Applied Computing and Intelligence

https://www.aimspress.com/journal/aci

Applied Computing and Intelligence
5(2): 168–190.
DOI: 10.3934/aci.2025011
Received: 04 June 2025
Revised: 01 August 2025
Accepted: 05 August 2025
Published: 21 August 2025

Research article

Forecasting daily customer flow in restaurants: a multifactor machine learning approach

Himat Shah[1,*], Niko Myller[2] and Carolina Islas Sedano[2]

[1] Robotics and Artificial Intelligence Group Centria University of Applied Sciences, Finland

[2] Department of Computing, Turku University, Finland

* **Correspondence:** Email: himat.shah@centria.fi; Tel: +358465840099.

Academic Editor: Pasi Fränti

Abstract: This paper presents a case study on predicting daily customer flow in a university's self-service restaurant. We conduct a systematic comparison of multiple machine learning techniques and diverse feature sets to identify the best-performing model within our experimental scope, aiming to improve forecasting accuracy in this dynamic environment. We analyze real-time data collected via RFID sensors from spring 2019 to 2024. To ensure high data quality, we apply a robust preprocessing pipeline, followed by careful feature engineering to select 10 distinct features, labeled M1 to M10. These features include temporal attributes such as day, month, year, and season, as well as external factors like local weather conditions, public holidays, and menu choices. We conduct a systematic comparison across all feature sets and identify M10 as the optimal combination. A key finding highlights the importance of handling missing data—particularly during the COVID-19 period—as one of the most critical steps in the preprocessing stage. To evaluate the predictive power of our selected features, we tested various machine learning models, including linear regression, random forest, extreme gradient boosting (XGBoost), and long short-term memory (LSTM). Our findings indicate that XGBoost achieves the lowest mean absolute error (MAE) and mean squared error (MSE) values. The XGBoost model outperforms other models across all the feature sets, M1 to M10. XGBoost is particularly effective because it uses past data and a technique called exponential smoothing to understand what customers will do in the short term. Our analysis identifies the most influential features as the previous day's customer count, exponential smoothing outputs, holidays, day of the week, and weather data. Overall, we recommend XGBoost as the most effective model for predicting daily customer numbers in similar contexts, given its superior accuracy across diverse feature sets.

Keywords: time series analysis; machine learning; forecasting; customers; restaurants

1. Introduction

Restaurants face ongoing challenges in optimizing operations and accurately estimating customer counts and food demand, especially as customers' behavior continues to evolve and sustainability becomes increasingly important in the service industry. After the COVID-19 pandemic, restaurants located in office buildings experienced a reduction of up to 30% in the average number of customers due to the rise in the amount of remote work [8]. Daily fluctuations in customer numbers may vary widely, ranging from 30% to 100% compared with pre-pandemic levels [8]. In this context, providing restaurant operators with automated tools to predict customer demand and behavior is essential for improving operational planning. Accurate forecasts can assist decision-makers in managing food preparation, inventory control, profit projections, and overall strategic planning [5]. Additionally, precise demand forecasting helps reduce food waste by ensuring that the right amount of food is prepared each day.

This case study introduces an approach for predicting daily customer demand in restaurants using daily data and advanced machine learning methods. This approach is essential for optimizing operations, managing inventory, and improving overall efficiency in the restaurant industry [22]. Our research focuses specifically on self-service restaurants (SSRs)—dining establishments where customers serve themselves from a buffet-food line and pay in advance. The data for this study are collected from Flavoria, a research-oriented restaurant located at the campus of the University of Turku, Finland (see Section 3.1 for details). The dataset spans the years 2019 to 2024 and includes a wide range of variables, such as customer visits, demographic information, food preferences, menu items, weather conditions, and public holidays [12, 19].

Our long-term objective is to optimize restaurant operations and improve efficiency, particularly as the industry moves towards more sustainable practices. To achieve this, we develop predictive models using machine learning techniques such as XGBoost, random forest, and long short-term memory (LSTM) networks. These models aim to accurately forecast key operational factors, including daily customer counts, food demand, and potential food waste across various stages of the service. This study also analyzes the relationships among the collected variables to enhance the accuracy of customer behavior forecasting and support data-driven decision-making.

We employ time series forecasting techniques commonly used in fields such as sales forecasting [20], climate analysis [1], stock market analysis [22], and healthcare demand prediction [10]. These techniques enable restaurants to proactively adjust inventory, anticipate food demand fluctuations, and reduce waste [16]. To achieve the best results, we assess various machine learning techniques for time series prediction. XGBoost, known for its strong performance in similar studies [15, 20], effectively handles large datasets by incorporating historical visitor data and factors like weather and demographics. We also explore the use of linear regression, random forest [31], and long short-term memory (LSTM) [22] models to predict restaurant customer counts. In the realm of restaurant sales forecasting, advanced models like LSTM and transformer-based forecasters (TFT) have shown promising results for long-term accuracy [20]. This study highlights the importance of using these predictive techniques to improve efficiency and customer satisfaction. By comparing various methods, we aim to identify the most effective machine learning models for real-time predictions based on our collected data.

Restaurant customer predictions combine internal and external factors, such as previous visitors, holidays, and local events. These significantly influence customer behavior. For instance, sunny weather may attract more customers to outdoor seating, while holidays and events can either increase or decrease the foot traffic. We aim to do the same in our current study and analyze different feature sets as well as methods to find the best possible combination of features for the SSRs settings. This will contribute to the field with new information on the applicability of different machine learning methods as well as different features that might have an impact on customer behavior.

This paper is organized as follows: In Section 2, we review the existing literature on forecasting the number of customers and related factors in restaurants and similar contexts, as well as time series analysis. Section 3 explains our research method, including how we construct features and select models. Section 4 presents our model evaluation and results. We discuss the results in Section 5. Finally, Section 6 summarizes our findings and explores the implications for restaurant operations and future research.

2. Predicting customer flow

In our case study, we begin with a literature review of relevant research published between 2015 and 2023. We sourced articles from major academic databases, including Google Scholar, institute of electrical and electronics engineers (IEEE) Xplore, and the ACM Digital Library. These studies were organized into categories aligned with our research focus, which we present in this section. Several of the reviewed works examine the use of time series forecasting techniques to predict daily customer numbers in the restaurant industry (see Tables 1 and 2).

Researchers widely utilize machine learning methods to forecast customer flow on the basis of historical data. Beyond customer prediction, studies also examine operational aspects like menu optimization, meal preparation times, and table availability. Forecasting intervals vary from hours to weeks, underscoring the need for robust models that can manage high-frequency data, seasonal variations, and missing data. Table 1 categorizes studies by their prediction targets in the restaurant industry and includes relevant statistics.

In recent years, machine learning techniques have revolutionized customer prediction in restaurants. Ma et al. [15] predict daily customer numbers using machine learning and big data. They also look at information like reservations, past visits, weather, local events, and online reviews. They compare K-nearest-neighbor (KNN), random forest, and XGBoost, finding that XGBoost excels in handling large datasets and complex relationships. An et al. [2] use linear regression to forecast daily customer counts, focusing on temporal, weather, and statistical features. Tanizaki et al. [25] use Bayesian linear regression to examine the influence of external variables on daily restaurant visitor predictions. They conclude that historical visit data and human behavior significantly impact forecasting accuracy, while weather variables like temperature, precipitation, and wind speed have surprisingly minimal effects. Building on these insights, our study integrates a range of commonly used features, including weather variables, to enhance customer number predictions. Our study uses commonly available information, such as weather details, to improve customer number predictions. We assess the importance of each feature to improve the reliability of our predictions.

Rarh et al. [18] aimed to optimize the daily operations of a restaurant using metrics like seat allotments, table counts, turnover rates, and average spending, offering insights for optimizing operations. Su et al. [24] and Siek et al. [23] use autoregressive integrated moving average (ARIMA) and computational intelligence for optimizing inventory management and forecasting customer demand. They emphasize the importance of historical data for predicting future customer needs and optimizing inventory levels. Zhao and Jayadi [32] use support vector regression (SVR) to forecast daily visitors and menu demand in Indonesian restaurant chains, demonstrating SVR's superiority over traditional methods for complex demand forecasting.

Chen et al. [4] and Fernandes et al. [6] apply sentiment analysis and regression techniques to explore how online reviews impact customer count forecasting. Insights from these studies help restaurants attract new customers by strategizing around customer sentiments and preferences. Technological advancements, like Kumaresan et al.'s [14] software reliability prediction model, highlight the potential of statistical time series approaches to enhance forecasting accuracy across domains.

Table 1. Summary of prediction targets in restaurant-related studies.

Category	References
Restaurant customer count or sales prediction	Daily predictions: [2, 3, 15, 23–26, 32] Hourly predictions: [26] Weekly predictions: [20]
Food demand and menu item forecasting	Food demand prediction: [9, 32] Menu item sales prediction: [9, 23] Ingredient waste prediction: [16]
Customer behavior and review analysis	Analysis of customer reviews: [4–6, 28] Dine-in and alone dining time prediction: [11, 13] Price sensitivity prediction: [7]
Other related topics	Food delivery time prediction: [16, 33] Place recommendation based on reviews: [29, 30]

Table 2. Summary of reviewed studies.

Study	Techniques	Results
Ahmad et al. [1]	support vector machine (SVM), artificial neural network (ANN), random forest (RF), k–nearest neighbor (KNN)	Strategies for climate-sensitive and non-sensitive conditions
Boomija et al. [3]	autoregressive integrated moving average (ARIMA)	Precise forecasting of visitor trends
Ma et al. [15]	extreme gradient boosting (XGBoost), KNN, RF	XGboost outperforms in handling large datasets and complex relationships
Schmidt [20]	XGBoost	Effective for sales forecasting; considers weather and seasonal trends
Zhao and Jayadi [32]	support vector regression (SVR)	Effective for demand forecasting in an Indonesian restaurant chain

Ahmad et al. [1] discuss challenges in forecasting, such as data quality issues, varying forecasting horizons, and regional differences. They propose machine learning models as solutions to these challenges.

In conclusion, reviewing the existing studies offers valuable insights into forecasting methods and their practical applications for improving restaurant operations. Building on these insights, our study introduces a predictive model designed to adapt to changes in customer behavior and shifting market conditions. This model, along with the proposed data-driven approach, aims to enhance the accuracy of forecasts, even in dynamic environments. Consequently, it can support more efficient restaurant operations.

3. Prediction method

Figure 1 illustrates the overall workflow of our approach in this case study, which includes, data preprocessing, feature selection, model training, testing, and customer prediction. Section 4 provides a detailed evaluation of the model's performance, showing the importance of each step in the process.

The data used in this study were collected from the Flavoria research restaurant, located on the University of Turku campus, as explained in Section 3.1. The machine learning model progresses through several stages—from data collection to the final prediction of daily customer counts. These phases are outlined in Figure 1 and further detailed in Sections 3.2 and 3.4.

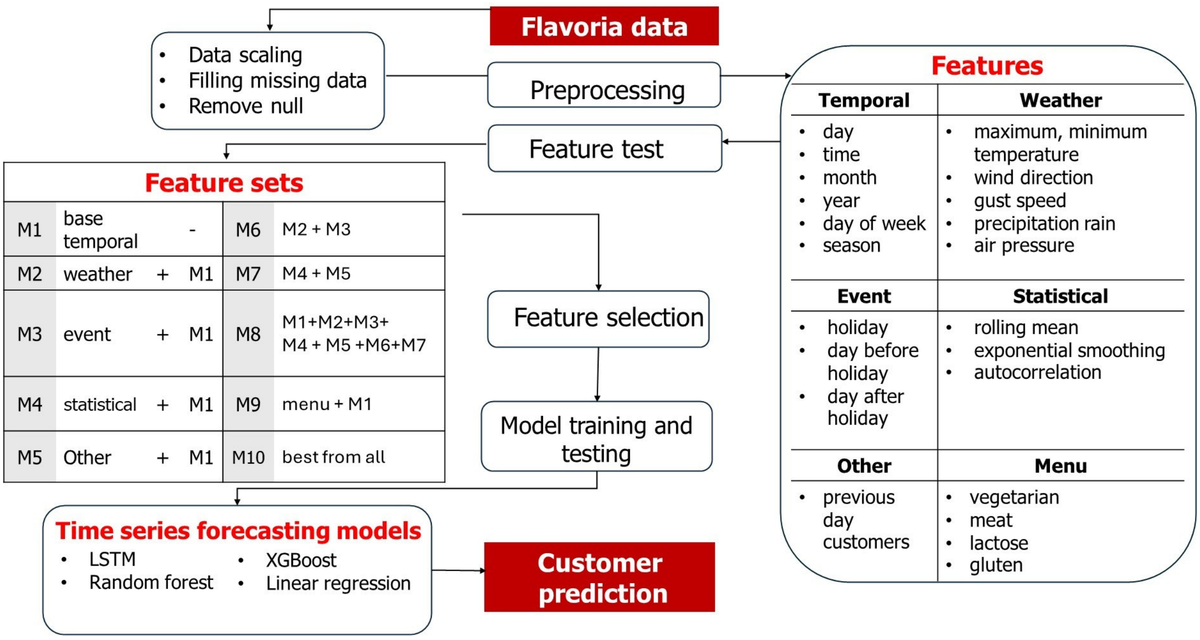

Figure 1. Workflow of the customer prediction model.

3.1. Flavoria data collection

Flavoria is a unique multidisciplinary research platform located on the University of Turku campus. It features a self-service lunch line restaurant (including a café and a grill) that is equipped with over 100 sensors, including 50 lunch component weighing scales. The lunch trays of the restaurant and the weighing scales are equipped with RFID tags and readers to collect data on the quantity of each food component taken by each dining customer. At the end of the dining, plate waste is measured automatically with the tray identifiers, allowing correlation of waste quantities with meal size. Dining behavior may also be tracked in the longer term via the use of a mobile app, *MyFlavoria*, that allows the normally anonymous lunch line dining sessions to be linked to the same customer each time they visit the restaurant. This linkage occurs at the cash register, where the tray, mobile app, and the cash register information are integrated. The Flavoria restaurant serves a wide range of diners, including university students, staff, and visitors. This setup allows us to collect both subjective data (through surveys) and objective data (via augmented sensors), resulting in a rich and comprehensive dataset for our study.

3.2. Data preprocessing

In real-world scenarios, data often contain inconsistencies and require cleaning. We address data cleaning and synchronization to prepare our data for machine learning models. Figure 2 presents the temporal distribution of unique customers per day over six years (2019–2024), providing insights into customer flow patterns, including peaks and troughs across different years and the notable impact of the COVID-19 pandemic in 2020. The x-axis represents the time, while the y-axis shows the count of unique customers. To enhance the quality and suitability of the data for our models, we remove outliers, address missing data, synchronize dates with weather data, and apply data scaling. Scaling is crucial to ensure that all features contribute equally to the model's performance, preventing features with larger ranges, like customer count, from dominating the learning process.

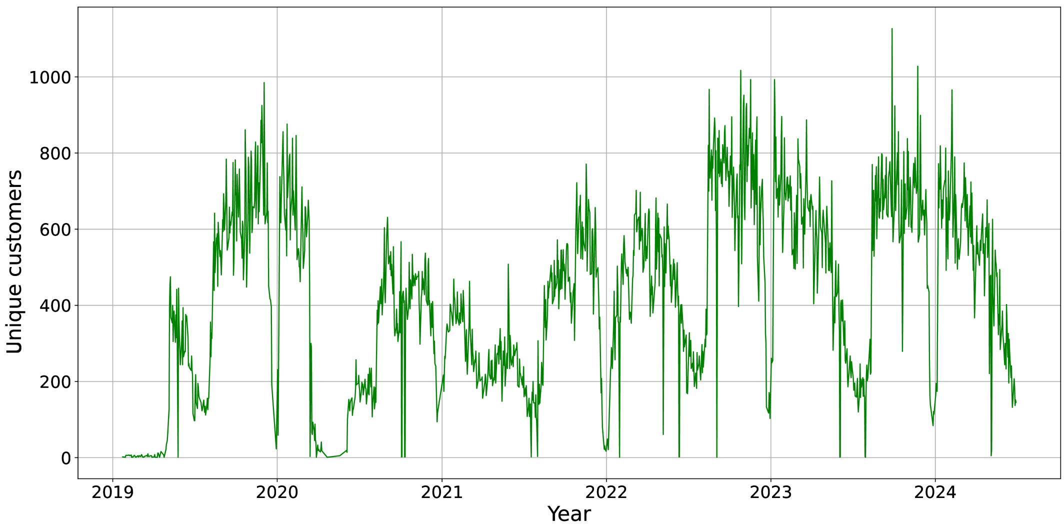

Figure 2. Temporal distribution of data points (2019–2024).

From the weather data, we preprocess the wind direction data to calculate the daily averages between 10:00 AM and 3:00 PM. Because wind direction is a circular variable, we use vector decomposition into Cartesian coordinates followed by recomposition to ensure accuracy. This approach avoids inaccuracies that could arise from the direct averaging of angular data.

We examine patterns and trends using visual tools, such as Figure 3, which shows the busiest hours of customer visits during the day. Understanding peak times helps us focus our preprocessing on the most active periods. To identify outliers, we rely on visual inspection, looking for gaps or points that clearly deviate from the overall pattern. These outliers may result from data entry errors or unusual events, and can affect the accuracy of our models if not addressed.

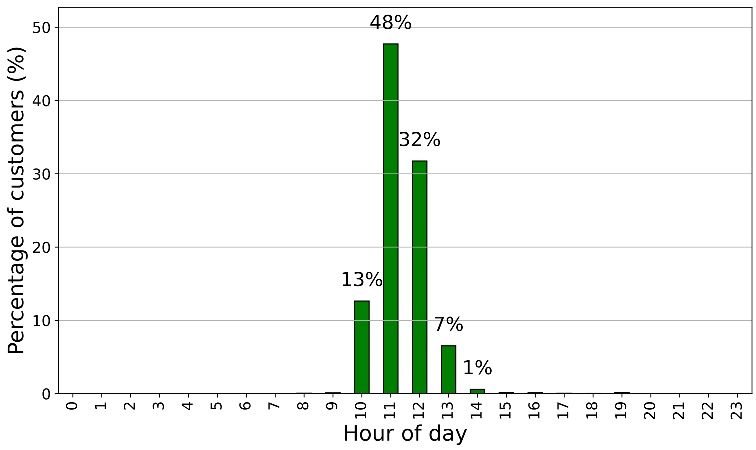

Figure 3. Customer visit peak times at the restaurant.

Our dataset, covering the period from 23 January 2019 to 1 February 2024, initially contained missing entries. These gaps may have resulted from sensor malfunctions or data transmission errors. To maintain the accuracy and reliability of our analysis, we applied data scaling and implemented appropriate methods to handle and impute the missing values where necessary.

For example, our original dataset, shown in Table 3, includes data from 2019 to 2024, with a total of 1,881,871 rows. However, after filtering the dataset to focus on data from 2021 to 2024, we narrowed it down to 1,347,356 rows. This filtering allows us to concentrate on the most recent and relevant data, potentially mitigating the impact of earlier outliers.

Table 3. Data overview (2019–2024).

Dataset	Number of rows
Original	1,881,871 rows
Filtered	1,347,356 rows

Removing null values is another essential step in data preprocessing to ensure the dataset's integrity and the reliability of our machine learning model. Null values can occur due to various reasons, such as sensor malfunctions or data entry errors.

Furthermore, we adjust for temporal consistency by localizing the time data to the Helsinki (Finland) time zone, where the restaurant is located. This ensures that all timestamps align with the local time, which is crucial for accurate analysis and modeling.

3.3. Feature test

We test feature combinations from M1 to M9, as described in Section 3.4, to find the best set for predicting customer visits. These features are common in prediction models according to the literature, including the baseline, weather, holidays, statistical features, and menu features. Our goal is to identify which features most significantly improve the accuracy of our predictions.

After testing all the combinations using machine learning methods, we chose the M10 feature set as the best, as explained in Section 3.5. This set includes important features such as exponential smoothing, the previous day's customers, and rolling mean, which capture key patterns in customer behavior.

3.4. Features

The selected features play a crucial role in how well our predictive model performs. In this section, we outline the features we use to predict the number of customer visits. We conduct feature engineering to enhance the predictive capabilities of the dataset. We start with the baseline features, which are the basic building blocks of our model. We then add other features to improve accuracy and better capture patterns in the data.

3.4.1. Baseline features

We utilize time-related features such as year, month, day, day of the week, and season (quarters of the year), extracting them from localized timestamps. These features provide additional context for analysis and help us understand temporal patterns in customer behavior. On the basis of Figure 3, we focus on the peak hours from 10:00 AM to 3:00 PM. Our goal is to predict the number of customers for future dates. Therefore, we focus our analysis on these hours, ensuring that the dataset captures the most significant patterns in customer traffic.

Table 4 presents the baseline features after time zone localization and feature extraction. The dataset includes columns for time, customer count, year, month, day, day of the week, and season. The customer count serves as the target variable, representing the number of customers for a single day. Our primary objective is to predict the number of customers for future dates.

Table 4. Time series data overview.

Time	Customers	Year	Month	Day	Day Of week
2021-01-04	217	2021	1	4	0
2021-01-05	174	2021	1	5	1
2024-01-31	626	2024	1	31	2
2024-02-01	608	2024	2	1	3

These baseline features provide a foundational framework for our predictive model, enabling us to analyze and forecast customer count effectively. We employ a cyclic representation of the day of the week, differentiating it from the day of the month. We represent the day of the week as a digit from 0 to 6, where 0 is Monday and 6 is Sunday. This allows us to capture its cyclical nature using the formulas (3.1) and (3.2) shown below:

$$\text{day of week}_{\text{sin}} = \sin\left(\text{dow} \times \frac{2\pi}{7}\right), \tag{3.1}$$

$$\text{day of week}_{\text{cos}} = \cos\left(\text{dow} \times \frac{2\pi}{7}\right). \tag{3.2}$$

These representations are given by Eqs (3.1) and (3.2), respectively, where *dow* denotes the numerical representation of the day of the week. Furthermore, considering that our restaurant is closed on weekends, we focus only on data from Monday to Friday. Seasonal features are derived from the quarter of the year, capturing potential seasonal and yearly trends in restaurant demand. The weekly, monthly, and yearly cyclical features provide information to the model from the repeated behavioral patterns of the customers.

3.4.2. Weather features

Local weather can influence consumers' decisions to dine out. Sunny days often attract customers, while rainy weather can deter them. To understand this impact, we employ weather data as a feature in our customer prediction models. We obtain hourly weather data for Turku, Finland, in CSV format from the Finnish Meteorological Institute (FMI) *.

Upon acquiring the weather data, we perform several preprocessing steps. First, we remove unnecessary columns and rename relevant columns for better readability.

Next, we convert these columns (temperature, wind speed) into numeric data types to facilitate analysis. Since our study focuses on daily predictions, we process the weather data to obtain daily averages. Specifically, we calculate the average values for each weather parameter from 10:00 AM to 3:00 PM. It is important to note that the historical weather data for Turku may contain missing values.

To address this issue, we adopt a simple imputation strategy. If a particular weather parameter is missing for a given day, we use the corresponding value from the previous day. This approach ensures continuity in the data and allows us to proceed with our analysis. By testing these commonly used weather features alongside other operational variables, our study provides a comprehensive evaluation of their impact on customer predictions in the restaurant industry.

3.4.3. Holidays

Holidays are an important factor, as they can significantly influence restaurant activity. We use the Python programming language library Holidays to find Finnish holidays from our dataset's time column. The library helps us create three binary features: "is holiday", "day before holiday", and "day after holiday".

Many use holidays and the day before as features, but we also include the day after. This gives the model more information on how holidays affect restaurant traffic. These features are stored as True or False (1 or 0) values, showing if the day, the previous day, or the next day is a holiday. This lets us add holiday data to our model easily, and our analysis shows that including holiday data improves our prediction accuracy.

In determining which days to mark as holidays, we consider both official holidays, such as Finland's Independence Day, Easter, Christmas, Boxing Day, and New Year, as well as unofficial holidays like Midsummer (Juhannes in Finnish), Father's Day, and Labor Day (Vappu in Finnish).

The Python library gives us the correct holiday dates for Finland. This lets us add holiday data to our model easily. Our analysis shows that including holidays improves our prediction accuracy. Since the restaurant closes on holidays, we pay close attention to customer behavior on the day before and after. This is important for us because we serve hospital workers, and it helps us plan better and keep customers happy.

*https://en.ilmatieteenlaitos.fi/local-weather/turku?forecast=short

3.4.4. Previous data on visitors

We consider the number of customers from the previous day as one of the key features. By looking at how many visitors came the day before, we can understand trends in customer traffic and make better predictions for the future. To include this information in our model, we simply look back one day to see how many people visited the restaurant. If there are any missing values (like days when the restaurant was closed), we fill them with 0 to avoid any problems with our analysis.

This "previous day visitors" feature is combined with our statistical features in Table 5 labeled as M4. By putting these together, we are able to capture both the history of customer visits and the statistical patterns in our data. This helps us build a stronger model that can make more accurate predictions about future customer numbers.

Table 5. Example of statistical feature computation.

Customers	**Rolling mean**	**Exponential smoothing**	**Autocorrelation**
217	462.6	217.0	0.82
174	462.6	202.6	0.82
0	462.6	135.1	0.82
626	462.6	178.7	0.82
608	183.6	206.1	0.82

3.4.5. Statistical feature

The statistical feature set comprises three distinct features: rolling mean, exponential smoothing, and autocorrelation. These features are designed to capture patterns and trends from historical data, enhancing the accuracy of our predictive model. Table 5 illustrates the combination of statistical and other features in the M4 feature section.

The rolling mean feature calculates the average number of customers over a window spanning the previous 5 days (window = 5). It employs a moving average technique on a specified window size to linearize the time series, commonly known as smoothing the series. Formula (3.3) demonstrates the computation of the rolling mean.

$$\text{Rolling mean} = \frac{1}{n} \sum_{i=1}^{n} X_i. \tag{3.3}$$

In exponential smoothing, customers are exponentially smoothed over 5 days to compute a smoothed average. Unlike simple rolling mean averages, exponential smoothing focuses on recent trends while considering past data. Formula (3.4) shows the computation of exponential smoothing. The exponential smoothing formula is given below.

$$\text{Exponential smoothing} = \alpha \times X_t + (1 - \alpha) \times \text{previous smoothed value}. \tag{3.4}$$

Autocorrelation measures the correlation of a time series with itself at a time lag. It helps identify patterns where past values might influence future demand. Formula (3.5) shows how to calculate autocorrelation. The autocorrelation formula is given as follows.

$$\text{Autocorrelation} = \frac{\sum_{t=1}^{n} (X_t - \bar{X})(X_{t+h} - \bar{X})}{\sum_{t=1}^{n} (X_t - \bar{X})^2}. \tag{3.5}$$

If there are any missing values (NaN) in the rolling mean column, we fill them with the overall average of the rolling mean series itself, ensuring coverage for all days. In exponential smoothing, missing values (NaN) are filled with the overall average of the series, ensuring continuity and comprehensiveness in the data representation. Autocorrelation fills missing values in autocorrelation columns with the average to ensure a complete feature.

Table 5 provides an example of how statistical features are computed from the customer data. It illustrates the calculation of rolling mean, exponential smoothing, and autocorrelation, based on the data presented in Table 4.

3.4.6. Menu feature

The menu feature comprises four binary indicators: vegetarian, meat, lactose-free, and gluten-free. Each feature represents the availability of specific food types on a given day, where a value of 1 indicates presence and 0 absence.

For example, if a day's data records `vegetarian = 1`, `meat = 0`, `lactose-free = 1`, and `gluten-free = 1`, it signifies that vegetarian, lactose-free, and gluten-free options are available, but no meat dishes are offered. By incorporating these variables into our predictive model, we analyze how different dietary options influence customer behavior. This analysis supports data-driven decisions in menu planning, optimizing customer satisfaction, and operational efficiency.

3.5. Feature selection

In this section, we explain the feature selection process and adjust it to optimize predictive performance. Table 6 ranks the features according to their importance in predicting daily customer flow, with values expressed as percentages for clarity.

Table 6. Top high-score features for predicting daily customer flow.

Feature	**Importance (%)**
Previous day customers	65.5
Exponential smoothing	9.2
Is holiday	6.1
Rolling mean (5-day)	5.2
Rolling mean (30-day)	2.5
Average air pressure	1.5
Week of year	1.5
Average wind direction	1.3
Day	1.0
Day of week	0.9

We determine the top-ranked features by analyzing their influence across multiple machine learning models. We call this selection the M10 best feature set, derived from M1 to M9. *Previous day customers* holds the highest importance score of **65.5%**, highlighting its critical role in predicting daily customer flow by reflecting recent trends. This feature dominates the ranking, making it indispensable for short-term forecasting, as it accurately captures fluctuations in customer behavior, such as responses to promotions or events.

Exponential smoothing ranks second with **9.2%**, emphasizing its value in capturing the influence of recent patterns while smoothing out irregularities. This feature is particularly effective in maintaining the integrity of trends necessary for accurate predictions. *Is holiday*, at **6.1%**, captures the significant effect of holidays on customer behavior. Holidays often lead to sharp spikes or dips in customer flow, making this feature essential for accurate modeling during these periods. *Rolling mean (5-day)* and *Rolling mean (30-day)* contribute **5.2%** and **2.5%**, respectively. These features encapsulate short-term and long-term trends, offering insights into both weekly fluctuations and broader seasonal patterns. Table 7 describes how we create the feature combinations.

Table 7. Feature categories and combinations.

Model	Features	Model	Features
M1	Base features: number of customers, time, day, month, year, season, day of the week	M6	M2 + M4 (base + weather + statistics)
M2	M1 + weather features: average; max, and min temperature (°C); wind speed/direction; gusts; precipitation	M7	M3 + M4 (base + holiday + statistics)
M3	M1 + holiday features: is holiday, day before holiday, day after holiday	M8	M1 + menu features: vegetarian, meat, lactose-free, gluten-free
M4	M1 + statistical features: rolling mean, exponential smoothing, autocorrelation	M9	M1 + M2 + M3 + M4 + M5 + M6 + M7 + M8 (All features combined)
M5	M2 + M3 (base + weather + holiday)	M10	Selected best features from M9: Previous day customers, exponential smoothing, holidays, weather

Other features, such as *average air pressure* (**1.5%**) and *average wind direction* (**1.3%**), reflect the minor but notable impact of weather conditions on customer flow. While their influence is less prominent, they provide valuable context for environmental factors.

Finally, *Day* (**1.0%**) and *Day of Week* (**0.9%**) capture temporal patterns in customer flow, indicating specific days of the month or weekdays that may consistently affect customer behavior.

By focusing on these top features, as listed in Table 6, the model effectively identifies the most influential factors for predicting the restaurant's daily customer flow.

Overall, the features extracted from various categories, including base, weather, holiday, statistics, and others, menu, and the selected best features from all categories provide additional insights into the raw daily customer count data. These features identify trends, smooth out noise, and reveal potentially hidden recurring patterns, thereby augmenting the accuracy of the forecasting model. The flexibility to adjust the window size for calculating statistical features allows for fine-tuning of the performance. We chose a window size of 5 days to align with our restaurant's operational schedule, which operates five days a week, with the weekends being closed.

Figure 4 illustrates the complete set of features used in our predictive models, showcasing how different features are combined to capture diverse patterns in the data. Overall, the selected features from various categories provide critical insights and improve the accuracy of our predictive model for restaurant customer counts.

Base		Weather			
Feature	Values	Feature	Values	Feature	Values
• Day	5	• Maximum temperature	-6.1	• Precipitation rain	0.08
• Month	1	• Minimum temperature	-6.5	• Air pressure	1031
• Year	21	• Wind direction	48.4		
• Day of week	1	• Gust speed	2.0		
• Season	1				
• customers	174				
Event		**Statistical**		**Menu**	
• Holiday	0	• Rolling mean	462.6	• Vegetarian	0
• Day before holiday	0	• Exponential smoothing	202	• Meat	1
• Day after holiday	1	• Autocorrelation	0.82	• Lactose	1
				• Gluten	1
		Other			
		• Prevous day	217		

Figure 4. Example of all features.

3.6. Model training and testing

Once the data had been prepared, we started the training, testing, and validation process of the predictive machine learning models. We began by splitting the dataset into two parts: One for training the models and the other for testing and validation to see how well they work. This approach enables us to evaluate how well the models generalize to unseen data, an essential factor for ensuring their effectiveness in real-world scenarios.

We use four different types of machine learning methods to predict customer visits to our restaurant. These models are applied to the preprocessed data to evaluate their predictive performance. We divided our dataset into two parts: 80% for training and 20% for testing. This approach allows the models to learn from the majority of the data while testing their ability to generalize to unseen data. Each model is described separately as follows.

Random forest is a popular forecasting technique based on multiple decision trees. It is an effective and easy-to-implement model with high accuracy, widely used in prediction and forecasting (for example, see [6, 7, 11, 13, 14, 27, 29]). It operates like a group of trees, each offering its opinion on how many customers will come. For instance, on a rainy day, a random forest may predict fewer customers because people might prefer to stay home. Conversely, on a sunny day, it might predict more customers because people may want to enjoy a meal outside. In general, random forest performs moderately well, providing predictions that are not too far off from the actual numbers. It handles tricky situations like noisy data and complex patterns adeptly. Additionally, it is not overly confused by outliers or unusual data points. However, understanding its decision-making process can be challenging, although the decision trees used can be visualized.

We use XGBoost, a decision-tree-based algorithm known for its high performance and scalability [13, 27, 29]. It applies gradient boosting by building a series of decision trees, where each tree corrects the errors of the previous one to reduce prediction errors. This model is trained to capture subtle patterns and interactions among the features that contribute to customer visit predictions. XGBoost operates like an excellent learner, continuously improving itself to make better predictions. It can handle complex relationships between different factors quite well. It requires a lot of fine-tuning to perform at its best and can overfit to the data.

An LSTM model, as used by [9,22], is a type of recurrent neural network (RNN) designed to handle sequential data. The LSTM model learns to recognize long-term dependencies and temporal patterns in the input sequences of features, which are crucial for predicting customer visit trends over time. LSTM can remember long-term patterns effectively and works best with data that change over time. However, it requires a significant amount of time and computational power to work with. It sometimes struggles with highly complex data.

Linear regression [2, 17, 20–22, 26] is a fundamental and straightforward model often used in predictive analytics. It works by fitting a straight line through a set of data points to predict future values on the basis of the relationships between variables. While linear regression is easy to understand and implement, it tends to perform best in scenarios where relationships between features are more linear and less complex.

The selection of predictive models for this case study was intentionally diverse to ensure a robust and comprehensive comparison. Our goal was to evaluate algorithms that represent different levels of complexity and underlying logic, from simple linear baselines to state-of-the-art nonlinear models. The chosen models were as follows.

The linear regression (LR) model was selected to serve as a fundamental baseline. It helps establish a performance benchmark and allows us to assess the extent to which a simple linear combination of features can predict customer flow. A strong baseline is essential for justifying the need for more complex models.

Random forest (RF) and XGBoost, two tree-based ensemble methods, were chosen because they are the consistently top performers for structured (tabular) data forecasting tasks [15, 31]. Random forest operates by building a multitude of decision trees (a technique known as bagging) and is known for its robustness and ability to prevent overfitting.

XGBoost is a powerful gradient boosting algorithm that builds a series of models, where each one learns from the mistakes of the last. It is widely recognized as a top-performing method for structured data, so we included it to benchmark its state-of-the-art accuracy in our forecasting context.

A long short-term memory (LSTM) network specifically addresses the time-series nature of our data, so we included an LSTM model. As a type of RNN, LSTMs are explicitly designed to recognize and learn long-term dependencies in sequential data, making them a natural choice for forecasting problems where past patterns heavily influence future outcomes [22].

By comparing the performance of these four distinct models, we can draw more reliable conclusions about which architectural approach—linear, tree-based ensemble, or recurrent deep learning—is best suited for the practical challenge of predicting customer counts in a self-service restaurant.

The key parameters for each machine learning model were selected to optimize forecasting performance. The final parameters, summarized in Table 8, result from a process of grid search optimization, empirical testing, and insights from prior studies.

Our LSTM model uses a single hidden layer with 50 memory units and a time step window of 3 to capture short-term data dependencies. We train the model for 50 epochs with a batch size of 1, using the Adam optimizer and a learning rate of 0.001. We do not apply dropout, as the model's low complexity sufficiently mitigates the risk of overfitting.

Table 8. Key parameters for machine learning models.

Model	Key parameters
XGBoost	max_depth=5, n_estimators=200, learning_rate=0.1, subsample=0.8, colsample_bytree=0.8
Random forest	n_estimators=100, max_features='sqrt'
LSTM	1 hidden LSTM layer with 50 units, no dropout, optimizer=Adam, learning_rate=0.001, epochs=50, batch_size=1
Linear regression	Default parameters (scikit-learn implementation)

For the random forest and XGBoost models, we perform a grid search to tune the key hyperparameters. We test the number of trees (from 100 to 200), learning rates (from 0.01 to 0.3), and maximum tree depths (from 3 to 7). The final selected configurations provide an effective balance between predictive accuracy and computational cost.

4. Results

In this section, we evaluate the performance of four machine learning methods, random forest, XGBoost, LSTM, and linear regression, in predicting daily customer counts in the Flavoria self-service restaurant setting. We assess the accuracy of these models using three key metrics: root mean squared error(RMSE), mean absolute error (MAE), and mean squared error (MSE). These metrics provide insights into both the magnitude and distribution of errors, helping us understand the strengths and weaknesses of each model across different feature sets.

We deliberately excluded the commonly used mean absolute percentage error (MAPE) because it is highly sensitive to zero and near-zero actual values, leading to instability and undefined results. Given that our dataset initially included days with zero customer counts, using MAPE would have introduced significant distortions. Although we performed data cleaning, it is still unclear whether removing all zero-customer days is justified, as doing so might exclude valid operational periods (e.g., public holidays or closures) that are relevant for prediction. To address these issues more rigorously, future work should explore alternative metrics such as the mean absolute scaled error (MASE), which is scale-independent and handles zero values more gracefully. MASE could provide a more robust basis for comparing models' performance, especially in contexts with intermittent demand.

The RMSE, MAE, and MSE results across feature sets (M1 to M10) for each model are presented in Tables 9–11, respectively, with an additional "Avg." column that provides the average error across all datasets for each model. Figures 5–7 visualize these metrics, with each model's results displayed as individual bars within the categories of RMSE, MAE, and MSE. These visual representations aid in understanding models' performance differences across various feature sets.

Table 9. RMSE for different models.

Model	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	Avg.
random forest	132	147	103	127	115	132	107	106	114	117	120
XGBoost	138	141	98	135	107	137	116	99	107	111	119
LSTM	144	144	144	144	144	144	144	144	144	144	144
Linear regression	308	313	270	148	273	150	123	270	125	128	211

Table 10. MAE for different models.

Model	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	Avg.
Random forest	96	107	84	89	90	90	80	84	82	83	88
XGBoost	99	100	77	93	85	95	84	79	79	76	86
LSTM	102	102	100	100	100	103	101	102	100	100	101
Linear regression	224	235	201	103	213	106	91	201	94	95	156

Table 11. MSE for different models.

Model	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	Avg.
Random forest	17,541	21,672	10,581	16,193	13,300	17,488	11,521	11,170	12,932	13,609	14,501
XGBoost	19,099	19,933	9,678	18,283	11,441	18,798	13,533	9,724	11,475	12,354	13,712
LSTM	21,709	21,055	21,085	20,829	20,924	21,784	21,208	20,934	21,003	20,815	21,035
Linear regression	94,689	97,679	72,797	21,747	74,472	22,501	15,194	72,797	15,661	16,426	50,396

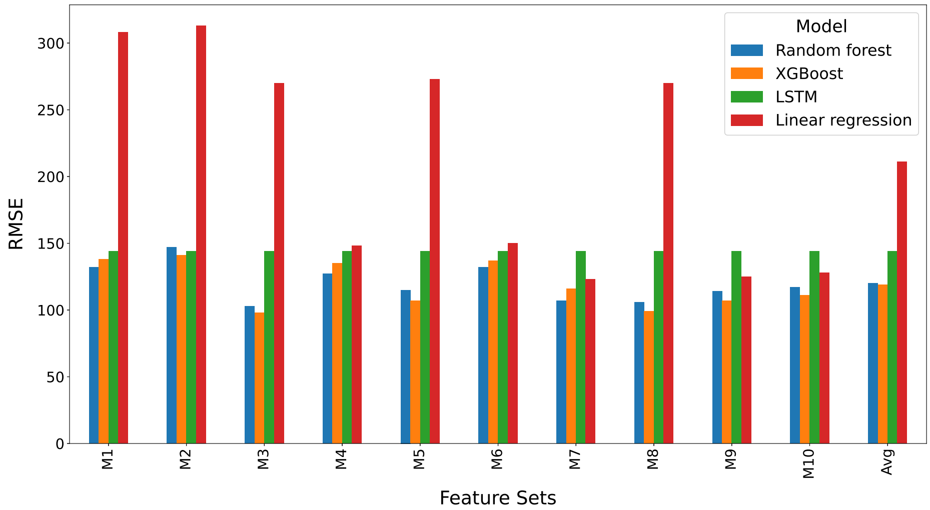

Figure 5. RMSE of different machine learning models across feature sets.

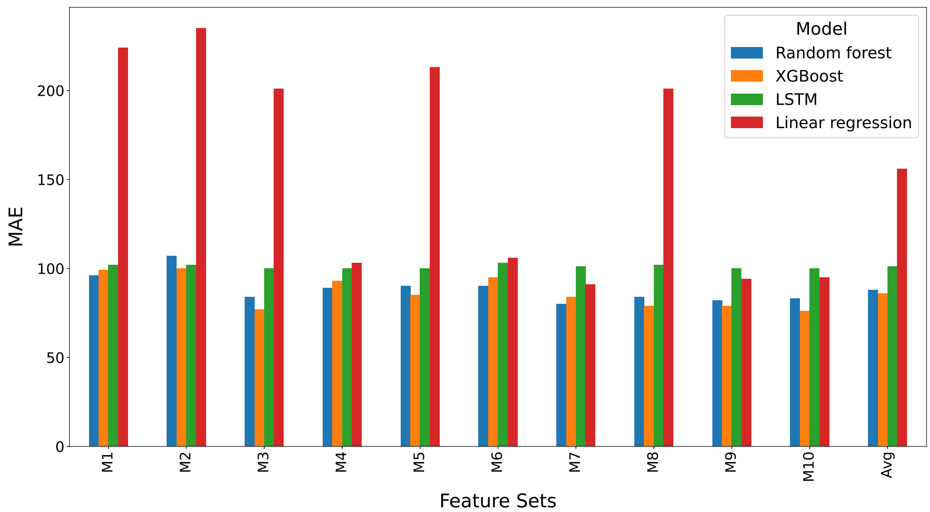

Figure 6. MAE of different machine learning models across feature sets.

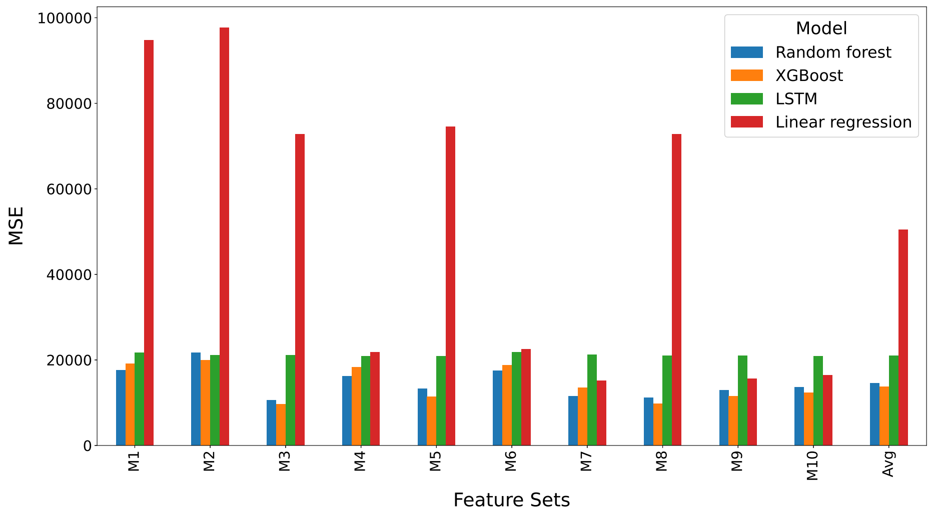

Figure 7. MSE of different machine learning models across feature sets.

Table 11 shows that the XGBoost model achieves the lowest MSE values on average across all feature sets, indicating its strong performance in minimizing prediction errors. With an average MSE of 13712 and an average MAE of 86, XGBoost proves to be particularly effective in capturing complex relationships within the data, making it suitable for predicting customer flow in dynamic environments.

Conversely, the linear regression model exhibits the highest average MSE value (50, 396) and an average MAE of 156 across most feature sets, except in cases with simpler linear relationships. This highlights the limitations of linear regression when applied to complex datasets with Nonlinear dependencies.

The LSTM model's performance remains consistent but with relatively high error values across all metrics, yielding an average MSE of 21, 035 and an average MAE of 101. This suggests that LSTM may not be the best fit for this dataset.

Figure 8 visualizes a comparison between the predicted customer flow values generated by the XGBoost model and the actual (original) customer flow data. The comparison helps in understanding how closely the model's predictions match the real-world data. Ideally, if the model is perfect, these points would align perfectly, indicating that every prediction exactly matches the actual value.

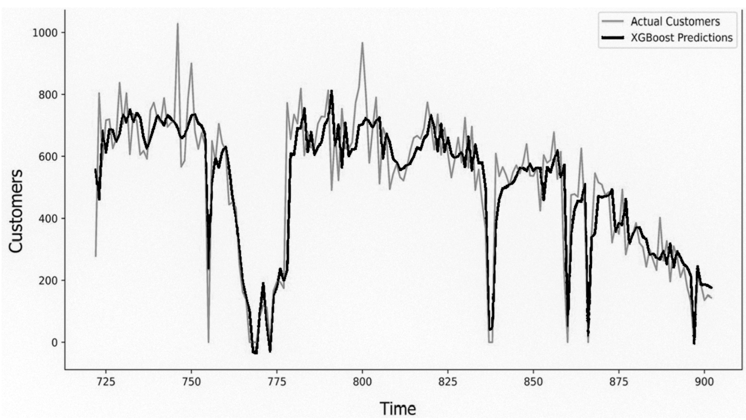

Figure 8. Comparison of the XGBoost model's performance to the actual number of customers each day.

5. Discussion

It is worth noting that previous research addressing the prediction of daily restaurant visitors utilizing weather variables such as temperature, precipitation, and wind speed has generally found only limited effects. In our case study, we also included weather-related features. Specifically, we observed that average air pressure and wind direction had only a minor influence on our prediction model (see Table 6).

This finding underscores the importance of including a dedicated feature evaluation step in the workflow (Figure 1) to assess the contribution of each variable. It also suggests that certain weather conditions, even if not consciously perceived by individuals, may subtly affect dining behavior. For instance, on rainy or windy days, customers may be less inclined to leave their office buildings, opting instead to dine on-site, particularly in self-service restaurants located within those buildings.

The results presented in Figures 5–7 offer a clear comparison of the strengths and weaknesses of each machine learning model across different performance metrics. XGBoost consistently emerges as the most reliable model, demonstrating strong performance across the RMSE, MAE, and MSE metrics, particularly with an average MAE of 86, which is the lowest among the models. Its ability to minimize both large and small errors makes it a highly effective tool for predicting daily customer counts, particularly in complex, Nonlinear datasets.

Random forest performs reasonably well, with an average MAE of 88, but its higher error values relative to XGBoost suggest that it may not fully capture the intricate relationships within the data. While random forest is a robust and reliable model, its limitations become apparent when faced with datasets that require more nuanced analysis.

LSTM's underperformance across all metrics, with an average MAE of 101, highlights a potential mismatch between the model and the dataset's characteristics. LSTM models excel in scenarios with strong temporal dependencies, but the relatively weak or Nonlinear temporal patterns in this dataset likely limit their effectiveness.

Our study improves on past research by using an LSTM model with added features. Although we tested the LSTM model in our research, its results were not as accurate as we had hoped. However, we include it in our review to provide a full picture of the techniques used in this area, recognizing both its limitations and potential when combined with other models.

linear regression surprises by performing well in specific feature sets like M10, where the relationships between features are more linear, achieving an average MAE of 156 despite having generally higher errors in Nonlinear datasets. This underscores the importance of feature engineering and suggests that simpler models can still be highly effective when the data are well-aligned with the model's assumptions.

These findings suggest that XGBoost, with its consistently low MAE and MSE values, should be the preferred model for most complex, Nonlinear datasets, especially when accuracy is critical. However, the strong performance of linear regression in certain feature sets highlights that simpler models should not be overlooked, particularly when data characteristics are well-understood. Future research could explore hybrid models that combine the strengths of different approaches or focus on further refining feature engineering techniques to enhance models' performance. Additionally, applying these models to other contexts, such as different restaurant environments, could provide valuable insights into their generalizability.

In addition to comparing machine learning methods and their fit with the data structure, our experiments provide some general insights. We found that including external factors such as weather, holidays, and the day of the week greatly improved prediction accuracy. Simple techniques like smoothing recent data (using rolling averages and exponential smoothing) helped stabilize predictions during unusual periods, such as the COVID-19 disruptions. The results also show that cleaning and preparing the data carefully are just as important as choosing the right model. While XGBoost and random forest performed best overall because they handle complex patterns well, simpler models like linear regression also gave good results when the data were more stable and the features were chosen carefully. Deep learning models like LSTM need more data and fine-tuning to work well, which makes them less practical in smaller datasets. These findings suggest that selecting the best approach depends not only on the model itself but also on the quality of data, the available features, and how easily restaurant managers can interpret and use the results.

In general, we see that there is still room for improvement, and one of the ways to increase the accuracy is to introduce new data sources to the model. An example of additional features is the menu offering better categorization to determine the differences in the lunch offerings and the ability to compare the menu offerings with restaurants nearby. Another example of an additional data source might be elevator usage and floor occupancy data from the same morning, which could be used to predict the customers for the same day.

Another promising direction for future research is the application of predictive data to improve kitchen operations. In this context, there is a need for three types of forecasts: one-week-ahead, next-day, and same-day predictions. These serve distinct operational purposes—advance forecasts support ingredient ordering, next-day predictions assist in ingredient preparation, and same-day forecasts guide the final preparation of dishes. Accurate predictions at each of these stages can lead to more informed decision-making and significantly reduce food waste throughout the kitchen workflow.

In practice, the model offers valuable guidance for managing inventory, preparing a suitable amount of food, and scheduling staff. For instance, it can predict a spike in customer traffic the day before major Finnish holidays, enabling managers to plan for these busy periods. These will help the restaurant to be both more sustainable and, at the same time, more profitable.

6. Conclusions

In conclusion, this case study set out to identify an effective and practical machine learning approach for predicting daily customer flow in a university's self-service restaurant. Through a systematic evaluation of four models, namely linear regression, random forest, XGBoost, and LSTM, we found that XGBoost consistently and significantly outperformed the other models across all 10 feature sets. However, linear regression delivers good results and is particularly effective when using the specific feature set M10, highlighting its ability to capture simpler, more linear relationships within the data.

Our findings indicate that XGBoost excels in leveraging recent historical data and applying exponential smoothing to capture short-term customer behavior trends. Additionally, our analysis underscores the significance of features like exponential smoothing, which effectively captures recent trends in customer visits, as well as the importance of features such as previous day customer count and the rolling mean of 5 days, which also highlight the influence of historical trends on future demand.

Our analysis revealed the most impactful features to be the previous day's customer count, exponential smoothing outputs, and external factors like holidays and weather. In light of these findings, this case study provides a clear recommendation for operators of similar establishments: Implementing an XGBoost model with a feature set that includes recent historical trends is the most robust strategy for accurate lunchtime customer forecasting. While this provides a strong blueprint, future work could explore ensemble methods to potentially achieve further incremental gains in accuracy.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

This research was supported by Business Finland, under the Veturi program with the Dining Flow project (6547/31/2022).

Conflict of interest

The authors declare no conflict of interest.

References

1. T. Ahmad, C. Huanxin, D. Zhang, H. Zhang, Smart energy forecasting strategy with four machine learning models for climate-sensitive and non-climate sensitive conditions, *Energy*, **198** (2020), 117283. https://doi.org/10.1016/j.energy.2020.117283
2. R. Fok, A. Lasek, L. Jiye, A. An, Modeling daily guest count prediction, *Big Data and Information Analytics*, **1** (2016), 299–308. https://doi.org/10.3934/bdia.2016012
3. G. Boomija, A. Anandaraj, S. Nandhini, S. Lavanya, Restaurant visitor time series forecasting using autoregressive integrated moving average, *J. Comput. Theor. Nanosci.*, **15** (2018), 1590–1593. https://doi.org/10.1166/jctn.2018.7345
4. J. Chen, H. Liu, J. He, Predicting the influence of group buying on the restaurant's popularity by online reviews, *Proceedings of 12th International Conference on Fuzzy Systems and Knowledge Discovery (FSKD)*, 2015, 1068–1072. https://doi.org/10.1109/FSKD.2015.7382090
5. M. Ching, R. de Dios Bulos, Improving restaurants' business performance using yelp data sets through sentiment analysis, *Proceedings of the 3rd International Conference on E-commerce, E-Business and E-Government*, 2019, 62–67. https://doi.org/10.1145/3340017.3340018
6. E. Fernandes, S. Moro, P. Cortez, F. Batista, R. Ribeiro, A data-driven approach to measure restaurant performance by combining online reviews with historical sales data, *Int. J. Hosp. Manag.*, **94** (2021), 102830. https://doi.org/10.1016/j.ijhm.2020.102830
7. Y. Fu, C. Soman, Real-time data infrastructure at uber, *Proceedings of the 2021 International Conference on Management of Data*, 2021, 2503–2516. https://doi.org/10.1145/3448016.3457552

8. *S. Haas, E. Kuehl, J. Moran, K. Venkataraman, How restaurants can thrive in the next normal*, McKinsey and Company, 2020. Available from: `https://www.mckinsey.com/industries/retail/our-insights/how-restaurants-can-thrive-in-the-next-normal`.
9. T. Iha, I. Kawamitsu, A. Ohshiro, M. Nakamura, An lstm-based multivariate time series predicting method for number of restaurant customers in tourism resorts, *Proceedings of 36th International Technical Conference on Circuits/Systems, Computers and Communications (ITC-CSCC)*, 2021, 1–4. https://doi.org/10.1109/ITC-CSCC52171.2021.9501432
10. S. Kazmi, A. Bozanta, M. Cevik, Time series forecasting for patient arrivals in online health services, *Proceedings of the 31st Annual International Conference on Computer Science and Software Engineering*, 2021, 43–52.
11. K. Y. Koay, D. L. T. Ong, I. Idris, N. S. B. Kamarulzaman, C. W. Cheah, The rise of solo dining: prediction and consumer profiling, *Journal of Foodservice Business Research* , in press. https://doi.org/10.1080/15378020.2024.2307682
12. L. Koivunen, S. Laato, S. Rauti, J. Naskali, P. Nissilä, P. Ojansivu, et al., Increasing customer awareness on food waste at university cafeteria with a sensor-based intelligent self-serve lunch line, *Proceedings of IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*, 2020, 1–9. https://doi.org/10.1109/ICE/ITMC49519.2020.9198571
13. P. Kongnim, N. Cooharojananone, T. Chavarnakul, Forecasting dining times in a full-service thai hotpot restaurant using random forest classifier, *Proceedings of the 2023 8th International Conference on Mathematics and Artificial Intelligence*, 2023, 15–18. https://doi.org/10.1145/3594300.3594303
14. K. Kumaresan, P. Ganeshkumar, Software reliability prediction model with realistic assumption using time series (S)ARIMA model, *J. Ambient Intell. Human Comput.*, **11** (2020), 5561–5568. https://doi.org/10.1007/s12652-020-01912-4
15. X. Ma, Y. Tian, C. Luo, Y. Zhang, Predicting future visitors of restaurants using big data, *Proceedings of International Conference on Machine Learning and Cybernetics (ICMLC)*, 2018, 269–274. https://doi.org/10.1109/ICMLC.2018.8526963
16. D. Dilkhush Mihirsen, J. T. Joseph, B. Renisha, Time series analysis for supply chain planning in restaurants, *Proceedings of 5th International Conference on Computing, Communication and Security (ICCCS)*, 2020, 1–5. https://doi.org/10.1109/ICCCS49678.2020.9276872
17. G. Pawar, A. Ganesan, R. Moghe, B. Nayak, T. Khanna, K. M. Adusimilli, Learning to predict two-wheeler travel distance, *Proceedings of the 5th Joint International Conference on Data Science and Management of Data (9th ACM IKDD CODS and 27th COMAD)*, 2022, 153–161. https://doi.org/10.1145/3493700.3493726
18. F. Rarh, D. Pojee, S. Zulphekari, V. Shah, Restaurant table reservation using time-series prediction, *Proceedings of 2nd International Conference on Communication and Electronics Systems (ICCES)*, 2017, 153–155. https://doi.org/10.1109/CESYS.2017.8321254
19. T. Sarapisto, L. Koivunen, T. Mäkilä, A. Klami, P. Ojansivu, Camera-based meal type and weight estimation in self-service lunch line restaurants, *Proceedings of 12th International Conference on Pattern Recognition Systems (ICPRS)*, 2022, 1–7. https://doi.org/10.1109/ICPRS54038.2022.9854056

20. A. Schmidt, W. U. Kabir, T. Hoque, Machine learning based restaurant sales forecasting, *Mach. Learn. Knowl. Extr.*, **4** (2022), 105–130. https://doi.org/10.3390/make4010006
21. H. Shah, P. Fränti, Combining statistical, structural, and linguistic features for keyword extraction from web pages, *Applied computing and intelligence*, **2** (2022), 115–132. https://doi.org/10.3934/aci.2022007
22. J. Shen, M. O. Shafiq, Short-term stock market price trend prediction using a comprehensive deep learning system, *J. Big Data*, **7** (2020), 66. https://doi.org/10.1186/s40537-020-00333-6
23. M. Siek, K. Guswanto, Developing accurate predictive model using computational intelligence for optimal inventory management, *Proceedings of 13th International Conference on Information and Communication Technology and System (ICTS)*, 2021, 218–223. https://doi.org/10.1109/ICTS52701.2021.9608492
24. C. Su, Passenger flow forecast of catering business based on autoregressive integrated moving average and smoothing index prediction model, *Proceedings of International Signal Processing, Communications and Engineering Management Conference (ISPCEM)*, 2020, 53–57. https://doi.org/10.1109/ISPCEM52197.2020.00016
25. T. Tanizaki, Y. Hanayama, T. Shimmura, Forecasting customers visiting using machine learning and characteristics analysis with low forecasting accuracy days, In: *Advances in production management systems: towards smart and digital manufacturing*, Cham: Springer, 2020, 670–678. https://doi.org/10.1007/978-3-030-57997-5_77
26. T. Tanizaki, S. Kozuma, T. Shimmura, Forecasting the number of customers visiting restaurants using machine learning and statistical method, In: *Advances in production management systems: artificial intelligence for sustainable and resilient production systems*, Cham: Springer, 2021, 189–197. https://doi.org/10.1007/978-3-030-85906-0_21
27. S. Wang, Y. Zhang, X. Lin, Y. Hu, Q. Huang, B. Yin, Dynamic hypergraph structure learning for multivariate time series forecasting, *IEEE Trans. Big Data*, **10** (2024), 556–567. https://doi.org/10.1109/TBDATA.2024.3362188
28. D. Wu, S. Zhong, H. Song, J. Wu, Do topic and sentiment matter? predictive power of online reviews for hotel demand forecasting, *Int. J. Hosp. Manag.*, **120** (2024), 103750. https://doi.org/10.1016/j.ijhm.2024.103750
29. Z. Yu, M. Tian, Z. Wang, B. Guo, T. Mei, Shop-type recommendation leveraging the data from social media and location-based services, *ACM Trans. Knowl. Discov. D.*, **11** (2017), 1–21. https://doi.org/10.1145/2930671
30. J. Zhang, C. Chow, Geosoca: exploiting geographical, social and categorical correlations for point-of-interest recommendations, *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2015, 443–452. https://doi.org/10.1145/2766462.2767711
31. X. Zhang, S. Yang, G. Srivastava, M. Chen, X. Cheng, Hybridization of cognitive computing for food services, *Appl. Soft Comput.*, **89** (2020), 106051. https://doi.org/10.1016/j.asoc.2019.106051

32. M. A. Zhao, R. Jayadi, Forecasting daily visitors and menu demands in an indonesian chain restaurant using support vector regression machine, *Proceedings of International Conference on Artificial Intelligence and Mechatronics Systems (AIMS)*, 2021, 1–6. https://doi.org/10.1109/AIMS52415.2021.9466036

33. L. Zhu, W. Yu, K. Zhou, X. Wang, W. Feng, P. Wang, et al., Order fulfillment cycle time estimation for on-demand food delivery, *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2020, 2571–2580. https://doi.org/10.1145/3394486.3403307