Flight Load Factor Predictions based on Analysis of Ticket Prices and other Factors

Miroslav Spak
EUROCONTROL, Training Developement and Delivery Unit
Rue de la Fusée 96, 1130, Bruxelles, Belgium
miroslav.spak@eurocontrol.int

Lorenzo Frigerio
Euranova
Rue Emile Francqui 4, 1435 Mont-Saint-Guibert, Belgium
lorenzo.frigerio@euranova.eu

Lenka Hanakova, Vladimir Socha
Czech Technical University in Prague, Faculty of Transportation Sciences, Department of Air Transport
Horska 2040/3, 12803 Prague, Czech Republic
hanakle1@cvut.cz, sochavla@cvut.cz

Rocio Barragan Montes, Vincent Treve
EUROCONTROL, Airports Unit
Rue de la Fusée 96, 1130 Bruxelles, Belgium
rocio.barragan-montes@eurocontrol.int, vincent.treve@eurocontrol.int

***Abstract*—The ability to forecast traffic and to size the operation accordingly is a determining factor, for airports. However, to realize its full potential, it needs to be considered as part of a holistic approach, closely linked to airport planning and operations. To ensure airport resources are used efficiently, accurate information about passenger numbers and their effects on the operation is essential. Therefore, this study explores machine learning capabilities enabling predictions of aircraft load factors. The rationale behind the logic used stems from the assumption that using past traffic statistics in a form of historic load factor may not be sufficient, especially at times of high traffic volatility such as during regional bank holidays. Therefore, exploration efforts were made to parameterize some novel predictive elements that could provide passenger demand predictions at different granularity levels. The investigation has been successful and through the use of gradient boosting technique, the model, including 9 significant predictors was created. The load factor predictions per flight perform highly accurately with an average mean absolute error around 10 percentage points. In principle, this achievement outscores any other related work conducted in this domain to date. On top of that, the model itself is scalable and can be applied to any airport in the network as applied to use cases within the presented paper.**

***Keywords*—operational planning; passenger demand; ticket price; price ratio; machine learning; gradient boosting; airport forecasting; load factor prediction.**

I. Introduction

Enhancing the ability to forecast load factors holds significant advantages for airlines and airports, as extensively analyzed by [1]. These benefits encompass improved resource management, enhanced passenger experiences, strengthened airport operational resilience, and heightened profitability. However, as advised by [2], aircraft load factors are crucial for landside resource allocation, including check-in counters, security lanes, and immigration counters, as emphasized by [3]. Inadequate resources can lead to process delays, affecting all stakeholders and resulting in financial losses and operational disruptions, as explained by [4]. Consequently, visibility into aircraft load factors becomes fundamental during the planning phase.

Additionally, load factors serve as a key metric [5] that investors scrutinize when evaluating airlines, signifying effective management and a robust customer base. However, airlines are typically reluctant to disclose planned, operational, or predicted load factors to third parties. Pricing policies, which often fluctuate seasonally, play a significant role. [6] elaborates on the multifaceted factors considered, underscoring the paramount importance of the strategies employed by airlines to safeguard these figures.

In summary, research on flight load factor prediction remains an active area of study in aviation management and operations research. Passenger load factor serves as a measure of an airline's passenger carrying capacity, prompting various studies on passenger demand forecasting and load factor prediction. However, to start with, one of them published by [7] served to make a good review of existing techniques and methodologies developed to provide such predictions, and as such they concluded that there is a lack of material developed in the field. The truth is that most of the load factor forecasts come from airline internal analysis. In spite of the latter, there is a number of journal articles and conference papers published that use statistical and machine learning techniques to provide load factor forecasting.

To start with, [8] conducted an analysis that delved into historical passenger load patterns. Utilizing a decision tree, a predictive model was formulated to forecast passenger loads based on specific criteria. [9] also focuses on passenger demand forecasting based on historical passenger flow data and macroeconomic indicators. Their work is based on time

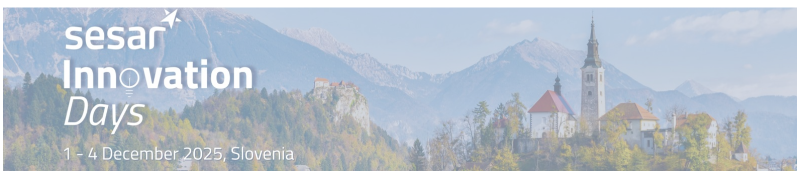

series prediction; however, it does not address forecasting at a granular level. Then, [10] addresses the estimation of the volume of passengers arriving at the airport with the objective of modeling the arrival time of passengers. Again, their work does not reach a flight based level granularity. Finally, [11] analyses both the per flight load and the airport passenger flows reaching a flight based level granularity. This constitutes the main reference for comparison for the performance of our work. Recognizing the continually evolving nature of this research, we subsequently identified flight ticket prices as a promising avenue to gain a more comprehensive understanding of how load factors evolve during a flight.

Airlines aim to optimize ticket prices to maximize revenue finding the right balance between prices and load factor. To achieve this, they employ sophisticated revenue management systems to analyze market demand and adjust ticket prices accordingly, as discussed by [12]. For instance, during peak travel seasons, such as holidays and summer vacations, airlines may charge higher prices due to increased demand. Another factor influencing ticket price oscillations is fuel costs, a significant expense for airlines. When fuel prices rise, airlines may increase ticket prices to offset these additional costs. Conversely, there are factors that may lower ticket prices. For example, competition on a particular route can lead airlines to lower prices to remain competitive and attract more customers, as highlighted by [13]. Similarly, less convenient flight times, such as overnight or early morning flights, are often priced lower than more desirable options. Booking in advance can result in cost savings, while last-minute bookings may lead to higher prices. All these factors interact in complex ways, making it challenging to predict precisely how ticket prices will change over time.

Furthermore, it is essential to recognize that different airlines may prioritize various factors when determining their prices. Therefore, it is advisable to compare prices across multiple airlines before making a purchase, as elaborated by [14]. A ticket's price is a crucial element that airlines regularly manage and adjust, considering customer purchasing power in relation to flight occupancies. For instance, [15] analyzed pricing strategies among European airlines and noted the prevalent use of dynamic pricing strategies that respond to changes in demand and competition.

Given the significance of ticket prices, we decided to explore their potential relevance for load factor predictions. Generally, lower ticket prices tend to boost demand, while higher prices tend to dampen it. However, this observation is not always accurate due to airlines' frequent use of dynamic pricing strategies, as discussed by [1], allowing for price adjustments based on factors like demand, competition, and seasonality. Moreover, as the number of tickets sold per flight increases, the price for remaining seats often rises, maximizing profits on high-demand routes.

Ticket prices can serve as a key indicator of demand for a specific flight. This ticket price analysis, combined with factors such as airline type, airport, and region, contributed to the development of a machine learning model capable of predicting flight load factors at selected airports over a specified time frame. These predictions were subsequently validated against real load factors on those routes to assess the model's performance. The validation was performed on data from five European airports from five different European regions. Notably, a distinctive feature of this work is the introduction of a novel synthetic variable known as the ”price ratio”, which feeds the machine learning model and enhances its accuracy in predicting future flight load factor.

This work aims at evaluating the effectiveness of a passenger demand predictive machine learning model based on ticket price information. The paper is organized as follows: In section 2, we explain the methodology followed to validate the relation between ticket prices and demand, and the architecture of the predictive model. In section 3, we present the results focusing on the model performance, and finally, in section 4, we discuss the results positioning our work in the existing state of the art.

II. Methods

Predicting flight load factors is a complex process that involves analyzing a wide range of data sources and factors. Machine learning models can be trained on historical data that includes information on aforementioned factors. Stemming from the literature review, there are many different machine learning algorithms that can be used for this task, including decision trees, random forests, and neural networks. The choice of algorithm depends on the specific characteristics of the data. In order to build an accurate prediction model, it is important to have reliable data. In this section, we explain the methodology used to prove the effectiveness of our predictive model and the impact of ticket prices. At first, we present the dataset used for the analysis; then, we introduce our load factor predictive model and how it can be further extended to predict the number of passengers. Finally, we explain the metrics that will be used to validate the results.

A. Dataset

For the purpose of this research, five European airports from five different European regions provided real historical load factor data. Data was collected for 72 818 flights between the 1st of January 2025 and the 29th of June 2025. Some days are missing due to data incompleteness for certain airports. For privacy considerations, as the load factors constitute sensitive information for airlines, we will denote the five airports as Airport 1, Airport 2, Airport 3, Airport 4 and Airport 5. An external flight ticket broker has provided data on the evolution of ticket prices for available flights. Prices refer to the lowest available fare for the given flight. The combination of actual load factor data provided by airports with ticket price information provided by the broker creates a comprehensive dataset that includes ticket price details for each flight, actual load factors, the count of available seats, and supplementary flight information such as the operating airline, departure date, and time.

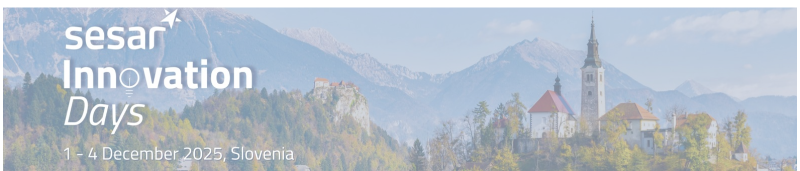

For each flight, we utilize the ticket price available 10 days before departure. The choice was made based on the trade off between the need of airport operators to evaluate the impact of demand in advance while using the best available predictions. In particular, certain airport operators expressed a preference for this time horizon. Prices in close proximity to the departure date offer more informative insights into demand levels. To conduct our exploratory research, we partitioned the dataset into training, validation, and test sets. The training set comprises data from the first 13 weeks of the year. Week number 14 served as the validation set, with the remaining 11 weeks forming the test set. This chronological division was executed to demonstrate the model's ability to forecast future demand based on historical data.

Notably, the selection of the test set is strategic; it encompasses Easter holidays and the different May and June bank holidays. This diverse composition presents a formidable challenge, creating a demanding scenario in which to rigorously validate our results. During this period, airports might experience high variations in passenger flow with an impact on airport operations. As far as the authors are aware, a similarly scrupulous validation has not been performed for comparable previous works in the literature.

B. Price Ratio

As ticket prices are expected to play an important role in the prediction of the load factor, this paper introduces an indicator that encapsulates the evolution of ticket prices for a certain route with respect to the past. The price ratio (PR) is defined as the ratio between each flight ticket price and the corresponding reference price (RP), i.e.:

$$PR_i = \frac{TP_i}{RP_i} \tag{1}$$

where PR_i is the price ratio, TP_i is ticket price and RP_i is reference price for a given flight i. The RP_i is the median price calculated over the last 56 days for the given flight i. Note that, in this context, the flight i is a uniform city-pair, i.e., a recurring event that can be classified based on airline, departure airport, and arrival airport. It is also clear from the above that not every flight i will contain the same number of ticket prices in a given time period, as this depends on its frequency.

In cases where there is a new route between the departure airport and the arrival airport and there is no available history for the past 56 days (8 weeks), the entire available history, including all routes and airports, is considered when calculating the median price RP.

Since RP evolves over time, it's not feasible to directly compare the PR values of flights operating in different periods of the year. Instead, the primary objective of this indicator is to provide insight into demand trends compared to the preceding 8 weeks. If PR exceeds 1, it indicates an increase in demand compared to the average of the previous 8 weeks. On the contrary, a value below 1 suggests a declining trend, thus no or minimum changes to a load factor predicted. The selection of an 8-week period strikes a balance, considering price variations throughout the year while ensuring sufficient stability. This consideration ignores external factors, such as fuel prices or airlines marketing strategies, which might affect the PR adding noise to the model. However, in this work, we prove a clear correlation between prices and demand even when this external noise is present. For the rest of this work, the PR_i of a flight i is intended, unless otherwise specified, to be computed using the TP_i available 10 days before departure.

While the PR indicator offers a valuable estimate of future passenger demand, combining it with additional flight information in a machine learning model can enhance our understanding of demand dynamics. However, in cases where historical load factor data is unavailable for training a machine learning model, PR stands as an effective means to qualitatively estimate future demand.

C. Airport Model

In general, one of the key advantages of using machine learning for load factor prediction is that it can analyze a wide range of variables and factors that may influence demand for air travel, including ticket prices, departure times, route distances, external events, size of an airport, and more. By identifying the most important factors and relationships in the data, machine learning algorithms can make accurate load factor predictions.

The choice of the most suitable machine learning model was based on historical data provided by partner airports. Historical data revealed how load factor distributions changed significantly between different airports and therefore it is necessary to train individual models for each airport splitting also departure and arrival flights. This means that two models are built for each airport, a departure airport model and an arrival airport model.

Gradient boosting regressor proved to be the best performing model in this context. The model is commonly used in the scientific literature thanks to its generalization capabilities and great performance on tabular data [16]. As the same architecture is used for different airports, generalization capabilities are fundamental to achieve the target results for all of them.

There are at least six load factor drivers in the airline industry according to [8]. The first driver is the industry's output decisions relative to demand growth. The second driver is pricing. Fare reductions generally stimulate demand. Load factors are affected depending upon capacity decisions. The third driver is the traffic mix. Historically, the higher the proportion of business travelers carried by an airline, the lower the average seat factor. That is, the random element in demand for business travels (highly volatile demand) suggests a lower average load factor in business and first class cabins [14]. The fourth driver is refund policies. A carrier taking non-refundable payment at the time of reservation is likely to have relatively fewer no-shows and a relatively higher seat factor than on selling a greater portion of tickets on a fully flexible basis. The fifth driver is commercial success. The success of product design, promotions, marketing communications,

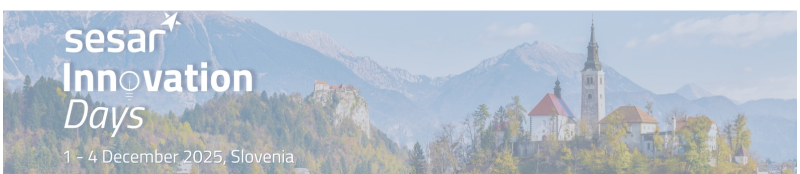

distributions, and service delivery will clearly influence current load factors. The sixth driver is revenue management. The effectiveness of revenue management systems (RMS) will influence load factors.

Each airport model was trained on the corresponding training set (based on airport and direction) extracted from the global training set. Before training, feature selection and hyper-parameter tuning were applied through k-fold cross validation on the validation set (with k equal to 5). The feature selection and the hyper-parameter tuning were done averaging the results of the different models. To avoid increased maintenance cost and improve the scalability of the solution every model used adopted the same set of features and the same hyper-parameters. This would allow to quickly add new airports to the system in the future. The feature selection process identified 9 relevant features after using GINI importance. GINI importance is a metric used in decision tree-based machine learning models to quantify the relative significance of each feature in predicting the target variable [17], helping us to understand which features contribute the most to the model performance. GINI importance was used due to its low computation time and capability of providing a ranking of the features based on their contribution to the model's predictive power [18]. Features with higher GINI importance are considered more important for making accurate predictions. The set N of relevant features for the model can be found (in order of importance) below:

1) Price Ratio: price indicator computed at D-10, see equation 1;
2) Price: ticket price of the lowest available fare at D-10 expressed in Euros;
3) Departure/Arrival airport: one-hot encoding of the departure/arrival airport;
4) Airline: one-hot encoding of the airline;
5) Seats: Total number of seats on the aircraft;
6) Time of day: morning, afternoon, evening or night;
7) Frequency: number of flights per day on a given route;
8) Low cost: true if the airline is a low cost airline;
9) Weekday: day of the week expressed in number.

The hyper-parameter tuning of the models was done applying a grid search on the validation set using as an error metric the mean absolute error (MAE). Thanks to the limited complexity and low training time of the model, grid search could be used to exhaustively explore the hyper-parameter space [19].

D. Passenger Model

Although the load factor metric well-summarises the demand, in several use cases the number of passengers on a flight constitutes a more useful metric. The number of predicted passengers is computed as:

$$PAXp_i = LFp_i \cdot Nseats_i \tag{2}$$

where $PAXp_i$ is the predicted number of passengers, the LFp_i is the predicted load factor and the $Nseats_i$ is the number of available seats on flight i as provided by the airport. Finally, the total daily number of predicted passengers $PAXp_j$ can be computed as:

$$PAXp_j = \sum_{i=1}^{n} PAXp_i \tag{3}$$

where j is a given day and n is the total number of flights for day j.

E. Analysis

As the goal is to analyze the effectiveness of the PR indicator, the analysis is split in two parts. The ticket price analysis is designed to prove the relation between the PR and the passenger demand. Then, the airport model analysis evaluates if the combination of the PR together with other features in a machine learning model can make accurate predictions. As a comparison, the airport model was evaluated in several tests against a baseline model where ticket information, and therefore the PR, is not available.

1) Ticket price analysis: The aim of this phase is to analyze the ticket price evolution of different flights depending on the demand. Furthermore, it explores the relation between the price evolution and the load factor. The objective is to prove that the PR effectively encapsulates the passenger demand. At first, the price evolution is analysed to compare how flights with different demands experience a different price evolution. Ticket prices are expected to be driven by two main factors: the demand and the number of days before the departure.

In this analysis the entire dataset that include 72 818 flights is used and the analysis is based on actual load factor data. The data is split in three categories based on the actual load factor, i.e. high demand (flights with a load factor >0.9), medium demand (flights with a load factor between 0.6 and 0.9) and low demand (flights with a load factor <0.6). Combining data-driven insights from passenger demand prediction distributions with expert judgments provided comprehensive basis for categorizing flights into these three categories. The analysis is designed to confirm the business understanding of the problem and how much airlines adapt their ticket price based on the demand compared to other marketing reasons.

Note that part of the bias might remain due to the fact that flights might not be uniformly distributed in the three categories as some routes are commonly busier than others and airlines adopt different pricing strategies.

2) Airport Model Analysis: Once the expected relation between ticket price and demand is proved, this second experiment analyzes the performance of the airport model on the data provided by the airports. It evaluates how well the airport model can predict the load factor for a single flight and how accurately it can predict the number of daily passengers for a given airport. Two metrics are used: one to analyze the performance at flight based level ($FlightBasedLevelError$) and one for the performance at daily level ($DailyLevelError$). For the $FlightBasedLevelError$, $RealLoadFactor_i$ and $PreidictedLoadFactor_i$ are compared, for every flight i in the test set, to analyse how accurately the model can predict the demand on a single flight. The adopted metric for this evaluation is the mean absolute error (MAE), i.e.

$$FlightBasedLevelError = \frac{1}{n}\sum_{i=1}^{n} |PreidictedLoadFactor_i - RealLoadFactor_i| \quad (4)$$

where n is the total number of flights in the test set. This metric is important if the focus is to have a precise forecast of the demand at low granularity, for example per hour or Schengen/non-Schengen. The metric is expressed in percentage points (pp) as computed as a difference between two percentages (load factors). The $FlightBasedLevelError$ can also be expressed using the Mean Absolute Percentage Error as reported by [10]. However, MAE is used as the preferred metric in this work as explicitly preferred by participating airports.

Secondly, the $DailyLevelError$ compares the $PredictedDailyPassengers_j$ and the $RealDailyPassengers_j$ for each day j in the test set. The metric used is the mean percentage absolute error, i.e.

$$DailyLevelError = \frac{1}{N}\sum_{j=1}^{N} \frac{|PredictedDailyPassengers_j - RealDailyPassengers_j|}{RealDailyPassengers_j} \quad (5)$$

where N is the total number of days in the test set. This second metric is important if the focus is to get an idea of the evolution of the demand over time and which dates will be the most critical. Moreover, this metric weights flights using bigger airplanes more than the others.

For both metrics, the airport model is compared with a baseline model with an identical architecture but without ticket price information to prove the importance of this feature especially during variable periods such as holidays. The accuracy of the airport model is computed using the real historical load factors provided by the airports. The final goal is to better understand the applicability of this solution in a real environment and the business value that it could create.

The experiments were executed using Python 3.11 on the Cloudera Machine Learning platform. The virtual machine selected was a 2vCPU machine with 20 GB of memory.

III. Results

In this section, we report the results obtained from the methodology described in the previous section. In particular we report the results of the two analysis described in section II-E: ticket price analysis and airport model analysis on the identified dataset.

A. Ticket Price Analysis

In Figure 1 the median price evolution computed n days before departure is shown. The median is used to exclude outliers with extremely low or high price that would bias the results. This helps to address ticket price volatility, that may be brought about by some rare promotional fares or unusually high last-minute prices. Days-to-Departure is used to compare the price evolution on the same scale independently from the real departure date. In Figure 1 is evident how flights with a high demand have a higher ticket price also 8 weeks before the departure and they keep an higher price until the departure date. Indeed, the initial median price for high demand flights is 143 euros compared to respectively 122 and 103 euros for medium and low demand flights, see Figure 1. Prices steeply increase for all the categories in the last 10 days before the departure. This constitutes ad additional justification to the choice of using 10 days as default window for prices.

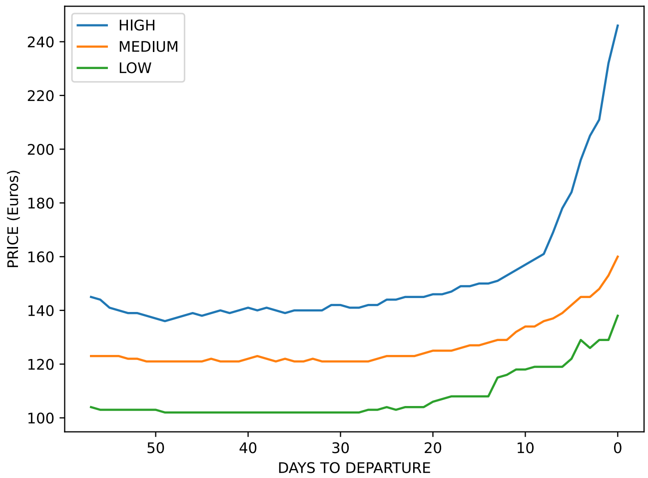

Figure 1. Median price evolution with respect to demand. Blue curve represent high demand flights, orange curve represent medium demand flights and green curve represent low demand flights.

Secondly, the correlation between price ratio (computed 10 days before departure) and the load factor is analyzed using real aggregated data provided by the 5 airports. The goal is to validate the correlation between the two indicators justifying the use of the price ratio when predicting the passengers demand. In Figure 2 a box plot of the distribution of the load factors for different price ratio intervals is shown. Price ratio values are grouped in 9 different intervals to perform the analysis. The median load factor passes from 77 % when the price ratio is low, up to 94 % for high price ratio values. A price ratio equal to 2 means that the ticket price is twice the price of its reference price in the past.

B. Airport Model Analysis

At a flight based level, the $FlightBasedLevelError$ is used as a metric to analyze the quality of the predicted load factor.

In order to analyze the importance of ticket prices, the airport model is compared to a baseline model (no ticket

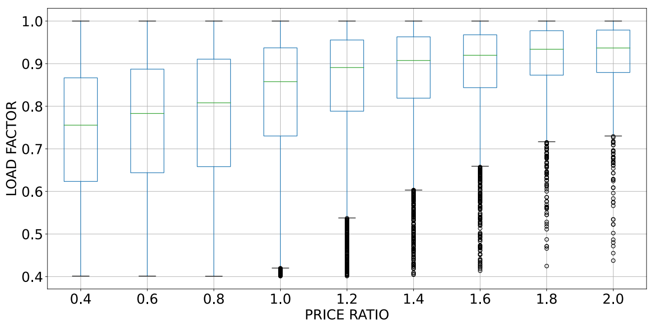

Figure 2. Graphical representation of the distribution of price ratios (computed 10 days before departure) with respect to the load factor in the form of a box plot.

information). Therefore, testing is performed in two setups: (1) Overall – Airport model using all the available features and (2) No ticket information (NTI) – Airport model without any access to ticket price information. The NTI model is comparable (in terms of features) to the model used by [10] with some additional features such as the frequency and the number of seats.

The average $FlightBasedLevelError$ of the 10 models on the test set is used in the tables. Results are shown in Table I and Table II. The model performance degrades when ticket information is not present, with a difference of more than 2 pp in the MAE and 3% in the MAPE. For completeness, we also report the predictions computed using ticket prices available at D-7 and D-14 in addition to D-10. The choice of using D-10 for the remaining experiments is based on the trade off between providing predictions to airport operators sufficiently in advance while maximizing performance although slightly better results could be achieved at D-7.

TABLE I. Mean absolute errors ($FlightBasedLevelError$) for test predictions when the airport model uses all the available features (Overall) and for case when airport model works without any access to ticket price information (NTI).

Type	MAE-7days	MAE-10days	MAE-14days
Overall	10.8 pp	11.2 pp	11.3 pp
NTI	13.7 pp	13.7 pp	13.7 pp

TABLE II. Mean absolute percentage errors for test predictions when the airport model uses all the available features (Overall) and for case when airport model works without any access to ticket price information (NTI).

Type	MAPE-7days	MAPE-10days	MAPE-14days
Overall	15.4 %	15.6 %	15.9 %
NTI	18.5 %	18.5 %	18.5 %

In Table III the same information is shown for the 5 different airports participating in the experiments. The $FlightBasedLevelError$ is strongly dependent to the given airport due to the specific characteristics of the different European regions. On the contrary, there is no major difference in the performances between departures and arrivals except for airport 5. Airport 4 has the best performances with an error below 8 pp, while Airport 2 error is about 13 pp. These differences are mainly due to different load factor distributions that make demand more or less predictable. Another factor is the drift of the test load factor compared to the one from the training set, due to the different considered periods of the year, which makes certain predictions less accurate.

TABLE III. Mean absolute errors ($FlightBasedLevelError$) for test predictions for each model (departure – D, arrival – A) and each airport (A1–A5) in case when airport model uses all the available features (Overall) and for case when airport model works without any access to ticket price information (NTI).

Type	D-A1	A-A1	D-A2	A-A2	D-A3	A-A3	D-A4	A-A4	D-A5	A-A5
Overall	10.2 pp	9.3 pp	13.7 pp	12.2 pp	12.2 pp	13.3 pp	7.9 pp	7.8 pp	10.2 pp	16.1 pp
NTI	12.2 pp	12.1 pp	16.3 pp	14.6 pp	14.8 pp	15.6 pp	9.8 pp	9.0 pp	13.3 pp	19.2 pp

As the $FlightBasedLevelError$ evolves over time, it is interesting to analyze its evolution during the 10 weeks test period. In Figures 3(a,c,e) and 4(a,c) the evolution of the daily predicted load factor of the two setups is compared with the real one. The demand is different in the different airports although some common trends could be identified, such as the peak in the demand during the Easter holiday and the several May and June regional holidays. From the figures it is clear how the baseline model is still capable of replicating the evolution of the demand during the week (e.g. the week-end peaks) but it cannot anticipate the overall shifts of the demand over time, especially during the holidays. On the contrary, the overall setup matches quite accurately the shifts, although smoothing the extreme values. In Figure 3(b,d,f) and Figure 4(b,d) the reason why price ratio is so relevant in load factor predictions is shown. The figures reveal a comparison between the evolution of the daily average load factor and the daily average price ratio. A moving average with window size of 7 days is used to smooth the curves as the goal is to focus on the ascending or descending trends. A similar evolution of the two indicators is clear for all airports. Regarding airport 1, the demand and the price ratio reach their peak at the end of the Easter holiday as the airport is located in a holiday destination Figure 3b. For airport 2, the descending trend after the Easter holidays is perfectly matched with the average load factor dropping to 66 % while the price ratio is below 0.8, see Figure 3d. The peaks in the price ratio at the end of May and mid-June perfectly match the peaks in the load factor at 76 % see Figure 3d. Finally, in the case of airport 4, after the Easter peak, the price ratio dropped to a stable value (between 1.2 and 1.3) suggesting for a stable demand that, in fact, is characterized by a steady load factor of 84 % see Figure 4b.

In contrast, for analysis focusing on daily predictions, the $DailyLevelError$ is used. In Table IV is clear how ticket

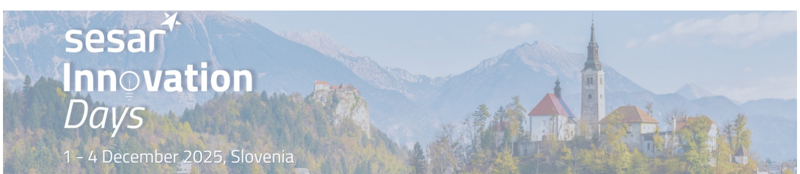

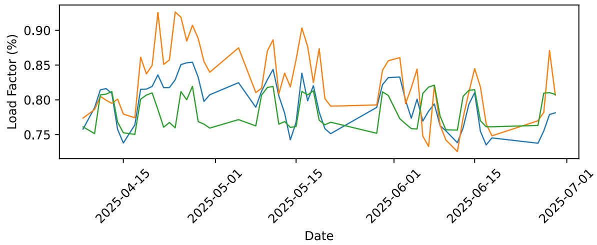

(a) Airport 1 airport model analysis, FlightBasedLevelError

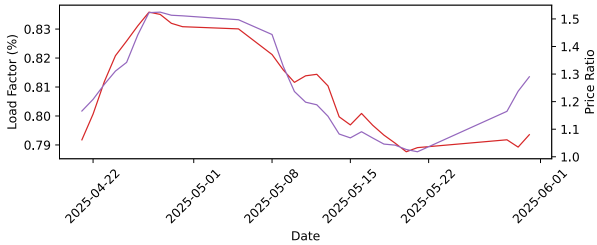

(b) Airport 1 airport model analysis, Price ratio evolution wrt load factor

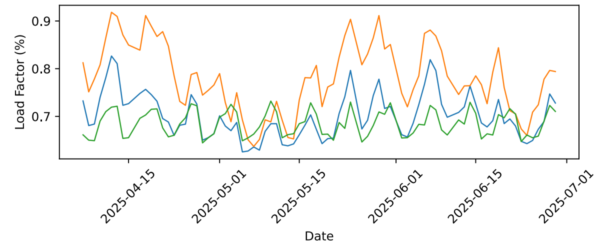

(c) Airport 2 airport model analysis, FlightBasedLevelError

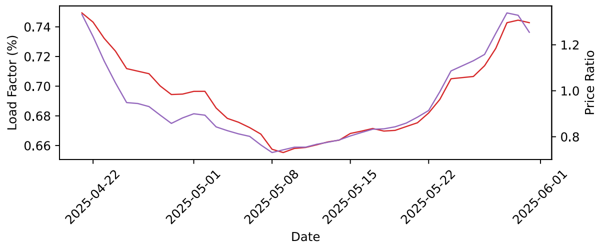

(d) Airport 2 airport model analysis, Price ratio evolution wrt load factor

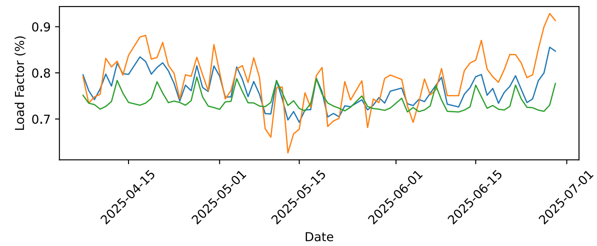

(e) Airport 3 airport model analysis, FlightBasedLevelError

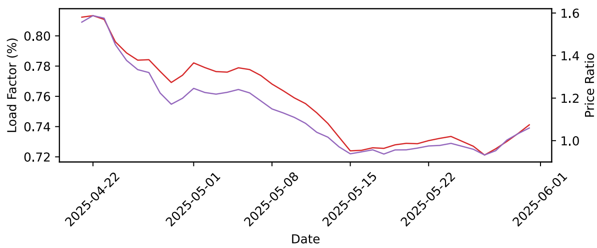

(f) Airport 3 airport model analysis, Price ratio evolution wrt load factor

Figure 3. Evolution of the average daily load factor for case when airport departure model uses all the available features (blue), for case when airport departure model works without any access to ticket price information (green) and for the real data (orange) for airport 1 (a), airport 2 (c), airport 3 (e). Evolution of the rolling average load factor (red) and price ratio (purple) for airport 1 (b), airport 2 (d), airport 3 (f).

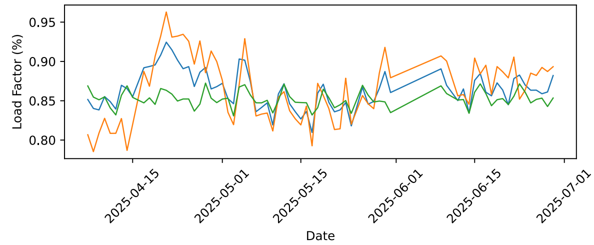

(a) Airport 4 airport model analysis, FlightBasedLevelError

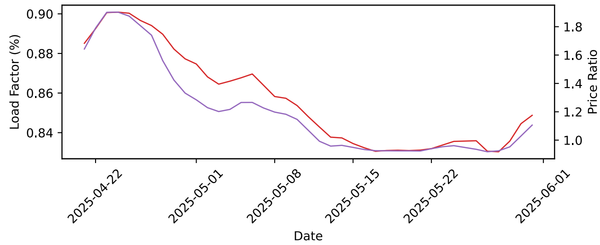

(b) Airport 4 airport model analysis, Price ratio evolution wrt load factor

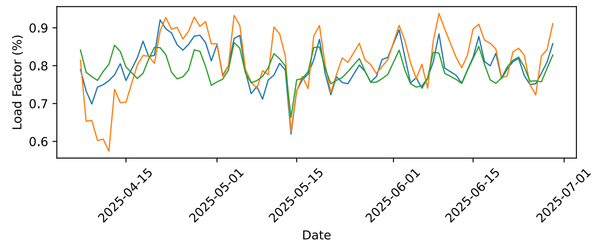

(c) Airport 5 airport model analysis, FlightBasedLevelError

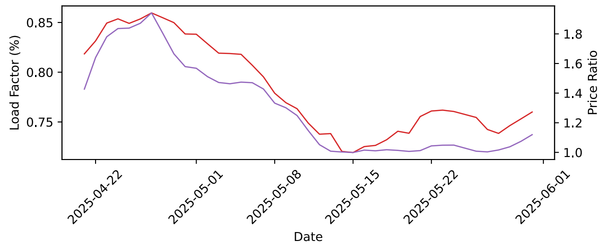

(d) Airport 5 airport model analysis, Price ratio evolution wrt load factor

Figure 4. Evolution of the average daily load factor for case when airport departure model uses all the available features (blue), for case when airport departure model works without any access to ticket price information (green) and for the real data (orange) airport 4 (a) and airport 5 (c). Evolution of the rolling average load factor (red) and price ratio (purple) for airport 1 (b), airport 2 (d).

prices provide a beneficial support in the predictions of the number of daily passengers reducing the average error from 8.5 % down to 5.75 %. The errors are overall lower compared to Table V because individual errors in single flights compensate each other during the day.

IV. Discussion

[8] analyzed passenger load from the past historical pattern and developed a predictive model using decision trees to forecast the passenger load based on certain criteria. Despite good results, it is not clear how well in advance of the

TABLE IV. Mean absolute percentage errors ($DailyLevelError$) for daily predicted passengers in case when airport model uses all the available features (Overall) and for case when airport model works without any access to ticket price information (NTI).

$DailyLevelError$	
Overall	5.75 %
NTI	8.5 %

TABLE V. Mean absolute percentage errors ($DailyLevelError$) for predicted passengers for each model (departure – D, arrival – A) and each airport (A1–A5) in case when airport model uses all the available features (Overall) and for case when airport model works without any access to ticket price information (NTI).

Type	D-A1	A-A1	D-A2	A-A2	D-A3	A-A3	D-A4	A-A4	D-A5	A-A5
Overall	4.4 %	6.6 %	8.5 %	6.9 %	4.0 %	4.6 %	2.4 %	4.3 %	5.1 %	11.4 %
NTI	6.3 %	8.9 %	11.4 %	9.3 %	7.5 %	6.2 %	4.6 %	5.9 %	7.8 %	17.2 %

operational day, predictions were generated. At the same time, the correlation between real passenger loads and those predicted was conducted for a single airport that makes the model somehow limited. Similarly, [11] addressed the problem of per-flight passenger predictions. Our work extended their approach by integrating new features, such as ticket prices and the price ratio, that significantly reduces the overall mean absolute error. Both the MAE and the MAPE of the overall model are significantly lower than those of the NTI. Note that the MAPE of the NTI model is higher than the one reported by [11]. This is due to different time horizons and different airports considered. Stable periods such as one month during the summer holidays are easier to predict due to the reduced variance in the load factor. In addition, our Airport Model was validated on multiple individual airports from different European regions, proving an excellent generalizability of the model performance.

We developed a predictive model that helps to generate anticipated load factors on respective routes operating into and from the European Civil Aviation Conference (ECAC) area using a novel theater in the modelling. The price of a flight ticket and its evolution was identified as a good candidate to better understand the volume of seats being sold to passengers on a given flight. In other words, knowledge of the price value of a ticket provides insight into the potential occupancy of an aircraft.

As such, the research objectives linked to development of a model being capable of generating passenger demand predictions on per flight basis were achieved. As explained previously, there is limited research conducted in this region; however, performance of the model achieved high accuracy as being assessed against real load factors received from subject airports. Furthermore, the work done has been further broadened and currently 23 airports actively subscribe to the load factor predictions provided. The latter being used as a primary or complementary input into their own predictions in order to optimize the resource planning. This fact underlines the significance of the results presented and supports the statements claiming the model to be scalable, which means it can be used at any European airport. MAE around 10 pp per flight constitutes credible input for airport staff when preparing the operational roster. At the same time, the performance of the model did not deteriorate when capturing demand for highly volatile periods, such as Easter or May and June holidays for different European regions. Additionally, the results presented clearly demonstrate that including ticket prices into the list of model features reduces the prediction error by 2.5 pp. In total and considering the current state of the art together with the exposure of the application to the outside world, we have confidence to state, that developed work presents a breakthrough in passenger demand predictions.

As explained, the predictions are highly dependent on availability of ticket prices, which allow for understanding of price vs load factor ratio. This is directly linked to revenue management policies across the airlines industry. However, pricing policies are constantly changing, and this causes deterioration of the prediction of certain routes, at certain times, at certain airports. In other words, although the predictions are accurate, the error is not uniformly distributed. The idea is to make the predictions stable and robust enough to avoid erroneous estimations of the load factors leading to false expectations of the airport operators that try to act upon the predictions when allocating resources. Operationally, the models should be enhanced with recent training data as soon as they become available. This to avoid the RP used to compute the price ratio in the training data from being excessively different from the ones used during inference. Finally, the more reliable the service is, the greater the likelihood for improved turnaround and departure predictability as highly occupied flights often impact turnaround durations.

V. Conclusion

In this study, we navigated the intricate task of predicting flight load factors by harnessing machine learning models. We emphasized the critical role of data quality and tailored our model selection to accommodate data nuances. Our innovative approach incorporated flight ticket prices as a pivotal factor in load factor predictions. While earlier research offered valuable insights by analyzing historical load patterns and utilizing decision trees, certain limitations called for further exploration. Notably, the lead time for predictions and the focus on a single airport constrained their model's scope. Nevertheless, their groundwork paved the way for our dynamic and forward-looking study.

We introduced a novel predictive model capable of generating accurate load factor predictions, with ticket prices taking center stage. Ticket prices emerged as a powerful indicator, shedding light on potential aircraft occupancy and enabling precise passenger demand predictions for individual flights.

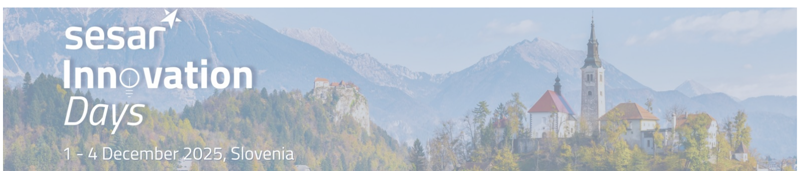

Our research achieved its primary objectives by delivering a model with exceptional predictive accuracy. Considering the limited prior research in this field, the model's performance against real load factor data from subject airports assumes particular significance. The widespread adoption of our predictions by 14 active airports, along with 9 additional ECAC airports expressing interest, underscores the model's scalability and practicality.

Our model demonstrated remarkable resilience, maintaining its accuracy even during periods of high volatility, such as the Easter travel season. By incorporating ticket prices into the predictive framework, we achieved a 2.5 pp reduction in prediction errors, marking a substantial improvement. Together, our work represents an original advance in passenger demand predictions. The main strengths of the presented work are linked to the prediction horizon (10 days in advance) and low prediction error during a test period that includes special events (Easter period and May and June bank-holidays) when the demand distribution is different from the rest of the year.

On the other side, we also acknowledge that our model's dependence on ticket prices introduces challenges tied to the dynamic nature of airline pricing policies. Although the model adeptly captures trends and reacts to market fluctuations, evolving pricing strategies can occasionally impact prediction accuracy for specific routes, times, and airports. Looking ahead, there are promising opportunities for further development in this field. Direct access to evolving ticket sales data per flight is proposed as a means to provide better prediction stability. Moreover, future research endeavors can explore additional predictive factors and their implications for load factor predictions, driving continued innovation and refinement in this critical domain. For example, we are conducting additional experiments using passenger intentions as an addition or replacement to ticket prices. Specifically, we are analysing whether the volume of online searches made by passengers intending to travel between Airport A and Airport B has a measurable impact on the model's prediction accuracy. The first results are positive, and the new indicator strengthens the overall reliability of the predictions.

In summary, our study has illuminated the potential of integrating ticket prices into the realm of load factor predictions, offering profound benefits for resource allocation, operational planning, and passenger experiences within the aviation industry. While pricing policy challenges persist, our model's accuracy, scalability, and resilience make it a powerful operational tool. As we forge ahead, the opportunities for expanding our understanding and improving our predictions remain boundless, ensuring that this field continues to evolve and thrive.

Acknowledgment

This work was carried out under the umbrella of the EUROCONTROL Air Transport Innovation Network (EATIN) initiative in collaboration with operational stakeholders throughout the European region.

References

[1] Avsar Zeynep Muge Selcuk Ahmen Melih. Dynamic pricing in airline revenue management. *Journal of Mathematical Analysis and Applications*, pages 1191–1217, 2019.

[2] Vojin Tošić. A review of airport passenger terminal operations analysis and modelling. *Transportation Research Part A: Policy and Practice*, 26(1):3–26, 1992.

[3] Jeffrey C Price and Jeffrey S Forrest. *Practical Airport Operations, Safety, and Emergency Management, Protocols for today and the future*, chapter 9 – Landside and Terminal Operations, pages 339–388. Elsevier, Amsterdam, Netherlands, 2016.

[4] S Mokhtarimousavi and A Mehrabi. Flight delay causality: Machine learning technique in conjunction with random parameter statistical analysis. *Journal of Tansportation Science and Technology*, pages 230–244, 2023.

[5] H Jenatabadi and N Ismail. The determination of load factors in the airline industry. *International Review of Business Research Papers*, 3:189–197, 2007.

[6] R Merkert and T Webber. How to manage seasonality in service industries - the case of price and seat factor management in airlines. *Journal of Air Transport Management*, pages 39–46, 2018.

[7] Nilabhra Banerjee, Alec Morton, and Kerem Akartunalı. Passenger demand forecasting in scheduled transportation. *European Journal of Operational Research*, 286(3):797–810, November 2020.

[8] Ma Nang Laik, Murphy Choy, and Prabir Sen. Predicting airline passenger load: A case study. In *2014 IEEE 16th Conference on Business Informatics*. IEEE, July 2014.

[9] Karina A Lundaeva, Zakhar A Saranin, Kapiton N Pospelov, and Aleksei M Gintciak. Demand forecasting model for airline flights based on historical passenger flow data. *Applied Sciences*, 14(23):11413, 2024.

[10] Jia He, Xiao-Qi Zhao, Bin Li, Jian Zhang, Tao Wang, and Tie-Qiao Tang. The short-time airport departure passenger volume prediction models based on aggregating information and flight feature similarity. *Expert Systems with Applications*, 272:126781, 2025.

[11] J Blasco-Puyuelo, J Burrieza-Galán, OG Cantú Ros, R Herranz, and D Mocholí. A machine learning approach for predicting airport passenger flows. *SESAR Innovation Days 2023*.

[12] Ioana C. Bilegan, Karim Achaibou, Carlos A.N. Cosenza, and Félix Mora-Camino. An adaptive approach for improved airline revenue management. *IFAC Proceedings Volumes*, 36(14):119–124, August 2003.

[13] S. Maertens. A metric to assess the competitive position of airlines and airline groups in the intra-european air transport market. *Research in Transportation Economics*, pages 65–73, 2018.

[14] Alberto A. Gaggero and Alexander Luttmann. The determinants of hidden-city ticketing: Competition, hub-and-spoke networks, and advance-purchase requirements. *Transportation Research Part E: Logistics and Transportation Review*, 173:103086, May 2023.

[15] Sveinn Vidar Gudmundsson. In search of sustainable strategies for low-cost long-haul airlines. *Case Studies on Transport Policy*, 12:100991, June 2023.

[16] Ravid Shwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *CoRR*, abs/2106.03253, 2021.

[17] Bjoern H Menze, B Michael Kelm, Ralf Masuch, Uwe Himmelreich, Peter Bachert, Wolfgang Petrich, and Fred A Hamprecht. A comparison of random forest and its gini importance with standard chemometric methods for the feature selection and classification of spectral data, 2009.

[18] Stefano Nembrini, Inke R König, and Marvin N Wright. The revival of the gini importance? *Bioinformatics*, 34(21):3711–3718, 2018.

[19] Tanay Agrawal. *Hyperparameter Optimization Using Scikit-Learn*, pages 31–51. Apress, Berkeley, CA, 2021.