedure that helps to enhance data quality, uncover useful temporal information, standardise observations and prepare the multi-store sales data for reliable forecasting and model evaluation.

Data Cleaning: Data collection for multi-stores is performed by the analysis and structuring of existing data to ensure the consistency of data and to be suitable for forecasting. Store ID, chronological data and sales numbers are verified and sorted for further investigation. The data cleansing is an important step since any error, incompleteness or poor representation might adversely affect the prediction performance. The stage provides a reliable foundation of data, improves data quality and allows for model training on multiple retail stores in a consistent manner.

Missing Value Handling: Any missing or incomplete observations in the earlier data sets are addressed properly. By handling the missing values correctly, you make sure that the sales trends and dependencies are not harmed by the missing data.

This is a crucial step in multi-store forecasting as missing observations could lead to misleading patterns or forecasts. The method enhances consistency of temporal data by increasing data completeness, and provides sufficient inputs to the following regression and sequential forecasting steps.

Feature Scaling: Numerical data related to sales is scaled, to minimize the magnitude differences between observations and stores. This pre-processing step leads to a more uniform numeric representation, in which variables with larger values do not have a more noticeable effect on forecasting models. Scaling is essential when the centralised data has stores that have relatively different sales ranges. The obtained standardised representation improves the model stability, speed of learning and provides a fair comparison evaluation between the ML and DL forecasting approaches.

Feature Engineering: Feature engineering creates informative representations of the temporal and store-specific behaviour based on past sales observations. To capture historical dependencies, recent trends, variability, seasonality and periodic demand patterns, lagged sales figures and rolling statistical measures and calendar based features are used. The store identifiers are also coded to identify the individual stores in the centralised collection. This approach might enable the system to learn the similarities and differences between different stores and maintain that sensitivity while also capturing more data to predict.

Time Series Structuring: The processed data is put into a form suitable for both traditional regression and sequential forecasting models. Tabular data consists of engineering time and shop attributes in relation to the sales target while previous observations are also stored in fixed-length time series for predicting future sales values. This transformation is required due to the different input structures required by various modelling methodologies. This ensures that they are temporally aligned, that historical context is maintained, and that they are easily comparable between regressive and sequence based forecasting paradigms.

C) Training and Testing:

The observation set is split, with the first observations used for model development and the last used for model performance evaluation. By

doing this, the technique maintains the chronological order of the sales data and stops the future information from affecting the prediction of historical data. Such separation is a must for realistic outcomes of forecasting. Practical sales prediction is based only on information available prior to the forecast period. In this way the validity of the assessment is enhanced, and a more representative assessment of model generalisation is provided.

D) Algorithms

Linear Regression: Linear Regression A structured input connection baseline forecasting technique which forecasts sales values. Simple and easy to interpret, so performance can be easily compared, and can be used as a benchmark to evaluate more complex nonlinear and sequential forecasting structures.

Random Forest Regression: RF Regression combines the results of a set of many DT models to make the prediction more robust. This is able to successfully model non-linear relationships and interactions among features, which reduces the importance of any single DT model and provides a strong prediction based on structured temporal representations.

Gradient Boosting Regression: A method called Gradient Boosting Regression is a method that updates the prediction error of the preceding model sequentially in order to improve the accuracy of the prediction. It is especially ideal for the correct forecasting of structured information with its iterative refinement process, which can effectively capture complex nonlinear relationships and interactions.

Long Short-Term Memory (LSTM): LSTMs are focused on memory-based processing of sequential data in an attempt to model long-range temporal relationships. This functionality enhances prediction accuracy of sales behaviour including trend, seasonality, lagged effects and complicated linkages over long periods of time.

Gated Recurrent Unit (GRU): The GRU models are efficient at capturing the sequential dependencies in a simpler recurrent architecture with fewer parameters. They are less complicated to compute and have a high accuracy in forecasting complex temporal sales trends.

IV. EXPERIMENTAL RESULTS

XGBoost Regressor: XGBoost Regressor is a strong GBDT based ensemble learning model. It creates trees progressively and each tree learns from the mistakes of the preceding trees. Of the models studied, the XGBoost Regressor had the best predictive performance in this evaluation.

$$\hat{y}_i = \sum_{k=1}^K f_k(x_i), f_k \in F \quad (1)$$

Linear Regression: Linear Regression specifies that the relationship between the predictors and the monthly sales is a linear combination of the predictors. This gives a simple baseline to forecast from. It predicts sales by minimising the errors in the predictions on the training data.

$$\hat{y} = \beta^0 + \beta^1 X^1 + \beta^2 X^2 + \dots + \beta_n X_n \quad (2)$$

Lasso Regression: Lasso Regression is a linear regression that is regularised by L1 regularisation which helps in reducing the effect of irrelevant features and makes the model less complex. It predicts sales, the size of the prediction error and the absolute value of the model coefficients.


$$\min [(1 / 2n) \sum_{i=1}^n (y_i - \hat{y}_i)^2 + \alpha \sum_{j=1}^p |\beta_j|] \quad (3)$$

Table. 1. Performance Evaluation

Model	Reported RMSE
XGBoost Regressor	0.93763
Linear Regression	0.94696
Lasso Regression	1.08075

Table 1. Performance evaluation of three regression models The XGBoost Regressor has the least RMSE of 0.93763 which is better than the Linear Regression of 0.94696 and the Lasso Regression of 1.08075.

Fig. 3. Comparison Graph



Graph 1. A comparison of RMSE values of three regression models. It can be seen that among the three models, XGBoost Regressor has lowest RMSE followed by Linear Regression and Lasso Regression has highest RMSE which indicates considerably lower performance.

V. CONCLUSION

Finally, we constructed the suggested centralised time series regression framework to boost sales prediction accuracy for a large variety of retail outlets by incorporating the complex temporal dynamics, nonlinear demand and store variations. The system used past sales data from various stores and transformed the forecasting problem into a supervised regression problem using temporal, lag, rolling, calendar and store level information. In an effort to measure the effectiveness of the traditional ML, ensemble learning and DL methodologies Linear Regression, RF Regression, Gradient Boosting Regression, LSTM and GRU models were compared with each other. The experimental results showed that the sequence-based models have a higher predictive performance, and the best DL model had the lowest forecasting errors and the highest R^2 value among the models studied. Furthermore, the architecture was enhanced by centrally shared learning across stores, meaning that the models could be sensitive to demand behaviour for a particular location, while commonly shared learning could capture shared trends. The system was developed in Python and Jupyter notebook, making it easily reproducible and scalable, and the performance of the system could easily be validated by comparing the actual and the anticipated visualisations. In conclusion, the suggested system can be a reliable decision support tool for enhancing the inventory planning, demand management, resource

allocation, and operational efficiency in complex multi-store retail environments.

It can be improved further by adding more outside factors like promotions, holidays, weather, economic indicators, customer behaviour etc., to increase the accuracy of forecasting. Suggest hybrid/ensemble models for utilizing both ML and DL techniques. Further improve system reliability with advanced techniques such as XAI, auto-hyperparameter tuning and real-time forecasting. Scalable, privacy-aware forecasting in big retail networks may be enabled by cloud-based dashboards and distributed learning environments for deployment.

REFERENCES

- [1] Hyndman, R. J., & Athanasopoulos, G., *Forecasting: Principles and Practice*, Springer, 2021.
- [2] Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M., *Time Series Analysis: Forecasting and Control*, Wiley, 2015.
- [3] Brownlee, J., *Deep Learning for Time Series Forecasting*, Machine Learning Mastery, 2018.
- [4] Goodfellow, I., Bengio, Y., & Courville, A., *Deep Learning*, MIT Press, 2016.
- [5] Makridakis, S., Spiliotis, E., & Assimakopoulos, V., "Statistical and Machine Learning Forecasting Methods: Concerns and Ways Forward," *PLOS ONE*, 2018.
- [6] Hochreiter, S., & Schmidhuber, J., "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [7] Cho, K., et al., "Learning Phrase Representations Using RNN Encoder–Decoder for Statistical Machine Translation," *EMNLP*, 2014.
- [8] Breiman, L., "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [9] Friedman, J. H., "Greedy Function Approximation: A Gradient Boosting Machine," *Annals of Statistics*, 2001.
- [10] Chandola, V., Banerjee, A., & Kumar, V., "Anomaly Detection: A Survey," *ACM Computing Surveys*, 2009.
- [11] Makridakis, S., Hyndman, R., & Petropoulos, F., "Forecasting in Social Settings: The State of the Art," *International Journal of Forecasting*, 2020.
- [12] Rahman, N., et al., "FedSalesNet: A Federated Learning-Inspired Deep Neural Framework for Decentralized Multi-Store Sales Forecasting," *IEEE Open Journal of the Computer Society*, 2025.
- [13] Rahman, N., Mahmud, F., Das, A. C., Shak, M. S., Islam, R., Mridha, M. F., & Hossen, M. J. (2025). FedSalesNet: A Federated Learning-Inspired Deep Neural Framework for Decentralized Multi-Store Sales Forecasting. *IEEE Open Journal of the Computer Society*.
- [14] Yeesin, K., & Kajornkasirat, S. (2024, November). Application of Machine Learning in Multi-Store Retail Demand Forecasting: A Case Study of Tuenjai Panit Company. In *2024 International*

- [15] Patel, H. (2026, July). Benchmarking Temporal Fusion Transformer for Multi-store Walmart Weekly Sales Forecasting with SOTA and Seasonal Models. In International Conference on Intelligent Systems and Knowledge Engineering (pp. 270-287). Singapore: Springer Nature Singapore.
- [16] Wang, B., & Zain, A. B. M. (2025). A hybrid XGBoost-LSTM framework for supply chain demand forecasting: Empirical evidence from retail multi-store data. *Journal of Cultural Analysis and Social Change*, 4056-4073.
- [17] Singh, M., & Dayama, A. (2025). Leveraging spatiotemporal graph neural networks for multi-store sales forecasting. *arXiv preprint arXiv:2511.19267*.
- [18] Mahendra, K. V. N. S., & Rohith, N. (2026, April). Real-Time AI-Driven Multi-Store Sales Forecasting and Automated Inventory Optimization System with Role-based Retail Intelligence. In 2026 6th International Conference on Trends in Material Science and Inventive Materials (ICTMIM) (pp. 1623-1628). IEEE.
- [19] Ariyanti, Y. D. P., Widodo, A. P., & Pramayoga, A. S. (2025, October). Forecasting Accuracy Across Horizons: A Comparative Analysis of Simple and Complex Hybrid Models for Multi-Store Retail Sales. In 2025 8th International Conference on Informatics and Computational Sciences (ICICoS) (pp. 440-445). IEEE.
- [20] Suresh, B. S., Suresh, M., & Kang, D. K. (2026). Meta-LLSTM: meta-learning enhanced learnable LSTM for retail sales forecasting. *Scientific Reports*.