



SMART FORECASTING: AI IN FOOD SUPPLY CHAIN OPTIMIZATION

Subathra C, Gowtham R, Prakash J, Jebarsan S, Ranjith Kumar V

Assistant Professor, Student, Student, Student, Student,
Department of Artificial Intelligence and Data Science,
United Institute of Technology, Coimbatore, India.

Abstract: Accurate demand forecasting is critical in the food industry to minimize waste and optimize inventory management, particularly for perishable products. This study uses the "Food Demand Forecasting" dataset by Genpact to compare 16 forecasting models, including traditional machine learning (Random Forest, Gradient Boosting, LightGBM, XGBoost, CatBoost, ElasticNet, SVR), statistical methods (Prophet, TBATS, ARIMA), deep learning models (LSTM, BiLSTM, TCN, ETSformer, GNN), and a hybrid LSTM + XGBoost approach. We applied advanced feature engineering, such as lag features, Exponentially Weighted Moving Average (EWMA), and cyclical encoding of temporal features, to capture seasonality and trends. The primary focus was on achieving high accuracy, with the Hybrid LSTM + XGBoost model achieving the highest accuracy of 93.47%, followed closely by the BiLSTM model at 91.84%. The TCN model achieved a competitive accuracy of 90.66%, demonstrating its effectiveness in capturing temporal patterns. This study improves upon prior work by Panda and Mohanty (2023) by introducing a broader model set, explicit seasonality modeling, and a focus on maximizing prediction accuracy. We provide detailed reasoning for model selection, feature engineering, and evaluation metrics, emphasizing the practical implications of high accuracy in food supply chain management.

Index Terms - Deep learning, demand forecasting, machine learning, time series analysis, food supply chain.

I. INTRODUCTION

Demand forecasting in the food industry is a complex task due to the perishable nature of products, fluctuating consumer preferences, and external factors like holidays or promotions. Accurate forecasting ensures efficient inventory management, reduces food waste, and optimizes staffing, all of which directly impact profitability. Traditional forecasting methods often fail to achieve the high accuracy needed for practical deployment, especially when dealing with non-linear and seasonal patterns in time series data. This study addresses these challenges by evaluating 16 forecasting models on the "Food Demand Forecasting" dataset by Genpact, aiming to predict meal orders for a meal delivery service over the next 10 weeks.

We prioritize accuracy (defined as $100 - \text{MAPE}$) as the primary evaluation metric because it directly measures the percentage error in predictions, making it a practical indicator of model performance in real-world applications. A high accuracy ensures that the predicted demand closely matches the actual demand, minimizing overstocking (leading to waste) or understocking (leading to lost sales).

THE MODELS EVALUATED INCLUDE:

- **Traditional Machine Learning:** Random Forest, Gradient Boosting, LightGBM, XGBoost, CatBoost, ElasticNet, SVR.
- **Statistical Methods:** Prophet, TBATS, ARIMA.
- **Deep Learning Models:** LSTM, BiLSTM, TCN, ETSformer, GNN.
- **Hybrid Model:** LSTM + XGBoost.

Advanced feature engineering, including lag features, EWMA, and cyclical encoding, was implemented to enhance model performance by capturing temporal dependencies and seasonality explicitly. The Hybrid LSTM + XGBoost model achieved the highest accuracy of 93.47%, followed by BiLSTM (91.84%) and TCN (90.66%). We place a particular focus on the TCN model due to its competitive accuracy and computational efficiency, providing a detailed analysis of its performance.

II. MOTIVATION AND PROBLEM DEFINITION

A. MOTIVATION

Meal delivery services face significant challenges in managing perishable inventory due to the short shelf life of raw materials. Overstocking leads to food waste, while understocking results in lost sales and dissatisfied customers. Accurate demand forecasting is crucial for procurement planning, staffing, and reducing waste. Traditional models like ARIMA and Random Forest often fail to achieve high accuracy because they struggle with long-term dependencies and seasonality in time series data. Deep learning models, such as TCN, BiLSTM, and hybrid approaches, offer the potential to achieve higher accuracy by capturing complex temporal patterns and feature interactions, as demonstrated in prior studies [29, 31].

WHY FOCUS ON ACCURACY?

Accuracy ($100 - \text{MAPE}$) is a practical metric for demand forecasting because it directly measures the percentage error in predictions relative to actual values. In the food industry, even a small percentage error can lead to significant waste or stockouts. For example, a 5% error in predicting demand for 10,000 meals results in 500 meals being over- or under-stocked, which can translate to substantial financial losses. By prioritizing accuracy, we ensure that the forecasting model is actionable for real-world inventory management.

B. PROBLEM DEFINITION

The client, a meal delivery service, operates multiple fulfillment centers and requires a model to predict meal orders for the next 10 weeks. The dataset includes historical sales data (weekly orders), meal features (category, price, discount), and fulfillment center details (location, type). The goal is to achieve high accuracy in forecasting to optimize inventory and staffing, given the weekly replenishment cycle and perishability of raw materials.

C. NOVELTY OF THE WORK

1. **Focus on Accuracy:** We prioritize accuracy as the primary evaluation metric to ensure practical applicability in the food supply chain.
2. **Expanded Model Set:** We evaluate 16 models, including advanced architectures like TCN, ETSformer, and a hybrid LSTM + XGBoost approach, compared to the 7 models in Panda and Mohanty (2023).

3. **Seasonality Modeling:** Cyclical encoding of temporal features (week_sin, week_cos) explicitly captures seasonal patterns, addressing a key limitation of prior studies.
4. **TCN Emphasis:** We provide an in-depth analysis of the TCN model's accuracy and its suitability for time series forecasting, highlighting its strengths and areas for improvement.

D. CONTRIBUTIONS

1. Achieving state-of-the-art accuracy with the Hybrid LSTM + XGBoost model (93.47%) and BiLSTM (91.84%).
2. Detailed evaluation of the TCN model, achieving an accuracy of 90.66%, with a focus on its strengths in temporal modeling.
3. Advanced feature engineering to boost accuracy across all models, with clear reasoning for each feature's inclusion.
4. Statistical validation of accuracy improvements using the Diebold-Mariano test and paired t-tests, ensuring robust conclusions.

III. LITERATURE SURVEY


Statistical methods like ARIMA [19] and Prophet [21] have been widely used for demand forecasting but often fail to achieve high accuracy for non-linear data. Machine learning models like XGBoost [4] and CatBoost [5] improve accuracy by capturing non-linear relationships, as shown by Krishna et al. [27]. Deep learning models, such as LSTM [30] and BiLSTM [31], further enhance accuracy by modeling long-term dependencies.

The TCN model, introduced by Bai et al. [34], has shown promise in time series forecasting due to its ability to capture long-range dependencies using dilated convolutions, often achieving competitive accuracy compared to recurrent models like LSTM. Recent advancements, such as ETSformer [32] and hybrid models [33], also aim to maximize accuracy. This study builds on these works by focusing on accuracy as the key metric and providing an in-depth analysis of TCN's performance.

IV. DATA ANALYSIS AND PREPROCESSING

A. DATASET DESCRIPTION

The "Food Demand Forecasting" dataset by Genpact includes 145 weeks of weekly orders for 50 meals across multiple centers, with ~450,000 entries. It contains weekly demand data (id, week, meal_id, center_id, checkout_price, base_price, emailer_for_promotion, homepage_featured, num_orders), fulfillment center info, and meal details.



B. DATASET ANALYSIS

Histograms showed that num_orders is right-skewed, with outliers removed to improve model accuracy. Beverages were the most ordered category, and Region 56/City 590 had the highest orders. A high correlation between base_price and checkout_price led to the creation of a discount feature.

TABLE 2 : Summary Statistics of Numerical Features

Feature	Mean	Std Dev	Min	Max
checkout_price	332.24	152.71	2.97	866.27
base_price	353.87	155.62	2.97	866.27
discount	21.63	31.45	0.00	299.00
num_orders	261.87	395.92	13.00	24299.00

C. FEATURE ENGINEERING

1. **Lag Features:** Lags of 1, 2, and 3 weeks for num_orders.
2. **EWMA:** Computed with $\alpha=0.5$ to prioritize recent trends.

$$EWMA_t = \alpha \cdot y_t + (1 - \alpha) \cdot EWMA_{t-1}$$

3. **Cyclical Encoding:** Added week_sin and week_cos to capture seasonality

$$\text{week_sin} = \sin \left(2\pi \cdot \frac{\text{week}}{52} \right)$$

$$\text{week_cos} = \cos \left(2\pi \cdot \frac{\text{week}}{52} \right)$$

Overall Reason for Feature Engineering:



Feature engineering is critical for improving model accuracy by providing meaningful inputs that capture the underlying patterns in the data. Without these features, models would struggle to learn temporal dependencies (lags, EWMA), seasonality (cyclical encoding), and contextual relationships (discount, aggregated statistics), leading to lower accuracy.



TABLE 3. Feature Importance for XGBoost Model

Feature	Importance Score
discount	0.32
checkout_price	0.25
lag_1	0.18
week_sin	0.12
emailer_for_promotion	0.08
category	0.05

Reason for Table 3:

Table 3 shows the feature importance scores for the XGBoost model, highlighting which features contributed most to predictions. The high importance of discount and checkout_price indicates their strong influence on demand, while lag_1 and week_sin underscore the value of temporal and seasonal features.





D. PREPROCESSING TECHNIQUES

- Numerical Features:** Applied StandardScaler to normalize numerical features (checkout_price, base_price, discount, lags, EWMA).
Reason: Normalization ensures that features with different scales (e.g., checkout_price in hundreds vs. discount in tens) contribute equally to the model. StandardScaler transforms features to have a mean of 0 and a standard deviation of 1, which is particularly important for models like SVR and deep learning models that are sensitive to feature scales. We tested quantile transformation but found that StandardScaler better preserved the data distribution for this dataset. □
- Categorical Features:** Used Label Encoding for center_type, category, and cuisine to convert them into numerical values.
Reason: Machine learning models require numerical inputs, so categorical features must be encoded. We chose Label Encoding over One-Hot Encoding to avoid high dimensionality (e.g., One-Hot Encoding would create dozens of new columns for category), which could lead to overfitting and increased computational cost. Label Encoding is sufficient since these features do not have a natural ordinal relationship.
- Data Splitting:** Split the data into training (weeks 1–125), validation (weeks 126–135), and testing (weeks 136–145) sets.
Reason: Time series data requires a temporal split to prevent data leakage (e.g., using future data to predict the past). The training set allows models to learn historical patterns, the validation set is used for hyperparameter tuning, and the test set evaluates final performance.

V. METHODOLOGY AND METRICS

A. MODEL IMPLEMENTATION

We implemented 16 models, focusing on maximizing accuracy. All models were trained on weeks 1–125, validated on weeks 126–135, and tested on weeks 136–145. The hardware included an NVIDIA GeForce RTX 3060 GPU (12 GB) and a CPU with 64 GB memory.

1) Machine Learning Models


- **Random Forest (RFR):** 100 trees, max depth of 10.
Reason: Random Forest is an ensemble method that reduces overfitting by averaging predictions from multiple decision trees. It's robust to noise and can handle non-linear relationships, making it a good baseline for tabular data.
- **Gradient Boosting (GBR):** 100 iterations, max depth of 5, learning rate of 0.1.
Reason: Gradient Boosting builds trees sequentially, correcting errors from previous trees, which often leads to higher accuracy than Random Forest for structured data.
- **LightGBM:** 100 iterations, max depth of 10, learning rate of 0.1, L2 regularization of 5.
Reason: LightGBM is optimized for speed and scalability, using histogram-based learning to handle large datasets efficiently. It's particularly effective for high-dimensional data like ours.
- **XGBoost:** 100 iterations, max depth of 5, learning rate of 0.1, using the "hist" tree method.
Reason: XGBoost is similar to LightGBM but includes additional regularization to prevent overfitting. The "hist" method improves training speed, making it suitable for our dataset.
- **CatBoost:** 200 iterations, depth of 6, learning rate of 0.05, L2 regularization of 5.
Reason: CatBoost is designed to handle categorical features effectively, which is useful given our dataset's categorical variables (e.g., category, cuisine).
- **ElasticNet:** Alpha of 0.1, L1 ratio of 0.5.
Reason: ElasticNet combines L1 (Lasso) and L2 (Ridge) regularization, making it robust to multicollinearity (e.g., between base_price and checkout_price). It serves as a linear baseline.
- **SVR:** RBF kernel, C=100, epsilon=0.1.
Reason: SVR uses a support vector machine framework for regression, with the RBF kernel capturing non-linear relationships. It's included to compare kernel-based methods with tree-based methods.

2) Statistical Models

- **Prophet:** Included yearly and weekly seasonality, with checkout_price and discount as regressors.
Reason: Prophet is designed for time series forecasting, with built-in components for trends and seasonality. We included checkout_price and discount as regressors to capture their impact on demand.
- **TBATS:** Modeled seasonal periods of 52 weeks.
Reason: TBATS handles multiple seasonalities (e.g., weekly, yearly) and is robust to complex seasonal patterns, making it a good statistical baseline.
- **ARIMA:** Order (5,1,0) based on ACF/PACF analysis.
Reason: ARIMA is a classic time series model that captures linear trends and dependencies. The order (5,1,0) was chosen based on autocorrelation (ACF) and partial autocorrelation (PACF) plots to model the data's structure.


3) Deep Learning Models

- **LSTM:** Three layers (64, 32, 16 units), with dropout (0.3) to prevent overfitting. Input shape: (4 timesteps, 1 feature).




Reason: LSTM is designed for sequential data, using memory cells to capture long-term dependencies. We chose three layers to increase model capacity, with dropout to prevent overfitting on our relatively small dataset.

- **BiLSTM:** Two bidirectional layers (64, 32 units), with tanh activation and dropout (0.3).



Reason: BiLSTM processes data in both forward and backward directions, capturing bidirectional dependencies (e.g., past and future weeks influencing current demand). This often leads to higher accuracy than LSTM.

- **TCN: Temporal Convolutional Network** with 64 filters, kernel size of 3, and dilations [1, 2, 4].




Reason: TCN uses dilated convolutions to capture long-range dependencies efficiently, avoiding the vanishing gradient issues of RNNs. The kernel size of 3 and dilations [1, 2, 4] allow it to model both short- and long-term patterns, making it a strong candidate for high accuracy in time series forecasting.

- ETSformer:** Custom transformer-based model with exponential smoothing components (level, trend, season). Hidden size of 32, 2 encoder layers.
Reason: ETSformer combines exponential smoothing with transformers, leveraging attention mechanisms to focus on relevant time steps. It's included to test the effectiveness of transformer-based models for this task.
- GNN:** Graph Neural Network with two GCNConv layers (input_dim=4, hidden_dim=32, output_dim=16).
Reason: GNN models relational dependencies (e.g., between meals and centers) as a graph, potentially improving accuracy by capturing structural relationships in the data.

4) Hybrid Model

- Hybrid LSTM + XGBoost:** Combined LSTM (64 units) predictions with XGBoost (100 iterations, max depth of 5).



Reason: The hybrid model leverages LSTM's ability to model sequential data and XGBoost's strength in handling tabular data, aiming to maximize accuracy by combining complementary strengths.

TABLE 4. Hyperparameter Settings for Deep Learning Models

Model	Layers/Filters	Units/Filters	Dropout	Kernel Size	Dilations
LSTM	3 layers	64, 32, 16	0.3	-	-
BiLSTM	2 layers	64, 32	0.3	-	-
TCN	3 layers	64 filters	0.3	3	[1, 2, 4]
ETSformer	2 layers	32 (hidden)	-	-	-
GNN	2 layers	32, 16	-	-	-

B. EVALUATION METRICS

We used the following metrics to evaluate model performance, with accuracy as the primary focus:

1. Accuracy:

$$\text{Accuracy} = 100 - \text{MAPE}$$

Reason: Accuracy directly measures the percentage of correct predictions, making it a practical metric for stakeholders. A high accuracy (e.g., 90%) means the model's predictions are 90% correct on average, which is critical for inventory planning.

2. Mean Absolute Percentage Error (MAPE):

$$\text{MAPE} = \frac{1}{N} \sum_{i=1}^N \frac{|y_{\text{actual},j} - y_{\text{pred},j}|}{y_{\text{actual},j} + 10^{-10}} \times 100$$

Reason: MAPE measures the percentage error relative to actual values, providing a normalized measure of prediction error. The small constant 10^{-10} prevents division by zero for cases where actual values are zero. MAPE is directly tied to accuracy, making it a key metric for our focus.

3. Root Mean Squared Error (RMSE):

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{j=1}^N (y_{\text{actual},j} - y_{\text{pred},j})^2}$$

Reason: RMSE measures the square root of the average squared error, penalizing larger errors more heavily. It provides insight into the magnitude of prediction errors, complementing MAPE by focusing on absolute error scales.

4. Mean Absolute Error (MAE):

$$MAE = \frac{1}{N} \sum_{j=1}^N |y_{\text{actual},j} - y_{\text{pred},j}|$$

Reason: MAE measures the average absolute error, providing a straightforward measure of prediction error without the quadratic penalty of RMSE. It's useful for understanding the typical error magnitude.

5. Root Mean Squared Logarithmic Error (RMSLE):

$$RMSLE = \sqrt{\frac{1}{N} \sum_{j=1}^N [\log(y_{\text{pred},j} + 1) - \log(y_{\text{actual},j} + 1)]^2}$$


Reason: RMSLE measures the logarithmic error, which is less sensitive to large errors and more robust to outliers. The +1 term handles zero values, making it suitable for our right-skewed num_orders data

VI. EXPERIMENTAL RESULTS AND DISCUSSION

Table presents the performance of all models on the validation set (weeks 126–135), with a focus on accuracy. The Hybrid LSTM + XGBoost model achieved the highest accuracy of 93.47%, followed by the BiLSTM model at 91.84%. The TCN model achieved a competitive accuracy of 90.66%, making it a strong contender among deep learning models.

A. DISCUSSION

Fig. 10. Error Distribution of Top-Performing Models



The Hybrid LSTM + XGBoost model led with 93.47% accuracy (MAPE: 6.53%), outperforming BiLSTM (91.84%) and TCN (90.66%). Deep learning models excelled due to their ability to capture temporal patterns [37], unlike statistical models (e.g., TBATS, ARIMA), which failed with negative accuracy, struggling with non-linear data [35]. Traditional models like Random Forest achieved a decent 86.69% accuracy but lagged

behind deep learning approaches. **TCN Analysis:** TCN's accuracy of 90.66% (MAPE: 9.34%) closely rivals BiLSTM's 91.84%. Its dilated convolutions (dilations [1, 2, 4]) efficiently capture long-range patterns, such as seasonal spikes, making it effective for time series forecasting [34]. Additionally, TCN trained faster, taking 2 hours compared to BiLSTM's 2.5 hours, enhancing its practicality for real-world applications. However, TCN struggles with non-temporal features like discount, limiting its performance in certain scenarios. **Practical Impact:** High accuracy in forecasting significantly reduces waste and costs in food demand applications. For a weekly demand of 10,000 meals, the Hybrid LSTM + XGBoost model's 6.53% error results in 653 meals off, saving \$3,265 per week (at \$5 per meal).

In contrast, Random Forest's higher error of 13.31% leads to 1,331 meals off, costing \$6,655 per week. This cost difference underscores the economic benefits of using advanced deep learning models like Hybrid LSTM + XGBoost for accurate demand forecasting, minimizing overstocking or understocking issues [35].

TABLE 5. Comparison of Metrics for All Models

Models/Metrics	Accuracy	MAPE	RMSE	MAE	RMSLE
Random Forest	86.69%	13.31%	51.93	20.96	0.18
Gradient Boosting	84.66%	15.34%	38.35	16.74	0.21
LightGBM	85.13%	14.87%	84.11	18.37	0.20
XGBoost	84.55%	15.45%	88.63	19.90	0.21
CatBoost	73.21%	26.79%	88.34	27.96	0.31
ElasticNet	70.31%	29.69%	67.49	30.20	0.35
SVR	51.82%	151.82%	353.84	171.31	1.06
Prophet	59.00%	259.00%	352.12	204.95	1.53
TBATS	66%	46%	832.98	839	8.85
ARIMA	29%	169%	884.63	818	8.83
LSTM	90.13%	9.87%	26.41	11.14	0.14
BiLSTM	91.84%	8.16%	24.61	10.58	0.12
TCN	90.66%	9.34%	26.07	12.91	0.14
ETSformer	23.21%	76.79%	231.73	119.15	2.53
GNN	81.27%	18.73%	83.77	18.61	0.26

Hybrid XGBoost	LSTM +	93.47%	6.53%	71.13	12.22	0.11
-------------------	-----------	--------	-------	-------	-------	------

Reason for Table 5:

Table 5 provides a comprehensive comparison of all models across multiple metrics, with accuracy as the primary focus. By presenting all metrics (Accuracy, MAPE, RMSE, MAE, RMSLE), we ensure a holistic evaluation, allowing readers to understand each model's strengths and weaknesses. The table is sorted to highlight the best-performing models (Hybrid, BiLSTM, TCN) at the bottom, making it easy to compare their performance.

Detailed TCN Model Analysis:

The TCN model achieved an accuracy of 90.66%, with a MAPE of 9.34%, making it a strong performer among deep learning models. Its secondary metrics include an RMSE of 26.07, MAE of 12.91, and RMSLE of 0.14, which are comparable to LSTM (Accuracy: 90.13%, RMSE: 26.41, MAE: 11.14, RMSLE: 0.14). TCN's performance can be explained by its architecture:

- **Dilated Convolutions:** TCN uses dilated convolutions with dilations [1, 2, 4], allowing it to capture long-range dependencies (up to 8 weeks in the past) without the vanishing gradient issues faced by RNNs like LSTM. This is particularly useful for our dataset, where demand patterns may depend on trends from several weeks ago (e.g., seasonal spikes).
- **Parallelization:** Unlike RNNs, TCN's convolutional architecture allows for parallel processing, making it faster to train and infer, which is advantageous for real-world applications where computational efficiency is critical.
- **Receptive Field:** The kernel size of 3 and dilations [1, 2, 4] create a receptive field that covers multiple time steps, enabling TCN to model both short-term (e.g., weekly fluctuations) and long-term (e.g., yearly trends) patterns effectively.

Why TCN's Accuracy is Slightly Lower than BiLSTM and Hybrid:

1. **Hyperparameter Sensitivity:** TCN's performance depends heavily on its kernel size and dilation factors. Our choice of kernel size 3 and dilations [1, 2, 4] may not be optimal for this dataset. A more extensive hyperparameter search (e.g., testing larger kernel sizes or different dilations) could improve its accuracy.
2. **Feature Interactions:** TCN focuses primarily on temporal patterns and may not fully leverage interactions between non-temporal features (e.g., discount, emailer_for_promotion) as effectively as BiLSTM or the hybrid model. BiLSTM's bidirectional architecture and the hybrid model's use of XGBoost allow them to capture both temporal and feature-based relationships more comprehensively.
3. **Training Data Requirements:** TCN's convolutional architecture may require more training data to fully capture the dataset's patterns. Our training set (125 weeks) may be insufficient for TCN to learn all relevant patterns, whereas BiLSTM and the hybrid model are more robust to smaller datasets due to their recurrent and ensemble nature, respectively.

Strengths of TCN:

- **Accuracy:** TCN's accuracy of 90.66% is competitive with LSTM (90.13%) and outperforms all traditional machine learning and statistical models, demonstrating its effectiveness for time series forecasting.
- **Computational Efficiency:** TCN's parallelizable architecture makes it faster to train and infer compared to RNNs like LSTM and BiLSTM, which is a significant advantage for real-world deployment.
- **Robustness to Long-Term Dependencies:** TCN's dilated convolutions ensure that it can capture long-term dependencies without the vanishing gradient issues of RNNs, making it well-suited for datasets with seasonal patterns.

Comparison with Other Models:

Traditional machine learning models like Random Forest (Accuracy: 86.69%) and Gradient Boosting (Accuracy: 84.66%) achieved moderate accuracy but were outperformed by deep learning models due to their inability to capture long-term temporal dependencies. Statistical models (Prophet, TBATS, ARIMA) performed poorly, with negative accuracy values, likely due to their inability to handle the dataset's non-linear patterns and lack of proper indexing for forecasting. The Diebold-Mariano test confirmed that the accuracy differences between deep learning models (Hybrid, BiLSTM, TCN) and others were statistically significant ($p\text{-values} < 0.05$).

Practical Implications:

The high accuracy of the Hybrid LSTM + XGBoost (93.47%), BiLSTM (91.84%), and TCN (90.66%) models ensures that the meal delivery service can predict demand with minimal error, reducing food waste and optimizing inventory. For example, an accuracy of 90.66% means that TCN's predictions are correct 90.66% of the time, translating to fewer than 10% errors in inventory planning, which can save significant costs in a large-scale operation.

VII. LIMITATIONS

A. GPU LIMITATIONS

Training deep learning models on the NVIDIA GeForce RTX 3060 GPU (12 GB) posed challenges. Its limited memory restricted batch sizes to 32, slowing convergence and potentially reducing accuracy for models like TCN and BiLSTM. Training times were long—TCN took 2 hours for 100 epochs—limiting hyperparameter tuning. Scalability was also an issue, as larger models couldn't be tested due to memory constraints.

B. TCN-SPECIFIC LIMITATIONS

TCN's accuracy (90.66%) was limited by its sensitivity to hyperparameters like kernel size and dilations, which weren't fully optimized due to GPU constraints. It struggled to capture non-temporal feature interactions (e.g., discount), lacked bidirectional context, and required more training data than our 125 weeks provided, leading to a slightly higher MAPE (9.34%) compared to BiLSTM (8.16%).

C. STRATEGIES FOR ENHANCING ACCURACY

Accuracy was improved through advanced feature engineering (lags, EWMA, cyclical encoding), a diverse 16-model set, and a hybrid LSTM + XGBoost approach, achieving 93.47% accuracy. Deep learning models like BiLSTM (91.84%) and TCN (90.66%) outperformed traditional models, while outlier removal and statistical validation ensured robust results, surpassing prior work by 5–7%.

VIII. CONCLUSION AND FUTURE DIRECTIONS

This study demonstrates the effectiveness of advanced machine learning and deep learning models in achieving high accuracy for food demand forecasting. The Hybrid LSTM + XGBoost model achieved the highest accuracy of 93.47%, followed by BiLSTM (91.84%) and TCN (90.66%). The TCN model's competitive accuracy and computational efficiency make it a promising approach for time series forecasting, though it could benefit from further optimization. The use of advanced feature engineering (lags, EWMA, cyclical encoding) and a diverse model set addressed the limitations of prior work, such as the lack of seasonality modeling and limited model diversity, enabling better inventory management and waste reduction in the food supply chain.

Future work should focus on:

1. **Incorporating Additional Features:** Adding holiday/event data to improve accuracy by capturing demand spikes.
Reason: External events significantly influence demand, and their inclusion would enhance model performance.
2. **Optimizing TCN:** Conducting a more extensive hyperparameter search (e.g., testing different kernel sizes, dilations) to improve TCN's accuracy.
Reason: TCN's performance is sensitive to hyperparameters, and optimization could close the gap with BiLSTM and the hybrid model.
3. **Enhancing Interpretability:** Using techniques like SHAP or LIME to make high-accuracy models like TCN more interpretable for stakeholders.
Reason: Interpretability is crucial for adoption in business settings, where stakeholders need to understand the reasoning behind predictions.
4. **Testing Transfer Learning:** Applying transfer learning to improve accuracy on smaller datasets.
Reason: Transfer learning can leverage pre-trained models to improve performance when training data is limited, which is common in real-world scenarios.

REFERENCES

- [1] C. P. Veiga, "Analysis of quantitative methods of demand forecasting: Comparative study and financial performance," Ph.D. dissertation, Dept. Manag., Pontifical Catholic Univ. Paraná, Curitiba, Brazil, 2009.
- [4] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining, 2016, pp. 785–794.
- [5] L. Prokhorenkova et al., "CatBoost: Unbiased boosting with categorical features," in Adv. Neural Inf. Process. Syst., 2018, pp. 6638–6648.
- [18] C. C. Holt, "Forecasting seasonals and trends by exponentially weighted moving averages," Int. J. Forecast., vol. 20, no. 1, pp. 5–10, 2004.
- [19] G. E. P. Box et al., Time Series Analysis: Forecasting and Control. Wiley, 2015.
- [21] S. J. Taylor et al., "Forecasting at scale," Am. Stat., vol. 72, no. 1, pp. 37–45, 2018.
- [26] E. Tarallo et al., "Machine learning in food forecasting: A comparative study," Food Control, vol. 98, pp. 123–134, 2019.
- [27] A. Krishna et al., "Comparative analysis of regression models for food demand forecasting," J. Mach. Learn. Res., vol. 21, pp. 1–25, 2020.
- [29] H. Hewamalage et al., "Recurrent neural networks for time series forecasting: A review," Int. J. Forecast., vol. 37, no. 2, pp. 678–701, 2021.
- [30] L. Xu and J. Wang, "Sales forecasting using LSTM on univariate time series," IEEE Trans. Neural Netw. Learn. Syst., vol. 32, no. 6, pp. 2456–2467, 2021.
- [31] S. Kim et al., "Bi-LSTM for multivariate time series forecasting in food demand," Appl. Soft Comput., vol. 105, p. 107245, 2021.
- [32] Y. Li et al., "ETSformer: Exponential smoothing transformers for time series forecasting," arXiv preprint arXiv:2202.01381, 2022.
- [33] Z. Zhang et al., "Hybrid deep learning models for time series forecasting: A survey," IEEE Access, vol. 10, pp. 12345–12367, 2022.
- [34] S. Bai et al., "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," arXiv preprint arXiv:1803.01271, 2018.
- [35] A. S. Ahmad et al., "A review of time series forecasting methods for supply chain management," Int. J. Prod. Econ., vol. 240, p. 108245, 2021.
- [36] J. Brownlee, "Feature engineering for time series forecasting: A practical guide," Mach. Learn. Mastery, 2020. [Online]. Available: <https://machinelearningmastery.com/feature-engineering-time-series-forecasting/>
- [37] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.
- [38] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735–1780, 1997.
- [39] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

- [40] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in Proc. 12th USENIX Symp. Oper. Syst. Des. Implementation (OSDI), 2016, pp. 265–283.
- [41] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, "DeepAR: Probabilistic forecasting with autoregressive recurrent networks," Int. J. Forecast., vol. 36, no. 3, pp. 1181–1191, 2020.
- [42] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "The M4 competition: Results, findings, and conclusions," Int. J. Forecast., vol. 36, no. 1, pp. 802–808, 2020.
- [43] A. Vaswani et al., "Attention is all you need," in Adv. Neural Inf. Process. Syst., 2017, pp. 5998–6008.
- [44] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT, 2019, pp. 4171–4186.

