

A Probability Density Metric for Automotive Three-way Catalyst Diagnostics

Jill S. Kirschman

Center for Nonlinear Dynamics and Control, Villanova University, Villanova, PA 19085

Jesse C. Frey

Department of Mathematical Sciences, Villanova University, Villanova, PA 19085

Kenneth R. Muske and James C. Peyton Jones

Center for Nonlinear Dynamics and Control, Villanova University, Villanova, PA 19085

Imad H. Makki and Michael J. Uhrich

Powertrain Research and Development, Ford Motor Company, Dearborn, MI 48126

Abstract

A statistical classification technique based on the fraction of time the catalyst gain is very close to zero is used as a diagnostic metric for on-board monitoring of an automotive catalyst. Preliminary results indicate that it is possible to perform very accurate discrimination between catalyst operation, even near the on-board diagnostic detection threshold, using this technique. Experimental vehicle tests with each of the different catalysts are used as the basis for comparison.

1. Introduction

Effective and robust on-board system monitoring is essential in order to ensure that the components of an automotive emission control system continue to operate properly with age. Because a small fraction of malfunctioning vehicles are believed to account for a large fraction of automotive emissions, on-board monitoring has the potential to significantly reduce these emissions. This observation, along with legislation that mandates on-board diagnostic (OBD) systems to monitor the health and performance of the catalyst system, has led to interest in monitoring strategies that are able to reliably indicate when an automotive three-way catalyst emission control system is no longer meeting specification or when a fault is present in the system.

Legislated targets aimed at improving air quality through the reduction of automotive emissions specify both the emission levels that must be achieved by new vehicles as well as the devi-

ation from these levels that must be detected by on-board diagnostic monitoring over the vehicle lifetime. These emission regulations are specified in terms of hydrocarbon and nitrogen oxide levels. However, there are no cost-effective and reliable automotive sensors for these components at present. This limitation has led to monitoring approaches that consider the oxygen storage capacity of the catalyst rather than tailpipe emissions. Although the catalyst oxygen storage capacity can not be measured directly, it can be inferred using exhaust gas oxygen (EGO) sensors and a suitable model of the catalyst. These sensors are already available in modern automotive emission control systems making this approach to catalyst monitoring attractive.

A model-based catalyst monitoring strategy, in which the catalyst oxygen storage capacity is modeled as a nonlinear integrator with a time-varying integral gain, is considered in this work. We present a diagnostic metric based on an estimate of the probability density of the catalyst gain around zero that is computed from a moving window of catalyst gain values.

2. Catalyst Gain Derivation

The monitoring strategy presented in this work is based on the use of the pre- and post-catalyst exhaust gas oxygen sensors to determine the catalyst gain. These sensors measure the excess or deficiency of oxygen in terms of the air fuel ratio

$$\lambda = \frac{1}{K_\lambda} \frac{\dot{m}_{O_2}}{\dot{m}_f} \quad (1)$$

in which $\lambda = 1$ represents stoichiometric air fuel ratio, \dot{m}_{O_2} is the oxygen mass flow rate, \dot{m}_f is the fuel mass flow rate, and K_λ is the stoichiometric air fuel ratio times the oxygen mass fraction of air. The rate of catalyst oxygen storage and release can be estimated using the relationship

$$\dot{\theta} = \dot{m}_f K_\lambda [(\Delta\lambda^{\text{pre}} - \Delta\lambda^{\text{post}}) + \epsilon_\lambda(\Delta H_2)] \quad (2)$$

where $\Delta\lambda = \lambda - 1$ is the difference from stoichiometry of the normalized air fuel ratio λ , $\Delta\lambda^{\text{pre}}$ is the pre-catalyst air fuel ratio deviation, and $\Delta\lambda^{\text{post}}$ is the post-catalyst air fuel ratio deviation. The term $\epsilon_\lambda(\Delta H_2)$ represents a dynamically varying measurement bias term due primarily to differences between the pre-catalyst and post-catalyst hydrogen concentrations [1].

Assuming that the measurement bias is negligible, the catalyst control system successfully maintains the post-catalyst air fuel ratio at stoichiometry ($\Delta\lambda^{\text{post}} = 0$), and oxygen storage and release rates are proportional to some function of temperature, Eq. 2 can be simplified to

$$\dot{\theta} = f_T(T) \dot{m} K_\lambda \Delta\lambda^{\text{pre}} \quad (3)$$

in which $f_T(T)$ is the temperature dependence. Integrating Eq. 3 with respect to time gives

$$\theta = \theta_o + \int f_T(T) \dot{m} K_\lambda \Delta\lambda^{\text{pre}} dt \quad (4)$$

where θ_o is the initial stored oxygen level. Further assuming that the post-catalyst EGO sensor voltage signal represents a measure of the catalyst stored oxygen level leads to the following relationship between the pre-catalyst air fuel ratio and the post-catalyst EGO sensor voltage

$$v_H = v_o + \int b \dot{m} \Delta\lambda^{\text{pre}} dt \quad (5)$$

where v_H is the post-catalyst EGO sensor voltage and b is the catalyst gain.

Although the exact relationship between the post-catalyst EGO sensor voltage v_H and the stored oxygen level θ is unknown, it can be assumed that this relationship is nonlinear. The effect of this unknown nonlinearity will be absorbed into the catalyst gain making b a function of θ in addition to temperature as follows

$$b = f_\theta(\theta) f_T(T) K_\lambda \quad (6)$$

where $f_\theta(\theta)$ represents the stored oxygen relationship. The catalyst gain is therefore a measure of both the EGO sensor voltage relationship to the oxygen storage level and the oxygen storage and release rate relationship to catalyst temperature. It is changes in the catalyst gain values due to the changes in these relationships as the catalyst systems degrades with age that are expected to indicate the catalyst system health.

An estimate of the catalyst gain can be computed from the derivative of Eq. 5 as follows.

$$\hat{b} = \frac{\dot{v}_H}{\dot{m}_f \Delta\lambda^{\text{pre}}} \quad (7)$$

A slowly varying post-catalyst EGO sensor voltage results in a small catalyst gain that is indicative of a healthy catalyst. Conversely, rapidly changing post-catalyst EGO sensor voltage signals coupled with small pre-catalyst stoichiometric deviations and low mass flow rates result in a large catalyst gain that is indicative of an unhealthy catalyst. A critical restriction on this metric, however, concerns the $\Delta\lambda^{\text{post}} = 0$ assumption leading to Eq. 3. If this assumption is violated to the extent that the post-catalyst EGO sensor voltage saturates from either a rich or lean emissions breakthrough, \dot{v}_H will approach zero and the resulting catalyst gain will indicate a healthy catalyst when the converse is the most likely cause of the breakthrough. In practice, monitoring for post-catalyst EGO sensor voltage saturation is necessary to avoid this scenario. A second restriction concerns the catalyst temperature. Low temperatures produce decreased catalytic activity and reduced oxygen storage capacity that result in an increased catalyst gain even for a healthy catalyst. In practice, a minimum entry condition on the estimated catalyst temperature is necessary for diagnosis.

Examination of Eq. 7 also reveals potential computational issues with the catalyst gain. The control system will typically cycle the pre-catalyst air fuel ratio $\Delta\lambda^{\text{pre}}$ across stoichiometry in a manner similar to relay feedback control. The value of $\Delta\lambda^{\text{pre}}$ will therefore go through zero at every cycle resulting in an undefined catalyst gain at that point. Assuming that the controller instantly achieves its pre-catalyst target as it is cycled, $\Delta\lambda^{\text{pre}}$ can be replaced by a gain that represents the magnitude of the stoichiometric de-

viation of the controller target. The use of this gain will prevent undefined values.

A second computational issue is the phase difference between \dot{v}_H and \dot{m}_f arising from the delay in the post-catalyst EGO sensor response. This phase difference can significantly effect the catalyst gain calculation at the relatively high switching frequency of the post-catalyst EGO sensor. Attempting to remove this phase difference by adding a delay to the \dot{m}_f value is problematic because the EGO sensor delay is a strong function of both speed and load. A simpler and more robust, although clearly approximate, alternative is to ignore the fuel mass flow rate dependence in the catalyst gain estimate. Variations in fuel mass flow rate will then propagate into the gain estimate, but the effect is likely to be more pronounced for unhealthy catalysts.

After implementing these computational modifications, a simplified catalyst gain estimate is

$$\hat{b} = \frac{\dot{v}_H}{K} \quad (8)$$

where K is the magnitude of the pre-catalyst controller target deviation from stoichiometry. The resulting catalyst gain metric is a scaled post-catalyst EGO sensor voltage time derivative where the magnitude of the scaling factor K is adapted by the catalyst control system based on the estimate of the catalyst gain. Larger values of the catalyst gain, indicating an unhealthy catalyst, result in a reduction in the stoichiometric deviation magnitude of the controller target.

3. Probability Density Metric

Experimental vehicle data from a 4.6 liter ULEV II gasoline engine operating over an EPA Federal Test Procedure (FTP) drive cycle is used to determine the ability of the test statistic to discriminate between catalysts of varying health. Three differently aged catalysts were used: healthy catalyst aged to useful life, critically aged catalyst, and unhealthy aged catalyst. The healthy catalyst was thermally aged to the baseline emission limit. The critically aged catalyst was thermally aged and poisoned with phosphorus to just meet the emission threshold limit for on-board diagnostic monitoring. The unhealthy aged catalysts was thermally aged and poisoned with phosphorus to just exceed the on-board diagnostic

threshold limit. Additional details on these catalysts are given in [2].

Previous work in [3] considered the mean, the variance to mean ratio, and the coefficient of variation of the absolute value of the estimated catalyst gain in Eq. 8 as test statistics. These values were computed from a data window consisting of the entire FTP drive cycle with a minimum temperature entry condition of 800° F. Similar test statistics computed from a slightly different catalyst gain for the same FTP drive cycle data are presented in [2]. Although these test statistics were able to distinguish the healthy catalyst from the aged catalysts, they were not able to discriminate between the two aged catalysts.

This difficulty in discriminating between the aged catalysts reflects the very close similarity in the aging process. The two catalysts are nearly identical. The only difference is the level of phosphorus poisoning (5% for the critically aged catalyst and 6% for the unhealthy aged catalyst). However, examination of the estimated cumulative distribution function of the catalyst gain for each of the three catalysts carried out in [3] did reveal a clear difference between the distributions for the two aged catalysts around zero. It is this difference that is exploited to develop a more sensitive OBD detection test statistic.

3.1. Density Estimate

In order to quantify the differences between the distributions around zero for the three catalysts, kernel density estimates were computed using the R statistical computing package [4]. These estimates were calculated using a Gaussian kernel density estimator with a bandwidth given by

$$h = 0.9 \frac{\min(s, \text{IQR}/1.34)}{N^{0.2}} \quad (9)$$

where s is the sample standard deviation of the catalyst gains in the data window, IQR is the interquartile range, and N is the sample size [5]. As shown in Figure 1, the three catalysts have different density estimates for catalyst gain values close to zero that allow the three catalysts to be clearly differentiated using FTP cycle data.

3.2. Fraction of Small Gain Values

A much more computationally efficient approximation to the kernel density around zero can be

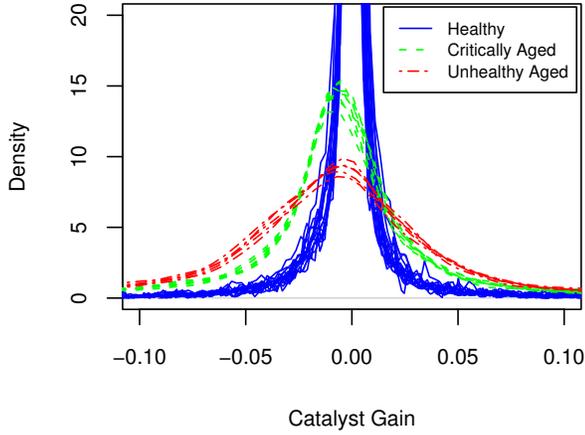


Figure 1: Density plot.

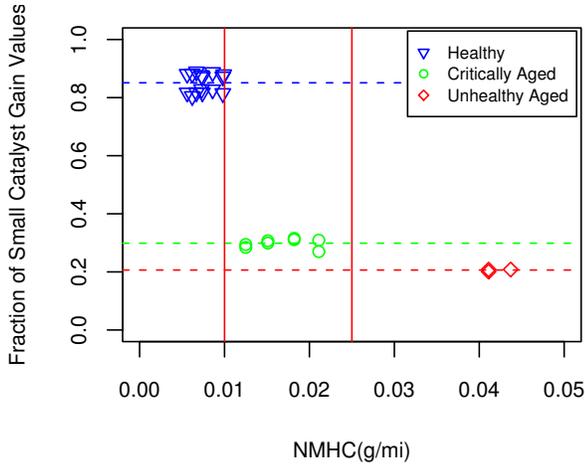


Figure 2: Fraction Δb vs. HC emissions.

obtained by considering the fraction of catalyst gain values that are within some small threshold around zero. This fraction of small catalyst gain values is defined as

$$\Delta b = \frac{\sum_{i=1}^N \eta_i}{N}, \quad \eta_i = \begin{cases} 1, & |\hat{b}_i| \leq \mathcal{T}_b \\ 0, & |\hat{b}_i| > \mathcal{T}_b \end{cases} \quad (10)$$

where N is the number of values in the moving window, \hat{b}_i is defined in Eq. 8, and $\mathcal{T}_b = 0.01$ is the threshold. Figures 2 and 3 present the fraction of small catalyst gain values in Eq. 10 for all of the FTP drive cycle data with the three differently aged catalysts as a function of hydrocarbon and NOx emissions. The entire FTP cycle is taken as the data window and a minimum temperature entry condition of 800° F is applied. The first vertical red line in each figure represents the

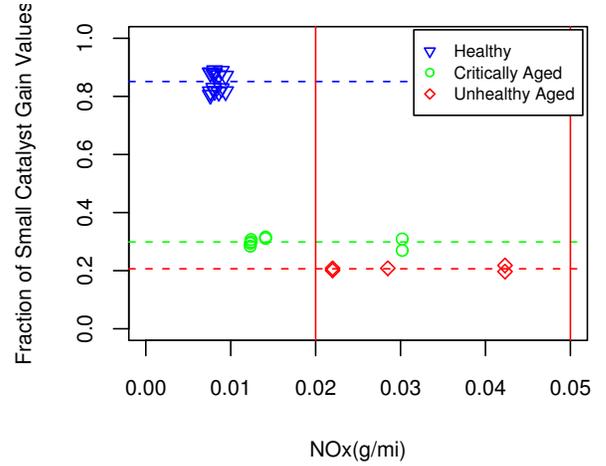


Figure 3: Fraction Δb vs. NOx emissions.

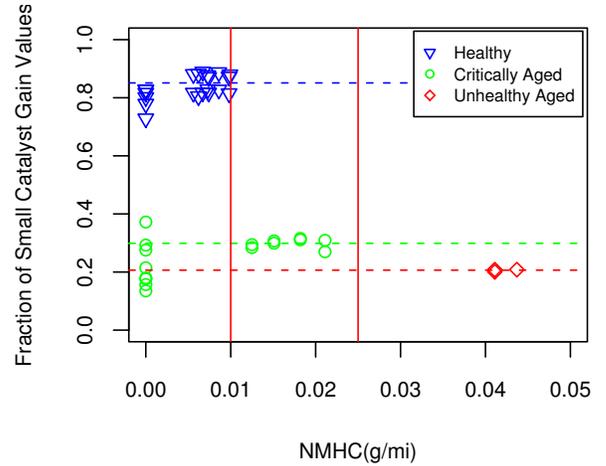


Figure 4: Fraction Δb vs. HC (all data sets).

baseline emission limit and the second vertical red line represents the on-board diagnostic detection emission limit. The dashed horizontal lines represent the average fraction over all of the FTP drive cycles for each catalyst. The fraction of small catalyst gain values, Δb , demonstrates good discrimination between the healthy and aged catalysts and can distinguish between the critically aged and unhealthy aged catalysts.

Figure 4 presents the same test statistic applied to street drive test data. Because there are no emission measurements for the street drive tests, the fraction of small catalysts gain values for this data were given a zero emissions value and incorporated with the results in Figure 2. The fraction of small catalyst gain values for street drive data also demonstrates good discrimination between the healthy and critically aged

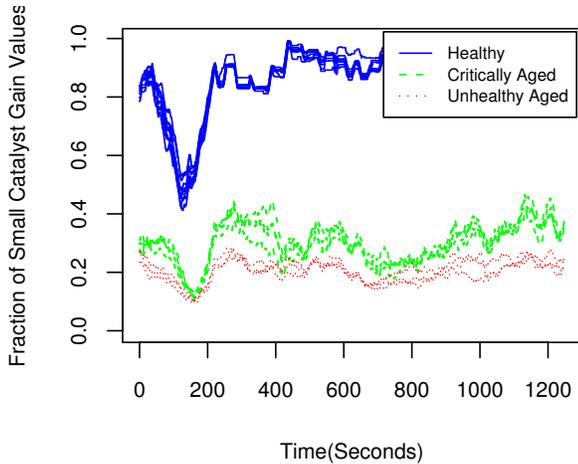


Figure 5: Fraction Δb time series (all FTP cycles).

catalysts. However, there is more variation between the Δb values for different tests and the average Δb for each catalyst is less than the FTP drive cycle average. There is no street drive data for the unhealthy aged catalyst.

The fraction of small catalyst gain values demonstrates good discrimination between the three differently-aged catalysts for the FTP drive cycles when the entire cycle is taken as the data window. However, if the detection method is to be effective for all types of driving, it should not be overly dependent on the particular drive cycle or data window chosen. Figure 5 presents the times series of Δb values for overlapping 100 sec windows and demonstrates that the Δb values vary considerably over the course of drive cycle for each catalyst. Therefore, no single diagnostic threshold is appropriate for *all* windows.

There is a negative linear correlation between the Δb values in Figure 5 and the mean air mass flow rate for the same window. For example, the linear correlation between Δb and the mean air mass flow rate for one FTP cycle of the healthy catalyst is -0.93 indicating a strong negative linear correlation. The linear correlation is -0.64 for one FTP cycle of the critically aged catalyst and -0.49 for the unhealthy aged catalyst indicating a more moderate negative linear correlation. These results are representative of all of the FTP cycle data where the correlation between the Δb values and the air mass flow rate decreases with catalyst age. This relationship between the Δb values and the air flow rate suggests that air mass flow could be included in the diagnostic monitor.

4. Two Dimensional Diagnostic Monitor

Because a relationship exists between the Δb values and the air mass flow rate, an on-board diagnostic detection threshold that incorporates both variables should be more discriminating. A sample of 100 sec non-overlapping windows from each FTP cycle was used to determine a detection threshold using three different criteria. The first criterion minimizes the total number of misclassifications subject to the number of critically aged catalysts misclassifications is equal to the number of unhealthy aged catalysts misclassifications. The second criterion minimizes the total number of misclassifications with no restrictions on type of catalyst. The last criterion minimizes the number of unhealthy aged catalysts misclassifications subject to no critically aged catalysts misclassifications. This criterion is of practical relevance because avoiding healthy catalyst misdiagnosis is critical for implementation.

Figures 6, 7, and 8 present the detection threshold line for each criterion. The slope and intercept for each of the detection threshold lines is given in Table 1. Data pairs that fall on or below the detection threshold line are indicative of an unhealthy catalyst. The total number of misclassifications using the first criterion is 18 resulting in nine misclassifications for both catalysts. There are 17 misclassifications using the second criterion with 11 for the critically aged and six for the unhealthy aged catalyst. There are 22 misclassifications using the final criterion all of which are for the unhealthy aged catalyst. For each criterion, there are no healthy catalyst misclassifications. This diagnostic metric can clearly distinguish between a useful life catalyst and a catalyst that is approaching the OBD detection limit. We also note that the data windows used in this study are much smaller than would be implemented in practice due to the limited FTP cycle data available in this study. Although acceptable performance is obtained with the 100 sec windows, better diagnostic discrimination would be possible with larger windows.

5. Conclusions

The results of this study indicate that the fraction of time the catalyst gain is close to zero is a sensitive detection metric for on-board diagnostic monitoring of an automotive catalyst. This

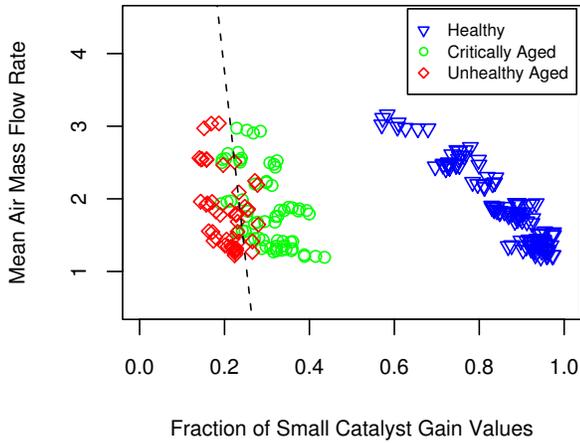


Figure 6: Equal number of misclassifications.

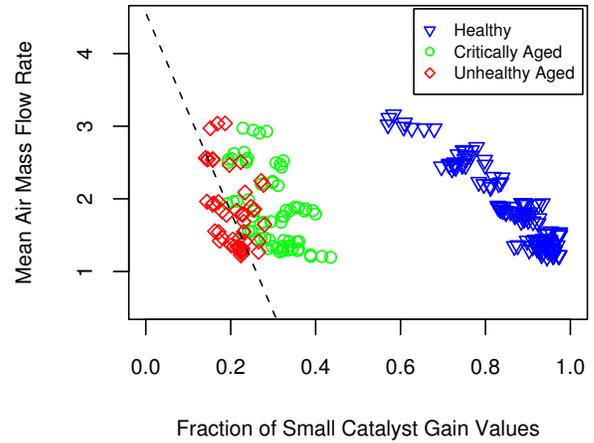


Figure 8: No critically aged misclassifications.

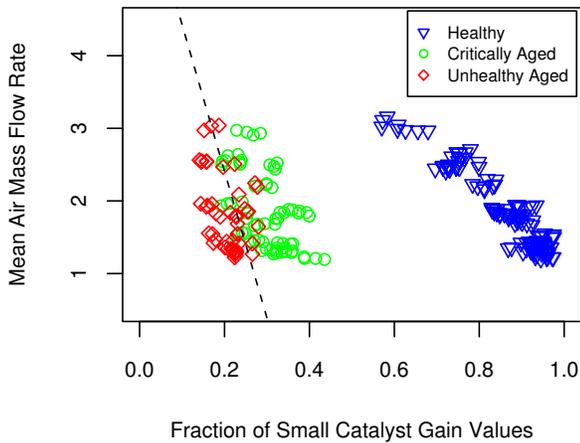


Figure 7: Minimum number of misclassifications.

metric is able to distinguish between the three differently aged catalysts over a series of FTP drive cycles. One of the major advantages to this metric is that it does not require any specific operating conditions and therefore is straightforward to implement. Future work includes testing this metric on additional vehicle platforms over a wider range of driving conditions.

Acknowledgment

Support for this work from Ford Motor Company, Johnson Matthey, and ExxonMobil is gratefully acknowledged.

References

[1] J. Peyton Jones and R. Jackson. Potential and pitfalls in the use of dual EGO sensors for 3-way catalyst monitoring and control. *Proc.*

Table 1: Detection threshold slope and intercept.

Criterion	Slope	Intercept
1	-52.5	14.2
2	-20.2	6.4
3	-13.7	4.5

Inst. Mech. Engin. Part D: J. Auto. Engin., 217:475–488, 2003.

[2] J. Peyton Jones, I. Makki, and K. Muske. Catalyst diagnostics using adaptive control system parameters. SAE Paper 2006-01-1070, SAE World Congress, 2006.

[3] K. Muske, J. Peyton Jones, I. Makki, M. Urich, and J. Howse. An evaluation of detection metrics for an integrated catalyst controller and diagnostic monitor. In *Proceedings of the Fifth IFAC Symposium on Advances in Automotive Control*, pages 191–198, 2007.

[4] R Development Core Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, 2007.

[5] B. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London, 1986.