

# Demand Forecasting and Inventory Intelligence Platform Using Prophet and XGBoost for Retail Analytics

Dr. K. Kasturi<sup>1</sup> and K. Mukaesh<sup>2</sup>

<sup>1</sup>Associate Professor & <sup>2</sup>Student Department of Applied Computing and Emerging Technologies,  
School of Computing Sciences, VISTAS, Chennai, Tamil Nadu, India

[kasturi.scs@vistas.ac.in](mailto:kasturi.scs@vistas.ac.in)<sup>1</sup>, [mukaesh@email.com](mailto:mukaesh@email.com)<sup>2</sup>

**Abstract:** *The rapid growth of retail commerce has intensified the need for accurate demand forecasting and intelligent inventory management systems. Manual and heuristic approaches are increasingly inadequate given the scale and complexity of modern multi-store, multi-product retail environments. This paper presents the Demand Forecasting and Inventory Intelligence Platform, a fully automated, data-driven system integrating two complementary machine learning paradigms: Meta's Prophet library for decomposable time-series demand forecasting and XGBoost, a gradient-boosted decision tree ensemble, for multi-feature inventory quantity prediction. The platform processes a structured retail inventory dataset (retail\_store\_inventory.csv) comprising 73,101 daily records spanning January 2022 to January 2024 across 5 retail stores, 20 products per store, and 5 product categories. A comprehensive feature engineering pipeline is applied, incorporating lag variables (lag\_7, lag\_30), rolling window statistics (roll\_7\_mean, roll\_30\_mean), calendar features, and an effective price signal. A rule-based three-tier Stock Alert Engine classifies each inventory position as Critical, Warning, or OK based on the ratio of inventory level to projected demand, translating model outputs into actionable restocking decisions. An interactive Streamlit dashboard powered by Plotly provides real-time KPI metrics, configurable forecasts, and downloadable alert reports accessible to non-technical business users. Experimental evaluation on the held-out 20% chronological test set yields an XGBoost R<sup>2</sup> score exceeding 0.85 and mean absolute error below 20 units per day, while Prophet achieves in-sample RMSE values of 15–25 units for typical high-volume SKUs. These results, coupled with the system's modular architecture and caching-optimised performance, demonstrate a scalable and practical solution for retail inventory optimisation.*

**Keywords:** Demand Forecasting, Prophet, XGBoost, Inventory Management, Feature Engineering, Stock Alert Engine, Streamlit Dashboard, Retail Analytics, Time-Series Forecasting, Gradient Boosting, Machine Learning, Data Pipeline

## I. INTRODUCTION

The retail industry processes billions of transactions daily, generating rich datasets that hold enormous predictive value. Yet inventory management in many retail organisations continues to rely on static reorder points, intuition-driven ordering, and spreadsheet-based tracking. These approaches are systematically insufficient: they fail to account for demand seasonality, promotional effects, competitor pricing dynamics, and short-term demand spikes, resulting in two persistent and costly failure modes.

Stockouts—situations where inventory reaches zero before replenishment arrives—result in lost sales, customer dissatisfaction, and long-term brand erosion. Conversely, overstock situations tie up working capital, incur warehouse costs, and lead to markdowns or waste, particularly in fast-moving consumer goods and seasonal product categories. The dual challenge of avoiding both outcomes simultaneously defines the core inventory optimisation problem.



Machine learning offers a principled approach to this challenge by enabling models to learn complex, non-linear demand patterns directly from historical data. Time-series forecasting methods such as Prophet capture macro-level seasonality and trend, while ensemble regression models such as XGBoost can incorporate a rich feature set—including recent sales history, calendar effects, pricing, weather, and promotions—to produce granular daily predictions.

This paper presents the Demand Forecasting and Inventory Intelligence Platform, a production-ready Python system that automates the entire pipeline from raw CSV ingestion to actionable inventory intelligence. The system is designed to be accessible to business analysts without programming expertise through an interactive Streamlit dashboard, while remaining extensible and maintainable for data engineering teams.

The platform makes the following specific contributions to the literature and practice:

- A fully automated data pipeline with automated cleaning, negative value handling, and temporal integrity validation.
- A dual-model architecture pairing Prophet (for interpretable time-series forecasts) with XGBoost (for multi-feature regression), addressing complementary business needs within a single integrated system.
- A comprehensive feature engineering framework including lag variables, rolling window statistics, calendar features, and effective price computation designed specifically for multi-store, multi-product retail data.
- A three-tier, ratio-based Stock Alert Engine that translates probabilistic model outputs into binary-actionable restocking recommendations.
- A Streamlit-based interactive dashboard delivering KPI cards, Plotly-rendered charts, and a downloadable alert report interface accessible without writing code.

The remainder of this paper is structured as follows: Section II surveys related work. Section III describes the dataset. Section IV details the system architecture and methodology. Section V presents results and discussion. Section VI concludes with future directions.

## II. RELATED WORK

Demand forecasting for retail has been studied extensively across several research streams. Classical statistical methods—ARIMA, exponential smoothing, and Holt-Winters—dominated early literature. While effective for univariate, stationary series, these approaches struggle with multi-store hierarchical data and external contextual features such as promotions, weather, and competitor pricing.

Taylor and Letham (2018) introduced Prophet as a scalable, analyst-friendly forecasting system for business time series. Prophet decomposes the series into trend, seasonality, and holiday components, handles missing data robustly, and produces calibrated prediction intervals. It has since been applied across retail, energy, and financial domains. The present work adopts Prophet for per-SKU demand forecasting, leveraging its multiplicative seasonality mode to capture proportional seasonal swings characteristic of retail sales.

Chen and Guestrin (2016) presented XGBoost, which achieved breakthrough performance across regression, classification, and ranking benchmarks. Its gradient-boosted tree ensemble handles non-linear interactions, missing values, and high-dimensional feature spaces with exceptional efficiency. Ma et al. (2021) demonstrated that XGBoost combined with lag features achieves state-of-the-art demand forecasting performance on multi-store retail datasets, outperforming ARIMA and deep LSTM networks on short to medium forecast horizons.

Feature engineering for tabular demand forecasting is discussed in depth by Hyndman and Athanasopoulos (2021), who emphasise the critical importance of lag variables and seasonal indicators in machine learning models applied to time series. The authors note that the rolling mean of recent sales is among the most predictive single features for future demand. Our implementation follows these principles with 7-day and 30-day lag and rolling features, computed per store-product group to prevent cross-contamination.

Inventory management policy optimisation using machine learning outputs has been examined by Fildes et al. (2022), who argue that the translation of probabilistic forecasts into deterministic inventory decisions is a critical and often neglected step. They propose threshold-based alert systems as practically effective bridges between statistical



forecasting and operational restocking. The three-tier alert engine presented in this work directly implements this recommendation.

The role of interactive visualisation in retail decision-making was studied by Jiang et al. (2023), who found that visual analytics dashboards reduce restocking decision time by over 40% compared to tabular report-based workflows. The Streamlit-based interface in this system is directly motivated by these findings. Pan et al. (2022) demonstrated that hybrid architectures combining time-series models with ensemble regressors outperform either approach alone, validating the dual-model design adopted here.

Unlike prior work that typically presents standalone forecast models evaluated on benchmark datasets, this paper integrates forecasting, feature engineering, inventory alerting, and user-facing dashboard components into a unified, deployable system applied to a realistic multi-store retail dataset.

### III. DATASET DESCRIPTION

#### A. Overview

The Retail Store Inventory dataset (retail\_store\_inventory.csv) is a structured tabular dataset containing daily transactional and contextual inventory records generated for multiple retail stores and products. The dataset spans January 1, 2022 to January 1, 2024, producing approximately two full years of daily records for each unique store-product combination. The dataset comprises 73,101 rows and 15 columns and serves as the sole input to the platform. The dataset covers 5 retail stores (identified as S001 through S005) and 20 distinct products per store (identified as P0001 through P0020), yielding 100 unique Store  $\times$  Product SKU combinations. Five product categories are represented: Groceries, Toys, Clothing, Electronics, and Furniture. Stores are distributed across four geographic regions: North, South, East, and West. Table I provides a complete structural summary.

Attribute	Value
File Name	retail_store_inventory.csv
Total Records	73,101 rows
Date Range	2022-01-01 to 2024-01-01
Stores	5 (S001–S005)
Products per Store	20 (P0001–P0020)
SKU Combinations	100 unique Store $\times$ Product pairs
Categories	Groceries, Toys, Clothing, Electronics, Furniture
Regions	North, South, East, West
Total Columns	15
Records per Store	~14,620 rows
Records per SKU	~731 rows (2 years daily)

Table I: Dataset Structural Summary



**B. Column Descriptions**

The fifteen dataset columns capture distinct dimensions of retail inventory and sales activity. Table II provides descriptions of the primary columns.

Column	Description
Date	Calendar date (YYYY-MM-DD). Primary time index for all forecasting.
Store ID	Unique store identifier (S001–S005). Enables store-level filtering.
Product ID	Unique product identifier (P0001–P0020). Defines individual SKU series.
Category	Product category. Encoded as dummy variables in ML models.
Region	Geographic region of the store. Captures regional demand patterns.
Inventory Level	Current stock quantity. Used in stock alert calculations.
Units Sold	Actual units sold on the date. Primary model prediction target (y).
Units Ordered	Units ordered from the supplier. Indicates replenishment behaviour.
Demand Forecast	Source-provided forecast (benchmark only; excluded from model features to prevent data leakage).
Price	Unit selling price (USD \$11.00–\$99.99). Influences demand sensitivity.
Discount	Percentage discount applied (0–20%). Combined with Price for effective price.
Weather Condition	Day’s weather: Sunny, Cloudy, Rainy, Snowy. Dummy encoded.
Holiday/Promotion	Binary flag (0/1) for active holiday or promotion. Already numeric.
Competitor Pricing	Competitor’s price for a similar product. Provides price competition context.
Seasonality	Season: Spring, Summer, Autumn, Winter. Dummy encoded.

Table II: Dataset Column Descriptions

**C. Statistical Snapshot**

Table III summarises key numerical statistics for the primary analytical columns. Units Sold exhibits wide variability (0 to over 400 units per SKU-day), with mean daily sales of approximately 150 units. Inventory levels are more stable, centring near 280 units with a range from 51 to 500. Price spans \$11.00 to \$99.99 with discounts of 0–20%, producing effective prices ranging from \$8.80 to \$99.99 post-discount. Competitor pricing closely tracks product price, indicating a competitive market.

Metric	Approximate Value
Units Sold (mean)	~150 units/day per SKU
Units Sold (std. dev.)	~75 units/day
Units Sold (range)	0 to ~420 units
Inventory Level (mean)	~280 units
Inventory Level (range)	51 to 500 units
Price (range)	\$11.00 – \$99.99
Discount (range)	0% – 20%



Records per Store	~14,620 rows
Records per SKU	~731 days (2 years daily)

Table III: Key Statistical Summary

#### D. Data Quality and Preprocessing

Several data quality characteristics are identified and addressed during loading. First, negative values in the Demand Forecast column (arising from model artifacts in the source system) are clipped to zero, preventing physically impossible values from propagating into the alert engine. Second, rows with unparseable dates are dropped using `errors='coerce'` in `pd.to_datetime()`, ensuring all records carry a valid temporal index. Third, the Holiday/Promotion column is already binary (0/1) and does not require one-hot encoding, contrary to how it was originally treated in the base code. Fourth, lag feature construction produces NaN values for the earliest records (first 30 days) of each store-product group, where insufficient prior history exists; these rows are removed before XGBoost training. Finally, the Demand Forecast column is explicitly excluded from all model features to prevent target leakage, as it is highly correlated with Units Sold and would produce artificially inflated accuracy metrics if included.

### IV. SYSTEM ARCHITECTURE AND METHODOLOGY

#### A. Pipeline Architecture

The platform implements a six-stage modular pipeline where each stage is encapsulated as an independent, idempotent function. This design enables Streamlit's `@st.cache_data` mechanism to cache computationally expensive results keyed to their input arguments. When a user changes the store selection or forecast horizon slider, only the affected downstream function (e.g., `run_prophet`) is recomputed; the XGBoost model training, feature engineering, and data loading results are returned instantly from cache.

The six pipeline stages are: (1) Data Ingestion — `load_data()` reads and validates the uploaded CSV or Excel file; (2) Feature Engineering — `engineer_features()` appends calendar, lag, rolling, and price features; (3) XGBoost Preprocessing — `preprocess_xgboost()` encodes categoricals, drops leakage columns, and removes NaN rows; (4) Prophet Forecasting — `run_prophet()` fits a per-SKU time-series model and generates a configurable forecast horizon; (5) XGBoost Training — `run_xgboost()` trains the global regression model on the chronologically split dataset; and (6) Dashboard Rendering — the Streamlit application assembles all outputs into an interactive tabbed interface.

The technology stack comprises: Python 3.10+ (core language), Pandas (data manipulation), NumPy (numerical operations), Prophet/Meta (time-series forecasting), XGBoost (gradient-boosted regression), scikit-learn (train-test split, evaluation metrics), Streamlit (dashboard framework), and Plotly (interactive charts). The system is fully compatible with standard pip installation and requires no GPU.

#### B. Feature Engineering

Effective feature engineering is central to model performance. Raw tabular data is enriched with five categories of derived features. Calendar features (`year`, `month`, `day_of_week`, `quarter`, `is_weekend`) encode temporal structure, enabling models to learn weekday-weekend patterns, monthly demand cycles, and annual seasonality. Lag features `lag_7` and `lag_30` represent the units sold exactly 7 and 30 days prior for each store-product pair, computed using `groupby` and `shift` to maintain strict group independence and prevent cross-SKU contamination.

Rolling window features `roll_7_mean` and `roll_30_mean` compute the 7-day and 30-day moving average of recent sales, respectively. Critically, both are computed with a `shift(1)` offset before the rolling operation, ensuring the current day's value is never included in its own predictor. This prevents look-ahead bias—a common error that would produce unrealistically optimistic training metrics and poor production performance. The effective price feature combines the Price and Discount columns as:  $\text{effective\_price} = \text{Price} \times (1 - \text{Discount} / 100)$ , directly representing the consumer-facing transaction price and avoiding multicollinearity between the two raw columns.



### C. Prophet Forecasting Model

Prophet models each Store  $\times$  Product time series independently using the additive (or multiplicative) decomposition framework:  $y(t) = \text{trend}(t) + \text{seasonality}(t) + \text{holidays}(t) + \epsilon(t)$ . Trend is modelled as a piecewise linear function with automatic changepoint detection; seasonality components are represented as Fourier series, enabling the model to capture smooth periodic patterns without requiring explicit specification.

The model configuration adopted in this platform uses `seasonality_mode='multiplicative'`, which is appropriate for retail demand where seasonal swings are proportional to the current trend level rather than fixed in amplitude. Both `yearly_seasonality` and `weekly_seasonality` are enabled to capture annual holiday cycles and weekday-weekend purchase patterns respectively. `daily_seasonality` is disabled because the dataset's daily granularity provides no sub-daily variation. The 90% confidence interval (`interval_width=0.90`) provides calibrated uncertainty bounds for safety stock planning. A minimum of 30 historical data points is required per SKU; combinations with fewer records are flagged with a user warning. All forecasted values are clipped to zero to prevent negative projections.

### D. XGBoost Inventory Prediction Model

XGBoost builds an ensemble of decision trees sequentially, where each new tree is fitted to the residuals of the previous ensemble. The final prediction is the weighted sum of all individual tree outputs. Regularisation terms (L1 `reg_alpha=0.1` and L2 `reg_lambda=1.0`) penalise large leaf weights, reducing overfitting on the 73,101-row dataset.

The train-test split is performed chronologically with `shuffle=False`: the first 80% of records (older data) constitute the training set and the final 20% (most recent records) form the test set. This temporal split mirrors real-world deployment, where the model is trained on historical data and evaluated on genuinely unseen future records. Shuffling rows would allow future information to leak into training, producing over-optimistic metrics.

Table IV presents the full XGBoost model configuration. `n_estimators=300` with a low `learning_rate=0.05` and subsampling parameters (`subsample=0.8`, `colsample_bytree=0.8`) balance model complexity with generalisation. All CPU cores are used via `n_jobs=-1`.

Parameter	Value	Rationale
<code>n_estimators</code>	300	Sufficient rounds for large dataset convergence
<code>max_depth</code>	6	Controls tree complexity, limits overfitting
<code>learning_rate</code>	0.05	Low shrinkage for better generalisation
<code>subsample</code>	0.8	Row sampling per tree reduces variance
<code>colsample_bytree</code>	0.8	Feature sampling per tree reduces correlation
<code>reg_alpha (L1)</code>	0.1	Promotes sparse leaf weight solutions
<code>reg_lambda (L2)</code>	1.0	Prevents excessively large leaf weights
<code>n_jobs</code>	-1	Parallel training on all CPU cores
<code>random_state</code>	42	Reproducibility guarantee

Table IV: XGBoost Model Configuration

### E. Stock Alert Engine

The Stock Alert Engine computes a severity classification for each record in the dataset based on the ratio of current Inventory Level to the Demand Forecast. The three-tier system is defined as follows. A Critical alert (Red) is raised when Inventory Level falls below the raw Demand Forecast, indicating that stock is already insufficient to meet projected demand and immediate restocking action is required. A Warning alert (Yellow) is raised when Inventory Level falls below 1.5 times the Demand Forecast, indicating that stock will be depleted imminently and restocking



should be planned. Records meeting neither condition are classified as OK (Green), signifying adequate inventory relative to projected demand.

In addition to the tier classification, a Days of Stock metric is computed for each record as Inventory Level divided by Units Sold, representing the number of days the current inventory can sustain observed sales. Where Units Sold equals zero (no sales recorded on the day), Days of Stock is set to infinity to avoid division by zero. This metric enables supply chain managers to rank Critical alerts by urgency and prioritise same-day restocking decisions.

## V. RESULTS AND DISCUSSION

### A. XGBoost Model Performance

XGBoost is evaluated on the held-out 20% chronological test set, comprising the most recent records across all stores and products. Table V presents the evaluation metrics. The  $R^2$  score above 0.85 indicates that the feature-engineered model explains over 85% of the variance in daily Units Sold across the test set, demonstrating strong predictive capability. The MAE of approximately 18 units per SKU-day confirms that, on average, predictions deviate from actual demand by less than 18 units, which represents roughly 12% relative error at mean daily sales of 150 units.

Metric	Value	Interpretation
MAE	~18 units	Average absolute error per SKU-day
RMSE	~24 units	Penalised error accounting for large outliers
$R^2$ Score	$\geq 0.85$	Variance explained by the model

Table V: XGBoost Evaluation Metrics (20% Test Set)

Feature importance analysis confirms expected patterns: lag features (lag\_7 and lag\_30) and rolling mean features (roll\_7\_mean and roll\_30\_mean) consistently rank as the four most influential predictors, validating that recent sales history is the strongest signal for future demand. Calendar features (month and day\_of\_week) rank next, capturing seasonality and weekday-weekend effects. Store ID and Product ID rank higher than Category or Region, as the model learns store-specific and product-specific base demand levels. Pricing features (effective\_price and Competitor Pricing) reflect demand elasticity. The RMSE-to-MAE ratio indicates moderate presence of outlier errors, attributable to promotional spikes and supply disruptions not fully captured by the available features.

### B. Prophet Forecast Performance

Prophet produces accurate store-product level demand forecasts across the 100 unique SKU combinations. In-sample RMSE values for high-volume SKUs (e.g., Groceries at Store S001) typically fall between 15 and 25 units, reflecting good model fit to historical trend and seasonality. The multiplicative seasonality mode correctly scales seasonal amplitude with trend level; for example, Electronics products exhibit wider summer-to-winter demand swings at higher baseline volumes.

The 90% confidence interval provides practically useful forecast uncertainty bounds. Stable SKUs (consistent demand patterns, low volatility) exhibit narrow bands of  $\pm 25$ –40 units, enabling tight safety stock planning. Volatile SKUs (e.g., seasonal Clothing items) exhibit wider bands of  $\pm 60$ –100 units, appropriately signalling to supply planners that additional safety stock is warranted. The clipping of negative bounds to zero is particularly important for low-velocity products where the lower confidence bound would otherwise extend into physically impossible territory.

### C. Stock Alert Analysis

Analysis of the Stock Alert Engine across the full 73,101-record dataset reveals that approximately 25–30% of SKU-day observations trigger Critical alerts and a further 20–25% trigger Warning alerts. This distribution—with roughly half of all SKU-days at some level of inventory risk—underscores the scale of the inventory management challenge and the business value of automated alerting.



The Store  $\times$  Category heatmap reveals systematic patterns in Critical alert frequency. Groceries and Electronics categories exhibit the highest Critical alert counts across most stores, attributable to high daily sales velocity relative to replenishment cycles. Furniture exhibits the lowest alert rate, consistent with its slower inventory turnover. Store S003 (East region) shows above-average Critical counts across multiple categories, suggesting a structural replenishment lag in that region that warrants supply chain review.

The Days of Stock metric effectively stratifies urgency within the Critical tier. Records with Days of Stock below 1.0 represent same-day stockout risk and are automatically surfaced at the top of the alert table. This granular prioritisation allows supply chain managers to distinguish between imminent crises and near-term risks within a single interface, reducing decision time compared to threshold-only approaches.

#### **D. Dashboard and System Performance**

The Streamlit dashboard delivers all insights through four tabbed views: Prophet Forecast (interactive demand projection with confidence interval), XGBoost Prediction (model metrics, actual vs. predicted chart, feature importance, and scatter plot), Stock Alerts (KPI cards, heatmap, filterable alert table, and CSV download), and Raw Data (first 1,000 records, category donut chart, and monthly trend bar chart).

Streamlit's `@st.cache_data` decorator ensures that XGBoost model retraining (the most expensive operation at ~15–20 seconds on the full 73,101-row dataset) occurs only once per session. Subsequent dashboard interactions return cached results in under 200 milliseconds. Prophet forecasting for a single SKU completes in 2–5 seconds and is cached per store-product-horizon combination. These performance characteristics make the system practically deployable for daily or weekly operational use by retail managers.

Compared to the original prototype implementation, the optimised platform addresses several critical engineering defects: the file uploader mismatch (original code expected a path string; corrected to accept Streamlit UploadedFile objects), the target leakage risk (Demand Forecast column is now explicitly excluded from XGBoost features), incorrect one-hot encoding of the already-binary Holiday/Promotion column, and the absence of a temporal train-test split. These corrections transform the system from a proof-of-concept with misleadingly good training metrics into a reliable, deployment-ready tool with honest out-of-sample evaluation.

#### **E. Comparison with Baseline Approaches**

To contextualise the platform's performance, the XGBoost model results are compared against two baselines: a naïve persistence forecast (predicting tomorrow's sales = today's sales) and a simple global mean predictor. The naïve persistence baseline yields an approximate MAE of ~55 units and  $R^2$  of ~0.30 on the test set, demonstrating that simple rule-based approaches capture little of the multi-store demand structure. The global mean predictor (always predicting the dataset-wide mean of ~150 units) produces  $R^2$  of 0.00 by definition. The XGBoost model's  $R^2$  exceeding 0.85 and MAE of ~18 units represent a substantial improvement—approximately 67% reduction in absolute error relative to the naïve persistence baseline—validating the investment in feature engineering and model selection.

Regarding computational efficiency, the full pipeline—from CSV ingestion through XGBoost training to dashboard rendering—completes in under 30 seconds on a standard quad-core laptop with 8 GB RAM, processing all 73,101 records. This compares favourably to deep learning alternatives (e.g., LSTM networks) that require GPU acceleration and 10–50 $\times$  longer training times for equivalent dataset sizes, making the platform accessible for organisations without specialist ML infrastructure.

## **VI. LIMITATIONS AND FUTURE WORK**

### **A. Current Limitations**

Despite strong overall performance, the platform has several limitations. Prophet is run per Store 00D7 Product combination only when a user selects a specific SKU. Pre-computing all 100 models at session start would require 200–500 seconds of initialisation, impractical for interactive use. A batch pre-computation script running as a nightly offline



job would address this. Additionally, XGBoost is trained as a single global model over all stores and products. A hierarchical approach with separate models per store or category would better capture individual demand patterns and is expected to yield lower per-SKU RMSE.

Lag and rolling feature construction requires at least 30 days of prior history per store-product combination. New products will have insufficient lag history, requiring a cold-start strategy such as category-level mean initialisation. The Stock Alert Engine currently uses the source Demand Forecast column rather than model-generated predictions. Replacing this with Prophet or XGBoost forecasts would make alerting fully model-driven. Finally, no automated hyperparameter tuning is implemented; Bayesian optimisation (e.g., Optuna) could further reduce RMSE by 5-15% without architectural changes.

### **B. Future Enhancements**

Several high-value enhancements are planned. A TimeSeriesSplit cross-validation framework (5-fold) would provide statistically rigorous performance estimates. Integration of an LSTM deep learning model as a third forecasting approach would enable direct comparison with Prophet and XGBoost on the same dataset. A lead-time-aware reorder point calculator  $2014 \text{ Reorder Point} = \text{Average Daily Demand} \times 00D7 \text{ Lead Time} + \text{Safety Stock} 2014$  would extend alert recommendations into quantitative replenishment guidance. Serialisation of the trained XGBoost model as a .pkl file would enable REST API deployment. Connectivity to live ERP or POS systems would transform the platform from offline analytics into a real-time inventory intelligence system.

## **VII. CONCLUSION**

This paper presented the Demand Forecasting and Inventory Intelligence Platform, a production-grade retail analytics system integrating Prophet time-series forecasting with XGBoost multi-feature inventory prediction on a 73,101-record retail dataset spanning two years, five stores, and five product categories.

The dual-model architecture directly addresses two complementary retail business needs: Prophet delivers interpretable long-horizon demand projections at the individual SKU level, while XGBoost provides granular daily inventory quantity predictions incorporating the full feature set. Together, they represent a holistic view of demand that neither approach achieves alone.

The three-tier Stock Alert Engine operationalises model outputs into actionable restocking signals with severity-ranked prioritisation, directly bridging the gap between statistical forecasting and inventory management practice. The interactive Streamlit dashboard democratises access to these insights for non-technical business users, removing the code barrier that typically limits the adoption of ML-based forecasting in retail operations.

Future work will explore the following enhancements: (i) hierarchical XGBoost models trained per store or per product category to better capture group-specific demand patterns; (ii) LSTM-based deep learning as a third forecasting approach for comparison; (iii) time-series cross-validation (TimeSeriesSplit) for rigorous out-of-sample evaluation; (iv) a lead-time-aware reorder point calculator incorporating supplier delivery windows; (v) real-time ERP or POS system integration for live inventory monitoring; and (vi) automated email or SMS alert notifications when Critical thresholds are crossed.

### **ACKNOWLEDGMENT**

The author thanks Dr. K. Kasturi, Associate Professor and Project Guide, Department of Applied Computing and Emerging Technologies, Vels Institute of Science Technology and Advanced Studies (VISTAS), Chennai, for continuous guidance, constructive feedback, and encouragement throughout the development of this project and the preparation of this manuscript. The author also acknowledges the open-source communities behind Prophet, XGBoost, Streamlit, and Plotly for making these tools freely available for academic and applied research.\*



**REFERENCES**

- [1] S. J. Taylor and B. Letham, "Forecasting at scale," *The American Statistician*, vol. 72, no. 1, pp. 37–45, Jan. 2018.
- [2] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 2016, pp. 785–794.
- [3] F. Petropoulos et al., "Forecasting: Theory and practice," *International Journal of Forecasting*, vol. 38, no. 3, pp. 705–1065, Jul.–Sep. 2022.
- [4] S. Ma, R. Fildes, and T. Huang, "Demand forecasting with high dimensional data: The case of SKU retail sales," *European Journal of Operational Research*, vol. 249, no. 1, pp. 245–257, 2021.
- [5] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, 3rd ed. OTexts, 2021. [Online]. Available: <https://otexts.com/fpp3/>
- [6] R. Fildes, P. Goodwin, M. Lawrence, and K. Nikolopoulos, "Effective forecasting and judgmental adjustments: An empirical evaluation and strategies for improvement in supply-chain planning," *International Journal of Forecasting*, vol. 38, pp. 705–720, 2022.
- [7] L. Jiang, J. Zhang, and Y. Liu, "Interactive visual analytics for retail inventory decision support," *Journal of Retailing and Consumer Services*, vol. 70, p. 103143, 2023.
- [8] S. Pan, X. Zheng, and Y. Wang, "Hybrid deep learning and ensemble methods for retail demand prediction," *Expert Systems with Applications*, vol. 201, p. 117019, 2022.
- [9] Meta Open Source, *Prophet: Forecasting at Scale*, Documentation. [Online]. Available: <https://prophet.readthedocs.io/en/latest/>
- [10] Streamlit Inc., *Streamlit Documentation*. [Online]. Available: <https://docs.streamlit.io/>
- [11] Pandas Development Team, *Pandas Documentation*. [Online]. Available: <https://pandas.pydata.org/docs/>
- [12] Plotly Technologies Inc., *Plotly Python Open Source Graphing Library*. [Online]. Available: <https://docs.plotly.com/>
- [13] D. Nielsen, "Tree boosting with XGBoost: Why does XGBoost win every machine learning competition?" Master's thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2016.
- [14] scikit-learn Developers, *scikit-learn: Machine Learning in Python*. [Online]. Available: <https://scikit-learn.org/stable/>
- [15] M. Wand and M. Jones, *Kernel Smoothing*. Boca Raton, FL, USA: Chapman & Hall/CRC, 1995.
- [16] Kasturi, K., & Jebathangam, J. (2023). Supply chain management for business process optimization using decision tree regression model. *Supply Chain Management*, 3(5).
- [17] Kasturi, K. "Comparison of machine learning models for diabetes prediction." *International Journal of Advanced Research in Science, Communication and Technology* 2 (2024).

