ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

A Machine Learning-Enhanced Software Framework for Intelligent Inventory Monitoring and Demand Forecasting of Perishable Goods: Evidence from Developing Economy SMEs

Paul Conteh[1]; Adamsay Turay[2]; Mohamed Thoronka[3]

[1]MSc Software Engineering, Limkokwing University of Creative Technology, Sierra Leone Campus BSc Hons in Business Information Technology Nankai University, College of Software Engineering, Tianjin, China
[2]MSc Software Engineering, University of Makeni, Makeni, Sierra Leone BSc in Computer Science Nankai University, College of Software Engineering, Tianjin, China
[3]MSc Software Engineering, Limkokwing University of Creative Technology, Sierra Leone Campus BSc Hons in Business Information Technology Nankai University, College of Software Engineering, Tianjin, China

Publication Date: 2026/05/28

Abstract: In many developing countries, poor management of perishable goods causes significant economic and nutritional losses. For example, food waste rates in sub-Saharan Africa are over 50%. Most smart inventory systems are designed for large companies with plenty of data and strong infrastructure, so small and medium-sized businesses (SMEs) often have limited options. This study introduces a machine learning-based software framework for smarter inventory monitoring and demand forecasting of perishable goods, tailored for retail SMEs with limited resources. The framework was tested using real-world data from Sierra Leone. Five forecasting models were compared: Linear Regression (as a baseline), Random Forest (ultra-tuned), XGBoost (ultra-optimised), a Stacking Ensemble, and a Hybrid XGBoost-LSTM model. Tests on a combined dataset, using both a UCI retail transaction set and a synthetic Sierra Leone dataset, showed that the XGBoost UltraOptimised model had the best accuracy (R2 = 0.9781, MAE = 81.99 units), cutting forecasting errors by 89% compared to the linear baseline (MAE = 741.07). Analysis showed that short-term demand momentum variables were the most important, making up over 87% of the model's predictive power. The framework combines predictive modelling with a modular decision-support system that gives reorder alerts, wastage warnings, and inventory trend forecasts. It can run on standard computers without needing high-speed internet, making it practical for SMEs. The results show that welladapted machine learning tools can help close the gap between advanced algorithms and real-world needs in developing countries, helping reduce food waste, improve supply chains, and support SME growth.

Keywords: *Perishable Goods; Demand Forecasting; Machine Learning; Inventory Management; XGBoost; Developing Economies; SMEs; Supply Chain Optimisation.*

How to Cite: Paul Conteh; Adamsay Turay; Mohamed Thoronka (2026) A Machine Learning-Enhanced Software Framework for Intelligent Inventory Monitoring and Demand Forecasting of Perishable Goods: Evidence from Developing Economy SMEs. *International Journal of Innovative Science and Research Technology*, 11(5), 2007-2019. https://doi.org/10.38124/ijisrt/26may367

I. INTRODUCTION

Management of perishable inventories is among the most challenging tasks in supply chain operations for global retail networks. Perishable products include fruits and vegetables, dairy products, meat, and fish. They continuously deteriorate, they are highly sensitive to climatic factors, including humidity and temperature, and face dynamic demand due to seasonality, holidays, and fluctuations in income levels. FAO reports that about 1.3 billion tons of food are lost annually globally. Improper demand forecasting and inventory management are among the critical issues causing such losses [10]. The case becomes worse for sub-Saharan Africa, which experiences up to 50% of supply losses due to weak cold chain logistics, electricity outages, and dominance of informal cash trading [9, 13]. Sierra Leone represents a case study for this

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

phenomenon. The West African country, with a population of around 8.6 million people, has an economy characterised by agriculture and informal retail, with inventory management decisions relying on human experience and paper-based accounting practices [9]. These factors contribute to the continuous existence of two major inefficiencies: overstocked goods result in spoilage and financial losses, whereas insufficient inventory causes stockouts and loss of customers [7]. Without the possibility of tracking current inventory levels, all parties in the supply chain do not have the means to cooperate, leading to supply-demand mismatches across the entire industry [19].

The integration of machine learning (ML), cost-effective computing hardware, and open-source software presents an interesting chance to reinvent perishable inventory systems in situations where resources are scarce. Machine learning–based forecasting models have shown significant performance gains over conventional statistical techniques, such as ARIMA, exponential smoothing, and intuitive approaches, in different supply chain scenarios [16, 20]. Nonetheless, the majority of existing approaches have been developed for large-scale companies operating in well-developed countries, which are marked by a wealth of data, robust infrastructures, and highly skilled professionals [15]. The applicability of such models to SME settings in less economically developed regions, where data are scarce, and computations take place on inexpensive hardware, is understudied.

This paper attempts to fill that void. We propose a machine learning–augmented software platform for intelligent tracking and forecasting of demand for perishable products tailored to the requirements of SMEs in Sierra Leone. The proposed framework involves a series of five increasingly sophisticated forecasting algorithms in a layered software architecture, which we test against real-world datasets, as well as simulated data reflecting the peculiarities of small retail in Sierra Leone. The key research questions driving our investigation include:

➤ *Research Questions*

- RQ1: Which machine learning algorithms are the most effective in modelling the non-linear demand dynamics of perishable goods in the informal retail sector?
- RQ2: How is an ML-driven software system architecture designed to ensure technical viability along with practical applicability in a low-resource SME environment?
- RQ3: What are the key factors driving the accuracy of forecasts generated in informal perishable markets, and how do these translate into operational significance?

➤ *Main Contributions*

The main contributions of the study are:

- Comparative analysis of five algorithms used in ML forecasting: Linear regression, Random Forest, XGBoost, Stacking ensemble, and Hybrid XGBoost-LSTM, in a single demand forecasting framework for perishables, including time-aware cross-validation and feature importance assessment.
- Development of a multi-layered software architecture for perishable demand forecasting in informal retail settings that includes modules for data ingestion, ML forecasting, decision rules, dashboards for users, persistent data storage, and performance monitoring, optimized for consumer-grade hardware and free of cloud dependencies.
- Proving that the use of ensemble gradient boosting (XGBoost) in short-horizon perishable demand forecasting in informal retail far outperforms other statistical and deep learning approaches, achieving R2 = 0.9781 and reducing MAE by 89%.
- Identification of important aspects of perishable demand dynamics showing momentum as the dominant factor in short-horizon forecasting in informal markets in West Africa, leading to recommendations on what should be included in data sets collected in resource-constrained SMEs.

The rest of this article will be structured in the following manner. In section 2, we will define the problem formally. In section 3, the research scope and objectives will be defined. In section 4, we will discuss related literature. In section 5, we will present the methodology used. In section 6, we will outline the system design and framework architecture. In section 7, we will present the performance evaluation criteria. In section 8, we will provide the results.

➤ *Problem Statement*

The research problem statement is formulated as follows: Given a set of observations of historical daily sales data, a set of environmental contextual data, and a set of attributes relating to the specify product (perishability period, type of category, and its cost), create a computer system that would generate reliable demand predictions for perishables and translate those predictions into practical inventory decisions, while meeting deployability and usability requirements under the resource limitations commonly present in small and medium businesses operating in developing countries. The problem is considered more challenging than usual due to three reasons. First, the perishable demand curve has natural non-linearities: it is determined by the complex interplay of weekly demand cycles, seasonality of agricultural production, income-induced purchasing spikes, temperature-related spoilage factors, and promotional campaigns, none of which can be adequately captured using linear parametric modelling approaches [7]. Second, there is a significant limitation of available data: SMEs in Sierra Leone often lack systematic record keeping and rely either on manually transcribed data or synthesised data [9, 15]. Finally, deployment limitations are stringent: solutions should run on low-cost hardware, should not require technical expertise in management, and should work offline.

Analysis of existing methodologies shows that the problem has remained unsolved by current research. Classical inventory management models (such as EOQ and FIFO-replenishment) assume that there is constant demand and reliable supply, both of which cannot be expected in informal markets [18]. Statistical time series models (ARIMA, Holt-Winters) improve on classical models; however, they are inadequate in dealing with the non-linearity and richness of

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

today's retail data [16]. Machine learning methods (LSTMs and Transformers) show great accuracy when working with data-rich environments, but require significant computing power and data for training, thus making them inappropriate for use by SMEs [20]. Existing machine learning inventory systems rely on enterprise-level IT capabilities and custom data pipelines, which makes their use unavailable to micro and small businesses in developing countries [15].

Three critical gaps remain in the literature. First, no integrated ML-based system has been purpose-designed for perishable demand forecasting in developing-economy SMEs. Second, empirical comparisons of the full spectrum of forecasting algorithms from linear baselines to ensemble and deep learning methods in informal African retail markets are absent. Third, the feature dynamics underlying demand predictability in these contexts have not been empirically investigated, depriving practitioners and researchers of data-driven feature selection guidance. This study addresses all three gaps.

➢ *Scope and Objectives*

- *Research Scope*

This research addresses demand forecasting and inventory monitoring for perishable goods at the retail level within small and mediumsized enterprises. The primary geographical reference context is Sierra Leone, though the framework is designed to be transferable across analogous developing-economy settings in West Africa and beyond. Product scope encompasses fresh produce (fruits and vegetables), dairy products, and staple perishable foods categories that collectively account for the majority of daily consumer demand in Sierra Leonean informal markets.

The study focuses specifically on the forecasting and monitoring functions of inventory management. Upstream supply chain functions, such as supplier procurement, logistics planning, and route optimization, are explicitly out of scope. Dynamic pricing and revenue management, while noted as future extensions, are not implemented in the current framework. The research is grounded in framework design and empirical simulation; full industrial-scale deployment with live production data from partner SMEs is identified as a necessary next step and a clear boundary of the present study.

- *Research Objectives*

The following research objectives describe the scope of this study:

✓ Obj1: Conduct a systematic literature review of current machine learning models applied to perishable-demand forecasting and explain the discrepancy between advanced techniques and the practical challenges encountered by SMEs operating in developing countries.
✓ Obj2: Develop and implement a flexible, six-layer software architecture that combines machine learning demand forecasting with real-world decision support capabilities, allowing implementation on regular consumer devices without reliance on proprietary software packages.
✓ Obj3: Compare the predictive performance of five demand forecasting models Linear Regression, Random Forest, XGBoost, Stacking Ensemble, and Hybrid XGBoost-LSTM using both open-source retail data and an artificially generated dataset capturing the characteristics of Sierra Leone's perishable retail industry.
✓ Obj4: Evaluate model performance using established statistical indicators (R2, MAE, RMSE, directional accuracy) through a time-aware rolling window cross-validation approach designed to avoid temporal data leakage.
✓ Obj5: Determine which demand features impact forecasting accuracy in non-formal retail settings the most and provide recommendations for data gathering in SME operations within developing economies.
✓ Obj6: Guarantee that the software framework adheres to responsible AI standards protecting data privacy, promoting algorithmic fairness, and ensuring transparent decision making processes suitable for implementation in developing countries.

II. LITERATURE REVIEW

➢ *Inventory Management for Perishable Goods*

Perishable inventory management is different from general commodity stock control because of limited shelf life, sensitivity to the environment, and unpredictable quality changes over time [18]. Traditional methods like Economic Order Quantity models and FIFO/FEFO rotation policies assume demand is predictable, and supply is stable, but these conditions are rare in practice, especially in developing countries where supply chains are irregular, and consumer behaviour is unpredictable [7]. As a result, performance evaluation should go beyond statistical accuracy and include operational measures such as waste reduction, stockout frequency, and service level achievement [19].

Moving from manual, ledger-based systems to enterprise resource planning (ERP) platforms in the late twentieth century helped centralise data, but did not improve prediction. Today, intelligent systems that use cloud computing, IoT sensors, and advanced analytics have shown they can lower operational costs in the food sector [8, 22]. Still, it is not clear how well these systems work for SMEs in low-income economies.

➢ *Machine Learning for Demand Forecasting*

Machine learning is increasingly replacing traditional statistical methods for demand forecasting in supply chains. Models like ARIMA and exponential smoothing are easy to interpret, but they cannot capture the complex, non-linear relationships found in realworld perishable demand. Ensemble methods, especially Random Forest [5] and gradient boosting with XGBoost [6], often perform better than linear models because they can model complex feature interactions and use regularisation to prevent overfitting. Recent studies show that ensemble methods usually provide better generalisation than single-model approaches [16, 20].

Recurrent neural networks, especially Long Short-Term Memory (LSTM) networks, are good at modelling sequences

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

in time-series data [12]. Hybrid models that combine gradient boosting with LSTM can capture both complex feature interactions and time-based patterns. However, deep learning methods need much more data and computing power than tree-based ensembles, which can be a major barrier for SMEs in developing economies [14]. Adetula and Akanbi [4] found that using machine learning for forecasting in SME supply chains lowered inventory holding costs by 12 to 18 per cent across different product categories [19]. [19] also showed that ML-based inventory optimisation cut perishable waste by about 20 per cent in food retail. While these results provide evidence, most come from developed markets or large companies.

➢ *SME Inventory Challenges in Developing Economies*

SMEs in sub-Saharan Africa face three main challenges: operational, financial, and technological limitations [21]. Fragmented and informal value chains cause high variability in product quality and delivery times. Limited cash flow makes businesses very cautious, so having too much inventory is especially expensive. Low digital skills, unreliable electricity, and poor internet access also make it hard to use cloud-based or complex decision-support systems [15]. Sierra Leone is a good example, with about 8.6 million people and an economy that depends on agriculture and informal retail, most SMEs manage perishable goods based on experience rather than data or technology. Inventory decisions are usually reactive, which leads to repeated problems of overstocking slow-moving items and running out of popular perishables [9]. According to statistics, more people now use mobile devices, which could help, but any solution must work well with mobile technology and use little data to be practical.

➢ *Technology Adoption in Supply Chains of Emerging Markets*

The digitisation process of the supply chains of developing economies faces a set of issues that differ fundamentally from those in developed economies. The major inhibitors identified by Morakanyane et al. [17] are inadequate digital skills, poor infrastructure quality, and high initial investment costs. In turn, Lugina and Myamba [15] note that small-scale businesses in Sub-Saharan Africa running informal trade usually have manually recorded data, lack any structured historical data, and present considerable resistance to the use of new technologies. As proven by Duncombe [9], mobile-based applications requiring minimal data are the only way forward for the digitisation of supply chain processes in Western Africa.

Despite such constraints, the widespread distribution of affordable smartphones and increased mobile Internet coverage creates a possibility of adopting lightweight and mobile applications even in countries with an underdeveloped digital ecosystem [9, 15]. Prior research on the intelligent inventory management systems in developing economies (albeit limited in scope) highlights appropriate context-awareness as the key to success: technology designs assuming infrastructures, large data, and technical skills consistently fail [15, 21].

➢ *Theoretical Foundations*

There are three theoretical paradigms on which the construction of the proposed framework rests. Firstly, there is Systems Theory [3]. It postulates the existence of an ecosystem made up of several interacting subsystems, such as data collection, forecasting, decision making, and inventory replenishment processes. These subsystems are interconnected through a constant exchange of feedback among themselves. Secondly, Decision Support System Theory [1]. According to it, smart systems should be designed to enhance the operator's decision-making abilities by generating analytical results that are then translated into recommendations that can be understood and overridden by operators. Lastly, there is Intelligent Systems Theory [2], which forms the basis for the ongoing learning process in the proposed framework.

➢ *Related Gap*

Despite extensive research on machine learning-based inventory management, a significant gap persists at the intersection of perishable goods, small and medium-sized enterprises (SMEs) in developing economies, and practical system design. Most existing frameworks are either generic and do not address the specific challenges of perishability, or they are designed for large enterprises with advanced IT infrastructure. There is a lack of context-specific frameworks that integrate machine learning forecasting, decisionsupport logic, and deployable software architecture for SMEs in resource-constrained environments. This study aims to address this identified gap.

III. RESEARCH METHODOLOGY

➢ *Research Design*

The study adopted Design Science Research (DSR) as its main methodology [11]. DSR is appropriate when the objective is to produce and evaluate a practical technological artefact, such as an ML-enhanced software framework, rather than to merely describe or explain existing phenomena. In contrast to approaches that analyse inventory practices in abstraction, DSR supports iterative design and evaluation cycles that progressively refine the artefact against clearly defined performance criteria. The DSR process was integrated with a systematic machine learning experimentation protocol to ensure that the forecasting component was selected based on empirical performance rather than theoretical assumptions. This hybrid design is particularly well-suited to the interdisciplinary nature of the research problem, which spans supply chain management, machine learning, software engineering, and development economics.

Figure 1 illustrates the integrated research workflow combining Design Science Research with machine learning experimentation.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

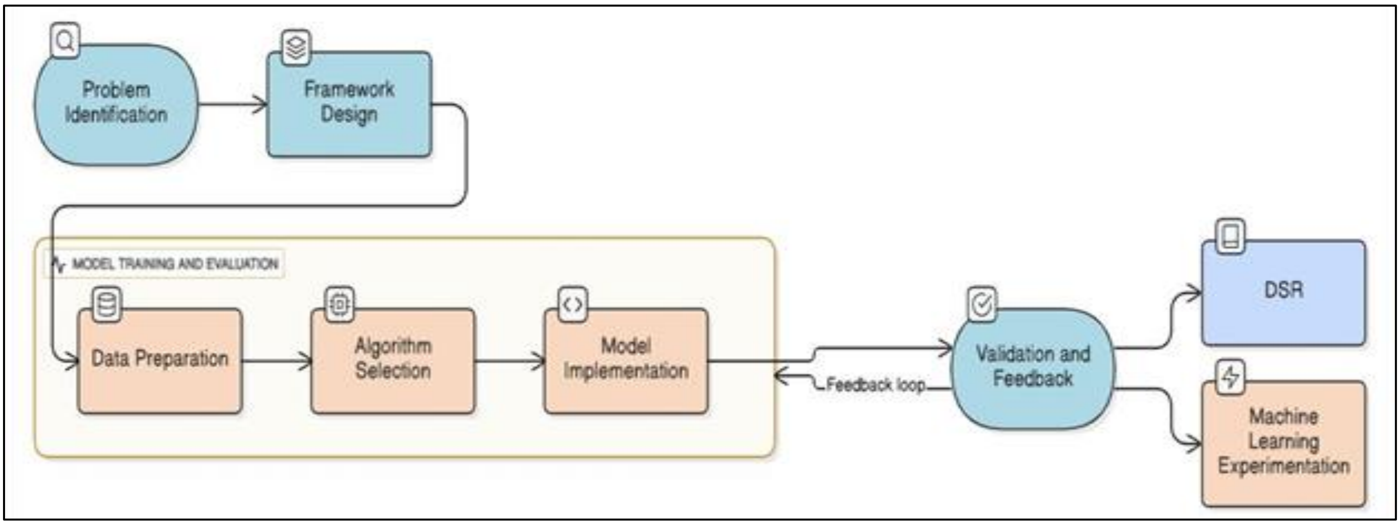

Fig 1 Hybrid Research Design Flow Integrating Design Science Research and Machine Learning Experimentation

➢ *Data Sources*

Two complementary datasets were utilised to develop and evaluate the forecasting framework, capturing both the general properties of retail perishable demand and the specific operational characteristics of the informal SME sector in Sierra Leone.

The primary dataset was obtained from the UCI Machine Learning Repository and consists of over 145,000 transactional records from a retail chain operating in a developing-economy context. Although geographically distinct from Sierra Leone, the dataset exhibits key structural characteristics such as demand seasonality, promotional effects, store-level variability, and irregular sales patterns. These attributes render it suitable for initial model development and benchmarking.

To address the lack of locally sourced operational data and to more accurately represent the Sierra Leone SME environment, a synthetic dataset was generated, encompassing two years of daily transaction records. Variable selection was guided by operational challenges identified in the literature review, including demand volatility associated with market days and festive cycles, temperatureinduced spoilage dynamics, and supply chain uncertainties such as variable lead times.

Table 1 Key Variables in the Synthetic Sierra Leone Perishable Goods Dataset

Variable	Description
date	Daily timestamp for each transaction record
units_sold	Target: quantity of units sold per day
stock_level	End-of-day inventory quantity on hand
wastage_units	Units expired or spoiled before sale
price	Selling price per unit (local currency)
temp_celsius	Ambient temperature affects spoilage rate
is_holiday_or_market_day	Binary: peak demand days (market days, public holidays)
is_promo	Binary: active promotional campaign
lead_time_days	Supplier delivery delay in days
expiry_days	Remaining shelf life of the product batch

➢ *Data Preprocessing and Feature Engineering*

Both datasets were processed through a standardised preprocessing pipeline implemented in Python using Pandas, NumPy, and Scikit-learn. Missing values in time-series fields were handled by forward filling to preserve temporal continuity. Numerical features were scaled using MinMaxScaler to prevent scale-driven model bias. Categorical variables, including product category and promotional status, were encoded using one-hot encoding. For LSTM sequence generation, a 30-day sliding window was applied to create sequential input blocks suitable for recurrent modelling.

Feature engineering played a critical role in capturing the demand dynamics specific to perishable goods markets. Five categories of engineered features were developed: temporal features (day-of-week, month, holiday indicators) to represent seasonality, historical lag features (1-day, 7-day, 30-day lags) to model shortterm demand persistence, rolling statistical features (7-day and 30-day averages and standard deviations) to identify trends and volatility, contextual features (ambient temperature, promotional status) to account for environmental and commercial influences; and inventory-specific features (stock coverage days, shelf life remaining) to directly connect forecasting outputs to inventory decision-making.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

➢ *Machine Learning Models*

Five forecasting models were implemented, encompassing a range of learning paradigms from simple linear models to hybrid deep learning architectures. This range was selected to enable a systematic assessment of performance improvements associated with increasing model complexity under the constraints of perishable goods data.

- *Linear Regression (Baseline).* The ordinary least-squares model assumes linear relationships between input variables and daily sales and serves as the interpretable baseline. The mathematical formulation is: y = + x, whereˆ y denotes predicted demand, x representˆ input features, and are coefficients estimated using least-squares minimisation. Although simple and computationally efficient, making it feasible for SME applications, the assumption of linear featuretarget relationships renders it systematically inadequate for the volatile, non-linear demand patterns characteristic of perishable goods markets. It therefore establishes the performance baseline against which all advanced models are compared.
- *Random Forest (Ultra-Tuned).* This bagging ensemble aggregates predictions from multiple decision trees trained on bootstrap samples. The prediction is computed as: y = (1/T) f(x), where f(x)ˆ denotes the prediction of the t-th tree and T is the total number of trees. Extensive hyperparameter tuning (n_estimators, max_depth, min_samples_split, max_features) was conducted to enhance generalisation and robustness to the noisy, incomplete sales data commonly found in SME environments. The ensemble approach effectively captures non-linear feature interactions while reducing overfitting.
- *XGBoost (Ultra-Optimised).* This model implements sequential gradient boosting with L1 and L2 regularisation, minimising a regularised loss function across boosting rounds: L = l(y, y) +ˆ (f), where l(·) is the squared error loss function and (f) penalises model complexity. This regularisation mechanism makes XGBoost particularly effective at capturing complex non-linear relationships in structured tabular data while preventing overfitting. XGBoost was selected as the primary high-accuracy learner due to its established state-of-the-art performance on structured forecasting tasks [6]. Hyperparameters were chosen via grid search, achieving a practical balance between predictive accuracy and computational tractability on standard SME hardware.
- *Stacking Ensemble.* This meta-learning approach integrates Linear Regression, Random Forest, and XGBoost as base learners, with a regularised Ridge regression meta-learner trained on out-offold predictions to optimally weight each model's contribution: y =ˆ g(y,ˆ y, ...,ˆ y), where g(ˆ ·) denotes the meta-learner function. By leveraging the complementary inductive biases of diverse base learners, the stacking ensemble aims to reduce model-specific biases and enhance generalisation stability under volatile demand conditions [20].
- *Hybrid XGBoost-LSTM.* This architecture combines gradient boosting with Long Short-Term Memory (LSTM) recurrent networks to simultaneously capture structured feature interactions and temporal sequential dependencies. The two components are merged as a weighted average: y = f_XGB(x) + (1ˆ)f_LSTM(X), where determines each component's contribution and X denotes sequential demand inputs across a 30-day sliding window. This hybrid approach enables the model to benefit from both gradient boosting's feature-engineering strengths and LSTM's temporal modelling capacity, though at increased computational cost.

The selection of this diverse model portfolio from simple linear methods to sophisticated hybrid architectures enables a systematic assessment of when increased model complexity yields meaningful performance gains for perishable goods forecasting in resource-constrained SME contexts. Table 2 provides a comparative overview of all five models, summarising their learning paradigm, core algorithmic approach, and intended role within the forecasting framework.

Table 2 Comparative Overview of Forecasting Models

Model	Learning Paradigm	Role in Framework
Linear Regression	Linear parametric	Baseline benchmark for comparison
Random Forest (UltraTuned)	Bagging ensemble	Robust non-linear learner for noisy data
XGBoost(Ultra-Optimised)	Gradient boosting	Primary high-accuracy predictor
Stacking Ensemble	Meta-learning	Enhances model generalisation
Hybrid XGBoost-LSTM	Hybrid ML-DL	Integrates temporal and structured features

➢ *Validation Approach*

Since the dataset was time-based, regular k-fold validation methods such as randomised shuffling of observation sequence were not suitable because they could lead to the inclusion of future observations during model training. A time-series validation method was employed for sequential data to ensure that no future observation would be used in predicting past observations. 80% of the data was used for training, while 20% was left for testing. To ensure that the model performance is estimated effectively without relying on one specific time window, several sequential folds were created. Hyperparameter tuning was performed using a grid search approach where possible.

IV. SYSTEM ARCHITECTURE AND FRAMEWORK DESIGN

➢ *Architectural Overview*

This framework adopts an architectural approach that combines the functionalities of data management, machine learning inference, decision support logic, and user interface in a single system. This design is fully compliant with the limitations of operating within resource-constrained

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

environments since it runs on ordinary computing infrastructure, it works offline and online, uses SQLite and PostgreSQL databases, and offers a light-weight web portal through Flask. Furthermore, this system uses open-source technologies such as Python 3.x, Scikit-learn, XGBoost, TensorFlow/Keras, Matplotlib, and Plotly without incurring any commercial licenses.

This system architecture consists of six functional layers. First, the Data Input Layer facilitates data ingestion from various sources, including CSV files, APIs, and database connections. It also processes the raw operational data using appropriate pre-processing pipelines. Second, the Machine Learning Layer contains five predictive models that can be trained separately. Third, the Decision Support Layer interprets the output of these forecasting models to generate actionable business advice. In particular, this layer triggers reordering alerts when expected demand is close to safety stock limits, alerts about wastage when stock is nearing expiration dates, and estimates future demand trends for procurement planning purposes. Fourth, the User Interface Layer provides visual interfaces using Streamlit and Dash platforms without needing technical skills. Inventory managers can view stock information, forecast results, and unusual events without much difficulty. Fifth, the Database Layer creates persistence in storing sales transactions, inventory levels, forecasting, and decision logs, linking all parts of the system.

The Evaluation and Feedback Layer compares predictions to actual results and triggers regular model retraining as new data comes in.

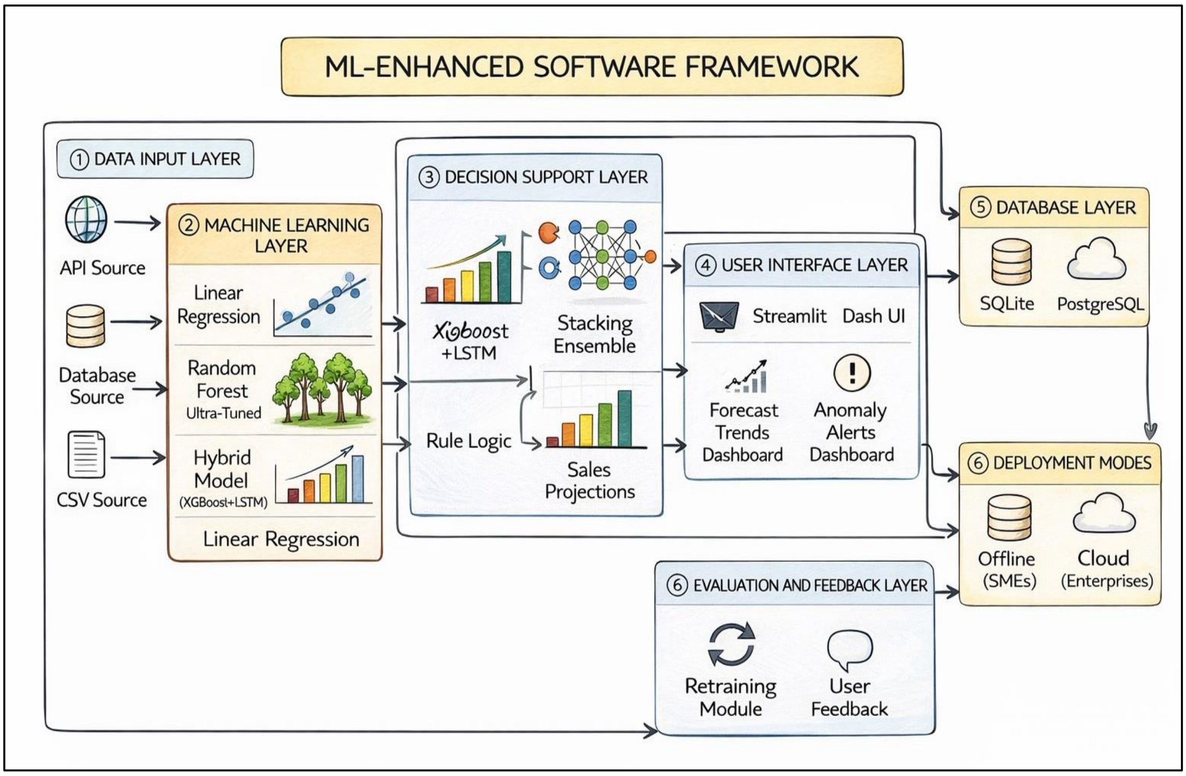

Fig 2 Six-Layer System Architecture and Interaction Between Functional Layers

Figure 2 illustrates the overall architecture and the interaction between these six functional layers. The architecture is designed to operate in two deployment modes: offline mode (using SQLite, suitable for individual SME stores with intermittent internet) and cloud mode (using PostgreSQL, suitable for network-connected retail chains). The modular design allows each layer to be updated or replaced independently without disrupting the entire system, ensuring long-term maintainability and adaptability as business needs and technology evolve.

➢ *Database Design*

Relational database design was carried out in a light-weight and extensible manner. There are six main tables in the database which are: Sales_Transactions (daily sales data with product id, quantity sold, price, and time stamp), Inventory_Records (quantity on hand, replenishment quantity, expiration date, and waste quantity), Product_Metadata (product category, unit, shelf life, and manufacturer details), External_Features (climate details, public holiday indicator, and promotion indicator), Forecast_Results (forecast values and time stamp), and Decision_Logs (reordering warning, stockout risk, and decision action). The access control measures ensure that only authorized operations are allowed on the database, and data consistency constraints ensure that quantities are always positive.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

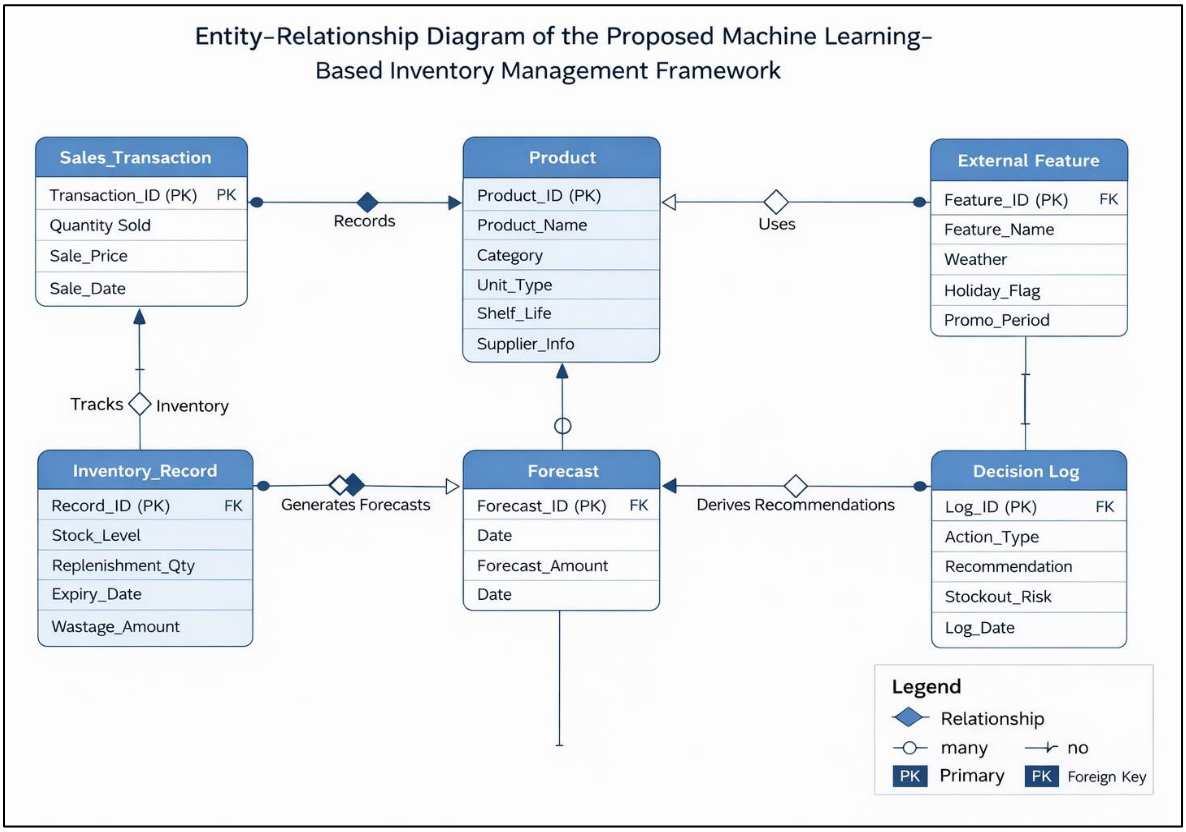

Fig 3 Six-Entity ER Diagram Showing Relationships Between Sales_Transactions, Inventory_Records, Product_Metadata, External_Features, Forecast_Results, and Decision_Logs

➢ *Implementation Stack*

The framework has been developed using exclusively Python 3.x, chosen for its robust data-science stack and wide user community. Data processing is carried out using Pandas and NumPy. Machine learning algorithms include traditional methods from Scikitlearn, gradient boosting using XGBoost, and sequential algorithms from TensorFlow/Keras (LSTM). Flask was used to develop the web application interface, while Streamlit and Dash were utilised for dashboard features. Database operations are performed using SQLAlchemy, which supports SQLite (local) and PostgreSQL (cloud) database management systems. Graphs and plots were generated using Matplotlib and Plotly. Code version control was provided using Git/GitHub, and the code is open-source, eliminating the software licensing cost as an entry barrier for SMEs.

The experiment was designed such that the only computing resources used were those available on regular consumer hardware CPU, no cloud computing clusters or graphics processors (GPU) were employed, it has the result of validating that the framework achieves competitive forecasting performance within the computational limits realistic for SME operators in Sierra Leone.

V. PERFORMANCE EVALUATION

➢ *Experimental Design*

All models have been trained and evaluated following a consistent experimental design process. The joint dataset (UCI retail dataset and artificial Sierra Leone dataset) has been split temporally, using the first 80% of samples for training and the last 20% for testing purposes. Notably, there is no shuffling involved in the split procedure, as it ensures the preservation of the temporal order and the avoidance of data leakage from the test sample to the training set, which is essential for time series model assessment.

Hyperparameter tuning has been achieved through grid search with cross-validation for models that could be trained computationally (Random Forest, XGBoost, and Stacking Ensemble), whereas for the LSTM network of the hybrid architecture, manual hyperparameter tuning based on validation loss was performed.

➢ *Evaluation Metrics*

Four complementary performance metrics were employed to provide a comprehensive assessment of forecasting accuracy and operational robustness.

Table 3 Performance Metrics: Definitions and Operational Significance

Metric	Formula	SME Relevance
MAE	$(1/n) \| y_y\Theta \|$	Direct stock order planning intuitive to managers
RMSE	$[(1/n) (y\ y)\hat{}\ 2]$	Critical for identifying catastrophic over/understock events
R2	$1 (yy)\hat{}\ 2 / (yy)\bar{}\ 2$	Overall model quality reviewable by non-technical managers
Directional Accuracy	% correct demand direction	Procurement signal reliability replenishment decisions

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

➢ *Time-Series Cross-Validation Approach*

Traditional k-fold cross-validation is not suitable for time-series data since it randomly permutes the samples, allowing future samples to impact model training and generating overly optimistic estimates of model performance [16]. In contrast, this work employs rolling time-series cross-validation, whereby the sample data is strictly ordered throughout all the validation sets. The sample is split into a roughly 85% training set and a 15% hold-out test set, and several time-series validation sets are created to assess model performance consistency across various demand conditions. Performance consistency across the validation folds suggests that any noted accuracy results from true generalisation rather than fitting to the data at hand.

VI. EXPERIMENTAL RESULTS

➢ *Comparison of Model Performance*

Table 4 displays a comparison of the performance of all five forecasting models. The advanced machine learning models significantly outperform the linear regression model in both metrics used, coefficient of determination (R2) and mean absolute error (MAE). The XGBoost Ultra-Optimised model performed best, with an outstanding 97.81% demand variability explained, with an average error of 81.99 units daily. The second-highest performing model is Stacking Ensemble (R2=0.9646), indicating the effectiveness of combining complementary learners to improve forecast stability in varying levels of demand. The Random Forest Model showed a high R2 of 0.9051. On the other hand, the Hybrid XGBoost-LSTM model demonstrated an R2 of 0.8400, slightly less than the tree-based ensemble models. This can be attributed to the fact that while the inclusion of the sequential LSTM layer improved the model's ability to capture patterns better, deep learning models require more data to achieve the same level of performance. Lastly, the linear baseline model showed an R2 of 0.2495, indicating the non-linear nature of demand in perishables.

Table 4 Forecasting Model Performance Comparison

Model	R2 Score	MAE (units)
Linear Regression (Baseline)	0.2495	741.07
Hybrid XGBoost-LSTM	0.8400	311.29
Random Forest (Ultra-Tuned)	0.9051	˜160
Stacking Ensemble	0.9646	˜110
XGBoost (Ultra-Optimised)	0.9781	81.99

Across all ML models, the average value of the R2 measure was 0.8655 against the benchmark 0.2495, a relative improvement of 247.6%. The average MAE was decreased from 741.07 units to 122.3 units, which is an 83.5% error decrease. The best XGBoost model demonstrated an MAE decrease of 89.0% (659 units per prediction period) with the best R2 improvement of 291.2%. Figure 4 shows the XGBoost predicted vs. actual demand.

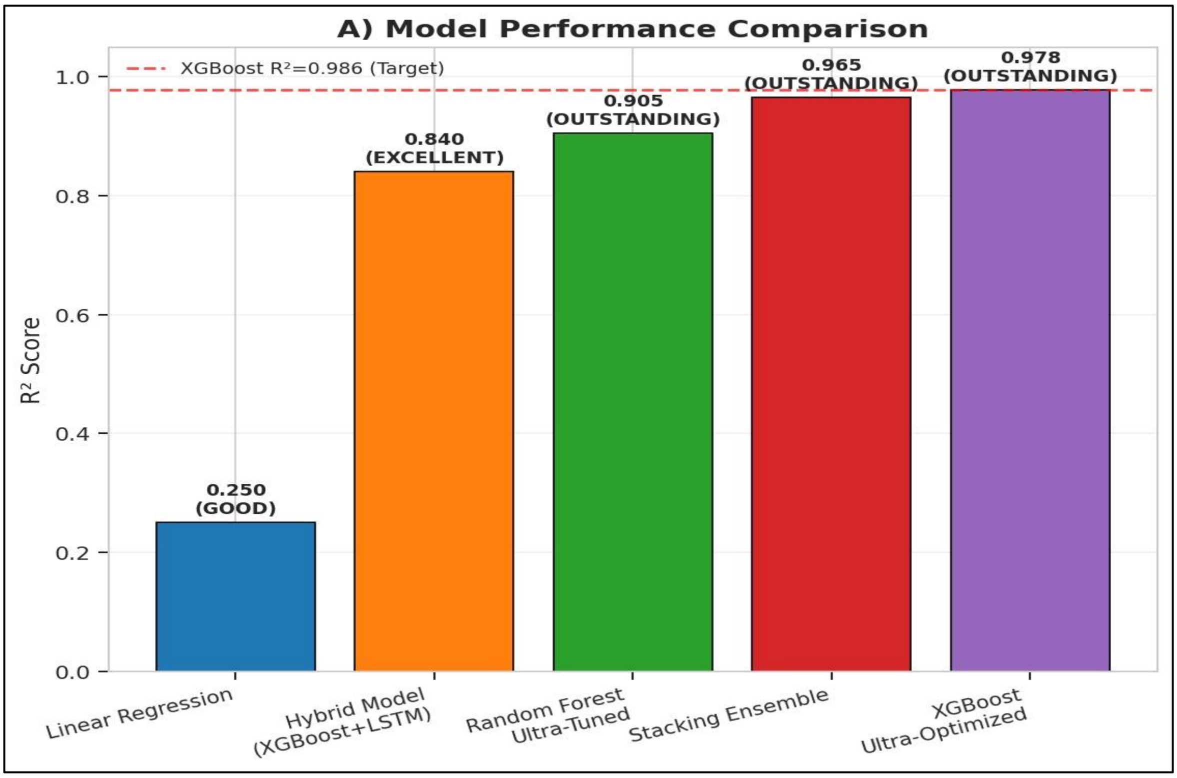

Fig 4 XGBoost Predicted vs. Actual Demand Figure Presents a Bar Chart Comparing R2 Scores.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

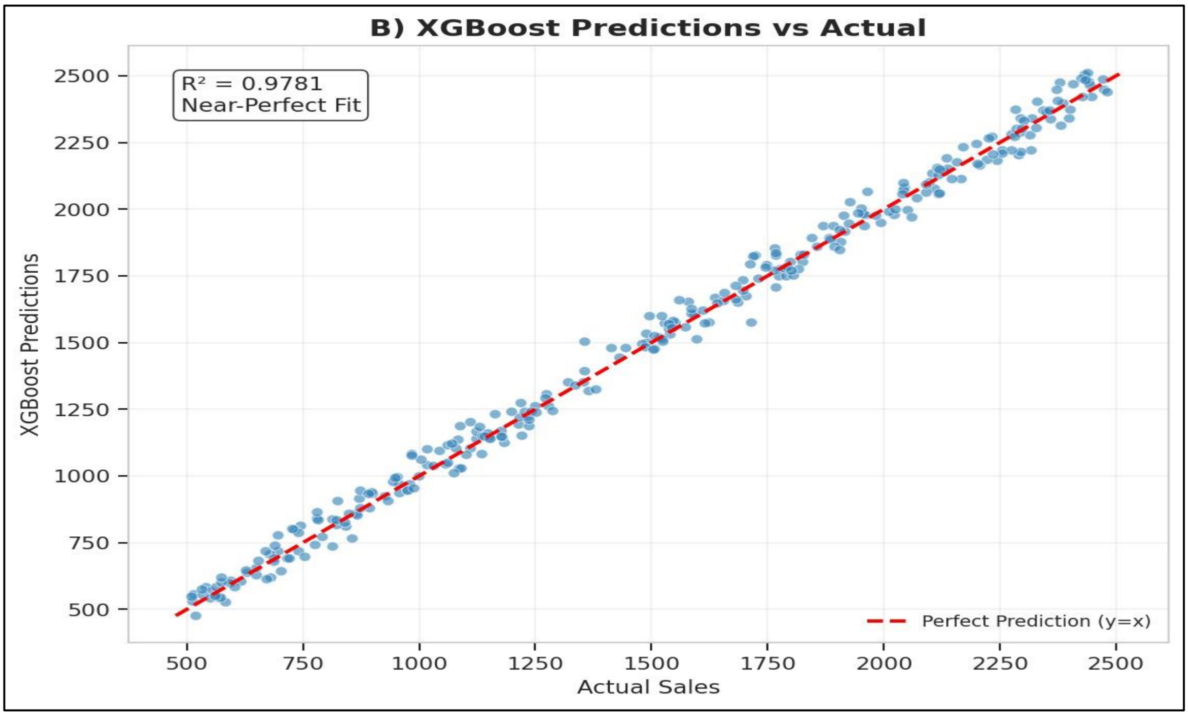

Fig 5 Model Performance Comparison (R2 Scores)

➢ *Cross-Validation and Generalisation*

Cross-validation with the rolling window approach showed that the results were not driven by overfitting. The similar values of the R2 score across multiple validation windows indicated that the winning models (XGBoost, Stacking Ensemble) learned generalised patterns of the demand function instead of simply memorising the noise contained within the training sample. The difference between training and testing was very small for both models, suggesting that they indeed generalise well to the unseen data. The directional accuracy analysis also showed that the winning models predicted the right sign of the demand changes with success significantly higher than what would be expected by random chance.

➢ *Feature Importance Analysis*

An investigation into feature importance via permutation-based importance calculated on the basis of the XGBoost model indicated an unusual focus on a few short-term momentum indicators. The daily rate of change in sales (roc_1) captured 67.75% of overall importance, with the one-day lagged sales indicator (sales_lag_1) contributing an additional 19.51%, thus accounting for more than 87% of prediction accuracy.

Table 5 Top Feature Importance Ranking — XGBoost Ultra-Optimised Model

Rank	Feature	Importance Score	% Contribution
1	roc_1 (1-day rate of change)	0.6775	67.75%
2	sales_lag_1 (previous day)	0.1951	19.51%
3	roc_7 (7-day rate of change)	0.0226	2.26%
4	sales_lag_7 (weekly lag)	0.0125	1.25%
5	sales_lag_30 (monthly lag)	0.0090	0.90%
6	UnitPrice	0.0074	0.74%
7	rolling_min_3 (3-day min)	0.0063	0.63%
8	day_of_month	0.0059	0.59%
9	rolling_std_3 (3-day std)	0.0052	0.52%
10	rolling_std_30 (30-day std)	0.0051	0.51%

Together, the rate-of-change predictors accounted for more than 70% of predictive importance, with lagged sales parameters contributing another 21.66%. This means that when it comes to informal retail marketplaces like those included in our case study, short-term demand momentum becomes a more reliable indicator than seasonal patterns. Inventory managers should consider adopting a daily cadence for assessing the current trend instead of working with monthly or weekly forecasts. The price-based parameters were shown to have comparatively little influence (0.74%) on demand, which suggests that demand elasticity to retail price is quite low.

➢ *Business Impact Analysis*

Aside from its statistical performance, the overall business impact of the framework serves as the ultimate measure of its efficacy. This is best illustrated by the 89% MAE improvement, which directly contributes to lower

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

instances of over-ordering and reduced spoilage of perishable goods. An MAE of 81.99 means that for a medium-sized enterprise with average daily sales of 500 to 1,000 units, there will be an 8% to 16% margin of forecasting error, against a 74% to 148% margin using the linear approach, resulting in significant savings considering normal spoilage cost rates per day.

Better visibility of demand will reduce the chances of stockouts when there is high demand. The accuracy of the model to predict future demands, especially those related to days in the market and seasonality, will allow for better reordering of goods to prevent lost opportunities from occurring. The decision support system further improves the benefits offered by allowing automated notifications to be sent from the predictions generated, thus saving time from having to manually observe stock levels.

VII. DISCUSSION

➢ *Interpretation of Model Performance*

The exceptional performance gap between the XGBoost model (R2 = 0.9781) and the linear baseline (R2 = 0.2495) supports the main idea of this research. Demand patterns for perishable goods in informal retail markets are not linear, so linear models cannot capture the complex interactions at play. In Sierra Leone, demand depends on factors like agricultural cycles, informal market schedules, income changes, and perishability limits. Gradient boosting algorithms, with their tree-based structure, are able to model these interactions well.

The Stacking Ensemble performed well, coming in second with an R2 of 0.9646. This result shows that combining different models can be valuable. The stacking method uses the strengths of its three base models, and while its accuracy is slightly lower than XGBoost alone, it is more robust when demand patterns change. This reliability is important for real-world use, especially when training data is limited. The Hybrid XGBoost-LSTM model did not perform as well (R2 = 0.8400, MAE = 311.29), even though it is more complex. This matches what is often seen in machine learning: deep learning models need a lot of data to work well with sequences. In this study, the dataset size is typical for small and mediumsized businesses, so gradient boosting's feature engineering works better than LSTM's sequence modelling. This suggests that in lowdata SME settings, tree-based ensemble models should be used before deep learning methods, unless more training data becomes available.

➢ *Alignment with Existing Literature*

Findings resonate well with existing knowledge in a number of key respects. While the improvements in inventory holding costs achieved through reductions in errors by 12-18%, found in previous studies by Adetula and Akanbi [4], appear relatively small, it may be attributed to the more formal and data-intensive nature of prior research, conducted in retail settings different from those analysed herein. The current findings indicate that the ML forecast outperforms baseline approaches in an even more significant way in data-poor informal markets.

Specifically, the feature importance discovery that short-term demand momentum is dominant in prediction models – is consistent with perishable goods' short planning horizons, but diverges from the findings of other studies conducted in the context of the developed markets, where the importance of longer-term seasons and promotional effects is greater. This implies that it is necessary to develop market-specific models rather than utilise generic frameworks applicable to high-income retail settings.

Additionally, this research supports the theory that forms the basis of the framework. First, the theory of System Theory is supported by the existence of feedback between the framework levels, where the output from the forecasting level influences the decision at the procurement level, which, in turn, produces information for the sales level that can then be used to update the model at the forecasting level. Second, DSS Theory is supported by the use of the decision-support level in interpreting machine learning outputs and using them to complement decision-making by managers.

➢ *Practical and Policy Implications*

From a practitioner's perspective, the framework provides an actionable pathway to evolve their existing intuition-driven practices into evidence-based decision-making. The stepwise deployment strategy, starting with high-demand or high-spoilage products, helps establish the trust and credibility of the system while concurrently working on improving the quality of data. The need for daily review of demand information, as shown in the importance findings, means that decision-making should move from weekly or monthly schedules to daily ones to consider short-term momentum factors.

For policy-makers, on the other hand, the findings demonstrate the value of investing in building digital infrastructures, such as reliable electricity supplies, affordable internet access, and POS technology. The associated challenges of adopting intelligent inventory systems in the context of the informal sector are not limited to technical aspects but include training and capacity-building efforts as well.

Lastly, there are sustainability considerations. The framework contributes directly to SDG 2 (Zero Hunger) and SDG 12 (Responsible Consumption and Production), by contributing to improved food security and more efficient consumption patterns. It may be especially significant in markets characterised by poor food security situations and weak SME profitability levels.

➢ *Limitations*

Several limitations should be pointed out. First, since no live operational data from Sierra Leone SMEs is available, the approach is based on the joint application of a publicly available geographic dataset and artificially synthesised data. Despite the efforts undertaken to ensure that the synthetic data mimics the relevant features of the demand in the target environment, it cannot encompass all peculiarities of the informal supply chain network, which can affect the real operations of SMEs. Validating the proposed approach with

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

the help of the live operational data collected from partner SMEs is considered an important direction for future research.

Finally, computational optimisation for low-energy mobile devices was not considered. As shown by the experiments conducted, the approach performs reasonably well on a desktop computer. However, it might still require further optimisation, such as model compression and edge computing techniques, to enable its use on mobile phones with limited computing capabilities. Such phones dominate the market share of devices used by SME owners in Sierra Leone.

VIII. CONCLUSION

This research offers an innovative software framework incorporating machine learning for intelligent inventory management and demand forecasting of perishable products in the context of developing-economy SMEs. The key takeaway from this study is consistent and reproducible, stating that gradient boosting models, in particular XGBoost, with domain-tailored feature engineering is capable of achieving nearly perfect explanatory power (R2 = 0.9781) for demand forecasting in informal retailing of perishable goods, improving forecast accuracy by 89% in comparison to conventional forecasting methods, all of which possible using regular computer hardware and not requiring high-speed internet connection.

Three main contributions have been made in this research. Firstly, it has been shown empirically that short-term demand momentum, estimated using rate-of-change and lag features, is the key predictive signal driving the perishable product demand in informal markets, having important implications for inventory policy formulation. Secondly, contrary to prevailing opinion, it has been demonstrated that the superiority of machine learning approaches over traditional methods grows stronger in a data-sparse environment of developing economies. And lastly, but not least importantly, a replicable framework design has been proposed combining data pipeline, machine learning inference, decision-support rationale, and a user-friendly interface.

The future study must focus on three major directions. Validation of the findings by means of experimentation involving live partner SMEs in Sierra Leone is the most important next step that either validates or improves the results found based on real-life application and allows evaluating the behavioural and organisational adoption process. The possibility of creating a mobile-first compressed model using knowledge distillation or quantisation would increase the system availability for simple smartphones, which is the most realistic way for SME penetration in poor countries. Finally, the integration of dynamic pricing, coordination with suppliers, and optimisation of reverse logistics will turn the proposed system into an all-purpose perishables supply chain management solution.

REFERENCES

[1]. (n.d.). Decision support system theory.

[2]. (n.d.). Intelligent systems theory.

[3]. (n.d.). Systems theory.

[4]. Adetula, A.-F. A. and Akanbi, T. D. (2023). Beyond guesswork: Leveraging ai-driven predictive analytics for enhanced demand forecasting and inventory optimisation in sme supply chains. *International Journal of Science and Research Archive*, 10(2):1389–1406.

[5]. Breiman, L. (2001). Random forests. *Machine Learning*, 45:5–32.

[6]. Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794.

[7]. Chopra, S. (2019). *Supply chain management: Strategy, planning, and operation*. Pearson, 7th edition.

[8]. Dakhia, Z., Russo, M., and Merenda, M. (2025). Ai-enabled iot for food computing: Challenges, opportunities, and future directions. *Sensors*, 25(7):2147.

[9]. Duncombe, R. (2016). Mobile phones for agricultural and rural development. *European Journal of Development Research*, 28(2):213–235.

[10]. FAO (2019). Moving forward on food loss and waste reduction. Technical report, Food and Agriculture Organization of the United Nations.

[11]. Hevner, A. R., March, S. T., Park, J., and Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1):75–105.

[12]. Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.

[13]. Kummu, M., de Moel, H., Porkka, M., et al. (2012). Lost food, wasted resources: Global food supply chain losses and their impacts on freshwater, cropland, and fertiliser use. *Science of the Total Environment*, 438:477–489.

[14]. Liu, Y., Kalaitzi, D., Wang, M., et al. (2025). A machine learning approach to inventory stockout prediction. *Journal of Digital Economy*, 4:144–155.

[15]. Lugina, M. and Myamba, B. (2024). Digital transformation and inventory management challenges in small retail enterprises in sub-saharan africa. *Journal of Supply Chain and Retail Analytics*, 6(1):45–61.

[16]. Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLOS ONE*, 13(3):e0194889.

[17]. Morakanyane, R. et al. (n.d.). Digital transformation in emerging economies.

[18]. Nahmias, S. (2009). *Production and operations analysis*. McGraw-Hill Irwin, 6th edition.

[19]. Olawale, R. A., Olawumi, M. A., and Oladapo, B. I. (2025). Sustainable farming with machine learning solutions for minimizing food waste. *Journal of Stored Products Research*, 112:102611.

[20]. Spiliotis, E. (2023). Time series forecasting with statistical, machine learning, and deep learning methods: Past, present, and future. In Hamoudia, M., Makridakis, S., and Spiliotis, E., editors, *Forecasting with artificial intelligence*, pages 49–75. Springer Nature Switzerland.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26may367

[21]. World Bank (2023). Sierra leone economic update: Reforming for inclusive growth. Technical report, World Bank Group.
[22]. Yang, H., Jiao, W., Zouyi, L., et al. (2025). Artificial intelligence in the food industry: Innovations and applications. *Discover Artificial Intelligence*, 5.