International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 11 | Nov - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

Food Demand Forecasting for Waste Reduction in Restaurants

Chinmay Kishor Naik[1], Darshan M S[2], Girish C K[3], Gourishankar Chapate[4], Dr. Prabhavathi K[5]

[1]Chinmay Kishor Naik, Student, Department of Computer Science Engineering,

RV Institute of Technology and Management (RVITM), Bengaluru,

[2]Darshan M S, Student, Department of Computer Science Engineering,

RV Institute of Technology and Management (RVITM), Bengaluru,

[3]Girish C K, Student, Department of Computer Science Engineering,

RV Institute of Technology and Management (RVITM), Bengaluru,

[4]Gourishanakr Chapate, Student, Department of Computer Science Engineering,

RV Institute of Technology and Management (RVITM), Bengaluru,

[5]Dr. Prabhavathi K, Assistant Professor, Department of Computer Science Engineering,

RV Institute of Technology and Management (RVITM), Bengaluru, India

Abstract :

Food wastage remains a persistent challenge in the restaurant industry, affecting both economic performance and the environment. This paper proposes the development of a machine learning-based system that accurately forecasts food demand, helping restaurants align production with actual consumption patterns. By analyzing historical sales, seasonal fluctuations, and external influences, the model aims to support more informed inventory planning. The anticipated outcome is a reduction in food waste, enhanced sustainability, and improved operational efficiency.

The main objective of this investigation is to minimize food waste and improve operational efficiency through accurate demand prediction. The model utilizes advanced algorithms such as Random Forest and Long Short-Term Memory (LSTM) networks to capture non-linear consumption patterns and dynamic changes in customer behavior. The outcomes of this research indicate a significant reduction in overproduction and enhanced sustainability within restaurant management systems. From the study, it is concluded that the application of machine learning in food demand forecasting contributes to better resource utilization, cost savings, and alignment with global sustainable development goals related to waste reduction.

Keywords : Food Waste Reduction, Demand Forecasting, Machine Learning Models, Optimization Methods

I. INTRODUCTION

In recent years, the issue of food waste has gained global attention due to its significant economic, environmental, and social consequences. According to the Food and Agriculture Organization (FAO), nearly one-third of all food produced globally is wasted, with a considerable portion of this waste originating from the hospitality and food service sectors, particularly restaurants. This wastage not only represents a severe misuse of natural resources such as water, energy, and labor but also contributes to greenhouse gas emissions, exacerbating climate change. In densely populated and rapidly urbanizing countries like India, the restaurant industry is expanding at an accelerated pace, which intensifies the need for sustainable operational practices, including food waste reduction.

Restaurants often struggle with accurately predicting customer demand on a daily or weekly basis. This uncertainty leads many establishments to prepare excess food to avoid stockouts and service failures. However, this safety margin frequently results in surplus food that cannot be reused or preserved, eventually ending up as waste. Traditional methods of forecasting, which rely on historical averages or manager experience, often fall short in capturing the dynamic and complex factors that influence food consumption, such as holidays, weather conditions, nearby events, or even local traffic patterns. In this context, the integration of data-driven technologies into restaurant operations has become increasingly important.

The emergence of artificial intelligence (AI) and machine learning (ML) technologies offers a promising solution to this challenge. These advanced techniques are capable of learning patterns from large volumes of historical data and making accurate predictions, even in the presence of non-linear relationships and external influences. Over the past decade, researchers and practitioners have explored a variety of ML models including Support Vector Machines (SVM), Decision Trees, Random Forests, Gradient Boosting algorithms, and deep learning architectures such as Long Short-Term Memory (LSTM) networks for the purpose of food demand forecasting. These models have demonstrated superior predictive accuracy compared to traditional statistical methods, particularly when external variables are incorporated into the training datasets.

This literature review aims to examine and synthesize existing research on demand forecasting techniques applied within the food service industry, with a particular focus on restaurant settings. The goal is to understand how different forecasting approaches perform under various conditions and how they contribute to reducing food waste. The review also explores how these technologies are being implemented in real-world scenarios, the types of data used, and the challenges faced during deployment. By critically analyzing and comparing different models and frameworks, this review identifies the current gaps in research and highlights opportunities for further innovation.

Furthermore, as sustainability becomes a critical focus area for businesses and policymakers, the adoption of intelligent forecasting systems is no longer just a tool for cost-saving but a strategic necessity. Improved demand prediction not only supports efficient inventory and kitchen management but also enhances customer satisfaction by reducing stockouts and ensuring menu availability. Importantly, it aligns with global sustainability goals, such as the United Nations Sustainable Development Goal 12.3, which aims to halve per capita global food waste by 2030.

In summary, this review seeks to provide a comprehensive overview of the state-of-the-art in food demand forecasting for restaurants, exploring the intersection of machine learning, predictive analytics, and sustainable food management. By shedding light on proven methodologies, practical case studies, and ongoing research challenges, the paper contributes to a growing body of knowledge aimed at promoting efficiency and sustainability in the food service industry.

II. LITERATURE REVIEW

A. Traditional Forecasting Techniques

Before the rise of machine learning and artificial intelligence, demand forecasting in the food industry predominantly relied on traditional statistical models. Techniques like Moving Averages, ARIMA (Autoregressive Integrated Moving Average), and Exponential Smoothing are popular because they are straightforward to understand, easy to apply, and provide clear, interpretable results.

The Moving Average technique involves computing the average of demand over a fixed time window to smooth out short-term fluctuations and reveal underlying trends. This approach is particularly useful for identifying consistent demand patterns but becomes less effective when there are abrupt changes in consumption behavior. Since it treats all past values equally within the chosen window, it may overlook the impact of more recent events or external factors.

Exponential Smoothing techniques address some of the limitations of Moving Averages by assigning exponentially decreasing weights to older data points. Variants like Simple Exponential Smoothing, Holt’s Linear Trend Method, and Holt-Winters Seasonal Method are commonly applied in forecasting scenarios with trend and seasonal components. While these models offer improved adaptability over time, their performance declines when facing highly volatile or non-linear datasets often seen in the restaurant environment, especially during holidays or local events.

The ARIMA model, on the other hand, combines autoregressive and moving average components while accounting for data stationarity through differencing. It has been widely used in time-series forecasting due to its mathematical complexity. However, ARIMA models assume a linear relationship in the data and require rigorous parameter tuning (p, d, q) and pre-processing steps, such as removing seasonality and ensuring stationarity. These constraints can limit their applicability in real-time restaurant scenarios where demand is influenced by external, non-linear variables.

Although these traditional models have laid the groundwork for forecasting practices, they fall short when applied to modern, data-rich environments like restaurants. Their inherent assumption of linearity, inability to incorporate multiple influencing factors (such as weather, promotions, and holidays), and limited scalability make them less suited for today's complex and dynamic food demand patterns. As a result, researchers and practitioners have increasingly turned toward more advanced machine learning models that can process high-dimensional data and uncover intricate patterns in consumer behavior..

B. Machine Learning Approaches

As traditional forecasting techniques began to show limitations in flexibility and predictive power, particularly in dynamic and data-intensive environments like the food service sector, machine learning (ML) emerged as a promising alternative. Unlike statistical models that often assume linear relationships and require rigid data preprocessing, ML models are capable of learning complex, non-linear patterns directly from raw or minimally transformed data. This adaptability makes them especially useful in scenarios where demand is influenced by a wide range of variables, such as time, menu composition, weather, holidays, and consumer behavior.

One of the most commonly used machine learning models in food demand forecasting is the Random Forest algorithm. It is an ensemble-based method that constructs multiple decision trees and aggregates their predictions to arrive at a final output. Random Forests are robust to overfitting and perform well with structured datasets. In various case studies involving restaurant and catering environments, this method has demonstrated high accuracy in predicting next-day or weekly demand when trained on historical sales data, reservation patterns, and external variables like calendar events.

Another popular technique is the Gradient Boosting Machine (GBM) and its advanced variants such as XGBoost, LightGBM, and CatBoost. These models build trees sequentially, each one correcting the errors of its predecessors. Their ability to handle missing values, categorical features, and outliers makes them highly effective in food industry applications. Studies have shown that these algorithms outperform simpler models in forecasting food demand and minimizing wastage by optimizing procurement and preparation schedules.

Support Vector Machines (SVM) have also been applied in the context of food supply prediction, particularly when the data set is not excessively large. SVMs work well with non-linear data by mapping input features into high-dimensional spaces, allowing them to separate complex relationships. Although they are less scalable to massive datasets, their accuracy can be impressive in small to medium-sized restaurant operations.

More recently, deep learning models such as Artificial Neural Networks (ANNs) and Long Short-Term Memory (LSTM) networks have gained attention. While LSTM models are particularly well-suited for time-series forecasting due to their ability to retain memory of past sequences, even simple feed-forward neural networks have shown promising results in capturing non-linear relationships in food consumption trends. These models can be further enhanced by integrating exogenous inputs like temperature, public holidays, and local events, which significantly affect dining behavior.

The key advantage of machine learning techniques lies in their capacity to generalize from historical patterns and adapt to new data in real time. Moreover, they can incorporate a wide range of structured and unstructured data, enabling restaurants to generate forecasts that are both timely and accurate. However, despite their high performance, challenges such as interpretability, model training time, and the need for clean, labeled datasets still persist. Nonetheless, the growing body of research and successful case studies suggest that machine learning offers a transformative path toward intelligent and sustainable food waste reduction in the hospitality sector.

C. Deep Learning Models

Deep learning, a subfield of machine learning, has recently become a leading approach for time-series forecasting, particularly in domains like food demand prediction where data patterns can be highly non-linear and influenced by multiple contextual factors. Unlike traditional or even basic machine learning models, deep learning methods are designed to automatically extract hierarchical features from raw data, which makes them highly effective for handling complex and unstructured datasets.

Among the most prominent deep learning techniques used in food demand forecasting is the Long Short-Term Memory (LSTM) network. LSTM is a specialized type of recurrent neural network (RNN) capable of learning temporal dependencies over long sequences, making it ideal for time-series problems. In restaurant settings, where demand can fluctuate based on time of day, day of the week, and even seasonality, LSTM models have been shown to outperform traditional forecasting models. Their ability to retain information from previous time steps allows them to model patterns like weekend peaks, holiday dips, or seasonal changes in consumption.

In addition to LSTM, other architectures such as Convolutional Neural Networks (CNNs) and Transformer models have also been explored for demand forecasting. CNNs, while originally designed for image processing, have proven effective in identifying localized patterns in temporal data when applied through 1D convolutions. Transformers, on the other hand, use attention mechanisms to weigh the

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 11 | Nov - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

importance of different time steps, enabling more flexible and scalable forecasting—especially useful in multi-variate and multi-step prediction scenarios.

Furthermore, hybrid models combining LSTM layers with dense (fully connected) layers, or CNN-LSTM architectures, have also been developed to improve forecasting accuracy. These architectures allow the model to learn both spatial (cross-feature) and temporal (sequence) relationships, leading to more robust performance. They are particularly effective in restaurant data, where relationships between dishes, pricing, weather, and holidays interact in complex ways.

One advantage of deep learning is its ability to work directly with raw or partially preprocessed data. For example, rather than manually engineering features like "holiday flag" or "weekday index", a well-trained LSTM or Transformer model can learn such patterns inherently. However, these models often require large volumes of high-quality data, and the training process can be computationally intensive. Additionally, they are less interpretable than traditional models, which can make them harder to justify to non-technical stakeholders such as restaurant managers.

Even with these hurdles, the food service industry is gradually embracing deep learning technologies. Studies across catering services, bakeries, and restaurants show significant improvements in demand forecasting accuracy—sometimes reducing food wastage by up to 50%. These outcomes demonstrate the real-world value of deep learning in building predictive systems that not only enhance operational efficiency but also contribute to environmental sustainability.

D. Integrating External Factors into Forecasting Models

While historical sales data provides a foundational basis for demand forecasting, real-world food consumption is influenced by a wide array of external variables. Relying solely on internal transaction data can often lead to underfitting or inaccurate predictions, especially in environments like restaurants where customer behavior is highly context-sensitive. To address this, recent advancements in food demand forecasting have emphasized the integration of external factors—such as holidays, weather conditions, local events, and social trends—into predictive models.

One of the most significant external variables in the food service sector is the calendar, particularly public holidays and weekends. These days often see spikes or drops in customer turnout depending on the type and location of the restaurant. For instance, fast food chains near offices may experience lower footfall during holidays, whereas family-style restaurants might see the opposite. Incorporating a “holiday flag” or “day-of-week” feature into models has consistently improved forecasting accuracy by allowing the model to adjust for routine behavioral patterns.

Weather conditions are another powerful influence on dining habits. Factors such as temperature, humidity, rainfall, and even air quality can affect consumer decisions—such as opting for indoor dining over takeout, or craving specific types of meals (e.g., hot soup in cold weather). Advanced models that include weather APIs or local meteorological data as inputs have demonstrated better alignment with real-world demand fluctuations. For example, predictive models in some studies showed significant accuracy improvements when weather data was included alongside transactional features.

Moreover, local events, such as concerts, sports matches, festivals, or nearby office conferences, can drive sudden and substantial changes in customer flow. While more difficult to quantify than calendar or weather data, these factors can be tracked using event calendars or social media trend analysis. Some modern models even incorporate social listening tools or sentiment analysis to anticipate demand surges linked to viral events or public sentiment.

In addition, promotions and pricing strategies—although internal decisions—function as external signals from the customer’s point of view. The introduction of discounts, combo offers, or limited-time menus can dramatically shift demand patterns. Including these variables in the model ensures that the forecasting system remains responsive to marketing activities and pricing dynamics.

Several machine learning and deep learning models have been extended to support multivariate forecasting by integrating these contextual factors. Tree-based algorithms like XGBoost and Random Forest are particularly effective in handling categorical features such as holiday types or weather categories. Similarly, LSTM and Transformer models can process time-series sequences enriched with exogenous variables, resulting in enhanced long-term and short-term forecasting precision.

Taking external factors into account helps make forecasts more accurate and the insights more practical for decision-making. Restaurant managers can plan inventory more effectively, reduce overproduction, and optimize staffing by anticipating busy or slow periods. Ultimately, this integration supports a more responsive and sustainable food management system, aligning with the dual goals of profitability and waste reduction.

E. Real-World Applications and Case Studies

The practical implementation of food demand forecasting models in real-world restaurant and catering environments has yielded valuable insights into their effectiveness. Numerous case studies from various countries and sectors reveal that predictive analytics particularly when powered by machine learning can lead to substantial reductions in food waste, improved inventory control, and enhanced customer satisfaction. These applications offer tangible proof that demand forecasting is not merely a theoretical exercise but a vital tool for sustainable operations.

One widely cited example is Foodforecast, a cloud-based platform developed to help bakeries minimize unsold inventory through demand prediction. The system uses machine learning algorithms trained on historical sales data, enriched with external variables such as day of the week, weather conditions, and nearby public holidays. The integration of a Life Cycle Assessment (LCA) approach in evaluating Foodforecast revealed both the direct environmental impact of running machine learning systems and their indirect benefits through waste reduction. While the energy and computational cost of these systems was non-trivial, the environmental gains from reduced wastage often outweighed the resource input.

In another notable case, a food manufacturing company - Kaleh Company applied artificial neural networks to forecast the demand for dairy products like milk and yogurt. By using multi-layer perceptron models trained in MATLAB, the company achieved significantly improved accuracy in predicting product-level demand. This allowed them to fine-tune both their production and purchasing decisions, leading to enhanced supply chain coordination and a decrease in surplus goods. The study emphasized the model's ability to reduce economic uncertainty and improve sustainability in a highly perishable product segment.

Supermarkets and large retailers like Walmart and Tesco have also successfully deployed machine learning systems to predict daily sales and automate inventory replenishment. In Tesco's case, a tailored ML system was used to forecast demand at the SKU level, enabling them to adjust restocking schedules dynamically. The implementation reportedly led to a substantial decrease in perishable waste, improved shelf availability, and better alignment between supply and demand.

Educational institutions and corporate canteens have become ideal environments for testing short-term demand forecasting systems. For instance, one case study involving two university dining halls and a company cafeteria implemented models such as Random Forest, LightGBM, LSTM, and Transformer neural networks. These systems incorporated real-time data on menu choices, reservations, and weather to predict the number of meals required for the following day. Results showed that the LSTM and Random Forest models reduced food wastage by up to 52%, while unmet demand was also lowered. This dual improvement in efficiency and service quality demonstrated the operational potential of machine learning in institutional food services.

In the context of fast-moving e-grocery platforms, deep learning models such as LSTM-based time-series forecasters have been deployed for product-level forecasting of fresh items. These platforms typically manage vast amounts of data spanning multiple product categories and geographic locations. Studies show that multivariate LSTM models, which incorporate variables like weekday trends and pricing data, outperformed both linear regression and random forest models, especially for food items.

Even smaller-scale operations have seen benefits. For example, startups in the food delivery and meal-kit sectors have adopted forecasting tools to optimize ingredient procurement and reduce unsold inventory. These systems rely on ML models trained on order history, seasonal trends, and promotional responses. Their implementation has led to greater cost control and more accurate meal portioning.

Across these case studies, one common insight is the importance of data availability and quality. The effectiveness of any forecasting system largely relies on how accurate, complete, and detailed the input data is. Moreover, many organizations highlight the need for user-friendly dashboards and visualizations to translate complex model outputs into actionable insights for kitchen staff and managers.

While each deployment scenario varies in terms of complexity, size, and objectives, the collective evidence suggests that real-world applications of demand forecasting are yielding measurable benefits. These include cost savings, reduced food wastage, improved inventory turnover, and enhanced customer service—all of which support the broader goal of building more sustainable and efficient food systems.

III. CHALLENGES IDENTIFIED IN LITERATURE

While machine learning and data-driven forecasting systems have shown great promise in the restaurant and food service industry, their implementation is not without challenges. A review of recent studies reveals that despite technical advancements and increasing accuracy in predictive models, several practical, technical, and ethical hurdles continue to limit their broader adoption and real-world effectiveness.

One of the most prominent challenges is data quality and availability. Predictive models depend heavily on accurate, consistent, and comprehensive historical data to learn meaningful patterns. In many restaurant settings, particularly small to mid-sized establishments, sales records may be incomplete, manually logged, or lack the level of granularity required for training robust models. Missing data points, inconsistent labeling of food items, and untracked variables such as walk-in traffic or special promotions can significantly reduce model reliability. As a result, much of the potential of machine learning remains untapped due to poor data infrastructure.

Another common issue is the integration of forecasting systems with existing operational workflows. Many restaurants and canteens rely on traditional processes and manual planning. Transitioning to a digital model-based system often requires changes in employee roles, retraining, and adjustments to kitchen routines. Without proper change management and user training, even the most accurate model may be underutilized or ignored. Studies note that unless the forecasting output is delivered in a clear, actionable, and easy-to-use format, it is unlikely to influence real-time decision-making effectively.

The cost of implementation also presents a significant barrier, especially for smaller restaurants or local food providers. Developing, training, and maintaining a predictive model involves both time and financial investment. This includes software tools, cloud storage, computational resources, and technical expertise. While some open-source frameworks reduce entry barriers, professional deployment often requires partnerships with data scientists or AI service providers. This financial commitment may be perceived as risky without clear short-term returns.

Moreover, model interpretability remains a technical concern. Advanced models like deep neural networks or ensemble algorithms, while highly accurate, often function as "black boxes," offering little transparency into how decisions are made. This lack of interpretability makes it difficult for restaurant managers to trust the model's recommendations, especially when outcomes are unexpected. In response, recent research has focused on explainable AI (XAI) tools, but these are still emerging and may not yet be practical for all users.

External factor unpredictability also introduces limitations. While models may be trained with weather data, holidays, and known events, spontaneous disruptions—such as last-minute cancellations, sudden shifts in consumer preferences, or pandemic-related closures—can render predictions obsolete. Adapting models to such real-time changes is still an area of active research, with techniques like reinforcement learning and adaptive forecasting being explored.

Additionally, ethical considerations surrounding data privacy and algorithmic fairness must not be overlooked. As forecasting systems collect and process consumer behavior data, there is a risk of infringing on customer privacy if adequate safeguards are not in place. Transparency in data usage and adherence to privacy laws such as the GDPR or India's DPDP Act are essential, especially when using loyalty programs, location tracking, or personalized promotions.

Lastly, many studies point to a lack of standardized benchmarking and real-world validation. While a model might perform well in a controlled research setting, its real-life performance may vary depending on operational complexity, staff cooperation, and customer behavior. There is a pressing need for more longitudinal studies that monitor forecasting tools over extended periods in diverse restaurant environments.

In summary, while food demand forecasting using machine learning presents exciting possibilities for reducing waste and improving operational efficiency, several challenges must be addressed for large-scale adoption. These include improving data quality, ensuring system usability, reducing deployment costs, enhancing interpretability, and managing ethical risks. Overcoming these barriers will require collaboration between technologists, restaurant managers, policymakers, and academic researchers.

IV. COMPARISON OF MODELS

When evaluating the effectiveness of various forecasting models used in the food industry, especially in restaurant settings, it becomes essential to compare them based on multiple performance dimensions. These include prediction accuracy, ability to handle external variables, scalability, ease of deployment, interpretability, and overall cost-efficiency. A clear comparison helps stakeholders select the most suitable model depending on the size of the business, data availability, and operational goals.

Statistical models such as Moving Averages, Exponential Smoothing, and ARIMA have historically been favored for their simplicity and interpretability. These models perform reasonably well in stable, low-variance environments with strong seasonal trends. However, their predictive capability diminishes in dynamic contexts where demand is affected by multiple non-linear factors. Moreover, they struggle to incorporate external variables like holidays or weather and require careful manual tuning of parameters.

In contrast, machine learning models particularly tree-based algorithms like Random Forest, Gradient Boosting, XGBoost, and LightGBM offer improved accuracy and flexibility. These models can capture complex patterns in data and are more resilient to outliers and missing values. In comparative studies, ensemble models often outperform traditional methods, especially when enriched with calendar, promotional, or weather-related data. Their main strength lies in their robustness and ability to generalize across different demand profiles. However, they can be resource-intensive during training and may require significant feature engineering.

Support Vector Machines (SVMs) are effective in scenarios with medium-sized datasets and perform well in separating non-linear patterns. However, they are less commonly used in large-scale restaurant forecasting due to scalability limitations and longer training times.

Moving to deep learning models, LSTM neural networks consistently rank among the top-performing approaches for time-series forecasting. Their ability to retain memory over long sequences makes them ideal for modeling food demand patterns that exhibit periodic or long-term dependencies. Studies have shown that LSTMs outperform both traditional statistical methods and basic ML models in most real-world restaurant forecasting scenarios, especially when dealing with multi-step or multi-variable predictions. On the downside, LSTM models require significant computational power and high-quality data to train effectively. Their interpretability is also limited, making it harder for non-technical users to understand their decision-making process.

More recently, Transformer-based models have been introduced for demand forecasting tasks. Their use of attention mechanisms allows them to assign different levels of importance to time steps, which is especially useful when working with multivariate data. These models have shown promising results in preliminary experiments, but they remain less commonly used in practice due to their complexity and data requirements.

When models are compared based on evaluation metrics such as Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and R-squared (R^2), deep learning models especially LSTM and XGBoost often deliver the lowest error rates. However, these models may not be the most practical choice for all restaurants, particularly small businesses that lack technical resources.

From a deployment perspective, simpler models like ARIMA or even decision trees are easier to implement and interpret, making them suitable for operations with limited technical capacity. More advanced models may require backend integration with POS systems, API interfaces, and cloud computing support for real-time forecasting.

In terms of cost-effectiveness, traditional models have the lowest entry barrier but deliver limited insights. Machine learning and deep learning models offer significantly more value when used with the right infrastructure and data quality. However, the trade-offs between accuracy and operational complexity must be carefully considered.

To summarize, each model type has its strengths and limitations. Traditional models are useful for basic trend forecasting, machine learning models offer a good balance between performance and scalability, and deep learning models deliver the highest predictive accuracy at the cost of greater computational and implementation complexity. The choice of model should align with the restaurant's forecasting needs, data readiness, and technological capability.

V. METHODOLOGY

The development of this project follows a well-defined methodology designed to combine real-world restaurant operations with modern data science techniques. The goal is to create a forecasting system that can predict food demand with high accuracy and support sustainable kitchen practices by minimizing unnecessary preparation and waste.

The first step in this process is problem identification. Restaurants frequently face the challenge of uncertain customer turnout, leading them to either overproduce food which results in wastage or underprepare, which affects customer satisfaction. To address this, our methodology revolves around building a machine learning–based model that uses both historical data and contextual factors to predict future demand.

We begin with data collection from multiple sources. The most critical data comes from the restaurant's point-of-sale (POS) system, which records item-wise daily sales. This is combined with external data such as weather forecasts, holiday calendars, and any local events that might influence customer behavior. These external influences are crucial for creating a robust and realistic forecasting model.

Based on insights from the literature review, Long Short-Term Memory (LSTM) neural networks were selected as the primary modeling approach. LSTM is a powerful deep learning model specifically designed for time-series data, making it well-suited for capturing sequential trends in food consumption. Its ability to "remember" past patterns over time makes it more accurate than traditional models for forecasting restaurant demand, especially when factors like seasonality and customer habits are involved.

The model is trained using a split of the prepared dataset (typically 70% for training and 30% for testing) and evaluated using industry-standard metrics such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-squared (R^2). These metrics help determine how closely the predictions match actual demand, and whether the model is suitable for practical deployment.

Following training and validation, the model is deployed through a backend API using tools like Flask or FastAPI. This allows seamless integration with the database and frontend. The relational database (MySQL or PostgreSQL) is used to store both historical data and real-time forecasts. The frontend dashboard is developed using Streamlit, which provides an intuitive and interactive interface for restaurant managers to visualize demand predictions, monitor performance metrics, and make informed decisions efficiently.

The final stage of the methodology focuses on the outcomes. By using predictive insights from the model, restaurants can plan meals more accurately, purchase ingredients more efficiently, and significantly reduce food waste. Our target is to achieve a 20–30% reduction in food waste, as well as improved operational efficiency, cost savings, and contribution to sustainability efforts.

This end-to-end methodology ensures that the system is not only technically sound but also practical and scalable for real world restaurant operations.

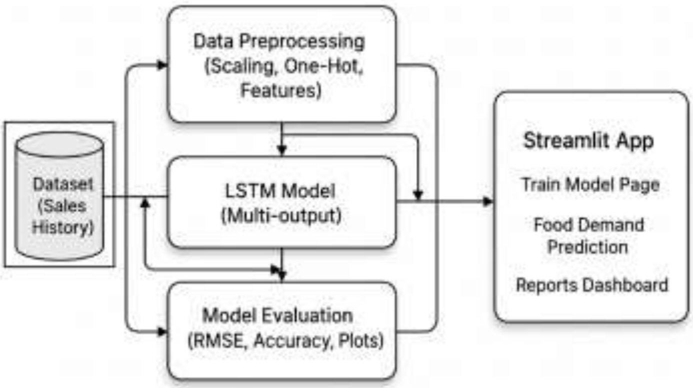

Architecture of the Proposed LSTM-Based Food Demand Prediction System

VI. CONCLUSION

Rising awareness about food waste in the restaurant sector has sparked increased interest in advanced forecasting methods. Accurate prediction of food demand is not only a pathway to reducing waste but also a strategic tool for optimizing inventory, improving operational efficiency, and supporting broader sustainability goals. Through this literature review, it is evident that advancements in data analytics particularly in machine learning and deep learning have introduced transformative possibilities for tackling this long-standing challenge.

Although traditional forecasting methods are fundamental and easy to understand, they often struggle to keep up with the complex demands of today's restaurant operations. Their limitations become evident in situations where demand is influenced by multiple dynamic factors, such as weather, public holidays, and promotional campaigns. Machine learning models like Random Forest, XGBoost, and SVM offer greater predictive accuracy and flexibility by capturing non-linear relationships and accommodating external variables. Meanwhile, deep learning techniques, especially LSTM and Transformer networks, push the boundaries further by enabling multi-step and multivariate forecasting with superior performance in time-series contexts.

A key takeaway from this review is the growing importance of integrating contextual variables—calendar effects, weather, events, and pricing data—into forecasting systems. These enrich the predictive power of models and bring forecasts closer to real-world behavior. At the same time, real-world applications demonstrate that technical accuracy alone is not sufficient; models must also be interpretable, cost-effective, and seamlessly integrated into restaurant operations to drive actual improvements.

Despite the clear advantages, significant challenges remain. Issues such as poor data quality, high computational costs, model complexity, and limited user-friendliness can hinder adoption. For small and medium-sized restaurants, the resources required to implement advanced AI systems may still be prohibitive. Furthermore, ethical concerns related to data privacy and transparency must be addressed as these technologies become more widespread.

Looking ahead, there is a strong need for interdisciplinary collaboration between technologists, restaurant managers, policy-makers, and researchers to make intelligent demand forecasting accessible and actionable. Future studies should focus not only on enhancing model accuracy but also on improving interpretability, reducing computational overhead, and validating solutions in diverse, real-world restaurant settings.

In conclusion, food demand forecasting is evolving into a powerful lever for change in the food service sector. When implemented effectively, these predictive systems can play a pivotal role in achieving operational efficiency, customer satisfaction, and environmental sustainability making them an essential component of the future restaurant ecosystem.

REFERENCES

[1] J. C. Silva, M. C. Figueiredo, and A. C. Braga, "Demand Forecasting: A Case Study in the Food Industry," Supported by FCT - Project UID/CEC/00319/2019.

[2] N. Hbner, J. Caspers, V. C. Coroam, and M. Finkbeiner, “Machine-learning-based demand forecasting against food waste: environmental implications and rewards of a bakery case study”, Journal of Industrial Ecology, 28, 5 (2024), pp. 1117-1131, DOI: 10.111/jiec.13528.

[3] A. Q. Olufemi-Phillips, O. C. Ofodile, A. S. Toromade, A. N. Igwe, and N. L. Eyo-

Udo, "Utilizing Predictive Analytics to Manage Food Supply and Demand in Adaptive Supply Chains," International Journal of Engineering Research and Development, vol. 20, no. 11, pp. 1041-1059, Nov. 2024.

[4] M. Pakbin, Y. R. Ghofrani, and K. Shahroodi, "Forecasting Demand and Optimizing the Ordering of Goods in the Supply Chain Using Artificial Intelligence in Kaleh Company," International Journal of Nonlinear Analysis and Applications, In Press, 2023, doi: 10.22075/ijnaa.2023.31695.4797.

[5] N. Nassibi, H. Fasihuddin, and L. Hsairi, "Demand Forecasting Models for Food Industry by Utilizing Machine Learning Approaches," International Journal of Advanced Computer Science and Applications (IJACSA), vol. 14, no. 3, 2023.

[6] S. Jayapal, "Food Demand Prediction using Statistical and Machine Learning Models," National College of Ireland, 2024.

[7] E. Ozcil, "Reducing Food Waste Using Machine Learning Models: Forecasting and Optimization Approaches," Journal of Data Science & Artificial Intelligence, vol. 2, no. 1, Article 136, 2024

[8] Gołąbek, M., Senge, R., & Neumann, R. (2020). Demand Forecasting using Long Short-Term Memory Neural Networks. Karlsruhe University of Applied Sciences and inovex GmbH. arXiv preprint arXiv:2008.08522

[9] Rodrigues, M., Miguéis, V., Freitas, S., & Machado, T. (2024). Machine learning models for short-term demand forecasting in food catering services: A solution to reduce food waste. Journal of Cleaner Production, 435, 140265.

[10] Tanizaki, T., Hoshino, T., Shimmura, T., & Takenaka, T. (2019). Demand forecasting in restaurants using machine learning and statistical analysis. Procedia CIRP, 79, 679–683.

[11] Ghosh, S., Ghosh, R., & Chakraborty, S. (2019)."Predicting Restaurant Sales Using Machine Learning Techniques."

International Journal of Computer Applications, 178(41), 1-5.

[DOI: 10.5120/ijca2019918805]

[12] Silva, A. L., Ferreira, D., & Cunha, P. (2021).

"The Influence of Weather Conditions on Food Demand in Restaurants: A Machine Learning Approach."Journal of Retailing and Consumer Services, 61, 102544.

[DOI: 10.1016/j.jretconser.2021.102544]

[13] Choi, J., Lee, S., & Ryu, K. (2020)."Deep Learning-Based Time Series Forecasting for Food Demand Prediction in Restaurants."IEEE Access, 8, 105944-105954. [DOI: 10.1109/ACCESS.2020.2999865]

[14] Food and Agriculture Organization of the United Nations (FAO). (2019).

"The State of Food and Agriculture 2019: Moving Forward on Food Loss and Waste Reduction."Rome: FAO.[ISBN: 978-92-5-131789-1]

[15] Agrawal, R., & Gupta, A. (2017)."Application of Predictive Analytics in Reducing Food Wastage: A Case Study of Restaurant Chains."International Journal of Advanced Computer Science and Applications (IJACSA), 8(10), 230-234.[DOI: 10.14569/IJACSA.2017.081030]

[16] E. J. Flies, B. W. Brook, L. Blomqvist, and J. C. Buettel, "Forecasting Future Global Food Demand: A Systematic Review and Meta-Analysis of Model Complexity," Environment International, vol. 120, pp. 93–103, 2018, doi: 10.1016/j.envint.2018.07.019.

[17] E. G. Sigala, P. Gerwin, C. Chroni, K. Abeliotis, C. Strotmann, and K. Lasaridi, "Reducing Food Waste in the HORECA Sector Using AI-Based Waste-Tracking Devices," Waste Management, vol. 198, pp. 77–86, 2025, doi: 10.1016/j.wasman.2025.02.044.

[18] N. Kozodoi, E. Zinovyeva, S. Valentin, J. Pereira, and R. Agundez, "Probabilistic Demand Forecasting with Graph Neural Networks," Preprint accepted at ECML PKDD 2023 ML4ITS Workshop, arXiv:2401.13096v1 [cs.LG], 2024.

[19] M. Cardenas, B. Schivinski, and L. Brennan, "Circular Practices in the Hospitality Sector Regarding Food Waste," Journal of Cleaner Production, vol. 472, 2024, Article ID 143452, doi: 10.1016/j.jclepro.2024.143452.