

STATISTICAL MODELING OF AIRCRAFT ENGINE FUEL FLOW RATE

Yashovardhan S. Chati*, Hamsa Balakrishnan*
***Massachusetts Institute of Technology, Cambridge, USA**

Keywords: *fuel flow rate, flight data recorder, statistical modeling, CART, boosting*

Abstract

Knowledge of the fuel flow rate is important for accurate fuel burn and emissions inventory generation and for accurate engine performance modeling. In this paper, machine learning techniques are applied to aircraft flight recorder data to model the fuel flow rate as a function of the aircraft altitude, ground speed, vertical speed, and takeoff mass in the airborne phases of flight. Models are built using the Classification and Regression Trees (CART) and the Least Squares Boosting algorithms. The models are found to be interpretable. They are also shown to predict fuel flow rate and its variability more accurately than other frequently-used methods for fuel flow rate estimation. A sensitivity analysis of the predicted fuel flow rate to takeoff mass is also conducted. The proposed nonparametric models can help set a benchmark for developing more sophisticated data-driven statistical models of engine fuel flow rate.

1 Introduction

The fuel flow rate of an aircraft engine is a key indicator of engine performance. Correct modeling of the fuel flow rate is important both for assessing engine performance, and for estimating aircraft emissions, since emissions are a direct consequence of fuel burn. Knowledge of fuel burn is also essential to estimate the direct operating costs to an airline.

Engine simulation software [1] model the idealized performance of an engine, which can differ from that of a real engine in flight. Moreover, the use of such software requires the knowledge of many internal engine parameters and component characteristics. Prior studies have looked into the problem of modeling aircraft engine fuel burn

[2, 3], but have typically used non-operational data (for example, from flight manuals, software or ground tests) to model the fuel burn. These studies have therefore not accounted for operational factors. For example, the International Civil Aviation Organization (ICAO) Aircraft Engine Emissions Databank (hereafter referred to as the ICAO Databank) provides values of the fuel flow rate for different aircraft engines in the Landing and Take Off (LTO) cycle based on ground tests [4], but these values have been shown to differ from those seen in actual operations [5, 6]. Moreover, methods based on the ICAO Databank and the Base of Aircraft Data (BADA) [7] give point estimates of the fuel flow rate and do not quantify uncertainties in its values. These problems can be overcome by the use of operational data from real flight operations.

There exist relatively few studies which have used operational flight data to model engine fuel burn [8, 9]. The incorporation of different types of operational data has been shown to improve estimates of aircraft fuel burn [10, 11]. Data-driven models of engine fuel flow rate can also give estimates of uncertainties in the fuel flow rate values. With these advantages in mind, in this paper, operational data from Flight Data Recorders (FDR) are used for modeling the engine fuel flow rate. Since, it records values of different aircraft and engine parameters during flight, the FDR provides a reliable archive of real flight data.

An aircraft engine is better modeled as a stochastic (and not a deterministic) system due to random variations in internal characteristics (arising from manufacturing errors, flow turbulence, component deterioration, etc.) and external ambient disturbances (ambient air temperature and pressure fluctuations, wind gusts, etc.)

[12]. These properties motivate the use of machine learning techniques to develop statistical models from aircraft engine data.

Since, all the variables involved in the modeling of fuel flow rate in this paper are metric and continuous, the machine learning problem is fundamentally a regression problem. Two nonparametric machine learning algorithms based on regression trees – the Classification and Regression Trees (CART) algorithm and the Least Squares Boosting (LSB) algorithm – are used to model the engine fuel flow rate in this paper. These algorithms have seen widespread application in diverse areas like solar radiation modeling [13], strength modeling of composites [14], modeling of signalling traffic in mobile networks [15], highway safety studies [16], health monitoring of engineering components [17], pollution studies [18], ecological studies [19], etc.

We start with giving a brief description of the FDR dataset in Sec. 2. The regression variables are explained in Sec. 3. Section 4 briefly explains the CART and the LSB algorithms and their advantages. We then describe the regression methodology employed in this paper in Sec. 5. Section 6 presents the main results of the regression analysis. In Sec. 7, the sensitivity of the analysis to takeoff mass is explored. We wrap the paper up with the main conclusions and directions for future research in Sec. 8.

2 Dataset

Our dataset comprises flight recorder data for 10 different aircraft types, as shown in Tab. 1. Each flight of each aircraft type is split into different phases, based on flight trajectory, velocity, and acceleration parameters [5]. In this paper, only the airborne phases of flight, that is, ascent (from takeoff upto top of climb/start of cruise), cruise, and descent (from end of cruise upto touchdown, including time spent in a holding pattern or loitering, if any), are considered. Prior to the statistical analysis, the fuel flow rate profiles in the dataset were studied as a function of time and altitude, and were found to qualitatively conform to a physical understanding of air-

Table 1. FDR data: aircraft types and engines.

Sr. No.	Aircraft Type	Engine	No. of Flights	No. of Points Across Flights
1.	A319-112	2 × CFMI CFM56-5B6/2 or 2P	130	815–1312
2.	A320-214	2 × CFMI CFM56-5B4/2 or P or 2P	169	753–1526
3.	A321-111	2 × CFMI CFM56-5B1/2 or 2P	117	783–1436
4.	A330-202	2 × GE CF6-80E1A4	84	965–1286
5.	A330-243	2 × RR Trent 772B-60	100	814–1392
6.	A340-541	4 × RR Trent 553	52	871–1532
7.	A340-313	4 × CFMI CFM56-5C4 or 5C4/P	76	974–2230
8.	B767-300	2 × GE CF6-80C2B7F	91	726–1296
9.	B777-3FX(ER)	2 × GE GE90-115B1	131	793–1261
10.	ARJ85/100	4 × LY LF507-1F	153	715–2071

craft and engine performance [20].

3 Regression Variables

The predictor/explanatory/input variables for regression are chosen by considering the underlying principles of flight [21]. These principles suggest that the fuel flow rate primarily depends on the ambient atmospheric density, the true airspeed, the aircraft mass, and the wing reference area. We assume International Standard Atmospheric conditions, no winds aloft, constant wing area for a particular aircraft type, and the takeoff mass to be a surrogate for the instantaneous aircraft mass. Thus, under these simplifying assumptions, the fuel flow rate functionally depends on the aircraft altitude, the ground speed, the vertical speed, and the aircraft takeoff mass. A key benefit of these simplifying assumptions is that with the exception of the aircraft takeoff mass, all of these predictor variables are derivable from aircraft trajectory data alone, which are more readily accessible than FDR data. In our analysis, the speeds are normalized by the design

cruise speed, the aircraft takeoff mass is normalized by its Maximum Take Off Weight (MTOW), and the fuel flow rate is normalized by the ICAO Databank [4] climb out fuel flow rate in ascent and cruise and the ICAO Databank approach fuel flow rate in descent for the appropriate engine type. This normalization is done with the aim of eventually developing generalized, aircraft type independent models of fuel flow rate. We choose the fuel flow rate as a measure of the engine fuel performance in place of another parameter (for example, the Thrust Specific Fuel Consumption (TSFC)) because the fuel flow rate is a directly measured quantity in the FDR dataset. Calculation of TSFC requires a knowledge of thrust values. Thrust is not a directly measured quantity and its values in our dataset are therefore, unreliable. Another point to note is that at the same altitude, speed and takeoff mass values, the value of the fuel flow rate in a particular phase of flight might be different due to secondary factors like turbulence, the environmental control system, component age and deterioration, etc. The effects of such factors cannot be explicitly modeled in this study due to lack of appropriate data in the FDR dataset, and hence, show up in the uncertainty estimates given by our statistical models.

4 Modeling Methods

This section explains the regression trees-based Classification and Regression Trees (CART) and Least Squares Boosting (LSB) algorithms for developing regression models.

4.1 Classification and Regression Trees (CART)

CART is a popular classification and regression algorithm developed by Leo Breiman et al. [22]. Starting with the root node containing all the data points, CART carries out recursive binary splitting of the data. The split criteria are always of the form $X < \chi$ where X is a particular predictor variable and χ is the split point. Points in a node satisfying the split criterion go into the

left child node and the others go into the right child node. This split at each node is a locally optimal split, chosen so as to maximally reduce the weighted sum of the mean squared errors of the resulting nodes. The sum is weighted by the fraction of the observations going in the left child and the right child nodes. The algorithm assigns the mean of the values of the dependent variable in a node as the value of that node. The recursive splitting continues till the algorithm reaches a stopping criterion. The nodes which do not further split are known as leaf nodes. The algorithm stops splitting a node when either the improvement in mean squared error due to a further split drops below a threshold, or when a further split drops the number of observations in a node below a pre-defined value, or when splitting is no longer possible as the predictors have the same distribution for all the points in the node, or when only one observation remains in the node. In this paper, a minimum number of 10 observations is required to be present in each node. In each leaf node, the algorithm assigns a constant value equal to the mean of the dependent variable for all the observations in that node. The tree growing procedure normally produces a very deep tree which overfits the data. Hence, the tree is pruned as a means of regularization. Pruning re-combines leaves (to reduce the total number of leaves and hence, the tree complexity) in a manner which keeps the increase in training mean squared error to a minimum. This successive pruning produces a sequence of subtrees. The mean squared error given by each subtree on out-of-sample data is chosen as the metric of its generalized performance. In this paper, 10 fold cross validation is used to calculate the out-of-sample mean squared error and its standard deviation. The most parsimonious subtree having its mean squared error within one standard deviation of the minimum error across all subtrees is chosen as the final tree to be used further for modeling and prediction of the fuel flow rate. There are several advantages of using CART [23]. CART is a fast and a relatively automatic algorithm, not requiring much intervention from the user. It can be easily scaled to large datasets. The trees gener-

ated are easily interpretable. CART is a nonparametric algorithm (as there are no regression parameters present, unlike, for example, the regression coefficients in ordinary least squares regression). Being a nonparametric method, the problem of having to choose a set of basis functions does not arise. The trees also capture interaction among different predictors (upto the depth of the tree). Regression trees are also robust to outliers in the input data and can deal with irrelevant data. The CART algorithm can easily handle missing values in the data. Lastly, the algorithm is invariant to monotonic transformations of the inputs.

4.2 Least Squares Boosting (LSB)

The CART algorithm generates only one tree, which may be unstable and can have low predictive power. A combination of several such ‘weak’ regression trees can be expected to yield models with better prediction capability. Boosting [23] is an ensemble method which combines several ‘weak’ learners to yield a ‘stronger’ model. In this paper, each weak learner is a regression tree generated by CART (as explained in Sec. 4.1). The boosting algorithm is run for many ensemble cycles. Based on the minimization of the cross validated mean squared error over the training dataset, the number of ensemble cycles in this paper is 100. The algorithm starts by building a CART-based weak learner to the entire dataset. In each successive cycle, the boosting algorithm fits a weak tree learner to the residuals from the previous cycle. The objective is to minimize the squared error loss and hence, the name Least Squares Boosting (LSB). The output predictions at a particular input from all the cycles are linearly combined to give the model prediction at that input. The contribution from each cycle is controlled by a manually set learning rate which, in this paper, is set to a value of 0.1. The learning rate thus, shrinks the contribution from each tree. The advantage of the LSB algorithm lies in its superior predictive ability compared to a single CART model. However, by combining many single CART models, the algorithm loses the interpretability seen in a single tree. The LSB al-

gorithm is also slower than the CART algorithm.

5 Regression Methodology

This section briefly explains the regression methodology employed for model building from the data. It also explains the different metrics used for model evaluation.

5.1 Model Training

For each aircraft type in each of the three phases, the dataset (formed by combining the data in that phase for all the flights of that type) is divided randomly into training and test sets in the ratio of 65:35. The training set is used to build (or train) the models and the test set is used to check the performance of the chosen models on data not seen during model training. For the A320-214, the dataset is randomly divided into training, validation, and test sets in the ratio of 65:25:10. The purpose of the validation set is to provide out-of-sample data to select few best models out of several models trained on the training data. The A320-214 is chosen for validation as it has the most number of flights in the FDR dataset.

To start with, the training dataset is used to build models by using different parametric regression techniques, such as ordinary and robust least squares, and ridge regression. However, the data are found to violate many underlying assumptions of the parametric methods. Hence, nonparametric methods are resorted to in order to model the fuel flow rate. Based on the modeling methods explained in Sec. 4, different CART and LSB models are built for each aircraft type in each of the three airborne phases of ascent, cruise, and descent. The aircraft altitude, h (in metre), the normalized ground speed, V_{GS} , the normalized vertical speed, \dot{h} , and the normalized takeoff mass, m_{TO} are the predictor/explanatory/input variables. The ICAO Data-bank normalized fuel flow rate per engine, \dot{m}_f is the predicted/dependent/output variable. Though all the variables constitute time series in the FDR dataset, in this paper, time series analysis is ignored for simplification and the ordering of the

observations is considered to be unimportant. The regression models trained using CART and LSB are further subjected to evaluation.

5.2 Model Evaluation

The main focus of the research is to build models which can predict well the fuel flow rate values on new data. The following metrics are used for model evaluation on the test dataset to quantify the generalized prediction performance of the models.

- **Mean Error (ME):** This is the mean relative prediction error on the test dataset. It is determined by Equation 1.

$$ME = \frac{1}{n_{test}} \sum_{l=1}^{n_{test}} \left| \frac{\dot{m}_{f_l} - \hat{m}_{f_l}}{\dot{m}_{f_l}} \right| \quad (1)$$

Here, \dot{m}_{f_l} is the actual normalized fuel flow rate value in the test dataset, \hat{m}_{f_l} is the mean predicted normalized fuel flow rate using the trained model, and n_{test} is the number of observations in the test dataset. A model with a lower ME is preferred.

- **Percentage Coverage (PC):** This is the percentage of the observations in the test dataset for which the 95% prediction intervals for the normalized fuel flow rate values include the actual normalized fuel flow rate values. A model with a higher PC is preferred. The 95% prediction intervals are determined by bootstrapping [24].

6 Results

In this section, some of the features of the regression models built, their interpretation, and their prediction performance are discussed. The model predictions on test data are also compared with those given by the ICAO Databank and the BADA methods, two methods commonly used for fuel flow rate estimation.

6.1 Model Interpretation

Figure 1 shows a typical regression tree built by CART and the partitions induced by it in predictor space for an A320-214 in ascent. For ease of representation and interpretability, only 1000 observations randomly chosen from the complete training dataset are used for model building in Fig. 1 (though for prediction, the complete training dataset upto the memory limitation of the software is used). Since, the partitions of the predictor space induced by a tree can be shown easily only for two predictors, the two most important predictors are chosen for tree building here (again, for prediction all the four predictors are used). The choice of the most important predictors is done using predictor importance methods [23]. Though all the models are trained using normalized variables, in Fig. 1 (and also in Fig. 2), the variables are re-converted to their dimensional values for ease of interpretation.

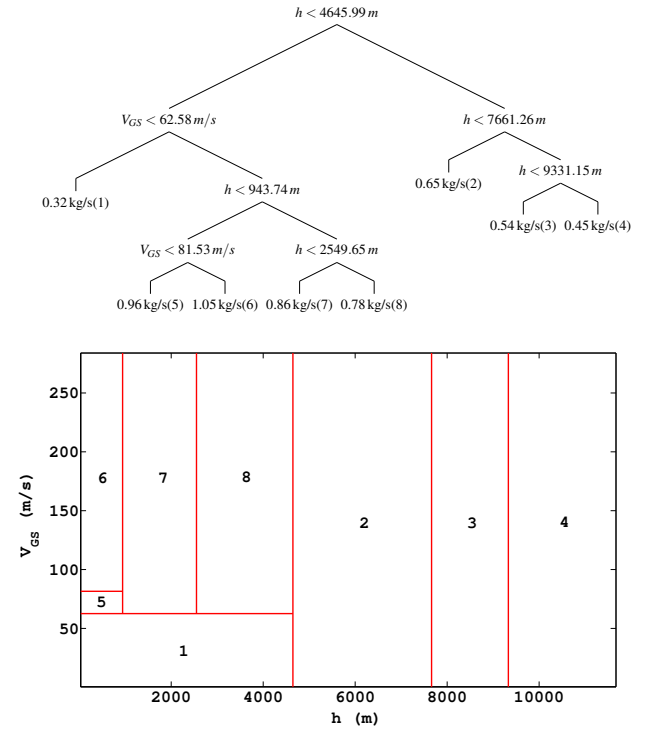


Fig. 1. A320-214: CART tree visualization in ascent: (top) the tree, and (bottom) the partitions of the predictor space induced by the tree. The numbers inside the partitions correspond to the leaf node numbers in the tree (written in parentheses for each leaf node).

The tree clearly shows that prediction is based upon a set of simple decision rules which direct each input to a leaf node and assign a node-

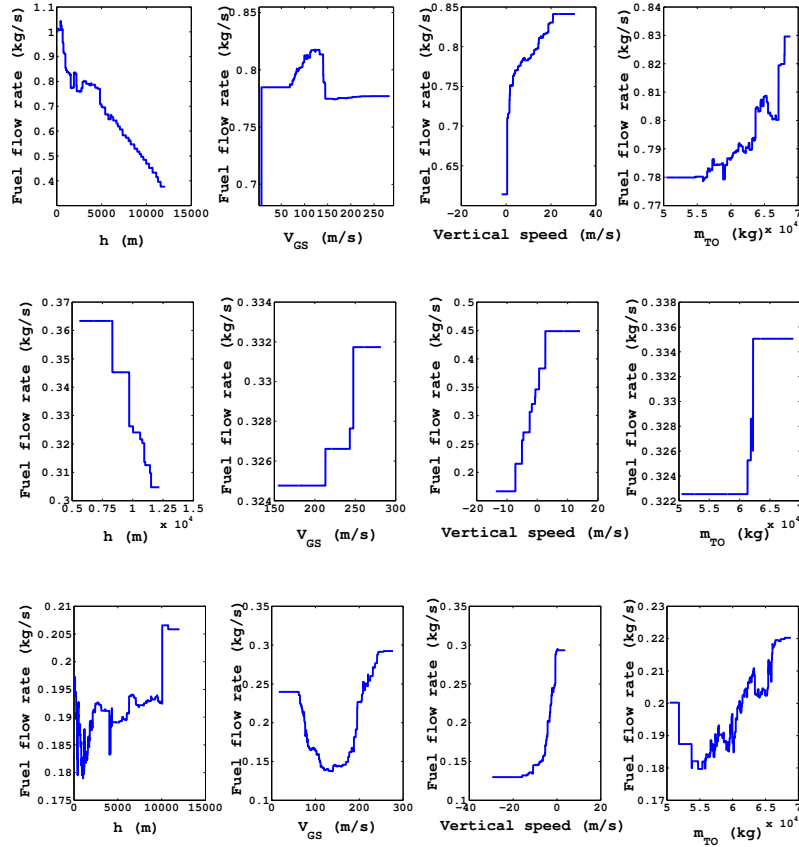


Fig. 2. A320-214: one variable partial dependence plots from the CART models: (top) ascent, (middle) cruise, and (bottom) descent.

specific value for the fuel flow rate. This simplicity enables fast prediction. It is also easy to see how the tree seamlessly captures interactions among predictors. This is evident from the fact that splitting variables seen in the left subtree of a node could be different from those seen in the right subtree. The tree breaks up the predictor space into different partitions, each corresponding to one leaf node. Each partition represents that region of predictor space where all points show homogeneity in fuel flow rate values. Thus, instead of fitting a global model to the entire dataset, CART (and also LSB) identifies homogenous regions of data and fits a separate model in each region. This has the capability of improving model accuracy as well as capturing the variability in data properly. The partitions diagram can also help make quick interpretations about the data. For example, according

to the partitions induced by the particular CART model for the A320-214 in Fig. 1, above about 4650 m (about 15260') altitude during ascent, the fuel flow rate shows no dependency on the ground speed and is solely determined by the altitude. The highest fuel flow rate during ascent is seen for altitudes below about 950 m (about 3120') at ground speeds greater than about 80 m/s (about 290 km/h). The lowest fuel flow rate is seen at ground speeds lower than about 60 m/s (about 220 km/h) and altitudes below about 4650 m (about 15260').

Figure 2 shows one variable partial dependence plots [23] developed using CART for the A320-214 in ascent, cruise, and descent. Partial dependence plots show the approximate dependence of the prediction function on each predictor variable, with the effect of all the other predictor variables averaged out (and not ignored). These

plots can be used to quickly and approximately interpret the models and find interesting trends in the predictions. The plots from the CART models built for the A320-214 here show that in ascent and cruise, the fuel flow rate decreases with altitude. This is expected as the drop in density with altitude reduces the engine thrust requirement. During descent, as the aircraft drops in altitude, the fuel flow rate also decreases. However, it increases at low altitudes to compensate for the increase in drag arising out of the deployment of auxiliary devices (flaps, slats) and the landing gear close to landing. The fuel flow rate is also seen to increase with the vertical speed as more energy is typically required for higher speeds. The fuel flow rate also increases with the takeoff mass as more thrust is needed to maintain a heavier aircraft in flight. Two interesting features are observed in the dependence of the fuel flow rate on ground speed. During ascent, the fuel flow rate maximizes around 125 m/s (450 km/h) ground speed. During descent, the fuel flow rate minimizes around 150 m/s (540 km/h) ground speed. Such knowledge derived from the partial dependence plots about the behavior of the fuel flow rate with each predictor variable as well as from the partition plots (Fig. 1) about the different regimes within each phase can help plan a flight trajectory from a fuel flow rate perspective. It should be noted that single variable partial dependence plots give only approximate interpretation as they do not capture interactions among different predictor variables. Another important point to note is that the plots do not appear to be smooth. This is because CART fits piecewise constant models which are discontinuous at the boundaries of the tree partitions (leaf nodes).

6.2 Model Prediction

The CART and the LSB models are used to predict the fuel flow rates for the test data. These predictions are compared with the target values in the test dataset using the metrics listed in Sec. 5.2. Table 2 shows the different prediction metrics for the models as ranges over the different

aircraft types. As one would expect, the LSB models give a better prediction performance (as indicated by the lower ME and higher PC) than the single tree CART models in all the airborne phases. However, the computational time for model building is higher for LSB as compared to CART. The descent phase has the largest mean error among all the phases. The lower prediction accuracy in descent could arise from the fact that the descent phase shows more variability compared to the other airborne phases. Part of this variability could be due to operational reasons as descent is a controlled phase.

6.3 Comparison with the ICAO Databank and the BADA Method

The ICAO Databank [4] maintains values of the fuel flow rate and emission indices for every certified engine at four thrust settings - take-off (100% thrust), climb out (85% thrust), approach (30% thrust), and idle (7% thrust). It is widely used for fuel burn and emission inventories. Since, this paper deals with only the airborne phases of flight, we compare model predicted fuel flow rate values in climb out and approach with those tabulated in the ICAO Databank. Climb out and approach are the part of ascent and descent, respectively, below 3000' Above Ground Level (AGL). The ICAO Databank enumerates values of fuel flow rates through ground tests done at sea level static conditions on an uninstalled engine. Hence, to do a comparison with the ICAO Databank values, the at-altitude fuel flow rate values (that is, the values at the actual flight altitudes in the FDR dataset) for an installed engine predicted by the regression models in the climb out and approach phases are converted into equivalent values at sea level static conditions for an uninstalled engine. This conversion is done using the equations of the Boeing Fuel Flow Method 2 (BFFM2) [5]. The Base of Aircraft Data (BADA) method [7] is another popular method used to determine aircraft and engine performance parameters, including the fuel flow rate. It is a total energy based method. It uses empirical equations to determine performance. The

Table 2. Model predictive performance on the test dataset across the different aircraft types and comparison with ICAO Databank and BADA values. It is important to note that the model results in ascent (takeoff to start of cruise), cruise, descent (end of cruise to touchdown), and the BADA values are the at-altitude results for an installed engine. However, the model results in climb out (takeoff to 3000' AGL in ascent) and approach (3000' AGL in descent to touchdown) have been converted to sea level static conditions for an uninstalled engine for comparison with the ICAO Databank results. The BADA results range over only 5 aircraft types. The model and the ICAO Databank results range over all the 10 aircraft types.

	CART		LSB		ICAO Databank		BADA	
	ME, %	PC	ME, %	PC	ME, %	PC	ME, %	PC
Ascent	1.4–4.5	50.3–60.3	0.7–2.5	67.6–77.0	–	–	5.9–22.4	0
Climb-out	1.0–4.8	55.3–63.8	0.3–2.5	69.8–79.1	6.2–33.8	0	–	–
Cruise	2.8–8.2	49.4–62.4	2.0–6.3	58.8–66.5	–	–	12.4–130.6	0
Descent	12.1–20.1	50.9–59.8	6.8–13.6	61.8–70.3	–	–	31.8–60.4	0
Approach	13.7–20.5	55.0–61.2	6.5–12.8	63.4–70.9	35.0–96.3	0	–	–

constants in the equations are enumerated in a database. In this paper, the comparison between model predictions and BADA results is made for five aircraft types: A321-111, A330-243, A340-541, A340-313, and Avro RJ85/100 (as these are the only aircraft types for which the engines in the FDR and the BADA databases match). Unlike the ICAO Databank, the BADA method estimates the fuel flow rates for at-altitude conditions itself. The comparisons show that the CART and the LSB models predict the test dataset fuel flow rates more accurately as compared to both the ICAO Databank and the BADA values in all the phases. More importantly, both the ICAO Databank and the BADA method give only point (and not interval) estimates of the fuel flow rates, thereby giving a percentage coverage of 0. This problem is overcome by the CART and the LSB models which give bootstrapped prediction intervals, resulting in a higher percentage coverage. Hence, fuel flow rate modeled by using operational data from several flights is more reflective of the variability observed in real flight conditions as compared to methods like the ICAO Databank or the BADA method.

7 Model Sensitivity to Takeoff Mass

Out of the four predictors being used to build the models, altitude, ground speed, and vertical speed can be obtained/derived from trajectory data which are easily accessible. Hence, accurate values of these three predictors can be known from trajectory data. However, it is dif-

ficult to know the true value of the takeoff mass for a flight. Hence, one would normally use an estimate of the takeoff mass for building the regression models. Thus, it is imperative to know how sensitive the models are to errors in the estimated takeoff mass. For this analysis, the values of the takeoff mass in the test dataset for each aircraft type are changed systematically by a fixed percentage of the true value to give a modified test dataset. All the other predictors are held at their original values (which is an approximation as the other variables might depend on takeoff mass too). Models trained using the true takeoff mass are run on the modified test dataset, and the percentage mean error in the predicted fuel flow rate is calculated. The variation of this error with the percentage deviation of the estimated takeoff mass in the test dataset from its true value indicates the sensitivity of the model predictions to the takeoff mass. Table 3 shows the percent increase in percentage mean error on the modified test dataset when there is a 1% deviation (in either direction) in the takeoff mass about its true value. The behavior of percentage mean error with deviation in takeoff mass for the B767-300 is shown in Fig. 3 for cruise. Other phases and aircraft types show similar behaviors.

As expected, the mean error is the minimum when the takeoff mass values in the modified test dataset equal the true values. The mean error increases as the magnitude of deviation from the true takeoff mass increases and stabilizes for large magnitudes of deviation in the takeoff mass. The model predictions are seen to be sensitive to

Table 3. Percent increase in percentage mean error in predicted fuel flow rate for a $\pm 1\%$ deviation in takeoff mass from its true value across the different aircraft types. The results in this table are from the LSB method.

	-1% Deviation	+1% Deviation
Ascent	57.3–223.2	57.1–213.4
Cruise	4.3–111.3	9.2–91.1
Descent	40.2–126.7	41.8–129.9




Fig. 3. B767-300: effect of deviation in takeoff mass on percentage mean error in predicted fuel flow rate on the modified test dataset in ascent. The results are from the LSB model.

even a 1% deviation (in magnitude) from the true takeoff mass.

8 Conclusions

This paper studied the application of machine learning techniques to the modeling of aircraft engine fuel flows. Two nonparametric methods – Classification and Regression Trees (CART), and Least Squares Boosting (LSB) – were employed. A physical understanding of the principles of flight helped in the selection of variables for regression. The regression methods were tested on real FDR data from 10 different aircraft types. CART gave simple and interpretable models of fuel flow rate, while the LSB models were found to have higher predictive ability than CART models, at the expense of increased complexity and some interpretability. Both methods performed better than the ICAO databank and the BADA methods used commonly for fuel flow rate estimation. The proposed models also provide uncertainty estimates (in the form of prediction intervals) for the fuel flow rates, which quantify the extent of variability seen in real flight operations as well as the effects of unmodeled factors. The sensitivity of the model to takeoff mass was

also evaluated.

The results of this paper can be used as a benchmark for more sophisticated models of fuel flow rate. A possible next step is the use of Bayesian regression trees, whereby a prior is placed on the distribution of trees to account for a priori beliefs about the problem. The advantage of Bayesian methods lies in their ability to directly give a predictive distribution for new data. Another future direction that is currently being investigated is the time-series analysis of flight data.

Acknowledgments

This work was supported in part by the US National Science Foundation under NSF Cyber-Physical Systems project FORCES, grant number 1239054. The views and conclusions contained in this document are those of the authors alone, and not the funding agency.

References

- [1] GasTurb 12 software. <http://gasturb.de/>, [Online; accessed 15 October 2015], 2015.
- [2] Senzig D.A., Fleming G.G., and Iovinelli R.J. Modeling of terminal-area airplane fuel consumption. *AIAA Journal of Aircraft*, Vol. 46, No. 4, pp. 1089–1093, 2009.
- [3] Trani A.A., Wing-Ho F.C., Schilling G., Baik H., and Seshadri A. A neural network model to estimate aircraft fuel consumption. *AIAA 4th Aviation Technology, Integration and Operations (ATIO) Forum*, Chicago, pp. 1–24, 2004.
- [4] International Civil Aviation Organization. ICAO aircraft engine emissions databank. <https://easa.europa.eu/document-library/icao-aircraft-engine-emissions-databank>, [Online; accessed 12 February 2014], 2014.
- [5] Chati Y.S., and Balakrishnan, H. Analysis of aircraft fuel burn and emissions in the landing and take off cycle using operational data. *6th International Conference on Research in Air Transportation*, Istanbul, pp. 1–8, 2014.
- [6] Patterson J., Noel G.J., Senzig D.A., Roof C.J., and Fleming G.G. Analysis of departure and arrival pro-

- files using real-time aircraft data. *AIAA Journal of Aircraft*, Vol. 46, No. 4, pp. 1094–1103, 2009.
- [7] Nuic A. User manual for the Base of Aircraft Data (BADA) revision 3.13. Eurocontrol Experimental Centre Technical/Scientific Report No. 15/04/02-43, 2015.
- [8] Khadilkar H., and Balakrishnan H. Estimation of aircraft taxi fuel burn using flight data recorder archives. *Transportation Research Part D: Transport and Environment*, Vol. 17, No. 7, pp. 532–537, 2012.
- [9] Stolzer A.J. Fuel consumption modeling of a transport category aircraft: a Flight Operations Quality Assurance (FOQA) analysis. *Journal of Air Transportation*, Vol. 8, No. 2, pp. 3–18, 2003.
- [10] Ryerson M.S., Hansen M., and Bonn J. Validating aircraft performance models with airline data. *5th International Conference on Research in Air Transportation*, Berkeley, pp. 1–8, 2012.
- [11] Yoder T. *Development of aircraft fuel burn modeling techniques with applications to global emissions modeling and assessment of the benefits of reduced vertical separation minimums*. Masters Thesis, Massachusetts Institute of Technology, 2007.
- [12] Kulikov G.G., and Thompson H.A. (eds.). *Dynamic modelling of gas turbines: identification, simulation, condition monitoring, and optimal control*. Springer-Verlag, London, 2004.
- [13] Mori H., and Takahashi, A. A data mining method for selecting input variables for forecasting model of global solar radiation. *Transmission and Distribution Conference and Exposition (T D), 2012 IEEE PES*, Orlando, pp. 1–6, 2012.
- [14] Carty D.M. *An analysis of boosted regression trees to predict the strength properties of wood composites*. Masters Thesis, University of Tennessee, 2011.
- [15] Hecker A., and Kürner T. Application of classification and regression trees for paging traffic prediction in LAC planning. *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th*, pp. 874–878, 2007.
- [16] Stewart J.R. Applications of classification and regression tree methods in roadway safety studies. *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 1542, pp. 1–5, 1996.
- [17] Vianna W.O.L., Gomes J.P.P., Galvão R.K.H., Yoneyama T., and Matsuura J.P. Health monitoring of an auxiliary power unit using a classification tree. *Annual Conference of the Prognostics and Health Management Society*, Montréal, pp. 1–7, 2011.
- [18] Cheng W., Zhang X., Wang K., and Dai X. Integrating Classification and Regression Tree (CART) with GIS for assessment of heavy metals pollution. *Environmental Monitoring and Assessment*, Vol. 158, No. 1–4, pp. 419–431, 2009.
- [19] Elith J., Leathwick J.R., and Hastie T. A working guide to boosted regression trees. *Journal of Animal Ecology*, Vol 77, No. 4, pp. 802–813, 2008.
- [20] Chati Y.S., and Balakrishnan H. Aircraft engine performance study using flight data recorder archives. *Aviation Technology, Integration, and Operations Conference, AIAA Aviation*, Los Angeles, pp. 1–12, 2013.
- [21] Collins B.P. Estimation of aircraft fuel consumption. *AIAA Journal of Aircraft*, Vol. 19, No. 11, pp. 969–975, 1982.
- [22] Breiman L., Friedman J.H., Olshen R.A., and Stone C.J. *Classification and regression trees*. CRC Press, New York, 1993.
- [23] Hastie T., Tibshirani R., and Friedman J. *The elements of statistical learning: data mining, inference, and prediction*. Springer-Verlag, 2009.
- [24] Efron B., and Tibshirani R.J. *An introduction to the bootstrap*. Springer - Science + Business Media, B.V., 1993.

Contact Author Email Address

mail to: yschat@mit.edu

Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.