ON-BOARD DIAGNOSTIC AND FAULT DETECTION STRATEGIES FOR AN AUTOMOTIVE THREE-WAY CATALYST

Kenneth R. Muske [*] **James C. Peyton Jones** [**]

[*] *Department of Chemical Engineering, Villanova University, Villanova, PA 19085*
[**] *Department of Electrical and Computer Engineering, Villanova University, Villanova, PA 19085*

Abstract: This work presents model-based strategies for on-board catalyst monitoring and fault detection using a simplified catalyst oxygen storage and reversible deactivation model. The performance of the catalyst is inferred from the error between the post-catalyst exhaust gas oxygen sensor measurement and the model predicted value. Each of these strategies is based on the use of some test statistic that is computed from a window of post-catalyst prediction error data updated in real-time. A fault is reported when the value of this statistic exceeds a threshold that is determined from some specified confidence level. We conclude with an evaluation of these strategies.

Keywords: Automotive Catalyst Monitoring, Automotive On-board Diagnostic, Moving Horizon Monitoring

1. INTRODUCTION

In order to ensure that the components of a catalytic automotive emission control system continue to operate properly with age, effective and robust on-board system monitoring is essential. Because emissions from a small fraction of malfunctioning vehicles are believed to account for a large fraction of the total automotive-related emissions, on-board monitoring has the potential to significantly reduce hydrocarbon and nitrogen oxide air pollutants. This observation, along with recent legislation such as SULEV and OBD-II that mandate 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 the emission control system is no longer meeting specification or when a fault is present.

On-board monitoring of the three-way catalyst component of catalytic automotive emission control systems is considered in this work. The specifications for catalyst operation are set based on the tailpipe hydrocarbon emissions. Under European legislation, these emissions must remain less than 0.4g/km during a specified test cycle. In the United States, OBD-II legislation requires hydrocarbon emissions to remain less than 1.5 times the 4000 mile emissions baseline. These requirements are particularly challenging due to the lack of a cost-effective and reliable automotive hydrocarbon sensor. This limitation has led to monitoring approaches that consider the oxygen storage dynamics of the catalyst rather than the hydrocarbon conversion efficiency. Although the oxygen storage capacity of the catalyst can not be measured directly, it can be inferred using pre- and post-catalyst exhaust gas oxygen (EGO) sensors. These sensors are often available as part of the engine air fuel ratio (AFR) control system making these monitoring approaches attractive.

2. CATALYST OXYGEN STORAGE

A key process in the reduction of tailpipe emissions by the three-way catalyst is the ability to store and release oxygen in response to the pre-catalyst exhaust air fuel ratio. Under rich (excess fuel) engine operation, the catalyst oxidizes the hydrocarbons and carbon monoxide present in the incoming engine exhaust by releasing previously stored oxygen. This oxygen release maintains stoichiometric combustion with commensurately low levels of hydrocarbon and carbon monoxide tailpipe emissions. Because of the finite storage capacity of the catalyst, however, this process cannot continue indefinitely. When the oxygen release rate of the depleted catalyst can no longer satisfy the demand, the post-catalyst air fuel ratio will decrease below stoichiometric and hydrocarbon breakthrough will eventually occur. A typical catalyst control system will therefore attempt to switch to lean (excess air) engine operation before this rich breakthrough condition is encountered. Under lean engine operation, the excess oxygen in the engine exhaust gas is now adsorbed onto the catalyst resulting in near-stoichiometric post-catalyst conditions and low tail-pipe emissions. As the oxygen storage capacity of the catalyst approaches its saturation condition, however, the post-catalyst oxygen concentration increases above stoichiometric and emission of nitrogen oxides will eventually occur. A typical catalyst control system will then attempt to switch back to rich engine operation before lean breakthrough. Through this cycling of the pre-catalyst air fuel ratio, the oxygen storage capacity of the catalyst is a buffer against tailpipe emissions by compensating for transient oxygen excess or deficiency.

Because of the nature of the three-way catalyst operation, a measure of the catalyst oxygen storage capacity would be expected to correlate well with the hydrocarbon conversion efficiency. In practice, however, the correlation is nonstationary and highly nonlinear resulting from differences in both the short-term and long-term dynamics between oxygen storage and hydrocarbon conversion. A further difficulty in the use of oxygen storage capacity for catalyst monitoring lies in the use of the post-catalyst EGO sensor. The sensor signal is not only a function of the exhaust gas air fuel ratio but also a function of the exhaust gas composition. In pre-catalyst applications, the dependency on gas composition can be largely removed by calibration. The composition of the post-catalyst exhaust, however, varies dynamically according to the reactions taking place on the catalyst. Static calibration in this case is insufficient to correct for composition dependent distortion. The resulting error is difficult to distinguish from oxygen storage and release effects and, unless treated explicitly, will propagate to the catalyst monitoring strategy.

Despite these difficulties, oxygen storage capacity (OSC) metrics for on-board monitoring have been proposed based on the assumption that the oxygen storage rate can be estimated from the difference between the pre-catalyst and post-catalyst EGO sensor measurements of the air fuel ratio. In order to measure catalyst OSC, the oxygen storage rate must be integrated over an appropriate time interval during which the catalyst oxygen storage state changes from an initial to a final state. It is the choice of the initial and final state of the integration that distinguishes these methods. The repeatability of the metric also depends on the accuracy with which the catalyst can be driven to the initial state and the determination of the point at which the catalyst has reached the final state. Measurement of these conditions is not straightforward because the level of stored oxygen cannot be observed directly.

A summary of these OSC metrics, along with the requirements on the engine and catalyst operation necessary to determine their value, is presented by Peyton Jones and Muske (2004). An inherent difficulty with diagnostic strategies based on these metrics is that they seek to encapsulate catalyst behavior using a single oxygen storage capacity determination method that is assumed indicative of catalyst health. Because catalyst dynamics are sufficiently complex, it is unlikely that such a metric can adequately describe hydrocarbon conversion efficiency for a variety of operating conditions. For this reason, model-based strategies are likely to be more accurate and discriminating.

A simple model-based monitoring strategy is to use the error between a model predicted and measured OSC metric as the catalyst diagnostic. When the catalyst is healthy, the model prediction and actual measurement should be in close agreement resulting in small prediction errors. As the catalyst performance degrades, the magnitude of these errors would be expected to increase. The effectiveness of this approach, however, depends on the model and the accuracy of both the measurement and the model prediction of the metric. A second consideration is the engine and catalyst operation that is necessary in order to provide sufficient information on the catalyst performance.

3. OXYGEN STORAGE MODEL

The oxygen storage model of Peyton Jones et al. (2000) is used in the monitoring strategies presented in this work. The catalyst oxygen storage is represented by the nonlinear integrating model

$$\dot{\phi} = K_\lambda \dot{m}_f \left(\Delta\lambda^{\triangleleft} - \mathcal{N}(\phi) \right) \tag{1}$$

in which ϕ is the oxygen storage state of the catalyst, $\Delta\lambda^{\triangleleft}$ is the pre-catalyst air fuel ratio deviation from stoichiometric, $\dot{m}_f$ is the fuel mass

flow rate to the engine, and $\mathcal{N}(\phi)$ is the catalyst oxygen capacity function. Here oxygen storage and release rates depend on the difference between the forcing function $\Delta\lambda^{\triangleleft}$, which promotes adsorption under conditions of oxygen excess, and the catalyst capacity $\mathcal{N}(\phi)$, which represents the desorption under similar conditions. This model describes stored oxygen relative to the equilibrium level when the pre-catalyst exhaust gas is stoichiometric, $\Delta\lambda^{\triangleleft} = 0$. We note that the pre-catalyst air fuel ratio deviation, $\Delta\lambda^{\triangleleft}$, is denoted using a left-facing triangle and that the post-catalyst air fuel ratio deviation, $\Delta\lambda^{\triangleright}$, is denoted using a right-facing triangle in this work.

The capacity function $\mathcal{N}(\phi)$ has a nonlinear spring characteristic. As ϕ increases from zero, it becomes progressively harder to store oxygen on the catalyst. As ϕ decreases from zero, it becomes progressively harder to remove oxygen from the catalyst. In this work, $\mathcal{N}(\phi)$ is parameterized using a polynomial expansion of ϕ.

$$\mathcal{N}(\phi) = a_1\phi + a_2\phi^2 + a_3\phi^3 + a_4\phi^4 + a_5\phi^5 \quad (2)$$

We note that the stored oxygen level does not hit hard saturation/depletion limits as in limited integrator models but instead attains an operating condition that is dependent on the steady-state pre-catalyst air fuel ratio. The oxygen saturation/depletion limits are approached asymptotically as would be expected from theory.

The post-catalyst air fuel ratio deviation from stoichiometric, $\Delta\lambda^{\triangleright}$, is determined from the capacity function by the following relationship.

$$\Delta\lambda^{\triangleright} = \begin{cases} 0, & (\Delta\lambda^{\triangleleft} < 0) \text{ and } (\phi > 0) \\ \mathcal{N}(\phi), & \text{otherwise} \end{cases} \quad (3)$$

In the case of rich pre-catalyst exhaust gas, $\Delta\lambda^{\triangleleft} < 0$, and an oxidized catalyst, $\phi > 0$, where stored oxygen is available for reducing the rich incoming exhaust gas, oxygen release is limited by the exhaust gas oxygen demand resulting in stoichiometric post-catalyst exhaust, $\Delta\lambda^{\triangleright} = 0$. This effect is clearly shown in Figure 1 where the post-catalyst air fuel ratio remains at stoichiometric for several seconds after the pre-catalyst air fuel ratio makes a lean to rich step change. As the stored oxygen becomes depleted from the catalyst, the post-catalyst air fuel ratio becomes rich until it eventually matches the pre-catalyst ratio. Oxygen storage on the catalyst is clearly shown after the pre-catalyst air fuel ratio rich to lean step change made at approximately seventy seconds.

The separation between the actual and model predicted air fuel ratios appearing after twenty seconds in Figure 1 is due to sensor distortion in the post-catalyst air fuel ratio sensor. As shown by Peyton Jones (2003), post-catalyst EGO sensor distortion under rich conditions can be related to

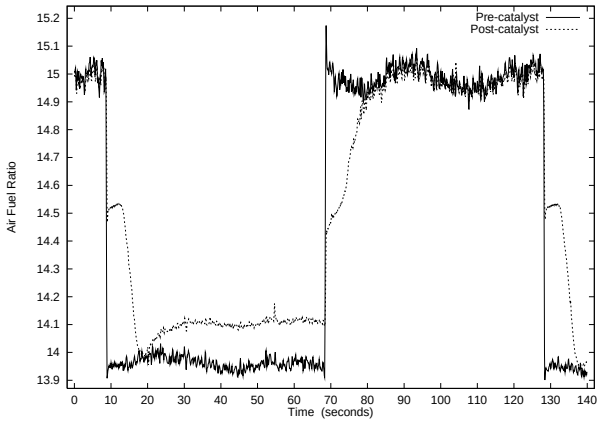

Fig. 1. Pre- and post-catalyst air fuel ratio.

the hydrogen generated by catalytic promotion of the water gas shift reaction.

$$\mathrm{CO} + \mathrm{H_2O} \rightleftharpoons \mathrm{CO_2} + \mathrm{H_2}$$

If the reaction proceeds in the forward direction, the ability of hydrogen to diffuse faster than hydrocarbons and carbon monoxide results in a sensor output that appears richer than the true value. Reduced levels of post-catalyst hydrogen due to a progressive inhibition of the water gas shift reaction result in a sensor output that appears leaner that the true value. The rise toward leaner values of the measured post-catalyst air fuel ratio in Figure 1 is due to this inhibition effect.

The reversible short-term catalyst deactivation effects that inhibit hydrogen production through the water gas shift reaction can also have an effect on conversion efficiency particularly under rich or rich-biased operation. This behavior suggests that the gas-composition dependent error in the post-catalyst air fuel ratio sensor measurement can be used as a measure of this reversible catalyst deactivation effect. If the degree of water gas shift reaction inhibition is assumed proportional to the deactivated fraction of the catalyst surface, ψ, then the apparent post-catalyst AFR, $\Delta\lambda_a^{\triangleright}$, can be related to the true post-catalyst AFR, $\Delta\lambda^{\triangleright}$, as

$$\Delta\lambda_a^{\triangleright} = \Delta\lambda^{\triangleright} + K_H\psi \quad (4)$$

where the constant K_H represents the sensor sensitivity to hydrogen concentration changes and the inhibition of the water gas shift reaction due to reversible catalyst deactivation.

Assuming that the rate of deactivation is proportional to the post-catalyst oxygen deficiency, $-\Delta\lambda^{\triangleright}$, and the fraction of the surface that is already deactivated, ψ, results in the following model for the deactivated fraction of the catalyst surface

$$\dot{\psi}_d = \dot{m}_f K_d(-\Delta\lambda^{\triangleright} - \psi) \quad (5)$$

where K_d is the deactivation constant of proportionality. The presence of pre-catalyst free oxygen, $\Delta\lambda^{\triangleleft} > 0$, reverses the deactivation process at

a rate proportional to the supply of pre-catalyst oxygen until the catalyst is reactivated, $\psi = 0$,

$$\dot{\psi}_r = -\dot{m}_f K_r \Delta\lambda^{\triangleleft} \tag{6}$$

where K_r represents the reactivation constant. The resulting reversible deactivation model is

$$\dot{\psi} = \begin{cases} \dot{\psi}_d, & (\Delta\lambda^{\triangleleft} < 0) \text{ and } (\Delta\lambda^{\triangleright} < 0) \\ \dot{\psi}_r, & (\Delta\lambda^{\triangleleft} > 0) \text{ and } (\psi > 0) \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

where $\Delta\lambda^{\triangleright}$ is the true post-catalyst AFR determined from the oxygen storage model in Eq. 3 and the sensor distortion of the measured post-catalyst AFR is determined from Eq. 4.

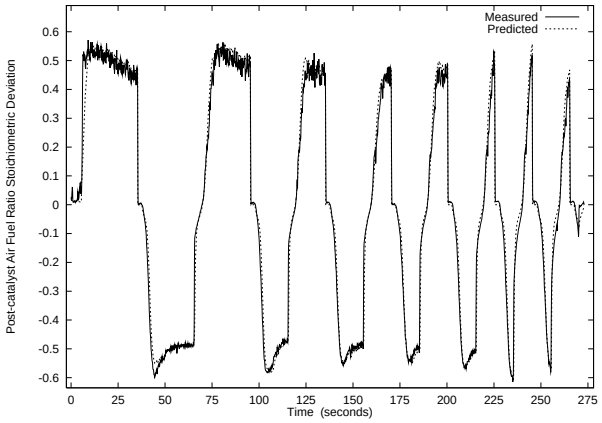

Fig. 2. Post-catalyst AFR model prediction.

Figure 2 presents a comparison between the measured post-catalyst air fuel ratio stoichiometric deviation and the combined oxygen storage-reversible catalyst deactivation rate model prediction for a series of step changes to the pre-catalyst air fuel ratio. As shown in this figure, the deviations at the end of each rich cycle are accounted for in the combined model. In particular, the combined model now characterizes the distortion that occurs in the post-catalyst sensor since there are no longer significant unmodeled dynamics. The close agreement between the predicted and actual measurements, especially during the periods of rich operation, demonstrates that the reversible catalyst deactivation dynamics are well described by this model.

4. STATISTICAL MODEL-BASED MONITORING STRATEGIES

In order to account for the uncertainty in both the measured and model predicted OSC metric, analysis of the prediction error is best carried out within a statistical framework. Previous work in statistical model-based monitoring strategies considered the Neyman-Pearson criterion to obtain a diagnosis at a given level of confidence. Analysis of the difference between the pre- and post-catalyst AFR measurements is considered by Arsie et al. (2000). The student t-test applied to multiple samples of the time to empty/fill the catalyst is presented by Brandt and Grizzle (2001). The disadvantage of these methods is the Gaussian distribution assumption for the metric which is not necessarily the case, especially given the biasing effects of sensor distortion. Catalyst deactivation effects may also change the distribution over time.

In this work, we consider the prediction error of the post-catalyst AFR as the basis for OSC metrics. The prediction error is easily computed and sensitive to both long-term catalyst deactivation effects and short-term emission control system device failures. It does not require catalyst operation that can result in increased tailpipe emissions or disruption in driveability as is possible during the direct determination of oxygen storage capacity. The monitoring strategies are based on a sample of the post-catalyst AFR prediction errors

$$\boldsymbol{e} = [e_1, e_2, \ldots, e_N]^T \tag{8}$$

where N is the sample population size, $e_k = \Delta\lambda^{\triangleright}_{m,k} - \Delta\lambda^{\triangleright}_{a,k}$ is the prediction error, $\Delta\lambda^{\triangleright}_{m,k}$ is the measured post-catalyst AFR, and $\Delta\lambda^{\triangleright}_{a,k}$ is the predicted post-catalyst AFR computed from Eq. 4 for samples $k = 1, \ldots, N$. The sample may be constructed using either a moving window or a batch approach.

We begin our discussion of monitoring strategies by analyzing the distribution of the post-catalyst AFR prediction errors for the engine operation in Figure 2. This distribution, shown in Figure 3, is nonsymmetric with a greater proportion of rich prediction errors but with larger lean prediction errors on average.

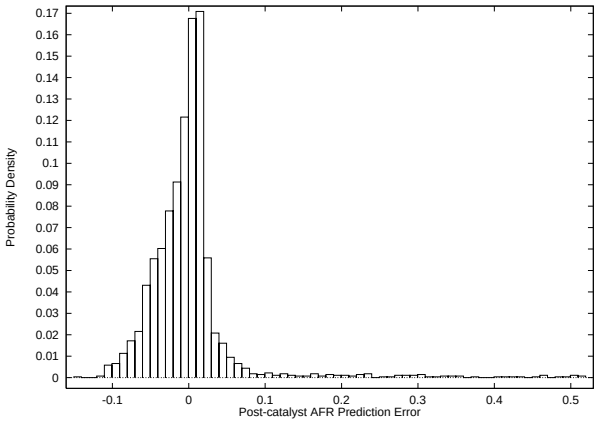

Fig. 3. Nominal prediction error distribution.

In order to estimate catalyst deactivation effects, the capacity function in Eq. 2 is compressed as

$$\tilde{\mathcal{N}}(\phi) = \frac{a_1}{\beta}\phi + \frac{a_2}{\beta^2}\phi^2 + \frac{a_3}{\beta^3}\phi^3 + \frac{a_4}{\beta^4}\phi^4 + \frac{a_5}{\beta^5}\phi^5 \tag{9}$$

where $0 < \beta \leq 1$ is the fraction of oxygen storage capacity remaining. Simulated post-catalyst AFR operating data for a deactivated catalyst is constructed by computing the model predicted post-catalyst AFR using the modified catalyst capacity

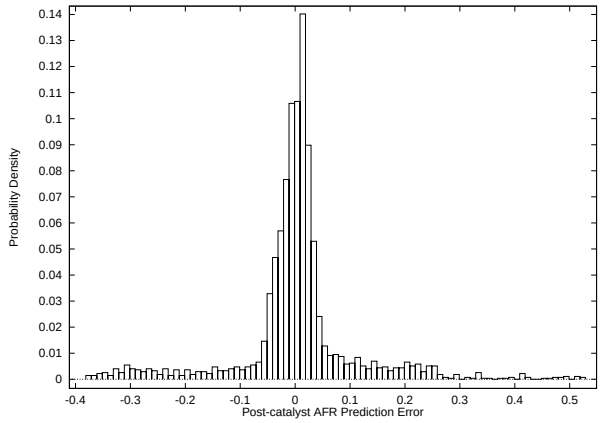

Fig. 4. Deactivated catalyst error distribution.

function in Eq. 9 with the same pre-catalyst conditions as Figure 2. The model prediction errors from Figure 2 are then inserted into this simulated data set. We choose these errors, as opposed to the addition of white noise, because the model prediction errors do show some correlation with the pre-catalyst AFR and we wish to retain this property. Figure 4 presents the prediction error distribution between the simulated deactivated operation with $\beta = 0.8$ and the model prediction with $\beta = 1$. As shown in this figure, there is a significant deviation from the nominal operation.

4.1 Mean Prediction Error Metric

In this section, we consider a statistical test on the mean of the population of prediction errors in the current sample as discussed by Christensen (1996). Provided that the sample population is large, a standard normal Gaussian distribution can be reasonably assumed for the test statistic

$$\mathcal{Z} = \frac{\bar{\boldsymbol{e}}_{\mathcal{Y}} - \bar{\boldsymbol{e}}_{\mathcal{X}}}{s_{\mathcal{Y}}^2/N_{\mathcal{Y}} + s_{\mathcal{X}}^2/N_{\mathcal{X}}} \tag{10}$$

in which $\boldsymbol{e}_{\mathcal{Y}}$ is the prediction error sample drawn from the current operation with mean $\bar{\boldsymbol{e}}_{\mathcal{Y}}$, variance $s_{\mathcal{Y}}^2$, and sample size $N_{\mathcal{Y}}$ and $\boldsymbol{e}_{\mathcal{X}}$ is the nominal operation prediction error sample with mean $\bar{\boldsymbol{e}}_{\mathcal{X}}$, variance $s_{\mathcal{X}}^2$, and sample size $N_{\mathcal{X}}$. The null hypothesis is $\bar{\boldsymbol{e}}_{\mathcal{Y}} = \bar{\boldsymbol{e}}_{\mathcal{X}}$ which indicates normal catalyst operation and is accepted if the test statistic $\mathcal{Z}$ is less than the threshold $\mathcal{T}_{\mathcal{Z}}(\alpha)$ chosen based on a significance level α. The alternative hypothesis is $\bar{\boldsymbol{e}}_{\mathcal{Y}} > \bar{\boldsymbol{e}}_{\mathcal{X}}$ which indicates improper catalyst operation. Because recursive estimates of the prediction error and its variance are easily computed, large sample populations are feasible which both reduces the sensitivity to outliers and improves the Gaussian distribution assumption made for $\mathcal{Z}$.

The symmetry of the deactivated distribution in Figure 4 suggests that the mean prediction error may not be an effective metric and, in this application, the mean absolute or squared error would be a better choice. Table 1 presents the value of the test statistic $\mathcal{Z}$ using each of these metrics for varying levels of catalyst oxygen storage capacity. A sample size of $N = 1000$ was used for each population where the samples for $\boldsymbol{e}_{\mathcal{X}}$ were obtained by random sampling of the prediction errors from the operation in Figure 2 and the samples for $\boldsymbol{e}_{\mathcal{Y}}$ were obtained by random sampling of the simulated deactivated operating data with varying values of β. As shown in this table, the metric based on mean absolute error is quite sensitive to small reductions in the oxygen storage capacity. Note that the threshold value is $\mathcal{T}_{\mathcal{Z}}(0.025) = 1.96$ for a 0.025 level of significance.

Storage	*Test Statistic* $\mathcal{Z}$		
Capacity	$\bar{\boldsymbol{e}}$	$\lvert\boldsymbol{e}\rvert$	$\boldsymbol{e}^2$
95%	0.55	2.42	1.21
90%	1.00	6.38	3.92
85%	1.33	10.6	7.64
80%	1.56	14.3	11.4
75%	1.69	17.2	14.5
70%	1.78	19.6	17.0
65%	1.83	21.7	19.0
60%	1.86	23.6	20.6
55%	1.90	25.3	22.2
50%	1.94	27.0	23.7

Table 1. Test statistic for mean error.

4.2 Prediction Error Distribution Metric

The monitoring strategy employed in this section considers the distribution of the prediction errors in the samples. The use of the error distribution, as opposed to the mean, can increase sensitivity in the presence of model inadequacy. Because of the simplicity of the catalyst model, large model errors may occur even when the catalyst system is functioning properly. The error magnitude may also not change significantly when a fault is present. In this case, changes in the distribution of the errors can provide additional information on the catalyst operation. The Kolmogorov–Smirnov test, a nonparametric test for comparing cumulative distribution functions discussed by Conover (1980), is used. Although no assumptions about the form of the underlying distribution are necessary, this test does assume that the sampling is random from a continuously distributed population with sufficient measurement precision to avoid tied observations. These assumptions are well satisfied for catalyst system monitoring.

The cumulative distribution function (*cdf*) is the mapping $F_{\mathcal{X}}(z)$ that assigns the probability that a random variable $\mathcal{X}$ is less than or equal to a specific value z for all possible z values.

$$F_{\mathcal{X}}(z) = \mathcal{P}\{\mathcal{X} \leq z\}, \quad -\infty < z < \infty$$

The nominal catalyst operation prediction error is denoted by the random variable $\mathcal{X}$ and the current operation prediction error by the random variable

$\mathcal{Y}$. There are three test statistics that can be used with the Kolmogorov–Smirnov test

$$\mathcal{K} = \sup_{-\infty<z<\infty} |F_{\mathcal{Y}}(z) - F_{\mathcal{X}}(z)| \quad (11)$$

$$\mathcal{K}^+ = \sup_{-\infty<z<\infty} \big(F_{\mathcal{Y}}(z) - F_{\mathcal{X}}(z)\big) \quad (12)$$

$$\mathcal{K}^- = \sup_{-\infty<z<\infty} -\big(F_{\mathcal{Y}}(z) - F_{\mathcal{X}}(z)\big) \quad (13)$$

where $F_{\mathcal{Y}}(z)$ is the *cdf* for the current operation prediction error and $F_{\mathcal{X}}(z)$ is the *cdf* for the nominal operation prediction error. The $\mathcal{K}$ statistic in Eq. 11 is the maximum difference between the two distributions. The null hypothesis is $F_{\mathcal{Y}}(z) = F_{\mathcal{X}}(z)$ for all z indicating normal operation and is accepted if the test statistic $\mathcal{K}$ is less than the threshold $\mathcal{T}_{\mathcal{K}}(\alpha)$ chosen based on a significance level α. The alternative hypothesis is $F_{\mathcal{Y}}(z) \neq F_{\mathcal{X}}(z)$ for at least one value of z. The null hypothesis for the $\mathcal{K}^+$ statistic in Eq. 12 is $F_{\mathcal{Y}}(z) \leq F_{\mathcal{X}}(z)$ with the alternative hypothesis $F_{\mathcal{Y}}(z) > F_{\mathcal{X}}(z)$ for at least one value of z. The null hypothesis for the $\mathcal{K}^-$ statistic in Eq. 13 is $F_{\mathcal{Y}}(z) \geq F_{\mathcal{X}}(z)$ with the alternative hypothesis $F_{\mathcal{Y}}(z) < F_{\mathcal{X}}(z)$ for at least one value of z. The threshold values for each statistic are available as a function of the significance level and number of samples in each population. The *cdf* for both the nominal and current operation is approximated by a normalized histogram of the prediction errors in the corresponding sample. The normalized histogram is computed by taking the proportion of the sample population less than or equal to a specified prediction error. These approximate distribution functions approach the true distribution only in the asymptotic limit as the number of samples goes to infinity when the true underlying distribution is continuous. For large sample size, however, the density function based on a finite sample should be a reasonable approximation to the asymptotic density function. Because the calculations required to perform the test are relatively trivial, each of the three statistics in Eqs. 11 through 13 can be tested. The requirement of computing normalized histograms, however, will limit the sample population sizes in practice.

Table 2 presents average values of the test statistics $\mathcal{K}$, $\mathcal{K}^+$, and $\mathcal{K}^-$ for varying levels of catalyst oxygen storage capacity. The nominal error sample $\boldsymbol{e}_{\mathcal{X}}$ was obtained by random sampling of the prediction errors from the operation in Figure 2. For each value of β, several current operation samples $\boldsymbol{e}_{\mathcal{Y}}$ were obtained by random sampling of simulated deactivated operating data. A sample size of $N = 100$ was used for each population and the average test statistic values were determined using 1000 tests. The threshold value for $\mathcal{K}$ is $\mathcal{T}_{\mathcal{K}}(0.025) = 0.121$ for a 0.025 level of significance. Because of the smaller sample size, this test is not as sensitive as the mean absolute error in this example. When $\beta = 0.9$, the test statistic average is just below the threshold with approximately half of the samples above the threshold value.

Storage	*Test Statistic*		
Capacity	$\mathcal{K}$	$\mathcal{K}^+$	$\mathcal{K}^-$
100%	0.058	0.041	0.050
95%	0.093	0.091	0.080
90%	0.115	0.108	0.081
85%	0.138	0.131	0.093
80%	0.164	0.159	0.094
75%	0.172	0.169	0.096
70%	0.208	0.208	0.099
65%	0.228	0.228	0.115
60%	0.247	0.247	0.117
55%	0.256	0.256	0.122
50%	0.274	0.274	0.134

Table 2. Kolmogorov–Smirnov statistic.

5. CONCLUSIONS

A model-based automotive catalyst diagnostic and fault detection methodology based on the error between the measured and model predicted post-catalyst air fuel ratio is developed. Hypothesis tests using the mean absolute error and the Kolmogorov–Smirnov nonparametric test for cumulative distribution functions are used to determine statistically significant changes between the current and nominal catalyst operation. This methodology is demonstrated using engine data with simulated oxygen storage capacity reduction and shows promise for on-board application.

ACKNOWLEDGMENTS

National Science Foundation support from grant CTS-0215920 is gratefully acknowledged.

REFERENCES

E. Brandt and J. Grizzle. Three-way catalyst diagnostics for advanced emissions control systems. In *Proc. of the 2001 American Control Conference*, pages 3305–3311, 2001.

R. Christensen. *Analysis of Variance, Design and Regression.* Chapman & Hall, London, 1996.

W. Conover. *Practical Nonparametric Statistics.* Wiley, New York, 2nd edition, 1980.

I. Arsie, et al. Confidence level analysis for on-board estimation of SI engine catalytic converter efficiency. In *Proc. of the SYSID 2000 IFAC Symposium on System Identification*, pages 137–142, 2000.

J. Peyton Jones, et al. A simplified model for the dynamics of three-way catalytic converter. Paper 2000-01-0652, SAE World Congress, 2000.

J. Peyton Jones. Modeling combined oxygen storage and reversible deactivation dynamics for improved emissions predictions. Paper 2003-01-0999, SAE World Congress, 2003.

J. Peyton Jones and K. Muske. Model-based OBD for three-way catalyst systems. Paper 2004-01-0639, SAE World Congress, 2004.