Improving Residual Generation for OBD in the Air Path of a Truck Engine

Antoine Berton

14th August 2003

Abstract

A method to improve model based diagnosis for the air path of a truck engine is presented. Originally, inaccuracies in both volumetric efficiency model and sensors limited the diagnosis performances. Two approaches are presented to improve the diagnosis performances. The first method is based on an elimination of part of the data according to the level of a certain operational variable. The second and preferred method make use of adaptive statistical charts built to generate residuals which show a reduced dispersion. Data taken from two different trucks working in various operating conditions were used to evaluate the proposed approaches. It was shown that the use of charts improved the diagnosis performances by approximately 25 %.

Contents

Introduction

This report presents the summary of the work achieved during a 4 months research cooperation to the Vehicular System Laboratory in University of Linköping, Sweden by PhD student Antoine Berton coming from University Laval, Canada. The work has also been achieved in cooperation with Scania, a multinational company and world leading truck manufacturer. This work follows in some ways the Masters thesis work of Petter Haraldsson [Har02].

The OBD problem

The emission requirements for heavy truck engines as well as for car engines are becoming stricter, especially in Europe and states like California. The incoming regulations of EURO 4 in 2005 includes restrictions on pollution and obligation of including an on-board diagnosis (OBD) system for truck enginesc̃iteeuro4. The goal of such a system is to make sure that the requirements on emissions are met during the whole truck operational life.

Gauges to measure all pollutant emission at the truck exhaust could allow to built a good OBD system. However, those gauges would be very expansive and are anyway not available. So it is necessary to infer emission problems by detecting other faults. Faults in the air path components of the engine is known to be linked to emissions problems. So it is precisely the kind of fault that, in the present study, will be aimed to detect with as good precision as possible.

The standard tools used in production vehicles [US93, Jur94] might not be sufficient to fulfill the requirements of the new regulations. One way to enhance the performance of fault diagnosis is to make use of more knowledge, using model-based fault diagnosis. The actual available model on the engine is a simple static model, based on the engine volumetric efficiency [Hey88].

A previous work on engine diagnosis [Har02] showed that using this raw model to generate residuals and perform fault detection is relatively inaccurate. The urgent need for a reliable engine diagnosis procedure then lead to the present attempt to improve the static model accuracy.

The Process

The engine of study is an experimental engine with exhaust gas recirculation (EGR). A schematic view of the engine air path is shown on Figure 1. The air entering the engine is measured by a air-mass flow-meter and the corresponding measurements is noted W_{in}. Then the air goes through a compressor and a

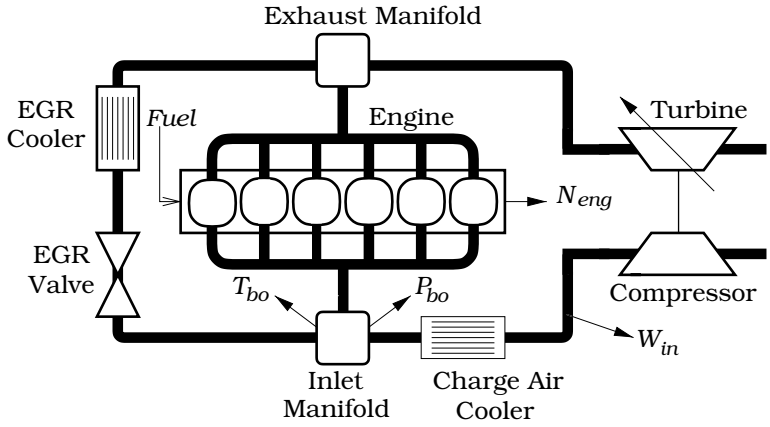

Figure 1: Schematic overview of the Scania test engine.

charge air cooler, enters the intake manifold where it is mixed with exhaust gases, and flows into the cylinders where the fuel is added. The exhaust gases are divided one part to drive the turbine and another part to be recycled via the EGR. The involved measurements for residuals generation will be the air mass flow (W_{in}), the boost pressure (P_{bo}) and boost temperature (T_{bo}) at the inlet manifold, as well as the engine speed (N_{eng}). Those measurements corresponds to the available sensors on the production version of the engine.

Studies are currently being made in order to accurately include the EGR path in the engine model. However, the present approved model does not allow to take into account the EGR component of the air path. Thus, the EGR valve will be shut for all the experiments presented in this paper.

Residual Generation for Diagnosis

The volumetric efficiency model infers the air-mass flow in the inlet based on other available measurements [Hey88]. Equation 1 illustrates this predicted air-mass flow (W_p).

$$W_p = \frac{P_{bo}\, N_{eng}\, n_{cyl}\, V_d\, \eta_{vol}(P_{bo}, N_{eng})}{120\, R_{air}\, T_{boost}} \tag{1}$$

In Equation 1, n_{cyl} is the number of cylinders, V_d is the displacement volume of the engine and R_{air} is the gas constant of air. η_{vol}, called the volumetric efficiency map, depends on both boost pressure and engine speed and must be calibrated from operational data. The air-mass flow is in kg/s, the engine speed is in RPM, the boost pressure is in Pascal and the boost temperature is in Kelvin.

In Equation 1, all variables are either known or measured, except for the volumetric efficiency. The volumetric efficiency follows a highly non linear behavior depending on P_{bo} and N_{eng}. In order to characterize this behavior, it is necessary to build a map (look-up table), from unfaulted data. The data that has been used to calibrate the map is the dataset *maverik*1. *Maverik* is the affective name of a truck, which with his little sister called *platina*, are the

devoted data providers for this study. Then, the fault detection will be based on the generation of a residual r:

$$r = W_{in} - W_p \tag{2}$$

The diagnosis procedure is not to be active all the time when the truck is working. The diagnosis starts under predefined conditions. These conditions must be so defined so that they appear in the standardized OBD test cycle. At a basic sampling rate of 20 Hz, an average diagnosis sample could contain about 15 000 measurements of each variable. Then, for each cycle, the test quantity for fault detection will be the average value of the residuals over this observation window.

Aimed strategy for fault detection

The design of fault detection for the truck engine is not an easy task. Both model and sensors contain a high level of innacuracy that create a lot of problems when aiming at a precise fault detection procedure. In his Masters thesis, Petter Haraldsson [Har02] partly solved those problems with *ad hoc* steps. Those different steps in the fault diagnosis method were for instance:

- Rejection of subsets of data where the static model seems to be inaccurate
- Noise reduction before residual generation
- Adaptive normalization of the residuals and determination of thresholds for statistical testing
- Outliers rejection

The aim of the present work will be to bring a more systematic approach for fault detection.

First, when dealing with fault detection, subsets rejection and outliers rejection is very delicate. It is easy to build a criteria that would bring the residuals distribution to look nicer but by the way, eliminate the fingerprint of any fault. Thus, it will be aimed in this work to be careful about rejecting data.

Then, it will be aimed to have a simple strategy that does not contain too much different steps before to bring a diagnosis.

Finally, more rigorous statistical contain will be introduce, as once the statistical properties of the residuals are well determined, thresholding problem doesn't exist anymore.

Outline

The first chapter presents a prospective work in which the data are assessed in different ways. The diagnosis problem in not yet discussed at this point. The different data available for residuals generation are shown in different ways. This chapter, that could be skipped by a diagnosis focussed reader, introduce how the ideas present in the second chapter emerged.

The second chapter illustrates the problem of residual generation. A first approach implying data selection based on fixed threshold of given operational

variables is illustrated. Then, the preferred method making use of statistical charts to improve the model is presented. Finally, the methods are compared and some expected results for fault detection are presented.

Chapter 1

Data analysis

Before to design any fault detection procedure, it is necessary to have a good knowledge of the available tools. In our case, as we intend to perform model-based fault detection, it is of prior importance to spend some time understanding the data and the process. Here is performed some analysis from the data set *maverik*1. This data set comes from the truck running on the Scania test track in Södertälje. The measurements are taken at a 20 Hz frequency during almost 90 minutes and are of course fault-free.

1.1 Data visualization

Before assessing the data with any statistical tool, it is interesting to have a quick look at it. This visualization operation aimed to have a rough idea of what the data is, in what range of values they are evolving, if there is redundant data, etc. For each figure, data will be shown for a short interval of around 8 minutes.

1.1.1 Data involved in the residual generation

The variables of interest in the purpose of residual generation are of course, those involved in the static model (P_{bo}, T_{bo}, N_{eng} and W_{in}). The data linked to those variables in the *maverik*1 dataset are called:

- eadc-dm-egr-s16
- eadc-dm-egrfilt-f32
- eadc-p-boostfilt-s16
- eadc-p-boostrawfilt-f3
- eadc-t-boostfilt-s16
- eess-n-avg10ms-s16

By simple observation, one can realize that the first two available data are measuring the same variable, with different filtering. Figure 1.1 and Figure 1.2 respectively show this air-mass flow measurements, for two different filtering. For both figures, the air-mass flow has been divided by 60 in order to have it in kg/s. The label of the data, as it seems related to EGR, has in fact nothing

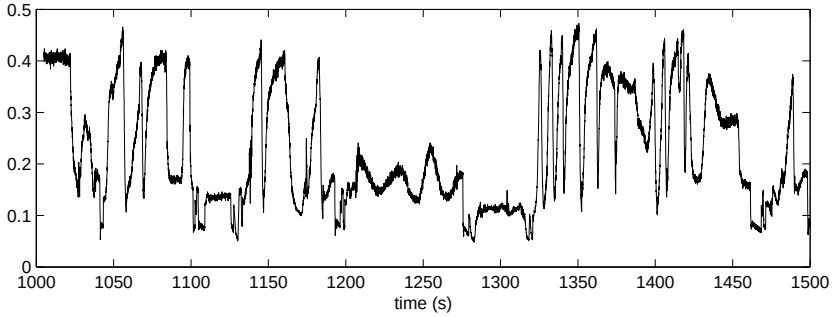

Figure 1.1: eadc-dm-egr-s16: air-mass flow measurement (kg/s)

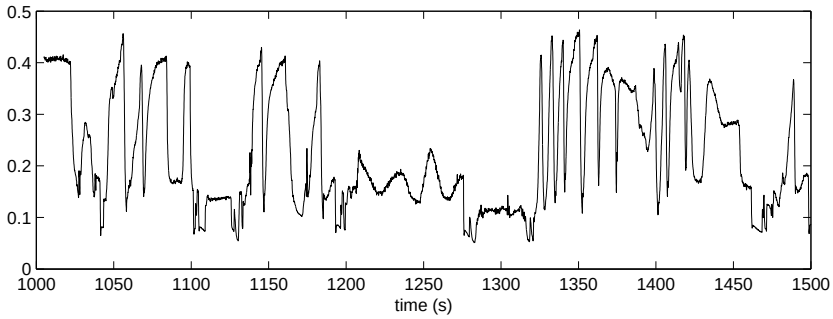

Figure 1.2: eadc-dm-egrfilt-s16: air-mass flow measurement with a stronger filter (kg/s)

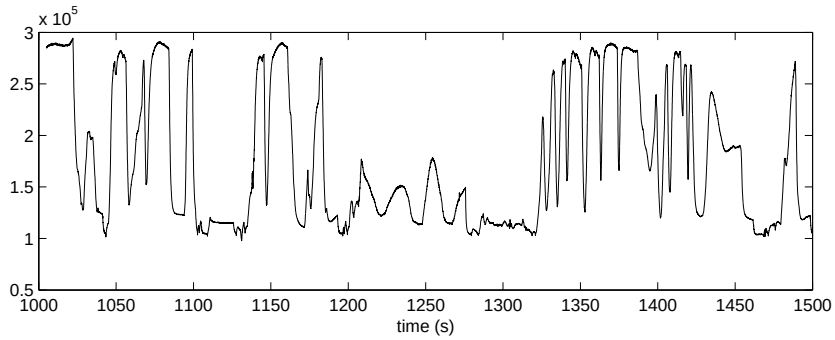

Figure 1.3: eadc-p-boostfilt-s16: boost pressure measurement (Pa)

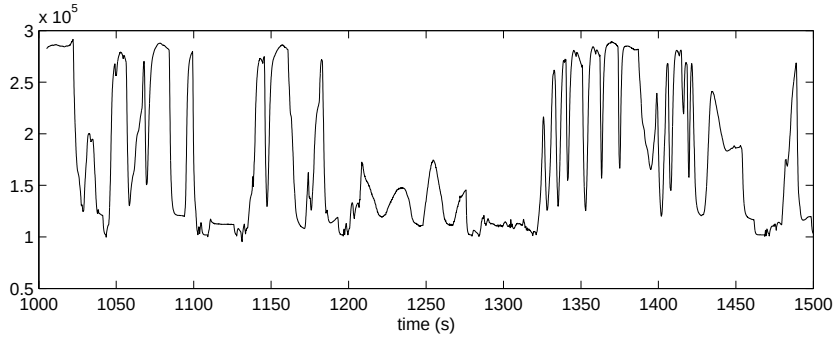

Figure 1.4: eadc-p-boostrawfilt-f32: boost pressure measurement with a stronger filter (Pa)

to do with the flowrate in the EGR but is really the measurement of W_{in}. This confusing EGR label stands because this sensor is used for diagnosing EGR faults.

Figure 1.3 and Figure 1.4 show the boost pressure measurement, for two different filtering. The data have been multiplied by 100, in order to have the measurement in Pascal.

On Figure 1.5, the measurement of the boost temperature is shown. 273 is added to the original data in order to have the temperature in Kelvin. Finally, Figure 1.6 shows the engine speed (N_{eng}) in RPM.

In the further sections, all the previous measurements will be used to generate residuals for fault detection. In order to be sure that no available information is lost, the less filtered signals will be used. This means that the data from eadc-dm-egr-s16 will be used for the measurement of air mass flow and eadc-p-boostfilt-s16 will be used as the measurement of the boost pressure.

1.1.2 Data derivatives

As only a static model is available to deal with dynamic data, some sort of filtering could be useful in order to minimize the effect of data coming from transitory regime on the residuals. To achieve this, time derivative of the original

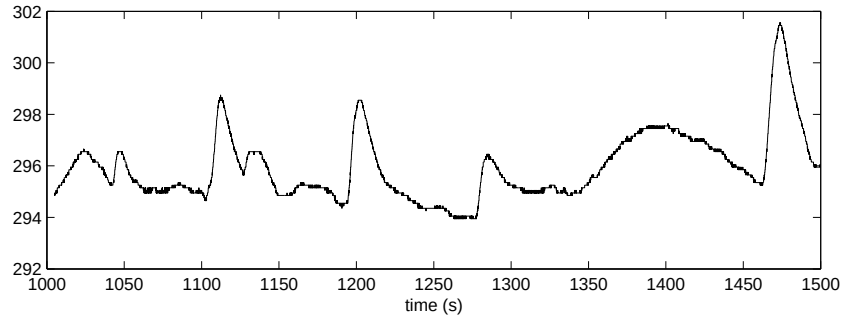

Figure 1.5: eadc-t-boostfilt-s16: boost temperature (Kelvin)

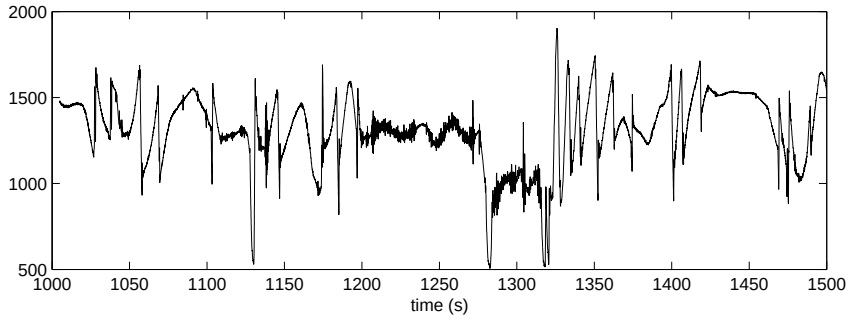

Figure 1.6: eess-n-avg10ms-s16: engine speed (RPM)

data will be computed to pinpoint the transitory regimes. This operation is done by priorly filtering the data through a low pass filter combined with a finite time response filter:

$$\frac{0.002(1+z^{-1}+z^{-2}+z^{-3}+z^{-4})}{(1-0.9z^{-1})^2} \tag{1.1}$$

Then, the differentiation was done by applying the *diff* Matlab function to the filtered data. Figure 1.7 shows the time derivative of the air-mass flow and Figure 1.8 shows the boost pressure time derivative.

Figure 1.9 shows the boost temperature time derivative and Figure 1.10 shows the engine speed time derivative.

All the previous figures are example of time derivatives measurements. They are shown in order to visualize how those variables can be used to localize transitory regimes.

1.2 Principal component analysis

In order to have a better understanding of the process, it is interesting to perform a principal component analysis on the data [WEG87]. This procedure is helpful to see what are the interactions between the different process variables and can also be used to evaluate a non measurable characteristic of the process. As we will see in the next chapter, the scores of the principal components can be used to either weight or threshold data, before performing fault detection.

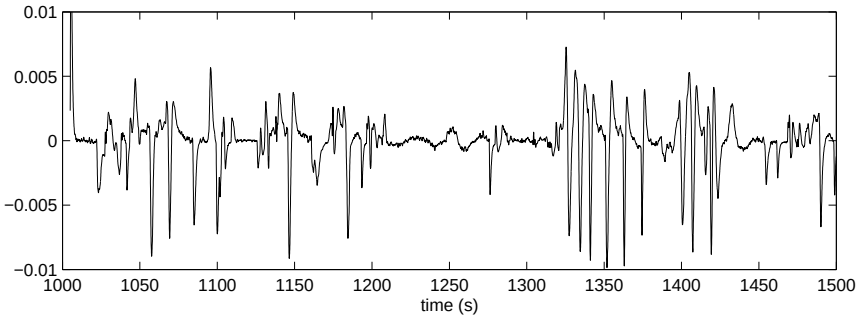

Figure 1.7: Air mass flow derivative

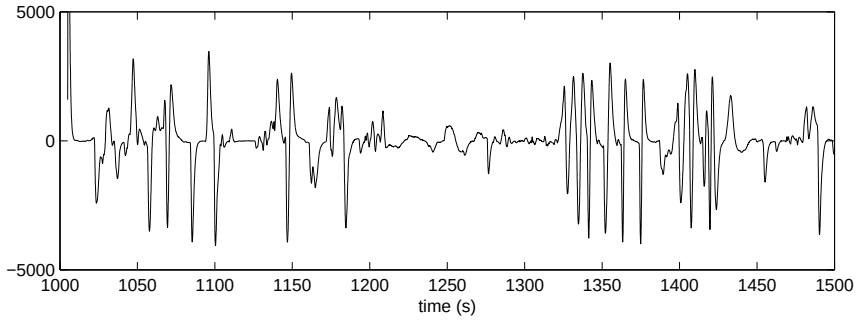

Figure 1.8: Boost pressure time derivative

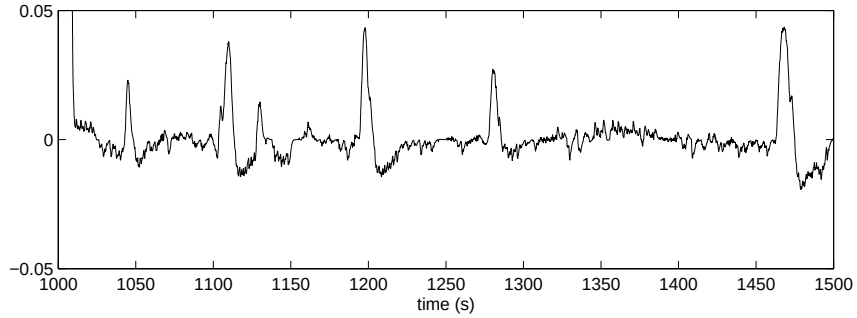

Figure 1.9: Boost temperature time derivative

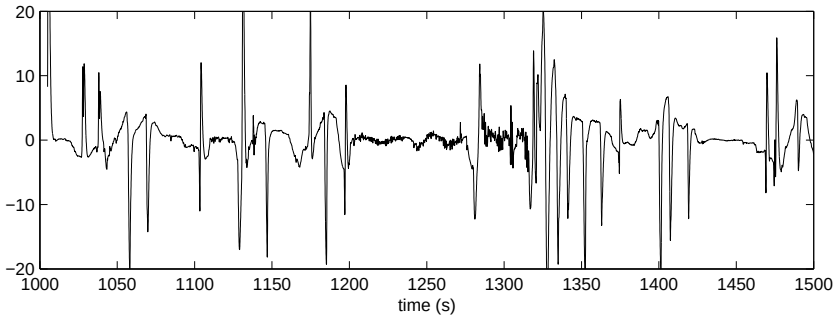

Figure 1.10: Engine speed time derivative

Figure 1.11 shows the first four principal components when PCA has been performed on the 4 raw variables and on their four derivatives. In the vertical

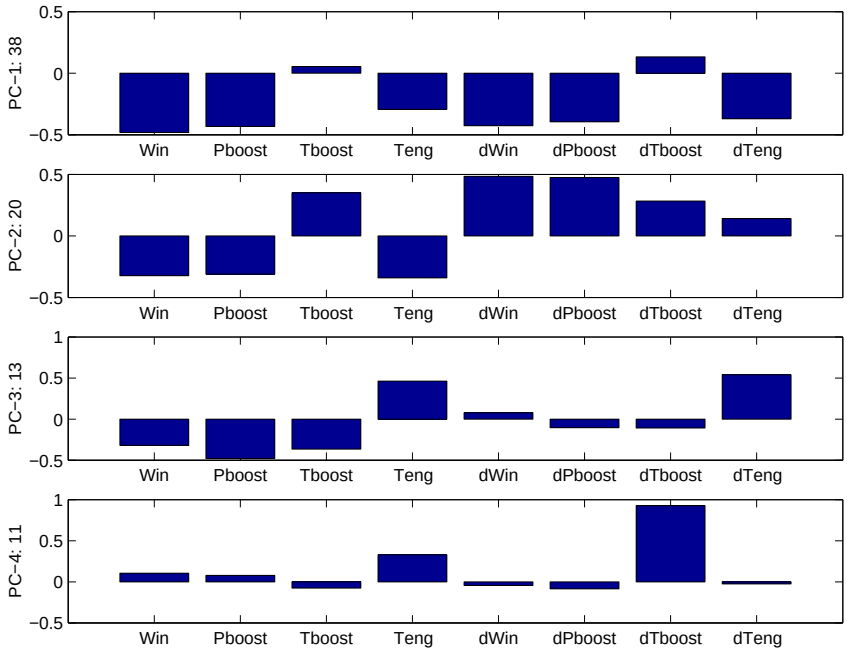

Figure 1.11: First four principal components of the data set *maverik*1, with explained variance (%)

axes, the explained variance is given. Thus, the first four principal components (out of 8) explained about 82 % of the total variance. This is actually quite difficult to assess the signification of the resulting principal components. But it will be possible to express the residuals with respect to the principal components scores, which could be an interesting tool to evaluate residuals statistical properties.

Figure 1.12 shows the four principal components of the PCA performed only to the derivatives of the four variables.

As the scores of the principal components are to be used to evaluate accuracy

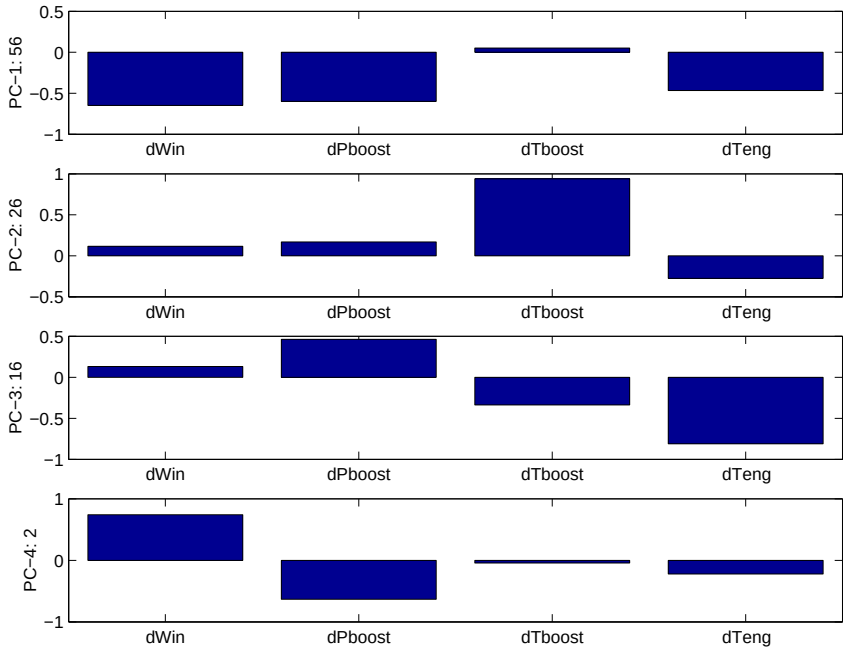

Figure 1.12: Four principal components of the data set *maverik*1, with explained variance (%)

of the data, it could be better to use only data derivative and not raw data. For fault detection purpose, using data derivative is not a problem but using the raw data can induce the exclusion of fault at a pre-treatment operation, which would definitely deteriorate the detection performances.

1.3 Residuals visualization

If we intend to improve the residual informative content, it is necessary to study what variables have a strong influence on the residual values. Therefore plotting the raw residuals against different variables of interest could give an hint about the most influent variables on the residual behavior. In the following, plots of residuals against some of those variables of operation are presented in order to show this kind of interaction. Ultimately, this study will be used in order to pinpoint variables that can be used, either as a threshold variable in case of data selection, either as weighting factor.

1.3.1 Time evolution of the residuals

In a first step, it is worth to look at the time evolution of the model residuals, in order to have an idea of the dynamics of the process. Figure 1.13 shows the time evolution of the residuals, as well as the evolution of the air-mass flow.

From a first view, it is already possible to associate high values of residuals with transitory regimes. It is then going to be an objective to find the best

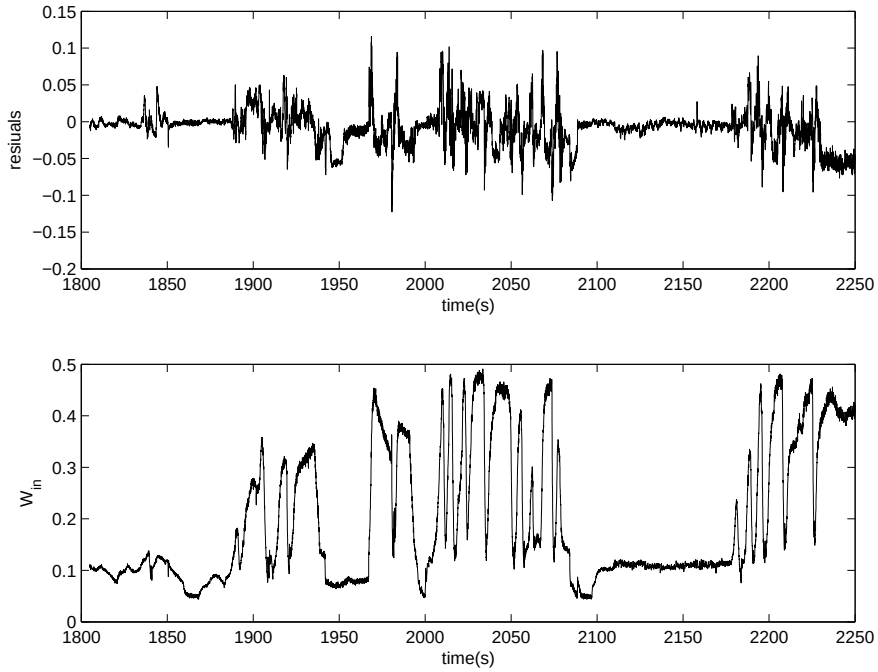

Figure 1.13: Temporal evolution of the residuals compared to the evolution of the air-mass flow.

way to improve the fault detection performances dealing with those transitory regimes.

1.3.2 Residuals plot against measured variables

It is already known that the accuracy of both the model and the air-mass flow sensor varies with the air-mass flow level. It is therefore interesting to plot the residuals with respect with W_{in} (shown in Figure 1.14). For all the following figures, data set *maverik*1 has been used and 15 000 points are plotted.

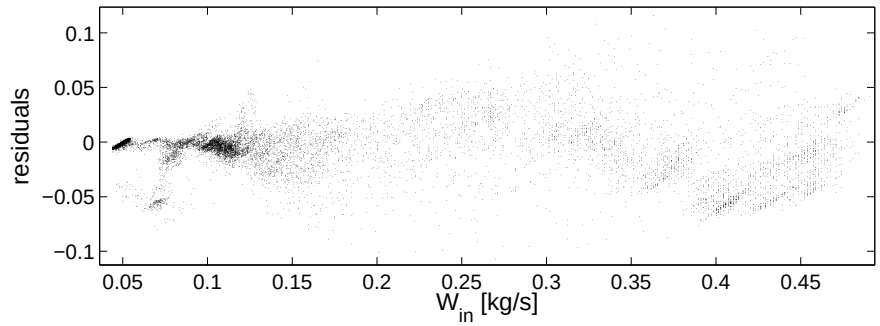

Figure 1.14: Residuals with respect to the air-mass flow level (W_{in})

Figure 1.14 shows that the residuals behavior seems to have different prop-

erties depending on W_{in}. At first, we can notice that for very low air-mass flow, there seems to be two different modes, generating two different types of residuals. It could look like a kind of hysteresis in the sensor, inducing a bias at some moment. Then, around $W_{in} = 0.1$ kg/s, the residual distribution seems to have a smaller dispersion. But for higher flows, the dispersion increases. Overall, the raw residuals have a mean of -0.0144 kg/s and a standard deviation of 0.0244 kg/s for this data set.

A similar bias on the mathematical expectancy of the residual was observed by Haraldsson in his Master Thesis. In a recurrent manner, he was multiplying the predicted air-mass flow by a factor 0.94, in order for the mean of the residuals to be zero.

In the present approach, many variables will be assessed in order to find the best way to generate reliable residuals. Ideally, it would be aimed to reduce significantly the variance of the remaining data, while approaching a zero mean for the residuals, and that for any dataset.

1.3.3 Residuals plot against variable derivatives

The principal assumption made when trying to discriminate more coherent data on a given observation window is that the model is accurate only when it is in a steady-state. Thus, it is interesting to plot the residuals against the derivatives of the four involved variables.

First, Figure 1.15 shows the plot of the residuals against the time derivative of the air flow. In Figure 1.15, we can notice that the dispersion of the data is

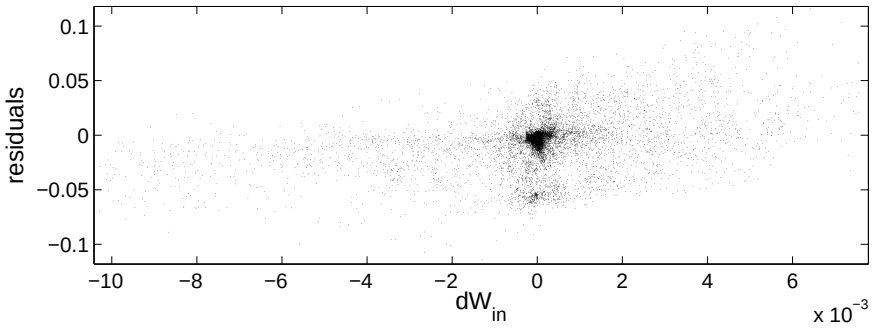

Figure 1.15: Residuals with respect to the air-mass flow derivative (dW_{in})

a little bit reduced when the derivative is close to zero. Also, it seems that for high positive derivative values, the residuals tend to be positively biased. This means that putting higher weight on low flow derivative absolute values of the data could probably improve the residuals.

Figure 1.16 shows the plot of the residuals against the the boost pressure time derivative. Just by a quick view of Figure 1.16, one can see that the time derivative of the boost pressure would not help as the dispersion of the residuals seems to be even at a maximum when the boost pressure derivative is close to zero.

Figure 1.17 shows the plot of the residuals against the time derivative of the boost temperature. As for boost pressure derivative, designing a threshold

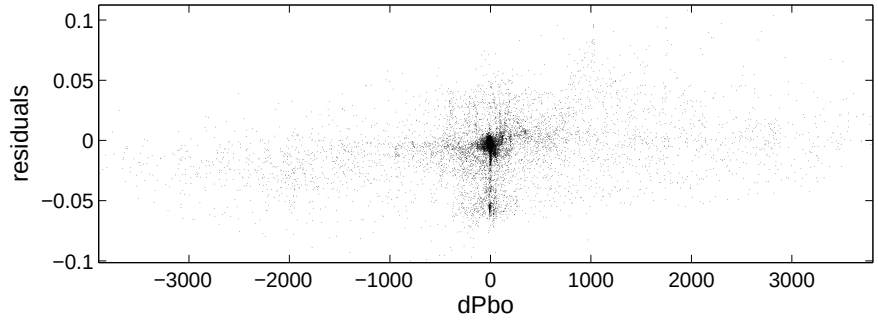

Figure 1.16: Residuals with respect to the boost pressure derivative (dP_{bo})

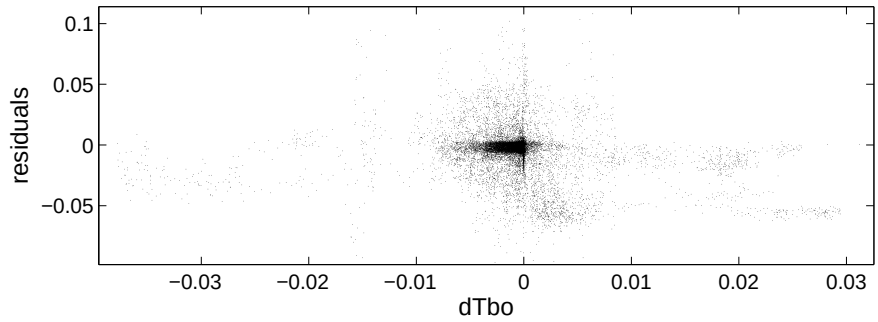

Figure 1.17: Residuals with respect to the boost temperature derivative (dT_{bo})

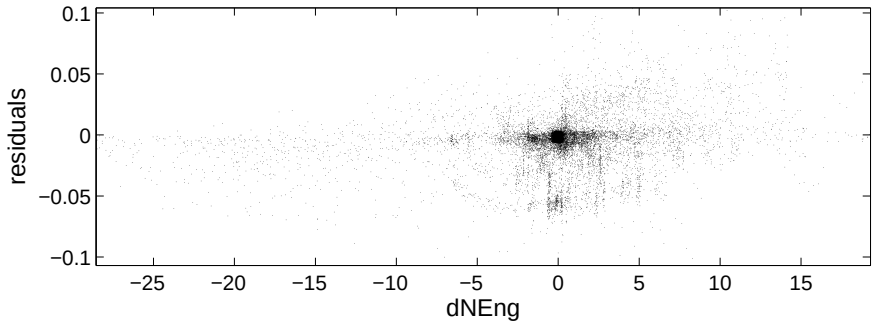

Figure 1.18: Residuals with respect to the engine speed derivative (dN_{eng})

on boost temperature derivative wont be very useful as the dispersion of the residuals is high when the time derivative is around zero. However, an interesting thing with this plot is that at high temperature derivative, there seems to have two residuals distribution, like observed in Figure 1.14 for low air flows. This could imply a correlation between those two behavior.

Figure 1.18 shows the plot of the residuals against the time derivative of the engine speed. As for the last two variable derivatives, the dispersion of the residuals around zero seems rather high to choose engine speed as a residual design variable.

A priori, the time derivative of the air-mass flow seems the best candidate to design a weighting or a threshold to discriminate model-coherent data.

1.3.4 Residuals plot against principal component scores

Plots of residuals can be made with respect to the four different measurement, as well as the four measurement derivative. But it is also possible to plot the residuals against a linear combination of the different available variables. The natural linear combination are the principal component of the different variables. Every new observation must be projected in the principal component directions. The result of this scalar product is the score of an given observation.

Figure 1.19 shows the plot of the residuals against scores of the first principal component given by the principal component analysis made over both raw and differentiated variables (see Figure 1.11). Around null scores of the first principal component, the dispersion of the residuals seems significantly reduced. This mean that designing a selection criteria based on the scores of the first principal component could be an interesting avenue.

Figure 1.20 shows the plot of the residuals against the scores of the second principal component (8 variables). On Figure 1.20, it is not that obvious that the scores of the second principal components are good candidate for residual discrimination. However, the analysis made just looking at the figures is very limited. For the third and fourth principal component, not much improvement was observed.

The principal component analysis made over only derivated variables can be assessed as well. Figure 1.21 shows the plot of the residuals against scores of the first principal component given by the principal component analysis made over only derivated variables (see Figure 1.12). The dispersion of the residuals

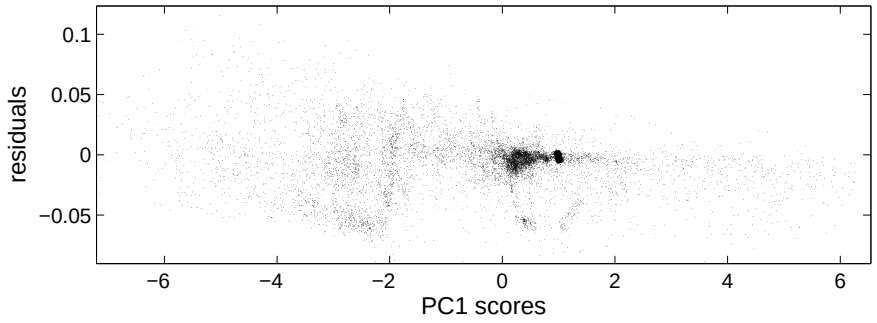

Figure 1.19: Residuals with respect to the scores of the first principal component (8 variables)

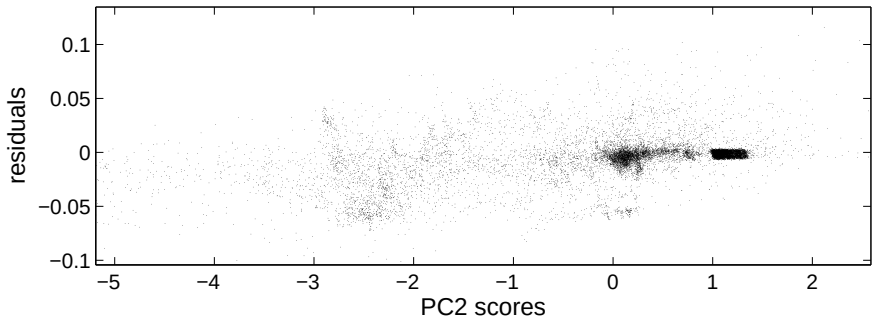

Figure 1.20: Residuals with respect to the scores of the second principal component (8 variables)

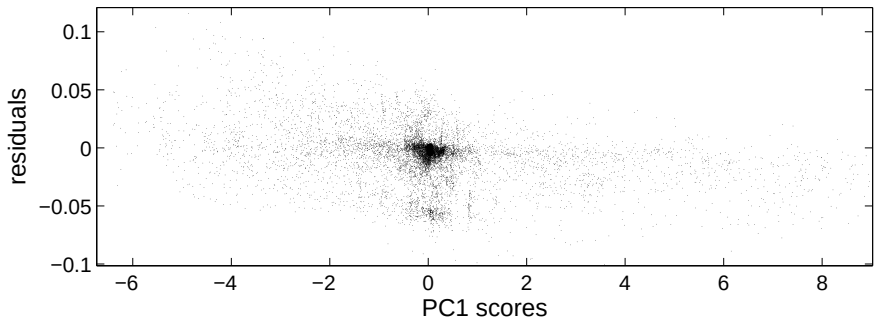

Figure 1.21: Residuals with respect to the scores of the first principal component (4 variables)

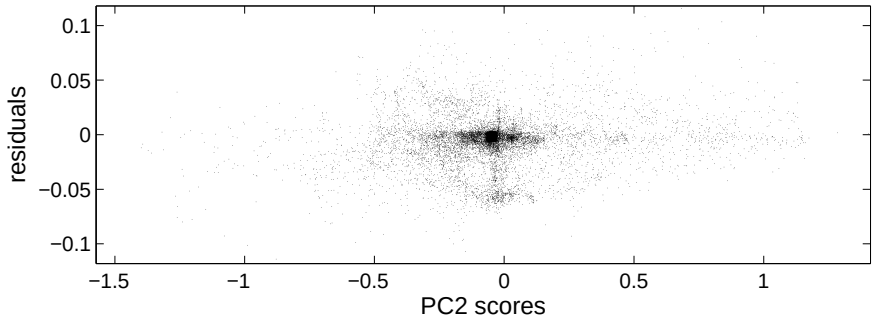

Figure 1.22: Residuals of the static model with respect to the scores of the second principal component (4 variables)

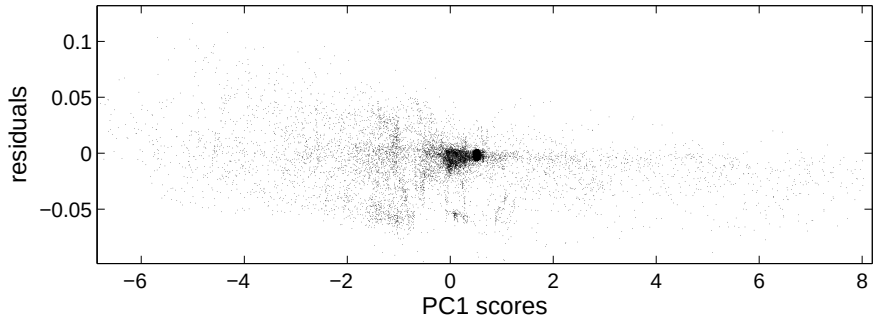

Figure 1.23: Residuals of the static model with respect to the scores of the first principal component (7 variables)

seems to be ligthly improved around zero, as high residuals values tend to be more associated with high negative score values.

Figure 1.22 shows the plot of the residuals against scores of the second principal component given by the principal component analysis made over only derivated variables. The dispersion of the residuals doesn't seem to be improved around zero, so this score will probably not be associated with data selection. The residuals projected against the third and fourth principal component are not giving interesting results either.

The selection of the first principal component in Figure 1.19 as threshold variable for selecting model-coherent data could seem straight forward. However, as this principal component is also made over the raw measurements, it could be influenced by the eventual presence of a sensor fault. For this reason, another principal component analysis has been done on all derivated and raw variables, except the air-mass flow measurement. This has been achieved in order find a variable that can reduce the dispersion of the residuals as the first principal componentscores do in Figure 1.19, but without using the measurement on which we intend to perform fault detection (W_{in}). Figure 1.23 shows the plot of the residuals against scores of the first principal component given by the principal component analysis made over all derivated variables, along with P_{boost}, T_{boost} and N_{eng}.

The dispersion around zero is significantly reduced, but maybe not as much

as in Figure 1.19. But probably, this variable will be less sensitive to eventual faults on W_{in}.

In the next chapter, the prospective study made here will be helpful in finding way to improve residual dispersion. The variables assessed here will be used in order to select or weight part of the data in the whole residual generation procedure.

Chapter 2

Residual Generation

If the process model we have described before was exact, the fault detection procedure would be quite simple. As a lot of measurements are available and supposing that the measurements are submitted to white noise, the measurement mean would have a much reduced variance (by the huge amount of data) and very low amplitude faults could then be detected. But as shown in the previous chapter, the process model is not exact and shows evident variability, even for very large samples.

In the first section, the behavior of the original model for residual generation will be assessed. Then, an attempt to improve the residuals selecting part of the data will be presented in the second section. The third section presents a way to generate residuals using statistical charts. This last technique is showing interesting results as the overall residual dispersion is significantly reduced. The two approaches will be compared in the fourth section. Finally, the corresponding expected inputs for fault detection will be evaluated.

In order to compare the different approaches for residual generation, it is necessary to have rigorous comparison tools. To do so, the different datasets will be always used in the same way. First, the first half of the dataset called *maverik*4, which includes 210 000 different observation, will be designated as the calibration set. All the other available datasets will be used for validation purposes. From the truck *Maverik*, the datasets *maverik*1, *maverik*2, *maverik*3 and *maverik*4*b* (second half) will be used. Also, coming from the similar truck *Platina*, the datasets *mplatina*1, *mplatina*3 and *mplatina*4 will be used for validation. The ultimate goal would be to have a residual generation procedure that when calibrated on a given dataset, still show good results on the different validation sets.

2.1 Residuals of the original model

It could be expected that the residual is simply the difference between the measured and the predicted air-mass flow, as in Equation 2. But due to innacuracy in both the model and the measurements, an important bias remain and must be corrected, even when testing the calibration dataset. Depending on the

calibration set used to evaluate the volumetric efficiency map and the bias, this bias will differ. Taking many different samples of 15 000 points from the calibration set, the bias has been evaluated and results in the following expression for the residual calculation:

$$r = W_{in} - W_p + 0.0129 \tag{2.1}$$

The corresponding units of the residuals is in kg/s, compared to a sensor operating range of 0.05 kg/s to 0.50 kg/s.

2.1.1 Residual distribution: calibration set

From this basic residual generation, it is interesting to plot an histogram in order to see what is the approximate distribution of the residuals. Figure 2.1 shows the distribution of the residual mean values over around 1000 different 15 000 points overlapping samples of the calibration set. From this only dis-

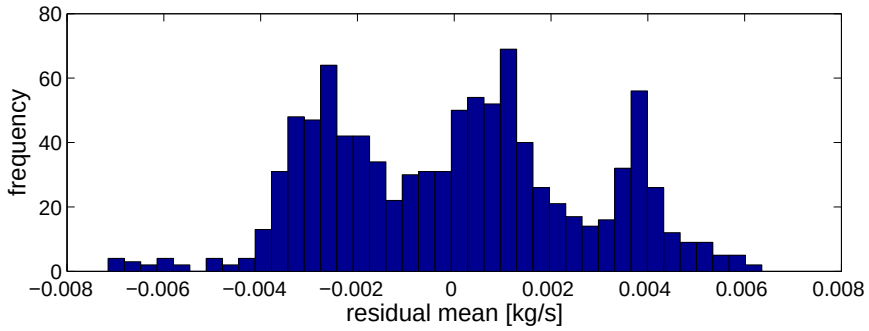

Figure 2.1: Averaged residual distribution for the calibration set.

tribution, it could seem that faults in the range of 0.01 kg/s could be easily detected.

2.1.2 Residual distribution: validation set

When using different dataset, the resulting residuals distributions shows huge differences. Figure 2.2 shows the residual mean distributions for the calibration set as well as for the seven different validation sets. Figure 2.2 illustrates the whole problem of fault detection in a real engine. The model calibrated with one data set is not reliable when applied to other datasets. The distribution biases observed in Figure 2.2 are significant as they appear to be higher than the dispersion of the original distribution.

For the dataset *maverik4b*, the behavior of the residuals is dramatically different. Some special event may have occurred for this particular dataset and the author have no idea if the whole dataset is reliable or not. The data are shown here, but maybe they wont be considered in the choice of the best residual generation method.

From the distributions observed in Figure 2.2, it would be possible to detect fault only on the range of 0.04 kg/s, at the best. From that point, the work will consist in finding a way to make the overall residual generation more reliable.

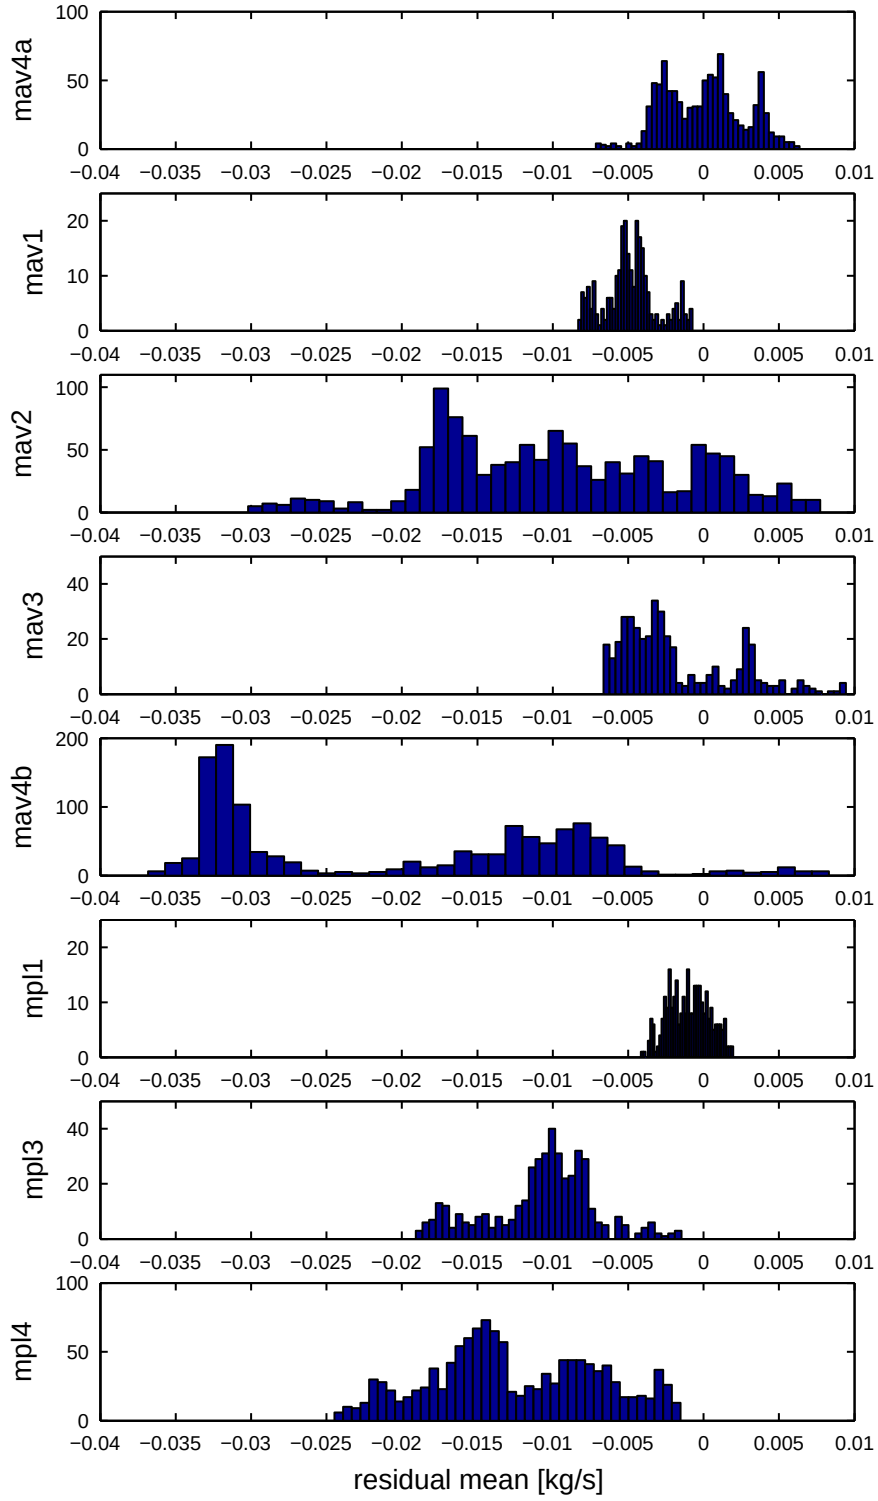

Figure 2.2: Averaged residual distributions for the calibration set and 7 validation sets.

threshold variable	$\bar{r}$	σ_r	best candidates
no threshold	-0.0144	0.0247	
dW_{in}	-0.0183	0.0194	$\surd$
dP_{boost}	-0.0177	0.0223	
dT_{boost}	-0.0125	0.0231	
dN_{eng}	-0.016	0.0244	
PC1 (8 var.)	-0.0089	0.0152	$\surd$
PC2 (8 var.)	-0.014	0.0157	$\surd$
PC3 (8 var.)	-0.0155	0.0277	
PC4 (8 var.)	-0.0116	0.0251	
PC1 (4 var.)	-0.0183	0.0207	
PC2 (4 var.)	-0.0169	0.0227	
PC3 (4 var.)	-0.0179	0.0215	
PC4 (4 var.)	-0.0152	0.0227	
PC1 (7 var.)	-0.0112	0.0159	$\surd$
PC2 (7 var.)	-0.0097	0.0193	
PC3 (7 var.)	-0.0162	0.0233	
PC4 (7 var.)	-0.0170	0.0230	

Table 2.1: Statistical properties of the remaining residuals after a selection has been made, based on a threshold on the given operational variable.

2.2 Residuals generation with data selection

At first, it was supposed that high residual values were associated with transitory regimes, as the process model could only be accurate with static data. Then, the first attempt to improve the residuals was to eliminate part of the data where those transitory behavior occurs. This is done by setting a threshold on a given operational variable to be determined.

2.2.1 Best thresholding candidates

All variables, filtered variables and linear combination of variables are candidates to become a threshold variable. But it is not possible to study all of them extensively. Then it is necessary to operate a pre-selection of the best thresholding candidates.

A way to evaluate which variable should be chosen to select model coherent data is to evaluate the statistical properties of the remaining residuals once a selection as been performed with the corresponding threshold variable. Table 2.1 shows the mean and the standard deviation of the remaining residuals, for a merge of many datasets. For each data selection, a threshold has been selected in order to keep approximately 20 % of the initial data. The first line shows the residuals properties prior to any selection. The fourth column indicates which variables seem to be the most appropriate for thresholding.Those are the properties of all residuals and not residual sample means.

The best performance when using only variable derivative as threshold is obtained by the air flow derivative, for which the standard deviation is reduced by about 20 %. When using the score of the first and second principal component

data	no thres.	dW_{in}	PC1 (8)	PC2 (8)	PC1 (7)
mav4a	0.0000	-0.0009	0.0078	0.0027	0.0001
mav1	-0.0048	-0.0053	0.0036	-0.0005	-0.0050
mav2	-0.0095	-0.0102	-0.0004	-0.0071	-0.0162
mav3	-0.0018	-0.0042	0.0020	-0.0009	-0.0047
mav4b	-0.0204	-0.0257	-0.0048	-0.0149	-0.0349
mpl1	-0.0010	-0.0078	-0.0013	-0.0031	0.0002
mpl3	-0.0107	-0.0151	-0.0108	-0.0105	-0.0053
mpl4	-0.0126	-0.0124	0.0001	-0.0071	-0.0061
tot.	0.0609	0.0818	0.0307	0.0469	0.0724

Table 2.2: Mean of averaged residual distributions, for different datasets and threshold variables.

(8 variables) for threshold, the standard deviation of the remaining residuals is reduced by almost 40 %. Finally, when using the score of the first principal component in the case when 7 variables are considered, the reduction of the standard deviation is around 35 %. Those 4 possible threshold variables will be considered in order to compare the residuals distribution between different datasets.

2.2.2 Threshold variable selection

The study made in Table 2.1 intend mostly to find which threshold variable could bring the most improvement in the dispersion on a big sample. Now, the datasets that have been previously used as calibration and validation sets will be used again to assessed if the selection of parts of the data using thresholds on specific variables can reduce the overall range of the residual distributions for different datasets.

For each data selection, the corresponding threshold has been set in order to eliminate around half of the data. As a lot of data is available to perform the diagnosis on each cycle, discarding half of it wouldn't have dramatic consequences on detection performances. In Table 2.2, the mean of each dataset residuals distribution is shown for the four different threshold variables. In order to compare values, the original residuals without data selection are also shown. The last line on Table 2.2 shows the added absolute biases of every dataset residual distribution for each threshold variable. As it could have been expected looking at Figure 1.19, the best results are obtained with a threshold on the scores of the first principal component (8 variables). For the corresponding remaining residuals samples, the total bias over the 7 validation samples is reduced by half, compared to the original residuals.

2.2.3 Selected residual distributions

Figure 2.3 shows the residual mean distributions for the 8 datasets when the threshold variable is the scores of the first principal component. The threshold was put to -1 and +1, which discard about half of the data. Comparing Figure 2.3 to Figure 2.2 which has the same scale, it is obvious that a significant improvement in the residual distributions has been achieved, as the various distributions are less spread.

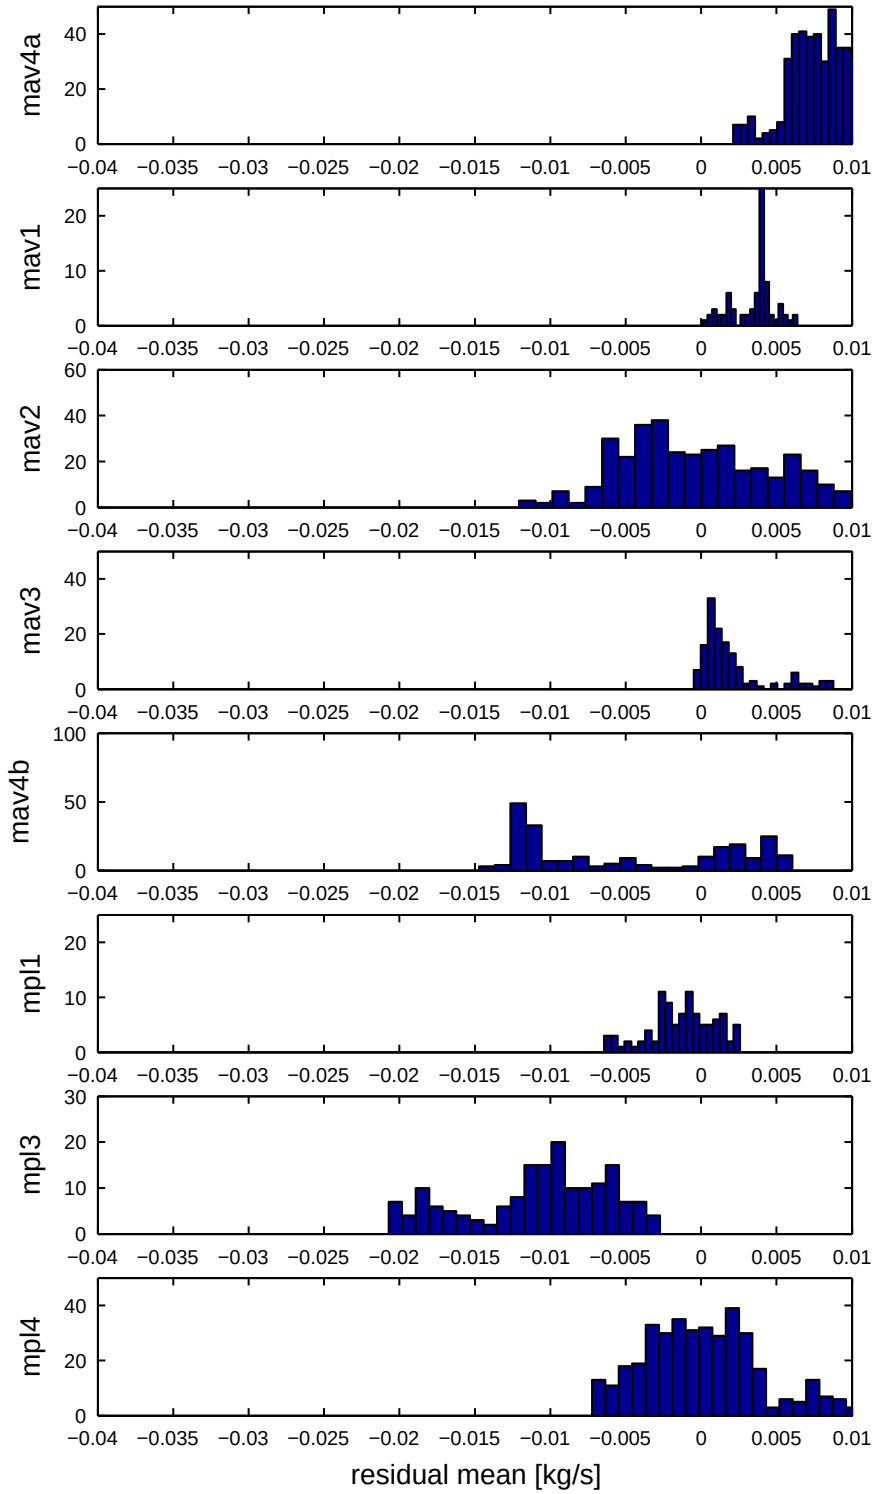

Figure 2.3: Averaged residual distributions for the calibration set and 7 validation sets, when a threshold has been put on the scores on the first principal component.

2.3 Residuals generation using statistical charts

As statistical properties of the residuals seem to be dependent on the operational level of the truck engine, it is natural to try to use this additional information in order to improve the residual generation.

2.3.1 Improved residual generation

The main idea of the new residual generation is to use adaptive statistical properties for the generation of each residual among the 15 000 residuals of a cycle sample. Therefore, depending on an operational variable to be determined, different bias (mean) and standard deviation will be assigned to each residual.

Instead of using Equation 2.1, the residuals will then be generated by Equation 2.2:

$$r = \bar{\sigma}\frac{W_{in} - W_p + b(x_i)}{\sigma(x_i)} \tag{2.2}$$

In Equation 2.2, b is a bias taken from a chart that depends on the observation i of the operational variable x. In a same way, σ is taken form a chart also depending on x_i. In order to bring back the original dimension of the residual, the quotient is multiplied by $\bar{\sigma}$, the average standard deviation of the residual over the sample. In that way, each residual of the sample will have its own bias and its own level of confidence, depending on its corresponding operational variable. Then, the test quantity for fault detection will again be the mean of all the 15 000 residuals of the sample.

2.3.2 Chart selection

As for the design of threshold variables, charts can be built on any operational variable, filtered variable, linear combination of variables, etc. It is then necessary to determine what chart would bring the most improvement.

The main problem observed in Figure 2.2 is the evident difference in the distribution mean of the different datasets. As for finding the best threshold variable, the performance index will be based on the difference between the calibration set distribution mean and the validation sets distribution means. Many different charts have been studied in order to find the one that minimize this difference. Table 2.3 shows, for the different data sets, the bias of the distribution mean, for some of the charts studied. The charts were done on the operational variables W_{in}, air-mass flow derivative (dW_{in}), scores of the first principal component (8 variables), scores of the second principal component (8 variables) and scores of the first principal component (7 variables). The last line of Table 2.3 shows the sum of the absolute biases for each chart.

The chart on W_{in} shows really good results on the first two validation sets, which were taken from the same truck as the calibration set. But when it comes to validation sets taken from the second truck engine, the results are no more reliable. In Table 2.3, the sum of the distribution absolute means show that the biases are generally minimized using a chart on the scores of the second principal component (PC2). Figure 2.4 illustrates the corresponding charts for residuals mean and standard deviation. This result is not an absolute one as

data	no chart	W_{in}	dW_{in}	PC1 (8)	PC2 (8)	PC1 (7)
mav4a	0.0000	0.0002	0.0002	-0.0007	0.0004	0.0002
mav1	-0.0048	-0.0019	-0.0049	-0.0041	-0.0016	-0.0038
mav2	-0.0095	-0.0041	-0.0098	-0.0065	-0.0058	-0.0090
mav3	-0.0018	0.0023	-0.0019	-0.0015	0.0018	-0.0007
mav4b	-0.0205	-0.0111	-0.0209	-0.0166	-0.0131	-0.0212
mpl1	-0.0010	-0.0038	-0.0018	-0.0041	0.0011	-0.0074
mpl3	-0.0107	-0.0139	-0.0116	-0.0115	-0.0075	-0.0134
mpl4	-0.0126	-0.0179	-0.0136	-0.0109	-0.0089	-0.0182
tot.	0.0609	0.0552	0.0647	0.0559	0.0402	0.0739

Table 2.3: Mean of averaged residual distributions, for different data sets and threshold variables.

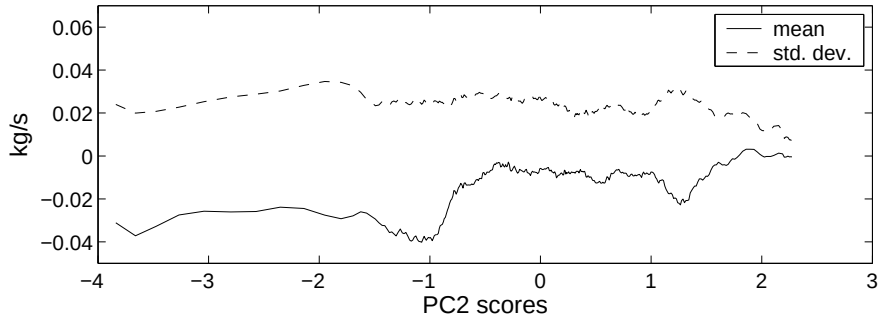

Figure 2.4: Residual mean and standard deviation of the calibration set against the scores of the 2nd principal component.

not all the different chart alternatives have been studied. But among all the ones studied, this chart is the best candidate to improve residuals reliability.

2.3.3 Dispersion reduction of the residuals

From the results of Table 2.3, it is already known that the use of the chart on PC2 will reduce the expanse of the different dataset residual distributions. Figure 2.5 shows exactly the same distributions as in Figure 2.2, but when the charts has been used to generate the residuals. Showing the same scale as Figure 2.2 and Figure 2.3 , Figure 2.5 demonstrates that the new residual distributions are less spread than the original ones. However, it is not yet clear if it is better to use a threshold or to use a chart on a given operational variable.

2.4 Comparison of the residual generation methods

The three different approaches for residual generations will be compared here. First, the raw residuals without any prior treatment, then the remaining residuals after a threshold has been set on the scores of the first principal component, and finally, the residuals generated using statistical charts on the scores of the second principal component.

As one of the dataset (*maverik4b*) shows totally different properties than the other, residuals distributions will also be assessed without this particular dataset. The engine and sensors during this test had a suspicious behavior, that could be interpreted as a fault. But as nothing is known about the particular experimental conditions of this test, nothing can be said and therefore, the two possibilities will be assessed: if this behavior is normal or if it is abnormal.

2.4.1 Raw residuals

Merging all residual distributions of Figure 2.2 gives the residual mean distribution of Figure 2.6. Depending on which dataset had been chosen for the calibration of the bias, the distribution would be slightly different, having a different mean. But the general dispersion of the residuals would be similar. When the dataset *maverik4b* is ignored, the resulting distribution is nicer (see Figure 2.7).

Depending if the dataset *maverik4b* is considered or not, the overall range of the raw residuals is 0.048 kg/s or 0.039 kg/s. This is huge comparing with the operating range of the air-mass flow sensor of 0.05 to 0.5 kg/s, and will therefore highly limitate the diagnosis performances.

2.4.2 Residuals after data selection

Merging all residual distributions after the thresholding of Figure 2.3 gives the overall residual mean distribution of Figure 2.8. The dispersion reduction can easily be seen comparing Figures 2.6 and Figure 2.8. No residual mean go further than -0.02 kg/s compare to -0.035 kg/s with the raw residuals.

Figure 2.9 shows the residual mean distributions when the dataset *maverik4b* is ignored. The range of the selected residual means goes from -0.021 kg/s to

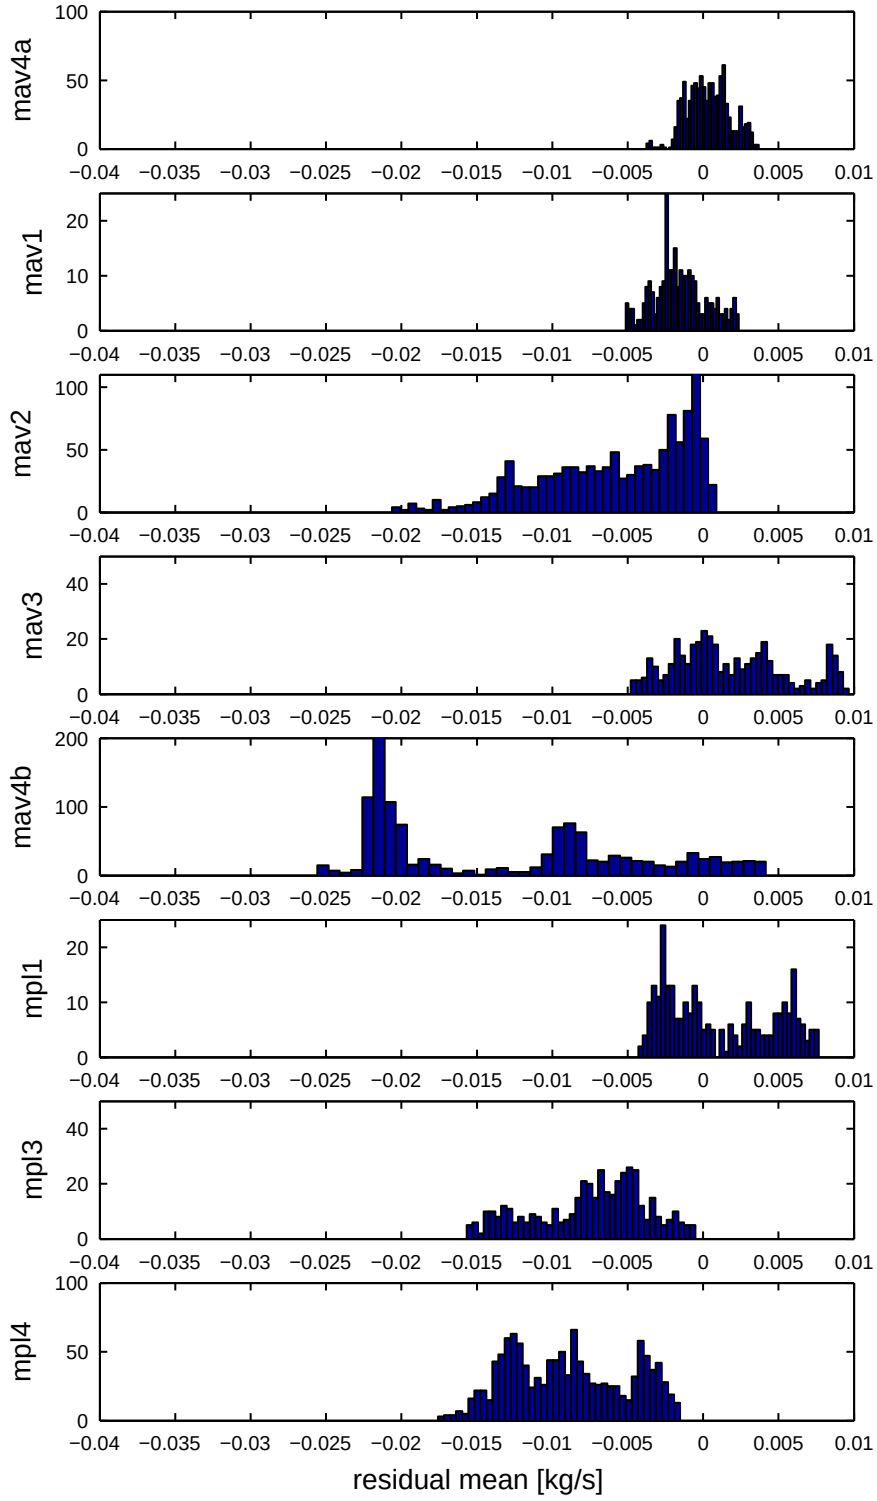

Figure 2.5: Averaged residual distributions for the calibration set and 7 validation sets, when a chart on PC2 scores has been used for residual generation.

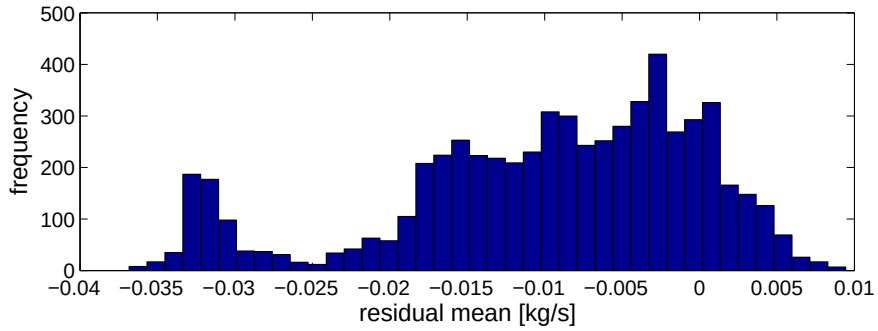

Figure 2.6: Averaged residual distribution for all datasets, using raw residuals.

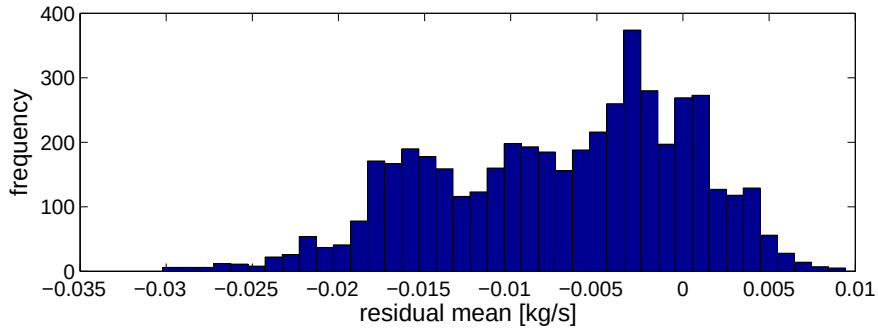

Figure 2.7: Averaged residual distribution for all datasets except *maverick4b*, using raw residuals.

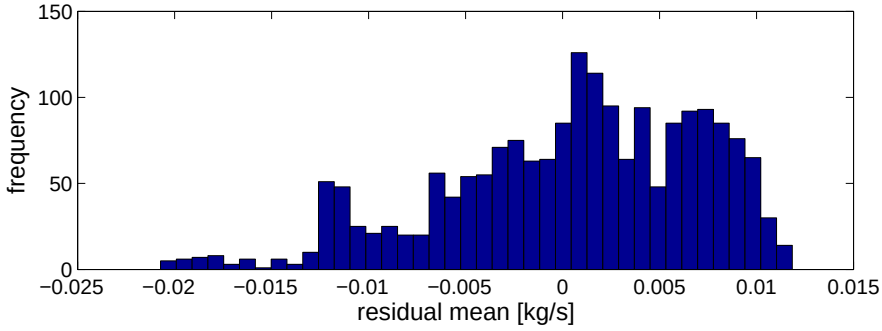

Figure 2.8: Averaged residual distribution for all datasets, using selected residuals.

0.012 kg/s, no matter if dataset *maverik4b* is considered or not. However, the distribution without *maverik4b* is slightly more compact than the first one.

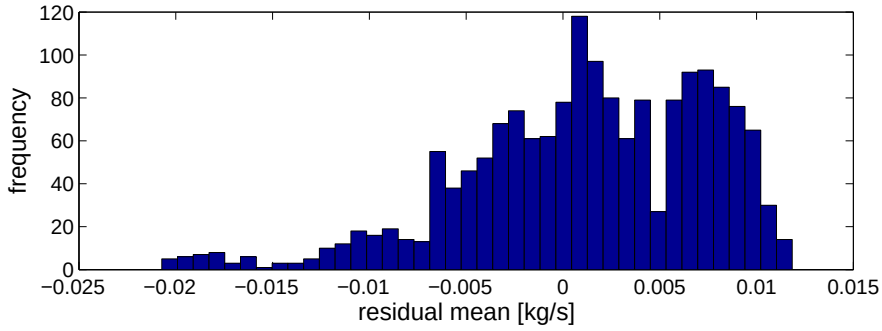

Figure 2.9: Averaged residual distribution for all datasets except *maverick4b*, using selected residuals.

2.4.3 Residuals after the use of a chart

Figure 2.10 shows the distribution of the residual mean from all dataset, when a chart on PC2 has been used in the residual generation. The suspicious behavior of dataset *maverik4b* is clearly present in Figure 2.10, compared to Figure 2.8 where it seemed to vanish. So depending if this behavior is considered abnormal or not, an approach would give better results than the other. The thresholding approach would most probably not detect this behavior as a fault as the chart approach will probably do.

Figure 2.11 shows the residual mean distribution of all datasets but *maverik4b*, when using charts to generate residuals. Depending if *maverik4b* is considered or not, the range of the residual mean is reduced from -0.025 kg/s to 0.009 kg/s or from -0.18 kg/s to 0.009 kg/s.

2.4.4 Best approach

In order to identify which approach would bring the best results for diagnosis

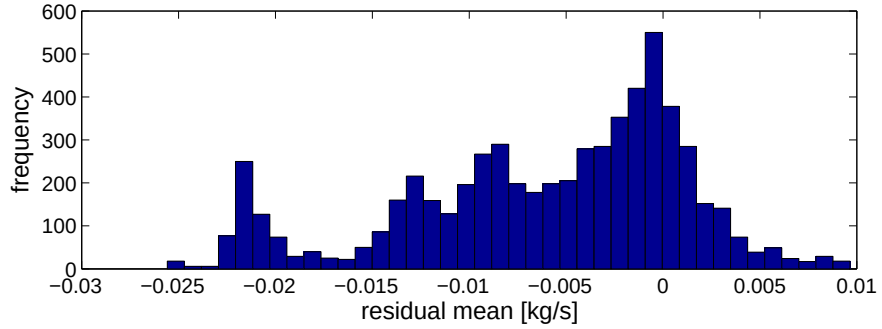

Figure 2.10: Averaged residual distribution for all datasets, using residuals generated with charts.

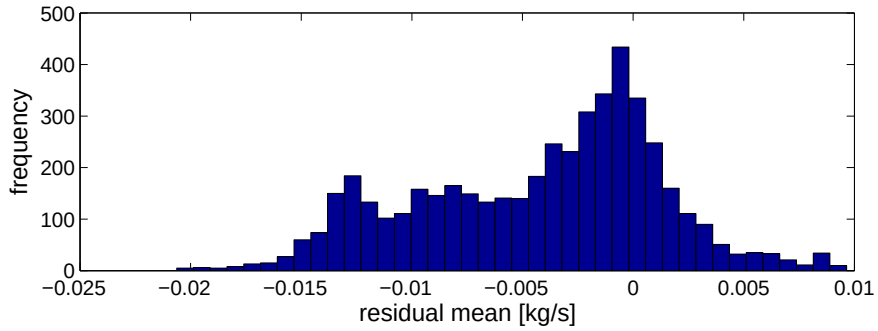

Figure 2.11: Averaged residual distribution for all datasets except *maverick4b*, using residuals generated with charts.

purposes, it is necessary to study the dispersion of the various residual mean distribution. But as the residual distributions of Figures 2.6 to 2.11 are not symmetrical, it is more appropriate to evaluate different dispersions for each sides of the distributions. Then, cutting each distribution at the most occurring residual mean (top of the distribution), a standard deviation is evaluated for the negative side and another for the positive side. That is repeated for each of the three approaches. Table 2.4 shows those dual standard deviations.

approach	σ_r-(all)	σ_r+(all)	σ_r- (no *mav4b*)	σ_r+ (no *mav4b*)
raw res.	0.0142	0.0042	0.0101	0.0041
selec. res.	0.0078	0.0067	0.0070	0.0059
chart	0.0104	0.0031	0.0076	0.0032

Table 2.4: Negative and positive standard deviations (kg/s) for the 3 residual generation approaches; on left: for all datasets; on right: for all datasets except *maverik4b* .

Adding both positive and negative standard deviation, the approach using a chart on PC2 to generate residuals is always better. When not considering *maverik4b*, the difference is more obvious. For that case, the dispersion is reduced by 24 % for the chart approach, comparing to a 10 % reduction for the threshold approach.

2.5 Diagnosis improvement

A dispersion reduction of the residuals will bring an improvement in the fault detection performances. Making some gross assumption, it is possible to evaluate this performance improvement. The next figures show the detection test power function [CB90] for a simulated sensor bias on the air-mass flow measurement. This power function has been roughly evaluated considering two different Gaussian distributions, on each sides of every residual distributions, and setting the false alarm rate to 1 %. Figure 2.12 shows the power functions of the three approaches when dataset *maverik4b* is consider as normal.

From Figure 2.12, it is not clear which shows the best performances as what's gained on the negative side seemed to be lost on the positive side. The results

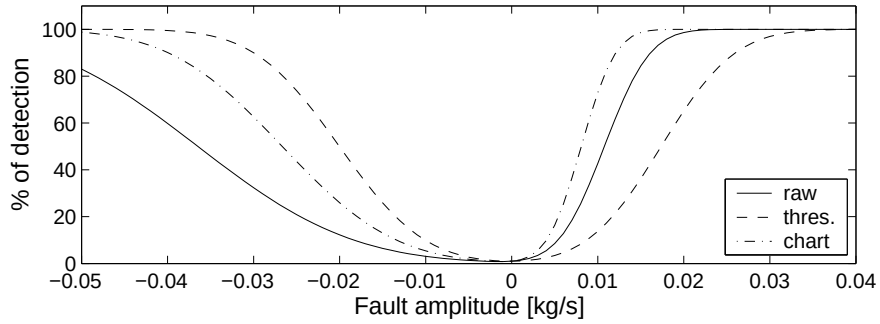

Figure 2.12: Power functions evaluation for a W_{in} sensor bias (considering *maverik4b* as normal).

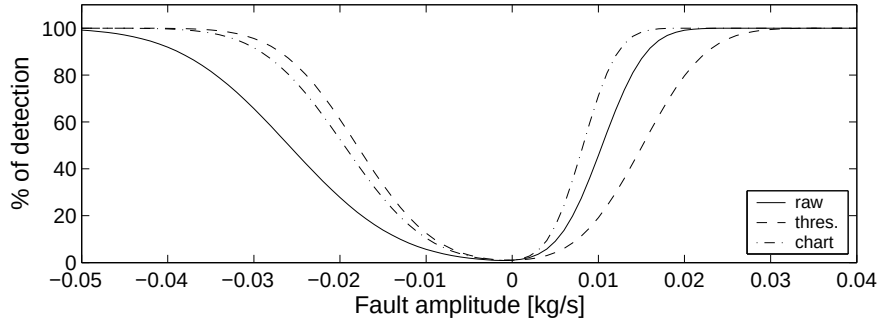

Figure 2.13: Power functions evaluation for a W_{in} sensor bias (considering *maverik4b* as abnormal).

are more obvious on Figure 2.13, which shows the same power functions, but when *maverik4b* is considered as abnormal.

From Figure 2.13, the power function of the approach using charts is definitely sharper. Thus, we can conclude that using charts will allow to detect smaller faults, for the same level of false alarm. Table 2.5 shows the approximate fault amplitudes where 98 % of detection is obtained in the power functions.

Generally, using a chart in the residual generation procedure is improving the detection performances by approximately 25 %.

approach	fault-(all)	fault+(all)	fault- (no *mav4b*)	fault+ (no *mav4b*)
raw res.	-0.065	0.019	- 0.046	0.019
selec. res.	- 0.036	0.031	-0.032	0.027
chart	- 0.048	0.015	-0.035	0.015

Table 2.5: Negative and positive fault amplitudes (kg/s) where 98 % of detection is reached; on left: for all datasets; on right: for all datasets except *maverik4b*.

2.6 Conclusion

With the incoming new regulations, the problem of diagnosis for truck engines need solutions in a near future. To match regulations, faults have to be detected as early as possible, without showing false alarm. From this viewpoint, it is necessary to improve model-based diagnosis, even having as starting points inaccurate model.

Having an inaccurate model of the truck engine air-path, this study first proposed a way to generate residuals of the model using selected parts of data, based on various operational variables. Then, another avenue is explored where the model for residual generation is empirically improved using adaptive statistical properties (mean and standard deviation) taken from empirical charts.

Those approaches, that could on a first view seem quite particular to the presented problem can be used for other mean. In fact, any problem where residuals are generated in large quantity is a candidate for that kind of model improvement.

In the present study, the best results were obtained by generating residuals making use of statistical charts based on the score of the second principal component. If the technique can be extended to other diagnosis problem, this results is particular to this problem. Doing so, the overall dispersion of the resulting residuals was significantly reduced compared to the original ones. This dispersion reduction in the residuals was improving the diagnosis performances by close to 25 %. Depending on the exact requirement of the environmental laws, this difference could be very helpful.

Bibliography

[CB90] George Casella and Roger L. Berger, *Statistical inference*, Duxbury Press, Belmont, CA, 1990.

[Har02] Petter Haraldsson, *Optimisation of a diagnosis test for a truck engine*, Master's thesis, ISY, Linköping Universitet, 2002.

[Hey88] John B. Heywood, *Internal combustion engine fundamentals*, McGraw-Hill International, 1988.

[Jur94] Ronald Jurgen, *Automotive electronics handbook*, McGraw-Hill, 1994.

[US93] A. Unger and K. Smith, *The OBDII system in the Volvo 850 turbo*, SAE Paper.(932665) (1993).

[WEG87] S. Wold, K. Esbensen, and P. Geladi, *Principal Component Analysis*, Chemometrics and Intelligent Laboratory Systems (1987), no. 2, 37–52.