MTU Aero Engines

Analysis of V2500 Engine Deterioration by Operator

Data Innovation Lab: Technische Universität München

Alexandra Barbu
Kashtin Bertoni
Jakub Mazur

Matthias Berger
Tobias Reinerth
Peter Schlick

Supervisor: Professor Dr. Massimo Fornasier
Project Lead: Dr. Ricardo Acevedo Cabra

February 9, 2018

Abstract

The focus of this report is to predict the effective scrap rates of V2500 engines at the operator level. Damage mechanisms of engines were studied in order to understand the important factors affecting engine deterioration. Alongside this, flight tracking data from Flightradar24 and maintenance reports were used as the primary datasets for modeling average engine scrap rates. Additional data regarding environmental factors was collected to support the existing data. Scrap rates were modeled in two approaches called the Fleet Approach and Single Engine Approach. While no definite predictions can be drawn, this report provides an understanding of some of the important factors that influence the deterioration of V2500 engines. Furthermore, it establishes a firm basis for future data collection and the possible connections that can be made by acquiring information to develop better predictions in the future.

Contents

1 Introduction

Airplane engine maintenance and repairs are a major source of income for engine manufacturers. On the other hand, scheduling maintenance visits and planning for spare part production requires a large amount preparation of on the manufacturer's behalf. It is possible to anticipate required parts for the repair by creating more sophisticated models of deterioration.

The goal of this report is to better understand engine deterioration and how the environmental conditions that an airline operates in can impact this. This research question will be explored through data and analysis of the V2500 engine. This particular engine is used on aircraft flying short to medium range flights. For MTU Aero Engines, the interest is on determining whether there exist differences in engine scrap rates between airlines, often referred to as operators, based on environmental conditions in their area of operation, or in their method of operation.

A variety of sources were consulted to determine factors affecting engine performance degradation. While there are many reasons leading engine performance degradation, for example defects in the components and random bird strikes, the main focus will be on normal wear of the components. To determine additional impacts on the normal wear of the components, literature regarding damage mechanisms was consulted. It suggests that the environment plays a considerable role in the degradation of the engine, especially during takeoff. Less damage is accumulated during cruise where the engines are exposed to less particles and active components.

The report's focus is on cities from which each flight departs. Flightradar24 was used as the source for completed flights by the operators. In addition, MTU experts were consulted to identify environmental factors for modeling. The factors selected for their perceived importance were NO_2, SO_2, the levels of particulate matter under 2.5 microns (PM2.5) and under 10 microns (PM10) as well as the temperature.

Publicly available datasets of different air quality factors were collected and used. During data preprocessing, the distributions of flights for each operator were analyzed and environmental data was connected to cities with flight departures. Airports had environmental factor scores developed and these were then aggregated to the operator level based on flight distributions.

MTU maintenance data of engine overhauls was used in the development of outcome factors. The scrap rates are considered the most important outcome, where the components are inspected and deemed either serviceable, meaning they can be reinserted into the engine, or scrap, meaning they must be discarded. The components of interest are located in the high pressure and low pressure turbine sets, abbreviated HPT and LPT. This is comprised of HPT blade sets 1 and 2, HPT vane sets 1 and 2 and LPT blade sets 3 through 7, often abbreviated HPTB*, HTPV* and LPTB* respectively. Different environmental factors are thought to impact each part of the engine differently, thus each component is analyzed separately with respect to the operating conditions.

The modeling was approached from two major perspectives, referred to as the 'fleet approach' and the 'single engine approach'. The main difference is in the fleet approach, each airline is a single data point of aggregated data and the prediction goal is scrap rates for each airline. In the single engine approach, each engine is a single data point and the scrap rate is being modeled individually.

There were issues in the analysis trying to improve predictive power and individualize engine performance so this report concludes with suggestions. These direct the focus to improve future modeling by collecting new data and making new connections in the data, which could make analysis significantly more precise and enhance the predictive power of the models.

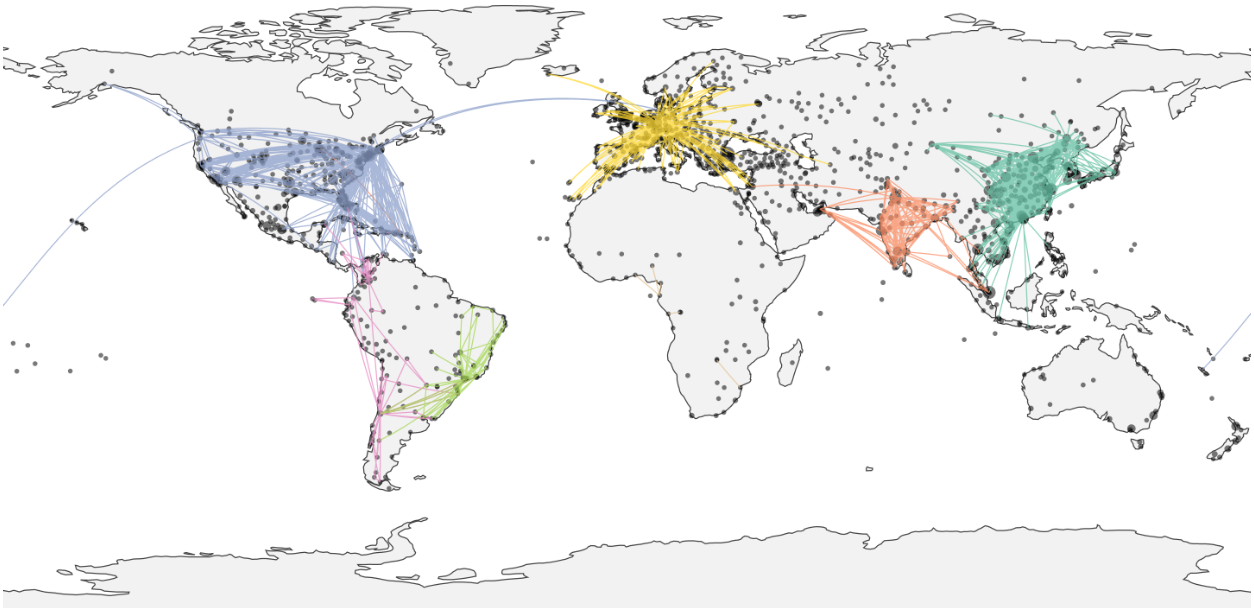

Figure 1: Visualization of flights powered with V2500 engines from several operators in 2017.

2 Preliminary Research

The initialization of the project required background knowledge on aircraft engines and the damage mechanisms that influence different engine components. There are subtle interconnections between different conditions that develop over time, causing damage to parts. While it cannot be determined exactly when it would no longer pass inspection, it is important to focus on the wear conditions that cause the components to deteriorate over time. This is often the reason for engines to be sent to the shop as the deteriorated elements cause the engine to no longer reach an operator's threshold for engine performance. The information on the airplane engines was obtained through discussion with experts and the analysis of "Untersuchungen zum Einfluss der Betriebsbedingungen auf die Schädigung und Instandhaltung von Turboluftstrahltriebwerken" by Matthias Müller.

2.1 Failure Classifications

When an engine comes in for an overhaul, each of the components is examined and classified as 'scrap', 'serviceable' or 'repair'. Serviceable can be put back into the engine with little maintenance. Those that are deemed repair will eventually end up in either the scrap or the serviceable categories after a second inspection. Certain operators have 100% scrap policies, where, regardless of their condition, all components being examined are scrapped. This led to the need to remove all 100% scrap rates from the data as it was not appropriate for modeling.

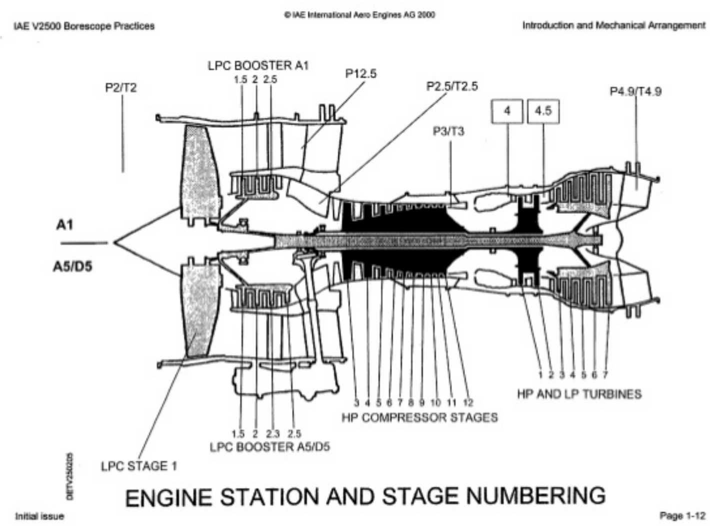

Figure 2: Diagram of V2500 engine where the HPT and LPT are located at the right end.

2.2 Engine Components

The focuses of this study were the parts within the V2500 engines' high and low pressure turbines (HPT/LPT). The HPT contains sets of blades 1 and 2 (HPTB1/2) and vanes 1 and 2

(HPTV1/2), while the LPT is composed of sets of blades 3 through 7 (LPTB3-7). Every component is arranged on a ring with its brethren, each set containing different numbers of blades or vanes. Additionally, certain operators handle spare components that they can choose to change based on internal decisions. These can cause a greater than normal number of components to be returned during overhauls and affects the modeling process used. Each set of components are fabricated according to engine needs, with their locations and function corresponding to different risk profiles in respect to different damage mechanisms.

2.3 Damage Mechanisms

Damage mechanisms are the methods by which the engine components decay. Many mechanisms contribute to component wear and the eventual performance reduction that brings the engines into the shop for an overhaul. Most generally, the components develop micro-cracks which increase in size over time into macro-cracks that reduce performance and can become unstable. The effects of the damage mechanisms occur over long periods of time, with engines having overhauls on average every four or five years. Some of the mechanisms are due to gradual deformation, such as;

- creep - where the part suffers gradual deformation under stress
- fretting - where wear occurs between two attached components through micro-movements
- abrasion - where wear occurs between rotating and stationary parts

In contrast, other mechanisms are more related to the environment that an engine is flying in, including;

- oxidation - surface deterioration linked to engine temperature
- corrosion - another form of surface deterioration linked to engine temperature
- erosion - the wear on surfaces caused by the intake of small airborne particles

The accumulation of these mechanisms over time will reduce engine performance and affect the scrap rates of the components.

2.4 Influence of Flight Segments

Each segment of a flight requires different things from the engines. The most significant portion of wear damage happens in lower altitudes, as cruising at altitude occurs in thin air with less damaging particles and molecules present. The cycle of a flight lasts from the departure from the gate, until the landing at the next airport. Engineers consider taxiing, take-off and climb as the most significant factors in long-term engine deterioration. These maneuvers require high thrust levels and are performed in more contaminated environments. As a result, many components of an engine are limited by law in relation to the number of cycles the engine completes instead of the number of flight hours. A key factor in the deterioration of components is the derate. This describes the reduction from the full potential power of the engine thrust based on a variety of factors at the airport and the weight of the plane. The greater the power reduction, the less damage the engine incurs from more detrimental flight phases, the most significant portion being the takeoff.

2.5 Relationship to Environment

The environment plays a key role in the development of engine damage. While there are connections between individual environmental factors and certain damage mechanisms, the relationships tend to be complicated with many interactions. Erosion is ostensibly connected to the concentration of particle matter found in the air, but this is not the case with many of the other mechanisms. For example, temperature often has a non-linear relationship involving external temperatures in conjunction with concentrations of compounds in the air. From the literature and professional advice, it was determined that the effects of temperature, concentrations of NO_2, SO_2 and particulate matter at sizes under 2.5 (PM2.5) and 10 (PM10) microns (μm) could have significant impacts on component deterioration over time.

2.6 Airport Environment Connections

Since it is accepted that the most detrimental portion of a flight cycle is centered around the airport of departure, these locations would be used to score a plane's operating environment. Each of the environmental conditions that were of interest were collected for the analysis from public sources and aggregated based on each operator's flight distribution. In this way, the environment of the average flight plan of an operator could be assessed during the modeling.

3 Datasets

The datasets of primary interest focused on flight distribution, maintenance results and environmental factors. They were a combination of purchased data, public information and internal records. The accuracy and granularity of data that needed to be collected depended extensively on what was available in existing datasets. Datasets for the project were collected, preprocessed and analyzed. This process helped clarify the final approaches to modeling scrap rates.

3.1 Shop Visit Results

The **Shop Visit** dataset came from MTU internal maintenance records. It contains information on all shop visits from **2011** until **2016** for different sets of engine components: HPTB1, HPTB2, HPTV1, HPTV2, LPTB3, LPTB4, LPTB5, LPTB6 and LPTB7.

This information covered shop visits of engines belonging to **26 operators**. The features from the overhaul data that were of greatest interest to the analysis consisted of:

- VCSN - number of cycles (flights) since production
- VCSO1 (and VCSO2) - number of flights since major (and minor) overhaul, as applicable
- VTSN - number of hours flown since production
- VH2CR.all - hour to cycle ratio since production
- VH2CR.run - hour to cycle ratio since last overhaul
- Operator - airline that owns the engine in the shop
- Scrap.Rate - percentage of parts rejected, per component

3.1.1 Preprocessing and Cleaning

Since every blade in the HPTB1/2 sets are serialized, the maintenance information was provided at the level of individual blades and had to be aggregated to engine-level data. The HPTV1/2 and low pressure turbine blades were already aggregated at the engine level.

The maintenance datasets demanded the most work in sense of data cleaning and processing. A significant issue was missing values. These occurred in data across all considered engine components. A method to reduce this was considered by noting that information regarding cycles and hours flown must be consistent across all parts inside of an engine. Reports on certain components had data recorded and could fill in what others missed, thus reducing the number of missing values. As HPT blades had information for individual blades, missing features from other records of the engine could be filled in from these values. Missing flight hour to cycle ratios could be calculated directly by division of the relative values. Two features that were vital to the analysis, Scrap.Rate (by component) and Operator, must be present to be able to use the data for modeling.

3.1.2 Connection to Further Data Sources

Each engine was exposed to different environmental conditions. These differences between engines of two airlines operating in distinct regions are evident, but differences also occur for the engines within an airline. The goal is to explain the observed scrap rates in an engine based on the amount of flight cycles and hours and the environmental conditions that a particular engine was flying in. Therefore it is important to check if environmental conditions among engines of one operator differ significantly or not.

3.2 Flightradar24 Data

The **Flights** dataset was purchased from Flightradar24, an internet service that shows real-time flight location and information. The company also records and stores information from past flights. This dataset contains flights from **January** through **September** in the year **2017**, for airlines who have not opted out of the tracking. Each row in the dataset is a unique flight with dozens of features, such as airline, destination and origin airports with their coordinates, airplane tail ID and the engine type.

airline <chr>	origin_airport <chr>	destination_airport <chr>	takeoff_date <chr>
Air Europa	Amsterdam Airport Schiphol	Adolfo Suárez Madrid–Barajas Airport	2017-01-09 10:44:56
Air France	Biarritz-Anglet-Bayonne Airport	Paris-Orly Airport	2017-01-22 11:23:33
Asiana Airlines	Gimpo International Airport	Jeju International Airport	2017-01-16 23:19:08

Figure 3: Sample of Flightradar24 data with several features displayed.

3.2.1 Preprocessing and cleaning

Since data about each airplane is available, it was possible to obtain the incidences of take-offs from each airport by aircraft. Dividing each departure count by the total number of all flights of this aircraft leads to the development of its **flight distribution**, or the frequency of its departures by airport. In the same manner, the flight distribution of the **whole fleet** is calculated by all plane departures by the operator.

Having obtained the flight distribution of every airplane within a fleet, it is possible to check if airplanes of this operator are flying in a balanced manner or not. This is done by comparing each plane's distribution to the distribution of the whole fleet. The more planes that have similar distributions to the overall flight distribution, the more balanced the fleet operates. A theoretical example of an unbalanced way of flying would be a North American operator whose planes are divided into two groups. The first operates on the West Coast and the second along the East. The fleet distribution would be an average of the two sets of flight distributions, which differ significantly. This would result in no planes following the overall flight distribution precisely.

In order to compare the individual airplane distributions against the fleet distribution, it is necessary to introduce a distance measure between two discrete distributions $p = (p_1, p_2, ..., p_m)$ and $q = (q_1, q_2, ..., q_m)$. Here there are four formulas for varied difference measures:

- Hellinger distance: $d_{Hellinger}(p,q) = \frac{1}{\sqrt{2}}\sqrt{\sum_{i=1}^{m}(\sqrt{p_i} - \sqrt{q_i})^2}$
- Bhattachryya distance: $d_{Bhattachryya}(p,q) = 1 - \sum_{i=1}^{m}\sqrt{p_i q_i}$
- L1 distance: $d_{L1}(p,q) = \frac{1}{2}\sum_{i=1}^{m} \mid p_i - q_i \mid$
- L2 distance: $d_{L2}(p,q) = \frac{1}{2}\sum_{i=1}^{m}(p_i - q_i)^2$

In the case of two identical distributions, distance is equal to **0**. The more the two distributions differ, the bigger distance gets. In perfect disagreement between distributions, each element i being compared has one of $p_i = 0$ or $q_i = 0$ causing distance to equal **1**. For each of the 26 operators of interest, the following steps were taken to check if the fleet operates in balance:

- the flight distributions of entire fleets were calculated
- the flight distribution of each plane in a fleet was calculated
- the four distance measures were calculated for each airplane and displayed in histograms

The two examples show how these procedures were applied to all 26 operators. In most cases, the fleets operated in balanced way, where a majority of the probability mass was concentrated on interval $[0, 0.25]$. There were also operators flying in non-balanced way, as shown in Figure 5. Therefore it would be beneficial if each engine had its own environmental scores.

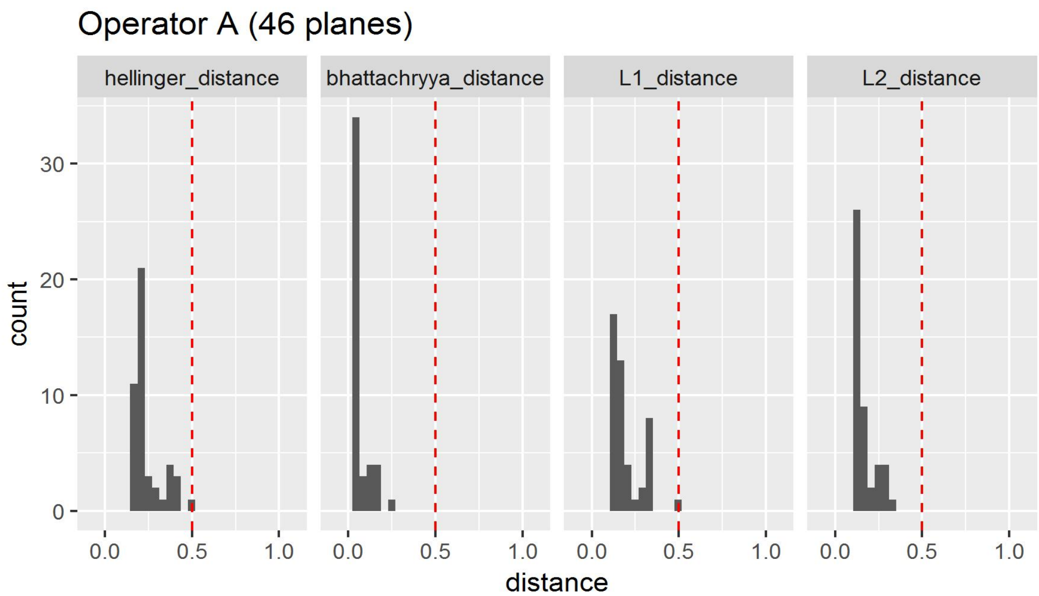

Figure 4: Histogram of balance measure - example of balanced operator.

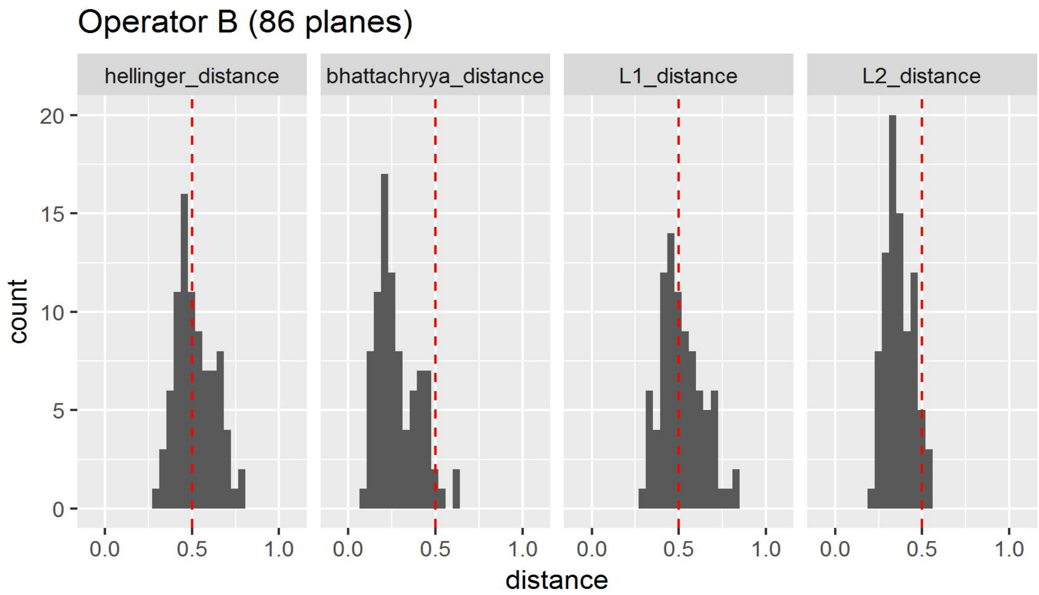

Figure 5: Histogram of balance measure - example of non-balanced operator.

3.2.2 Connection to Further Data Sources

The next task was collecting environmental data for airports of interest. Since it is very hard to find consistent and detailed data that cover the globe, some actions were taken to make this process less difficult. This included ignoring airports with fewer take-offs by taking the most traveled airports by operator such that **80%** of all flights of this operator were covered. This removal filtered out less frequented airports and was also useful during data collection as it was difficult to find environmental data for the less popular airports.

The data does not provide a way to track the flight pattern of a single engine. While there is history of flights by airplane, engines can be rotating between aircrafts and there is no available link between engines and airplane. Another data source is needed to potentially track a single engine.

3.3 Engine Trend Monitoring

Another dataset is **Engine Trend Monitoring (ETM)**. This is composed of continuous in-flight monitoring of key engine readings from standard engine and aircraft instrumentation. The most relevant elements of this dataset were the aircraft ID, engine ID, altitude, outside air temperature, derate and exhaust gas temperature margin (EGT Margin). It can provide a link between a particular engine and the aircraft it is attached to, and hence the engine's flight route.

Unfortunately, the dataset that is currently available only covers one airline and 2 months of flights in 2017. This dataset could not be used in final approach, but it is mentioned as it has potential to improve future analysis if more operators and a broader time period could be added.

As it is currently impossible to track flights by engine, aggregated environmental conditions must be applied across the entire fleet.

3.4 Environmental Data

Collection of environmental data aimed to have standardized findings in order to compare values between airports with confidence. Some collected environmental data was connected to the city that the airport services. Other data used global readings that were then isolated to the latitude and longitude in the region around the airport.

Based on the preliminary research regarding engine damage mechanisms, the most important environmental factors were determined to be temperature, fine particulate matter at the levels of 2.5 and 10 microns, SO_2 and NO_2. The initial idea was to collect environmental data that characterized every airport in the network of the operators of interest. The first method was to use environmental information from government environment agencies around the world. After analyzing the data from the American and Chinese environmental agencies, it was determined that this would not be feasible for every relevant airport. The level of granularity in this data was extremely fine, with hourly measurements for every county in the United States. In comparison, the Chinese environmental data measured these features daily. Expanding to global country-specific data sources would have been increasingly difficult to standardize all measurements to allow for comparison. Issues also arose with regards to language barriers and state manipulation of data.

This changed the focus of the collection method to finding globally comparable information for each environmental factor. Cities serviced by an airport were used connect data to an airport and simplified the data collection process. Where it was not possible to measure at the city level, measurements were estimated by taking nearby cities. In other cases, the data was connected by latitude and longitude coordinates.

3.4.1 Temperature Data

Temperature data was compiled from multiple sources. It came primarily from a travel website **Current Weather Results** and missing values were filled in with information from the weather conditions provider **Weather Base**. Temperatures on Current Weather Results were reported by various national meteorological organizations as the average monthly high and low temperatures in Celsius at the city level. Temperature was reported as averages measured from the 1990s to 2015. Missing values were filled in with measurements obtained from Weather Base, which reports current and historic average monthly temperatures in degrees Celsius. Weather Base collects its data from public domain sources, including the National Climatic Data Center in the United States. The data from Current Weather Results was converted from the high and low temperature to a single monthly value based on the following formula:

$$Temp = \alpha * Temp_{high} + (1 - \alpha) * Temp_{low} \quad \alpha \in [0, 1]$$

For each month, α was calculated from cities which reported average temperature, as well as the average monthly highs and lows to interpolate a weighted average as above. Once a single monthly average was obtained for each city, these temperatures were averaged over 12 months in order to have a single value for each airport.

3.4.2 PM Datasets

Particulate matter is the concentration of small particles in the air with commonly measured levels of those under 2.5 microns and those under 10 microns in size. As these particles are deemed a vital portion of the air quality index, there is a compilation of government data by the World Bank in support of population health analysis. The values are reported in micrograms (microns) per cubic meter ($\mu g/m^3$). The data is self-reported by country for cities and is often averaged over a number of stations in a city a year between **2012** and **2015**. Some stations only report either PM2.5 or PM10, so the missing value is estimated by World Bank using linear extensions that are region dependent. For cities that did not report the data directly, two methods were attempted. If there was another city nearby that had values reported and similar geographic features, this could be used. If such a neighbour could not be found, the World Health Organization's satellite observations of the PM levels at the latitude and longitude of the airports were used.

3.4.3 SO_2 Data

SO_2 was provided by Socioeconomic Data and Applications Center (SEDAC). The data provides an estimation of anthropogenic SO_2 emissions in kilograms from **2005** on a country-level. Data was created with the use of the bottom-up mass balance method, an application of the law of conservation of mass. This data is coarse as SO_2 is not monitored precisely on a global scale. The advantage is that this source provides **comparable** data from around the world. Since the goal is to assign these values to particular airports, the country level data needs to be scaled so it is representative. To reach appropriate emission levels for airports, the country values were divided by number of inhabitants, ending with unit **kg/person**. If this correction had not been done, bigger countries would tend to have higher scores, and this would not have been appropriate.

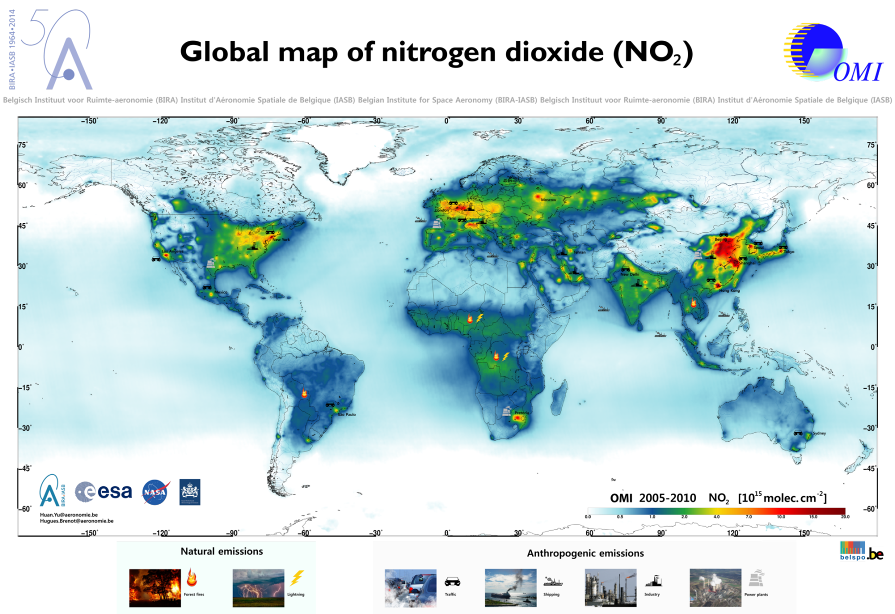

Figure 6: NO_2 levels captured by OMI using slant technology

3.4.4 NO_2 Data

The concentrations of NO_2 were determined based on satellite passes with slant technique. They come from the Tropospheric Emission Monitoring Internet Service's (TEMIS) Ozone Monitoring Instrument (OMI), which performs global scans of NO_2 levels. The data is provided as **monthly averages** with some potential negative results due to noise and the difference taken between the slant and vertical measurements. The measurements are by the number of molecules by 10^{15} per square centimeter and the monthly values for the **past year** were aggregated to yearly averages by each latitude and longitude information square. The airport values were taken from the latitude and longitude area that contained the airport and averaged with the eight areas surrounding it.

3.5 Final Datasets

Inputs for the modeling incorporate environmental scores and the hour and cycles features. This data needs to be aggregated in order to be compatible with fleet and single engine approaches.

Environmental Scores

The aggregation of environmental data for each operator used their flight distribution for weighting the individual scores measured for airports of interest. Weights for each airport were proportional to frequencies of take-offs from this airport by the operator. It can be interpreted as the expected pollutant levels that an engine by this operator would experience every take-off. For example, if the plane belonging to an operator departs from airport A 30% of the time, the pollutant value measured at airport A is taken with weight 0.3. Data in this format is used as an input for both approaches.

Shop Visit Data

For the single engine approach, the preprocessed data is ready for modeling. For the fleet approach, the means of flight times and cycles of an operator's engines were calculated, providing averages as features.

4 Modeling

Building on the understanding of the balance measure, the goal of modeling was to predict the average scrap rate. The two approaches are referred to as the **fleet approach** and the **single engine approach**. In both methods, models were developed for each stage of the engine. This was done to account for the fact that each stage is examined individually in the shop visits and different stages have different factors affecting the scrap rates.

4.1 Fleet Approach

The first approach to modeling average engine scrap rates involved looking at operator averages for each stage of the engine: HPTB1/2, HPTV1/2 and LPTB3-7. This was done to evaluate the significance of environmental factors on average scrap rates at the operator level.

For each operator that had a shop visit recorded in the 5 year time frame, averages over the scrap rates, flight hours and flight cycles were taken. While this approach loses the details about scrap rates of each engine, its advantage is the focus on the significance of the environmental factors. It attempts to effectively differentiate between the impacts of different types of regions. Differences can already be seen between airlines operating in North America when compared to one flying in the Middle East simply by comparing average scrap rates. On average the Middle Eastern airline will experience higher scrap rates than the North American one. From the data collection phase, it was observed that the North American operator will experiences a range of temperatures while the operator in the Middle East will, on average, experience hotter climates, leading to different average temperature. This, and significant ranges in the environmental values, is an indicator that the operator oriented approach is worth exploring.

The low number of data points limited the models that could be used and linear regression was in the end the most sensible choice. While there were around 960 shop visits in a five year time period, not every stage of the engine got checked in every shop visit. Furthermore, since the operator approach involves averaging each feature by operator, we have at most 26 data points for each of HPTB1 and HPTB2, one per operator. In HPTV1 there are 23 airlines, in LPTB3 21, HPTV1 and LPTB5-7 each have 19 and in LPTB4 there are 17 that have any maintenance performed. Engines which had 100% scrap rate were omitted in the average scrap rate calculation, since 100% scrap rate can be due to the operator's policy, not damage.

For each HPT and LPT stage, correlations were looked at to determine which features to use in the linear regression. There were high correlations between VH2CR.all and VH2CR.run, only the former was used because it averages over a longer time period. VCSN, VCSO1, VTSN and VH2CR.all covered the most of the variability in the deterioration rates along with the environmental factors. The model used the average scrap rate as the response variable.

The linear regression equation is as follows:

$$\begin{aligned} reject.rate = {} & \beta_0 + \beta_1 * VH2CR.all + \beta_2 * VCSN + \beta_3 * VTSN + \beta_4 * VCSO1 \\ & + \beta_5 * NO_2 + \beta_6 * SO_2 + \beta_7 * PM2.5 + \beta_8 * PM10 + \beta_9 * temperature \end{aligned}$$

After fitting the linear regression model for every stage of the engine, feature selection was performed to avoid overfitting. This was done by selecting the linear regression model that maximizes the Akaike Information Criterion (AIC).

Leave one out cross validation (LOOCV) was performed in order to see how the scrap rate would be predicted for a new operator. For each stage of the engine, one operator would be omitted and a new linear regression model would be fitted. This would then be used to predict the scrap rate for the operator that was left out.

4.2 Single-Engine Approach

The second method was developed to overcome the low number of operators and is based on individual engine maintenance results. This greatly increased the number of records available and improved the statistical significance of the factors in the models. A downside of this method is it used the same levels for each environmental factor across all engines within an operator. The only data that can be used to differentiate between different engines within an operator are the number of flight hours, flight cycles and their ratio. As a result, many models showed only slight variations from their operator's average scrap rate. Additionally, interaction terms were formed with the number of flight cycles and the environmental factors to represent the relative difference between engines. In this way, there was a better distinction between engines that flew only 2,000 cycles and those that flew 6,000. Cycles were used as the interaction factor instead of flight hours since they better represent the exposure of the engine to the low lying ground level particles.

Each model was developed and evaluated through the AIC but also minimize the average sum of squares prediction error. The prediction error comes from 15-fold cross validation. Minimizing the prediction error was most important, but AIC was needed to help reduce models. It was a balance of finding the AIC suggested model and reintroducing terms to improve the prediction error. The summaries of the final models for each component can be found in the Appendix. To evaluate the effectiveness of different probability distributions in the regression, the scrap rates and the scrap counts were modeled through logistic regression and regression with count data respectively.

4.2.1 Binomial Regression

A regression model with a binomial function was used to predict the scrap rate of the components. For most models, the logit was the link function chosen. The C-Log-Log link had better predictions for some components and was also used. The probit link was not compared due to its similarity to the logit link. The binomial regression models had more issues with attaining statistical significance for components than the counting regression models, although they often performed better with respect to prediction errors.

High Pressure Turbine Components

These models had difficulty matching the quality of the counting regression predictions. While most of the models had many of the factors as statistically significant at $\alpha = 0.05$, the HPTB2 did not have any significant factors. Across the HPT components, the most commonly included environmental factors were the PM2.5 level, NO_2 and the temperature. They also appeared to be stronger influences based on the predicted coefficients and their presence in viable interaction terms. As would be expected for degradation over time, flight cycles or flight hours were included as factors in every final model. Interaction terms appear in the final blade models, but not the vanes. Additionally, the vanes tend to have very low scrap rates, averaging around 0.13% for vane

1 and 0.035% for vane 2, which explain why the models have lower prediction errors. These will be further discussed in the next chapter.

Low Pressure Turbine Components

The binomial models for this portion of the engine performed better using the C-Log-Log link for blades 3, 4 and 7, while the logit was better for blades 5 and 6. These differences are likely due to chance, but the models that performed best are used in the analysis. Final models for the components in the LPT often included a factor for either PM2.5 or PM10. In contrast, temperature was never in a final model. The regressions with the best performance tended to be simple and often had low statistical significance. Some of this is due to the reduced number of records for overhauls on the LPT, but it could also suggest that counting regression is more stable.

4.2.2 Regression with Count Data

A different approach to examine models that could prove more accurate was using the scrap counts for each part. Some operators occasionally carry spare engine components that they switch internally and can send in for repairs with an engine. This produces artificially high scrap and serviceable quantity counts, but occurs rarely enough such that omitting these misleading amounts, did not eliminate too much data. There were also cases where less than the normal amount of blades or vanes were assessed. When there were only minor differences from the normal count, these observations were included.

Poisson Regression

Regression using a Poisson distribution was done to model the number of scrap components in each set. The data was used more effectively than in the binomial approach since Poisson models use the number of failed components, instead of the failure rate. Although the observed scrap counts are limited by the number of components in each set, and the Poisson is theoretically unbounded, the Poisson models proved superior in some of the predictions. In contrast to the binomial regression, these models often showed more statistical significance for the factors that were included in the final model.

High Pressure Turbine Components

The Poisson models yield high levels of interaction with three or four term interactions while maintaining statistical significance for the factors. The predictions were superior to the binomial models in most cases of the HPT components. The models suggest that SO_2 was not as statistically significant as the other factors. Due to the high number of interaction terms, it is difficult to distinguish the effects that a change in a single factor could have on the number of scrapped components. Given more time, it would be preferable to reduce the number of terms and interactions to help identify potential environmental causes that impact the number of scrapped blades and vanes.

Low Pressure Turbine Components

In the LPT portion of the engine, the binomial models regularly outperformed the Poisson models. Once again, the Poisson models tend to have strong statistical significance while maintaining many factors in the model. Choosing to remove some factors and repeating the predictions to check for improvements could provide superior models relative to the current AIC testing. In the final Poisson models, the NO_2 and temperature factors were often included in fourth-level interaction terms, which often suggests some importance. SO_2 is also involved in smaller interactions and is often connected to the number of flight cycles. This could imply that there is a buildup of damage from SO_2 over a number of cycles. At least one particulate matter term should be included in the model, but both PM2.5 and PM10 show differing amounts of significance throughout. This is likely due to their high correlation. The Poisson models appear to excessively complicate many of the interactions and this could be due to the mean-variance restriction.

Negative Binomial Regression

As an attempt to improve on the restraints present when using the Poisson distribution in the model, negative binomial regression was attempted. This relaxes the mean-variance equivalence restriction. Overall, it allowed improvements on the prediction error for some components. The gains in performance for the counting model were mostly present in the HPT. The negative binomial models were able to improve upon some of the Poisson models and provide the best predictions among the single engine models for HPTB1 and HPTV2. The models often showed the same significant factors as in the Poisson regression.

4.2.3 Regression Trees

The relationship between scrap rates and the independent variables is not necessarily linear. For instance, cold temperatures are said to be the least harmful for engines, but warm climates are considered to be worse that hot ones. Therefore, nonlinear models were taken into account as well. Since the datasets were only between 300 and 900 observations, a complicated method could not be applied. For this reason, regression trees were attempted as well. After building trees with all of the features of interest, they were pruned so that cross-validated mean squared error is minimized. Additionally, it was possible to measure the importance of each individual feature by consecutively testing each model that had exactly one feature left out at each iteration.

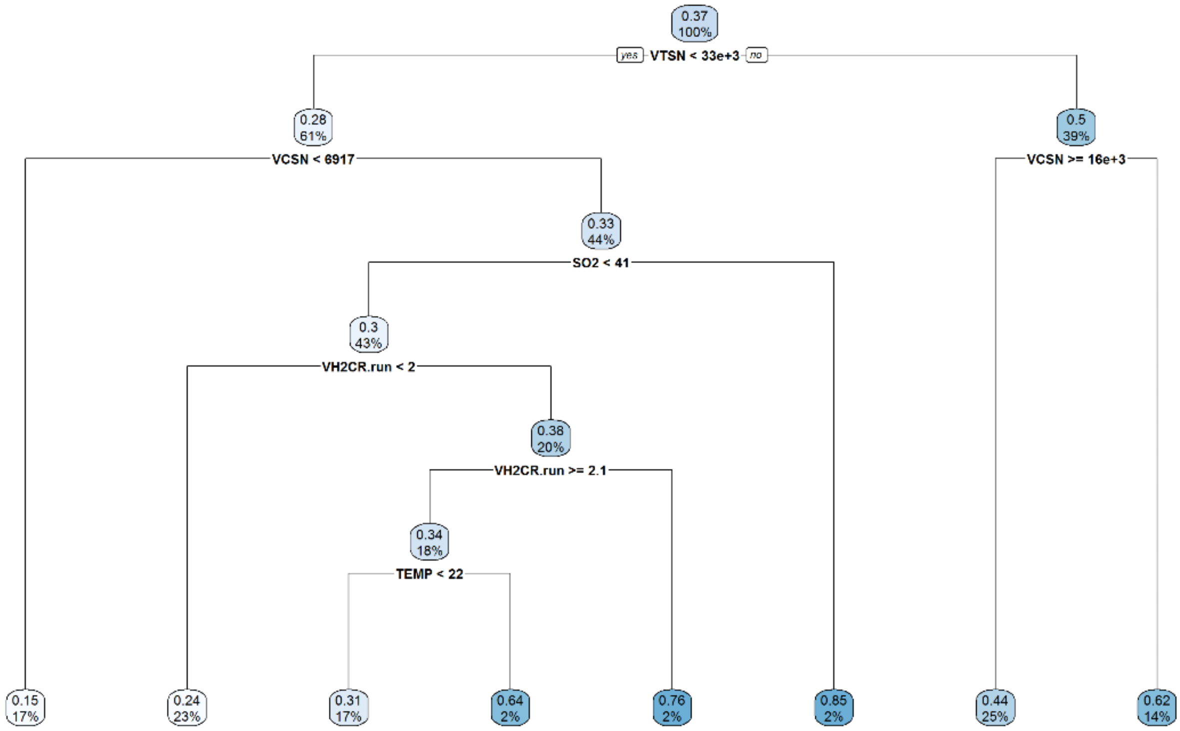

Figure 7: Example of a modeled regression tree.

5 Discussion of Results

The models suffered from a lack of precise data. It was not possible to connect an engine with the specific flight plan that it flew. This was due to both temporal issues and the missing information that would connect the data for each engine to the airplane it was attached to. The temporal issue is centered around the time periods that each of the data sources cover. The maintenance data was provided for 2011 until 2016, while the flight data was only from 2017, thus reducing the ability to have distinct cause and effect conclusions. The information regarding specific engines being attached to specific planes needs to come from operators themselves as it is not in the public domain. Operators have their own internal policies and can switch engines between their own planes. Additionally, some operators have policies that only one engine should be used while taxiing, causing an unbalanced relationship between two engines on the same plane. These differences make definite conclusions difficult to draw, but some general relationships can be suggested.

5.1 Environmental Factors

The environmental factors have their own distinct distributions of values that differ across the 26 operators. When interpreting the potential impact a factor has on the scrap rate, the range of these values must also be considered. Those with a larger range, such as PM10 and NO_2, can have more impact even with a smaller estimated coefficient. In comparison, the weighted average temperature varies by at most 18 degrees and needs a higher coefficient to make an impact.

A disadvantage of using average environmental scores for each operator is the fact that the true measures of the environment are not reflected in these means. Using temperature as an example, it is fairly constant throughout the year in the Middle East, but North America experiences a larger range of temperatures throughout the year and between different regions. For example, a plane that flies between New York and Florida during the winter can experience a difference in temperature of 30 degrees, but the average temperature for and operator in this region is 15 degrees. Unfortunately, these extremes experienced by the engines cannot be accounted for without the ability to identify the specific flight locations of engines.

These connections should be considered during analysis to determine the true importance of each factor in the model.

	PM2.5	**PM10**	**NO2**	**SO2**	**TEMP**
Min.	8.274	12.4	11.83	6.233	9.629
1st Qu.	14.872	23.28	16.7	20.019	15.927
Median	23.301	30.99	38.71	24.88	18.375
3rd Qu.	45.332	70.34	63.72	38.517	21.778
Max.	82.191	128.42	182.15	58.211	27.452
Range	73.917	116.02	170.32	51.978	17.823
Mean	32.235	49.67	48.36	27.362	18.785

Table 1: Distribution of Environmental Factors

5.2 Fleet Approach

Feature selection for the fleet approach found that, on average, the most important features are VCSN, PM2.5 and PM10. VCSO1 was never included in the final model. There are some differences between the HPT and LPT stages in terms of the features that were selected. In the HPT stages, VTSN was most often selected. In contrast, in the LPT stages, VCSN and temperature were chosen most frequently. Testing the significance at the level $\alpha = 0.05$, for the HPT stages, there were no patterns in statistical significance between the four sets. It is remarked that PM2.5 was selected for HPTB1/2 and PM10 was selected for PM10 but was not necessarily statistically significant.

In LPTB3-5, none of the features selected were statistically significant. This is likely due to the small number of data points. In the LPT stages 6 and 7, VCSN was the only statistically significant feature that they had in common. The statistical significance of VCSN indicates that it matters more how often an engine has flown and not the amount of time it has flown, supporting the intuition that takeoff is the flight segment where the most damage is incurred.

The ability to take average environmental measures and compare them to average operator values is the main advantage of the fleet approach. In this manner, there is variety between each data point with respect to all of the features for every single operator. By taking average values for the engine data including the measures of scrap rates and hour to cycle ratio, there is a correction for some misleading values. Such values include the uncharacteristically high scrap rates from operators that have a specific contract with MTU. Some airlines bring in spare parts such as vanes or blades to get them repaired or replaced during the shop visit. Others have the strategy to scrap most of the parts due to contract obligations or their own internal policies. These are not related to actual part deterioration and averaging remedies some of this.

On the other hand, there are downsides to the fleet approach. Since there was at most only 26 data points, more sophisticated models could not be chosen. Averaging across airlines removed the ability to maintain statistical significance, despite the fact that there were relatively many shop visits. In addition, it must be kept in mind that these averages are not indicative of true operator behaviour. As previously stated, in the calculation of the average scrap rate, the engines which had 100% scrap rate were omitted. Moreover, these averages do not represent the fact that some airlines had almost 200 shop visits during the considered five year period and other airlines had only one and lacked credibility. This variation in shop visits weakens some of the conclusions that can be drawn from the data. It is thus misleading to predict average scrap rates for operators when there is significant difference between the number of operator shop visits.

	High Pressure Turbine				Low Pressure Turbine				
Stage	**Blade 1**	**Blade 2**	**Vane 1**	**Vane 2**	**Blade 3**	**Blade 4**	**Blade 5**	**Blade 6**	**Blade 7**
Cross Validation Error	343,473	186,3725	17,95271	10,43099	121,8654	77,15165	227,6863	13,04222	33,66967

Figure 8: LOOCV mean square error

Mean square errors were calculated to be able to assess the models. In order to make the errors comparable with those of the single engine approach, the errors of each stage was multiplied by the number of parts in that stage squared. The goal of leave one out cross validation (LOOCV) was to determine how new operator scrap rates would be predicted. This in turn would aid MTU in devising maintenance contracts with new potential customers. This would allow the accurate rates to be set for the new contracts and help with part fabrication estimates. In the LOOCV, the prediction error increased because the model was generalized and fitted with less data points. In figure 9, the points represent how the model would predict the scrap rate for a new MTU customer. In the middle, the error is relatively small, however when operators have very low (almost zero) or relatively high (greater than 0.5) scrap rates, the model does not predict the rate as well. This can be attributed to the low number of data points with which the model was fitted. In cases where negative scrap rates are predicted, the value zero should be used instead.

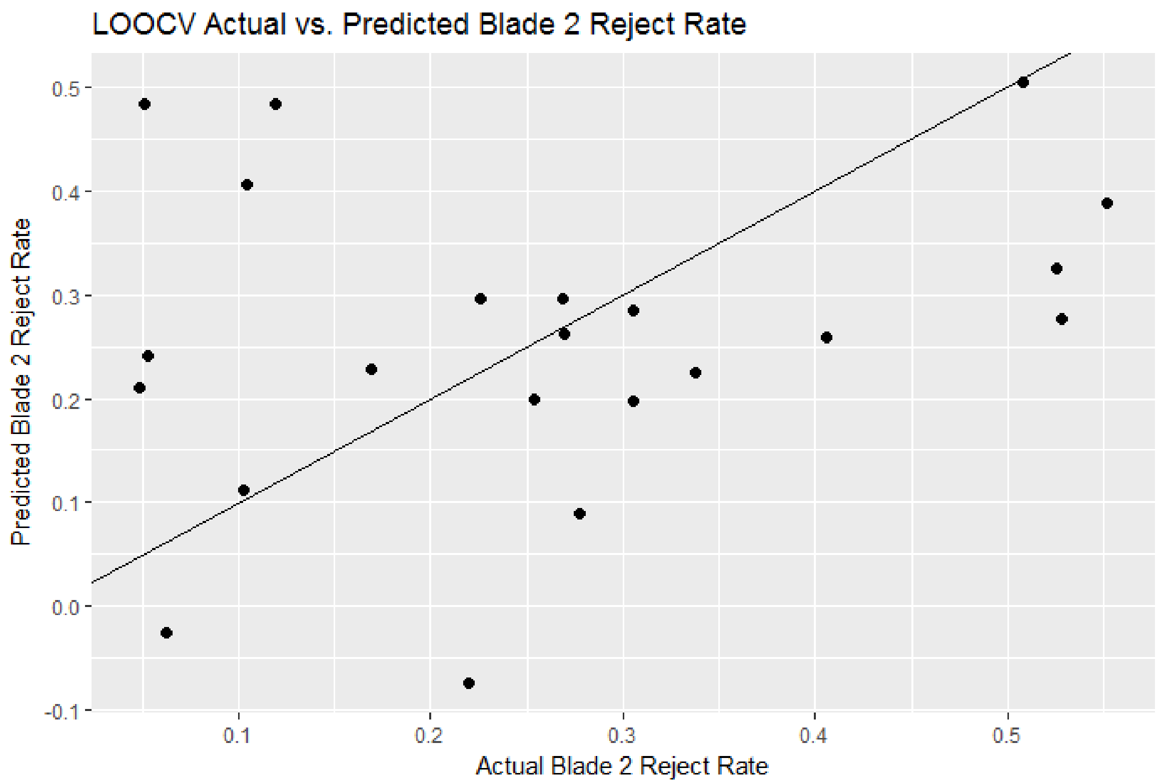

Figure 9: LOOCV for HPTB2

5.3 Single Engine Approach

The single engine approach compared different forms of regressions and regression trees to identify the best fit for the results. Regression used some feature selection from AIC and these models were compared to others from the same regression family using the prediction error. After this, the best models for each distribution were compared in the table below.

	High Pressure Turbine				Low Pressure Turbine				
Model Type \ Part	b1	b2	v1	v2	b3	b4	b5	b6	b7
Binomial	351.60	358.49	23.86	36.08	310.89	371.05	41.53	5.07	18.49
Poisson	299.22	534.54	19.16	17.72	498.02	386.51	176706.59	4.73	53.27
Negative Binomial	290.32	1637.92	44.85	16.92	578.16	388.86	125.83	6.62	216.80
Regression Tree	430.77	372.30	71.16	14.13	540.52	769.34	283.61	295.75	207.00

Figure 10: Comparison of single engine models by average of the squared prediction error in 15-fold cross-validation

The best model for each component is highlighted in green. This is based on the 15-fold cross-validation average prediction errors. The binomial and regression tree outputs were modified to be comparable to the counting regressions by multiplying the average prediction error by the total number of components in a set squared, as described below.

$$\sum_{k=1}^{15}\left(\frac{\sum_{i=1}^{n_k}(\hat{y}_i - y_i)^2}{n_k}\right)/15 = \sum_{k=1}^{15}\left(\frac{\sum_{i=1}^{n_k}(C \times (\hat{p}_i - p_i))^2}{n_k}\right)/15 = C^2 \times \sum_{k=1}^{15}\left(\frac{\sum_{i=1}^{n_k}(\hat{p}_i - p_i)^2}{n_k}\right)/15$$

where y_i represents the scrap count and p_i represents the scrap rate for engine i, n_k describes the number of records in the 15-fold cross-validation and the specified part set has normal quantity C

This method was chosen as it evaluates the models in order to find the best predictor of the scrap count or scrap rate. The counting models performed better for components in the HPT, but binomial regression excelled in the LPT. Some components have very low prediction errors, most notably the vanes and blade 6. These have very low scrap rates, with the third quartile being about 13% for vane 1 and under 5% for vane 2 and blade 6, making the predictions less significant. This is why it is safer to use the prediction errors for comparisons between models of the same component but not between components.

High Pressure Turbine

For the high pressure turbine blade and vane sets, counting regression proved more useful in the predictions. The models suggest that at least one of the particulate matter terms is significant as well as the NO_2, SO_2 and temperature values. As for regression trees, features that were the most important were SO_2, PM2.5, VTSN and VCSN. The counting models generally maintain more statistical significance for their factors than the binomial models and this supported many interaction terms between the flight cycles and environmental factors.

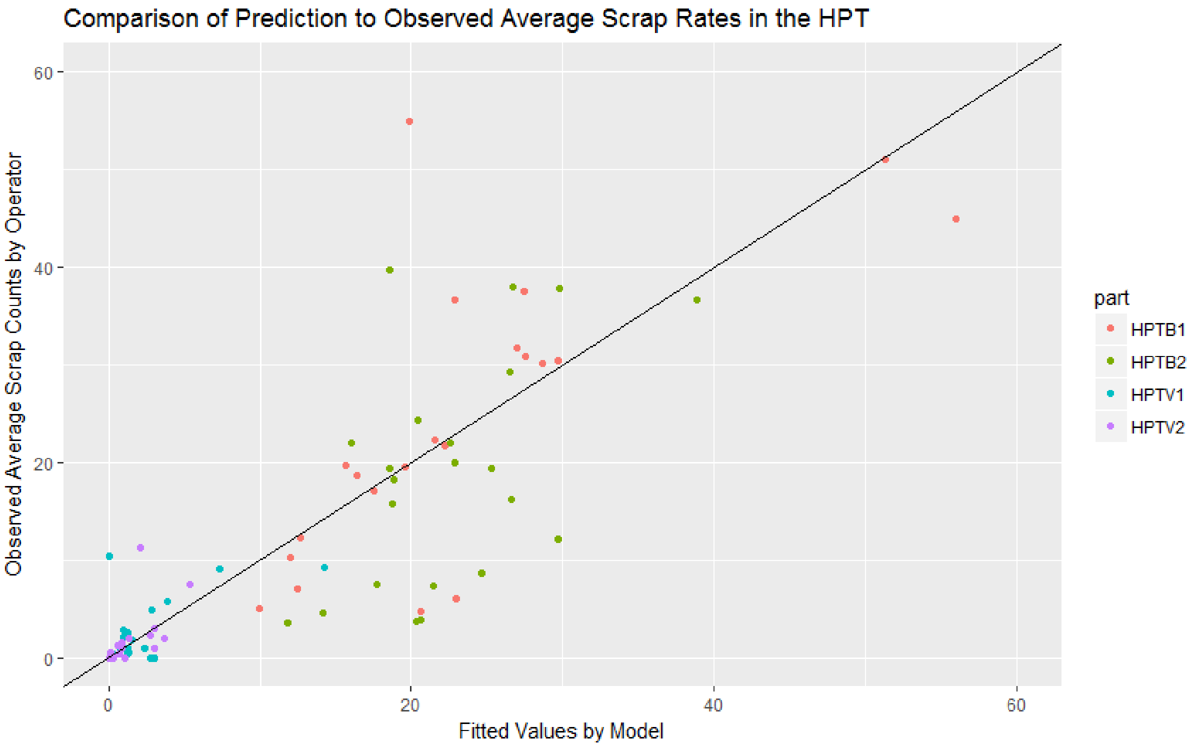

Figure 11: Comparison between model-predicted HPT values and fitted values by operator.

Using the average environmental and engine values by operator, the scrap counts were predicted and then compared with the actual averages in figure 11. Note that each component occurs in different quantities in an engine. However, the graph provides a relatively better visualization for comprehension than if scrap rates were used. It can be seen that very few results lie exactly on the line of perfect prediction. Some of this can be attributed to differences from the aggregation and the individual engines, but much of it comes from being unable to directly link the engine to the environmental factors it operated in. There are also differences between blades and vanes where it can be seen that vanes tend to have much lower scrap counts.

Low Pressure Turbine

The LPT sets were better predicted using regression with binomial models than counting regression. The C-Log-Log link also performed better for many of these models. As a result of poorer performance from the counting models, the final models tended to have fewer statistically significant factors. The particulate matter factors suggested PM10 was superior to PM2.5 most of the time and should be included. The other environmental factors were often not significant, however, NO_2 and SO_2 are present in some of the final models.

For regression trees, there were two stages, LPTB3 and LPTB6, where cross validation suggested maximal pruning, meaning the average of all scrap rates should be taken. This means that regression trees cannot explain these datasets in a sufficient way. For other components, the most important feature was the VCSN. For the LPT, parts tended to have lower scrap rates than HPT, even though they have more components in each set. This is likely because HPT blades are subject to greater stress than the other components. Additionally, while they are subjected to stricter and more expensive design conditions, operation tends to cause them to deteriorate more quickly. The model misjudges the some of the operators with higher scrap rates more so than those with lower. Once again, the variation within operators can make it difficult to compare with the aggregated values.

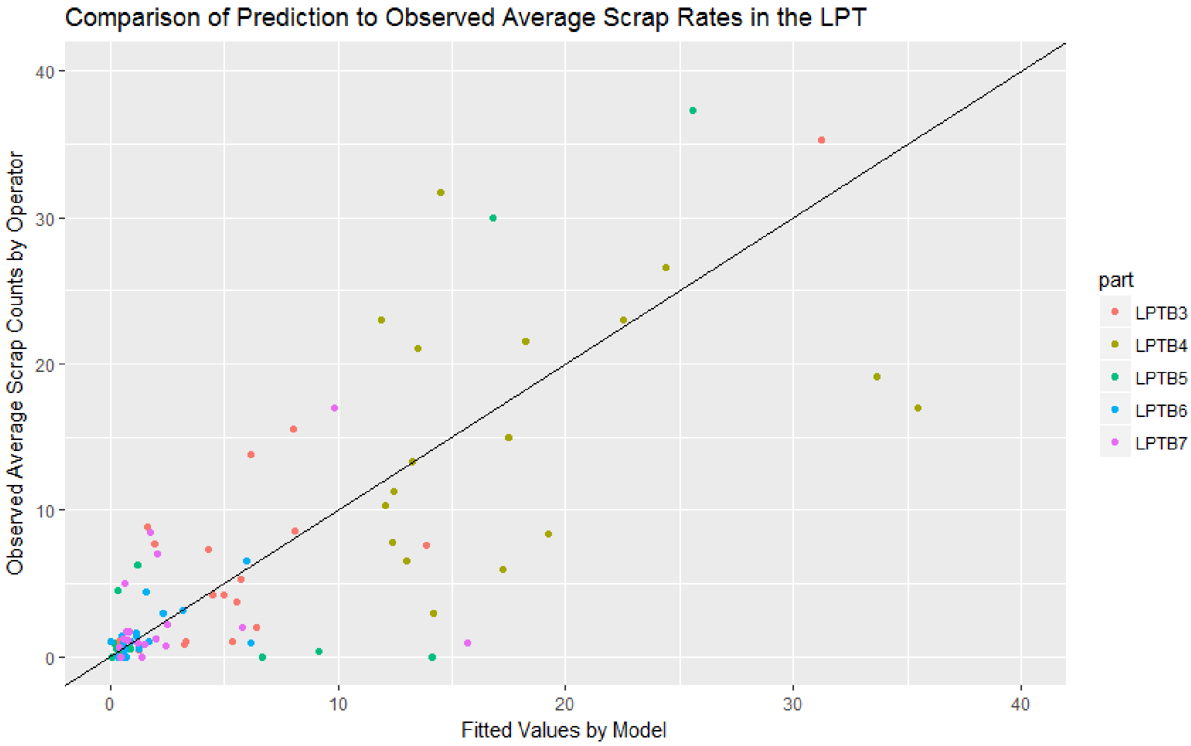

Figure 12: Comparison between model-predicted LPT values and fitted values by operator.

5.4 Model Assessment

The basis for model assessment lies in comparing mean square errors of the cross validation. The mean square error of prediction for the single engine approach is lower than that of the flight approach half of the time, and higher for the rest. The single engine approach had lower cross validation error for HPTB1, LPTB5, LPTB6 and LPTB7 while the fleet approach had lower cross validation errors for the rest: HPTB2, HPTV1, HPTV2, LPTB3 and LPTB4. These variations are due somewhat to chance and it is difficult to say what approach is the best. They all have benefits and drawbacks, and moving forward, they can provide insight on different aspects of damage mechanisms.

The benefit of the fleet approach is that it aims to predict average scrap rates for a new operator. Should this approach be used, it would be an accurate portrayal of how the scrap rate for new operators would be assessed by MTU. It effectively makes use of the variability in the environmental factors in order to distinguish between different types of operators. Unfortunately, due to the low number of data points, it is difficult to provide accurate predictions.

The biggest advantage of the single engine approach to modeling is the increase in the amount of available data for comparison. Currently, maintenance results of each engine for a certain component set can be used. This allowed for interaction terms to be used in the models while maintaining statistical significance. The models have more data to regress on and provide better predictions for factor value and significance. It also allowed precise use of the number of flight hours and cycles per engine, whereas the fleet approach was restricted to the aggregated values. There are a couple of restrictions and notes that need to be considered when using this method. As the data regarding flight paths of each engine could not be obtained, the operator average had to be used for each of their engines. This meant that the differentiation between engines of the same operator was dependent solely on engine-specific values, such as the flight hours or cycles it had conducted over its entire lifetime.

The enhanced distinction provided by the single engine approach outweighs its drawbacks and should be used in future predictions. Ideally, the record of each engine would allow better analysis of the effect of the environment on part deterioration. More shop visit data would also prove useful in order to enhance significance. Continuing with this modeling, it would be good to remove some of the interaction terms and try to untangle the importance of each factor. While this could negatively impact the prediction errors, the understanding of the underlying process would be beneficial for future research and part development.

6 Recommendations

The models presented cannot cover all of the complexities that exist in the operation of the engines. For the development of future models, there are some recommendations that could improve performance and accuracy to allow more for more meaningful and actionable conclusions. The first set of suggested improvements focuses on how improving connections between data sources could more precisely differentiate between operators and the flights that they conduct. The second focuses on new data sources that can be leveraged to enhance analysis.

6.1 New Data Connections

Through better connections in the data, engines can have more individually identifiable reasons for the damages that they incur. This can be done through their flight paths, their method of operation and knowledge around airline operating policies.

The initial idea was to connect the Flightradar24, ETM and maintenance data in order to track specific engines. This would have been possible since Flightradar24 provides the tail ID of each aircraft and the maintenance data tracks the engine serial number. The ETM data would have been the link between the two because it connects tail ID and engine serial number. While the current ETM data only allows this for one operator, if more information from other operators is collected in the future, this would lead to a deeper understanding of damage mechanisms. In this manner, it would be possible to assign precise environmental values to each engine, instead of using fleet averages.

For environmental data considerations, the time periods of the data could also lead to a better understanding of engine damage. For analysis, it would be more advantageous to know before the shop visits where the engine flew, instead of using 2017 Flightradar24 paths and assuming the operators flew these same routes throughout the last five years.

The existing maintenance data has large potential for future analysis. One of the limitations was the fact that the "Removal Reason" was not standardized and made it hard to classify the different reasons that engines would come in to be repaired. In addition, when each engine is taken in for a shop visit, the main inspection classifies some components as repair, instead of just serviceable or scrap. During secondary inspection and restoration, each part is then designated as scrap or serviceable. The current analysis does not distinguish between repair-scrap and scrap and only the end results are used. This distinction may provide insight for future analysis of whether or not parts are actually repaired and not eventually scrapped.

Finally, one of the most important considerations that was only partially explored, was the engine derate. As previously mentioned, the engine incurs the most damage at takeoff due to the thrust needed for liftoff, this level can be reduced by the pilot's selection of the derate. A greater understanding around how this derate is determined, or specific information about operator's regulations on its calculation, would be advantageous and allow for better predictions. Flightradar24 data is potentially useful with respect to this since it gives time and location snapshots of the aircraft. From this, derate could be estimated from the calculated speed, airport features and with an assumption on the weight of the aircraft.

6.2 New Data Sources

Increasing and enhancing the data that was used could allow for better predictions and more accurate interpretation of factor significance. For the environmental data, a new satellite, Sentinel 5P, will begin providing NO_2 and SO_2 measurements this summer. It will record levels of other compounds as well, allowing the potential for more new factors to be added to the model. Another new factor that could be included is the airport size. This could help represent an expansion factor on the measured environmental factors, as the greater amount of traffic on a runway, the relatively greater amount of waste compounds each successive airplane will be subjected to. This could be measured in either the number of flights per day, or the number of passengers serviced. By the number of passengers, heavier airplanes are penalized more strongly, which should be appropriate as they produce more toxic emissions that can more adversely affect other aircraft.

More policy specific knowledge about each MTU customer could lead to better predictions. For example, knowing the obligations MTU has to each airline in their maintenance contract would also be beneficial. This could be useful in distinguishing between engines that had true 100% scrap and those whose policy indicated that the customer can scrap all their parts when a certain condition is met. If not possible at the operator level, it would be interesting to know if airlines operating in similar regions, have similar considerations in their contracts.

Following up on the environmental data, it would be beneficial to access country-level data through government regulatory departments. This could be more granular and precise than the global level data. Expanding on environmental factors, it could be useful to track the changes over time, either monthly or a trend over years, to connect to respective flight distributions over time.

7 References

Images

NO2 Map from - http://uv-vis.aeronomie.be/news/20141027/
V2500 image from - https://image.slidesharecdn.com/v2500bsiissue01-150709152816-lva1-app6892/95/v2500-bsi-issue-01-21-638.jpg?cb=1436455774

Preliminary Analysis
"Untersuchungen zum Einfluss der Betriebsbedingungen auf die Schädigung und Instandhaltung von Turboluftstrahltriebwerken" by Matthias Müller

Environmental Data
Temperature Data
Current Weather Results (Travel Website - https://www.currentresults.com/Weather/index.php
World Bank - http://sdwebx.worldbank.org/climateportal/index.cfm?page=downscaled_data_download&menu=historical
Weather Base - http://www.weatherbase.com/
PM Data
World Bank - http://apps.who.int/gho/data/view.main.AMBIENTCITY2016?lang=en
WHO Satellite Data Set - http://www.who.int/entity/phe/health_topics/outdoorair/databases/airqualit_dimaq_pm25.csv
SO_2 **Data**
Socioeconomic Data and Applications Center (SEDAC) - http://sedac.ciesin.columbia.edu/data/collection/haso2
NO_2 **Data**
TEMIS – Tropospheric Emission Monitoring Internet Service from OMI Satellite http://www.temis.nl/airpollution/no2col/no2regioomimonth_v2.php?Region=9&Year=2016&Month=11

8 Appendix: Model Outputs

Fleet Approach Linear Regression Models

HPTB1	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	4,58E+01	1,61E+02	0.285	0.7791
VCSN	1,58E-02	9,58E-03	1.648	0.1158
PM2.5	4,68E+00	2,22E+00	2.109	0.0484

HPTB2	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-1,04E+02	1,07E+02	-0.978	0.34028
VTSN	1,12E-02	3,29E-03	3.400	0.00301
PM2.5	2,54E+00	1,41E+00	1.806	0.08682

HPTV1	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-7,83E+01	5,73E+01	-1.367	0.186847
VTSN	2,28E-03	1,69E-03	1.354	0.190983
PM10	1,77E+00	4,00E-01	4.414	0.000267

HPTV2	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-1,39E+02	4,19E+01	-3.312	0.00514
VTSN	2,61E-03	1,12E-03	2.331	0.03521
PM10	8,30E-01	2,60E-01	3.188	0.00658
NO2	3,14E-01	1,97E-01	1.594	0.13325
SO2	1,59E+00	8,02E-01	1.977	0.06803

LPTB3	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-5,81E+02	3,18E+02	-1.828	0.0852
VCSN	3,36E-02	1,84E-02	1.822	0.0861
vh2cr.all	2,55E+02	1,37E+02	1.868	0.0791
VTSN	-1,28E-02	8,28E-03	-1.541	0.1416

LPTB4	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	0.028085	0.077132	0.364	0.721
TEMP	0.006993	0.004053	1.725	0.105

LPTB5	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	7,83E+01	2,10E+02	0.373	0.7148
VCSN	-1,42E-02	6,62E-03	-2.140	0.0505
vh2cr.all	9,19E+01	6,43E+01	1.429	0.1749
SO2	-4,16E+00	2,27E+00	-1.831	0.0884
TEMP	7,53E+00	4,92E+00	1.530	0.1482

LPTB6	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	1,00E+02	2,04E+01	4.911	0.000188
VCSN	-4,65E-03	1,06E-03	-4.385	0.000533
PM2.5	-5,56E-01	3,59E-01	-1.548	0.142534
PM10	3,26E-01	2,28E-01	1.427	0.174097

LPTB7	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	2,12E+02	4,05E+01	5.242	0.000159
VCSN	-1,17E-02	1,62E-03	-7.257	6.39e-06
PM2.5	-1,94E+00	5,30E-01	-3.659	0.002887
PM10	1,03E+00	3,32E-01	3.094	0.008552
SO2	-6,75E-01	5,46E-01	-1.236	0.238321
TEMP	2,68E+00	1,77E+00	1.510	0.155047

Binomial Models

HPTB1	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-5.250	1.703	-3.083	0.0021
TEMP	0.222	0.089	2.497	0.0125
PM2.5	0.178	0.063	2.838	0.0046
NO2	-0.022	0.016	-1.355	0.1754
VCSN	-7.10E-05	3.27E-05	-2.169	0.0301
VTSN	3.04E-05	1.33E-05	2.292	0.0219
TEMP:PM2.5	-0.008	0.003	-2.616	0.0089
NO2:VCSN	1.87E-06	1.01E-06	1.86	0.0628

HPTB2	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-3.448	4.323	-0.798	0.4251
PM2.5	0.021	0.022	0.936	0.3492
SO2	-0.098	0.129	-0.759	0.4476
VCSN	-1.60E-04	1.83E-04	-0.872	0.3830
SO2:VCSN	9.91E-06	5.79E-06	1.713	0.0866

HPTV1	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-9.725	1.659	-5.861	4.61E-09
NO2	0.018	0.006	2.843	0.00448
TEMP	0.292	0.067	4.339	1.43E-05
VCSN	4.87E-05	2.65E-05	1.837	0.06622

HPTV2	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-8.252	1.410	-5.851	4.88E-09
PM10	0.021	0.013	1.635	0.102
VTSN	1.08E-04	2.54E-05	4.263	2.02E-05

LPTB3 cloglog	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-6.455	2.230	-2.894	0.0038
NO2	-0.043	0.042	-1.014	0.3108
VCSN	9.12E-05	6.88E-05	1.326	0.185
VH2CR_Total	0.982	0.466	2.107	0.0351
NO2:VCSN	2.81E-06	1.97E-06	1.431	0.1525

LPTB4 cloglog	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-2.188469	0.324	-6.747	1.51E-11
PM10	0.010799	0.007	1.552	0.121

LPTB5	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-1.022	2.731	-0.374	0.7083
PM10	0.037	0.018	2.04	0.0414
SO2	-0.114	0.065	-1.764	0.0777
VCSN	-4.54E-04	2.24E-04	-2.026	0.0428
VTSN	1.72E-04	9.70E-05	1.776	0.0758

LPTB6	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-6.182	5.948	-1.039	0.299
PM2.5	0.092	0.124	0.748	0.454
VCSN	1.03E-04	3.01E-04	0.342	0.732
VH2CR_Total	0.029	1.275	0.023	0.982
PM2.5:VCSN	-5.26E-06	8.73E-06	-0.602	0.547

LPTB7 cloglog	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-8.0076	2.773	-2.887	0.00389
PM10	0.0220	0.013	1.709	0.08753
VH2CR_Total	1.4564	1.091	1.335	0.18186

Poisson Models

HPTB1	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	24.590	2.63E+00	9.363	< 2e-16
TEMP	-1.032	1.38E-01	-7.506	6.08E-14
NO2	-0.352	5.31E-02	-6.638	3.19E-11
PM2.5	-0.290	6.33E-02	-4.575	0.000
SO2	-0.068	9.93E-03	-6.834	0.000
PM10	-0.004	3.04E-03	-1.471	0.141
VCSN	-1.23E-03	1.39E-04	-8.797	< 2e-16
VTSN	4.19E-05	4.44E-06	9.427	< 2e-16
VH2CR_Total	-0.344	6.58E-02	-5.237	0.000
TEMP:NO2	0.019	3.06E-03	6.321	0.000
TEMP:PM2.5	0.016	3.48E-03	4.559	0.000
NO2:PM2.5	0.004	1.42E-03	3.130	0.002
TEMP:VCSN	5.42E-05	7.23E-06	7.498	0.000
NO2:VCSN	1.70E-05	2.79E-06	6.086	0.000
PM2.5:VCSN	2.00E-05	3.48E-06	5.755	0.000
SO2:VCSN	3.54E-06	6.36E-07	5.564	0.000
TEMP:NO2:PM2.5	-2.35E-04	7.77E-05	-3.027	0.002
TEMP:NO2:VCSN	-8.75E-07	1.63E-07	-5.370	0.000
TEMP:PM2.5:VCSN	-9.38E-07	1.87E-07	-5.018	0.000
NO2:PM2.5:VCSN	-2.07E-07	7.42E-08	-2.784	0.005
TEMP:NO2:PM2.5:VCSN	9.28E-09	4.05E-09	2.291	0.022

HPTB2	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	0.662	3.75E-01	1.766	0.077
NO2	0.029	5.52E-03	5.199	0.000
PM2.5	-0.012	1.22E-02	-0.989	0.323
PM10	-0.003	4.26E-03	-0.754	0.451
SO2	0.004	4.58E-03	0.886	0.375
TEMP	0.108	2.16E-02	5.013	0.000
VCSN	-1.63E-04	3.43E-05	-4.748	0.000
VTSN	5.44E-05	6.19E-06	8.778	< 2e-16
VH2CR_Total	-0.220	6.59E-02	-3.332	0.001
NO2:PM2.5	-0.001	1.84E-04	-3.689	0.000
NO2:VCSN	2.00E-06	5.36E-07	3.738	0.000
PM2.5:VCSN	9.84E-06	1.26E-06	7.833	0.000
PM10:VCSN	-8.61E-07	3.67E-07	-2.349	0.019
SO2:VCSN	2.14E-06	6.07E-07	3.521	0.000
TEMP:VCSN	-6.92E-06	1.76E-06	-3.939	0.000
NO2:PM2.5:VCSN	-9.24E-08	1.76E-08	-5.238	0.000

HPTV1	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-3.384	5.43E+00	-0.623	0.533
NO2	0.008	3.88E-03	2.056	0.040
PM2.5	-0.030	1.16E-02	-2.562	0.010
PM10	-0.480	9.11E-02	-5.266	0.000
SO2	-0.437	1.92E-01	-2.272	0.023
TEMP	0.534	2.59E-01	2.064	0.039
VCSN	-1.99E-04	3.10E-04	-0.641	0.521
VH2CR_Total	0.528	1.07E-01	4.941	0.000
PM10:SO2	0.023	3.14E-03	7.238	0.000
PM10:TEMP	0.013	3.32E-03	3.802	0.000
SO2:TEMP	0.011	1.05E-02	1.041	0.298
NO2:VCSN	4.32E-07	1.91E-07	2.257	0.024
PM10:VCSN	2.94E-05	7.42E-06	3.966	0.000
SO2:VCSN	2.80E-05	1.03E-05	2.708	0.007
TEMP:VCSN	-1.45E-05	1.46E-05	-0.998	0.318
PM10:SO2:TEMP	-7.77E-04	1.14E-04	-6.798	0.000
PM10:SO2:VCSN	-1.31E-06	2.78E-07	-4.703	0.000
PM10:TEMP:VCSN	-7.69E-07	2.60E-07	-2.953	0.003
SO2:TEMP:VCSN	-7.63E-07	5.47E-07	-1.394	0.163
PM10:SO2:TEMP:VCSN	4.57E-08	1.03E-08	4.443	0.000

HPTV2	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-8.397	1.53E+00	-5.487	0.000
PM10	0.019	5.11E-03	3.634	0.000
PM2.5	-0.057	1.80E-02	-3.157	0.002
NO2	-0.015	1.98E-02	-0.766	0.443
TEMP	0.247	9.08E-02	2.718	0.007
VCSN	0.000	7.39E-05	2.402	0.016
VTSN	0.000	6.03E-06	17.045	< 2e-16
NO2:TEMP	0.002	1.20E-03	1.6	0.110
PM2.5:VCSN	0.000	8.93E-07	2.666	0.008
TEMP:VCSN	-1.15E-05	4.45E-06	-2.578	0.010

LPTB3	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-19.450	2.49E+00	-7.818	0.000
PM10	0.039	4.92E-03	7.89	0.000
PM2.5	0.222	4.60E-02	4.819	0.000
SO2	0.271	4.50E-02	6.021	0.000
NO2	0.209	4.91E-02	4.252	0.000
TEMP	0.401	7.14E-02	5.613	0.000
VCSN	1.21E-03	1.47E-04	8.216	< 2e-16
VTSN	9.15E-05	5.33E-06	17.183	< 2e-16
PM2.5:SO2	-0.005	9.90E-04	-5.271	0.000
NO2:TEMP	-0.014	3.00E-03	-4.686	0.000
PM2.5:VCSN	-1.45E-05	2.74E-06	-5.291	0.000
SO2:VCSN	-1.72E-05	2.53E-06	-6.819	0.000
NO2:VCSN	-1.54E-05	2.76E-06	-5.566	0.000
TEMP:VCSN	-3.26E-05	4.62E-06	-7.042	0.000
PM2.5:SO2:VCSN	2.66E-07	6.14E-08	4.332	0.000
NO2:TEMP:VCSN	1.03E-06	1.69E-07	6.121	0.000

LPTB4	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-60.660	8.95E+00	-6.778	0.000
TEMP	2.875	4.57E-01	6.299	0.000
NO2	0.840	1.90E-01	4.428	0.000
PM10	1.351	1.17E-01	11.535	< 2e-16
PM2.5	-0.181	5.82E-02	-3.102	0.002
SO2	0.033	4.69E-02	0.705	0.481
VCSN	0.004	4.68E-04	7.512	0.000
VTSN	2.74E-05	4.84E-06	5.665	0.000
TEMP:NO2	-0.028	1.10E-02	-2.492	0.013
TEMP:PM10	-0.059	6.78E-03	-8.688	< 2e-16
NO2:PM10	-0.027	2.62E-03	-10.213	< 2e-16
TEMP:VCSN	-1.61E-04	2.45E-05	-6.557	0.000
NO2:VCSN	-4.84E-05	9.88E-06	-4.898	0.000
PM10:VCSN	-7.85E-05	6.85E-06	-11.468	< 2e-16
PM2.5:VCSN	1.01E-05	3.53E-06	2.853	0.004
SO2:VCSN	-3.87E-06	2.51E-06	-1.542	0.123
TEMP:NO2:PM10	0.001	1.55E-04	7.428	0.000
TEMP:NO2:VCSN	1.62E-06	5.92E-07	2.742	0.006
TEMP:PM10:VCSN	3.52E-06	4.07E-07	8.646	< 2e-16
NO2:PM10:VCSN	1.57E-06	1.49E-07	10.549	< 2e-16
TEMP:NO2:PM10:VCSN	-6.95E-08	9.26E-09	-7.499	0.000

LPTB5	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-54.870	6.85E+00	-8.005	0.000
TEMP	2.622	3.52E-01	7.443	0.000
PM2.5	0.362	1.47E-01	2.466	0.014
PM10	0.109	4.46E-02	2.456	0.014
SO2	-0.160	5.21E-02	-3.067	0.002
NO2	0.326	6.70E-02	4.865	0.000
VCSN	0.002	2.39E-04	8.226	< 2e-16
VTSN	1.59E-04	2.63E-05	6.05	0.000
TEMP:PM2.5	-0.023	4.69E-03	-4.834	0.000
SO2:NO2	-0.005	1.71E-03	-3.061	0.002
TEMP:VCSN	-1.17E-04	1.46E-05	-8.026	0.000
PM2.5:VCSN	9.17E-06	5.96E-06	1.539	0.124
PM10:VCSN	-5.16E-06	2.99E-06	-1.723	0.085
SO2:VCSN	1.23E-05	2.61E-06	4.716	0.000
NO2:VCSN	-1.21E-05	2.48E-06	-4.878	0.000

LPTB6	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	2.799	1.62E+00	1.732	0.083
PM10	-0.146	4.35E-02	-3.348	0.001
PM2.5	0.302	7.72E-02	3.908	0.000
SO2	-0.096	2.09E-02	-4.595	0.000
NO2	-0.113	2.67E-02	-4.224	0.000
TEMP	-0.177	6.45E-02	-2.738	0.006
VCSN	1.40E-04	4.04E-05	3.472	0.001
VH2CR_Total	0.595	2.04E-01	2.919	0.004
PM2.5:SO2	0.002	5.14E-04	3.579	0.000
NO2:TEMP	6.84E-03	1.59E-03	4.303	0.000
PM10:VCSN	1.16E-05	2.99E-06	3.88	0.000
PM2.5:VCSN	-2.77E-05	5.77E-06	-4.79	0.000

LPTB7	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-9.4620	3.87E+00	-2.448	0.014
PM10	0.0383	1.36E-02	2.815	0.005
TEMP	0.1270	1.97E-01	0.644	0.519
SO2	0.1153	1.42E-01	0.81	0.418
NO2	-0.6955	2.24E-01	-3.1	0.002
PM2.5	-0.0846	2.97E-02	-2.845	0.004
VCSN	0.0003	2.14E-04	1.545	0.122
VTSN	6.37E-05	1.52E-05	4.192	0.000
TEMP:SO2	0.0016	4.26E-03	0.371	0.711
TEMP:NO2	0.0602	1.48E-02	4.069	0.000
SO2:NO2	-0.0064	4.44E-03	-1.451	0.147
TEMP:VCSN	1.70E-07	1.18E-05	0.014	0.989
SO2:VCSN	-1.01E-05	4.92E-06	-2.042	0.041
NO2:VCSN	3.57E-05	1.06E-05	3.364	0.001
TEMP:SO2:NO2	-0.0002	1.34E-04	-1.543	0.123
TEMP:NO2:VCSN	-3.22E-06	7.12E-07	-4.52	0.000
SO2:NO2:VCSN	5.94E-07	1.89E-07	3.152	0.002

Negative Binomial Models

HPTB1	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	7.599	5.25E-01	14.462	< 2e-16
TEMP	-0.182	1.98E-02	-9.200	< 2e-16
NO2	-0.033	6.85E-03	-4.827	1.38E-06
PM2.5	9.91E-04	8.77E-03	0.113	0.910
SO2	-0.026	7.68E-03	-3.394	0.001
VCSN	-3.59E-04	3.00E-05	-11.937	< 2e-16
VTSN	2.18E-05	3.81E-06	5.740	0.000
VH2CR_Total	-0.109	5.98E-02	-1.819	0.069
NO2:PM2.5	6.87E-04	2.31E-04	2.968	0.003
TEMP:VCSN	1.33E-05	1.15E-06	11.607	< 2e-16
NO2:VCSN	2.48E-06	4.05E-07	6.125	0.000
PM2.5:VCSN	8.48E-07	5.34E-07	1.588	0.112
SO2:VCSN	1.30E-06	4.65E-07	2.792	0.005
NO2:PM2.5:VCSN	-5.01E-08	1.35E-08	-3.710	0.000
Theta estimated at 1.418				

HPTB2	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-15.510	4.61E+00	-3.365	0.001
PM2.5	0.285	8.72E-02	3.274	0.001
TEMP	0.869	2.23E-01	3.891	0.000
NO2	0.027	7.52E-03	3.611	0.000
SO2	0.089	2.32E-02	3.845	0.000
VCSN	0.001	3.27E-04	3.431	0.001
VTSN	9.38E-05	2.20E-05	4.269	0.000
VH2CR_Total	-6.91E-01	2.30E-01	-3.008	0.003
PM2.5:TEMP	-0.016	4.45E-03	-3.568	0.000
PM2.5:VCSN	0.000	7.59E-06	-2.675	0.007
TEMP:VCSN	-5.50E-05	1.59E-05	-3.468	0.001
NO2:VCSN	-1.56E-06	4.48E-07	-3.495	0.000
SO2:VCSN	-7.23E-06	2.09E-06	-3.464	0.001
PM2.5:TEMP:VCSN	1.13E-06	3.69E-07	3.055	0.002
Theta estimated at 1.48				

HPTV1	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	1.985	7.10E+00	0.28	0.780
NO2	0.010	4.13E-03	2.483	0.013
PM10	-0.369	1.47E-01	-2.517	0.012
SO2	-0.571	2.20E-01	-2.592	0.010
TEMP	0.154	2.83E-01	0.543	0.587
VCSN	0.000	2.59E-04	0.109	0.913
VH2CR_Total	7.73E-01	2.79E-01	2.77	0.006
PM10:SO2	0.015	5.36E-03	2.723	0.006
PM10:TEMP	0.010	4.63E-03	2.203	0.028
SO2:TEMP	0.020	8.84E-03	2.281	0.023
PM10:VCSN	0.000	3.87E-06	2.172	0.030
SO2:VCSN	1.18E-05	5.61E-06	2.11	0.035
TEMP:VCSN	-2.49E-05	1.02E-05	-2.447	0.014
PM10:SO2:TEMP	-5.13E-04	1.85E-04	-2.773	0.006
PM10:SO2:VCSN	-1.58E-07	9.82E-08	-1.611	0.107
Theta estimated at 0.334				

HPTV2	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-8.698	1.07E+00	-8.122	0.000
NO2	0.020	4.05E-03	4.94	0.000
TEMP	0.204	4.66E-02	4.389	0.000
VCSN	3.54E-05	1.96E-05	1.802	0.072
VTSN	1.10E-04	1.23E-05	8.986	< 2e-16
Theta estimated at 0.3579				

LPTB3	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-8.152	1.67E+00	-4.883	0.000
TEMP	-0.021	4.48E-02	-0.459	0.646
PM10	0.031	7.75E-02	0.397	0.691
SO2	-0.008	4.33E-02	-0.196	0.845
NO2	-0.024	2.24E-02	-1.074	0.283
PM2.5	0.028	1.44E-01	0.197	0.844
VCSN	3.91E-04	8.51E-05	4.59	0.000
VTSN	-5.88E-05	4.09E-05	-1.438	0.150
VH2CR_Total	2.626	6.67E-01	3.935	0.000
PM10:VCSN	0.000	5.04E-06	-0.163	0.871
SO2:VCSN	-4.71E-07	2.63E-06	-0.179	0.858
NO2:VCSN	1.73E-06	1.21E-06	1.432	0.152
PM2.5:VCSN	-2.51E-06	9.39E-06	-0.267	0.789
Theta estimated at 0.6338				

LPTB4	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-1.583	3.01E+00	-0.527	0.598
PM10	0.015	6.04E-03	2.487	0.013
TEMP	0.315	1.54E-01	2.047	0.041
VCSN	2.09E-04	1.62E-04	1.294	0.196
VTSN	8.65E-05	3.28E-05	2.641	0.008
VH2CR_Total	-1.524	6.04E-01	-2.523	0.012
TEMP:VCSN	-1.78E-05	7.95E-06	-2.245	0.025
Theta estimated at 0.7394				

LPTB5	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-23.370	7.20E+00	-3.248	0.001
TEMP	1.041	4.34E-01	2.399	0.016
PM10	0.094	4.76E-02	1.967	0.049
SO2	-0.253	9.49E-02	-2.667	0.008
NO2	0.292	6.94E-02	4.206	0.000
VCSN	0.001	4.24E-04	2.079	0.038
SO2:NO2	-0.003	1.42E-03	-2.133	0.033
TEMP:VCSN	-4.22E-05	2.48E-05	-1.704	0.088
PM10:VCSN	0.000	2.67E-06	-2.099	0.036
SO2:VCSN	0.000	4.97E-06	3.449	0.001
NO2:VCSN	-1.13E-05	3.09E-06	-3.658	0.000
Theta estimated at 0.4913				

LPTB6	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-0.089	2.18E+00	-0.041	0.968
PM10	0.032	1.83E-02	1.77	0.077
PM2.5	0.080	4.29E-02	1.856	0.064
SO2	-0.051	2.17E-02	-2.35	0.019
NO2	-0.084	4.05E-02	-2.074	0.038
TEMP	-0.170	1.03E-01	-1.658	0.097
VCSN	1.95E-04	6.92E-05	2.817	0.005
VH2CR_Total	0.589	3.43E-01	1.719	0.086
NO2:TEMP	0.005	2.43E-03	2.038	0.042
PM2.5:VCSN	-7.71E-06	2.30E-06	-3.356	0.001
Theta estimated at 0.7005				

LPTB7	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-27.3300	8.38E+00	-3.262	0.001
TEMP	1.4020	4.40E-01	3.189	0.001
PM2.5	0.2373	1.04E-01	2.277	0.023
NO2	0.0827	4.26E-02	1.942	0.052
VCSN	0.0009	3.91E-04	2.205	0.027
VTSN	0.0001	2.22E-05	3.6	0.000
TEMP:PM2.5	-0.0109	4.76E-03	-2.288	0.022
TEMP:NO2	-4.47E-03	2.58E-03	-1.731	0.083
TEMP:VCSN	-5.26E-05	2.06E-05	-2.56	0.010
Theta estimated at 0.379				