SMU Data Science Review

Volume 1 | Number 2 Article 11

2018

Fuel Flow Reduction Impact Analysis of Drag Reducing Film Applied to Aircraft Wings

Damon Resnick
Southern Methodist University, daresnick@yahoo.com

Chris Donlan
Southern Methodist University, cdonlan@smu.edu

Nimish Sakalle
Southern Methodist University, nsakalle@smu.edu

Cody Pinkerman
Edge Aerodynamix, cpinkerman@edge.aero

Fuel Flow Reduction Impact Analysis of Drag Reducing Film Applied to Aircraft

Damon Resnick[1], Chris Donlan[1], Nimish Sakalle[1], Cody Pinkerman[2]

[1] Southern Methodist University, Department of Data Science,
6425 Boaz Lane, Dallas, TX 75205
{dresnick, cdonlan, nsakalle}@smu.edu

[2] Edge Aerodynamix, 415 Richard Jackson Blvd,
Panama City Beach, FL 32407
cpinkerman@edge.aero

Abstract. In this paper, we present an analysis of flight data to help better determine whether the application of the Edge Aerodynamix Conformal Vortex Generator (CVG), applied to the wings of aircraft, reduces fuel flow during cruising conditions of flight. The CVG is a special treatment and film applied to the wings of an aircraft to protect the wings and reduce the non-laminar flow of air around the wings during flight. It is thought that by reducing the non-laminar flow or vortices around and directly behind the wings that an aircraft will move more smoothly through the air and provide a safer and more fuel-efficient flight. Analysis was performed on over a year's worth of flight data collected from four different aircraft corresponding to nearly 100,000 total flight hours. The analyses presented in this paper are intended to add to and complement the previous analysis done by Edge Aerodynamix. It focuses on modeling the data using a multiple linear regression technique. The results from this analysis compare well with the results previously obtained by Edge Aerodynamix. On average, the difference in fuel flow achieved using multiple linear regression were similar to the results achieved by Edge Aerodynamix. We found that the model to best fit the data was the simple multiple linear regression (MLR) model while the model to produce the best results the random forest regression (RFR) model.

1 Introduction

Even though fuel efficiencies have improved substantially in the past decade for large aircraft, the increase in the size of the airline industry has outpaced these effects and the overall contributions of aircraft to greenhouse gas emissions has increased. Improvements in efficient traffic patterns, hybrid aircraft, and aerodynamics have helped to reduce fuel consumption for single aircraft, but they are unable to keep up with the growth of passenger air travel and therefore other methods should be explored.

In this paper we analyze flight data to help better determine the effectiveness of the CVG product for increasing fuel efficiency. This method hinges on the application of a patented CVG product developed by Edge Aerodynamix (EA) [1]. The CVG is a special elastomeric adhesive backed film that can be applied to the wing of an aircraft to help reduce the drag on the aircraft by producing a series of smaller counter-rotating vortices behind the wing.

Our focus in this paper is to extend EA's analysis by introducing additional statistical methods such as multiple linear regression, dimensionality reduction, classification machine learning algorithms, random forest regression, and gradient boosting regression. We also focused on finding completely new methods for quantifying the product's effectiveness in reducing fuel flow.

Our analysis used data provided by EA's airline partners. Data from Boeing 737s, both classic and next-generation variants, were used in the analysis. Raw data from the aircraft's flight data recorders were found to be too noisy for use in our analysis. The data was sampled to find periods of flight when the aircraft was in stable cruise conditions. These periods consisted of window sizes, ranging from 3-15 minutes, when the aircraft was operating in stable flight at cruise altitudes and within certain thresholds for indicators such as speed, pitch, heading, wind direction and velocity, temperature, and altitude. As the variance within the original data was too large and varied between individual flights, these cruise samples where vital in factoring out portions of a flight that commonly vary between flights like taxi, takeoff, departure, approach, and landing. All the analysis in this paper was done on these cruise samples.

The analysis performed on this data in this paper is an extension of an analysis performed by EA. One key component of EA's work is dimensional analysis. Dimensional analysis is a method for reducing the number and complexity of experimental variables that affect a given physical phenomenon, by using a compacting technique [2]. As we are dealing with a physical object, attributes of key features share some dependency with one another. This dependency is due to the basic principle of flight, which includes: thrust, lift, weight, and drag. The relationship these features have with each other, provide insight into how data between flights can be and modeled.

EA used a common fluid dynamics technique involving dimensional analysis to relate features such as the weight, pressure, speed, and temperature, to the rate of fuel-flow to the engines of an aircraft. Combining these features in a unique way the fuel-flow can be modeled with simple linear regression. This model was used to calculate the difference in fuel-flow between flights within similar operating conditions and at various power settings. Finding other accurate and meaningful ways of calculating this fuel flow difference was the primary task for this paper.

The use of multiple linear regression is a key differentiator between our analysis and EA's dimensional analysis. Multiple linear regression allows the variables to operate independent of each, allowing for additional flexibility when it comes to feature selection. Similar to EA's dimension analysis, transformations were performed on features used in the analysis. The dimensional analysis performed by EA, and in turn the transformations performed on the data for our analysis, are proprietary and not explained in detail. The term *Phi*, ϕ, is used to represent this combination of these features into a single interaction feature. Multiple linear regression models using *Phi* as an interaction term along with the other transformed variables are discussed in later section. Models including the *Phi* term, along with aircraft power setting, produced the best fitting models.

Cross validation, using a *k*-fold technique, was used to help with feature selection during the development of the multiple linear regression model. By holding data out in the training phase, we were able to reduce the risk of over-fitting the model. With *k*-fold validation, the original data was split into *k* number of folds. Each of these test sets were taken from different parts of the data so that the entire data set could be used to both train and test the data *k* different ways. To calculate the difference in fuel flow between two flights, each of the cruise samples were split into multiple folds and each fold was modeled. Difference in fuel-flow was calculated between the two data sets.

In addition to linear modeling, ensemble learning methods, such as random forest regression and gradient boosting regression were also used. These ensemble techniques utilize random decision trees to fit the data. By fitting multiple models and combining techniques, these models are less likely to overfit the data. To further reduce overfitting, cross validation was also used to compare these models with the linear models. These ensemble models produced more favorable results, less fuel consumption, than EA's original dimensional analysis and our multiple linear regression models but don't appear to be quite as accurate as the multiple linear regression (MLR) models.

Beyond the transformations performed on the data to mimic the original dimensional analysis, we also performed additional logarithmic and exponential transformations to better fit the data. The use of the aircraft power setting was also incorporated into the models. The original EA dimensional analysis failed to do this.

Details on model performance for each aircraft can be found in Table 4. The table shows the performance of all the models and shows that the percentage difference in fuel flow using the random forest regression model produced the largest difference -0.95%.

In addition to evaluating the difference in fuel flow, the fitness of models was also compared with root-mean-square error (RMSE) and mean absolute error (MAE). Based on these metrics, multiple linear regression inclusive of operating condition term performed the best. A comparison summary of all models used is found in Table 4.

2 Economic and Environmental Impact

In order to understand the societal and industry impacts of the CVG we will first discuss some of the details of the aviation industry. The aviation industry accounts for 3.5% or $2.7 trillion of the world's gross domestic product. Specifically, in the U.S., civil aviation accounts for more than 5.4% of the gross domestic product, or $1.5 trillion. Airlines make up a significant portion of the aviation industry, and each year airliners move over three billion people and 50 million tons of cargo. It's estimated that the aviation industry directly creates 9.9 million jobs.[1]

Beyond the direct contribution of the many airlines to the world's economy, the indirect benefit, albeit hard to measure, of fast long-distance travel plays a significant role in the modern business world. According to the U.S. Bureau of Transportation Statistics, Americans take over 405 million long-distance business trips per year, accounting for 16% of the world's long-distance travel. Even with technological advances in communication tools, the large percentage of business travel indicates the

[1] ATAG Inc., "Value to the economy", https://aviationbenefits.org/economic-growth/value-to-the-economy

importance and necessity of the person to person meeting.[2] This network of connectivity provided by airlines is a significant contribution to a dynamic infrastructure, that keeps the US in particular, along with other nations, competitive on the global stage.[2]

3 CVG Product

The EA CVG is designed to increase fuel efficiency on transport category aircraft by reducing aerodynamic drag on the upper surface of the aircraft wing. The CVG is a thin strip of adhesive-backed PVC material that is applied to the wing just aft of the slat trailing edge. The CVG runs the full length of the slat outboard of the engine nacelle. The device uses a patented shape of its leading and trailing edges to duct high-energy airflow into the boundary layer as a series of small counter-rotating vortices.

4 Data and Data Capture

Flight data for the analysis was provided by two of EA's airline partners. Throughout this study, these airlines will be referred to as Airline 1 and Airline 2, to keep the names of these companies private as the partnership does not allow for the names to be published.

Table 1. Sample of flight data attributes. Names of a few selected features as well as a description of them is provided.

Name	Description
Mach	Mach number, which is the ratio of True Airspeed and speed of sound.
Engine N1 Setting	RPM of the Low-Pressure Shaft
Engine N2 Setting	RPM of the High-Pressure Turbine Shaft
Altitude	Flight Altitude (ft)
Total Air Temperature (TAT)	Temperature if the air was brought adiabatically to rest (K)
True Airspeed (TAS)	Averaged measured True Airspeed (Kt) during a cruise duration
Pitch	Measured Aircraft pitch angle (degrees)
Weight	Gross weight of an Airplane (kg)

[2] Bureau of Transportation Statistics, "America On The Go", https://www.bts.gov/archive/publications/america_on_the_go/long_distance_transportation_patterns/entire

Flight data recorders interpret and record those flight parameters impacting the aircraft's flight and engines for the entire duration of the flight [3]. There can be upward of 70 – 80 different parameters in these data sets depending on the aircraft. These parameters are recorded every second of the flight and the flight data recorders are generally turned on before taxing and turned off after landing.

For example, Airline 1 provides a data file spanning several weeks of flights and maintenance activity, but the data has been converted to the appropriate units and is a selected subset of the flight data recorder parameters. Airline 1 uses a miniQAR MKII made by Avionica and the Avionica Avscan ground station software to provide data converted to the appropriate units [4].

The time interval size and evaluation criteria are clearly defined in the raw data. Currently four different time interval sizes are available: 3-minute, 5-minute, 10-minute and 15-minute. Since the data was recorded every second, in an average 2-hour flight a flight data recorder produces 7200 data points (one data point for every second) for 70 to 80 different parameters [3].

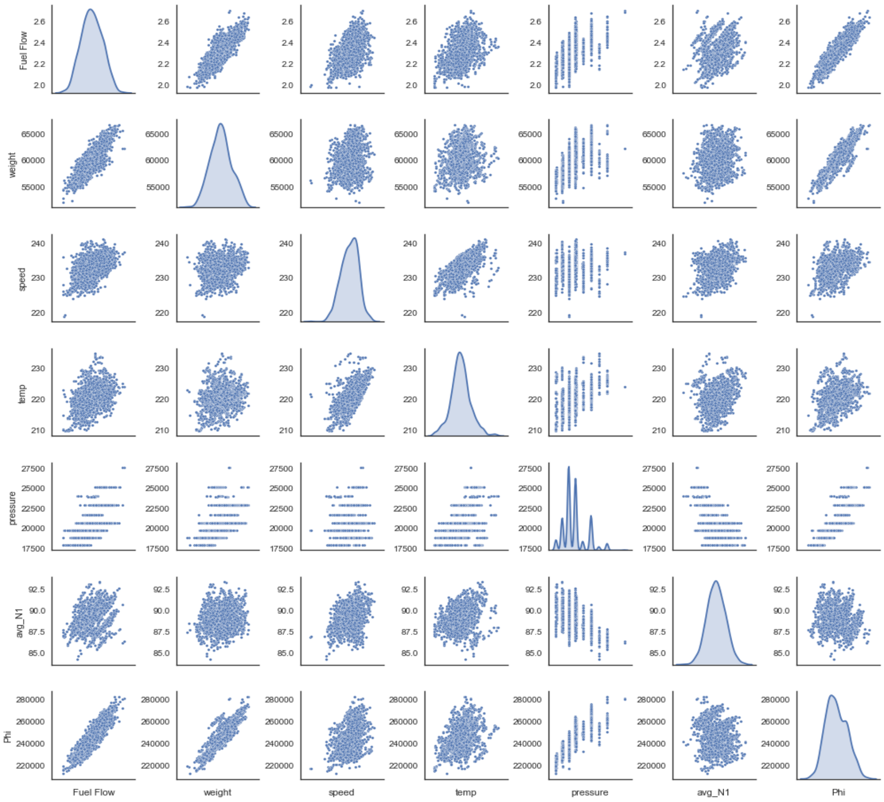

Fig. 1. Matrix of scatter plots for most of the relevant features used in the analysis. Most have very normal distributions and are linearly proportional to fuel flow.

Airline 2 provided a separate data file for each flight. Each Airline 2 data file has been converted to the appropriate units and is a selected subset of the flight data recorder parameters and has been scrubbed for data errors. Airline 2 contracted with Flight Data Services (FDS) to manage their flight data recorder data using Flight Data Monitoring (FDM) program (also known as the Flight Data Analysis Program) to identify, quantify, assess and address operational risks [5]. It uses a systematic, pro-active and non-punitive approach to digital flight data from routine operations to improve aviation safety. Airline 2 used a FDS provided data transfer station to upload the data to the FDS data server.

Data was obtained from over a hundred different flights with a year or so window before and after the CVG was applied to four different aircraft. For the analysis performed and discussed in this paper EA provided data files from both Airline 1 and Airline 2, from before and after application of the CVG.

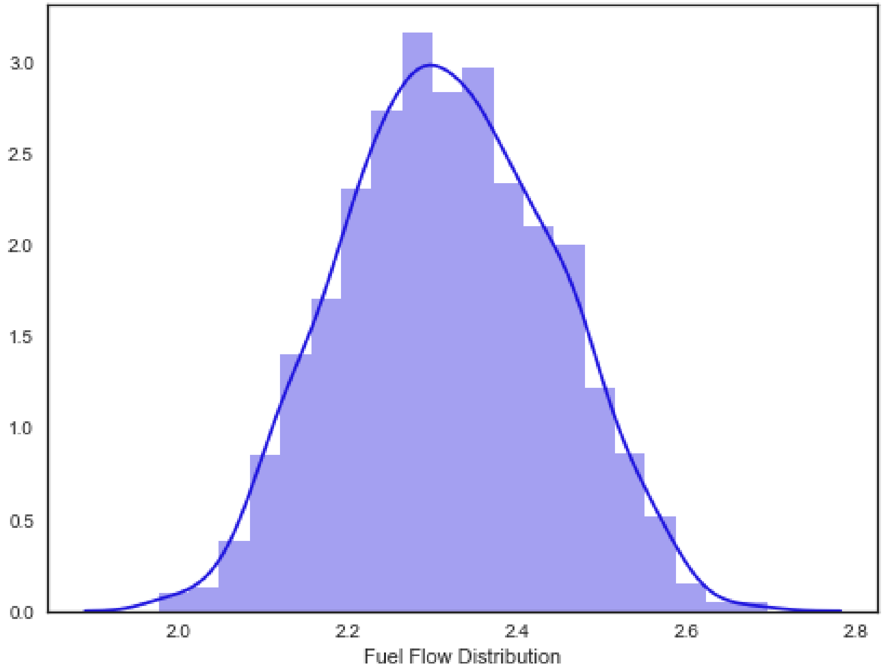

Fig. 2. The distribution of Fuel Flow from flight steady state data for three-minute sliding windows. This shows that Fuel Flow has a normal distribution for three-minute sliding windows for numerous flights after application of the CVG product.

A subset of the attributes captured by the flight data recorder and used in the analysis are found in Table 1. Mach number is calculated from airspeed. The engine RPMs of both engines of a 737 are recorded, as well as the RPMs from the low and high-pressure turbine shafts. Altitude is also recorded and with other quantities used to calculate pressure and temperature.

Data exploration analysis is performed on steady state data to check data quality, data distribution, data features, linear relationships with fuel flow and multicollinearity checks to identify features selected for analysis.

Scatter matrix plot Fig. 1 on selected features shows data distribution and correlation among attributes. Data distribution of response feature Fuel Flow and other exploratory features is normal, apart from Pressure. Log transformation was applied on pressure to slightly improve normality in some case but was not used in the final model.

Most of exploratory features Weight, Speed, Temperature, Phi have strong linear relationship with response feature Fuel Flow. Phi is function of Weight, Speed, Pressure and Temperature.

Most of the features have normal distributions. Fig. 2 shows this for the Fuel Flow, however a log transformation was applied to the pressure variable to bring it closer to normality in some cases. Fig. 3 shows the Fuel flow versus the weight of the aircraft.

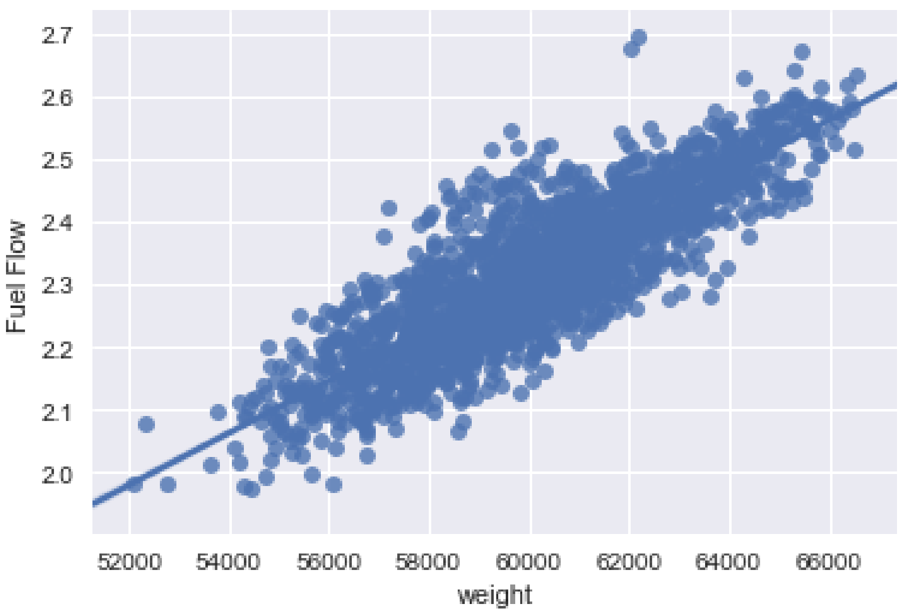

Fig. 3. The Fuel Flow is plotted versus the weight of AC-001. This graph clearly shows the linear nature of the relationship between the weight of the aircraft and the fuel flow.

The Pearson Correlation coefficient is a measure of the strength of the linear relationship between two variables as well as a check for multicollinearity. Multicollinearity can increase the variance of the coefficient estimates and make the estimates more sensitive to minor changes in the model. Fig. 4 shows a heatmap with correlation coefficients.

Features such as Weight, Pressure, Power and Phi have strong correlation and have linear relationship with response feature Fuel Flow. This linear relationship is also shown in Fig. 1.

Fig. 4. Correlation coefficients of selected features. The function Phi has a high correlation with Weight, Pressure, and Fuel flow.

Other features in datafile have a high collinearity like Fuel Flow, Engine 1 N1 RPM and Engine 2 N1 RPM, Stabilizer, True Air Temperature, MACH and Pitch. These features are excluded from the analysis.

The summary statistics for some of the most important features used in our models are provided in Table 2. Most of the features have normal distributions and no extreme outliers. They have fairly even distributions but not all values are clustered around the mean values. This is one of the reasons an advanced modeling technique is needed to determine the difference between the fuel flow. A simple statistical test would not be powerful enough to determine if the means of the fuel flow pre and post treatment are significantly different.

Table 2. Summary statistics of features selected from flight steady state data for three-minute sliding window on the data from the Airline 1 before the CVG was applied.

Feature	Mean	STD Dev	Min	Max
Fuel Flow	2.316866	0.123461	1.976131	2.69446
Weight	60110.33	2523.653	52082.57	66564.49
Speed	232.7033	3.052487	218.6499	241.1515
Temperature	220.0188	3.772163	209.5371	234.4683
Pressure	20428.3	1574.466	17826.67	27514.03
Avg N	90.61706	1.025013	87.17865	94.05486
Phi	245048.8	12217.58	211892.1	282428.2

5 Modeling Methods

The engineers at Edge Aerodynamix have performed an initial analysis that involves the dimensional analysis of relevant flight variables as an aircraft flies during stringent cruising conditions. This dimensional analysis is used to analyze fuel flow in different stable cruise conditions after the installation of their CVG product. The analysis is based on data recorded by flight data recorders during the duration of hundreds of flights both pre and post product application.

The type of dimensional analysis EA used is a modified fluid dynamics technique that works well during the stable cruise operation of an aircraft [1]. It estimates the fuel flow efficiency only during this narrow type of cruise performance. It should be noted that for typical flights the cruising portion of the flight utilizes only roughly half of the total fuel during a flight and that there is no reason to assume that the CVG is only suited to help increase fuel efficiency during this portion of a flight. In fact, based on the underlying physics of the CVG it may be that the CVG could have an even greater effect on the drag on the aircraft during the non-cruising portions of the flights. It is expected that the CVG should produce an even greater increase in fuel efficiency during more dynamic portions of a flight including the climbing portions of a flight where the most fuel is expended. However, due to the variance and noise associated with non-cruise data, both EA's and our analysis will be performed only on sampled cruise data.

5.1 Edge Aerodynamix Dimensional Analysis

A dimensional analysis approach works well to model flight parameters during steady state conditions like the stable cruising portions of a flight [1]. Because of this, stringent requirements are needed to determine if a continuous portion of the flight data occurred during cruise. Portions of flight not in cruise contain excessive variance and noise, in turn, producing inconsistent modeling results. To take this into account a data selection is performed using a sliding window function that evaluates the flight data for stable cruise conditions. Standard algorithms are used that automatically identifies and extracts steady-state engine operating points from engine flight data. The algorithms calculate the mean and standard deviation of select parameters contained in the incoming flight data stream. If the standard deviation of the data falls below defined constraints, the engine is then assumed to be at a steady state of operation, and the mean measurement data at that point are archived for subsequent condition monitoring purposes.

5.1.1 Details of the Dimensional Analysis

The details of the dimensional analysis are outlined in a private EA report from 2017 [1]. We will summarize the methods and results of that report briefly; however, some of the details in the report cannot be outlined due to the private nature of the report. The goal is to summarize the report to put it into context and to help give a better understanding of the additional complementary analyses in this paper.

A dimensional analysis technique was used to derive a formula that relates the fuel flow rate of the aircraft to the fuel efficiency multiplied by certain operating parameters of the aircraft in flight [1]. This relationship between these parameters depends on the assumption that the fuel flow is a function of speed, mass, and flight altitude at a constant power setting (N1) of the engine. This is a standard fluid dynamics technique that requires a steady state [2]. This method also equates altitude to atmospheric density and assumes that the aircraft is in steady and level cruise operation.

This EA dimensional analysis results in four parameters having dimensions of mass (weight of the aircraft, M), length (L), and time (t). The fuel flow rate has dimensions of M/t, density as M/L^3, while speed (U) has dimensions of L/t.

The parameters are then combined and transformed in a way to dimensionally match the fuel flow to the engines. The relationship of these four parameters produced a one term scaling law that, by definition, must be equal to a constant at a constant power setting (N1), where the constant is C_{ff}, the fuel flow coefficient. Using the ideal gas law and solving for the fuel flow (*FF)*, the formula becomes a simple relation between W_t (weight of the aircraft), P (air pressure), U (speed of the aircraft), T (air temperature), and R (the ideal gas constant).

The relationship between the terms is proprietary to Edge Aerodynamix. The function called *Phi* (ϕ), is used to represent this relationship. Equation 1 illustrates the complete relationship between *FF*, the fuel flow coefficient (C_{ff}), and the *Phi* function.

$$FF = C_{ff} \cdot \phi(W_t, P, U, T) \tag{1}$$

The fuel flow coefficient C_{ff} is then determined by plotting the fuel flow rate versus these operating conditions and finding the slope of the line that best fits the data from a single aircraft over multiple flights and multiple cruising windows. This coefficient is directly proportional to *FF* (i.e. a 1% increase in C_{ff} results in a 1% increase in *FF*).

Using this formula, one can compare the fuel flow for aircrafts before and after the application of the CVG [1]. The decrease of needed fuel to keep an aircraft at a constant power has been determined to be anywhere between 0.8% to 1.22%, depending on the model of aircraft [1]. The EA dimensional analysis and reported results also take into account the degradation of flight performance over time and other maintenance effects on fuel efficiency.

5.2 Simple Linear Regression

Based on the initial analysis and current understanding of the data several different analyses were performed. The main commonality of the different analyses is the application of techniques that are easily reproducible by anyone with access to similar data. Towards this end our analyses will be done using simple python scripts in a Jupyter notebook. The full code and output can be provided with a request to Edge Aerodynamix. Basic details and results are provided and summarized below.

Our initial analysis consists of taking the data from all flights from both pre and post CVG application and comparing the average total fuel per second used. A simple regression technique was used to fit the data. It should be noted that even when using a normalizing metric, the data for the individual weights of each aircraft and the data for the weather and wind patterns throughout each flight are not easily available. This makes it difficult to make comparisons between the pre and post applications in most cases.

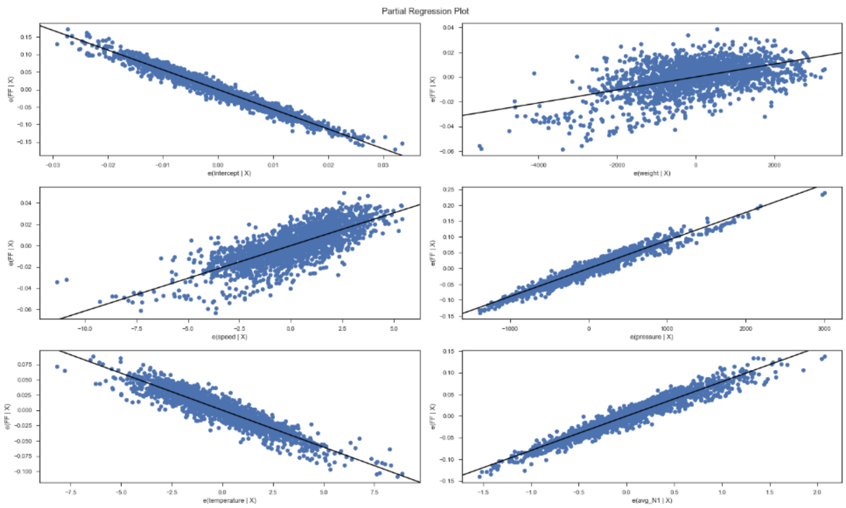

Fig. 5. Partial regression plots for each parameter in the model shown in Eq. 2 for AC-101 before the CVG was applied.

This first basic analysis employs a simple linear regression of the raw data by examining the relationship of a rough estimation of distance traveled for an aircraft. The average speed during basic cruise was multiplied by the total time of over 100 flights to give a rough estimate of the distance traveled. This distance was then plotted versus the fuel used during those periods before and after the application of the CVG, and then linear regression was used to estimate the slope and intercepts of the linear fits. The results of this analysis proved to be inconclusive and misleading.

It was found that using simple metrics that make sense for automobiles, like miles per gallon, don't make sense for aircraft data like that used in these analyses. The variation in aircraft weight and weather changes make it too difficult to apply a simple miles per gallon metric. This leads to a more complementary model to the dimensional analysis instead of a completely different one.

5.3 Multiple Linear Regression

A MLR technique was also used to analyze the data. This method is significantly different than the dimensional analysis approach as it assumes that the variables are marginally independent of each other in the narrow cruising range. This is significant as it allows a decoupling of the variables so that a change in fuel flow can be determined for different ranges of specific variables.

The specific MLR model used to fit FF. In this analysis consists of the four variables used in the dimensional analysis: W, U, P, T, and N.

$$FF = \beta_0 + \beta_1 W + \beta_2 U + \beta_3 P + \beta_4 T + \beta_5 N \quad (2)$$

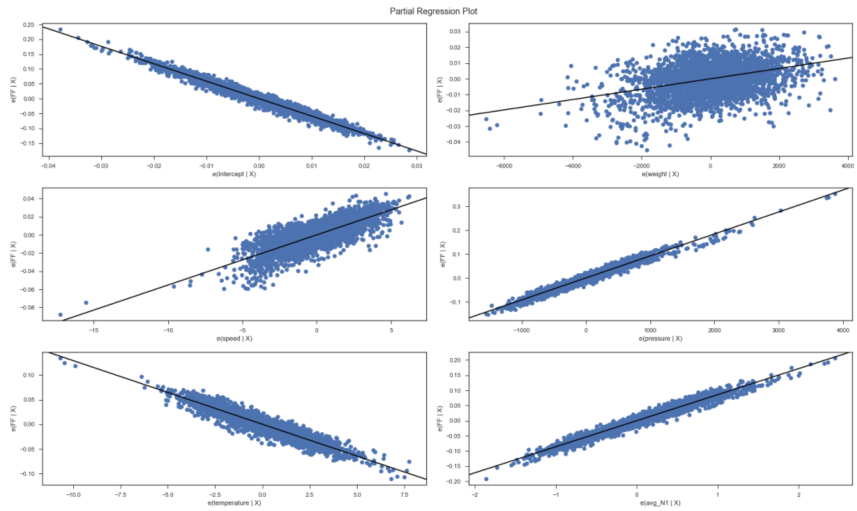

Fig. 6. Partial regression plots for each parameter in the model shown in Eq. 2 for AC-101 after the CVG was applied.

The data includes the values for the engine power of both engines of the 737 aircraft as well as the power for the low and high-pressure shafts of the engine. N1 corresponds to the power setting of the low-pressure shaft while N2 for the high-pressure shaft. In this analysis, we use the average N1 between both engines as the power variable.

The chosen model in Eq. 2 was used to fit the data for all four aircraft. This model was then used to predict the *FF* for the average value of each variable over the range of values of each variable for both data sets. Only the parts of the data that overlapped were used. This was done so that only the range of data both sets had in common were compared. Fig. 5 shows the relative fits for each variable on AC-001 before the CVG was applied, while Fig. 6 shows the relationship after application. Considering the linear shape of the partial regression plots a linear model is appropriate to fit the data. The figure also gives you an idea of the range over which the data is fit for each variable. Fig.6 shows the partial regression plots after the CVG was applied.

The model in Eq. 2 was applied to all four aircraft and a percentage change in fuel flow was determined for each aircraft. Table 3 details the percentage change in *FF* from before and after the CVG was applied for the simple dimensional analysis and the MLR. DA stands for percentages obtained with the dimensional analysis approach while the MLR column summarizes the values obtained using Eq. 2.

Table 3: Summary of the percentage change in fuel flow. DA refers to the dimensional analysis while MLR for the multi-linear regression approach.

Aircraft	DA	MLR
AC-001	-0.89%	-0.38%
AC-003	0.76%	-0.01%
AC-004	-0.19%	-0.12%
AC-101	0.20%	0.14%
Avg.	-0.03%	-0.09%

The overall percent change in fuel for these aircraft was very small. However, it should be noted that the data points taken before the CVG was applied were taken up to a year before. The consensus in the aircraft industry is that these types of aircraft generally show an annual average degradation of fuel flow of over +0.5%. This means that applying the CVG counteracted this annual degradation factor and reduced the fuel

flow even more. Therefore, it would make sense to think that an aircraft should see a 0.5 – 1% decrease in fuel flow after the CVG is applied to an aircraft of these types.

5.4 Feature Selection and Parameter Tuning using Cross Validation

We used a basic shuffled k-fold cross validation technique to perform feature selection and model tuning [6]. The data was randomly ordered then broken up into 10 sets of 90% training and 10% testing sections so that each test set did not use any of the same data and all the data is used to both test and train. This produced 10 different train and test sets generating 10 different R^2, MSE, and MAE values. These values were then averaged to produce one set of R^2, MSE, and MAE values. The averaged MSE and MAE values were then used to perform feature selection and parameter tuning. The features that gave the smallest MSE or MAE value were chosen for the final model.

We found that the best MLR model was to use the previous model from Eq. 2 with one additional interaction term. The interaction term that gave the best MSE was the *Phi*, ϕ, variable as seen in Eq. 1. This variable is a combination of all the variables in Eq. 2 except for the engine power N1 which we call *N*. The addition of the ϕ variable decreased the best MSE value from Eq. 2 by 10%. Eq. 3 shows this new model.

$$FF = \beta_0 + \beta_1 W + \beta_2 U + \beta_3 P + \beta_4 T + \beta_5 N + \beta_6 \phi \tag{3}$$

The k-fold cross validation technique was also used to tune the parameters in the random forest regression (RFR) algorithms as well as the gradient boosting regression (GBR) algorithms. The parameter values that gave the smallest average MSE across all the folds were the ones used to perform the RFR and GBR analysis.

5.5 Random Forest Regression

A random forest regression (RFR) technique was also used to model the data. The model that minimized the average MSE was Eq. 2. Taking into account the engine power seemed to be the one constant to improve the fits no matter which model was used. Random Forest is an ensemble technique that utilizes random decision trees to fit the data that generally to a good job of not overfitting the data [7]. It does this by trying to reduce variance. In this case we used the *RandomForestRegressor* function from the python sklearn library.

The RFR function works well after tuning the hyper-parameters of the function. The hyperparameter that we tuned are, the number of estimators, maximum depth, and the minimum samples split. The number of estimators details the number of trees in the forest, the maximum depth details the maximum depth of the trees, and the minimum number of samples details the minimum number of samples to split an internal node.

Table 4 includes the values for the percentage change in *FF* using the RFR technique. You can see that the values have more variation but are on average show a larger reduction in fuel flow.

5.6 Gradient Boosting Regression

A gradient boosting regression (GBR) was also used to model the data. The model that minimized the average MSE for GRB again was Eq. 2. Like RFR, GBR is an ensemble technique that utilizes random decision trees to fit the data that generally to a good job of not overfitting the data [8]. However, unlike RFR the trees are shallower, and it seeks to reduce the error by reducing the bias instead of the variance. In this case we used the *GradientBoostingRegressor* function from the python scikit-learn library.

The GRB function works well after tuning the hyper-parameters of the function. The hyperparameter that we tuned are, the number of estimators, learning rate, and the minimum samples split. The number of estimators details the number of trees in the forest, the learning rate details the contribution of each tree, and the minimum number of samples details the minimum number of samples to split an internal node.

Table 4 includes the values for the percentage change in *FF* using the GBR technique. You can see that the values have more variation but are on average show a larger reduction in fuel flow than the MLR model but not the RFR model.

6 Analysis of the Results

Goodness of fit and difference in fuel flow are the two key metrics used to compare our model. Goodness of fit is measured with root-mean-square error (RMSE) and mean absolute error (MAE). RMSE is used to measure actual observation values versus the model's predicted model. Similarity, MAE is the difference between the predicted value and actual value. For our analysis, MLR provided the best fit without overfitting the data. Table 4 shows the results of all the models averaged over all four aircraft.

Table 4. Summary of average RMSE and MAE for the dimensional analysis, multiple linear regression approach, multiple linear regression with Phi interaction term, random forest regression, and gradient boosting regression models based on 10-fold cross validation. The MLR model does the best without overfitting the data.

Model	RMSE	MAE
DA	0.0572	0.0453
MLR	0.0122	0.0095
MLR+	0.0112	0.0087
RFR	0.0242	0.0173
GRG	0.0169	0.0121

Difference in fuel flow before and after application of the CVG product was also evaluated. Table 5 shows the results of the percentage change in fuel flow for the five models developed. You can see that the better fit from the MLR+ model gives slightly different values for the change in the fuel flow. However, these results are not significantly different than the MLR numbers.

6.1 Different Values Change the Fuel Flow

The results in Table 5 are calculated using the average values of each feature in each model. This focuses on the difference in fuel flow for average operating conditions for each individual aircraft. However, it may be useful to look at how the effect of the CVG is different for non-average values of weight, speed, and power. With this in mind we passed different values of these features to our models to get results that corresponded to different regions of the features.

Table 5. Summary of the percentage change in fuel flow. DA refers to the dimensional analysis, MLR for the multiple linear regression approach, MLR+ for the MLR model plus the Phi interaction term, RFR for the random forest regression on model in Eq. 2, and GRB for the gradient boosting regression on model in Eq. 2.

Aircraft	DA	MLR	MLR+	RFR	GBR
AC-001	-0.89%	-0.38%	-0.17%	+0.74%	+0.59%
AC-003	0.76%	-0.01%	-0.31%	-0.67%	+0.14%
AC-004	-0.19%	-0.12%	-0.25%	-2.73%	+0.74%
AC-101	0.20%	0.14%	+0.9%	-1.14%	-0.17%
Avg.	-0.03%	-0.09%	0.00%	-0.95%	-0.41%

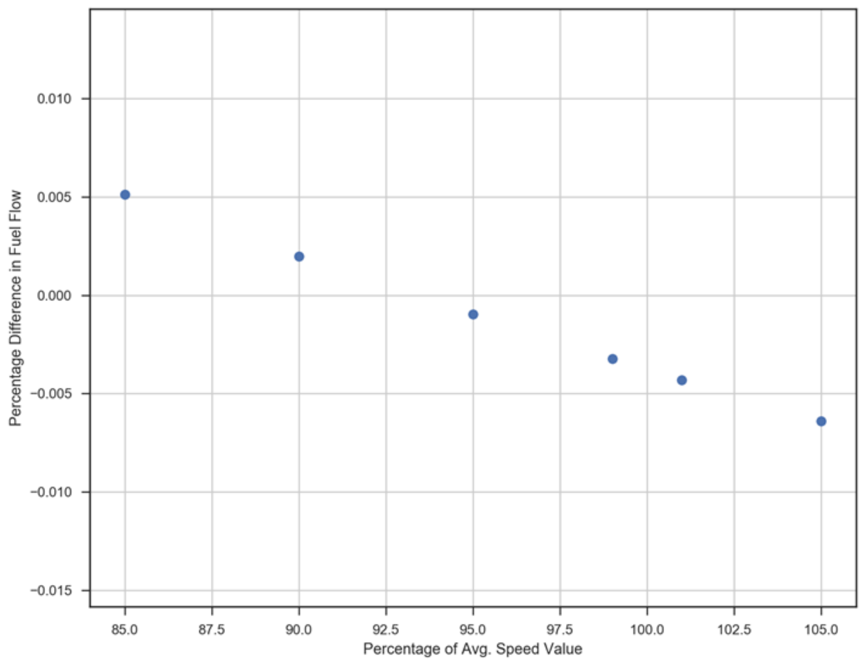

Fig. 7. Changes in fuel flow for different values of speed. This shows improved efficiency in fuel flow with smaller values speed for AC-001.

Different values of weight, speed, and power were passed to our models. We saw an increase in the change of fuel flow for increased values of weight and speed and a decrease in the change for increases in power. Fig. 7 shows the increase in the change of fuel flow with 0.2% to 0.5% increase in the average speed passed to the models. This indicates that the CVG is better at reducing drag at higher speeds when the airflow over the wings is ore turbulent.

7 Conclusions

Four additional analyses were performed on flight data from before and after the application of EA's vortex reducing coating. Data for the analyses came from over a year's worth of flight data from four different aircraft corresponding to over 100,000 flight hours.

The models we developed in this study had only marginally different results from the original dimensional analysis model developed by EA. The multiple linear regression model with the interactive term estimated no difference in fuel flow, but this could have been due to overfitting. Multiple linear regression without the interactive term was closest to EA's original DA analysis, with a difference of -0.09% in fuel flow versus -0.03%. The ensemble models, random forest regression and gradient boost regression, resulted in the largest differences with -0.95% and -0.41% respectively.

While the analyses presented in this study was not complete, the main conclusion that can be drawn is that the Edge Aerodynamix CVG reduced the fuel flow in the aircraft studied by 0.5% to 1.0%. More work needs to be done to better specify the error and accuracy associated with the models used in this study.

Acknowledgments. We would also like to thank Buddy Black, Bivin Sadler, Robert Slater, and Daniel Engels for their help and contributions to this work.

References

1. Edge Aerodynamix, Inc., "Airline Fuel Performance Data Reduction and Analysis", Proprietary and Company Confidential, EA21038-11 rev 1, October 2017.

2. Anderson Jr, John David, "Fundamentals of Aerodynamics", Fourth Edition. Tata McGraw-Hill Education. 2010.

3. Neil A. H. Campbell, "The Evolution of Flight Data Analysis", MO3806. https://www.skybrary.aero/index.php/Flight_Data_Recorder_(FDR)

4. Avionica Inc., "Digital Flight Data Analysis Software", Part Number 900-0004, http://avionicainc.com/products-3-4-2/

5. "Flight Data Monitoring A Guide to Implementation" -16 February 2009. https://www.flightdataservices.com/

6. Scikit-learn: Machine Learning in Python, Pedregosa et al., JMLR 12, pp. 2825-2830, 2011. http://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf

7. Ho, Tin Kam (1995). Random Decision Forests (PDF). Proceedings of the 3rd International Conference on Document Analysis and Recognition, Montreal, QC, 14–16 August 1995. pp. 278–282. http://ect.bell-labs.com/who/tkh/publications/papers/odt.pdf

8. Friedman, J. H. "Greedy Function Approximation: A Gradient Boosting Machine." (February 1999). http://statweb.stanford.edu/~jhf/ftp/trebst.pdf