Three-Way Catalyst Diagnostics for Advanced Emissions Control Systems

E. P. Brandt[1] and J. W. Grizzle[2]

Abstract

Automotive emissions are stringently regulated. Since 1980, a three-way catalyst (TWC) has been used to convert harmful emissions of hydrocarbons, carbon monoxide, and oxides of nitrogen into less harmful gases in order to meet these regulations. The TWC's efficiency of conversion of these gases is primarily dependent on the mass ratio of air to fuel (A/F) in the mixture leaving the exhaust manifold and entering the catalyst. This paper develops a method by which a dynamic TWC model can be used for diagnostic purposes. This diagnostic method is analyzed in the context of a hypothesis test that is based on the oxygen storage capacity of the TWC. The Neyman-Pearson criterion is used as the basis for this hypothesis test. It is initially applied in the case of a single sample where the variance of the data is assumed to be known. This is then expanded to a multiple-sample case through the use of Student's t test. The improved fidelity of the t test is demonstrated, and it is shown that larger sample sizes provide further improvement in the quality of the hypothesis test.

1 Introduction

Part of the increasingly tight automotive emissions control standards includes a requirement to monitor the condition of the catalyst. In addition to the requirement that a specific tailpipe emissions level be met, automobile manufacturers are required to diagnose deficiencies in the performance of the emissions control system.

This paper proposes a method for diagnosis of a three-way catalyst (TWC), the major component of the emissions control system. Using a dynamic TWC model, a hypothesis test is proposed that provides a basis for statistical confidence in the condition of the TWC. Section 2 provides a brief description of the TWC model that is used for this study. Sections 3–5 provide a basic description of the general hypothesis testing problem and its relevance to this work, applying the hypothesis testing framework to the TWC diagnostics problem. Section 6 enhances this work by showing the improvement that results through the use of multiple samples.

[1]E. P. Brandt is with the Automotive and Industrial Electronics Group, Motorola, 15201 Mercantile Drive, Dearborn, MI 48120-1236, USA. E-mail: Erich.Brandt@motorola.com .

[2]J. W. Grizzle is with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109-2122, USA. E-mail: grizzle@eecs.umich.edu . The work of J.W. Grizzle was supported in part by an NSF GOALI grant, ECS-9631237.

2 Oxygen storage model

The model that is summarized here was initially proposed in [3] and was developed further in [1, 2], based on initial work in [8]. The reader is referred to these works for a more detailed discussion.

Oxygen storage and release, the property of oxygen attaching to metal and cerium sites in the catalyst under lean conditions, thereby decreasing the A/F (or enriching the mixture), and the release of oxygen under rich conditions, thereby enleaning the mixture, is an important feature of modern catalytic converters for vehicle applications. The goal of the oxygen storage model is to capture this property in a concise and sufficiently accurate manner.

Let $0 \leq \Theta \leq 1$ be the fraction of oxygen sites occupied in the catalyst, alternately denoted ROL (which stands for relative oxygen level). The oxygen storage capacity is modeled as a limited integrator in the following way:

$$\dot{\Theta} = \begin{cases} \frac{1}{C(MAF)} \times \rho(\lambda_{FG}, \Theta) \times 0.21 \times MAF \times \left(1 - \frac{1}{\lambda_{FG}}\right) & 0 \leq \Theta \leq 1 \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where

- $\dot{\Theta}$ represents $\frac{d\Theta}{dt}$;
- MAF denotes the mass air flow rate, used to approximate the flow rate of the mixture entering the TWC;
- C represents the effective catalyst "capacity," or the volume of active sites for oxygen storage, expressed in terms of the mass of oxygen that can be stored in the catalyst;
- ρ describes the exchange of oxygen between the exhaust gas and the catalyst;
- and λ denotes the relative air-fuel ratio, or equivalence ratio, with stoichiometry at $\lambda = 1$ (the subscript FG refers to the feedgas, or pre-catalyst measurement).

The effective TWC volume parameter, C, is expressed as a function of MAF in order to account for an observed increase in effective volume at high flow rates. The breakthrough times at medium flow and high flow rates are the same, but the breakthrough times at low flow rates are longer. This would seem to indicate that the effective volume increases as flow rate increases past a certain point. This effect is discussed in detail in [1, 2]. One possible physical explanation is that at higher flow rates, the air-fuel mixture is able to deposit or extract oxygen from more of the catalytic material than is accessible at the lower flow rates. This can be handled mathematically by setting C to a constant at low flow, then past a specific threshold (found to be 80 lb_m/hr), C would be increased by a quantity sufficient to maintain the same breakthrough time. In order to simplify the model structure, this effect is combined with the MAF term by applying a saturation function to the MAF input value. A function $f_{sat}(MAF)$ is created that is exactly equal to MAF below 80 lb_m/hr but maintains the value of 80 lb_m/hr when $MAF > 80$ lb_m/hr.

The function ρ is modeled as

$$\rho(\lambda_{FG}, \Theta) = \begin{cases} \alpha_L f_L(\Theta) & \lambda_{FG} > 1 \\ \alpha_R f_R(\Theta) & \lambda_{FG} < 1 \end{cases}, \qquad (2)$$

with $0 \leq f_L \leq 1$ representing the fraction of oxygen (combined or free) from the feedgas sticking to a site in the catalyst, and $0 \leq f_R \leq 1$ representing the fraction of oxygen being released from the catalyst and recombining with the feedgas. In (2), f_L and f_R vary with the percentage of occupied oxygen sites. In the model, f_L is assumed to be monotonically decreasing, with value one at $\Theta = 0$ and zero at $\Theta = 1$, and f_R is assumed to be monotonically increasing, with value zero at $\Theta = 0$ and one at $\Theta = 1$. The parameters α_L and α_R are included to represent the fact that the catalyst's storage and release rates of oxygen are different, with the release rate normally being higher than the storage rate. The difference in breakthrough times can be handled by allowing C to take different values on different sides of stoichiometry. The differences in adsorption and desorption rates are due to the difference between the processes by which oxygen is adsorbed and released [4].

The specific equations that have been selected for f_L and f_R, are as follows:

$$f_L(\Theta) = \frac{1 - e^{6\Theta}}{e^6 - 1} + 1 \qquad (3)$$

$$f_R(\Theta) = \frac{e^{-9\Theta} - 1}{e^{-9} - 1} \qquad (4)$$

The quantity $0.21 \times MAF \times (1 - \frac{1}{\lambda_{FG}})$, which can be rearranged to $0.21 \times MAF \times \frac{\Delta\lambda_{FG}}{\lambda_{FG}}$, represents the differential mass flow rate of oxygen (combined or free) in the feedgas with respect to stoichiometry. When multiplied by ρ, it gives the mass flow rate of oxygen that is deposited in (or released from) the catalyst. After some manipulations, one arrives at the following equation for direct computation of tailpipe A/F:

$$\lambda_{TP} = \lambda_{FG} - \rho(\lambda_{FG}, \Theta) \times (\lambda_{FG} - 1). \qquad (5)$$

The A/F of the exhaust feedgas (pre-catalyst) is a well-defined quantity since mass is conserved during the combustion process. The notion of the "A/F" of the tailpipe exhaust (post-catalyst) is less clear because mass is not instantaneously conserved through the TWC; indeed, oxygen is stored and released in the catalyst. By A/F for a given volume of exhaust gas at the tailpipe is meant the mass ratio of oxygen to hydrogen and carbon, whether free or combined, divided by 0.21. When applied to the feedgas for standard gasoline, this yields the standard measurement. If the fuel is oxygenated (reformulated gasoline), a one- to two-percent correction to this would need to be added.

For this oxygen storage model, the effects of feedgas and catalyst temperature are not included. The block diagram representation of the oxygen storage submodel is shown as Fig. 1.

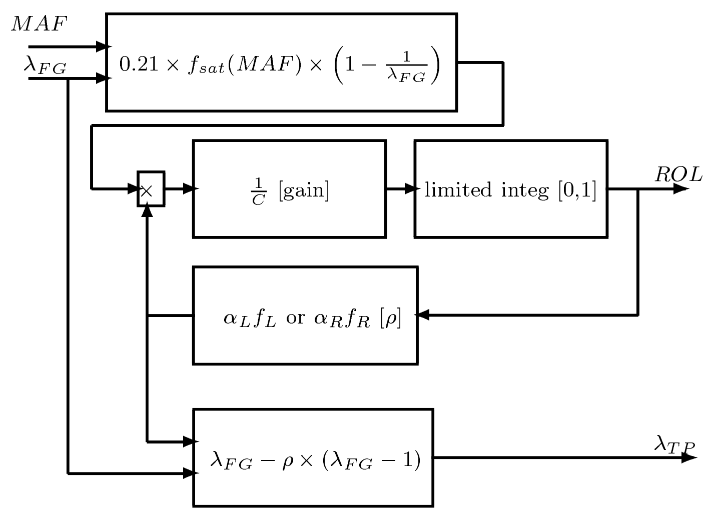

Figure 1: Structure of the oxygen storage submodel.

3 Basics of hypothesis testing

In a typical hypothesis testing problem, a decision has to be made regarding the source of an observation or a set of observations. A hypothesis can be thought of as a statement about the source of the observations. In this paper, the problem will be limited to a decision between two possible sources of observations, known as a *binary* hypothesis test [9]. In this case, the two hypotheses will be represented as H_0 and H_1. H_0 is assumed to be the original assertion, or the *null hypothesis*. H_1 is called the *alternate hypothesis*. In general, H_0 is assumed to be true until sufficient statistical evidence is gathered to reject H_0, correspondingly accepting H_1, given the observations.

A common use of hypothesis testing is in radar detection. H_0 corresponds to the assertion that a target is not present, and H_1 corresponds to the assertion that a target is present. In this paper, H_0 represents the assertion that the TWC is functioning properly, and H_1 represents the assertion that the TWC is not functioning properly. Just as a target is assumed to be absent until evidence "proves" one to be present (with a certain degree of statistical confidence), the TWC is assumed to be properly functioning until statistical evidence shows a deficiency in the TWC's operation.

Suppose that corresponding to each hypothesis there is an *observation*, a random variable, Z, that is generated according to some probabilistic law. The hypothesis testing problem is one of deciding which hypothesis is the correct one, based on a single measurement, z, of this random variable. The range of values that z takes constitutes the *observation space*. The binary decision problem essentially consists of partitioning this one-dimensional space into two regions, R_0 and R_1, such that whenever z lies in R_0, it is decided that H_0 was the correct hypothesis, and whenever z lies in R_1, it is decided that H_1 was the correct hypothesis. The regions R_0 and R_1 are known as *decision regions*. Whenever a decision does not match the true hypothesis, it is said that an *error* has occurred. The problem lies in choosing the decision regions such that the fewest errors are obtained with several realizations of Z. [9]

More generally, one may have a set of observations, $\mathbf{z} = (z_1, z_2, ..., z_n)$. In this case, the observation space is n-dimensional. The binary hypothesis testing problem is essentially the same as with a single measurement. However, this paper will highlight different methods for handling the cases of single and multiple observations.

In the decision-making process, there are two possible errors that can be made. If the alternate hypothesis is chosen when the null hypothesis is true, this is called a *false alarm*. The opposite type of error is called a *miss*. The typical measurements of the fidelity of a hypothesis test are P_F, the probability of false alarm, and P_D, the probability of detection (correctly choosing the alternate hypothesis). Ideally, P_F should be kept as small as possible, with P_D as large as possible. Realistically, a choice is made for P_F, and P_D is maximized for this P_F. The receiver operating characteristic (ROC) is a plot of P_D vs. P_F, which graphically shows the relationship between these two quantities.

4 Application to TWC diagnostics

In order to formulate a hypothesis testing problem, a specific hypothesis test is required. In the most general problem statement, it is desired that the null hypothesis be that the TWC is functioning properly and the alternate hypothesis be that the TWC is malfunctioning. The idea is then to ascertain a degree of confidence in the assertion of a TWC failure and, based on that, decide whether to declare a failure of the TWC.

The work in this paper focuses on the TWC failure mode of depleted oxygen storage capacity. In reality, deterioration of a TWC would show up in more than just the oxygen storage capacity. For example, it is very likely that TWC deterioration would show up in the steady state conversion efficiency. However, for the purposes of this study, it is assumed that a depletion in oxygen storage would accompany or precede any other TWC failure modes, and thus it is sufficient to detect this failure mode.

Since the effective TWC volume cannot be measured directly, another quantity will need to be measured, one that adequately represents the effective TWC volume. Assume that feedgas and tailpipe air-fuel ratios can be measured with linear exhaust gas oxygen sensors. One can then pursue a measurement of the time required to empty or fill the oxygen storage capacity of the TWC, using input test signals such as those in Figs. 2 and 3. Since test data was not available for "marginal" TWC's (catalysts with degraded oxygen storage capacity), for all analysis performed here, data was generated with the SIMULINK model of Section 2, for a given effective TWC volume.

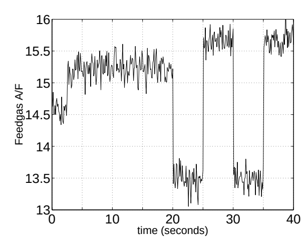

Figure 2: Input square wave signal.

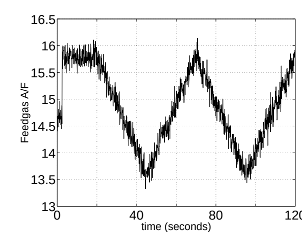

Figure 3: Input triangle wave signal.

Figures 2 and 3 show the input signals used for the square and triangle wave signals, respectively. Both of them use an initial lean period to set the initial condition of Θ, the relative oxygen level of the TWC. The modulation is begun after the oxygen storage capacity is filled by the initial enleanment. In each case, a nominal amount of noise is added to the air-fuel ratio signal to represent realistically observed air-fuel ratios; specifically zero-mean, white, Gaussian noise with a standard deviation of 0.01 λ (0.145 A/F) is assumed.

For a sufficiently slow square wave input signal, the empty/fill time could be represented as the rise or fall time of the tailpipe A/F response. Figure 4 shows the 10-90% rise times for square wave tests at different effective TWC volume points. The plot is based

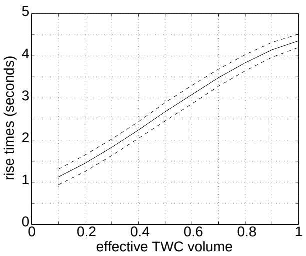

Figure 4: Square wave response rise times vs. effective TWC volume.

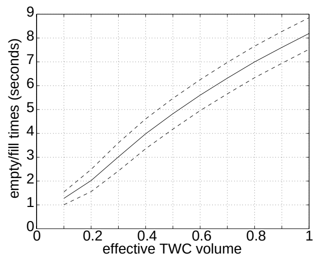

Figure 5: Triangle wave empty/fill times vs. effective TWC volume.

on 1000 simulations of the TWC model, with the noisy square wave of Fig. 2 at the input and a similar zero-mean, white, Gaussian noise with standard deviation of 0.01 λ (0.145 A/F) added to the output (i.e., the computed tailpipe air-fuel ratio). The solid line represents the mean of the rise times at each TWC volume point, while the dashed lines represent one standard deviation above and below the mean. It can be seen that the relationship between square wave response rise times and effective TWC volume is fairly linear, certainly reasonable enough to establish a hypothesis test based on these results.

A square wave test might be deemed to be too drastic, because it would require sudden changes in A/F, and hence in torque. Thus, a triangle wave test, such as the one plotted as Fig. 3, was also pursued as a less intrusive alternative. Since the rise time is not a clearly distinguishable measurement for a triangle wave response, a slightly different approach was pursued. The feedgas and tailpipe estimates were subtracted from each other, and the duration of the departure from zero of this difference was measured. Figure 5 shows this approach to be reasonable, highlighting a fairly linear relationship with the effective volume. The issues of the wider standard deviation shown will be explored later in this paper as part of the trade-off with the square wave test.

Since the empty/fill time measurements have a direct correspondence with the effective TWC volume, the evaluation of the hypotheses is based on the observations of the time required to empty or fill the oxygen storage capacity of the TWC. In the next section, the Neyman-Pearson criterion is outlined. It is then used to evaluate the usefulness of the hypothesis test, under a variety of conditions, on the basis of a single sample. After this, Student's t test is introduced and used to improve the fidelity of the hypothesis test through the use of multiple samples.

5 Neyman-Pearson criterion

In some hypothesis testing problems, there are specific costs associated with each decision, most notably with misses and false alarms. This can lead to an optimization problem in which a threshold of detection is chosen that minimizes the total cost over the observation space. If it is known, prior probability information is used to set the threshold of detection. However, there are many problems for which *a priori* probability information is not available and the costs of misses and false alarms are not clearly known. In addition, it may be desired to have more direct control over the false alarm probability. This is the case with the radar detection problem as well as the TWC failure detection problem. For these cases, an alternate method is used.

The *Neyman-Pearson criterion* provides the framework for the hypothesis test used here. It is based on a binary test of simple hypotheses, but its results can be extended to composite hypothesis tests, as will be shown in the next section. The Neyman-Pearson criterion employs a constrained optimization framework which seeks to maximize P_D subject to a chosen P_F. This section will outline the basic idea of the application of the Neyman-Pearson criterion to the TWC failure detection problem. A more detailed discussion of this and other methods for hypothesis testing can be found in [9] or any other textbook on detection theory.

In this initial application, the observed signal (the empty/fill time) is assumed[1] to be distributed as a normal random variable with *unknown* mean m and *known* variance σ^2. The mean depends on the TWC effective volume, as shown in Figures 4 and 5. If the variance is not assumed known, then multiple samples are required, and Student's t test can be used for the hypothesis testing problem, as will be discussed in Section 6. The null hypothesis, denoted H_0, is that the TWC is functioning properly, meaning that its effective volume is greater than a specified quantity. The alternate hypothesis, H_1, is that the TWC is malfunctioning, meaning that its effective volume is smaller than what is required to meet emissions regulations.

The hypothesis testing problem can be stated as follows:

$$H_0 : m \leq m_0 \quad \text{or} \quad m = m_0 \quad \text{or} \quad Z = m_0 + V \qquad (6)$$

$$H_1 : m > m_0 \quad \text{or} \quad m = m_1 \quad \text{or} \quad Z = m_1 + V \qquad (7)$$

where $-m$ is the actual (uncorrupted measurement of the) empty/fill time, $-m_0$ is the "bad" threshold for the empty/fill time, $-m_1$ is the empty/fill time of an alternate TWC for which P_D will be determined, V is a random variable normally distributed with zero mean and variance σ^2, and Z is the random variable representing the fill/empty measurement corrupted by V, and thus Z is normally distributed with mean m and

[1] In order to simplify the analysis, specifically for construction of detection probabilities, it is desirable to assume that the empty/fill times for each effective TWC volume point are reasonably approximated by Gaussian random variables. An analysis of this can be found in [1], where is it concluded that the Gaussian approximation, while imperfect, is reasonable enough for the work in this paper.

variance σ^2. The reason for the "$-$" sign is to state the problem in such a way that $m_1 > m_0$, keeping it in line with standard statements of hypothesis testing problems. A short empty/fill time would be the failure condition, so use of the additive inverse keeps the inequality in the "standard" direction.

In order to completely apply the Neyman-Pearson criterion, both H_0 and H_1 must be *simple* hypotheses (equalities). If H_1 is composite (inequalities), the Neyman-Pearson criterion can still be applied with a specified P_F constraint, but it becomes impossible to calculate a corresponding P_D, unless a choice of m_1 is made. The choice of m_1 does not affect the choice of the detection threshold, since that calculation is based on the choices of m_0 and P_F, independent of m_1. The application of the Neyman-Pearson criterion provides a maximum P_D for any choice of $m_1 > m_0$.

The following is a step-by-step procedure that can be used to obtain probabilities of failure declaration for any given TWC effective volume. A derivation of these equations can be found in [1] or any standard textbook on statistics. For calculations in MATLAB, the integrals are converted to be in terms of the error function (erf), defined as follows:

$$\operatorname{erf}(x) = \int_0^x \frac{2}{\sqrt{\pi}} e^{-\frac{t^2}{2}}\, dt. \tag{8}$$

Step 1: Choose m_0 (based on the TWC volume threshold) and P_F (the allowable "false alarm" probability). This P_F is the probability that a TWC with the given empty/fill time, and hence, effective volume, will be declared bad.

Step 2: Use these choices to calculate γ, the threshold of detection, as follows:

$$P_F = \int_\gamma^\infty \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(z-m_0)^2}{2\sigma^2}}\, dz \tag{9}$$

$$= \frac{1}{2} - \frac{1}{2}\operatorname{erf}\left(\frac{\gamma - m_0}{\sqrt{2}\sigma}\right). \tag{10}$$

$$\gamma = m_0 + \sqrt{2}\sigma \operatorname{erf}^{-1}(1 - 2P_F). \tag{11}$$

Threshold: (z is the observed empty/fill time)

$$z \underset{H_0}{\overset{H_1}{\gtrless}} \gamma \tag{12}$$

Step 3: Use γ to calculate P_D for a choice of $m_1 > m_0$:

$$P_D = \int_\gamma^\infty \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(z-m_1)^2}{2\sigma^2}}\, dz \tag{13}$$

$$= \frac{1}{2} - \frac{1}{2}\operatorname{erf}\left(\frac{\gamma - m_1}{\sqrt{2}\sigma}\right). \tag{14}$$

In (14), if m_1 is chosen less than m_0, then P_D is no longer the *detection probability*, but rather the probability that a TWC with an empty/fill time of $-m_1$ will be declared bad. In fact, for the TWC application, for any value of m_1, P_D then becomes the probability that a TWC with empty/fill time $-m_1$ is *discarded* (i.e., declared bad) on the basis of the hypothesis test, (7).

The empty/fill times for the square and triangle waves are measured as described in Section 4. The sample means and standard deviations are calculated from the 1000 simulations that were performed at each data point. Based on these numbers and the assumption of normality, probabilities of failure declaration are tabulated. The threshold of failure declaration is then calculated based on the chosen probability of detection/false alarm, and the other probabilities are calculated based on this detection threshold. A graphical representation of this probability spectrum is known as the receiver operating characteristic (ROC). Two sample ROC plots can be found as Figs. 6 and 7. In those plots, the volume threshold for a good catalyst is chosen to be 0.5 (i.e., only half of the oxygen storage sites are usable).

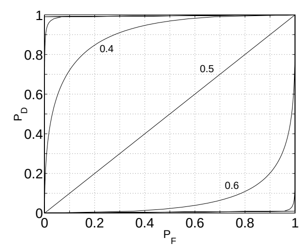

Figure 6: ROC for square wave rise times (centered at 0.5).

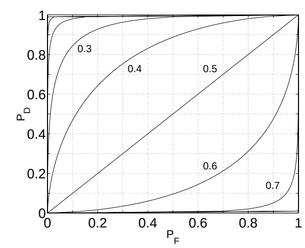

Figure 7: ROC for triangle wave empty/fill times (centered at 0.5).

The graphs show the results of the hypothesis tests for a range of true catalyst volume from 0.1 to 1.0, in increments of 0.1. The curves closest to 0.5 are labeled with their corresponding volumes. It is observed in Fig. 6 that if m_0 corresponds to a volume of 0.5 and a P_F of 0.3 is tolerated (i.e. it is acceptable to discard 30% of the TWCs with oxygen storage volume of 0.5), then the P_D for a 40% TWC is greater than 0.9 (i.e. greater than 90% of the 40% TWCs will be discarded), while only about 1% of the 60% TWCs will be declared bad. On the other hand, if one wishes to be more aggressive in discarding questionable TWCs, it may be deemed acceptable to set P_F to 0.9 (so one now discards 90% of the 50% TWCs). Again, from Fig. 6, one can see that this will discard about 99% of the 40% TWCs, but it will also discard 20% of the 60% TWCs. A way to improve upon these results it to use multiple samples; this is pursued next.

6 Student's t test

Around the turn of the century, W. S. Gossett (publishing under the pseudonym "Student") obtained results for the distribution of sampled Gaussian data. An important test statistic, t, is calculated on the basis of N independent observations with sample mean $\bar{x}$ and variance s^2 as follows:

$$t = \frac{(\bar{x} - m_0)\sqrt{N}}{s} \tag{15}$$

where

$$\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x_i \tag{16}$$

$$s^2 = \frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x})^2 \tag{17}$$

One such hypothesis test for which this is useful is the following:

$$H_0 : \theta = m_0 \tag{18}$$

$$H_1 : \theta > m_0 \tag{19}$$

If H_0 is true, t has the Student-Fisher t distribution with $N-1$ degrees of freedom; and the null hypothesis is rejected for a test of size α if $t > t_{\alpha;N-1}$, where $t_{\alpha;N-1}$ is the upper 100α percentage point of the t distribution. These values are typically available in tables, and they can also be found through use of a statistical software package, such as the Statistics Toolbox in MATLAB.

The shape of the Student-Fisher t distribution is independent of m_0 and σ. It is dependent only on the so-called degrees of freedom, the number of "free" observations. In this example, since one parameter is being estimated, the number of degrees of freedom, ν, is one less than the number of samples, N, so $\nu = N - 1$. The difference between the t distribution and the standard normal distribution tends to zero as ν increases.

This test also follows the Neyman-Pearson criterion, but its use of information from multiple samples allows for improvement in test fidelity, even without *a priori* knowledge of σ. The t tests, like the single sample Neyman-Pearson tests of Section 5, provide a statistical measure of the likelihood of the alternate hypothesis. The statistical significance is the probability that the measured event would occur by chance under the null hypothesis. This significance is calculated from the Student-Fisher t distribution, and it is analogous to the chosen P_F and calculated P_D from the Neyman-Pearson criterion.

The statistical significance provides the basic decision mechanism for rejection of the null hypothesis by providing a measurement of confidence in the null hypothesis. For example, if the criterion is to be 95% sure that the null hypothesis is false before rejecting it and accepting the alternate hypothesis, then the alternate hypothesis is accepted when the significance of the null hypothesis is less than 0.05. Similarly, a 99% confidence requirement would mean that the significance of the null hypothesis would have to be less than 0.01 for the alternate hypothesis to be accepted.

Another possible idea is to use a 50% confidence interval. This might be useful in cases where it is desired to reject half of the items that are close to the threshold, for example when the null and alternate hypotheses are inequalities on opposite sides of the threshold and are assumed to have equal probability of being true. In this application, such a threshold would declare a failure in half of the TWC's with an effective volume of 0.5 (or any other specified borderline volume). The detection probabilities at other TWC volumes can be calculated based on this, increasing the probability of detection at smaller TWC volumes and decreasing the false alarm probability at larger TWC volumes.

Figures 8 and 9 show the probabilities of failure declaration for varying numbers of samples for TWC effective volumes of 0.4, 0.5, and 0.6. These plots are centered at 0.5, indicating that the threshold of detection is set based on the statistical significance of the TWC performance at half of its effective volume. They are based on the cumulative distribution function (CDF) of this significance, which, for this data, is also the ROC. Figures 8 and 9 show the results for the square and triangle wave test signals of Figs. 2 and 3, respectively. All of these plots are created from t test results of 3, 5, and 10 samples. These can be compared to the Neyman-Pearson ROC's of Figs. 6 and 7, to see the level of improvement in detection fidelity that is achieved through the use of the t test.

Figure 10 is based on the same square wave input signal as Fig. 8; but for this plot, the failure threshold is 0.9, indicating a greater constraint on TWC performance. Similarly, Fig. 11 shows the detection probabilities based on the triangle wave test signal. Figures 10–11 show results for tests based on 1, 3, 5, and 10 samples. The single sample test is a Neyman-Pearson test based on Section 5, while the multiple sample tests are t tests.

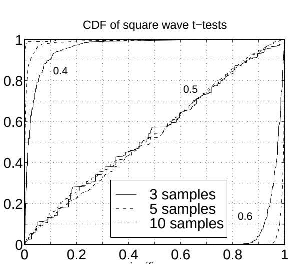

Figure 8: Approximate CDF/ROC for square wave rise time significances.

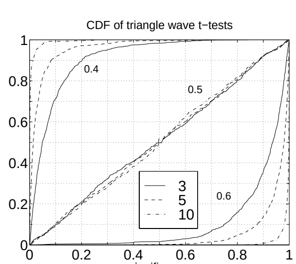

Figure 9: Approximate CDF/ROC for triangle wave empty/fill time significances.

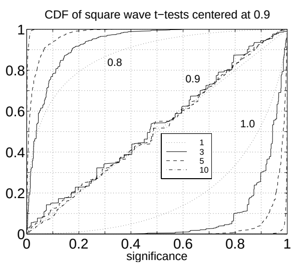

Figure 10: Approximate CDF/ROC for square wave rise time significances, centered at 0.9.

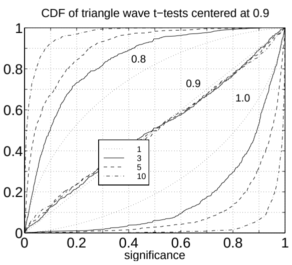

Figure 11: Approximate CDF/ROC for triangle wave empty/fill time significances, centered at 0.9.

In all of these figures, the increased fidelity from the use of multiple samples is highlighted. Larger sample sizes bring the ROC curves closer to their ideal of a perfect (unity) probability of detection for any TWC effective volume less than the threshold and a zero probability of detection (false alarm) for any TWC volume greater than the threshold.

7 Conclusion

The TWC diagnostics problem was approached from the perspective of a hypothesis test, based on the time required to empty or fill the catalyst's oxygen storage capacity. All of the figures in this paper are based on simulated data. Unfortunately, marginal TWC's, ones which actually have depleted oxygen storage capacity, were not available for testing. This prevents hardware validation of the results in this paper. However, the model's demonstrated performance against actual data [2, 3] and the physical relevance of the model parameters suggests that this approach is still a reasonable one. The work in this paper can be performed using any desired TWC model, as long as the model includes an adjustable parameter for effective storage capacity and A/F as an output.

When only a single data sample was collected, *a priori* knowledge of variance was needed. It was then demonstrated that the additional information from multiple samples could be used in the framework of Student's t test to improve the fidelity of the test as the number of available samples increases. In addition, Student's t test does not require *a priori* knowledge of the variance.

Another benefit of the hypothesis test presented here is that it can be applied directly with a switching HEGO sensor. In that case, one could simply measure the time between feedgas and tailpipe switching in a square wave test. This would correspond roughly to a 0-50% rise time instead of a 10-90% rise time, but it would still lend itself to a hypothesis test through use of the methodology presented in Sections 5 and 6.

Finally, in [1], the robustness of the hypothesis testing problem to the addition of an unknown constant bias (a normally distributed random number, again with mean zero and standard deviation 0.01 λ) to the input A/F signal (square or triangular) was studied. This was done to account for the possibility of a bias in the feedgas A/F sensor. This biased A/F signal was run through the TWC model, but the distribution functions and diagnostic information were generated based on the assumption of an unbiased input signal. The resulting degradation in performance was quantified and shown to be "moderate."

References

[1] E. P. Brandt. *Modeling and Diagnostics of Three-Way Catalysts for Advanced Emission Control Systems.* PhD thesis, University of Michigan, August 1998.

[2] E. P. Brandt, Y. Wang, and J. W. Grizzle. Dynamic modeling of a three-way catalyst for SI engine exhaust emission control. *IEEE Transactions on Control Systems Technology*, 8(5):767–776, September 2000.

[3] E. P. Brandt, Yanying Wang, and J. W. Grizzle. A simplified three-way catalyst model for use in on-board SI engine control and diagnostics. In G. Rizzoni, editor, *Proceedings of the ASME Dynamic Systems and Control Division*, volume 61, pages 653–659. American Society of Mechanical Engineers, November 1997. 1997 ASME International Congress and Exposition, Sixth ASME Symposium on Advanced Automotive Technologies.

[4] Alan Brenner. Principles and practice of heterogeneous catalysis. Course notes, 1997.

[5] Bradley Jones. *Statistics Toolbox User's Guide.* The MathWorks, Inc., Natick, MA, March 1996.

[6] Donald F. Morrison. *Multivariate Statistical Methods.* McGraw-Hill, second edition, 1976.

[7] Lothar Sachs. *Applied Statistics: A Handbook of Techniques.* Springer-Verlag, 1982.

[8] E. Shafai, C. Roduner, and H. P. Geering. Indirect adaptive control of a three-way catalyst. SAE paper 961038, 1996 SAE International Congress, Detroit, MI, February 1996. included in *SAE Special Publication SP-1149, Electronic Engine Controls.*

[9] M. D. Srinath, P. K. Rajasekaran, and R. Viswanathan. *Introduction to Statistical Signal Processing with Applications.* Prentice Hall, Englewoord Cliffs, New Jersey, 1996.

[10] Robert G. D. Steel, James H. Torrie, and David A. Dickey. *Principles and Procedures of Statistics: A Biometrical Approach.* McGraw-Hill, Inc., third edition, 1997.