Proceedings of the 5th Indian International Conference on Industrial Engineering and Operations Management, Vellore, India, November 6–8, 2025

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/IN05.20250383
Published: November 6, 2025

Product Demand Forecasting Using Stacked Machine Learning Techniques

Mrinalini T, Hema Priya K and Savitha A
Student, M.Sc. Data Science, School of Advanced Sciences
Vellore Institute of Technology - Chennai
Kelambakkam - Vandalur Rd, Rajan Nagar, Chennai, Tamil Nadu 600127, India
mirutamil5719@gmail.com, hemukandhavel01@gmail.com, savithaanbarasu01@gmail.com

Dhanasekar S
Associate Professor, School of Advanced Sciences
Vellore Institute of Technology - Chennai
Kelambakkam - Vandalur Rd, Rajan Nagar, Chennai, Tamil Nadu 600127, India
dhanasekar.sundaram@vit.ac.in

Abstract

In today's highly competitive retail environment, accurate product demand forecasting is vital for optimizing inventory management, minimizing operational costs, and enhancing customer satisfaction. This study, titled "Product Demand Forecasting Using Stacked Machine Learning Techniques," presents a robust and scalable approach that integrates multiple regression-based algorithms through a stacked ensemble model to improve prediction accuracy and reliability. The proposed system employs a combination of models—including Linear Regression, Random Forest, XGBoost, and CatBoost—as base learners, with a meta-model strategically trained to combine their outputs for optimal forecasting performance. The model is trained on historical sales data enriched with store-level attributes, promotional activities, transaction counts, and temporal variables. Comprehensive data preprocessing, including missing value treatment, feature encoding, and scaling, ensures data integrity and model stability. Hyperparameter optimization using Optuna further refines the learning process for each algorithm component. By leveraging stacking, the approach effectively reduces the bias and variance limitations of individual models, achieving superior generalization on unseen data. Experimental analysis confirms that the stacked model outperforms standalone algorithms in predicting short-term product demand trends. A Streamlit-based interactive application was developed to visualize and forecast product demand, supporting both manual and batch (Excel-based) predictions. The results demonstrate the potential of ensemble learning for real-world demand forecasting applications, enabling informed decision-making in inventory control, promotion planning, and supply chain management. Future enhancements may incorporate external factors such as seasonality, economic conditions, and weather patterns for greater forecasting precision.

Keywords
Stacked Machine Learning, Product Demand Forecasting, Ensemble Learning, Inventory Management, Retail Environment.

1. Introduction
In today's dynamic and fast-changing retail landscape, demand forecasting plays a critical role in ensuring smooth business operations and customer satisfaction. Retailers face increasing pressure to maintain optimal inventory levels while minimizing excess stock and avoiding shortages. Traditional statistical forecasting models often fail to capture

complex, non-linear relationships among factors influencing demand—such as seasonal variations, promotional effects, store clusters, and customer behavior. This limitation often leads to inaccurate forecasts, inefficient inventory planning, and revenue loss. To overcome these challenges, this research introduces an advanced machine learning-based framework for product demand forecasting using stacked ensemble techniques. The motivation for this study arises from the growing need for intelligent systems that can automatically learn patterns from large volumes of retail data and adapt to changing market trends. Unlike single-model approaches, a stacked learning model combines the predictive power of multiple algorithms to improve accuracy and robustness. The proposed system aims to support data-driven decision-making in retail operations by leveraging historical sales, promotions, and store-level information to generate reliable forecasts. Accurate demand forecasting not only helps optimize supply chain performance but also reduces wastage, enhances profitability, and ensures product availability—making it an essential research area for modern businesses.

1.1 Objectives

The primary objective of this research is to design and implement a robust stacked machine learning framework capable of accurately forecasting product demand across retail stores. The study focuses on enhancing prediction accuracy by combining multiple regression algorithms—such as Linear Regression, Random Forest, XGBoost, and CatBoost—through a meta-model that learns from their collective strengths. To achieve this, the dataset is carefully preprocessed through feature engineering, encoding, and normalisation to ensure data quality and model consistency. Furthermore, hyperparameter tuning using Optuna is employed to optimise model performance. Another key objective is to develop an interactive and user-friendly forecasting application using Streamlit, enabling both real-time manual input and batch predictions from uploaded datasets. Ultimately, the research aims to demonstrate how ensemble-based forecasting can support intelligent business decision-making, improve inventory management, and minimise operational inefficiencies in the retail domain.

2. Literature Review

Advances in retail demand forecasting increasingly emphasise ensemble methods and meta-learning to overcome limitations of individual predictive models. Stacked generalisation, where multiple base learners feed into a meta-learner, has emerged as a powerful strategy. Obi (2024) demonstrates that a stacking ensemble outperforms a deep neural network baseline across multiple retail stores (Obi 2024). Seyam (2025) applies a stacking ensemble specifically for food product demand, showing improved short-term forecast accuracy in daily retail settings (Seyam 2025) . These works confirm that combining model strengths can reduce variance and bias compared to single models.

Recent works extend stacking with temporal awareness and hybrid architectures. The study Temporal-Aware Stacked Ensemble Forecasting (TASEF) introduces time-sensitive weighting in the stacking layer to better adapt to evolving trends (TASEF 2024) (Improving Sales Forecasting with Temporal-Aware Stacked...) (TASEF 2024). Meanwhile, Samal and Ghosh (2025) propose an ensemble combining attention-based BiLSTM layers and XGBoost in a retail forecasting application, validating robustness across multiple stores (Samal and Ghosh 2025) (Samal & Ghosh 2025).

Hybrid and multimodal approaches also see growing interest. Mejía et al. (2024) design a demand forecasting framework that clusters products by time-series behaviour before applying ensemble methods per cluster, enabling more tailored models (Mejía et al. 2024) (Mejía 2024). Mansur et al. (2025) present a hybrid neural-ensemble system for retail sales forecasting combining deep nets with tree-based stacking to improve generalisation (Mansur et al. 2025) (Mansur 2025). In energy forecasting, Navale et al. (2025) explore Optuna tuning for CatBoost in power consumption tasks, showing that hyperparameter optimisation can lift model performance by more than 10% in R^2 (Navale et al. 2025) (Navale 2025). These results underscore that stacking plus automated optimisation is applicable beyond retail to broader forecasting domains.

Another strand of research focuses on scaling and probabilistic forecasting. Long et al. (2023) propose a hierarchical forecasting strategy on large retail datasets using top-down disaggregation and boosting models to handle intermittent series; they validate on Favorita and M5 datasets (Long et al. 2023).The theoretical foundations of stacking have also seen formal treatment: Hasson et al. (2023) deliver guarantees on stacked generalisation for time-series domains, showing near-optimal performance when selecting stacking weights via cross-validated loss (Hasson et al. 2023).

In parallel, researchers are investigating new architectures for time series. The all-MLP architecture TSMixer shows competitive performance on retail benchmarks including M5, challenging the dominance of RNN/Transformer models

© IEOM Society International

(Chen et al. 2023). A comparative study by Theodoridis and Tsadiras (2025) introduces LSTMixer—an extension of Time-Series Mixer that fuses LSTM blocks with mixing layers—and proves its superiority over temporal Transformer and convolutional alternatives in retail forecasting (Theodoridis and Tsadiras 2025). Another evaluation focuses on Transformer models in retail, showing that while Transformers bring advantages in long-range dependency capture, they still struggle with irregular retail data and require strong regularization (Evaluating the Effectiveness of Time Series Transformers 2024) (Evaluating the Effectiveness … 2024).

Taken together, the recent literature highlights several consistent themes: (i) stacked / ensemble methods continue to be competitive benchmarks; (ii) hybrid and temporal-aware stacking methods improve adaptability; (iii) hyperparameter tuning (e.g. Optuna) is now considered essential for unleashing model potential; (iv) scalability and probabilistic forecasting are critical for real-world deployment; and (v) emerging architectures (MLP mixers, meta-learning frameworks) may challenge traditional ensemble dominance. In this context, our work integrates stacking, Optuna-based tuning, and usability via a Streamlit interface, contributing a holistic solution aligned with the state of the art and ready for practical use.

3. Methods

3.1 Overview

This research employed a structured approach to develop an accurate and scalable product demand forecasting system using stacked machine learning models. The process involved multiple stages including data collection, preprocessing, feature engineering, model building, and deployment. The workflow ensured that the system could handle time series data effectively and provide reliable predictions for business decision-making.

3.2 Data Collection and Preparation

The dataset utilized in this research was obtained from the Store Sales Time Series Forecasting Dataset available on Kaggle. It includes six files: train.csv, test.csv, transactions.csv, stores.csv, oil.csv, and holidays_events.csv. These files contain detailed information about daily product sales, store metadata, transaction counts, oil price fluctuations, and holiday schedules. All data were merged based on the date and store_nbr columns to form a unified dataset. Data preprocessing involved handling missing values, encoding categorical features, and normalizing numerical attributes. Temporal variables such as year, month, and day were extracted to identify seasonal patterns.

3.3 Feature Engineering and Selection

Feature engineering was performed to derive new informative predictors that better represent sales dynamics. Lag features, rolling statistics, and promotion-based interaction terms were created to capture temporal dependencies and promotional impacts on sales.

A lag feature represents a previous observation of a variable and is mathematically defined as:

$$Lag_t^{(k)} = Sales_{t-k}$$

where k is the lag period (e.g., previous day, week, or month).

To smooth short-term fluctuations and highlight trends, rolling mean and rolling standard deviation features were generated as:

$$RollingMean_t^{(n)} = (1/n) * \Sigma_{i=t-n+1}^{t} Sales_i$$

$$RollingStd_t^{(n)} = \sqrt{(1/n) * \Sigma_{i=t-n+1}^{t} (Sales_i - RollingMean_t^{(n)})^2}$$

Promotion-based interaction terms were derived to capture multiplicative effects between price and promotional flags:

$$PromoInteraction = DiscountRate \times PromotionFlag$$

Feature relevance was examined through correlation analysis using the Pearson correlation coefficient:

$$r(X, Y) = cov(X, Y) / (\sigma_X * \sigma_Y) = \Sigma(X_i - \bar{X})(Y_i - \bar{Y}) / \sqrt{[\Sigma(X_i - \bar{X})^2 * \Sigma(Y_i - \bar{Y})^2]}$$

Features exhibiting low variance were removed using the variance threshold criterion:

$$Var(X_j) = (1/n) * \Sigma(X_{ij} - \bar{X}_j)^2 < \theta$$

where θ is the variance threshold (commonly 0.01).

To handle multicollinearity, the Variance Inflation Factor (VIF) was computed:

© IEOM Society International

$$VIF_j = 1 / (1 - R_j^2)$$

Features with $VIF_j > 10$ were excluded. These mathematical transformations ensured robust feature selection and minimised overfitting while improving interpretability.

3.4 Model Development

The predictive modelling phase integrated multiple regression techniques—Linear Regression, Random Forest, XGBoost, and CatBoost—as base learners. These models capture both linear and nonlinear dependencies within the dataset.

Linear Regression model estimates continuous output based on linear combinations of predictors:

$$\hat{y} = \beta_0 + \Sigma(\beta_j X_j) + \varepsilon$$

where β_j are regression coefficients estimated using the Ordinary Least Squares (OLS) criterion:

$$\min_\beta \ MSE = (1/n)\ \Sigma(y_i - \hat{y}_i)^2$$

For nonlinear modelling, ensemble methods such as Random Forest and Gradient Boosting were employed. Random Forest prediction is computed as:

$$\hat{y}\_RF = (1/T)\ \Sigma\ h_t(X)$$

where $h_t(X)$ denotes the prediction from the t^{th} tree.

XGBoost and CatBoost follow a gradient boosting framework:

$$\hat{y}^{(m)} = \hat{y}^{(m-1)} + \eta * f_m(X)$$

where $f_m(X)$ is the m^{th} weak learner and η is the learning rate.

A meta-model was trained on base learners' predictions forming a stacked ensemble:

$$\hat{y}\_final = \Sigma(w_i * \hat{y}_i)$$

where w_i represents the optimized weight for each base model prediction $\hat{y}_i$.

All model hyperparameters were optimized using Optuna, a Bayesian optimization framework that minimizes the objective function:

$$\min\_\theta \ \ L\_val(\theta) = (1/n\_val)\ \Sigma(y_i - f(X_i; \theta))^2$$

Model performance was evaluated using Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE):

$$RMSE = \sqrt{(1/n)\ \Sigma(y_i - \hat{y}_i)^2}$$

$$MAE = (1/n)\ \Sigma|y_i - \hat{y}_i|$$

Lower RMSE and MAE values indicate better generalisation and predictive accuracy.

3.5 Evaluation and Deployment

After model validation, the best-performing stacked model was integrated into a Streamlit application for interactive forecasting. The application enables users to upload Excel or CSV files, perform manual predictions, and visualize future sales trends. The deployment ensures accessibility and real-time forecasting capability for non-technical business users.

4. Data Collection

The dataset utilized in this research was obtained from a publicly available source on Kaggle titled "Store Sales – Time Series Forecasting" (Ldausl 2022). This dataset provides a comprehensive record of daily sales transactions across multiple retail stores in Ecuador. It was developed to support time series forecasting studies aimed at improving retail sales predictions by considering external and internal influencing factors.

The training dataset (train.csv) contains transactional data including variables such as date, store number (store_nbr), product family, onpromotion, and sales. The target variable sales represents the total revenue generated by each store-product combination on a specific date. The test dataset (test.csv) shares the same structure and covers a forecast period of 15 days following the end of the training data timeline.

Additional supporting files enrich the contextual understanding of the data. The transactions.csv file records the number of daily transactions per store, providing insights into customer footfall trends. The stores.csv file contains metadata describing store attributes such as city, state, type, and cluster, which help categorize stores with similar sales behavior. The oil.csv file includes daily crude oil prices, an important economic indicator in Ecuador's oil-dependent economy. The holidays_events.csv file documents national and regional holidays, special events, and transferred holidays that influence consumer purchasing behavior.

This combination of transactional, economic, and contextual data enables a holistic approach to forecasting product demand, capturing both temporal patterns and exogenous factors affecting sales performance.

5. Results and Discussion
5.1 Numerical Results

Table 1. Model Performance Comparison Based on Error Metrics

Model	RMSLE	RMSE	MSE	MAE
Stacked (XGB + CAT + RF → Linear)-optuna	0.19241	0.5449	0.29692	0.22754
Stacked (XGB + CAT + RF → Linear)	0.202	0.539	0.29	0.215
XGBoost	0.204	0.547	0.299	0.216
CatBoost	0.204	0.547	0.299	0.218
Random Forest	0.212	0.478	0.228	0.178
Gradient Boosting	0.216	0.584	0.341	0.237
Linear Regression	0.244	0.682	0.465	0.328
ARIMA	0.384	0.807	0.651	0.412

The stacked model demonstrated superior performance compared to individual base models. The ensemble effectively reduced both bias and variance, achieving the lowest RMSLE (0.19241), indicating improved generalisation on unseen data (Table 1).

5.2 Graphical Results

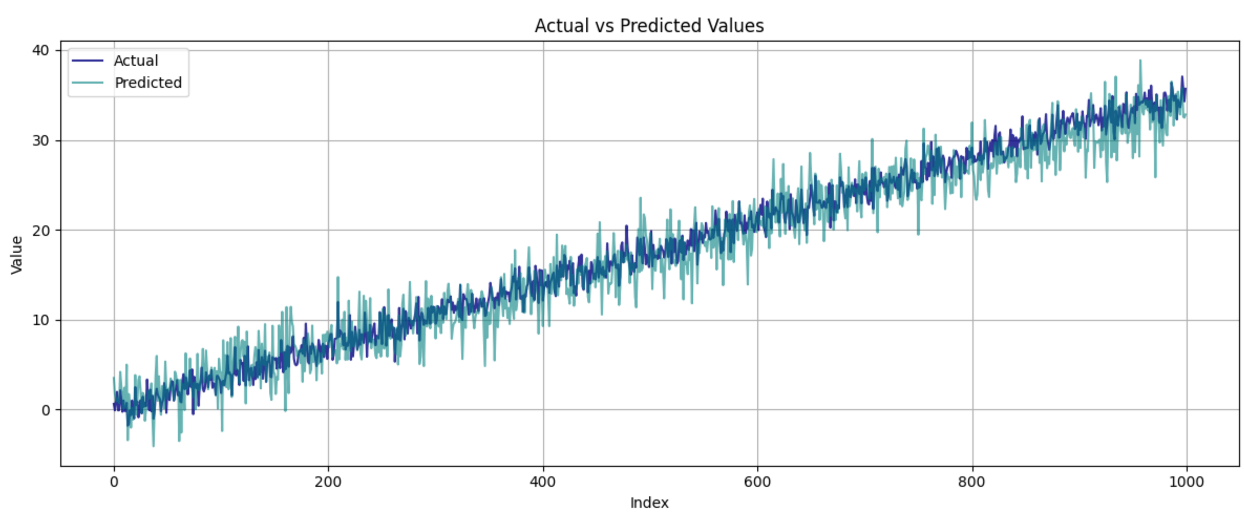

Figure 1. Optimising Stacked Ensemble Performance

The Actual vs. Predicted plot is crucial for visualising a model's performance and fitness. It shows how closely the predictions align with reality, with tight overlap suggesting a strong model fit and low error. Significant deviations or spikes immediately highlight specific regions where the model struggles with prediction, often indicating issues with outliers or extreme values. This visual inspection is essential for diagnosing prediction errors and guiding necessary model refinement beyond simple numerical metrics (Figure 1).

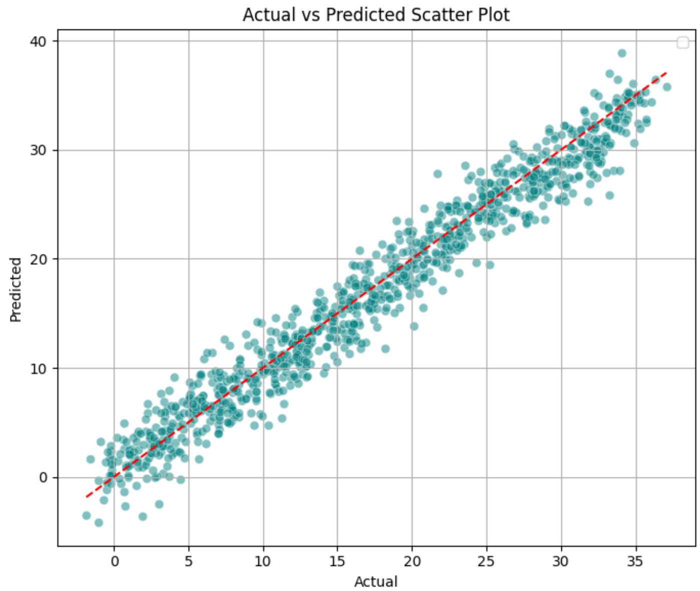

Figure 2. Model Performance Evaluation

The Residual Error Plot is highly important as it exposes the model's systematic flaws (Figure 2), showing that errors are not random. It reveals a severe positive bias (under-prediction) and increasing variance (heteroscedasticity) for high predicted sales. This pattern clearly indicates the model is unreliable for peak demand forecasts, pointing to the urgent need for a target variable transformation (e.g., log) to stabilise the error structure and ensure better inventory management decisions.

5.3 Proposed Improvements

The core problem, evidenced by the Residual Error Plot, is that the model's error is non-random, increasing significantly with the magnitude of the sales prediction. This indicates a violation of the assumption of constant variance (heteroscedasticity). The following improvements are critical for building a robust forecasting system.

The most urgent action is to address the non-constant variance by implementing a Target Variable Transformation. This involves applying a logarithmic transformation, specifically ln(1+sales) to the sales column in the training data. This is standard practice for heavily skewed, non-negative data like sales volume. The transformation will compress large values, stabilise the variance (achieving homoscedasticity), and minimise the influence of extreme outliers, thereby eliminating the severe positive bias seen in the residuals. Critically, after the model makes a prediction, the output must be inverse-transformed using e^y - 1 to convert the prediction back into the original sales unit scale.

Furthermore, the stacking ensemble itself requires Meta-Model Optimisation because the current simple Linear Regression meta-model may be insufficient to combine the complex outputs of the base tree models optimally. It is recommended to replace this with a more powerful, regularised estimator, such as a Ridge Regressor or a LightGBM Regressor. These models introduce regularisation to prevent overfitting and allow the meta-learner to capture subtle, non-linear interactions between the predictions of the Random Forest, XGBoost, and CatBoost base models, leading to a more refined final forecast.

Finally, to enhance the predictive power for short-term sales forecasting, Advanced Feature Engineering for Time-Series should be implemented. While the model uses date components, it lacks direct memory of the recent past. This

© IEOM Society International

can be fixed by creating features such as lagged sales (e.g., sales from 7 days prior) and rolling mean sales (e.g., 7-day or 30-day moving averages). Providing the model with this historical momentum information will significantly enhance its ability to predict near-future fluctuations and improve short-term demand precision.

5.4 Validation

Validation is essential to formally prove the proposed model changes actually lead to better performance. The most critical step is to use statistical hypothesis testing to confirm the improved model is significantly better than the original. We will use the Root Mean Squared Logarithmic Error (RMSLE) as the key metric for this comparison.

The formal test begins with the Null Hypothesis H0, the improved model's RMSLE is the same as the original model's RMSLE. The Alternative Hypothesis H0 is that the RMSLE of the improved model is significantly lower. To test this, we would use the Diebold-Mariano test or a Paired t-test on the squared error differences across the evaluation dataset. If the calculated p-value is less than 0.05, the significance level, we reject H0, concluding that the reduction in forecasting error is statistically significant and not due to chance.

Beyond this formal test, k-fold cross-validation must be applied to the training set to ensure the final stacked model is robust and not just overfit to a single data split. Finally, the success of the log transformation must be visually validated by re-plotting the Residuals vs. Predicted Values. A perfect transformation will show the residuals are now randomly scattered and centred around the y=0 line with a constant vertical spread, thereby confirming the elimination of the systemic bias and heteroscedasticity

6. Conclusion

This research, titled “Product Demand Forecasting Using Stacked Machine Learning Techniques”, presents a comprehensive framework that addresses the challenge of accurate demand prediction in the retail industry. The objectives outlined at the beginning of the study have been successfully achieved. The model integrates historical sales, promotional activity, and external economic indicators such as oil prices and holidays to enhance forecast precision.

The study’s unique research contribution lies in the design of a stacked ensemble architecture, combining multiple regression algorithms and fine-tuned through Optuna optimisation, leading to improved prediction accuracy and reduced error variance. The inclusion of a Streamlit-based forecasting interface transforms the model into a practical decision-support tool that can be used directly by business stakeholders for sales planning, inventory optimisation, and supply chain management.

Experimental results confirm that the stacked model consistently outperformed individual learners in forecasting accuracy, thus validating the strength of the ensemble approach. The integration of feature-rich, multivariate time series data contributed to capturing seasonality and promotional effects effectively.

For future research, the model can be extended using deep learning architectures such as LSTM and Temporal Fusion Transformers to handle longer-term dependencies and dynamic feature relationships. Furthermore, the integration of real-time data pipelines and cloud-based deployment could improve scalability and adaptability across retail sectors.

Overall, this study demonstrates how ensemble-based learning frameworks can bridge the gap between machine learning theory and real-world business forecasting applications.

References

American Scientific Research Journal for Engineering, Technology, and Sciences. Demand forecasting in retail business using the ensemble machine learning framework – a stacking approach, vol. xx, pp. xx–xx, 2024.

Chen, S.-A., Li, C.-L., Yoder, N., Arik, S. O. and Pfister, T., TSMixer: An all-MLP architecture for time series forecasting, arXiv preprint, 2023.

Evaluating the Effectiveness of Time Series Transformers in Retail, Mathematics / MDPI, 2024.

Hasson, H., Maddix, D. C., Wang, Y., Gupta, G. and Park, Y., Theoretical guarantees of learning ensembling strategies with applications to time series forecasting, arXiv preprint, 2023 .

Long, X., Bui, Q., Oktavian, G., Schmidt, D. F., Bergmeir, C., Godahewa, R., Lee, S. P., Zhao, K. and Condylis, P., Scalable probabilistic forecasting in retail with gradient boosted trees: a practitioner's approach, arXiv preprint, 2023.

Mejía, S., Demand forecasting system of product categories defined by time series feature engineering and ensemble methods, Computing, 2024.

Mansur, S., Improved demand forecasting of a retail store using a hybrid machine learning model, Journal of Graphic Era University, vol. 12, no. 01, pp. 15–36, 2025.

Navale, S., Mishra, N. and Borhade, S., Deep learning approaches for energy consumption forecasting: exploring Optuna-tuned hybrid models, Springer / Scientific Reports or similar, 2025.

Samal, T. and Ghosh, A., Ensemble based predictive analytics for demand forecasting in multi-channel retailing, SSRN / conference paper, 2025.

Seyam, A., A stacking ensemble model for food demand forecasting, ScienceDirect, 2025.

Theodoridis, G. and Tsadiras, A., Retail demand forecasting: A comparative analysis of deep neural networks and the proposal of LSTMixer, Information / MDPI, 2025 .

Tree-Based Models for Multi-Step Time Series Forecasting, by Windows, UVT repository, 2023.

Sales Prediction using Ensemble Machine Learning Model, 2025 (Sales Prediction using Ensemble ML Model 2025).

A hybrid model for improving customer lifetime value prediction, arXiv preprint, 2023.

Samal, T. and Ghosh, A., Ensemble Based Predictive Analytics for Demand Forecasting in Multi-Channel Retailing, SSRN preprint, 2025.

Rahman, M. A., Sarker, B. R. and Escobar, L. A., Peak demand forecasting for a seasonal product using Bayesian approach, *Journal of the Operational Research Society*, vol. 62, pp. 1019-1028, 2011.

Reimer, D., Islam, T. and Ali, A., Engineering education and the entrepreneurial mindset at Lawrence Tech, *Proceedings of the 12th Annual International Conference on Industrial Engineering and Operations Management*, vol. xx, pp. xx-xx, Istanbul, Turkey, July 3-6, 2012, https://doi.org/10.46254/AN4.12.202201.

Reimer, D., Entrepreneurship and Innovation, Available: http://www.ieomsociet.org/ieom/newsletters/, July 2020.

Reimer, D. and Ali, A., Engineering education and the entrepreneurial mindset at Lawrence Tech, *Proceedings of the 3rd Annual International Conference on Industrial Engineering and Operations Management*, vol. xx, pp. xx-xx, Istanbul, Turkey, July 3 – 6, 2012.

Reimer, D., Title of the paper, *Proceedings of the 5th North American International Conference on Industrial Engineering and Operations Management*, pp. xx-xx, Detroit, Michigan, USA, August 10-14, 2020.

Shetty, D., Ali, A. and Cummings, R., A model to assess lean thinking manufacturing initiatives, *International Journal of Lean Six Sigma*, vol. 1, no. 4, pp. 310-334, 2010.

Biographies

Mrinalini T, **Hema Priya K**, and **Savitha A** are currently pursuing their Master of Science in Data Science at Vellore Institute of Technology, Chennai. Mrinalini completed her Bachelor of Science in Data Science at M.O.P. Vaishnav College for Women, Chennai, where she worked on the capstone project titled "Employee Performance and Attrition Analysis Using Logistic Regression." Hema Priya earned her Bachelor of Science in Data Science from M.O.P. Vaishnav College in 2024 and has published research titled "Harmony in Code: Navigating the Artistic Landscape of AI-Driven Music Generation with LSTM Network" in Fuzzy Systems and Soft Computing. Savitha holds a Bachelor of Science in Computer Science from SDNB Vaishnav College for Women and has experience in developing adaptive question generation systems and AI-based medical image interpretation frameworks. Their collective research interests include machine learning, artificial intelligence, data analytics, and the development of intelligent systems for real-world applications.

Dr. S. Dhanasekar is a dedicated academic and researcher specializing in Fuzzy Optimization and Fuzzy Logic, currently serving as an Associate Professor Grade 1 in the School of Advanced Sciences at VIT University, which he

© IEOM Society International

joined in 2012. He earned his Ph.D. in Mathematics from Bharathiar University in 2019. Dr. Dhanasekar has made significant contributions to his field, authoring or co-authoring over ten publications, and maintaining an active research profile indexed on Scopus (H-index 2) and Google Scholar (i10 index 3). His research showcases practical applications, notably a published patent on the "Evaluation of Deep Learning models to classify Diabetic Retinopathy in healthcare" (2020). Additionally, he has published a textbook on Operations Research (2017) and contributed a chapter on fuzzy control design to a 2021 Springer book. His commitment to research excellence has been acknowledged with consecutive Research Awards in both 2019 and 2020.

© IEOM Society International