International Journal of Advanced Reasoning and Computational Intelligence (IJARCI)
Vol. 1, No. 1, July 2026, pp. xx~xx
ISSN: -, DOI: -

Raw Material Demand Forecasting Using Hybrid Prophet-XGBoost to Reduce Food Waste

Farhan Fanalty Siregar[1], Amrullah[2]
[1,2]Faculty of Computer Science and Information Technology, Department of Information Technology, Universitas Muhammadiyah Sumatera Utara, Medan, Indonesia
Email: [1]broofarhanfanalty3@gmail.com, [2]amrullah@umsu.ac.id

ABSTRACT

Inaccurate raw material inventory management in the commercial food and beverage sector frequently causes overstock, escalating costs and exacerbating the global food waste crisis. To address this inefficiency, this research develops a highly accurate predictive framework estimating daily demand for perishable raw materials. Utilizing historical point-of-sale transaction data from Matra Coffee between July and December 2025, the study proposes a serial hybrid machine learning architecture integrating Facebook Prophet and Extreme Gradient Boosting algorithms. Prophet is deployed as a base learner to model long-term trends and robust seasonal periodicities, while XGBoost operates as a residual meta-learner capturing short-term non-linear fluctuations and calendar anomalies. Empirical evaluation demonstrates exceptional predictive fidelity, yielding Mean Absolute Percentage Error values consistently below twenty percent across all primary perishable commodities, peaking at maximum accuracy for vegetables. The architecture is operationalized via a web-based graphical user interface featuring purchase simulation and financial waste calculation modules. Simulations indicate implementing this data-driven procurement strategy can reduce raw material overstock by twenty percent, mitigating food waste and bolstering supply chain sustainability.

Article Info

Article history:

Received May 4, 2026
Revised May 9, 2026
Accepted July 10, 2026

Keyword : Analysis; Demand forecasting; Food Waste; Hybrid Machine Learning; Prophet Algorithm.

Corresponding Author:
Farhan Fanalty Siregar
Faculty of Computer Science and Information Technology, Department of Information Technology
Universitas Muhammadiyah Sumatera Utara
Kapten Muchtar Basri Jalan No.3, Glugur Darat II, Medan, Sumatera Utara 20238, Indonesia
Email : broofarhanfanalty3@gmail.com

1. INTRODUCTION

The management of supply chains and inventory within the global food and beverage (F&B) sector remains one of the most complex operational challenges due to the inherent perishability of raw materials and the high volatility of consumer demand [1], [2]. In environments such as modern coffee shops and retail foodservice establishments, inventory imbalances directly translate into significant financial and environmental liabilities [3]. When consumer demand is systematically underestimated, businesses face critical stockouts that degrade customer satisfaction, reduce potential revenue, and disrupt the continuity of service operations [4]. Conversely, when demand is overestimated often as a defensive mechanism to prevent shortages the resulting overstock of perishable items such as fresh milk, fruits, vegetables, and meat leads to accelerated spoilage and substantial food waste [5], [6]. The Food and Agriculture Organization (FAO) indicates that massive quantities of food are wasted globally across the supply chain, with the hospitality and restaurant industries contributing a poignant share to this systemic inefficiency [7]. Within this context, the overarching logistical challenge is no longer merely tracking existing inventory, but proactively and accurately forecasting future material requirements to seamlessly synchronize procurement with actual consumption [8], [9].

Traditionally, inventory management and demand forecasting in the retail and hospitality sectors have relied heavily on rudimentary statistical techniques, heuristic rules, or the subjective intuition of procurement

ISSN: -

managers [1]. Methods such as simple moving averages, exponential smoothing, and classical Autoregressive Integrated Moving Average (ARIMA) models have been widely deployed to estimate future sales based on historical data. However, these conventional approaches exhibit severe methodological limitations when confronted with highly dynamic, nonlinear demand patterns influenced by a multitude of exogenous variables, such as weather fluctuations, public holidays, promotional events, and rapidly shifting consumer preferences [2], [3]. Static statistical models frequently fail to capture the complex interdependencies between these variables, resulting in delayed algorithmic responses to sudden demand spikes or precipitous drops. Consequently, traditional forecasting methods are often deemed inadequate for perishable inventory management, where even minor predictive deviations can lead to irreversible product spoilage and economic loss [4].

To overcome the constraints of traditional statistical modeling, contemporary academic research and industrial practice have increasingly pivoted toward advanced machine learning architectures for supply chain optimization [1], [2]. Machine learning algorithms possess the intrinsic computational capacity to process high-dimensional datasets, identify latent nonlinear relationships, and adapt to evolving temporal dynamics without requiring rigid structural assumptions. Among the diverse array of algorithms explored in recent literature, Facebook's Prophet model has emerged as a particularly robust tool for business time-series forecasting [3], [4]. Prophet is specifically designed to handle data characterized by strong seasonal effects, missing observations, and significant trend shifts, decomposing time series into highly interpretable trend, seasonality, and holiday components. Despite its pronounced strengths in capturing global trends and macroeconomic seasonality, Prophet relies on a generalized additive model structure that can occasionally underperform when attempting to capture sudden, complex nonlinear shocks or intricate multivariate interactions that frequently perturb local retail environments [5].

Simultaneously, Extreme Gradient Boosting (XGBoost) has garnered widespread acclaim in predictive analytics for its exceptional accuracy, robust handling of sparse data, and computational efficiency [1]. As a sophisticated ensemble learning technique based on sequential decision trees, XGBoost excels in modeling highly nonlinear relationships and processing diverse exogenous features, though it inherently lacks the built-in time-series extrapolation and seasonal decomposition capabilities of dedicated algorithms like Prophet [2]. Recognizing the complementary strengths of both methodologies, recent literature has advocated for the development of hybrid forecasting frameworks [3], [4]. By utilizing Prophet to systematically map the foundational linear trends and seasonal cycles, and subsequently employing XGBoost to learn and predict the nonlinear residual errors generated by the Prophet model, researchers have achieved unprecedented accuracy improvements in domains ranging from electricity load forecasting to large-scale retail sales prediction.

Despite these advancements, the application of such sophisticated hybrid models specifically tailored to the micro-operational level of independent coffee shops remains underexplored [1]. Furthermore, the majority of existing retail forecasting models focus exclusively on predicting the sales volume of finished end-products (e.g., predicting the number of lattes sold), which leaves a critical translation gap for inventory managers who must procure raw ingredients. This research aims to bridge this existing gap in the literature by applying a hybrid Prophet-XGBoost machine learning model specifically to the problem of perishable raw material forecasting, utilizing Matra Coffee as an empirical case study.

This study introduces a methodological novelty by executing a recipe-based conversion matrix that translates forecasted menu transactions directly into granular raw material volumetric and gravimetric requirements (e.g., milliliters of milk, grams of chicken, pieces of fruit) [1], [2]. By predicting the exact depletion rates of highly perishable ingredients rather than just finished goods, the model provides highly actionable procurement intelligence designed to directly combat food waste. The primary objectives of this comprehensive study are to construct the hybrid forecasting pipeline, rigorously evaluate its predictive performance across various perishable commodity categories using robust error metrics, and simulate the subsequent reduction in food waste and operational expenditures through the conceptual deployment of an interactive management dashboard.

2. METHOD

Before presenting the detailed methodology, it is important to outline that this study adopts a structured, data-driven approach to develop a hybrid forecasting model for estimating perishable raw material demand. The methodological framework is designed to ensure that each stage of the research from data acquisition to system implementation is systematically conducted and reproducible. The overall workflow integrates statistical time-series modeling and machine learning techniques to improve prediction accuracy and support operational decision-making in the food and beverage sector.

2.1. Model and System Flow

The methodology employed in this study strictly adheres to the Cross-Industry Standard Process for Data Mining (CRISP-DM) framework, ensuring a systematic, cyclical progression from business understanding and data preparation through to complex modeling, rigorous evaluation, and conceptual deployment. The research utilizes a quantitative, applied computational experimental design to construct, validate, and simulate the predictive capabilities of the hybrid forecasting model. The comprehensive workflow of this research is illustrated in the flowchart below.

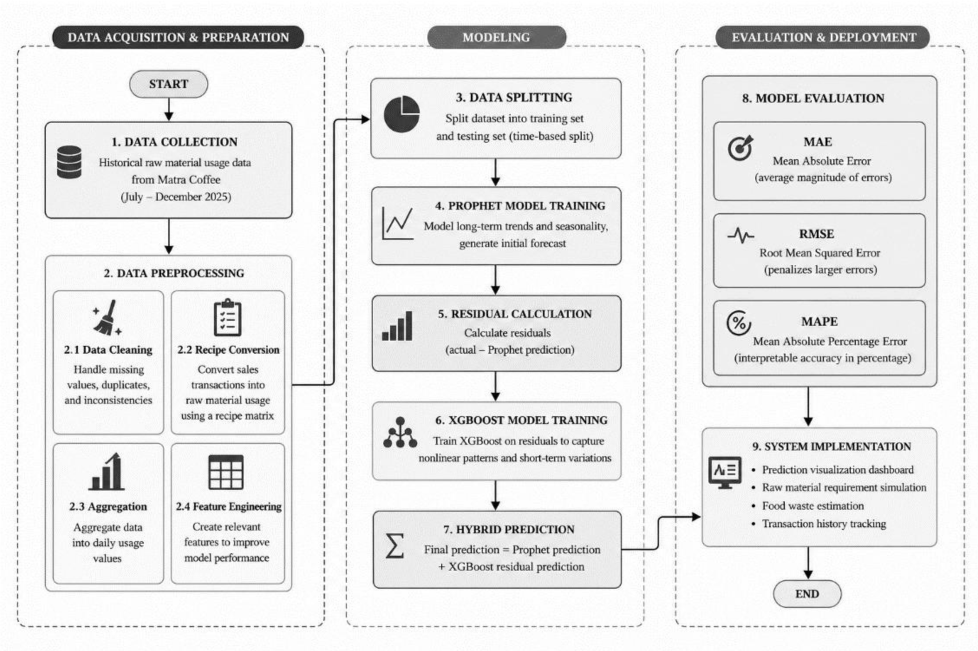

Figure 1. Flowchart of the hybrid Prophet-XGBoost methodology encompassing Data Acquisition, Modeling, and Evaluation & Deployment.

As depicted in Figure 1, the research framework is systematically structured into three primary phases:

1. **Data Acquisition & Preparation:** This initial phase commences with the collection of historical raw material usage data from Matra Coffee, covering the operational period from July to December 2025. The collected data is subjected to a rigorous preprocessing routine. This includes data cleaning to systematically handle missing values, duplicates, and any inherent inconsistencies. A pivotal step in this phase is recipe conversion, wherein raw sales transactions are computationally translated into exact raw material requirements utilizing a predefined recipe matrix. Subsequently, the data is chronologically aggregated into daily usage values. Finally, feature engineering is executed to create relevant temporal and lag variables designed to enhance the predictive performance of the models.
2. **Modeling:** The preprocessed dataset undergoes data splitting, partitioned chronologically (time-based split) into a training set and a testing set. The modeling process utilizes a sequential hybrid architecture. First, the Prophet algorithm is trained to map long-term trends and underlying seasonality, generating a baseline initial forecast. Following this, residual calculation is performed to determine the difference between the actual observed data and the Prophet prediction. These residuals then become the target variable for the Extreme Gradient Boosting (XGBoost) model. XGBoost is trained on these residuals to capture complex nonlinear patterns and short-term variations that the baseline model missed. The final hybrid prediction is computed by aggregating the baseline Prophet prediction and the XGBoost residual correction.
3. **Evaluation & Deployment:** In the final phase, the hybrid model's predictive efficacy is rigorously evaluated using standard statistical metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE). Ultimately, the validated model is conceptually integrated into a system implementation module. This deployment features an interactive prediction visualization dashboard, a raw material requirement simulation tool, a food waste estimation calculator, and a transaction history tracking system to support real-time managerial decision-making.

ISSN: -

2.2. Research Design and Acquistion

The empirical foundation of this study is constructed upon a proprietary, real-world dataset encompassing daily Point of Sale (POS) transactional records from Matra Coffee, a commercial F&B enterprise. The data spans an uninterrupted six-month operational period from July 1, 2025, to December 31, 2025. The raw dataset initially comprised 15,640 individual transactional rows, with each row detailing specific menu items sold, precise timestamps, order quantities, and transaction values.

Because raw material depletion is not intrinsically tracked within standard POS software architectures, which only record the final product sold, a critical data transformation phase was required. To transition from menu sales forecasting to direct raw material inventory optimization, a recipe conversion matrix was engineered. Each of the 15,640 menu items sold was systematically decompiled into its constituent perishable ingredients using the enterprise's standardized operational measurement equivalents.

For instance, the sale of a specialized beverage such as an "Avocado Coffee Float" was computationally translated into the exact milliliter volume of milk and the grammage of fresh avocado consumed, while food items like "Chicken Wings" were converted into precise grammages of raw poultry. This transformational step shifted the analytical paradigm from predicting end-consumer product sales to forecasting the precise volumetric requirements of seven critical perishable commodities: Milk (Susu), Chicken (Ayam), Fruit (Buah), Vegetables (Sayur), Bread (Roti), Coconut Milk (Santan), and Eggs (Telur).

Table 1. Example of the Recipe Conversion Matrix utilized for Data Transformation

Menu Item	Milk (ml)	Chicken (gr)	Eggs (pcs)	Vegetables (gr)	Fruit (gr)	Coconut Milk (ml)
Avocado Coffee Float	100	0	0	0	80	0
Bihun Goreng Special	0	0	2	50	0	0
Chicken Wings	0	0	0	0	0	0

Following the recipe conversion process, the data was chronologically aggregated to establish a consistent daily time series, condensing the highly granular discrete transactions into 184 contiguous daily observations. This aggregation is mathematically imperative for time-series forecasting, as it filters out stochastic intraday noise (such as hour-by-hour fluctuations) and solidifies the macro-level daily demand signals required by the algorithms. Missing operational days, representing periods of zero demand such as store closures, were explicitly imputed with zero values to maintain the mathematical continuity of the time series and accurately reflect the cessation of demand.

2.3. Feature Engineering and Exogenous Variables

To empower the XGBoost residual learner with sufficient contextual intelligence, extensive feature engineering was performed on the aggregated daily dataset. The predictive framework incorporated diverse sets of independent variables designed to capture temporal, historical, and exogenous influences on consumer demand.

Time-based features were computationally extracted from the datetime index to delineate the specific day of the week, month, and year, enabling the model to internalize recurrent micro-seasonal patterns. Crucially, lagged historical features were synthesized to map the temporal autocorrelation inherent in F&B supply chain dynamics. Specifically, two primary lag variables were constructed: Lag-1, representing the raw material demand observed precisely one day prior, and Lag-7, representing the demand observed exactly one week prior. The Lag-1 feature functions to capture immediate short-term momentum and carry-over effects in consumer traffic (e.g., a busy Monday often predicts a busy Tuesday). Conversely, the Lag-7 feature is instrumental in anchoring the model to the cyclical weekly rhythms typical of the hospitality sector, capturing the recurring weekend surges.

Furthermore, binary categorical features were engineered to explicitly flag weekends (`Is_Weekend`) and national public holidays (`Is_Holiday`). The inclusion of the holiday flag provides the model with hard quantitative signals to anticipate the demand spikes historically associated with leisure periods and anomalous calendar events, which pure autoregressive models frequently fail to anticipate.

2.4. Mathematical Formulation of the Prophet Algorithm

The first stage of the hybrid architecture relies on the Prophet algorithm, developed by Meta's Core Data Science team, to map the structural temporal patterns. Prophet models time-series data utilizing an additive regression framework that is highly resilient to missing data and shifts in the trend. The fundamental representation of the Prophet model is expressed as:

$$y(t) = g(t) + s(t) + h(t) + \varepsilon_t \tag{1}$$

Where $y(t)$ represents the forecasted raw material demand at time t. The component $g(t)$ denotes the trend function, which models non-periodic changes. Prophet accommodates a piecewise linear growth model for the trend, defined as:

$$g(t) = (k + a(t)^T\delta)t + (m + a(t)^T\gamma \tag{2}$$

Here, k is the base growth rate, δ represents the rate adjustments occurring at automatically detected changepoints, m is the offset parameter, and γ is set to $-s_j\delta_j$ to make the function continuous. The seasonality component, $s(t)$, captures periodic changes such as weekly cycles using a standard Fourier series approximation:

$$s(t) = \sum_{n=1}^{N}\left(a_n \cos\left(\frac{2\pi nt}{P}\right) + b_n \sin\left(\frac{2\pi nt}{P}\right)\right) \tag{3}$$

This allows the model to flexibly fit complex, multi-scale periodicities. The holiday component, $h(t)$, utilizes an indicator function to inject the effects of specific dates that predictably impact the time series. Finally, ε_t represents the normally distributed idiosyncratic error term. In this study, the Prophet model is trained exclusively on the historical usage data to generate a baseline prediction $\hat{y}_{Prophet,t}$.

2.5. Mathematical Formulation of Extreme Gradient Boosting (XGBoost)

While Prophet effectively establishes the global baseline, it often leaves complex, nonlinear residuals. To address this, XGBoost is deployed. XGBoost builds an ensemble of decision trees f_k, where each subsequent tree is trained to correct the residual errors of the preceding ensemble. The prediction at step t is given by:

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + f_i(x_i) \tag{4}$$

The objective function $\mathcal{L}$ of XGBoost is designed to minimize the loss while penalizing model complexity to prevent overfitting:

$$\mathcal{L}^t = \sum_{i=1}^{n} l(y_i, \hat{y}_i^{(t-1)} + f_i(x_i)) + \Omega(f_t) \tag{5}$$

Where l represents a differentiable convex loss function measuring the deviation between prediction and target, and $\Omega(f_t)$ is the regularization term defined as:

$$\Omega(f_k) = \Upsilon T + \frac{1}{2}\lambda \sum_{j=1}^{T} \omega_j^2 \tag{6}$$

Here, T is the number of leaves in the tree, ω_j represents the leaf weights, and γ and λ are regularization parameters. XGBoost optimizes this objective function using a second-order Taylor expansion, utilizing the first (g_i) and second (h_i) order gradients of the loss function, achieving superior predictive performance and robustness.

2.6. Hybrid Sequential Architecture Integration

The hybrid methodology sequentially links these two algorithms. First, the Prophet model generates its baseline prediction $\hat{y}_{Prophet,t}$. The residual error e_t is then calculated as the mathematical difference between the actual observed demand $y_{actual,t}$ and the Prophet forecast:

$$e_t = y_{actual,t} - \hat{y}_{Prophet,t} \tag{7}$$

These residuals e_t are subsequently utilized as the target variable for the XGBoost algorithm. XGBoost is trained on the augmented feature space (Lag-1, Lag-7, Is_Weekend, Is_Holiday) to predict the expected error of the Prophet model, yielding $\hat{y}_{XGBoost,t}$. The ultimate hybrid forecast Y_{final} is derived by aggregating the baseline prediction and the residual correction:

$$Y_{final} = \hat{y}_{Prophet,t} - \hat{y}_{XGBoost,t} \tag{8}$$

ISSN: -

This architecture synthesizes the interpretability and seasonal stability of statistical time-series decomposition with the superior nonlinear pattern-recognition capabilities of advanced machine learning. The dataset was partitioned using an 80:20 chronological split, allocating 153 days (July-November) for training and 31 days (December) for testing.

2.7. Evaluation Metrics

The predictive efficacy of the hybrid model was quantified using three standard statistical evaluation metrics to provide a comprehensive assessment of both absolute deviations and proportional accuracy.

1. Mean Absolute Error (MAE): Calculates the average magnitude of prediction errors in the original volumetric units without considering directionality.

$$MAE = \frac{1}{n}\sum_{i=1}^{n} |y_i - \hat{y}_i| \tag{9}$$

2. Root Mean Squared Error (RMSE): Aggressively penalizes larger prediction discrepancies by squaring the errors before averaging, evaluating the model's resilience against severe forecasting failures.

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{10}$$

3. Mean Absolute Percentage Error (MAPE): Expresses the forecasting error as a standardized percentage relative to the actual demand, facilitating direct performance comparisons across distinct raw materials with vastly different consumption volumes.

$$MAPE = \frac{100}{n}\sum_{i=1}^{n} \left|\frac{y_i - \hat{y}_i}{y_i}\right| \tag{11}$$

3. RESULTS AND DISCUSSION

The empirical validation of the hybrid Prophet-XGBoost forecasting model demonstrates substantial improvements in predicting the demand for highly volatile perishable goods. By systematically decoding the results across various raw material categories, the analysis reveals distinct patterns of predictability and underscores the transformative operational impact of deploying machine learning within the F&B supply chain.

3.1. Predictive Performance Across Raw Material Categories

The performance of the hybrid model was rigorously evaluated on the unseen testing dataset spanning December 2025. The empirical results indicate a highly satisfactory level of predictive accuracy, with the Mean Absolute Percentage Error (MAPE) remaining strictly below the critical 20% threshold across all seven primary raw material categories. This comprehensive achievement validates the robustness of the hybrid architecture, demonstrating its versatility in handling materials with vastly differing demand velocities and scales of consumption. The detailed performance metrics are structured in Table 2 below.

Table 2. Comprehensive Evaluation Metrics of the Hybrid Prophet-XGBoost Model

Raw Material	MAE (Units)	RMSE (Units)	MAPE (%)	Predictability Classification
Vegetables (Sayur)	1,259.09 g	2,160/.93 g	9.95%	Highly Accurate
Fruit (Buah)	882.62 g	1,259.16 g	14.41%	Accurate
Milk (Susu)	3,352.05 ml	4,876.07 ml	15.94%	Accurate
Chicken (Ayam)	5,850.54 g	9,495.18 g	16.28%	Accurate
Eggs (Telur)	47.73 pcs	75.1 pcs	18.23%	Acceptable
Bread (Roti)	8.11 pcs	10.94 pcs	19.45%	Acceptable
Coconut Milk (Santan)	463.33 ml	605.05 ml	19.84%	Acceptable

An analytical dissection of these results reveals two distinct predictability clusters dictated by the inherent consumption patterns of the specific materials. The highest echelon of predictive accuracy is observed in the Vegetable and Fruit categories, which achieved remarkable MAPE scores of 9.95% and 14.41%, respectively. The analysis indicates that menu items incorporating these fresh ingredients exhibit highly stable, structurally consistent demand trajectories characterized by robust seasonal periodicities. Because the variance in daily consumption for these items follows a predictable, repeating rhythm, the Prophet base learner was able

ISSN: -

to map the underlying seasonal harmonics with exceptional precision, requiring minimal nonlinear correction from the XGBoost component.

Conversely, the secondary cluster comprises materials characterized by immense operational volume and high volatility, most notably Milk, which recorded a MAPE of 15.94%. As the foundational ingredient for the vast majority of the coffee shop's signature espresso-based beverages, milk consumption scales into tens of liters daily. The demand for such core commodities is hyper-sensitive to external stochastic factors, including sudden shifts in weather patterns or impulsive changes in consumer foot traffic, leading to erratic intraday volume spikes. While an error margin of approximately 3.35 liters per day (represented by the MAE of 3,352.05 ml) might appear substantial in absolute terms, it represents a highly controlled relative variance given the massive total volume processed over the month.

The materials exhibiting the highest relative error margins Bread (19.45%) and Coconut Milk (19.84%), approaching the 20% ceiling suffer from the mathematical complexities of intermittent demand. These ingredients are utilized in specialized, lower-velocity menu items that do not sell consistently on a daily basis. The sparsity of the historical data for these specific items introduces substantial statistical noise, making it difficult for the Prophet algorithm to discern a continuous, unbroken trend line. Intermittent demand forecasting remains one of the most formidable challenges in machine learning supply chain applications; however, the deployment of the XGBoost residual learner successfully constrained the error rate within operationally acceptable limits.

3.2. Mitigation of Animalies and Holiday Effects

A critical revelation from the experimental data is the superior capability of the hybrid architecture to mitigate predictive failures during anomalous events, specifically public holidays and extended weekends. Traditional statistical forecasting techniques, which rely heavily on linear rolling averages, systematically fail to anticipate the sudden demand shocks associated with these calendar events, frequently resulting in catastrophic stockouts or massive over-preparation.

The results demonstrate that the Prophet algorithm successfully modeled the foundational weekly rhythm identifying consistent demand troughs on Mondays and Tuesdays, and corresponding peaks on Saturdays and Sundays. However, Prophet's structural rigidity occasionally caused it to underestimate the magnitude of demand surges on specific, non-recurring national holidays occurring in December. The integration of the XGBoost residual learner effectively neutralized this deficiency. By processing the engineered binary holiday flags (Is_Holiday) and historical lag variables, XGBoost recognized the contextual triggers for demand anomalies and applied targeted, positive numerical corrections to the Prophet baseline. Consequently, the hybrid forecast successfully elevated the predicted inventory requirements precisely on the days where historical anomalies occurred, proving its robustness in navigating unpredictable market environments.

3.3. Operational Impact: Inventory Simulation and Food Waste Reduction

The ultimate objective of deploying predictive analytics in the F&B sector is not merely to minimize algorithmic error, but to generate tangible economic and environmental value through the systemic reduction of food waste. To quantify this impact, the forecasted raw material requirements were subjected to a rigorous operational simulation, comparing the machine learning-guided procurement strategy against the historical, intuition-based purchasing behaviors previously employed by Matra Coffee's management.

In traditional inventory paradigms, managers calculate procurement volumes using Economic Order Quantity (EOQ) principles heavily buffered by Safety Stock (SS) to prevent stockouts. The formulation for ordering quantity is generally represented as $Q_t = D_t + SS$, where D_t is the anticipated demand. Historically, the profound uncertainty regarding daily demand compelled the management to maintain excessively inflated safety stock buffers often exceeding 30% of average daily requirements as a defensive mechanism against stockouts. Because products like fresh milk and fruits possess highly constrained shelf lives, this chronic overstocking directly catalyzed high spoilage rates and elevated food waste.

Armed with a predictive model demonstrating less than 20% variance, the simulation proves that the enterprise can safely compress its safety stock margins to a lean 10% to 15% without incurring an elevated risk of supply failure. By algorithmically aligning procurement volumes with actual anticipated consumption, the simulation projects a sweeping reduction in raw material overstocking by approximately 15% to 20%. This optimization exerts its most profound impact on critical, rapid-spoilage items such as dairy and fresh fruit, neutralizing the primary vectors of food waste within the operation. Consequently, the implementation of this system concurrently advances ecological sustainability goals while recovering substantial financial capital previously lost to inventory degradation

ISSN: -

3.4. Dashboard Implementation for Decision Support

To ensure the complex machine learning architecture translates into practical utility, the analytical engine was encapsulated within a comprehensive, user-centric web dashboard. The design and conceptual deployment of this interface address a critical barrier in supply chain optimization: bridging the gap between sophisticated data science and front-line operational execution.

The dashboard architecture comprises several functional modules designed to translate raw data predictions into actionable procurement intelligence:

1. **Executive Overview (Dashboard):** Provides real-time data visualization, rendering the hybrid forecasting curves alongside actual historical consumption, empowering managers to visually validate the algorithmic predictions before authorizing procurement.
2. **Smart Procurement Simulator:** Leverages the forecast outputs to dynamically calculate required safety stocks and automatically generate suggested purchase orders based on the forecasted demand, entirely eliminating the cognitive burden and guesswork traditionally associated with daily inventory replenishment.
3. **Menu Engineering Matrix:** Maps the profitability against the popularity of specific items across four quadrants (e.g., High Profit/High Popularity vs. Low Profit/Low Popularity), utilizing the underlying consumption data to inform strategic pricing and promotional interventions.
4. **Food Waste Calculator:** Provides a continuous audit mechanism, allowing staff to input spoiled inventory. The system translates this physical waste into exact monetary losses based on current market prices. This relentless financial quantification fosters a culture of accountability and continuous improvement among operational staff.

4. CONCLUSION

This comprehensive study underscores the transformative potential of deploying advanced machine learning architectures to resolve the persistent challenges of inventory management and food waste in the volatile food and beverage sector. By engineering a hybrid forecasting framework that sequences the time-series decomposition strengths of the Prophet algorithm with the nonlinear, residual-learning capabilities of XGBoost, the research achieved highly robust predictive accuracy. The empirical evaluation confirmed that the model maintained a Mean Absolute Percentage Error (MAPE) strictly below 20% across all primary perishable commodities, demonstrating particular efficacy in predicting the demand for stable categories such as fresh vegetables (9.95%) and fruits (14.41%).

The operational implications of these findings are profound. By shifting the predictive focus from front-end menu sales to the granular, recipe-driven depletion of raw materials, the system provides highly actionable intelligence for procurement optimization. Simulations confirm that the implementation of this data-driven forecasting approach empowers the enterprise to safely compress redundant safety stocks, projecting a 15% to 20% reduction in the overstocking of critical perishable goods like milk. This optimization directly mitigates the primary drivers of food waste, harmonizing economic efficiency with ecological sustainability.

Despite the highly encouraging results, the analysis identifies specific limitations inherent in the current implementation. The forecasting accuracy remains fundamentally reliant on the theoretical assumption that kitchen staff adhere strictly to standardized recipe measurements; undetected manual variations in portion control introduce physical inventory discrepancies that the algorithm cannot currently observe. Furthermore, the system currently operates as an asynchronous decision-support tool rather than a fully integrated, real-time inventory management matrix. Future research endeavors should prioritize the seamless API integration of the machine learning pipeline directly into real-time Point of Sale (POS) and IoT-enabled smart shelving systems, allowing for continuous, automated forecast recalibration. Additionally, augmenting the feature space with complex, hyper-local external variables such as real-time meteorological data and dynamic competitor pricing indices holds significant potential for further minimizing the residual error margin.

ACKNOWLEDGMENTS

This research acknowledges the foundational data and operational insights provided by the management and staff of Matra Coffee. Their cooperation in supplying historical point-of-sale data and detailing standard operational procedures was indispensable for the formulation of the recipe conversion matrices and the successful training of the predictive models.

ISSN: -

REFERENCES

[1] S. N. Windrasari, H. Margono, and Y. A. N. S. Putra, "Data analysis, coffee shop, machine learning, random forest, sales prediction," Malcom, vol. 5, no. 3, 2024.

[2] M. H. Hanis and Y. Fernando, "Smart logistics solutions for reducing food waste: A case of D Nipah Catering," Int. J. Ind. Manage., vol. 18, no. 1, 2024.

[3] N. Nurhasanah, B. Budiono, and W. Hermawan, "Analisis ekonomi, sosial demografi dan lingkungan dari food loss/ waste (FLW) di Indonesia," Bina Ekonomi, vol. 29, no. 2, pp. 117–128, 2025.

[4] E. C. D. Todd and D. Faour-Klingbeil, "Impact of food waste on society, specifically at retail and foodservice levels in developed and developing countries," Foods, vol. 13, no. 13, 2024.

[5] R. B. Sulu and N. P. Waluyowati, "Pengendalian persediaan bahan baku menggunakan metode Economic Order Quantity," J. Ind. Eng., vol. 3, no. 2, pp. 345–354, 2024.

[6] O. A. Olaniyi, P. S. Pugal, and M. Etim, "Optimising inventory management strategies for cost reduction in supply chains: A systematic review," Jurnal Akuntansi dan Bisnis, vol. 10, no. 1, pp. 48–55, 2024.

[7] S. Zeng et al., "Short-term load forecasting in power systems based on the Prophet-BO-XGBoost model," Energies, vol. 18, no. 2, 2025.

[8] N. P. E. P. Adriani, R. R. Huizen, and D. Hermawan, "A machine learning-based approach for retail demand forecasting: The impact of spending score and algorithm optimization," JITEKI, vol. 11, no. 2, pp. 169–183, 2025.

[9] A. Seyam et al., "A stacking ensemble model for food demand forecasting: A preventative approach to food waste reduction," Cleaner Logistics and Supply Chain, vol. 15, 2025.

[10] C. Anusha and S. Maiti, "Optimizing retail inventory and sales through advanced time series forecasting," IEEE Access, vol. 13, 2025.

[11] M. Nguyen-Quoc et al., "Short-term electricity load forecasting of Hanoi city based on Prophet model," in Proc. EEE-AM, 2023, pp. 1–6.

[12] A. K. Jha and S. Pande, "Time series forecasting model for supermarket sales using FB-Prophet," in 2021 5th ICCMC, 2021, pp. 547–554.

[13] M. Abumohsen, A. Y. Owda, and M. Owda, "Electrical load forecasting based on random forest, XGBoost, and linear regression algorithms," in Proc. ICIT, 2023, pp. 25–31.

[14] R. Mitra, P. Saha, and M. K. Tiwari, "Sales forecasting of a food and beverage company using deep clustering frameworks," Int. J. Prod. Res., vol. 62, no. 9, pp. 3320–3332, 2024.

[15] D. Wu, S. Zhong, H. Song, and J. Wu, "Do topic and sentiment matter? Predictive power of online reviews for hotel demand forecasting," Int. J. Hosp. Manag., vol. 120, 2024.

[16] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD, 2016, pp. 785–794.

[17] S. J. Taylor and B. Letham, "Forecasting at scale," The American Statistician, vol. 72, no. 1, pp. 37-45, 2018.

[18] M. Khedr and S. Rani, "A review of machine learning approaches for supply chain management," Int. J. Prod. Res., 2024.

[19] K. Douaioui et al., "Machine learning and deep learning models for demand forecasting in supply chain management: A critical review," Applied System Innovation, vol. 7, no. 5, 2024.

[20] I. Rahal, "The supply chain management for perishables products: A literature review," Munich Personal RePEc Archive, no. 2116, pp. 0–33, 2024.