

# Automated Failure Diagnosis in Aviation Maintenance Using eXplainable Artificial Intelligence (XAI)

Sophie ten Zeldam<sup>1,2,3</sup>, Arjan de Jong<sup>1</sup>, Richard Loendersloot<sup>2</sup> and Tiedo Tinga<sup>2,3</sup>

<sup>1</sup> *Netherlands Aerospace Centre (NLR), Amsterdam, the Netherlands*  
Arjan.de.Jong@nlr.nl

<sup>2</sup> *University of Twente, Dynamics based Maintenance group, Enschede, the Netherlands*  
r.loendersloot@utwente.nl  
t.tinga@utwente.nl

<sup>3</sup> *Netherlands Defence Academy, Military Technical Sciences, Den Helder, the Netherlands*  
sg.t.zeldam@mindef.nl

## ABSTRACT

An incorrect or incomplete repair card, typically used in aviation maintenance for reporting failures, may result in incorrect maintenance and make it very hard to analyse the maintenance data. There are several reasons for this incomplete reporting. Firstly, (part of) the information is often unknown at the moment the maintenance crew fills in the card. Also, the findings on repair cards are generally filled in as free-form text, making it difficult to automatically interpret the findings. An automatically assessed failure description will lead to more complete and consistent repair cards. This will also improve the efficiency of troubleshooting since this failure diagnosis can add information which would otherwise not be at the disposal of the maintenance crew at that time. This research utilises a data driven approach combining maintenance and usage data. The model will be based on Artificial Intelligence (AI) such that it is no longer necessary to completely understand the physics of a (sub)system or component. A newly proposed XAI (eXplainable AI) methodology, Failure Diagnosis Explainability (FDE), will be added to the model to provide transparency and interpretability of the assessed diagnosis. The assessed diagnosis is explained by checking whether a new failure matches the expected values of a certain diagnosis (class). On the other hand, when a failure is dissimilar to the expected values of a certain diagnosis (class), it is unlikely to be the actual diagnosis. The different steps towards this failure diagnosing model are applied to a case study with a main wheel of the RNLAF

Sophie ten Zeldam et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

(Royal Netherlands Air Force) F-16. This feasibility study already showed the value of this automated failure diagnosis model with an achieved accuracy of 81% of classifying a diagnosis. The proposed XAI methodology was able to explain the diagnosis assessed by the failure diagnosis model.

## 1. INTRODUCTION

A repair card is used in aviation maintenance to report a failure or anomaly and register it in the maintenance management system. An incorrect or incomplete repair card may result in incorrect maintenance and make it very hard to analyse the maintenance data. An example from practise is a helicopter Main Gear Box (MGB) removal due to a leakage found during a 500 hours inspection. The maintenance crew described the complaint as ‘defect, 500 hrs’. The component shop carried out an overhaul when a small repair could also have solved the problem. The overhaul, however, was unforeseen and there was no spare MGB available which resulted in a grounded helicopter. So the consequences of this incomplete repair card were additional maintenance costs and a decrease in availability.

This example is not unique in the aviation sector, it is rather common and there are several reasons behind this. Firstly, (part of) the information is often unknown at the moment the maintenance crew fills in the repair card. Also, the findings on repair cards are generally filled in as free-form text. As a result, repair cards may contain incorrect information and can be incomplete. A consequence of the incorrect and incomplete repair cards is that they are not practical for data analysis because it is difficult to automatically interpret the findings.

Therefore a model will be generated which can automati-

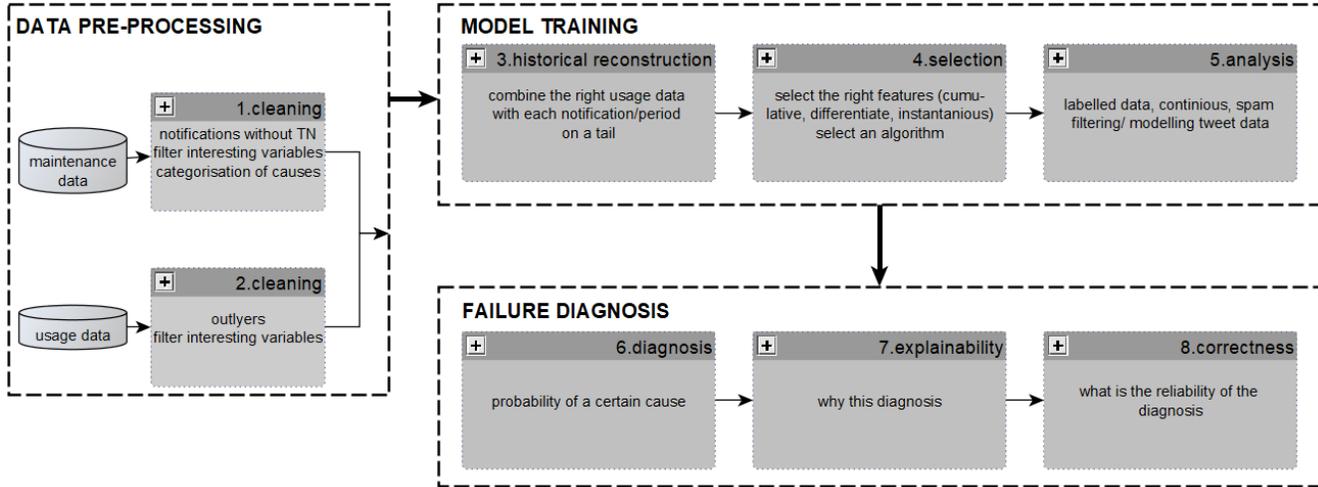


Figure 1. Functional diagram of an automated failure diagnosis model

cally assess a failure diagnosis based on usage data. An automatically assessed failure description will lead to more complete and consistent repair cards. This will improve the efficiency of troubleshooting since this failure diagnosis can add information which would otherwise not be at the disposal of the maintenance crew at that time. Conventional ways to link failures to the usage are physical models, but this research utilises a data driven approach combining maintenance and usage data, and Artificial Intelligence (AI) into a failure diagnosing model. AI is capable of recognising more patterns and relations than humans can. With this model, it is no longer necessary to completely understand the physics of a (sub)system or component. This data driven approach makes it difficult to establish causal relations between features. To convince the users of the model, a plausible explanation is needed to understand the cause of the failure. XAI (eXplainable AI) techniques will be implemented in the model to provide transparency and interpretability of the resulting diagnosis.

Figure 1 proposes the steps to achieve an automatic failure diagnosis based on usage data (sensor information from the flights) with XAI. This methodology is newly proposed in this research and will be discussed in the remainder of this paper. Section 2 describes which data have to be collected and in which form, to develop an automated failure diagnosing model (blocks 1,2). In Section 3, the training of the model is described (blocks 3,4,5). A historical reconstruction will be made and features and algorithms will be selected. The failure diagnosis will be discussed in Section 4 (blocks 6,7,8). The diagnosis will consist of an assessment of the possible causes, the explainability of this assessment and the correctness of the diagnosis. Finally, all these steps will be demonstrated in a case study in Section 5 where the feasibility of the model will be tested.

## 2. DATA PRE-PROCESSING

Before the model can be trained, the data need to be pre-processed so that it can be implemented in a ML (Machine Learning) algorithm. This raises the question which data, in which form and how much is required for proper training. The model will be based on historic maintenance and usage data which is sensor data from historic flights. In this research, the maintenance data will be used to label the failures. The usage data will be added to incorporate the historic usage of a component in the assessment of the diagnosis.

This section will discuss the algorithm prerequisites, the required variables and types of data. The requirements will differ per system and per component. This section provides some guidelines. Later, in the case study, a specific example will be elaborated on.

### 2.1. Algorithm prerequisites

There are no specific requirements for the dataset that is applied to ML algorithms. Actually Hua et al. (2005) mention that one should be wary of rules-of-thumb generalised from specific cases. The optimum sample size will differ per situation. But there are some general guidelines for the data:

- The optimal amount of features relative to the sample size depends on the algorithm and the feature-output distribution. Hua et al. (2005) show the relation between the number of samples, number of features and the error rate for various algorithms.
- The performance of an algorithm can be greatly influenced by the number of features. Therefore the amount of features used should be close to the optimal amount (Hua et al., 2005).
- A balanced labelled dataset (each class has about the same amount of samples) is preferable to avoid overfit-

ting on the class which contains more examples (Jain & Chandrasekaran, 1982).

Since the guidelines are not directional, the data have to be processed by reasoning in which format the data can be implemented in the ML algorithms. Engineering judgement is needed for the amount of features and data to be used.

## 2.2. Variable selection procedure

This research uses two data sources, namely usage data and maintenance data. The usage data comes from sensors in the airplane and consist of continuous, numeric values. This dataset contains information such as altitude, velocity, acceleration and outside temperature. The maintenance dataset consists mainly of text and categorised variables. This dataset is filled by the maintenance crew and contains information about the failure of a component. Besides the selection of variables from both datasets, the data probably will need cleaning. There are two types of errors in the datasets, namely measurement faults (e.g. an altitude of -100 m) and incompleteness (e.g. when a maintenance notification does not have a tail number (identification number) such that it cannot be linked to the usage data of an airplane).

Maintenance data entails information about the date of installation/removal of a part on a tail, preventive and corrective maintenance and the reason of a notification. The data have to be labelled before training. In order to limit the amount of labels, the failures with similar failure descriptions (and the same failure cause) will be grouped. E.g. 'removal due to being worn out' and 'component  $x$  is worn' should be in the same group. This grouping is also done because ML algorithms always require multiple events per group in order to train the model.

Variables are used to describe the usage. More variables can result in a more accurate description, but not all variables contain relevant information on the usage of the system. The usage variables will be different for each system, in each situation. In the case of a new system that has not been built yet, there is more freedom of selecting variables than in the case of an existing system. In general, to select the usage variables, it should be checked whether there is a relation between that variable and the loads on the system or the performance of the system (Tinga, 2010).

Also the features, i.e. the specific parts or details retrieved from a variable or signal, have to be determined. Is data required from the entire flight or only from take-off and landing or only during the flight? From all flights or only from the last one or last 10? This need will depend on the type of failure mechanism i.e. for corrosion, fatigue and wear all flights are required but an overload can probably be seen in (one of) the last flights. For the landing gear failures only the take-off and landings are interesting but for structural failures the entire flight (from take-off to landing) is needed.

## 2.3. Training and input data

A distinction is made between training data and input data. When this model is trained, data can be entered in the model and the model provides a diagnosis. Subsequently, the maintenance crew validates the diagnosis. The input data are now labelled and can be used to enrich (i.e. train) the model.

Since the usage data will be combined with the maintenance data (repairs, removals etc.), usage data of a certain period can be both training and input data. The usage data of the last year for a specific tail will be used as input data, e.g. when a failed component has been on a tail for one year. If another component, on the same tail failed four months earlier, then the usage data for the months before this failure can be used for training the model for that specific component.

## 3. MACHINE LEARNING MODEL

After pre-processing the data, the data from the different sources have to be combined and the model has to be trained for failure diagnosis. This section will discuss the data requirements for the ML algorithms and the procedure to select the variables from both maintenance data and usage data. Finally, the distinction between training data and input data will be discussed.

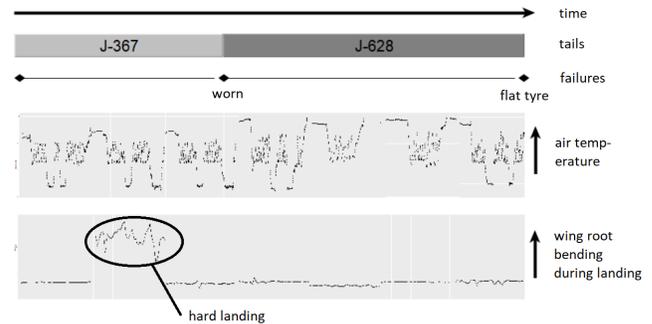


Figure 2. Reconstruction of a historical timeline for one component, installed on two different tails during a certain period of its operational life

### 3.1. Historical reconstruction

As can be seen in the functional diagram in Figure 1, the maintenance and usage data will be combined to come to a historical reconstruction where both sources are aligned to one timeline as is shown in Figure 2.

Most air forces and airlines exchange components between airplanes (tails), which is not common in every industry (e.g. process industry). Making a historical timeline of a component may show that a component has been installed on different tails. The usage data, in this case, need to be collected from different tails. Visualisation of the historical timeline is very important. This will give the maintenance

crew a quick overview of the history of a component which helps to judge the final diagnosis assessed by the model. Figure 2 shows the air temperature and wing root bending during landing to which the airplane was exposed during the lifetime of a specific component. As can be seen, there was a high wing root bending during one of the landings in the beginning of the wheel's life which probably indicates a hard landing.

### 3.2. Feature selection

All ML algorithms can only handle single values per case, therefore the usage of the component has to be expressed in features. Features are individual measurable properties or characteristics of a phenomenon being observed (Bishop, 2006). The features will be selected depending on the failure mechanism and the ML algorithm. E.g. when the velocity during take-off is one of the selected variables it is likely that the amount of data points is not equal for every landing and the data can therefore not be implemented directly in a ML algorithm. To get data strings of the same size, several options are possible: choose the lowest amount of data points and delete the superfluous points for the other take-offs; select the largest amount of data points and extra- or interpolate between data points for the other take-offs; or determine characteristics such as average, standard deviation and kurtosis for each take-off. Neural networks prefer to work with the whole data string (extra- or interpolation), but more simple algorithms can utilise characteristics.

### 3.3. Algorithms

During the training process, the dataset contains supervised (labelled) data, i.e. each notification is classified in a failure diagnosis category. Diagnosing a new failure will be a supervised multi-classification problem, for which there are some suitable algorithms (Wang & Xue, 2014; Kotsiantis, 2007; Wu et al., 2008):

- Support Vector Machines (SVM)
- $k$ -Nearest Neighbour (kNN)
- Naive Bayes
- Random forest
- Neural networks

The above mentioned algorithms are all suitable for the failure diagnosing model and will be compared after applying them individually to the data from the case study.

## 4. FAILURE DIAGNOSIS

The failure diagnosing model (from Figure 1) is trained with the chosen algorithm and pre-processed data. When a test example (or a new failure) is given to the model as input, it comes up with a probable failure diagnosis. This section will

discuss how this diagnosis can be presented to the maintenance crew and the possible ways to explain the diagnosis to help the maintenance crew to verify the reliability of the diagnosis. Finally, the correctness of the model will be discussed.

### 4.1. Diagnosis

Visualisation is a strong method to transfer data from the model to the user (here the maintenance crew). Therefore the output of the model will be visualised. Commonly, probabilities are shown as simple function plots, with either probability versus data value or value versus cumulative probability. The ubiquity of these representations make them easy to read and interpret, even if the user is unfamiliar with the subject (Potter et al., 2012). This is why the assessment results will be presented in a bar graph as shown in Figure 3.

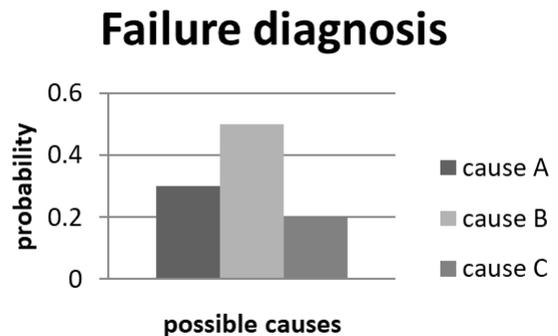


Figure 3. Example of a bar graph showing the outcome of the failure diagnosis model

### 4.2. Explainability

ML algorithms are often quite reliable, but sometimes not very explainable. An example is cancer diagnosis. Some algorithms are able to predict whether the patient has cancer or not, more accurately than doctors. But as long as humans do not understand how this algorithms made the assessment, it will not be used in practise since they do not trust them (Holzinger et al., 2017). Trust can be based on deterrence, on knowledge or on identification according to Lewicki & Bunker (1995). This research will strive to achieve trust based on knowledge. In order to give the maintenance crew this knowledge, it is very important to make the diagnosis assessment explainable.

#### 4.2.1. XAI methodologies proposed by DARPA

DARPA (Defense Advanced Research Projects Agency), the agency of the United States Department of Defense (Gunning & et al., 2016), is currently researching the explainability of various AI applications (Gunning & et al., 2016). Depending

on the chosen algorithm, some options are suggested which will be discussed below. To refer these options to a failure diagnosis, an example of a tyre failure of an airplane will be used. The three possible diagnoses in this example are wear, impact (overload) and other (for the remaining causes).

**Local Interpretable Model-agnostic Explanations (LIME)** Ribeiro et al. (2016) state that ‘LIME is an algorithm that can explain the assessments of any classifier or regressor in a faithful way’. One of the examples in this paper is the explanation of an image classification assessment. The image is a person, with the head of a dog, playing an acoustic guitar. The top three classes predicted were electric guitar, acoustic guitar and labrador. The model also shows which part of the picture is used for each assessment. Although this research is focussing on the explanation of image recognition, it can be applied to airplane component failure diagnosis as well. The original image will be a new failure and in the example of the tyre a possible top three of predicted classes is low profile, high amount of landings/take-offs and moderate or low landing/take-off velocities. This can help the maintenance crew to validate the diagnosis of wear. The model will not perform the diagnosis, only the explainability. The classes are pre-defined but the importance of each class can vary for each failure.

**Explanatory text** In the paper of Hendricks et al. (2016) explanatory text sentences are used to justify an assessment. The research is also focussed on explaining deep visual models. One example is a picture of a bird. The model gives: ‘This is a yellow breasted chat because this is a bird with a yellow breast and a grey head and back’. For the tyre example the explanation could be, this tyre is worn because it has lasted for over 200 landings/take-offs and now has a profile below  $x$  mm.

**Bayesian teaching** Bayesian teaching is the optimal selection of examples for machine explanation. So examples of the training data will be selected to explain the assessment of the model. In Gunning & et al. (2017) a picture of a child is used as input for the model. The output is that the face is angry because it is similar to certain examples (examples of kids with angry faces are given) and dissimilar to other examples (examples of kids with sad and happy faces are given). For the failed tyre, training notifications will be selected with similar feature values. E.g. the tyre had an impact because the amount of landings is (more or less) equal to the ones from these failures (followed by showing these similar notifications).

**Neural networks** Neural networks are used for classification but if the decisions towards this classification can be made visible it can also be used as an explanation for a diagnosis. A neural network contains an input layer, one or more hidden layers and an output layer. The input layer is fed with training data and the unlabelled input and the output gives

the assessment. But the hidden layers in between are seen as a black box. In Gunning & et al. (2017) a neural network is shown which predicts the type of animal. The training data contains images of all kinds of animals, the input is an image of a dog. They explain that the first layer neurons respond to simple shapes, the second layer neurons to more complex structures (a tail, paws, head). This will further increase until the  $n^{\text{th}}$  layer were the neurons respond to highly complex, abstract concepts. The output of the model was 10% wolf and 90% dog. For the example of the tyre, the simple shapes will probably be rough and clear separations between data (tail numbers, number of landings/take-offs) more complex layers will include values of velocity, variation in velocity etc.

**Decision trees** Decision trees are very powerful, but also simple and efficient for extracting knowledge from data. They are easy to interpret, understand and control by humans (Ertel, 2009). In Figure 4 an example is given for the tyre failure. The decision tree determines the nodes (number of landings/take-offs, surface of landing strip) and the limit values (15 and smooth/rough) by computing the information gain. For each layer in the tree, the feature with the highest information gain will become a node.

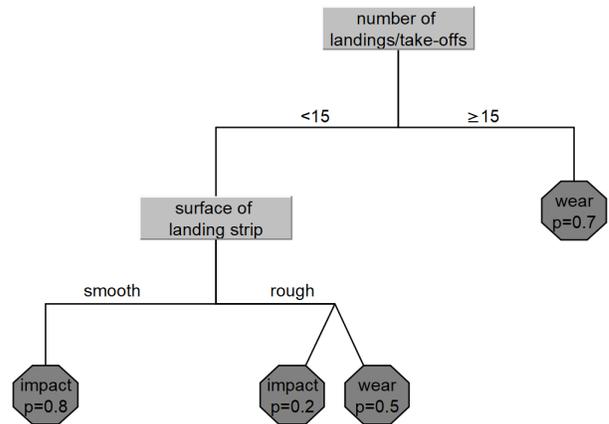


Figure 4. Simplified explanation of a tyre failure diagnosis by a decision tree

**4.2.2. Proposed XAI methodology: Failure Diagnosis Explainability (FDE)**

As discussed in Section 4.2.1, there are various methodologies to make a ML decision more explainable for the user (in this case the maintenance crew). As these methodologies are still in development and no detailed information has been available yet, a new XAI methodology will be proposed here: Failure Diagnosis Explainability (FDE). This methodology is based on the methods used for the different options suggested by DARPA. A common method to describe on what grounds the decision was based is with characteristics. In the example of the guitar playing dog, parts of the image were highlighted which were characteristic for a specific decision. For the

bird, the characteristics were described in text form and for the example of the child, pictures with similar and dissimilar characteristics were shown. The dog from the neural network was explained by its characteristic shapes and finally in a similar manner, the decision tree from Figure 4 explains tyre failures based on usage characteristics.

The failures are described by different features (as discussed in Section 3.2) so these will be used to explain the failure diagnosis assessed by the model. Since many features are implemented in the model, one can wonder which ones are most important: which features characterise a certain diagnosis the most. E.g. a certain value for feature  $f$  is very typical for a worn tyre. The variable importance of a model gives the (relative) importance of each variable (feature) for a trained model. But since this model includes various diagnoses (classes), the importance of each feature per class is not determined: the variable importance shows which feature is most important in deciding which class a failure belongs to, which is not equal to the characterisation of a certain class. To achieve this, new models will be trained in which, in each model, one class is appointed as 1 and the others as 0 (binary). E.g. class A is 1 and classes B and C are 0. Still, this does not precisely determine the variable importance per class, but it does yield the variable importance of a specific class relative to the other classes. A top five of most important features is determined per class (diagnosis) which can be found in Figure 6 in the case study. Before the model is trained, all values will be normalised using feature scaling. Note, the variable importance can only be determined for so called white box models as decision tree, random forest and Naive Bayes.

Finally, this methodology will show the user how much (and on which points) a new failure match with each diagnosis. The more the features matches the expected values of a specific diagnosis, the more likely it is that this failure belongs to that class. Also, if the features are very dissimilar to the values of a certain class, it is very unlikely that this failure belongs to that specific class. So, the reasoning of this methodology is similar to the Bayesian teaching methodology from Section 4.2.1. A failure is from class (diagnosis) A, because the features have similar values as class A. Contrary, this failure is not from classes (diagnosis) B and C, because the features have dissimilar values as class B and C (Gunning & et al., 2017).

The expectations for each feature are expressed by boxplots. If the value of a new failure falls within the 50% range (recognised by the box) of the boxplot, there is a good match. If the value falls within the 95% range (recognised by the whiskers of the plot), the values still matches but less compared to the 50% range. In case the value falls outside the 95% range, it is unlikely the failure belongs to that class, based on a 95% confidence interval; 2 sigma limit. The boxplots of the top five features of each class were represented in one graph. The

values of a new failure will be added to the graph to give the maintenance crew an indication of which variables match with the expectations of a certain diagnosis and which do not. A fit is made through these value points. If this line is equal to the  $x$ -axis, the new failure corresponds completely with the expected values. An example of such a plot can be found in Figures 7 and 8 in the case study.

For a better visual presentation, the medians of all boxplots were aligned with the same  $x$ -axis. Also, all values are, per variable, multiplied with their scaled importance  $w(f_j)$ . This means that a deviation on the most important variable will be enlarged compared to a deviation of the fifth most important variable. Equation 1 shows the computation of all absolute explanation values  $S_{i,j}$  (both for the boxplot and for a new failure), where the fraction normalises the value  $f_{i,j}$  where  $f$  is the sensor data from feature  $j$  and failure  $i$ . Following the median will be subtracted and finally, multiplied with the scaled importance. This methodology will be shown and tested in the case study in Section 5.3.2.

$$S_{i,j} = \left( \frac{f_{i,j} - \min(f_j)}{\max(f_j) - \min(f_j)} - \text{median}(f_j) \right) * w(f_j) \quad (1)$$

### 4.3. Correctness

The number of correctly classified examples is a performance measure for the diagnosing model. To measure this, part of the data is randomly separated before training and assigned as the test data. If the training data would be used for performance measurements, overfitting of the model cannot be noticed. After the initial training, when the model is in use, it will continuously enrich itself after every new example. The maintenance crew validates the diagnosis before it is added to the training data, hence the model will learn by updating itself. This process is also called incremental learning (Ertel, 2009). Incremental learning will improve the correctness of the model, unless only failures from the same class are added to the training data. In which case the dataset can become very unbalanced.

The correctness will be expressed in the accuracy of the algorithm, which means the percentage of correctly classified failures. A common practice within ML is to compare the achieved accuracy with a certain baseline. One common baseline for classification problems is the ‘most frequent’ which always classifies the most frequent label in the data set. When the accuracy of a ML algorithm is below this baseline, the algorithm will not be of any value.

## 5. CASE STUDY

In this section, the techniques proposed in this paper are applied to a case study to test the feasibility of these techniques. For the case study, the data from RNLAf (Royal Netherlands

Air Force) F-16s (fighter airplane) main wheels are used. The F-16 is equipped with a nose landing gear and a main landing gear. The nose landing gear consists of one nose wheel, the main landing gear of two main wheels. The RNLAf has a total of 68 F-16s using these main wheels. The tyres are removed for corrective maintenance or in case of other maintenance e.g. the wheel will be removed for when say a shock strut replacement is conducted. After the replacement of the shock strut, the original wheel will be installed again on the same tail. Common failures include the tyre being worn out, a flat tyre and a replacement due to contamination (mainly oil). There is no preventive maintenance on the wheels. The RNLAf does not retread their wheels, so a worn tyre will be directly disposed of. Flat tyres will be repaired if possible and contaminated tyres will be cleaned. After repair, the wheels will be mounted on an arbitrary tail. Most likely a different one, but occasionally it could be the same. Since the data from the used case is not public, the steps in data processing will be elaborated to provide more transparency.

### 5.1. Data pre-processing

The data has to be pre-processed before the model can be trained. The maintenance and usage data are prepared separately prior to being combined. For the pre-processing, the steps described in Section 2 will be followed. The main wheel is chosen for the case study since this component had one of the most removals and has few failure modes which are easy to understand.

#### 5.1.1. Maintenance data

The maintenance data is obtained from the Computerised Maintenance Management System (CMMS) SAP (Systems, Applications and Products). A repair card is filled in manually by the maintenance crew. The data is extracted from SAPs database and compiled to a CSV (Comma Separated Values) file. Thereafter the data is filtered by tail and by number of notifications per SN (Serial Number). Failure registrations without a specified tail are deleted since the tail is needed to link a failure to the usage data. If there is only one notification of a specific SN, the component is still in use (only the installation is reported, while no removal is reported). The installations do not mark a failure, therefore only the removals are kept. For this case study it is assumed that all removals are failures since cannibalisation (i.e. removal of sound components to be used on another tail) of wheels is not common for these F-16s. Also, for the case study, only one removal is eliminated due to these filters.

The data is labelled in three categories; 'flat', 'worn' and 'other and unclassifiable' (when the cause is unknown, different from the other two or the failure did not have enough information to be labelled). For labelling, several data fields were checked for words as 'worn', 'due', 'flat', 'deflating'

etc. This labelling is done automatically, but it is based on words which are found manually by going through the failure descriptions. Text mining may be an option to fully automate this process. There is no contamination class since these failures could not be traced from the repair cards or were not present in the data set. The labelled notifications (242) are combined with a dataset containing information about the tail and installation and removal date of each unique component. While combining the maintenance data with the usage data, only 144 removals are kept (elaborated in Section 5.2). So the amount of data is big (notifications and usage data together), but there are only a few failure cases which can also be combined with usage data. For data mining, 144 notifications is generally considered as a relatively small number, but in industries such as the aviation sector and in particular in the military sector, these numbers are rather common. A better registration of failures can easily lead to more notifications. Also, the class 'other and unclassifiable' is now considerably high (82 of the 144 failures), this is mainly due to many failures which did not have enough information to classify the notification. Preferably, this class is the smallest of the three, but this is a limitation of the data from current practise.

#### 5.1.2. Usage data

Sensor data from almost 7,000 flights are used. The data is extracted via a spool file, as it is a relational database, requiring a considerable amount of computing effort. The variables, which are appointed as related to the loads on the wheels, are: CAS (Calibrated Air Speed), strain from several strain gauges, weights, longitudinal and lateral accelerations, pitch and air temperature. There are more variables available but these are (for example) related to the engines or electronic components. Also, the phase of the flight (take-off, flight, landing) is assigned to each point. The variable selection is done by logical reasoning. In case of doubt, the variable was selected. This selection is done to avoid that the model will fit on unrealistic relations and to gain more causality in the model. E.g. the model could find a strong relation between maximum altitude and flat tyres but it later turns out that there were nails on the landing strip where flights with high altitude were trained. These type of unrealistic correlations can be prevented by the manual variable selection. The drawback could be, however, that a real but unexpected correlation is not revealed.

### 5.2. Machine Learning model

The ML algorithms require a dataset in which each row is a notification and the columns are the different features. After the data pre-processing, the maintenance and usage will be combined to put the data in the right format. Every notification should be combined with the usage related to this notification. No usage data was available for 98 of the 242 failures, leaving 144 failures for training and testing. The

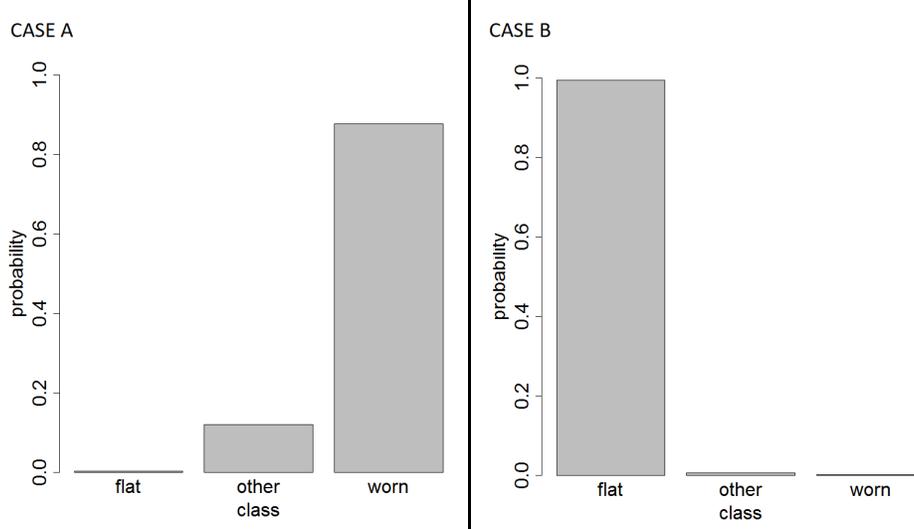


Figure 5. Probability of each class plotted for two specific cases, A and B

unavailability of the usage data can have various reasons e.g. the recorder was already full or the data are not added to the database due to the mission location of the airplane.

A notification can only include single values and not a (varying) range of data points. Therefore the variables of the usage data were translated to characteristics (as discussed in Section 3.2) with single values such as number of flights, average speed etc. The selected features are:

- number of flights
- average CAS during 1) take-off and 2) landing
- average and maximum strain gauge (called FS325) measuring the wing root bending during 1) take-off and 2) landing
- average and maximum strain gauge (FS374) measuring the fuselage bending during 1) take-off and 2) landing
- average and maximum strain gauge (BL120) measuring the wing tip bending during 1) take-off and 2) landing
- average lateral acceleration during 1) take-off and 2) landing
- average longitudinal acceleration during 1) take-off and 2) landing
- average total weight during 1) take-off and 2) landing
- average air temperature during 1) take-off and 2) landing
- the amount of calendar days between failures (i.e. the component age)

All these characteristics are calculated over the period of time since the first installation of the considered component. Now that the datasets are combined, a historical reconstruction can be made. This example has already been shown in Figure 2.

Since the algorithms can only handle numerical and categorical values, all text fields were deleted. Since these fields

Table 1. Accuracies obtained from various ML algorithms

Algorithm	Accuracy
Naive Bayes	62%
Neural network	43%
Support Vector Machine	69%
Random forest	81%

were already used for the labelling of the notifications, these data were already taken into account. Several algorithms (from Section 3.3) were applied to the training data. Table 1 shows the accuracies of the different algorithms. The first algorithm, Naive Bayes, probably performs moderately since this algorithm is simply too weak or too simple to detect the patterns of this data set. The neural network and SVM are probably overfitting on the training set. Overfitting can be seen when the accuracy on the test set is remarkably lower than on the training set. Since random forest gave the highest accuracy, this algorithm is used for the remaining steps.

### 5.3. Failure diagnosis

The model is now trained with the training set. After that, one failure from the test set will be fed into the model to assess a diagnosis for this failure. This single failure from the test set simulates a new failure. The results (diagnosis) of two new failures and how they should be interpreted (explained) are discussed in this subsection.

#### 5.3.1. Diagnosis

As discussed in Section 4.1 the assessed diagnosis is represented in a bar graph. Figure 5 shows the results for two specific failures which are most likely a ‘worn’ tyre (case A) and a ‘flat’ tyre (case B). This figure only shows the

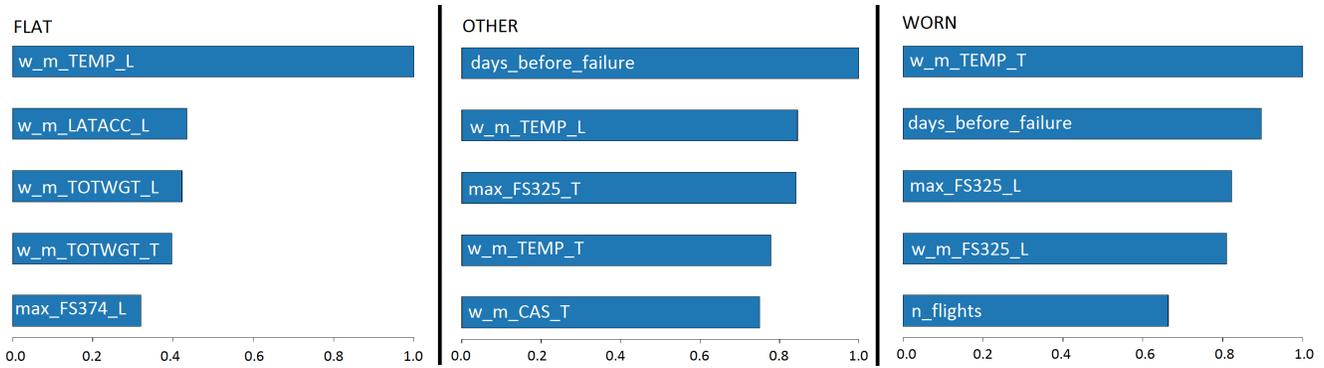


Figure 6. Variable importance plot for the different features

probability of each diagnosis yet does not explain the user why each diagnosis has a certain probability.

### 5.3.2. Failure Diagnosis Explainability (FDE)

As discussed in Section 4.2.2, the variable importances per class relative to the others had to be determined. The top five of scaled importances per class are shown in Figure 6. As can be seen, the weighted mean air temperature during landing is the most important feature to determine whether the tyre is ‘flat’ or not and the maximum fuselage bending during landing is the fifth.

Figure 7 and Figure 8 present the results of a ‘worn’ (case A) and a ‘flat’ tyre (case B) (as shown in Figure 5). As can be seen in Figure 7, the values of this new failure (case A) fall within the 50% range for four of the five variables for the ‘flat’ diagnosis. The first variable is on the edge of the 95% interval. For the diagnosis ‘other defect’, three values fall within the 50% range. Therefore it is unlikely that the new failure belongs to the class ‘flat’ or ‘other and unclassifiable’. For the diagnosis of a ‘worn’ tyre, the new failure falls within all 50% ranges and two variables even match the median (the third and the fourth variable). So it is likely that this new failure is ‘worn’ because all values correspond with the expectations. Even the fourth variable (the weighted average wing root bending during landing), which has a very tight range, matches.

In Figure 8, the values of case B completely correspond with the expectations of a ‘flat’ tyre failure. This makes it very likely, the assessed class will be ‘flat’. On the contrary, the new failure does not correspond with most of the variables from the classes ‘other and unclassifiable’ and ‘worn’.

In conclusion both, from case A and B, assessed diagnoses (shown in Figure 5) can be explained with the proposed XAI methodology FDE.

### 5.3.3. Correctness

After training the model, it achieved an accuracy of 81% on the test set which means that the notifications from the test set were classified to the correct diagnosis (‘flat’, ‘other and unclassifiable’ or ‘worn’) in 81% of the cases. As discussed in Section 4.3 this result will be compared to the ‘most frequent’ baseline, which is the percentage of occurrence of the biggest class. In this case the class ‘other and unclassifiable’ is the biggest class. 28 of the in total 42 cases from the test set are labelled as ‘other and unclassifiable’. If the model would (without any knowledge or training) assess every case as ‘other and unclassifiable’, it has an accuracy of 67%. The model showed its value already with the first preliminary test within this feasibility study, with a relative improvement of 21% compared to the baseline. For further research, the features can be changed, the data can be more balanced (optimal is equal amount per class) and more data can be used (also with incremental learning). These changes may lead to improvement of the model. Other features which could be taken into consideration are the number of exceedances of certain values and the variance or distribution of the data. By using other labelling techniques, consulting more data sources or filling in the repair cards more complete will decrease the number of unclassifiable diagnoses.

## 6. CONCLUSION

The combination of maintenance and usage data with (X)AI showed its value. An accuracy improvement of 21% (to a total of 81%) for classifying a diagnosis compared to the baseline is achieved already with this feasibility study. Even though the model was trained with a small amount of failures, which is very common in this industry. The usefulness of automatically diagnosing a failed tyre based on usage data can be questioned. After all, given that the tyre is failed, the distinction between a flat and worn tyre is easy to make. However, this case study served as a feasibility study. The tyres were one of the components with most data available and the failures were easy to understand which made validation

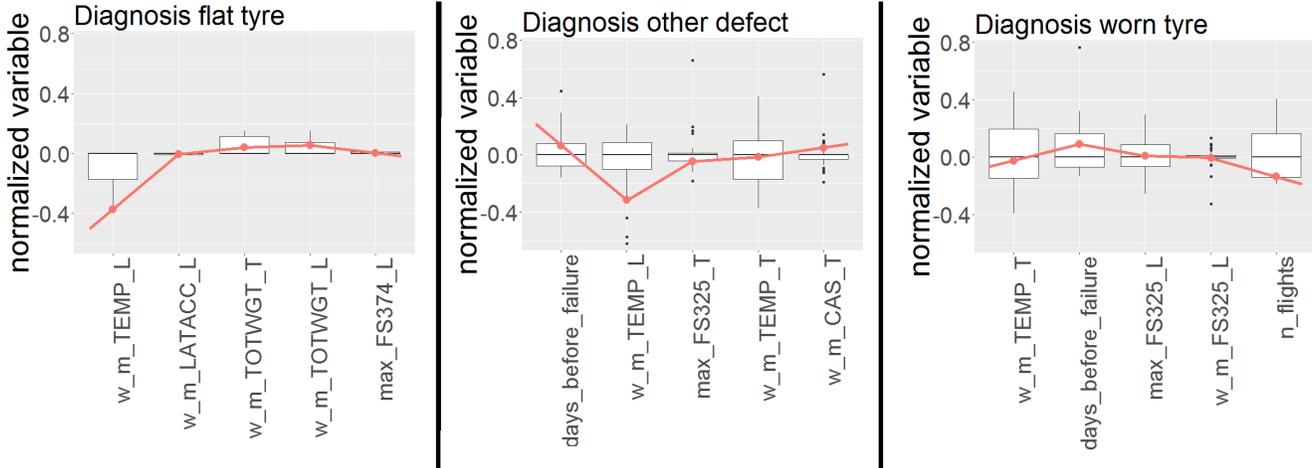


Figure 7. Boxplots explaining a specific case of a ‘worn’ tyre (case A)

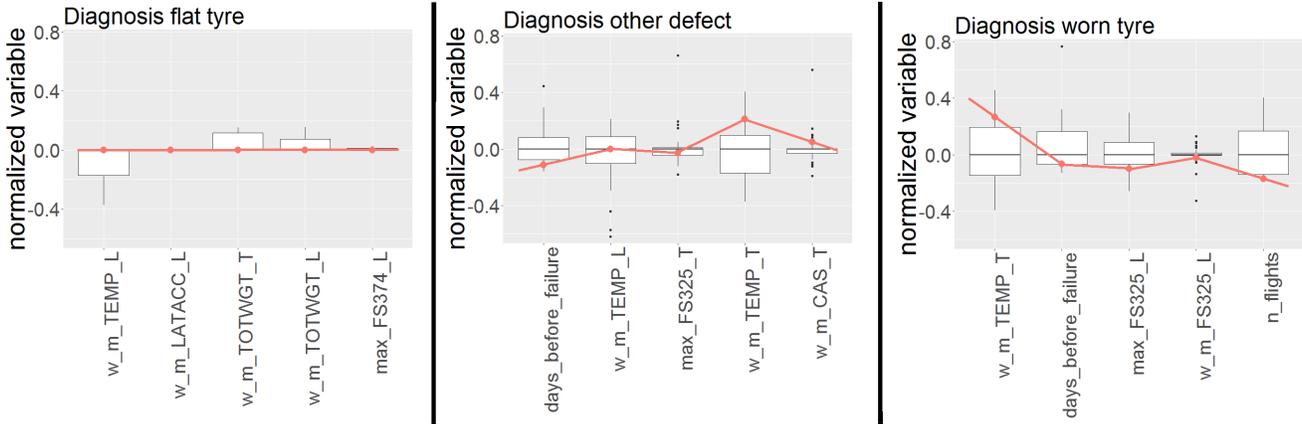


Figure 8. Boxplots explaining a specific case of a ‘flat’ tyre (case B)

of the model possible. Applying this model to components which are hard to diagnose (e.g. when the component is physically hard to reach) can really improve the efficiency of troubleshooting. With this research the requirements for the data and for such an analysis have become clear.

**7. RECOMMENDATIONS**

With the implementation of the case study, many difficulties arose. Most of them concerned the preparation of the data in such a way that the data could be implemented in the ML algorithm. Since many of these difficulties will be common for data-driven approaches in this industry, some general conclusions can be drawn. First of all, when the data are stored in a database, it does not mean they are also available for the preferred analysis. In case it is not available it should be made clear which steps need to be taken. The sensor data from the case study were stored in a relative database. Therefore it took significant effort to extract the data from the database. Secondly, the maintenance and usage data have to

be combined. To do this, there must be a possibility to link these datasets with each other (by tail, by date, by pilot, by airport or anything else). Thirdly, the data should be prepared for the algorithm (cleaning, labelling, inventing features etc.). This process can take a lot of effort and logical reasoning by selecting the features can add causality to the model. The way the data is recorded and stored will partly have influence on this effort. E.g. for cleaning and labelling a certain way of recording and storing can help, but it does not have influence on the effort to select the features.

**REFERENCES**

Bishop, C. (2006). *Pattern recognition and machine learning*. Springer.  
 Ertel, W. (2009). *Introduction to artificial intelligence*. Springer. (Section 8.4, 9.4)  
 Gunning, D., & et al. (2016). *Explainable artificial intelligence XAI* (Tech. Rep. No. DARPA-BAA-16-53). Defense

- Advanced Research Projects Agency (DARPA).
- Gunning, D., & et al. (2017). *Explainable artificial intelligence XAI ' - program update* (Tech. Rep. No. DARPA-BAA-16-53). Defense Advanced Research Projects Agency (DARPA). (slide 4, 15)
- Hendricks, L., Akata, Z., Rohrbach, M., & et al. (2016). Generating visual explanations. In *European conference on computer vision*.
- Holzinger, A., Biemann, C., Pattichis, C., & Kell, D. (2017). *What do we need to build explainable AI systems for the medical domain?* (Tech. Rep.).
- Hua, J., Xiong, Z., Lowey, J., Suh, E., & Dougherty, E. (2005). Optimal number of features as a function of sample size for various classification rules. *Bioinformatics*, 21(8), 1509-1515.
- Jain, A., & Chandrasekaran, B. (1982). Dimensionality and sample size considerations in pattern recognition practice. *Handbook of Statistics*, 2, 835-855.
- Kotsiantis, S. (2007). Supervised machine learning: A review of classification techniques. *Frontiers in Artificial Intelligence and Applications*, 160, 3-24.
- Lewicki, R., & Bunker, B. (1995). *Trust in relationships* (Tech. Rep. No. RS 95-7). Business Research College of Business.
- Potter, K., Kirby, R., Xiu, D., & Johnson, C. (2012). Interactive visualization of probability and cumulative density functions. *International Journal for Uncertainty Quantification*, 4(2), 397-412.
- Ribeiro, M., Singh, S., & Guestrin, C. (2016). why should i trust you? explaining the predictions of any classifier. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*.
- Tinga, T. (2010). Application of physical failure models to enable usage and load based maintenance. *Reliability Engineering and System Safety*(95), 1061-1075.
- Wang, Z., & Xue, X. (2014). *Support vector machines applications*. Springer. (Chapter 2)
- Wu, X., Kumar, V., Quinlan, J., Ghosh, J., & et al. (2008). Top 10 algorithms in data mining. *Knowledge and Information Systems*, 14(1), 1-37.

## BIOGRAPHIES

**Sophie ten Zeldam** obtained her Bachelor degree Military Systems & Technology, research group Maintenance Technology, at the Netherlands Defence Academy in 2016. Currently she is working on a Master degree Mechanical Engineering, specialization in Maintenance Engineering &

Operations, at the University of Twente. This research is part of her Master degree's final assignment which she will finish in July 2018.

**Arjan de Jong** received his BEng in Aeronautical Engineering from Haarlem Polytechnic and his Master Degree in Aviation Management from the University of Southampton. He worked for the Schreiner Aviation Group in the Netherlands and the CHC Helicopter Corporation in Vancouver, BC, Canada. He held various operational and strategic positions in maintenance, engineering and asset management. Arjan received his PhD in Aerospace Engineering & Management from the Technical University Delft. He now works at the Netherlands Aerospace Centre - NLR as Maintenance Engineering, Management & Technology team leader.

**Richard Loendersloot** obtained his Master degree in Mechanical Engineering, research group Applied Mechanics, at the University of Twente in 2001. He continued as a PhD student for the Production Technology group of the University of Twente, researching the flow processes of resin through textile reinforcement during the thermoset composite production process Resin Transfer Moulding. He obtained his PhD degree in 2006, after which he worked in an engineering office on high end FE simulations of a variety mechanical problems. In 2008 he returned to the University of Twente as part time assistant professor for Applied Mechanics. From September 2009 on he holds a fulltime position. Since then, his research started to focus on vibration based structural health and condition monitoring, being addressed in both research and education. He became part of the research chair Dynamics Based Maintenance upon its initiation in 2012. His research covers a broad range of applications: from rail infrastructure monitoring, to water mains condition inspection and aerospace health monitoring applications, using both structural dynamics and ultrasound methods. He is involved in a number of European and National funded research projects.

**Tiedo Tinga** received a Master degree in Applied Physics (1995) and a PDEng degree in Materials Technology (1998) at the University of Groningen. He did his PhD research during his work with the National Aerospace Laboratory NLR on the development of a multi-scale gas turbine material model. He received his PhD degree in 2009 from Eindhoven University of Technology. In 2007 he was appointed associate professor at the Netherlands Defence Academy and in 2016 became a full professor Life Cycle Management. In 2012 he also became a part-time full professor in dynamics based maintenance at the University of Twente. He now leads a research program on predictive maintenance and life cycle management in both institutes.