

On-board diagnosis of car catalytic converters using statistical pattern recognition

A. Boatas */** and B. Dubuisson * and M.A. Dillies-Peltier **

* Laboratoire HEUDIASYC, UMR CNRS 6599, UTC

BP 20529, F-60205 Compiègne cedex

e-mail : aboatas@hds.utc.fr, bdubuisson@hds.utc.fr

** PSA Peugeot Citroën, DINQ/DRIA/RPAI/CSH

2 Route de Gisy, F-78943 Vélizy-Villacoublay cedex

Abstract

This paper introduces statistical pattern recognition techniques, applied to the monitoring of dynamic systems. Usually, distance rejection options enable to deal with incomplete knowledge about classes. A new technique, which extends the possibilities of distance rejection, is presented in order to detect partially unknown classes. These techniques have been applied in this paper to a very important legislative problem : the monitoring of car catalytic converters.

Introduction

Pattern recognition aims at classifying patterns. It can be easily applied to the monitoring of dynamic systems where the goal is to detect and identify the current operating mode. When a learning set representing all the classes is available, the classification problem is rather simple. It can be achieved using statistical tests (Fukunaga 1990).

Some new operating modes, which cannot be learnt off-line, could appear on-line. They are usually detected using distance rejection options (Dubuisson 1990), (Dubuisson & Masson 1993). Sometimes, some classes are partially unknown : no training set is available but the location in the feature space is approximately known. However, ordinary distance rejection does not enable to distinguish patterns of this kind of classes. Therefore, we propose local distance rejection in order to be able to discriminate patterns into such classes.

Firstly, one of the most usual decision rule, Bayes' test, will be quickly reminded to the reader. Then, our local distance rejection technique will be presented. This study has been initiated as part of our researchs on On-Board Diagnosis (OBD) of car emission critical components. In this paper, we will concentrate on the real-time monitoring of car catalytic converters. Finally, on-board diagnosis of car catalytic converters will be presented as an example of statistical pattern recognition with incomplete learning set.

Copyright ©1999, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Diagnosis of dynamic systems using pattern recognition

Bayes' decision rule

Let us suppose we have to decide between M classes $\omega_1, \omega_2, \dots, \omega_M$ in a d -dimensional feature space. The patterns \underline{x} (the observable measurements $\underline{x} \in \mathbb{R}^d$) can be interpreted as a realization of a d -dimensional random vector \underline{X} . The decision function $d(\underline{x})$ is set to i if \underline{x} is classified into the class ω_i . Let us define a cost function $C(\underline{x})$:

$$C(\underline{x}) = \sum_{i=1}^M C_{\omega_{d(\underline{x})}/\omega_i} P(\omega_i|\underline{x}) \quad (1)$$

where C_{ω_j/ω_i} is the cost of deciding to classify \underline{x} into ω_j when \underline{x} belongs in fact to ω_i ($i = 1, \dots, M; j = 1, \dots, M$) and $P(\omega_i|\underline{x})$ is the posterior probability. $P(\omega_i|\underline{x})$ can be computed using the bayes' rule

$$P(\omega_i|\underline{x}) = \frac{f(\underline{x}|\omega_i)P(\omega_i)}{\sum_{j=1}^M f(\underline{x}|\omega_j)P(\omega_j)} \quad (2)$$

The mean cost C is the expectation of $C(\underline{x})$:

$$\int_{\mathbb{R}^d} C(\underline{x})f(\underline{x})d\underline{x}, \quad (3)$$

where $f(\underline{x}) = \sum_{i=1}^M f(\underline{x}|\omega_i)P(\omega_i)$ is the mixture density. If no decision is more important than another one, one chooses the $\{0, 1\}$ costs

$$C_{\omega_i/\omega_i} = 0 \quad i = 1, \dots, M \quad (4)$$

$$C_{\omega_i/\omega_j} = 1 \quad i, j = 1, \dots, M \quad i \neq j \quad (5)$$

Then, $C(\underline{x})$ becomes

$$C(\underline{x}) = \sum_{i=1, d(\underline{x}) \neq i}^M P(\omega_i|\underline{x}) = 1 - P(\omega_{d(\underline{x})}|\underline{x}) \quad (6)$$

The Bayes' decision rule minimizes the mean cost C (C^* is the minimum of C and is called the Bayes' cost or the Bayes' risk) of deciding to classify each pattern \underline{x} into a class ω_i (Fukunaga 1990), (Dubuisson 1990).

C reaches its minimum C^* when $C(\underline{x})$ is also minimum for every feature vector \underline{x} :

$$C^*(\underline{x}) = \min_{d(\underline{x})} C(\underline{x}) = 1 - \max_{j=1,\dots,M} P(\omega_j|\underline{x}). \quad (7)$$

Therefore, the optimum decision function is :

$$d^*(\underline{x}) = i \quad \text{if} \quad P(\omega_i|\underline{x}) = \max_{j=1,\dots,M} P(\omega_j|\underline{x}). \quad (8)$$

For instance, for $M=2$ and $d=1$, figure 1 shows the gaussian classes ω_1 and ω_2 decision area Ω_1 and Ω_2 . Bayes' decision rule is optimum if $f(\underline{x}|\omega_1)$, $f(\underline{x}|\omega_2)$, $P(\omega_1)$ and $P(\omega_2)$ are known.




Figure 1: Ω_1 and Ω_2 decision area for $M=2$, $d=1$

In practice, the probability distributions are unknown, but are estimated by generative (parametric or non-parametric density function estimation) or discriminative models (Smyth 1994a).

The diagnosis technique of a dynamic system presented here is local : the decision relies on only one pattern presented at an instant t . Most of the time, it is valuable, when the observation is sequential, to use patterns sequences for diagnosis purpose : more and more information become available as time proceeds. It is particularly true when the dynamic system operating modes evolve slowly, comparatively to the features sampling period. Sequential diagnosis can be achieved using, for instance, sequential statistical tests (Fukunaga 1990) or hidden Markov models (Smyth 1994b), (Smyth 1994a).

Rejection options

Ambiguity and distance rejection options were already detailed in (Dubuisson 1990), (Dubuisson & Masson 1993), (Fukunaga 1990) and (Denoeux, Masson, & Dubuisson 1997). This section will sum up ordinary techniques and will introduce a way of defining distance rejection in only one area of the feature space.

Ambiguity rejection Ambiguity rejection, as defined by Chow (Chow 1957), consists in taking no decision if the cost of deciding $C^*(x)$ is bigger than a reject cost C_r which is an arbitrarily set threshold. When ambiguity rejection is decided, the pattern \underline{x} is classified

into a class ω_0 ($d(\underline{x}) = 0$). Bayes' decision rule becomes

$$d(\underline{x}) = i \quad \text{if} \quad \begin{cases} P(\omega_i|\underline{x}) = \max_{j=1,\dots,M} P(\omega_j|\underline{x}) \\ P(\omega_i|\underline{x}) \geq 1 - C_r \end{cases} \quad (9)$$

$$d(\underline{x}) = 0 \quad \text{if} \quad 1 - C_r \geq \max_{j=1,\dots,M} P(\omega_j|\underline{x}). \quad (10)$$

In order to enable ambiguity rejection, because $\max_{j=1,\dots,M} P(\omega_j|\underline{x})$ belongs to the interval $[1/M; 1]$, C_r must satisfy the constraint :

$$0 \leq C_r \leq 1 - \frac{1}{M}. \quad (11)$$

It means that if the reject cost C_r is too high, there is no more ambiguity rejection because it costs too much.

Thus, when equation (11) is true, the ambiguity rejection rule introduces a new class ω_0 between the other classes $\omega_1, \omega_2, \dots, \omega_M$.

For instance, for $M=2$ and $d=1$, figure 2 shows the ω_0 decision area Ω_0 .

Distance rejection Bayes' decision rule with ambiguity rejection does not enable to detect unknown classes and to deal with incomplete learning set. Distance rejection (Ω_D) occurs when the mixture density $f(\underline{x})$ is lower than a threshold C_D .

$$\Omega_D = \{\underline{x} : f(\underline{x}) \leq C_D\} \quad (12)$$

For $M=2$ and $d=1$, figure 2 shows the rejected patterns area Ω_D .




Figure 2: Ω_0 and Ω_D rejection areas for $M=2$, $d=1$

To deal with an incomplete learning set, one can check (on-line or off-line) if the distance rejected patterns can be organized into classes, for instance using clustering algorithms. Thus new classes can be learnt, but it is no longer our business here.

Local distance rejection Distance rejection sets up boundaries around the classes. However, it makes no difference in which direction the rejection occurs. We propose to define an area in the feature space where we can detect distance rejected patterns. This area is defined as a new class with an uniform distribution. Here, only one area is defined as a local distance rejection class. However, more than one area can be defined, as long as they do not overlap each other. The

$\omega_1, \omega_2, \dots, \omega_M$ classes define a class $\omega_{\{1, \dots, M\}}$ for which the probability density can be written

$$f(\underline{x}|\omega_{\{1, \dots, M\}}) = \sum_{i=1}^M f(\underline{x}|\omega_i)P(\omega_i). \quad (13)$$

The local distance rejection area is defined as a new class ω_{M+1} . Its probability distribution is uniform over a volume V :

$$f(\underline{x}|\omega_{M+1}) = \begin{cases} 1/V & \text{if } \underline{x} \in V \\ 0 & \text{otherwise} \end{cases}. \quad (14)$$

Comparatively to the distance rejection option, local distance rejection area can be defined as follows :

$$\Omega_{M+1} = \{\underline{x} : f(\underline{x}|\omega_{1, \dots, M}) \leq C_D, \text{ and } f(\underline{x}|\omega_{M+1}) > 0\} \quad (15)$$

However, as ω_{M+1} can be interpreted as a uniform distribution, Ω_{M+1} can be written :

$$\Omega_{M+1} = \{\underline{x} : P(\omega_{M+1}|\underline{x}) \geq P(\omega_{1, \dots, M}|\underline{x})\}. \quad (16)$$

The posterior probabilities can be computed using the Bayes' rule. The boundary is set up by the prior probabilities $P(\omega_{\{1, \dots, M\}})$ and $P(\omega_{M+1})$. According to the first definition of Ω_{M+1} (see equation (15)):

$$P(\omega_{M+1}) = \frac{C_D}{1/V + C_D}, \quad (17)$$

$$P(\omega_{1, \dots, M}) = 1 - P(\omega_{M+1}). \quad (18)$$

It is still a rejection rule because there is no learning set for this class. Moreover, we have got here an insight into the spatial location of the distance rejection area. Therefore, the difference with the distance rejection rule is that we have defined a location where patterns are rejected.

For $M=2$ and $d=1$, figure 3 shows how an area Ω_3 , where patterns are locally distance rejected, can be defined.




Figure 3: Ω_3 local distance rejection area for $M=2, d=1$

This technique presents some limitations : a volume shape for the class ω_{M+1} has to be chosen in order to define its uniform distribution. It is easier to choose an analytical definition of the volume (like a hypercube). However, it becomes difficult to set the volume when the dimension of the feature space is high. Moreover, by choosing a uniform distribution, we

have no "fuzziness" for this class. Finally, by defining $f(\underline{x}|\omega_{M+1})$ as an uniform distribution, we have made a wrong assumption on the shape of the ω_{M+1} probability distribution. The best thing to do, when patterns are locally distance rejected, is to learn the new class probability distribution.

Monitoring of car catalytic converters

These techniques have been applied to the monitoring of three way car catalytic converters. This function is mandatory on all cars since 1994 in north America and will be mandatory after 2000 in Europe. We aim at showing how pattern recognition based methods can be implemented to solve this real life problem. The conversion efficiency of the catalytic converter is monitored by measuring the oxygen storage capacity of the catalyst using the dual oxygen sensors method (Koupal, Sabourin, & Clemmens 1991), (Clemmens, Sabourin, & Rao 1990), (Hepburn & Gandhi 1992), (Hepburn & others 1994), (Koltsakis & Stamatelos 1995).

The aim is to detect an aged catalyst which would lead to exceed an emission threshold. The learning set is composed of patterns acquired on a vehicle during a reference driving cycle (ECE+EUDC cycle in Europe, FTP cycle in north America), for two reference catalysts :

- a fresh catalyst,
- an aged catalyst which corresponds to the emission threshold.

This is a typical case of incomplete knowledge about classes : only two reference catalysts patterns are available. More catalysts would be needed to learn the entire aging process. Nevertheless, this would be too expensive because aged catalysts can only be obtained after long aging cycles on engine benches.




Figure 4: Principal Components Analysis of the learning set

Signals from the engine control unit and from the catalyst oxygen sensors are computed to build a 4-dimensional feature space. A Principal Components

Analysis of the learning set (see figure 4) shows that the good and the bad catalysts patterns are well separated.

We have to deal with the unknown intermediate aging states of catalysts. When a catalyst is aged but located between the two reference catalysts, it is still a good catalyst (see figure 5). We cannot estimate the probability distribution of these intermediate catalysts because it would be too expensive to obtain a complementary learning set.

Modeling of intermediate catalysts

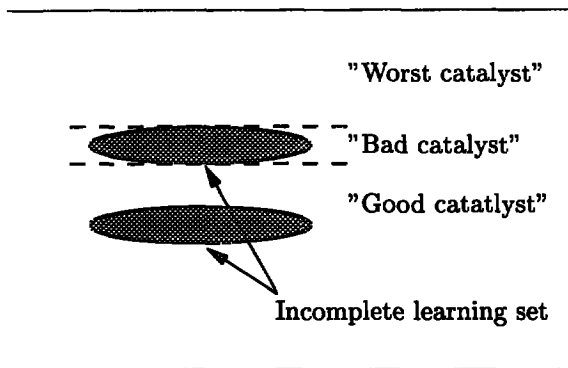


Figure 5: Extension of the learning sets

Therefore, it is necessary to extend the “good catalyst” class by setting the boundary just below the “bad catalyst” class. This could be achieved with the Bayes’ decision rule by tuning the prior probabilities. However, an intermediate catalyst is in fact unknown and has to be rejected. It is better to decide distance rejection than to classify it into the “good catalyst” class. Here distance rejection means that the catalyst is unknown, but makes no difference between a good or a bad one. Local distance rejection enables to define the area before the bad catalyst patterns as a “good catalysts” rejection area.

Modeling of severely aged catalysts

When the catalyst conversion efficiency is worse than this of the reference “bad catalyst”, it can be useful to detect it. As said in the previous section, ordinary distance rejection does not make any difference between a severely aged catalyst and an intermediate catalyst for which no learning set was available. Here again, local distance rejection can be used : it enables to define an area behind the “bad catalyst” class where the catalysts are assumed to be worse.

Algorithm implementation, results and discussion

Six classes have been defined :

- the “good catalyst” class (ω_1), which represents the fresh catalyst,

- the “bad catalyst” class (ω_2), which is the aged catalyst that leads to exceed the emission threshold,
- the “unknown good catalyst” class (ω_3) for which no training set is available and that is a priori good, represented by a local distance rejection area,
- the “worst catalyst” class (ω_4) defined by a local distance rejection area behind the “bad catalyst” class,
- the ambiguity rejection class (ω_0),
- the distance rejection class (ω_D).

The probability distribution of random patterns \underline{X} for classes ω_1 and ω_2 can be estimated from the learning set. At the moment, it has been done using Parzen window estimators. However, for better on-board performance, other estimators could be used.




Figure 6: Decision areas of catalyst classes with $C_r = 0.3$ and $C_D = 2.10^{-2}$

The figure 6 shows a 2-dimensional representation (first plan of a Principal Components Analysis) of the decision areas Ω_1 , Ω_2 , Ω_3 and Ω_4 . The ambiguity and distance rejection areas (Ω_0 and Ω_D) are also represented. In this example, the ambiguity reject cost is set to 0.3 which means that a pattern is classified into ω_1 or ω_2 if the corresponding posterior probability is greater than 0.7. If the density estimation $f(\underline{x}|\omega_{\{1,2\}})$ is smaller than 2.10^{-2} , the pattern is distance rejected :

- locally if the pattern is located into the Ω_3 or the Ω_4 volume,
- globally if the pattern is located outside the local distance rejection areas.

Here, designing a local distance rejection area is rather simple because the boundary is set by a straight class. It could be more difficult if the class shape was much more complex.

We have got some reference fresh catalyst and aged catalyst patterns. The error probabilities for these classes are well known (we have estimated the probability distribution) and depends of the chosen threshold. However, more tests are needed to measure the false alarm rates and to draw some final conclusions. For

instance, some road and endurance tests will be necessary to check the interest of the algorithm for a car manufacturer.

Therefore, we have showed how the incomplete training set problem can be bypassed. However, the diagnosis is still "local" from a temporal point of view. Because the aging states are stable during time (a catalyst aging can go on more than 80000 kilometers), it is sensible to take decisions from sequences of patterns. Several possibilities like sequential tests (Fukunaga 1990) or hidden Markov models (Smyth 1994b), (Smyth 1994a), will be tested on this practical issue in the near future.

Conclusion

This paper shows how a real life diagnostic application with incomplete knowledge about classes can be treated. Today, the in-car implementation of catalytic converters monitoring functions is usually achieved without pattern recognition. The advantages of the pattern recognition methods described here for on-board catalytic converters diagnosis are :

- to get optimum diagnosis,
- to achieve automatic calibration by using the learning skills of pattern recognition algorithms.

Finally the pattern recognition techniques presented here should enable, in this context :

1. to lower the development costs,
2. to enhance performance : lower false alarm rate and increased detection rate.

The second point is highly important, because it means that the legislation requirements are fulfilled (detection rate) as well as the customers' demands for high robustness.

The algorithm has been proposed in order to monitor usual three way converters by using the dual binary oxygen sensors. However, it can be adapted to other catalyst and sensor technologies.

These techniques could probably be efficient for other real examples, where the learning set cannot be completed, but where one can get an insight into the shape and the spatial location of the classes.

Acknowledgments

We wish to thank Mr Arnaud TURLAN, Jacques BERNARD and Pierre LIGOURET from PSA Peugeot Citroën for having provided the data necessary to this study.

References

Chow, C. 1957. An optimum recognition error and reject tradeoff. *IEEE Trans. Inform. Theory* 16:41-46.

Clemmens, W. B.; Sabourin, M. A.; and Rao, T. 1990. Detection of catalyst performance loss using on-board diagnostics. *SAE Technical Paper Series* (900062).

Denoeux, T.; Masson, M.; and Dubuisson, B. 1997. Advanced pattern recognition techniques for system monitoring and diagnosis : a survey. *Journal Européen des Systèmes Automatisés* 31(9-10):1509-1539.

Dubuisson, B., and Masson, M. 1993. A statistical decision rule with incomplete knowledge about classes. *Pattern Recognition* 26(1):155-165.

Dubuisson, B. 1990. *Diagnostic par reconnaissance des formes*. Paris: Hermès.

Fukunaga, K. 1990. *Introduction to statistical pattern recognition*. San Diego: Academic Press.

Hepburn, J. S., and Gandhi, H. S. 1992. The relationship between catalyst hydrocarbon conversion efficiency and oxygen storage capacity. *SAE Technical Paper Series* (920831).

Hepburn, J. S., et al. 1994. A review of the dual EGO sensor method for OBD II catalyst efficiency monitoring. *SAE Technical Paper Series* (942057).

Koltsakis, G. C., and Stamatelos, A. 1995. Detection of automotive catalysts failure by use of on-board diagnostics. *Proc Instn Mech Engrs* 209.

Koupal, J.; Sabourin, M. A.; and Clemmens, W. B. 1991. Detection of catalyst failure on-vehicle using the dual oxygen sensor method. *SAE Technical Paper Series* (910561).

Smyth, P. 1994a. Detecting novel fault conditions with hidden Markov models and neural networks. In *Pattern Recognition in Practice IV*. Elsevier Science B.V. 525-536.

Smyth, P. 1994b. Hidden Markov models for fault detection in dynamic systems. *Pattern Recognition* 27(1):149-164.