Irish Interdisciplinary Journal of Science & Research (IIJSR)
Volume 8, Issue 3, Pages 39-55, July-September 2024

In-Flight Sales Prediction Using Machine Learning

Habtamu Awoke[1*], Lemlem Kassa (PhD.)[2] & Tarikwa Tesfa[3]

[1-3]AI and Robotics Center of Excellence, Addis Ababa Science and Technology University, Addis Ababa, Ethiopia.
Corresponding Author (Habtamu Awoke) Email: habtamu.awoke@aastustudent.edu.et[*]

DOI: https://doi.org/10.46759/IIJSR.2024.8305

Copyright © 2024 Habtamu Awoke et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Article Received: 02 July 2024 *Article Accepted: 07 September 2024* *Article Published: 18 September 2024*

ABSTRACT

This study introduces an advanced machine learning approach for predicting in-flight sales, aiming to transform the airline retail landscape. By harnessing advanced algorithms and historical sales data, the proposed model offers precise predictions of passenger preferences, allowing airlines to optimize their in-flight inventory and tailor product offerings more effectively. This data-driven strategy is designed to enhance inventory management, increase ancillary revenue, and improve passenger satisfaction. The research explores six distinct machine learning models: Linear Regression (LR), Support Vector Regression (SVR), Extreme Gradient Boosting (XGBoost), Multilayer Perceptron Neural Network (MLP), Recurrent Neural Network (RNN), and Deep Neural Network (DNN). Among these, XGBoost has proven to be the most effective model for sales forecasting, demonstrating exceptional accuracy and reliability. The study also highlights key features influencing sales, with ItemCategory_ TOBACCO emerging as the most significant factor among all features, and TotalPassenger as the most influential numerical feature. This research not only advances the field of predictive analytics but also provides practical insights for implementation in the aviation industry. As airlines continue to pursue digital transformation, this machine learning solution offers a strategic opportunity to deliver measurable economic benefits and enhance the in-flight retail experience through personalized and relevant product offerings.

Keywords: In-flight sales; Machine learning; Sales forecasting; Aviation industry; XGBoost; Feature importance; Predictive analytics; Ancillary revenue; Inventory management; Passenger preferences.

1. Introduction

In the dynamic realm of duty-free and travel retail, where global commerce intersects with the needs of international travelers, in-flight sales emerge as a pivotal segment. The duty-free industry, marked by tax exemptions on goods sold to travelers, offers a unique shopping experience for those crossing borders. In-flight retail serves not only as a way to increase airline revenue but also as a means to enhance customer satisfaction (Kucerova, 2021; Marisa Garacia, 2017). Passengers view shopping as an enjoyable experience, making it a more favorable spending option compared to baggage and change fees, which are often seen as burdensome (Marisa Garacia, 2017).

This research will address a central challenge within the duty-free landscape, the lack of effective prediction for in-flight sales item, leading to suboptimal loading practices and potential revenue loss for airlines (Eric Leopold, 2019; Kucerova, 2021). Ancillary revenue (additional services and products airlines offer beyond the base airfare) (OAG, 2023), statistics compiled by IdeaWorks and CarTrawler show that with a few exceptions in-flight sales mostly account for less than a quarter of airlines' total ancillary revenue (Marisa Garacia, 2017). In their most recent report, combined airline ancillary revenue for 2023 was projected at US$117.9bn (IdeaWorks, 2023). Deciding the exact number of in-flight items to load, especially for the fresh food is always a problem (Hast, 2019) and the crux of the problem lies in the absence of a reliable predictive model that can anticipate the items to be sold during flights. To bridge this gap, the study focuses on the development and application of machine learning algorithms to enhance the accuracy of in-flight sales prediction.

The significance of this research extends beyond operational improvements. By harnessing machine learning to predict and optimize in-flight sales, the study aims to contribute to a more personalized and enticing shopping experience for every traveler (Marisa Garacia, 2017). The outcomes of this research will hold importance for airlines, duty-free operators, and the broader aviation industry, fostering innovation, sustainability, and customer-centric practices. It is essential for effective inventory management, supply chain optimization, and revenue maximization (Jadhav, 2023) lowering cost of ticket and increase revenue by selling complementary In-flight items for passenger (Hast, 2019) page 19. Ultimately, this research is poised to shape the future of in-flight retail, ensuring a seamless and enriching journey for passengers worldwide.

1.1. Background of the Study

In-flight shopping is the commercial activity of purchasing products or services while on an airplane in-flight. It is a way for passengers to purchase items to enjoy during the flight or purchase items to be delivered later (*Global Inflight Shopping Market*, 2024). It allows passengers to purchase goods and services such as food, beverages, merchandise, and entertainment during the flight. This trend is rapidly gaining traction among customers across the globe (*Global Inflight Shopping Market*, 2024).

The demand for convenience and comfort during long-distance travel has driven the growth of the Global In-flight Shopping Market. Technological advancements, particularly the implementation of e-commerce platforms, have facilitated online purchases of goods and services during flights. Major players in the market, such as Lufthansa Systems, Airbus, and Honeywell, contribute to the expansion of In-flight shopping options (*Global Inflight Shopping Market*, 2024).

In the broader context of duty-free and travel retail, in-flight sales hold a unique position as a pivotal segment. Duty-free shopping, characterized by tax exemptions on goods sold to travelers, is intricately linked to international travel settings, including airports, sea terminals, cruise ships, and airline flights. The duty-free industry offers a distinctive shopping experience, emphasizing exclusivity and premium products (What Is Duty Free & Travel Retail? n.d.).

The allure of duty-free shopping extends beyond the practicality of tax exemptions. Advertisements for duty-free outlets emphasize a substantial price advantage, with claims of prices ranging from 10% to 50% lower than domestic prices. The duty-free experience is not merely transactional; it often involves the sale of upscale tourist items that serve as cultural ambassadors, showcasing the heritage of the host country (*Duty Free Facts*, 2022).

The In-flight shopping market, valued at USD 6.5 billion in 2022, is projected to grow significantly, reaching USD 10.48 billion by 2032 with a compound annual growth rate of 5.45%. Duty-free shops, offering premium items such as Liquor, Tobacco Products, Perfumes, Electronics, and Fashion, contribute to this growth by delivering exclusive products to passengers (Sejal Akre, 2024).

The conventional methods of setting sales and marketing objectives are becoming less effective in today's fast-paced and competitive market, as they often lack insights into customer purchasing behaviors. Significant changes have occurred in the sales and marketing sector due to advancements in machine learning. These technological advancements enable businesses to analyze critical elements like consumer buying habits, identify target audiences, and forecast future sales trends, ultimately empowering sales teams to devise strategies that enhance business growth (Bajaj et al., 2020).

In-flight sales are undergoing a transformation marked by the integration of new technologies, departing from traditional models adopted by an increasing number of airlines (Jason Holland, 2019; Kucerova, 2021). Marian Fagbemiro, Vice President at Gate retail, emphasizes the urgency for in-flight retail to adapt swiftly to changing shopping behaviors: "The way we shop and what we shop for has changed and continues to do so, and in-flight retail needs to act quickly to keep up" (Jason Holland, 2019). A noticeable trend in in-flight retail is the move towards carrying fewer products onboard, facilitated by apps and services enabling onboard payment but fulfilling orders on the ground (Jason Holland, 2019). Fagbemiro suggests that retailers must leverage technology and data to shift from a 'sense and respond' approach to a 'predict and act' strategy, involving anticipating consumer demands, understanding their preferences, and responding rapidly to stay ahead in the evolving market.

The research for in-flight sales prediction using machine learning becomes crucial and beyond enhancing the shopping experience, offers other benefits such as fuel saving through effective prediction of items to be sold and optimize loading of goods, heightened customer satisfaction, streamlined warehouse management, and improved supply chain arrangements (Mark Finlay, 2020).

1.2. Statement of the Problem

The absence of an effective in-flight sales prediction system is the root cause of operational inefficiencies, financial burdens, and diminished customer satisfaction within the airline industry (Kucerova, 2021). This deficiency leads to inaccuracies in estimating and loading items, resulting in inadequate provision for passenger preferences, unnecessary fuel consumption costs, and complications in warehouse management. Consequently, airlines face compromised revenue potential, heightened customer dissatisfaction, and a restricted ability to optimize in-flight sales strategy. Furthermore, the lack of predictive analytics hampers the ability to adapt to changing market dynamics and passenger preferences.

Despite the critical importance of optimizing in-flight sales, there is a notable lack of comprehensive predictive models tailored specifically for in-flight retail items. To date, there has been limited research conducted on predicting in-flight sales for items excluding meals (Young-Chan Lee, 2001). While there is one research study focused on in-flight meal sales prediction, it significantly differs from predicting sales of a wider range of in-flight retail items such as duty-free products, and beverages. These categories involve different consumer behavior patterns and logistical considerations. The existing research on meal sales (Young-Chan Lee, 2001) does not account for the varied nature of in-flight retail items, leaving a substantial gap in the literature.

The solution lies in the development of an in-flight sales predictive model for optimization, mitigating these challenges and fostering a more streamlined and customer-centric approach to in-flight retail. By addressing this research gap, the proposed model aims to enhance operational efficiency, reduce financial burdens, and improve customer satisfaction through precise and adaptive sales predictions, specifically tailored for the diverse range of in-flight retail items.

1.3. Research Question

RQ1: Which ML models perform better accuracy in in-flight sales prediction?

RQ2: Which attributes are relevant for in-flight sales prediction?

RQ3: Which hyperparameters optimize ML model performance for in-flight sales prediction?

1.4. Study Objectives

This study aims to address key challenges in predicting in-flight sales by utilizing machine learning models to enhance operational efficiency and customer satisfaction. The specific objectives of this study are:

- To identify key factors influencing in-flight sales predictions by reviewing existing literature and industry practices.
- To collect, organize, and preprocess relevant data needed for model development and analysis.
- To apply advanced data preprocessing techniques to ensure high-quality, reliable datasets by removing outliers and optimizing data integrity.
- To select and implement appropriate machine learning models for accurate prediction of in-flight sales quantities.
- To train and optimize the selected models using the pre-processed data to achieve high predictive accuracy.
- To evaluate and compare the performance of different machine learning models based on accuracy metrics such as RMSE, R^2, and MAE.

2. Literature Review

In this study, a literature review has been carried out to identify an appropriate prediction model for forecasting sales. This review involved examining past research that utilized various machine learning algorithms to determine their effectiveness and applicability for sales prediction. Research on customer demand prediction of in-flight meals (Hast, 2019), assesses the performance of four Machine Learning Algorithms (MLAs) – Linear Regression, Support Vector Regression, Extreme Gradient Boosting, and Multilayer Perceptron Neural Network – in predicting customer demand for in-flight meals using data from a single airline, it extracted new features called sold load level by dividing summation of sold food categories by total passenger which does not eat special meal categories. The training was done in 850, 887 records of data and 58 features. Despite minimal differences in prediction accuracy, Support Vector Regression underperforms in model fitting and prediction time. Notably, scheduled flight duration emerges as the most influential predictor for in-flight meal demand. The study recommends Extreme Gradient Boosting and Multilayer Perceptron Neural Network as the most suitable models, considering their superior accuracy, acceptable efficiency, and broad hyperparameter options.

Research on Big Mart sales prediction (Jadhav, 2023) explored many machine learning regression algorithms like Linear Regression, Decision Trees, Random Forests, Gradient Boosting, and Neural Networks to predict sales. The study leveraged a dataset containing numerous features related to product categories, store characteristics, and item attributes. Among the methods evaluated, Gradient Boosting and Random Forest performed particularly well, demonstrating robust predictive capabilities. However, the paper did not focus on the detailed comparison of the models' feature importance or provide a comprehensive evaluation of factors beyond sales data, such as seasonal variations or customer demographics, which could have enriched the findings. The study remains important for demonstrating the efficacy of ensemble methods in sales prediction, echoing findings from other works like (Premnath S P, 2022).

A study focusing on restaurant sales prediction (Nazmuz Sakib, 2023) evaluated the effectiveness of Support Vector Regression (SVR), Light Gradient Boosting Machine (LightGBM), Linear Regression, and Decision Tree models in forecasting sales. LightGBM outperformed all other models, achieving the lowest Mean Absolute Error (MAE). This success was attributed to the algorithm's efficiency in managing large datasets and its capability to capture complex, non-linear relationships within the data. However, one limitation noted in the study was the absence of customer feedback or satisfaction metrics in the dataset, which could have provided deeper insights into sales patterns and customer behavior. The finding aligns with other literature (Hast, 2019), (Premnath S P, 2022), which also highlight the advantages of gradient boosting techniques in sales forecasting.

The research in (Zhao & Keikhosrokiani, 2022) focused on sales prediction and product recommendation, utilizing XGBoost, Random Forest, Decision Tree algorithms, along with RFM analysis and the Apriori algorithm for product recommendation. The study found that XGBoost outperformed the other models, achieving an F1-score of 0.789, which marked it as the most effective approach for predicting sales. The use of the Apriori algorithm for product recommendation further enriched the study by providing valuable insights into customer purchasing behavior. While the research demonstrates strong performance from the XGBoost model, it does not offer detailed insights into feature importance or the effects of various data preprocessing techniques, which limits the scope of the findings. This result is consistent with other research (Hast, 2019), (Premnath S P, 2022), that showcases the superior performance of gradient boosting methods in predicting sales.

In (Ms. M. Anitha, 2023), the researchers emphasized the importance of AI-based sales forecasting for traditional retailers, comparing the performance of Linear Regression and Random Forest. The study highlights that machine learning-based sales predictions offer a significant advantage in improving inventory management and business strategy. While Random Forest was found to be more effective than Linear Regression in this context, the paper also recommended exploring more sophisticated time-series models like ARIMA for more accurate forecasting of long-term trends. The study's limitation lies in its lack of consideration for external variables such as promotions and discounts, which are critical in influencing sales but were not incorporated into the model. Similar gaps are identified in other works like (Hast, 2019), and (Jadhav, 2023), where external factors are also underexplored.

A sales prediction study conducted for Big Mart (Premnath S P, 2022) utilized Regression, Gradient Boosting, and Random Forest algorithms. The study revealed that Gradient Boosting achieved the highest accuracy (0.69) and the lowest Root Mean Squared Error (RMSE) of 10.343. The researchers analyzed various factors influencing sales using the dataset, but the study's time frame limited its ability to capture long-term sales trends or seasonality. Despite these constraints, the research aligns with findings from other studies (Hast, 2019), (Jadhav, 2023; Zhao & Keikhosrokiani, 2022), reinforcing the effectiveness of gradient boosting techniques for sales prediction tasks.

In general, gradient boosting is selected as the best algorithm across multiple research papers (Hast, 2019; Jadhav, 2023; Nazmuz Sakib, 2023; Zhao & Keikhosrokiani, 2022). These papers consistently demonstrate its superior performance in sales prediction tasks, showcasing its ability to handle complex relationships and provide accurate forecasts.

3. Research Methodology

The research methodology for predicting in-flight sales adopts a structured approach, utilizing machine learning algorithms and data-driven insights. This section details the entire process, encompassing data collection, preprocessing, feature selection, model selection, training, and evaluation. Accurate sales prediction is grounded in the effective utilization of historical sales data (Jason Holland, 2019). By applying a range of machine learning algorithms and identifying the most effective models, this research aims to substantially improve the accuracy of in-flight sales predictions.

3.1. Dataset overview

Before delving into data preprocessing, it's essential to discuss the dataset utilized in this study. The original dataset contains over 6,244,000 records of in-flight sales transactions from Anonymous Airlines, providing detailed information on various aspects of each sale and the corresponding flight details. The data spans multiple flights, dates, and routes, offering a comprehensive view of sales activities in the airline industry. It has in-flight sales history from March 2022 to May 2024. Below is a detailed description of each column included in the dataset:

Table 1. Dataset columns

Column	Description	Type	Example
FlightNumber	The unique identifier for a particular flight. Each flight number is associated with a specific route and date.	Alphanumeric	AC0123
FlightDate	The date on which the flight is scheduled or occurred. This helps in tracking sales across different days.	Date	04/12/2023
Routing	The flight path or route taken by the aircraft, typically represented as a combination of the origin and destination airport codes.	Alphanumeric	ABC-BCD-ABC
Origin	The airport code where the flight originated. This helps in identifying the starting point of the flight.	Alphanumeric	ABC
Destination	The airport code where the flight is destined to land. This helps in identifying the endpoint of the flight.	Alphanumeric	BCD
ItemCategory	The category of the item sold during the flight, such as beverages, fragrance, or merchandise.	String	Fragrance
ItemCode	The unique code identifying a specific item. This allows for the tracking and analysis of sales by	Alphanumeric	FR021

	item.		
ItemName	The name or description of the item sold, providing a more detailed understanding of what was purchased.	String	YSL Black Opium Eau de Parfum
Quantity	The number of units of the item sold, indicating the volume of sales for each item.	Integer	2
Price	The price per unit of the item sold.	Decimal	11.99
TotalPrice	The total price for the quantity of items sold (calculated as Quantity * Price). This gives the total revenue from each sale.	Decimal	23.98
Ecocomy	The number of economy class passengers on the flight. This can be used to analyze sales per passenger class.	Integer	182
Business	The number of business class passengers on the flight. This helps in understanding the distribution of passengers by class.	Integer	18
Infant	The number of infant passengers on the flight. Infants typically do not purchase items, but this data can help in demographic analysis.	Integer	5
TotalPassenger	The total number of passengers on the flight, including economy, business, and infants. This is critical for analyzing sales per passenger.	Integer	205
FlightNature	Indicates whether the flight is an outbound (departure), inbound (return) flight, or unknown which can affect sales patterns.	String	Departure

3.2. Data Preprocessing

Data preprocessing involves preparing the raw dataset for machine learning algorithms (Hast, 2019). Effective data preprocessing significantly influences the predictive performance of machine learning models. This crucial step involves a series of operations aimed at eliminating noise and generating clean data suitable for model training. The primary focus during this stage has been on achieving a balanced dataset to mitigate skewness, performing data transformations, and removing irrelevant data, such as rows containing null values (Gebremedhin et al., 2022) page 29.

In our research, we employ a series of preprocessing techniques to prepare the data for effective machine learning modeling. This includes outlier handling to remove extreme values that could skew the results, and data cleaning to address missing or inconsistent values. We also perform feature extraction to derive new, meaningful features that enhance the model's predictive power. Categorical encoding transforms categorical variables into numerical formats that machine learning algorithms can process (SP774RVUQ, 2024). Additionally, feature scaling is applied to normalize the range of numerical values (Ahmed Ouameur et al., 2020), ensuring that each feature contributes equally to the model. These preprocessing steps are crucial for enhancing data quality and boosting the accuracy and robustness of machine learning models. The specifics of these processes will be elaborated upon in the subsequent sections

3.2.1. Handling outlier

Outliers are values that deviate significantly from the overall range or trend of the dataset, often manifesting as extremely large or small values (Gebremedhin et al., 2022) page 29. In this research, the identification of outliers involves a two-step process: data filtering and subsequent visual inspection methods.

Initially, data filtering is employed to pinpoint values that deviate substantially from the general data distribution (Cousineau & Chartier, 2010). In this research, a visual inspection method is applied to further scrutinize and confirm the presence of outliers. Records associated with values deemed excessively large are then excluded from the analysis, ensuring a more accurate representation of the dataset in subsequent research and analysis (Cousineau & Chartier, 2010).

In this study, out of 6 million records, the following outliers were identified and excluded:

- 132,800 records show a total passenger count greater than 393, which is unrealistic as no aircraft can carry more than 393 passengers.
- 430,801 records with a total passenger count of less than 51. According to the dataset description, the minimum value of the 25th percentile is 112 passengers. Records with fewer passengers typically result in no sales and are therefore not useful for our analysis.
- 17 records have a total sales quantity per flight greater than 30, the maximum is 61.
- 475,000 records show business class passengers greater than 28 per flight, whereas the maximum business class capacity is 28.
- 236,202 records indicate more than 6 infants per flight, as seen in dataset overview table most of the record has non to 2 infants, with maximum 67 which deviates, and looks unrealistic.
- 105,900 records have flight dates after May 1, 2024, although the complete data is only available until May 1, 2024.

3.2.2. Data cleansing

After removing outliers, we cleaned the data by removing unnecessary columns and converting certain feature data types to formats more suitable for machine learning. In our research features such as Item Name and Total price are removed at first, as item name is not required as long as item Code is available and total price is multiplication of quantity of sold item and price of each item, this in turn, helps mitigate overfitting by eliminating data redundancy, reducing training time, and improve model accuracy (Vineeth, 2019).

In addition, to ensure the dataset is in an optimal format for machine learning algorithms, several columns were converted to appropriate data types. The Quantity column was converted to integer to accurately represent the number of items sold as discrete values. The Price column were converted to float. Similarly, the Economy, Bussiness, and Infant columns were converted to integer to denote the count of economy class passengers, business class passengers, and infant passengers as discrete values. Finally, the TotalPassenger column was also converted to integer to capture the total number of passengers as discrete values. These conversions ensure that our data is well-structured and appropriate for machine learning models, thereby enhancing both processing efficiency and prediction accuracy.

3.3. Feature engineering

After data cleaning, one of the primary works in prediction is to extract valuable features from a variety of datasets (Gebremedhin et al., 2022). As feature selection is a machine learning process that involves identifying and choosing the most relevant features from your data that contribute significantly to the prediction output of interest (Nikolaj Buhl, 2023).

In our research, several feature engineering techniques were also applied to preprocess the dataset effectively for predicting the quantity of items loaded on flights. These techniques include creating new features from existing ones and applying transformations to capture the cyclic nature of certain features.

Creating New Features: New features were derived from existing ones to provide additional information to the model. For example, from FlightDate column, new features such as Year, Month, Day, DayOfWeek, and WeekOfYear were extracted. These features capture temporal aspects of the data that may influence the quantity of items loaded.

Cyclic Feature Transformation: Certain features, such as months and days, have a cyclic nature (e.g., January and December are close in time). To capture this cyclic nature, sine and cosine transformations were applied. These transformations convert the cyclic features into two components, effectively capturing their cyclical behavior. This approach was applied to the Month and Day features.

3.4. Correlation Analysis

After completing feature engineering, we conducted correlation analysis. Correlation describes the statistical relationship between features and can indicate mutual dependencies among them or with the chosen target variable (Vineeth, 2019) (quantity of sold item in our case). A correlation can be negative, indicating that as one feature's value rises, the value of another feature falls. Conversely, a positive correlation signifies that as one feature's value increases, the value of another feature also increases (Vineeth, 2019).

Several methods can be used to calculate the correlation coefficient. The Pearson correlation coefficient evaluates the linear relationship between continuous variables, whereas the Spearman correlation coefficient examines the association between variables based on a monotonic function (Vineeth, 2019). In this study, the Spearman correlation coefficient has been preferred, since we want to determine the target based on other features. In this section, we examined how the target variable Quantity correlates with other features in the dataset. This analysis helps us understand which features have the strongest relationships with Quantity, which is crucial for predicting this target variable.

3.5. Numerical Feature Scaling

Scaling is a vital preprocessing step in machine learning that ensures all numerical features in a dataset have a consistent range and distribution (Ajaykumar Dev, 2024). This is particularly important especially for algorithms like neural networks that are sensitive to input data scaling. In this research, MinMaxScaler was used to scale the numerical features such as `Price`, `Economy`, `Bussiness`, `Infant`, `TotalPassenger`, `Month_sin`, `Month_cos`, `Day_sin`, `Day_cos`, 'DayOfWeek', 'WeekOfYear', and 'Year' to a specified range, between 0 and 1. This process enhances the performance and convergence speed of machine learning models by normalizing features, ensuring that no single feature with a larger scale disproportionately influences the model (Ajaykumar Dev, 2024).

The key benefits of scaling include:

1. Improved Model Performance: Scaling helps algorithms converge faster and achieve better performance. In this study, MinMaxScaler was applied to numerical columns to ensure they fall within the same range (Ajaykumar Dev, 2024).

2. Enhanced Model Training: Scaling aids in stabilizing the training process, especially for neural networks. This is because unscaled data can lead to issues where the model might have difficulty learning effectively, as the gradients can become excessively large or small, resulting in slow learning or even causing the model to fail (Ajaykumar Dev, 2024).

3. Fair Feature Contribution: By scaling the features, each one contributes equally to the distance calculations in algorithms that rely on distance metrics, such as neural networks. This ensures that no single feature can dominate the model training due to its larger magnitude.

This method improved the model's capacity to predict the target variable, Quantity, accurately, while keeping the target variable unscaled.

3.6. Categorical Feature Encoding

In this study, categorical features were transformed into numerical values using one-hot encoding, making them compatible with machine learning algorithms. This encoding process was performed on features such as FlightNature, FlightNumber, Routing, Origin, Destination, ItemCategory, and ItemCode. The one-hot encoding was implemented with the pd.get_dummies function and the drop_first=True parameter. This approach effectively avoids multicollinearity by removing the first category of each feature, thus creating binary columns that retain the necessary information for model training.

One of the most widely used methods for encoding categorical features is one-hot encoding. This approach involves comparing each level of a categorical variable to a fixed reference level. It creates a binary column for each category, converting the categorical data into a dense array of binary values. In simpler terms, it binarizes each category and includes these binary columns as features in the dataset (Vineeth, 2019).

3.7. Final Dataset Preparation

In our research, to prepare the final dataset after pre-processing, we aimed to ensure that the dataset used for modelling was complete and free of missing data by attempting to remove any rows containing NaN values (Saturn

Irish Interdisciplinary Journal of Science & Research (IIJSR)
Volume 8, Issue 3, Pages 39-55, July-September 2024

Cloud, 2023), even though none were found. Despite our thorough pre-processing steps and the minimal presence of null values, we applied the `dropna` function to the DataFrame as an extra precaution. This step ensured that any potential rows with missing data were removed, thus preventing issues during analysis, and improving the reliability and accuracy of our results. The final dataset comprises 4,985,885 records and 167 features.

3.8. Train-Test Split

In this research, the dataset was divided into training and testing sets to assess the performance of the models. To simulate a real-world scenario where predictions are made on future data based on historical trends, the latest 20% of the data was reserved for testing (Idrees et al., 2019).

To implement this temporal split method, we first converted the `FlightDate` column to DateTime format and sorted the dataset by flight date. We then defined the split index to separate the most recent 20% of the data for testing and the older 80% for training. After that, features and the target variable were defined, with X containing all columns except Quantity and FlightDate, and y representing the target variable Quantity. The data was divided into training and testing sets by determining the split index according to a specified test size of 20%. The split index was determined by multiplying the total length of X by 0.8. The training set consisted of data from the beginning up to this index, while the testing set comprised the remaining data. This approach ensures that the model is trained on historical data and evaluated on the most recent observations, aligning with the temporal nature of real-world predictions.

3.9. The Proposed Models

After splitting our dataset, we must have a model to train and test our results. Model selection is a critical step in developing a predictive analytics framework, involving the choice of the most suitable machine learning algorithms based on their characteristics and performance. The selection of models significantly impacts the accuracy and generalization capability of the predictive model (Manika, 2023). In this study, we employed Linear Regression (LR), Support Vector Regression (SVR), Extreme Gradient Boosting (XGBoost), Multilayer Perceptron Neural Network (MLP), Recurrent Neural Network (RNN), and Deep Neural Network (DNN) models after summarizing and determining which models are good for our research based on a review of related papers.

These models were chosen for their proven effectiveness in handling various types of regression tasks and their ability to capture different aspects of the data. Linear Regression offers simplicity and interpretability, making it a useful baseline model. Support Vector Regression is known for its ability to model complex relationships with a focus on minimizing prediction errors. XGBoost is a powerful ensemble method that has shown superior performance in many machine learning competitions, particularly for structured data. MLP offers flexibility in capturing non-linear patterns, making it well-suited for more complex datasets.

Additionally, RNN was included due to its strength in modeling sequential data, which can be advantageous when there are dependencies across time or order in the dataset. DNN was chosen for its deep learning architecture, which allows for capturing intricate patterns in data, offering high predictive accuracy. By evaluating and comparing the performance of these models, we aim to identify the one that best suits our specific problem and data characteristics, ensuring that our predictive analytics framework is both robust and effective.

3.10. Model Training

Model training is a crucial stage in the data science development lifecycle where practitioners work to optimize the weights and biases of a machine learning algorithm in order to minimize a loss function across predictions. This process aims to create the most accurate mathematical model that captures the relationship between data features and a target label in supervised learning, or the relationships among features in unsupervised learning (*Model Training*, 2024).

Model training for Linear Regression (LR) involves optimizing coefficients using methods like Ordinary Least Squares (OLS) or gradient descent (Manish Kumar, 2023). Support Vector Regression (SVR) aims to find a hyperplane minimizing errors through a cost function (Alakh, 2024). Extreme Gradient Boosting (XGBoost) sequentially trains weak learners, optimizing a regularized objective function (guest_blog, 2024). Multilayer Perceptron Neural Network (MLP) employs backpropagation and stochastic gradient descent to iteratively adjust weights (Saarathi Anbuazhagan, 2021). These tailored approaches ensure effective learning from diverse perspectives, contributing to a robust predictive model for in-flight sales data.

In our research, model training process involved applying different distinct machine learning algorithms: Linear Regression (LR), Multilayer Perceptron (MLP), Extreme Gradient Boosting (XGBoost), Support Vector Regression (SVR), Recurrent Neural Network (RNN), and Deep Neural Network (DNN) selected in proposed model. Each model was initialized using a set of optimized hyperparameters, which is done previously, referred to as `best_params`, to enhance performance and ensure the best possible predictions.

➢ Linear Regression (LR): The LR model was configured with standard settings but adjusted using the optimal parameters derived from previous experimentation. This approach ensures that the model can effectively capture linear relationships between the features and the target variable.

➢ Multilayer Perceptron (MLP): For the MLP model, hyperparameters such as the number of hidden layers, activation functions, and learning rates were fine-tuned based on the `best_params`. This configuration allows the neural network to capture complex, non-linear patterns in the data.

➢ Extreme Gradient Boosting (XGBoost): The XGBoost model was set up with the best parameters for `max_depth`, `n_estimators`, and `subsample`. This configuration optimizes the model's performance by balancing complexity and generalization, leveraging XGBoost’s gradient boosting framework to improve predictive accuracy.

➢ Support Vector Regression (SVR): The SVR model was trained with optimal settings for the kernel type and regularization parameters. This configuration helps the model to efficiently capture both linear and non-linear relationships in the dataset, enhancing its predictive capabilities.

➢ Recurrent Neural Network (RNN): The RNN model was configured with optimal settings for the number of layers, hidden units, and learning rate. Specific attention was given to the sequence length and batch size to handle the temporal dependencies in the dataset. By capturing sequential patterns in the data, the RNN model can effectively predict outcomes where the order of data points is crucial.

IIJSR PROGRESS THROUGH RESEARCH
Irish Interdisciplinary Journal of Science & Research (IIJSR)
Volume 8, Issue 3, Pages 39-55, July-September 2024

➢ Deep Neural Network (DNN): The DNN model was optimized by fine-tuning hyperparameters such as the number of layers, neurons per layer, activation functions, and dropout rates. This deep architecture allows the model to learn multiple levels of abstraction, capturing intricate and complex patterns in the data. The use of dropout helps in preventing overfitting, ensuring the model generalizes well to unseen data.

Each model was trained on the pre-processed dataset, with a focus on ensuring robust performance and generalization to future data.

4. Result

This study aimed to forecast in-flight sales by applying machine learning models. Our key objectives were to identify the most accurate models, determine the most relevant features for prediction, and optimize model hyperparameters. Below is a summary of the model performance results.

Table 2. Model performance result

Model	**RMSE**	**R^2**	**MAE**
LR	0.3382564828	0.4202785618	0.0641330699
MLP	0.6673749235	0.7263696384	0.2846567912
XGBoost	0.1945461364	0.8082337379	0.0172848906
SVR	1.2403041264	-0.0388065325	0.4110197360
RNN	0.2701371887	0.8070925474	0. 0608975738
DNN	0.3231852963	0.8145098686	0.0689975843

Below, we discuss the findings in relation to each research question posed.

RQ1: Which ML models perform better accuracy in in-flight sales prediction?

Through our analysis, we evaluated the performance of six machine learning models: Linear Regression, Multi-Layer Perceptron (MLP), Extreme Gradient Boosting (XGBoost), Support Vector Regression (SVR), Recurrent Neural Network (RNN), and Deep Neural Network (DNN). The performance metrics—Root Mean Squared Error (RMSE), R-squared (R^2) score, and Mean Absolute Error (MAE)—revealed that XGBoost and DNN outperformed the other models.

RQ2: Which attributes are relevant for in-flight sales prediction?

Feature importance analysis and correlation results reveal several key insights into the prediction of in-flight sales. ItemCategory_TOBACCO emerged as the most influential factor, with a strong positive correlation of 0.369446 with the quantity of loaded items, while TotalPassenger also identified as the most significant from numeric features, showing a correlation of 0.099722 with the quantity of loaded items. This suggests that while the total number of passengers has some effect on sales, it is less influential compared to other factors like ItemCategory_TOBACCO.

These results underscore the importance of pricing and specific item categories in predicting in-flight sales, while also acknowledging that other numeric features have a relatively minor role in influencing sales quantities.

RQ3: Which hyperparameters optimize ML model performance for in-flight sales prediction?

Hyperparameter tuning was crucial for enhancing the performance of the machine learning models. To identify the optimal hyperparameters for each model, we utilized GridSearchCV, which systematically explored different configurations to find the best-performing set of parameters.

5. Conclusion

The objective of this study was to develop and evaluate machine learning models for predicting in-flight sales, with the aim of enhancing revenue and customer satisfaction within the airline industry. We assessed six regression models—Linear Regression (LR), Multi-Layer Perceptron (MLP), Extreme Gradient Boosting (XGBoost), Support Vector Regression (SVR), Recurrent Neural Network (RNN), and Deep Neural Network (DNN)—to determine their effectiveness in forecasting the quantity of loaded items. The performance metrics revealed that XGBoost and DNN were the top performers in terms of prediction accuracy. Specifically, XGBoost achieved a Root Mean Squared Error (RMSE) of 0.1945, an R^2 score of 0.8082, and a Mean Absolute Error (MAE) of 0.0173, while DNN yielded a RMSE of 0.2701, an R^2 score of 0.8071, and an MAE of 0.0609. These results highlight that both models provide high accuracy in predicting in-flight sales. The LR model, with its more straightforward approach, achieved decent performance but was less accurate compared to XGBoost and MLP. The SVR model, on the other hand, showed less favorable results with an RMSE of 1.2403, an R^2 score of -0.0389, and an MAE of 0.4110, indicating that it was less effective in capturing the underlying patterns in the data. On the other hand, feature importance analysis revealed that ItemCategory_TOBACCO was the most significant feature among all features. The study applied rigorous data preprocessing methods, including feature scaling and encoding, to ensure robust model performance. The combination of training and testing data for visualization purposes illustrated how well the models predicted the quantities over time. This study establishes a basis for further exploration and advancement in predictive modeling for in-flight sales. It offers valuable insights into model selection and feature importance, contributing to improved predictive accuracy within the airline industry.

6. Future Suggestions

In light of the findings and limitations identified in this research, several avenues for future work are recommended.

1. Advanced Modeling Techniques: Exploring deep learning models or hybrid architectures could further improve performance, especially given the large dataset of over 5 million records. Neural networks, combined with traditional machine learning approaches, could provide enhanced predictive accuracy.

2. Feature Expansion and Engineering: Incorporating additional features, such as flight duration, seasonal variations, or economic indicators, could provide a richer context for the model. Advanced feature selection techniques would help identify and integrate the most relevant predictors, improving the model's accuracy and reducing overfitting.

3. Data Enrichment: Expanding the dataset to include a wider time range and data from more diverse flight routes and airlines would improve model robustness and generalizability. In addition, external data sources like customer demographics or historical sales trends could add valuable insights.

4. Real-time Data and Operational Integration: Integrating predictive models into airline operations, such as inventory management and real-time decision-making systems, is a crucial next step. Developing APIs and interfaces for seamless data flow between models and operational systems would enhance their practical applicability.

5. Scalability and Generalization: Future work should focus on developing scalable models that can handle not only in-flight sales but also inventory management across various airline operations, ensuring that the solution can be generalized to different contexts or operational environments.

Declarations

Source of Funding

This research did not receive any grant from funding agencies in the public, commercial, or not-for-profit sectors.

Competing Interests Statement

The authors declare no competing financial, professional, or personal interests.

Consent for Publication

The authors declare that they consented to the publication of this research work.

Authors' contributions

All the authors took part in literature review, analysis and manuscript writing equally.

References

Ahmed Ouameur, M., Caza-Szoka, M., & Massicotte, D. (2020). Machine learning enabled tools and methods for indoor localization using low power wireless network. Internet of Things, 12: 100300. https://doi.org/10.1016/j.iot.2020.100300.

Ajaykumar, D. (2024). Data Science: Understanding Feature Scaling in Machine Learning. Medium. https://medium.com/@nikaljeajay36/data-science-understanding-feature-scaling-in-machine-learning-6b290f76668f.

Alakh (2024). Support Vector Regression in Machine Learning. Analyticsvidhya. https://www.analyticsvidhya.com/blog/2020/03/support-vector-regression-tutorial-for-machine-learning/.

Bajaj, P., Ray, R., Shedge, S., Vidhate, S., & Shardoor, N. (2020). Sales Prediction Using Machine Learning Algorithms. International Research Journal of Engineering and Technology.

Cousineau, D., & Chartier, S. (2010). Outliers detection and treatment: a review. International Journal of Psychological Research, 3(1): 58–67. https://doi.org/10.21500/20112084.844.

Duty Free Facts (2022). Duty Free Facts. https://dutyfreefacts.com/duty-free-facts/.

Eric, L. (2019). Airline Industry Retailing (AIR) Think Tank. https://www.iata.org/contentassets/2d997082f3c84c7cba001f506edd2c2e/air_tt_whitepaper_2019.pdf.

Gebremedhin, S., Lemma, S., & Ababa, A. (2022). Forecasting Ethiopian Agricultural Commodity Price Using Time Series Features and Technical Indicators. https://etd.aau.edu.ethandle/123456789/31732?show=full.

Global Inflight Shopping Market (2024). Verified Market Report. https://www.verifiedmarketreports.com/product/inflight-shopping-market/.

Guest_blog (2024). What is the XGBoost algorithm and how does it work?. Analyticsvidhya. https://www.analyticsvidhya.com/blog/2018/09/an-end-to-end-guide-to-understand-the-math-behind-xgboost/.

Hast, M. (2019). Evaluation of machine learning algorithms for customer demand prediction of in-flight meals. In Degree Project Computer Science and Engineering. https://kth.diva-portal.org/smash/get/diva2:1337269/fulltext01.pdf.

IdeaWorks (2023). Airline Ancillary Revenue Reaches Record $117.9 Billion Worldwide for 2023. https://ideaworkscompany.com/.

Idrees, S.M., Alam, M.A., Agarwal, P., & Ansari, L. (2019). Effective Predictive Analytics and Modeling Based on Historical Data. Pages 552–564. https://doi.org/10.1007/978-981-13-9942-8_52.

Jadhav, A. (2023). Machine Learning for Sales Prediction in Big Mart. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4530624.

Jason, H. (2019). The future of inflight retail. The Moodie Davitt EZine. https://ezine.moodiedavittreport.com/ezine-265/inflight-retail-report-gateretail/.

Kucerova, V. (2021). Unlock your cabin crew's full potential and boost your in-flight sales. https://www.sita.aero/globalassets/docs/brochures/crewtab-cabin-digitalization-brochure.pdf.

Manika (2023). Your 101 Guide to Model Selection In Machine Learning. ProjectPro. https://www.projectpro.io/article/model-selection-in-machine-learning/824.

Manish, K. (2023). Understanding Linear Regression Optimization. Medium. https://medium.com/@kmrmanish/understanding-linear-regression-optimization-gradient-descent-and-ols-approaches-ece8048f7ccb.

Marisa Garacia, R.H. (2017). Inflight Retail. https://static1.squarespace.com/static/59f9aaa6914e6bf45841fd9e/t/5a2006f2ec212d3c970a3257/1512048388582/Onboard-Retail.pdf.

Mark, F. (2020). How Do Onboard Duty Free Sales Work?. Simple Flying. https://simpleflying.com/duty-free-sales/.

Model Training (2024). C3.Ai. https://c3.ai/glossary/data-science/model-training/.

Anitha, M., et al. (2023). Sales Prediction Using Machine Learning Techniques. International Research Journal of Modernization in Engineering Technology and Science. https://doi.org/10.56726/irjmets43135.

Nazmuz Sakib, S.M. (2023). Restaurant Sales Prediction Using Machine Learning. https://doi.org/10.4018/978-1-6684-7105-0.ch011.

Nikolaj, B. (2023). Data Cleaning & Data Preprocessing for Machine Learning. ENCORD. https://encord.com/blog/data-cleaning-data-preprocessing/.

OAG (2023). Shaping Airline Retail: The Unstoppable Rise Of Ancillaries. OAG. https://www.oag.com/blog/shaping-airline-retail-unstoppable-rise-ancillaries.

Premnath, S.P., et al. (2022). Sales Prediction Using Machine Learning. http://ymerdigital.com.

Saarathi, A. (2021). A Complete Guide to train Multi-Layered Perceptron Neural Networks. Medium. https://paarthasaarathi.medium.com/a-complete-guide-to-train-multi-layered-perceptron-neural-networks-3fd8145f9498.

Saturn Cloud (2023). Python Pandas: How to remove nan and inf values. Saturn Cloud Blog. https://saturncloud.io/blog/python-pandas-how-to-remove-nan-and-inf-values/.

Sejal, A. (2024). Global Inflight Shopping Market Overview. Market Research Future. https://www.marketresearchfuture.com/reports/inflight-shopping-market-8272.

SP774RVUQ (2024). Encoding Categorical Data in Sklearn. GeeksforGeeks. https://www.geeksforgeeks.org/encoding-categorical-data-in-sklearn/.

Vineeth, V.S. (2019). Machine Learning Approach for Forecasting the Sales of Truck Components. www.bth.se.

What is Duty Free & Travel Retail? (n.d.). Duty Free World Council. Retrieved January 14, 2024, from https://dfworldcouncil.com/what-is-duty-free-travel-retail/.

Young-Chan, L. (2001). Airline In-flight Meal Demand Forecasting with Neural Networks and Time Series Models. The Journal of Information System, 10(2).

Zhao, X., & Keikhosrokiani, P. (2022). Sales prediction and product recommendation model through user behavior analytics. Computers, Materials and Continua, 70(2): 3855–3874. https://doi.org/10.32604/cmc.2022.019750.