Proceedings of the Twenty-Ninth AAAI Conference on Innovative Applications (IAAI-17))

Predicting Fuel Consumption and Flight Delays for Low-Cost Airlines

Yuji Horiguchi, Yukino Baba, Hisashi Kashima
Kyoto University
horiguchi@ml.ist.i.kyoto-u.ac.jp, {baba,kashima}@i.kyoto-u.ac.jp

Masahito Suzuki, Hiroki Kayahara, Jun Maeno
Peach Aviation Limited
{M.Suzuki,H.Kayahara,J.Maeno}@flypeach.com

Abstract

Low-cost airlines (LCAs) represent a new category of airlines that provides low-fare flights. The rise and growth of LCAs has intensified the price competition among airlines, and LCAs require continuous efforts to reduce their operating costs to lower flight prices; however, LCA passengers still demand high-quality services. A common measure of airline service quality is on-time departure performance. Because LCAs apply efficient aircraft utilization and the time between flights is likely to be small, additional effort is required to avoid flight delays and improve their service quality. In this paper, we apply state-of-the-art predictive modeling approaches to real airline datasets and investigate the feasibility of machine learning methods for cost reduction and service quality improvement in LCAs. We address two prediction problems: fuel consumption prediction and flight delay prediction. We train predictive models using flight and passenger information, and our experiment results show that our regression model predicts the amount of fuel consumption more accurately than flight dispatchers, and our binary classifier achieves an area under the ROC curve (AUC) of 0.75 for predicting a delay of a specific flight route.

1 Introduction

Airline deregulation eliminated the government restrictions on airline fees and resulted in the rise and growth of a new category of airlines called *low-cost airline* (LCA). LCAs compete with traditional full-service airlines by offering low flight fares. The success of LCAs lies in their commercial strategy of no-frills service, high fleet utilization, and fleet homogenization to reduce maintenance costs. LCAs had a global market share of more than 20% in 2013 and continue to gain popularity.

Significant effort is required to reduce the operating costs of LCAs in order to reduce their flight prices and compete with other airlines. One of the main expenses for LCAs is fuel costs. LCA companies apply various approaches to reducing aircraft weight, which affects fuel costs. For example, many LCAs impose strict restrictions on baggage weight, and some companies replace in-flight entertainment devices with lighter ones. An airline reported that it saved £600,000 in annual fuel costs by descaling the toilet pipes on their aircraft. [1]

Although the customers of LCAs are attracted by the low fares, they still demand high-quality services. As the market has become crowded, pressure has been placed on LCAs to improve their services. A common measure of airline service quality is on-time departure performance. Because the business model of LCAs relies on efficient aircraft utilization, the time between flights is likely to be small. To provide high-quality services and increase customer satisfaction, LCAs require additional efforts to avoid flight delays.

LCA companies continually collect data, including flight and passenger information, which are expected to provide useful insights in designing operating strategies for cost reduction and service quality improvement. In this paper, we apply predictive modeling approaches to real airline datasets and investigate the feasibility of machine learning methods for cost reduction and service quality improvement in LCAs. The datasets were collected over three years of operations by Peach Aviation, an LCA company providing flights between more than 20 airports in Asia.

We address two prediction problems: *fuel consumption prediction* and *flight delay prediction*. This paper first tackles the problem of predicting the amount of fuel consumption per flight. We use passenger information and flight information, including flight dates, routes, and aircraft IDs, to build regression models for predicting the amount of fuel consumption. This prediction would be helpful to optimize the amount of fuel on board, which would make it possible to reduce the weight of the aircraft and save fuel costs. Additionally, because fuel prices vary throughout the day and week, LCA companies desire an effective fuel purchasing strategy that considers the balance between the fuel price and demand. In order to explore the feasibility of supporting the purchase decision, we construct models to predict the weekly fuel consumption several months in advance.

We also address the flight delay prediction problem. Using flight and passenger information, we construct a binary classifier that predicts whether a flight will be delayed or not. If we can predict flight delays several days in advance, we will be able to assign a convenient gate or deploy more

[1] http://www.mirror.co.uk/news/uk-news/flush-with-cash-british-airways-saves-740383

ground crews to a flight that is likely to be delayed. If we can predict a delay the day before the flight, we can instruct staff to shorten the operation time such as for fueling. Because we can take different actions according to the timing of the prediction, we train predictive models using different sets of features and investigate the prediction accuracies; the features available one day before the flight, those available one week before the flight, and those available five months before the flight.

The results of our study to apply predictive modeling approaches to LCA datasets are summarized as follows:

- **Fuel consumption prediction**: Our regression model predicts the amount of fuel consumption for each flight with a relative root mean squared error (RMSE) of 8.8%. Our model performs better than human flight dispatchers and reduces their prediction errors by 39.3%. We verify our regression model is less likely to underestimate the amount of fuel consumption as well.
- **Flight delay prediction**: Using flight and reservation information, our binary classifier achieves an area under the ROC curve (AUC) of over 0.75 for several routes in the one-day-before predictions of the flight delays.

2 Related work

Although there have been many attempts to apply machine learning in the aviation industry (Caplener and Janku 2012; Mathur 2002; Li et al. 2011; Das et al. 2010; Zhang and Wang 2012; Melnyk et al. 2016; An et al. 2016; Ayhan and Samet 2016), there are only several studies which focus on predictions of fuel consumption and flight delays. A few reseachers attempted to estimate the fuel consumption using information given by quick access recorders, which provide quick and easy access to raw flight data (Hong, Gan-Xiang, and Xiao-Dong 2014; Jiaxue and Tao 2015). Marques and Leal created some flowcharts and generated profiles to estimate fuel consumption (2012). Haifeng, Xu-Hui, and Xin-Feng trained a support vector machine model to estimate the fuel consumption of a certain aircraft using the data from the route and an aircraft performance manual (2015). From the perspective of flight operational quality assurance (FOQA), an explicit aircraft performance model was used to analyze fuel consumption (Stolzer 2002). Different from these studies, we employ information about the passengers of each flight to predict the amount of fuel consumption.

In order to avoid flight delays, a few studies aimed to minimize the flight route length by formulating this problem as an optimization problem based on flight trajectory data (AhmadBeygi, Cohn, and Lapp 2010; Hu et al. 2016). In contrast, we apply predictive modeling approaches to predict when a flight will be delayed by using a variety of information such as information about the passengers and the reservations of the flight.

3 Airline datasets

We apply predictive modeling approaches to real airline datasets provided by Peach Aviation, an LCA company operating flights among more than 20 airports in Asia.

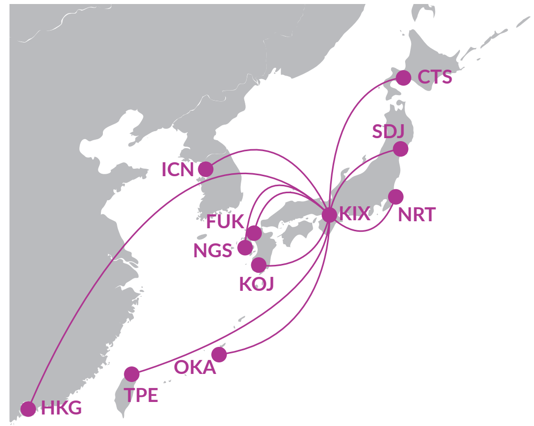

Figure 1: Target flight routes for our experiments

3.1 Data description

We have three datasets called *flight dataset*, *passenger dataset*, and *reservation dataset*. These data were recorded from July 2012 to March 2015, and contain $54{,}000$ flights and $9{,}900{,}000$ passengers in total. The flights were operated on 12 different airports in our datasets. Figure 1 shows the routes of the flights in our datasets.

- **Flight dataset**: This dataset contains general information about each flight, such as scheduled departure date and time, arrival date and time, airports of departure and arrival, and airframe ID. Such general information is fixed several month before each flight. We also have information about flight plans. A flight plan includes the scheduled amount of fuel on board and estimated time of the flight. These values are calculated by human flight dispatchers on the day before each flight.
- **Passenger dataset**: We have a list of passengers for each flight with their ages and genders.
- **Reservation dataset**: This dataset provides the reservation completion date of each passenger on each flight, which enables us to figure out how many people have already booked the flight at a given timing.

3.2 Flight feature representation

For addressing two problems, namely, fuel consumption prediction and flight delay prediction, we construct feature vectors for predictive modeling from the flight, passenger, and reservation datasets. We consider three prediction timings: one day before the departure date, one week before the departure date, and five months before the departure date. The available information is different depending on the timing of the prediction. Only the time and date information is available at the timing of five months before the departure date. A week before the departure date, in addition to the time and date, the reservation data and airframe ID are also available. The amount of fuel on board, standby position ID, and scheduled flight time can be used as additional features on the day before the departure. In addition, the number of

Table 1: List of features. We use different feature sets according to the time of the prediction: one day before the flight (D), one week before the flight (W), and five months before the flight (M).

Feature	Prediction timing D	W	M	Numerical or categorical	Num. of categories	Example
Flight dataset						
Departure year	✓	✓	✓	Numerical	-	*2014*
Departure month	✓	✓	✓	Categorical	12	*November*
Departure day of week	✓	✓	✓	Categorical	7	*Monday*
Departure day of month	✓	✓	✓	Numerical	-	*10*
Departure day of year	✓	✓	✓	Numerical	-	*314*
Scheduled departure time (in minutes of day)	✓	✓	✓	Numerical	-	*1110*
Scheduled arrival time (in minutes of day)	✓	✓	✓	Numerical	-	*420*
Airport of departure	✓	✓	✓	Categorical	11	*KIX*
Airport of arrival	✓	✓	✓	Categorical	10	*TPE*
Domestic or international flight	✓	✓	✓	Categorical	2	*Domestic*
Airframe ID	✓	✓	✓	Categorical	16	*XI000*
Standby position ID	✓	-	-	Categorical	20	*No.1*
Scheduled fuel on board (in litres)	✓	-	-	Numerical	-	*12000*
Estimated time enroute (in minutes)	✓	-	-	Numerical	-	*100*
Reservation dataset						
Num. of reservations	-	✓	-	Numerical	-	*150*
Num. of reservations by age	-	✓	-	Numerical	-	*0–11: 1, 12–19: 4, ...*
Num. of reservations by gender	-	✓	-	Numerical	-	*male: 46, female: 49*
Passenger dataset						
Num. of passengers	✓	-	-	Numerical	-	*170*
Num. of adult passengers	✓	-	-	Numerical	-	*140*
Num. of child passengers	✓	-	-	Numerical	-	*30*
Num. of passengers by age	✓	-	-	Numerical	-	*0–11: 3, 12–19: 5, ...*
Num. of passengers by gender	✓	-	-	Numerical	-	*male: 90, female: 55*

reservations observed on the day before the departure is almost equivalent to the actual number of passengers. This means that the actual number of passengers can be assumed to be available on the previous day of the flight. The list of features is shown in Table 1.

We count the numbers of passengers and those of reservations for each flight , and use these numbers as features. Additionally, since our datasets contain the age and gender of each passenger, we employ the numbers of passengers and reservations by age and gender. We classify passengers into 14 age intervals: "0–2", "3–11", "12–19", "20–24", "25–29", $\cdots$, "65–69", and "over 70". Then, the number of passengers (or reservations) in each age interval is used as a feature. We use the number of passengers who are at least 18-years old (called "adult passengers"), and that of passengers who are under 18-years old ("child passengers") as well.

We encode a categorical feature value, such as "airport of departure", into a feature vector by using one-of-K encoding. In order to take the periodicity into account, we transform the value of *departure day of year* by using trigonometric functions; let d be the departure day of the year of a flight and we incorporate $\sin\left(2\pi d/365\right)$ and $\cos\left(2\pi d/365\right)$ into a flight feature vector. This transformation enables us to handle New Year's Day and New Year's Eve as sequential dates. We also apply the transformation to "scheduled departure time" and "scheduled arrival time". We incorporate $\sin\left(2\pi m/1440\right)$ and $\cos\left(2\pi m/1440\right)$ into a feature vector where m is the scheduled departure (or arrival) time of a flight in minutes of day. Finally, we apply min–max normalization to the features to transform the each value of feature into the range of zero to one.

4 Fuel consumption prediction

We first address the problem of predicting the fuel consumption. Reducing fuel costs is one of the most important issues in LCAs because they account for the largest portion of the expenditures of airline companies. We consider two problem settings: predicting the fuel consumption for each flight and for each week.

4.1 Baseline

Airline companies decide on the amount of fuel on board based on the estimates made by domain experts (called flight dispatchers). The experts usually use some empirical rules such as those depending on the flight altitude or weather information. The amount of fuel on board is determined by considering both the expert estimations and the legal constraints. We use the estimations by the experts as our baseline, which we denote by EX-D. The experts always estimate the amount of fuel one day before the departure date.

4.2 Prediction models

We define the fuel consumption prediction as a regression problem. We apply three prediction models: random forests,[2] XGBoost (Chen and Guestrin 2016),[3] and deep neural networks. We apply a standard four-layer feed forward neural network trained with the AdaDelta optimization method (Zeiler 2012). Rectified linear unit (ReLU) and the sigmoid function are employed as the activation function of the hidden layers and that of the output layer, respectively. We use Keras[4] for implementing this model.

4.3 Fuel consumption prediction for each flight

Our first problem setting aims to predict an amount of fuel consumption for each flight. We consider two prediction timings: one day before and one week before the departure date of the flight. The predicted values are expected to be used for determining the amount of fuel on board or in storage. The prediction made one day before the flight is valuable for making a final decision of the amount of fuel on board, and the prediction made one week before the flight is useful for foreseeing the degree of future decrease in the storage. We denote the predictions by random forests, XGBoost, and deep neural networks at the timing of one day before the flight by RF-D, XGB-D, and DL-D, respectively, and those at the timing of one week before the flight by RF-W, XGB-W, and DL-W, respectively. The acronyms are summarized in Table 2.

We use the 47,000 flights from July 2012 to December 2014 for training, and the 7,000 flights from January 2015 to March 2015 for testing. The feature vectors are constructed from the information available at each prediction timing. We also observe that we can obtain accurate predictions of the number of passengers to some extent by using the information available one week before the flight, and therefore we incorporate the predicted numbers of passengers to build the fuel consumption prediction models of one week before the flight. The predictions of the number of passengers are given by deep neural networks.

We use the *relative RMSE* for our evaluation measure, which is the root mean squared error normalized by the mean of observed values. The actual amount of fuel consumption of each flight is used as the groundtruth. From the perspective of safety, we should avoid underestimation cases because an underestimation of the fuel consumption can result in a crash. Thus, we also evaluate an underestimate ratio, which is the frequency of the cases where the predicted value of a method is lower than 97% of the actual amount of fuel consumption.

Table 3 shows the relative RMSE of each method. We find that all the prediction models outperform EX-D even though RF-W, XGB-W, and DL-W produce predictions at the timing of one *week* before the flight and the experts estimate the values one *day* before the flight. It is also notable that XGB-D reduces the errors of EX-D by 39.3%. Figure 2 shows the relative RMSE according to the airport of departure or arrival. We observe that XGB-D outperforms EX-D in 10 out of the 12 airports, and RF-D is comparable to XGB-D in most of the airpots. Figure 3 shows the relative RMSE according to the week. XGB-D and RF-D consistently perform better than EX-D. Table 4 shows the underestimate ratios. We observe that the underestimate ratio of XGB-D is very small although it outputs accurate predictions. Figure 4 shows the underestimate ratios according to the airport of departure or arrival, and Figure 5 shows those according to the week. It is shown that XGB-D consistently achieves small underestimate ratios in all the airports and all the weeks. In contrast, the predictions given by EX-D are more likely to be lower than the actual values.

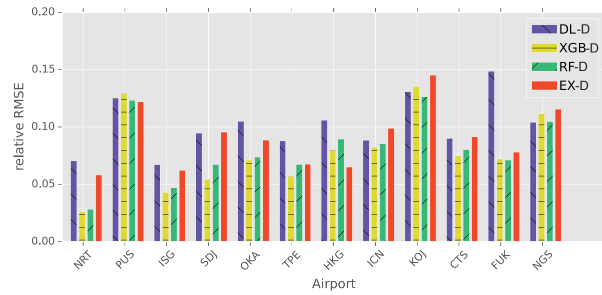

Figure 2: Relative RMSE of fuel consumption prediction for each flight according to the airport of departure or arrival. XGB-D outperforms EX-D in 10 out of 12 airports.

We find effective features by analyzing the model built by XGBoost. The trained XGBoost model indicates that "departure day of year", "cosine of departure day of year", and "number of passengers" are useful features for fuel consumption prediction. The first two features are helpful to incorporate seasonality to the prediction, and the last one is related to the weight of aircrafts. We also discover that "scheduled departure time" and "scheduled arrival time" are important features. They affect the amount of fuel consumption because the congestion of runways and that of airways depend on time, and they cause extra time for take-off and landing, which consumes an amount of fuel. Another possible reason is that selection of diversion airports depends on time; aircrafts have to carry an enough amount of fuel for the flights to diversion airports, and therefore the selection of diversion airports affects to the aircraft weight.

4.4 Fuel consumption prediction for each day

Because fuel prices are often highly volatile, airline companies want to hedge their risks in several ways, including bulk purchasing. Accurate estimates serve as valuable basic information supporting such risk hedging decisions. Thus, we next consider the problem of predicting the total fuel consumption for a particular day and route.

We use the same training and testing datasets for the fuel consumption prediction for each flight. We obtain the feature vectors from the information available five month before the date to build the prediction models. We denote the

[2] http://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html

[3] https://github.com/dmlc/xgboost

[4] https://keras.io

Table 2: Acronyms of each method according to the prediction timing

Method	Prediction timing		
	One day before the flight	One week before the flight	Five months before the flight
Random forests	RF-D	RF-W	RF-M
XGBoost	XGB-D	XGB-W	XGB-M
Deep neural network	DL-D	DL-W	DL-M
Human experts	EX-D	-	-

Table 3: Relative RMSE of fuel consumption prediction for each flight. All the prediction models outperform EX-D even though DL-W, XGB-W, and RF-W produce predictions at the timing of one week before the flight.

EX-D	DL-D	XGB-D	RF-D	DL-W	XGB-W	RF-W
0.145	0.119	**0.088**	0.092	0.120	0.094	0.100

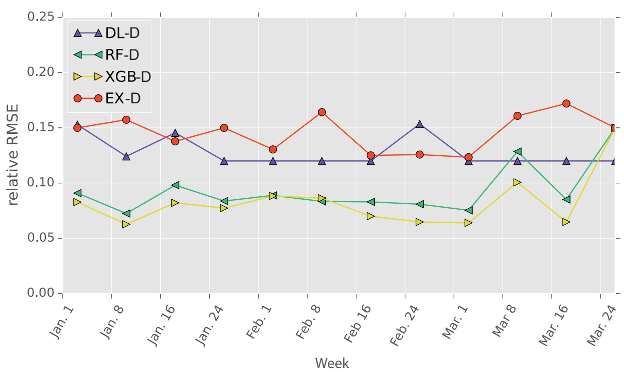

Figure 3: Relative RMSE of fuel consumption prediction for each flight according to the week. XGB-D and RF-D outperform EX-D in almost all the weeks.

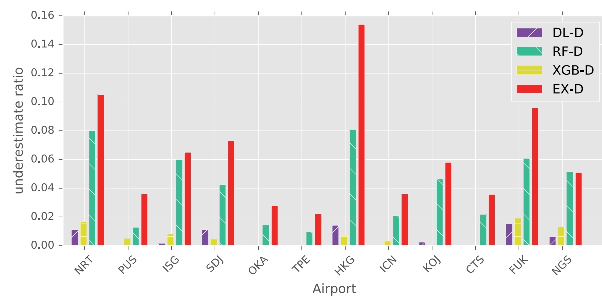

Figure 4: Underestimate ratios of fuel consumption prediction for each flight according to the airport of departure or arrival. The underestimate ratio of XGB-D is less than 1% for almost all the airports.

predictions by random forests, XGBoost, and deep neural networks at the timing of five month before the date by RF-M, XGB-M, and DL-M, respectively.

The relative RMSEs are shown in Table 5. Even five months before the date, XGB-M and RF-M outputs more accurate predictions than EX-D. The performance of DL-M is inferior to that of the other methods; we consider deep neural networks would be too complicated for this setting because we do not have many features at the timing of five month before the date. Figure 6 shows the relative RMSE according to the week; the prediction performance of RF-M and XGB-M are comparable to EX-D in all the weeks. Figure 7 indicates the relative RMSE according to the airport of departure or arrival. XGB-M outperforms the other methods for almost all the airports. We conclude boosting methods can be applied to predict day's use of fuel with better performance.

5 Flight delay prediction

The punctuality of operations is one of the most important indicators in the service quality of airline businesses. In particular, on-time departures are very important for LCAs because their flight schedules are very busy in comparison with traditional full-service airlines and LCAs have very short turnover times of aircrafts for cost saving; for example, an aircraft that has just arrived is immediately inspected and prepared for departure to another airport. Such overloaded schedules easily make flight delays and they propagate to subsequent flights to cause further delays.

An important first step to alleviate this problem is to predict flight delay in advance; in this section, we formulate the problem as a binary classification problem. It is reasonable to define this problem as a classification problem because the on-time performance rating, a common measure of airline service quality, is defined as the ratio of flights where the delay is lower than a certain threshold, and this rating does not concern the amount of delay time. In this paper, delayed flights are defined as "flights whose departure time is 15 minutes behind the schedule," based on the criteria of Federal Aviation Administration (FAA).

We apply random forests, XGBoost, and deep neural networks to the flight delay prediction problem. The settings of deep neural networks for this problem are the same as those for the fuel consumption prediction problem described in Section 4. We use the same training and testing datasets for the fuel consumption prediction. We compare three prediction timings, namely, one day before, one week before, and

Table 4: Underestimate ratios of fuel consumption prediction for each flight. The underestimate ratio of XGB-D is very small although it outputs accurate predictions.

EX-D	DL-D	XGB-D	RF-D	DL-W	XGB-W	RF-W
0.082	**0.002**	0.007	0.040	0.030	0.041	0.093

Table 5: Relative RMSE of fuel consumption prediction for each day. Even five months before the date, XGB-M and RF-M output more accurate predictions than EX-D.

EX-D	DL-M	XGB-M	RF-M
0.071	0.120	**0.064**	0.070

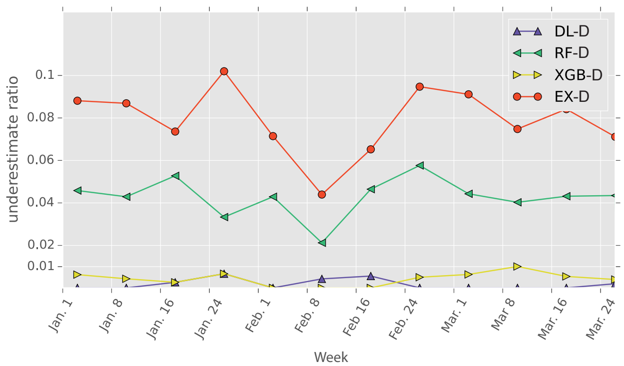

Figure 5: Underestimate ratios of fuel consumption prediction for each flight according to the week. XGB-D achieves smaller underestimate ratios than other methods in most of the weeks.

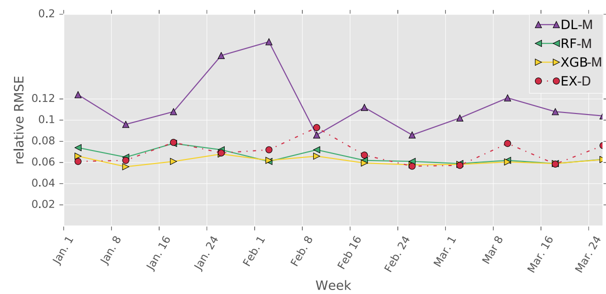

Figure 6: Relative RMSE of fuel consumption prediction for each day according to the week. RF-M and XGB-M demonstrate comparable performance to EX-D in all the weeks.

five months before the departure date, and obtain the feature vectors from the information available at each timing. As we did for the fuel consumption prediction, we incorporate the predicted number of passengers as features for the one-week-before prediction. We denote the predictions by random forests, XGBoost, and deep neural networks at the timing of one day before the flight by RF-D, XGB-D, and DL-D, respectively, those at the timing of one week before the flight by RF-W, XGB-W, and DL-W, respectively, and those at the timing of five months before the flight by RF-M, XGB-M, and DL-M, respectively. We apply an area under the ROC curve (AUC) as the evaluation measure.

Table 6 shows the AUC score of each method at each prediction timing; although the AUC scores of DL-D, XGB-D, and RF-D are over 0.6, those of the other methods are below 0.6. This result indicates that the information available at the timing of one month before or five months before the flight is not enough to predict the flight delays.

By investigating the trained models, we find that "fuel on board" and "number of passengers" are important features, which are available only at the prediction timing of one day before the flight. The lack of them would be one of the reasons for the poor performance of the predictions at the other timings. It is also observed that "scheduled departure minute of day" is an efficient feature for flight delay prediction. We plot the relationships between the "fuel on board," "the number of passengers," and "scheduled departure minute of day" on Figure 10. From this figure, we observe that the delayed and the on-time flights are clearly separated by these features.

Figure 8 shows the AUC scores of DL-D, XGB-D, and RF-D according to the airport of departure of arrival. We find that our prediction models achieve an AUC over 0.75 for NRT and NGS. Figure 9 illustrates the AUC scores for NRT and NGS according to the week. Our methods achieve an AUC of over 0.8 in several weeks. We thus conclude that our flight delay prediction models using the information available at one day before the flight are effective for such specific airports and weeks.

6 Conclusion

We conducted some studies on the feasibility of applying predictive modeling methods for low-cost airline companies to achieve cost saving and improve service quality. We predicted the amount of fuel consumption to save cost, and flight delay to enhance service quality. Because the airline companies can take different actions according to the timing of the prediction, we consider three prediction timings: the day before the departure reflected the state where all the reservations are almost perfectly collected, the week before the departure reflected the state where some of reservations have not yet been collected, and five months before the departure reflected the state where there were no reser-

Table 6: AUC scores of the flight delay predictions. The predictions at the timing of one day before the flights (DL-D, XGB-D, and RF-D) outperform the predictions at the other timings.

DL-D	XGB-D	RF-D	DL-W	XGB-W	RF-W	DL-M	XGB-M	RF-M
0.647	0.634	0.604	0.584	0.573	0.560	0.500	0.542	0.534

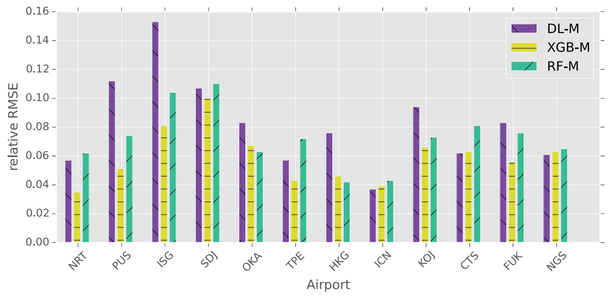

Figure 7: Relative RMSE of fuel consumption prediction for each day according to the airport of departure or arrival. XGB-M outperforms the other methods for almost all the airports.

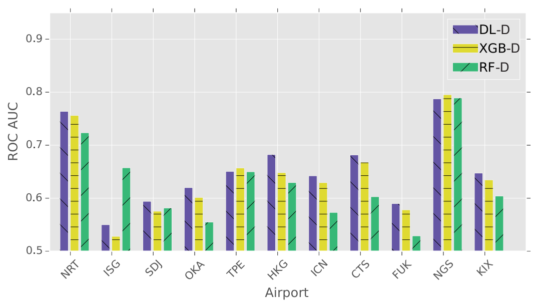

Figure 8: AUC scores of the flight delay prediction according to the airport of departure or arrival. The prediction models achieve an AUC over 0.75 for NRG and NGS.

vations. Our regression model predicted the amount of fuel consumption for each flight with a relative RMSE of 8.8%. Our model performed better than human flight dispatchers and reduced their prediction errors by 39.3%. Using the flight and the reservation information, our binary classifier achieved an AUC of over 0.75 for several routes in the one-day-before predictions of the flight delays.

Some promising directions for future work are described below. To predict the fuel consumption and flight delay, the flight data recorder data, meteorology information, and flying would be helpful to be used in combination with the passenger and the reservation information. With the goal of increasing sales, it would also be worthwhile to study a method for recommending seats or in-flight purchases using individual customer information.

References

AhmadBeygi, S.; Cohn, A.; and Lapp, M. 2010. Decreasing airline delay propagation by re-allocating scheduled slack. *IIE Transactions* 42:478–489.

An, B.; Chen, H.; Park, N.; and Subrahmanian, V. 2016. MAP: frequency-based maximization of airline profits based on an ensemble forecasting approach. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*.

Ayhan, S., and Samet, H. 2016. Aircraft trajectory prediction made easy with predictive analytics. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*.

Caplener, H. D., and Janku, J. A. 2012. Data mining applied to aviation data. *Morris Undergraduate Journal* 2.

Chen, T., and Guestrin, C. 2016. XGBoost: a scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data mining (KDD)*.

Das, S.; Matthews, B. L.; Srivastava, A. N.; and Oza, N. C. 2010. Multiple kernel learning for heterogeneous anomaly detection : algorithm and aviation safety case study. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data mining (KDD)*.

Haifeng, Z.; Xu-Hui, W.; and Xin-Feng, C. 2015. Support vector with ROC optimization method based fuel consumption modeling for civil aircraft. *Procedia Engineering* 99:296–303.

Hong, S.; Gan-Xiang, H.; and Xiao-Dong, W. 2014. An evaluation model for civil aircraft routes average fuel consumption. *Industrial Engineering Journal* 5.

Hu, Y.; Song, Y.; Zhao, K.; and Xu, B. 2016. Integrated recovery of aircraft and passengers after airline operation disruption based on a GRASP algorithm. *Transportation Research Part E: Logistics and Transportation Review* 87:97–112.

Jiaxue, L., and Tao, M. 2015. A method of aircraft fuel consumption performance evaluation based on relax signal separation. In *Proceedings of the 2015 IEEE International Conference on Cyber Technology in Automation, Control, and Intelligent Systems (CYBER)*, 1433–1437.

Li, L.; Gariel, M.; Hansman, R. J.; and Palacios, R. 2011. Anomaly detection in onboard-recorded flight data using cluster analysis. In *Proceedings of the 2011 IEEE/AIAA 30th Digital Avionics Systems Conference (DASC)*, 4A4–1–4A4–11.

Marques, B., and Leal, N. 2012. Airline fuel savings estimation based on segmented fuel consumption profiles. In

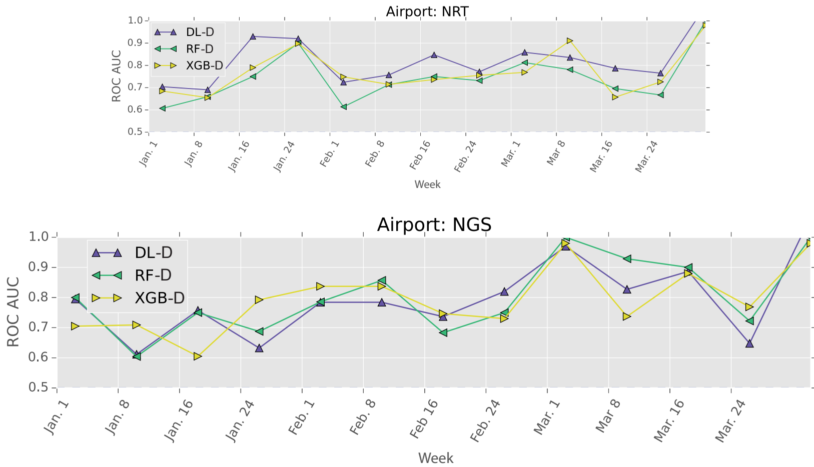

Figure 9: AUC scores of flight delay prediction for NRT and NGS according to weeks. The AUCs of all the methods vary among weeks, and the machine learning methods achieve an AUC of over 0.8 in several weeks.

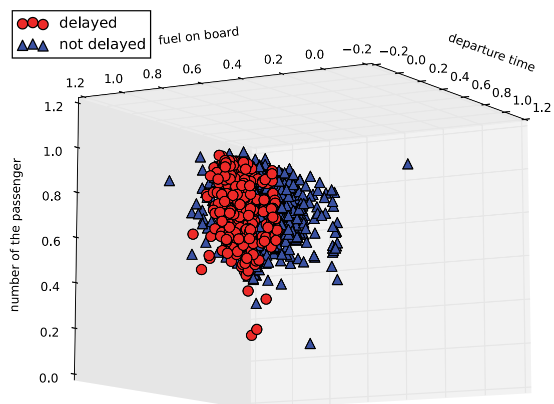

Figure 10: Relationships between the amount of fuel on board, scheduled departure time, and the number of passengers. The axis values are normalized. Each point represents a flight. The delayed and on-time flights are clearly separated by the values.

Proceedings of the 1st International Workshop on Information Technology for Energy Applications (IT4ENERGY).

Mathur, A. 2002. Data mining of aviation data for advancing health management. In *Proceedings of SPIE 4733, Component and Systems Diagnostics, Prognostics, and Health Management II*, 61–71.

Melnyk, I.; Banerjee, A.; Matthews, B.; and Oza, N. 2016. Semi-markov switching vector autoregressive model-based anomaly detection in aviation systems. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD).*

Stolzer, A. J. 2002. Fuel consumption modeling of a transport category aircraft using flight operations quality assurance data. *Journal of Air Transportation* 7(1):93–102.

Zeiler, M. D. 2012. Adadelta: An adaptive learning rate method. In *arXiv preprint arXiv:1212.5701, 2012*.

Zhang, C., and Wang, N. 2012. Aero-engine condition monitoring based on support vector machine. *Physics Procedia* 24:1546–1552.