



HAL
open science

Adaptive mixture discriminant analysis for supervised learning with unobserved classes

Charles Bouveyron

► **To cite this version:**

Charles Bouveyron. Adaptive mixture discriminant analysis for supervised learning with unobserved classes. 2009. hal-00392297v1

HAL Id: hal-00392297

<https://hal.science/hal-00392297v1>

Preprint submitted on 6 Aug 2009 (v1), last revised 23 Jun 2010 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Supervised learning with unobserved classes

Charles BOUVEYRON

SAMOS-MATISSE, CES, UMR CNRS 8174
Université Paris 1 (Panthéon-Sorbonne), Paris, France

Abstract

In supervised learning, an important issue usually not taken into account by classical methods is the possibility of having in the test set individuals belonging to a class which has not been observed during the learning phase. Classical supervised algorithms will automatically label such observations as belonging to one of the known classes in the training set and will not be able to detect new classes. This work introduces a model-based discriminant analysis method, called adaptive model-based discriminant analysis (AMDA), which is able to detect unobserved groups of points and to adapt the learned classifier to the new situation. Two EM-based procedures are proposed for the parameter estimation. Experiments on artificial and real data demonstrate the ability of the proposed method to deal with complex and real word problems. The proposed method is also applied to the detection of novel species in DNA barcoding.

1 Introduction

The usual framework of supervised classification assumes that all existing classes in the data have been observed during the learning phase and does not take into account the possibility of having in the test set individuals belonging to a class which has not been observed. In particular, such a situation could occur in the case of rare classes or in the case of an evolving population. For instance, an important problem in Biology is the detection

of novel species which could appear at any time resulting from structural or physiological modifications. Unfortunately, classical supervised algorithms, like support vector machines (SVM) or Linear Discriminant Analysis (LDA), will automatically label observations from a novel class as belonging to one of the known classes in the training set and will not be able to detect new classes. It is therefore important to find a way to allow the supervised classification methods to detect unobserved situations and to adapt themselves to the new configurations.

1.1 Related works

In statistical learning, the problem of classification with unobserved classes is a problem which has received very few attention. Indeed, both supervised and unsupervised classification contexts have been widely studied but intermediate situations have received less attention. We would like however to mention two related topics in statistical learning called semi-supervised classification and novelty detection.

Semi-supervised classification Semi-supervised classification focuses on supervised classification with partially labeled data. Usually, unlabeled data are added to the learning data in order to improve the efficiency of the final classifier. Such an approach is particularly useful when only few labeled observations are available for learning (applications with a high supervision cost). A good review on semi-supervised classification can be found in [16]. However, semi-supervised classification methods are not able to detect unobserved groups of points and, more importantly, will use them to re-estimate the model parameters of known classes and the estimates of known classes parameters will be deteriorated.

Novelty detection Novelty detection focuses on the identification of new or unknown data for which the learned classifier was not aware during the learning phase. This approach has become very popular in several application fields such as finance (fault and fraud detection), medical imaging (mass detection in mammograms) or web mining. In the last years, many

methods have been proposed to deal with this problem which can be split into two main categories: statistical and neural network based approaches. An excellent review on both categories of novelty detection methods can be found in [17] and [18]. However, even though all these methods are able to detect new or unobserved groups of points, no one of them is able to adapt the classifier to the new situation for classifying future observations.

1.2 The proposed approach

This work introduces an approach based on the mixture model which combines unsupervised and supervised learning for detecting in the test set unobserved groups of observations and for adapting the supervised classifier to the new situation in order to correctly classify new observations in the future. In the learning phase, a generative classifier of C classes is learned on the learning dataset in a classical way. The classification of the test set is done through a two-step approach which first searches for unobserved groups of points and then adapts the supervised classifier to the new situation. The discovery step which identifies new classes relies on an unsupervised modelling of the test data through EM-based strategies and on Bayesian model selection tools. Finally, the classification of new observations can be then done through the *maximum a posteriori* rule.

1.3 Organization of the paper

The paper is organized as follows. Section 2 introduces an adaptive discriminant analysis method based on the mixture model able to detect unobserved classes. Section 3 presents experimental results highlighting the main features of the proposed method on simulated and real datasets. Section 4 is devoted to the application of the proposed approach to the detection of novel species in DNA barcoding. Finally, Section 5 proposes some concluding remarks and discusses further works.

2 Adaptive model-based discriminant analysis

We introduce in this section an adaptive model-based classifier able to detect novel classes which have not been observed during the learning phase. Parameter estimation and model selection will be discussed as well.

2.1 The mixture model

Let us consider a classical parametric mixture model of K components: the observations $\{x_1, \dots, x_n\} \in \mathbb{R}^p$ are assumed to be independent realizations of a random vector $X \in \mathbb{R}^p$ with density:

$$f(x; \Theta) = \sum_{k=1}^K \pi_k f_k(x; \theta_k), \quad (1)$$

where $\pi_k \geq 0$ for $k = 1, \dots, K$ are the mixing proportions (with the constraint $\sum_{k=1}^K \pi_k = 1$), $f_k(x; \theta_k)$ is the density of the k th component of the mixture parameterized by θ_k and finally $\Theta = \{\pi_1, \dots, \pi_K, \theta_1, \dots, \theta_K\}$. We discuss briefly below the choice of the mixture densities.

Mixture of Gaussians Among all parametric densities, the Gaussian model is probably the most used in classification. The Gaussian mixture model has been extensively studied in the last decades and used in many situations (see [3] and [19] for a review). Therefore, if the Gaussian model is chosen, $f_k(x; \theta_k)$ will denote the density of a multivariate Gaussian density parameterized by $\theta_k = \{\mu_k, \Sigma_k\}$ where μ_k and Σ_k are respectively the mean and covariance matrix of k th component of the mixture.

Mixture of parsimonious Gaussians In some situations, modelling the data with a full covariance matrix can be too expensive in terms of number of parameters to estimate. In such a case, it is possible to make additional assumptions on the structure of the covariance matrix. For example, in the well-known Linear Discriminant Analysis (LDA) method, the covariance matrices of the different components are supposed to be equal to a unique

covariance matrix. It is also possible to assume that the covariance matrix of each mixture component is diagonal or proportional to the identity matrix. These models are known as parsimonious Gaussian models in the literature since they require to estimate less parameters than the classical Gaussian model. Celeux and Govaert proposed in [9] a family of parsimonious Gaussian models based on an eigenvalue decomposition of the covariance matrix including the previous models. These parsimonious Gaussian models were then applied in [5] to supervised classification.

Mixture of HD Gaussians Nowadays, many scientific domains produce high-dimensional data like medical research, image analysis or Biology (see an application to DNA barcoding in Section 4). Classifying such data is a challenging problem since the performance of classifiers suffers from the *curse of dimensionality* [4]. Classification methods based on Gaussian mixture models are directly penalized by the fact that the number of parameters to estimate grows up with the square of the dimension. However, parsimonious models are usually too constrained to correctly fit the data in a high-dimensional space. To overcome this problem, Bouveyron *et al.* proposed recently in [8] a family of Gaussian models adapted to high-dimensional data. This approach, based on the idea that high-dimensional data live in low-dimensional spaces, assumes that the covariance matrix of each mixture component has only $d_k + 1$ different eigenvalues where d_k is the dimension of the subspace of the k th mixture component. These Gaussian models were then used in [7] for high-dimensional data clustering.

Mixture with a noise component Banfield and Raftery have introduced in [3] a mixture model with a noise component in order to improve the robustness of the cluster analysis on noisy datasets. The original work proposed to add to the mixture model a uniform distribution over the convex hull of the data as an additional component. Good results of the use of this specific mixture model were observed in different situations. Hennig and Coretto [14] proposed recently to use an improper uniform distribution that does not depend on the data for improving the robustness and provide a

better approximation of the likelihood than the one proposed in the original work. An application of noise detection is proposed in Section 3.

2.2 Parameter estimation: inductive approach

This paragraph focuses on the estimation of the mixture parameters in the specific situation where one or several classes can not be observed during the learning phase. In the mixture model framework, the maximum likelihood (ML) estimation method is usually used for estimating the model parameters. For the mixture model (1), the log-likelihood has the following form:

$$\mathcal{L}(x_1, \dots, x_n; \Theta) = \sum_{i=1}^n \sum_{k=1}^K s_{ik} \log(\pi_k f_k(x_i; \theta_k)),$$

where $s_{ik} = 1$ if x_i belongs to the k th class and $s_{ik} = 0$ otherwise. However, this work considers a specific learning situation in which one or several classes are not represented in the learning dataset. Therefore, the mixture parameter estimation can not be done using the classical way and a two step approach made of a learning phase and a discovery phase is proposed below.

The learning phase For this first phase of the parameter estimation, let us assume that only C classes are represented in the learning dataset $\mathcal{X} = \{x_1, \dots, x_n\}$ with $1 \leq C \leq K$. Since the data of the learning set are complete, *i.e.* a label $z_i \in \{1, \dots, C\}$ is associated to each observation x_i of the learning set ($i = 1, \dots, n$), we fall into the classical estimation framework of model-based discriminant analysis. In such a case, the maximization of the likelihood reduces to separately estimate the parameters of each class density by maximizing the associated conditional log-likelihood \mathcal{L}_k , for $k = 1, \dots, C$:

$$\mathcal{L}_k(\mathcal{X}; \Theta) = \sum_{i=1}^n s_{ik} \log(\pi_k f_k(x_i; \theta_k)).$$

The maximization of the conditional log-likelihood $\mathcal{L}_k(\mathcal{X}; \Theta)$, for $k = 1, \dots, C$, conduces to an estimation of π_k by $\hat{\pi}_k = \frac{n_k}{n}$ where $n_k = \sum_{i=1}^n s_{ik}$ is the number of observations of the k th class and to an estimation of θ_k by $\hat{\theta}_k$

which depends on the chosen component density. For instance, in the case of a Gaussian density, the maximization of $\mathcal{L}_k(\mathcal{X}; \Theta)$ conduces to an estimation of $\theta_k = \{\mu_k, \Sigma_k\}$ by:

$$\begin{aligned}\hat{\mu}_k &= \frac{1}{n_k} \sum_{i=1}^n s_{ik} x_k, \\ \hat{\Sigma}_k &= \frac{1}{n_k} \sum_{k=1}^n s_{ik} (x_i - \hat{\mu}_k)(x_i - \hat{\mu}_k)^t,\end{aligned}$$

for $k = 1, \dots, C$. We refer respectively to [9] and [8] for parameter estimation in the case of parsimonious and HD Gaussian models, and to [3] in the case of a mixture with a noise component.

The discovery phase Usually, in discriminant analysis, the classification phase consists only in assigning new unlabeled observations to one of known classes. However, in this work, it is assumed that all the classes have not been observed during the learning phase. It is therefore necessary to search for new classes before to classify the new observations for avoiding to misclassify observations of an unobserved class (by assigning them to one of the observed classes). Using the model and the notations introduced above, it remains to find $K - C$ new classes in the set of n^* new unlabeled observations $\mathcal{X}^* = \{x_1^*, \dots, x_{n^*}^*\}$. Since these new observations are unlabeled, we have to fit the mixture model in a partially unsupervised way. In this case, the log-likelihood has the following form:

$$\mathcal{L}(\mathcal{X}^*; \Theta) = \sum_{i=1}^{n^*} \left(\sum_{k=1}^C s_{ik}^* \log(\pi_k f_k(x_i^*; \theta_k)) + \sum_{k=C+1}^K s_{ik}^* \log(\pi_k f_k(x_i^*; \theta_k)) \right),$$

where the parameters θ_k for $k = 1, \dots, C$ have been estimated in the previous phase and the parameters θ_k for $k = C + 1, \dots, K$ remain to estimate. Due to the constraint $\sum_{k=1}^K \pi_k = 1$ on the parameters π_k , it is not possible to keep the proportions estimated during the learning and the mixture proportions have to be re-estimated on the new sample $\{x_1^*, \dots, x_{n^*}^*\}$. However, the test set $\{x_1^*, \dots, x_{n^*}^*\}$ is an incomplete dataset since the labels z_i^* are missing and

the s_{ik}^* are consequently unknown for all observations of this set. In such a situation, the direct maximization of the likelihood is an intractable problem and the Expectation-Maximization (EM) algorithm [11] is usually used to estimate the mixture parameters by iteratively maximizing the likelihood. We propose below a modified EM algorithm for estimating the parameters of the $K - C$ unobserved classes which alternates between the following E and M steps at each iteration q :

- **E step:** the conditional probabilities $t_{ik}^{*(q)} = P(Z = k | X = x_i^*)$, for $i = 1, \dots, n^*$ and $k = 1, \dots, K$, are updated according to the mixture parameters as follows:

$$t_{ik}^{*(q)} = \frac{\hat{\pi}_k^{(q-1)} f_k(x_i^*; \hat{\theta}_k^{(q-1)})}{f(x; \hat{\Theta}^{(q-1)})},$$

where $\hat{\pi}_k^{(q-1)}$ and $\hat{\theta}_k^{(q-1)}$ are the mixture parameters estimated in the M step at the step $(q - 1)$.

- **M step:** the parameters of the $K - C$ unobserved classes are estimated by maximizing the likelihood conditionally to the probabilities $t_{ik}^{*(q)}$ whereas the estimated parameters of the observed classes remain fixed to the values obtained in the learning phase except for the proportions which are re-estimated. Therefore, this step only updates the estimates of parameters π_k for $k = 1, \dots, K$ and θ_k for $k = C + 1, \dots, K$. In the case of the Gaussian mixture, for instance, the update formulas for the parameter estimates are, for $k = 1, \dots, K$:

$$\hat{\pi}_k^{(q)} = \frac{n_k^{*(q)}}{n^*},$$

and for $k = C + 1, \dots, K$:

$$\hat{\mu}_k^{(q)} = \frac{1}{n_k^{*(q)}} \sum_{i=1}^{n^*} t_{ik}^{*(q)} x_i^*$$

$$\hat{\Sigma}_k^{(q)} = \frac{1}{n_k^{*(q)}} \sum_{i=1}^{n^*} t_{ik}^{*(q)} (x_i^* - \hat{\mu}_k^{(q)})(x_i^* - \hat{\mu}_k^{(q)})^t,$$

where $n_k^{*(q)} = \sum_{i=1}^{n^*} t_{ik}^{*(q)}$.

2.3 Parameter estimation: transductive approach

The previous paragraph proposed an EM-based algorithm which assumes that model parameters of the C observed classes have been estimated in the past (during the learning phase) and the modelling of new classes depends naturally on the quality of these estimates. However, in some situations, the cost of supervision is very high and only few labeled observations are available for learning. In such cases, the model parameters of observed classes could be badly estimated and the discovery phase could be consequently less efficient in detecting new classes. As the learning sample $\mathcal{X} = \{x_1, \dots, x_n\}$ and the test sample $\mathcal{X}^* = \{x_1^*, \dots, x_{n^*}^*\}$ are assumed to come from the same population, both samples could be used in the discovery phase to improve the model parameters while searching for unobserved classes in the test set. Such an approach should mainly benefit to the parameter estimation of the C observed classes and, consequently, should benefit as well to the detection of unobserved groups of observations. In this case, the log-likelihood has the following form:

$$\mathcal{L}(\mathcal{X}, \mathcal{X}^*; \Theta) = \sum_{i=1}^n \sum_{k=1}^C s_{ik} \log(\pi_k f_k(x_i; \theta_k)) + \sum_{i=1}^{n^*} \sum_{k=1}^K s_{ik}^* \log(\pi_k f_k(x_i^*; \theta_k)).$$

An alternative version of the EM-based algorithm proposed in the previous paragraph is presented below to jointly estimate model parameters while searching for new classes. The joint estimation procedure alternates between the following E and M steps at each iteration q :

- **E step:** on the one hand, the conditional probabilities $P(Z = k | X = x_i)$ remain fixed for the learning observations $\{x_1, \dots, x_n\}$ and are equal to s_{ik} , for $i = 1, \dots, n$ and $k = 1, \dots, K$, where $s_{ik} = 1$ if x_i belongs to the k th class and $s_{ik} = 0$ otherwise. On the other hand, the conditional

probabilities $t_{ik}^{*(q)} = P(Z = k | X = x_i^*)$ are updated for the test sample $\{x_1^*, \dots, x_{n^*}^*\}$, *i.e.* for $i = 1, \dots, n^*$ and $k = 1, \dots, K$, according to the mixture parameters as follows:

$$t_{ik}^{*(q)} = \frac{\hat{\pi}_k^{(q-1)} f_k(x_i^*; \hat{\theta}_k^{(q-1)})}{f(x_i^*; \hat{\Theta}^{(q-1)})},$$

where $\hat{\pi}_k^{(q-1)}$ and $\hat{\theta}_k^{(q-1)}$ are the mixture parameters estimated in the M step at the step $(q-1)$.

- **M step:** the parameters of the C observed classes and of the $K - C$ unobserved classes are estimated by maximizing the likelihood conditionally to the probabilities s_{ik} and $t_{ik}^{*(q)}$. Therefore, this step updates now the estimates of parameters π_k and θ_k for $k = 1, \dots, K$. In the case of the Gaussian mixture, for instance, the update formulas for the parameter estimates are, for $k = 1, \dots, K$:

$$\begin{aligned} \hat{\pi}_k^{(q)} &= \frac{n_k^{(q)} + n_k^{*(q)}}{n + n^*}, \\ \hat{\mu}_k^{(q)} &= \frac{1}{n_k^{(q)} + n_k^{*(q)}} \left(\sum_{i=1}^n s_{ik} x_i + \sum_{i=1}^{n^*} t_{ik}^{*(q)} x_i^* \right), \\ \hat{\Sigma}_k^{(q)} &= \frac{1}{n_k^{(q)} + n_k^{*(q)}} \left(S_k^{(q)} + S_k^{*(q)} \right). \end{aligned}$$

where $S_k^{(q)} = \sum_{i=1}^n s_{ik} (x_i - \hat{\mu}_k^{(q)}) (x_i - \hat{\mu}_k^{(q)})^t$, $S_k^{*(q)} = \sum_{i=1}^{n^*} t_{ik}^{*(q)} (x_i^* - \hat{\mu}_k^{(q)}) (x_i^* - \hat{\mu}_k^{(q)})^t$, $n_k^{(q)} = \sum_{i=1}^n s_{ik}$ and $n_k^{*(q)} = \sum_{i=1}^{n^*} t_{ik}^{*(q)}$.

2.4 Model selection: determining the number of components

In the usual case of supervised classification, the number of classes is known and the model selection consists only in choosing the most adapted densities for the considered dataset. In the context of the studied situation, the total number K of classes is assumed to be unknown and has to be chosen as well as the conditional densities of the mixture model. Classical tools for model selection in the mixture model framework are penalized likelihood criterions

and include the AIC [1], BIC [20] and ICL [6] criterions. The Bayesian Information Criterion (BIC) is certainly the most popular and consists in selecting the model which maximizes the quantity:

$$BIC(\mathcal{M}) = \mathcal{L}(x_1, \dots, x_n; \Theta) - \frac{\nu(\mathcal{M})}{2} \log(n),$$

where $\nu(\mathcal{M})$ is the the number of parameters in model \mathcal{M} and n is the number of observations. For instance, the number of parameters for the full Gaussian mixture model, *i.e.* full and different covariance matrices, is equal to $(k - 1) + kp + kp(p + 1)/2$. The AIC criterion penalizes the log-likelihood by $\nu(\mathcal{M})$ and the ICL criterion add to the BIC criterion the quantity $\sum_{i=1}^n \sum_{k=1}^K t_{ik} \log(t_{ik})$ in order to favour well separated models. An evaluation of both criterions in the context of unobserved class detection is presented in the next section.

2.5 Classification with the adapted classifier

The previous paragraphs introduced a model-based discriminant analysis method which adapts its mixture model to a new situation. Therefore, the adapted model can be used to classify new observations in the future. In the classical discriminant analysis framework, new observations are usually assigned to a class using the *maximum a posteriori* (MAP) rule. The MAP rule assigns a new observation $x \in \mathbb{R}^p$ to the class for which x has the highest posterior probability. Therefore, the classification step mainly consists in calculating the posterior probability $P(Z = k|X = x)$ for each class $k = 1, \dots, K$. In the case of the model described in this section, this posterior probability can be expressed classically using the Bayes' rule as follows:

$$P(Z = k|X = x) = \frac{\pi_k f_k(x; \theta_k)}{f(x; \Theta)},$$

where $f(x; \Theta) = \sum_{k=1}^K \pi_k f_k(x; \theta_k)$. Therefore, the posterior probabilities of the new observations depend on both the classes observed in the learning phase and the classes discovered in the discovery phase.




Figure 1: Classification with AMDA of the Iris dataset: the classes “setosa” (red triangles) and “versicolor” (green plus-es) have been observed during the learning phase whereas the class “virginica” (blue crosses) has not.

3 Experimental results

This section presents experiments on toy and simulated datasets in order to highlight the main features of the method introduced in the previous section.

3.1 An introductory example: the iris dataset

The dataset considered in this first experiment is a classical one: the iris dataset made famous by its use by Fisher in [12] as an example for discriminant analysis. This dataset, in fact collected by Edgar Anderson [2] in the Gaspé Peninsula (Canada), is made of three classes corresponding to different species of iris (*setosa*, *versicolor* and *virginica*) among which the classes *versicolor* and *virginica* are difficult to discriminate (they are at least not

linearly separable). The dataset consists of 50 samples from each of three species and four features were measured from each sample. The four measurements are the length and the width of sepal and petal. This dataset is used here as a toy dataset because of its popularity and its biological nature. Indeed, let suppose that botanists are studying iris species and have only observed the two species *setosa* and *versicolor*. For this experiment, the dataset has been randomly split into a learning dataset without *virginica* examples and a test dataset with several *virginica* examples. The top-left panel of Figure 1 shows what the botanists are supposed to have observed in the past. The top-center panel of the same figure presents a sample of new observations of iris for which the botanists are asked to classify. However, as the top-right panel indicates, this new sample contains individuals from a class which has not been observed by the botanists in the past and the iris experts will very likely classify all these new observations as belonging to either the class *setosa* or the class *versicolor*. The bottom-left panel of Figure 1 shows the result of such a *scenario*, using Quadratic Discriminant Analysis (QDA) in place of the iris experts, which yields to the classification of all *virginica* observations in the class *virginica*. Remark that, even though this results is disappointing from our point of view, it is understandable both for an human expert and a classification method since the classes *versicolor* and *virginica* are indeed very difficult to discriminate. The strategy proposed in the previous section, hereafter referred to by Adaptive Model-based Discriminant Analysis (AMDA), was applied to this dataset. The bottom-center and right panels of Figure 1 presents the results provided by AMDA. On the one hand, it turns out that Bayesian model selection criterions (BIC here) succeed in identifying a new group of points in the test set. On the other hand, once the number K of mixture component chosen, AMDA classifies almost perfectly (only 2 errors on this example) the observations of the unobserved class *virginica*.




Figure 2: Classification with AMDA of simulated data: 3 observed classes and 1 unobserved class (light blue diamonds) in \mathbb{R}^2 .

3.2 Novelty detection with a multi-class classifier

This paragraph presents a Monte Carlo experiment on simulated data in order to both evaluate Bayesian model selection criteria in the context of unobserved class detection and to determinate the breakdown group size for the detection of new classes. For this experiment, data were simulated according a Gaussian mixture model made of 4 groups (3 observed groups and one unobserved group) in a 2-dimensional space. Means and covariance matrices were chosen in order to obtain separated enough groups. Figure 2 presents the simulated learning dataset (left panel), the test dataset (center panel) and the true labels of the test observations (right panel). For each of the 50 replications of the Monte Carlo study, 750 observations were simulated according to a Gaussian mixture model (250 obs. for each of the observed classes) in order to form the learning set and $250 + \eta$ observations were simulated for the test set where η is the number of observations of the unobserved class. For each replication, the number η varied from 2 to 50. The left panel of Figure 3 presents the number of times where the correct number of groups was selected by the three selection model criteria AIC, BIC and ICL. It appears that the three studied selection model criteria select always the correct number of groups when the unobserved group size if larger than 10. For smaller sizes of the unobserved group, AIC turns out to be the more stable criterion since it selects the correct number of groups




Figure 3: Selection of the number of classes and recognition rates according to the size of the unobserved class with AMDA on simulated data: 3 observed classes and 1 unobserved class in \mathbb{R}^2 .

more frequently than the two other criteria. We therefore recommend the use of AIC as model selection criterion in such a context. The right panel of Figure 3 shows three recognition rates averaged on the Monte Carlo replications: total recognition rate, true positive rate and false positive rate. The total recognition rate measures the overall correct classification rate for the four classes (the three observed classes and the unobserved one). The true positive rate measures the correct classification rate for observations of the unobserved class (class #4). Conversely, the false positive rate evaluates how many observations of the three observed classes are classified as belonging to the new class. In a satisfying situation, the total recognition rate and the true positive rate should be close to 1 whereas the false positive rate should be close to 0. Both recognition rates were computed on the test dataset. The right panel of of Figure 3 shows that the three recognition rates are very good for sizes of the unobserved class larger than 10. We observe as well that for sizes of the unobserved class smaller than 5 the true positive rate is very unstable and this means that the unobserved class is not well modeled. To summarize, this experiment demonstrates that AMDA in com-




Figure 4: Classification with AMDA of simulated data: 3 observed classes and 1 unobserved noise class (light blue diamonds) in \mathbb{R}^2 .

bination with AIC succeeds in detecting and modelling an unobserved class in the test set for unobserved classes as small as 5–10 observations.

3.3 Detection of an unobserved noise class

This new experiment aims to evaluate the ability of AMDA to detect an unobserved non Gaussian class of noise. For this, data were simulated in a 2-dimensional space according a mixture model made of 4 components: 3 Gaussian components and one uniform noise component. Means and covariance matrices of Gaussians were again chosen in order to obtain separated enough groups. The learning set was made of 750 observations from the three Gaussian classes. The top-left panel of Figure 4 shows the observations of the learning set. The test set was made of 250 observations from the three Gaussian classes (observed during the learning) and 31 observations

Truth Classif.	1	2	3	4
1	75	2	0	0
2	0	78	0	10
3	0	0	65	21
-	-	-	-	-

(a) Confusion matrix for QDA

Truth Classif.	1	2	3	4
1	75	2	0	0
2	0	77	0	0
3	0	0	64	0
4	0	1	1	31

(b) Confusion matrix for AMDA

Table 1: Confusion tables for QDA and AMDA on the test dataset for the simulated data with one unobserved noise class (class #4).

from the unobserved uniform noise class. The top-center panel of this figure presents the unlabeled test observations and the top-left panel indicates the true labels of these observations. The bottom-left panel of Figure 4 shows the classification of the test observations with the supervised classifier Quadratic Discriminant Analysis (QDA). Unsurprisingly, QDA classifies all the observations from the noise class to one of the three known Gaussian classes. Table 1 presents confusion tables for QDA and AMDA on the test dataset and shows that all noise observations were classified into the classes #2 and #3. The bottom-center and right panels of Figure 4 show respectively the BIC values for AMDA with and without a noise component and the classification provided by AMDA with a detected noise component (as indicated by the largest BIC value). We can observe on this quite complex example that AMDA succeeds in both detecting the unobserved noise class and modelling it through a uniform component. Table 1.b confirms that AMDA recognizes all noise observations as belonging to one unobserved class in the past and makes only 2 false noise detections which is very satisfying.

4 Application to DNA barcoding

Determining to what species an organism belongs is probably the most common problem in Biology. The answer concerns many areas of practical importance such as protecting endangered species, sustaining natural resources, stopping disease vectors or monitoring environmental quality. Created in

2003, the Consortium for the Barcode of Life¹ is an international initiative devoted to developing DNA barcoding as a standard tool to identify species. Its purpose is to provide a simple and automatic method to correctly identify the species, with no or limited recourse to taxonomic expertise. The 5' half of the mtDNA gene COI has been chosen as the barcode locus for most animals, and gene markers with similar barcoding properties are investigated in plants, fungi, and protists. Traditionally, the barcoding procedure is based on an algorithm combining k -NN with neighbour-joining trees [15]. Several alternatives to this method were quite successfully applied to various kinds of organisms, although several major problems remain. One of them is: how to detect new or unobserved species? This section aims to demonstrate that AMDA can be used in such an application to detect new or unobserved species.

4.1 The data

The data considered for this application come from a study on neotropical bats within Guyana [10]. The original dataset contains DNA barcodes of 840 bat specimens representing 87 species. Each of the 840 DNA sequences has a length of 657 variables and where each variable is coding for the “A”, “C”, “G”, “T” nucleotides. Among the 840 observations of the original set, we only kept for our study the 471 observations belonging to the 6 most represented species (most of the original species only contain few observations). The 657 qualitative variables were also transformed to quantitative variables using Multiple Correspondence Analysis (see [13] for details) and the 176 axes associated to the largest eigenvalues (explaining more than 90% of the total variance) were kept for the remaining of the experiment. Finally, duplicates observations were removed from the dataset since many observations were not unique due to the nature of the data (DNA sequences). The final dataset contains 199 observations from 6 bat species described on 176 quantitative variables.

¹<http://www.barcodingoflife.org>

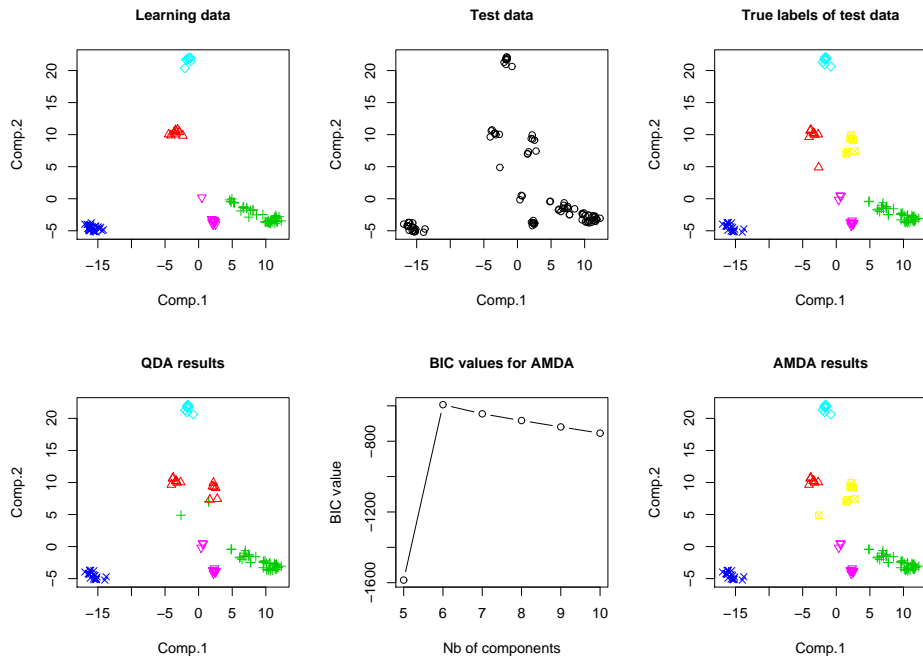


Figure 5: Classification with AMDA of the DNA barcode data: 5 observed bat species and 1 unobserved bat species represented on the two first principal axes.

4.2 Experimental results

For this experiment, the preprocessed dataset was split into a learning set of 100 observations without observations of the 6th class and a test set of 99 observations containing 8 observations of the 6th class (assumed to be unobserved). The top panels of Figure 5 show the learning and test data on the two first principal axes. We can first observe that the different species are globally well separated and that some classes have inhomogeneous distribution. In particular, the classes #1 (red triangles) and #5 (purple inverse triangles) seem to be made of several sub-species. It appears as well that the observations of the unobserved class (yellow squares) are very close to observations of classes #1, #2 and #5 and this could complicate the discovery and classification tasks. Unsurprisingly, the supervised classifier QDA

Truth Classif.	1	2	3	4	5	6
1	8	0	0	0	0	6
2	1	42	0	0	0	2
3	0	0	20	0	0	0
4	0	0	0	9	0	0
5	0	0	0	0	11	0
-	-	-	-	-	-	-

(a) Confusion matrix for QDA

Truth Classif.	1	2	3	4	5	6
1	8	2	0	0	0	0
2	0	42	0	0	0	0
3	0	0	20	0	0	0
4	0	0	0	9	0	0
5	0	0	0	0	11	0
6	1	0	0	0	0	8

(b) Confusion matrix for AMDA

Table 2: Confusion tables for QDA and AMDA on the the DNA barcode test dataset: 5 observed bat species and 1 unobserved bat species (class #6).

assigned the observations of the new class to the known classes #1, #2 and #5 as the bottom-left panel of Figure 5 shows. AMDA was also applied to these data and BIC indicates (*cf.* bottom-center panel) that the most adapted model is a mixture model with 6 components which means that the test set contains 1 unobserved class. The bottom-right panel of Figure 5 shows the final classification of the test dataset provided by AMDA. As we can see, AMDA has correctly detected and classified the 8 observations of the unobserved species but misclassified 1 observations from class #1 (*cf.* Table 2). However, the belonging to the class #1 of this observation could be discussed regarding its DNA sequence. It is indeed natural to suspect a labelling or a sequencing error in this case. To summarize, this experiment has demonstrated the ability of AMDA to detect new species in a complex real-world context and to adapt the supervised classifier to the new situation.

5 Conclusion and further works

This work has focused on the problem of learning a supervised classifier with unobserved classes. An adaptive model-based discriminant analysis method has been presented in this paper which is able to both detect unobserved groups of points in a new set of observations and to adapt the supervised classifier to the new situation. A two step approach made of a learning phase and a discovery phase has been proposed for parameter estimation

and two EM-based procedures have been presented for the discovery phase. Experiments on simulated and real datasets have shown that the proposed method is able to detect different kinds of unobserved classes (Gaussian, uniform noise, ...). The proposed strategy has been also applied with success to an important problem in Biology: the detection of novel species in DNA barcoding. It remains however to deal with the problem of label switching when $C - K > 1$. A way to solve this problem could be to ask domain experts to classify some observations of the new detected groups in order to associate a class name with the detected groups. Finally, it could be very interesting to study the evolution of the proposed strategy in the context of dynamic classification.

References

- [1] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.
- [2] E. Anderson. The irises of the gaspé peninsula. *Bulletin of the American Iris Society*, 59:2–5, 1935.
- [3] J. Banfield and A. Raftery. Model-based Gaussian and non-Gaussian clustering. *Biometrics*, 49:803–821, 1993.
- [4] R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- [5] H. Bensmail and G. Celeux. Regularized Gaussian discriminant analysis through eigenvalue decomposition. *Journal of the American Statistical Association*, 91:1743–1748, 1996.
- [6] C. Biernacki, G. Celeux, and G. Govaert. Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(7):719–725, 2000.
- [7] C. Bouveyron, S. Girard, and C. Schmid. High-Dimensional Data Clustering. *Computational Statistics and Data Analysis*, 52(1):502–519, 2007.

- [8] C. Bouveyron, S. Girard, and C. Schmid. High Dimensional Discriminant Analysis. *Communications in Statistics: Theory and Methods*, 36(14):2607–2623, 2007.
- [9] G. Celeux and G. Govaert. Gaussian parsimonious clustering models. *Pattern Recognition*, 28(5):781–793, 1995.
- [10] E. Clare, B. Lim, M. Engstrom, J. Eger, and P. Hebert. DNA barcoding of neotropical bats: species identification and discovery within guyana. *Molecular Ecology Notes*, 7:184–190, 2007.
- [11] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 39(1):1–38, 1977.
- [12] R. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7:179–188, 1936.
- [13] M. Greenacre and J. Blasius. *Multiple Correspondence Analysis and Related Methods*. Chapman & Hall, 2006.
- [14] C. Hennig and P. Coretto. *The Noise Component in Model-based Cluster Analysis*, pages 127–138. Data Analysis, Machine Learning and Applications. Springer, 2008.
- [15] R. Kelly, I. Sarkar, D. Eernisse, and R. Desalle. DNA barcoding using chitons (genus mopalialia). *Molecular Ecology Notes*, 7:177–183, 2007.
- [16] B. Krishnapuram, D. Williams, Y. Xue, A. Hartemink, L. Carin, and M. Figueiredo. On semi-supervised classification. In *NIPS*, 2004.
- [17] M. Markou and S. Singh. Novelty detection: A review - part 1: Statistical approaches. *Signal Processing*, 83(12):2481–2497, 2003.
- [18] M. Markou and S. Singh. Novelty detection: A review - part 2: Neural network based approaches. *Signal Processing*, 83(12):2499–2521, 2003.
- [19] G. McLachlan and D. Peel. *Finite Mixture Models*. Wiley, New York, 2000.

- [20] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6:461–464, 1978.