Error rate control for classification rules in multiclass mixture models

Tristan Mary-Huard, Vittorio Perduca, Marie Laure Martin-Magniette, Gilles Blanchard

► To cite this version:

Tristan Mary-Huard, Vittorio Perduca, Marie Laure Martin-Magniette, Gilles Blanchard. Error rate control for classification rules in multiclass mixture models. The international journal of biostatistics, 2021, 10.1515/ijb-2020-0105 . hal-03357461

HAL Id: hal-03357461

https:

//hal-universite-paris-saclay.archives-ouvertes.fr/hal-03357461

Submitted on 28 Sep 2021

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.

Error rate control for classification rules in multiclass mixture models

Tristan Mary-Huard

GQE - Le Moulon and UMR MIA-Paris
INRAE, CNRS, AgroParisTech, Université Paris-Saclay

Vittorio Perduca

Laboratoire MAP5 (UMR CNRS 8145),
Université Paris Descartes, Université de Paris

Marie-Laure Martin-Magniette

Institute of Plant Sciences Paris-Saclay (IPS2) and UMR MIA-Paris
Université Paris-Saclay, Université de Paris,
AgroParisTech, CNRS, INRAE, Univ. Evry,

Gilles Blanchard

Laboratoire de Mathématiques d'Orsay,
Inria, CNRS, Université Paris-Saclay

Abstract

In the context of finite mixture models one considers the problem of classifying as many observations as possible in the classes of interest while controlling the classification error rate in these same classes. Similar to what is done in the framework of statistical test theory, different type I and type II-like classification error rates can be defined, along with their associated optimal rules, where optimality is defined as minimizing type II error rate while controlling type I error rate at some nominal level. It is first shown that finding an optimal classification rule boils down to searching an optimal region in the observation space where to apply the classical Maximum A Posteriori (MAP) rule. Depending on the misclassification rate to be controlled, the shape of the optimal region is provided, along with a heuristic to compute the optimal classification rule in practice. In particular, a multiclass FDR-like optimal rule is defined and compared to the thresholded MAP rules that is used in most applications. It is shown on both simulated and real datasets that the FDR-like optimal rule may be significantly less conservative than the thresholded MAP rule.

1 Introduction

Consider a sample $X_1, \dots, X_n$ of independent observations in space $\mathcal{X}$ stemming from P populations. The distribution of such a sample can be modelled as a finite mixture of distributions

$$f(x) = \sum_{p=1}^P \pi_p f_p(x)$$

where each class is described by its own probability density distribution f_p and a weight π_p so that $0 \leq \pi_p \leq 1$ and $\sum_p \pi_p = 1$. For each observation a label variable Z can be introduced, that equals p if the observation belongs to class p . In the following, the labels $Z_1, \dots, Z_n$ are supposed to be independent, and the mixture model is assumed to be completely known, i.e. the number of classes P , the class distributions $f_1, \dots, f_P$ and the weights $\pi_1, \dots, \pi_P$ are known.

In most practical situations the labels are unobserved, and the goal is to find a suitable clustering of the sample. This is typically obtained by applying a classification rule, i.e. a function ψ that maps $\mathcal{X}$ into $\{1, \dots, P\}$. The key quantity for building such a function is the posterior probability for observation x to belong to class p :

$$\tau_p(x) = \Pr(Z = p | X = x) = \frac{\pi_p f_p(x)}{\sum_{p'=1}^P \pi_{p'} f_{p'}(x)}. \quad (1)$$

The Maximum A Posteriori (MAP) classification rule, defined as

$$\psi^{MAP}(x) = \operatorname{argmax}_{1 \leq p \leq P} \tau_p(x)$$

is by far the most popular classification rule. In particular, it is known to minimize the classification error rate $P(\psi(X) \neq Z)$, McLachlan and Peel (2000).

While optimality in terms of classification error rate is a desirable property, two drawbacks of the MAP rule should be mentioned. First, optimality does not prevent against a high level of misclassification. When the classification task is difficult, as much as half of the observations can be misclassified by the MAP rule. Second, the MAP rule does not account for the asymmetry that may exist between classes: in some real applications only a small number among the P classes may be of major importance for the experimenter. One should then focus on the misclassification rate on these classes of interest. This situation arises in most cases involving comparisons between different conditions, where one is interested in identifying units having different behaviors in different conditions, but not units whose behavior is unchanged across conditions. For instance in Bérard, Martin-Magniette, Brunaud, Aubourg, and Robin (2011), a methylation analysis was conducted to identify regions of the genome that are differentially methylated between different organs of *Arabidopsis thaliana*. This problem can be cast into a clustering problem where units are probes (locations) spread over the *Arabidopsis* genome, and classes are defined according to the methylation status of these probes in the different organs. Only classes corresponding to differentially methylated behaviors are then of interest.

Different solutions have been proposed to circumvent these two problems. One possible approach is to take into account the asymmetry between classes by selecting unequal misclassification cost functions that emphasize the cost of misclassification $c_{k\ell}$ from a non-interesting class ℓ into a class k of interest. This has been investigated in Friedman, Hastie, and Tibshirani (2009) for instance.

On the other hand, a high level of misclassification for the MAP rule may come from the fact that it classifies *any* observation x , regardless of the uncertainty of the classification of point x . This uncertainty may be quantified through the probability to not belong to the class x was classified into. A possible solution is to perform a partial classification, using e.g. the thresholded classification rule

$$\psi^\alpha(x) = \begin{cases} p & \text{if } p = \operatorname{argmax}_{1 \leq p' \leq P} \tau_{p'}(x) \text{ and } \tau_p(x) > 1 - \alpha, \\ 0 & \text{(not classified) otherwise,} \end{cases} \quad (2)$$

where $0 < \alpha < 1$ is a parameter to be chosen by the experimenter. The condition $\tau_p(x) >$

$1 - \alpha$ can be reformulated as $1 - \tau_p(x) < \alpha$ which amounts to focusing on the classification of observations for which uncertainty is not too high.

The design of partial rules has been considered through the use of classification rules with a reject (or abstention) option. An often considered setting there is that the reject option has a fixed cost (less costly than a misclassification error); see Chow (1970), Herbei and Wegkamp (2006), Pillai, Fumera, and Roli (2013). This setting has been also considered in conjunction with specific machine learning methods such as large margin classifiers (Bartlett and Wegkamp, 2008, Grandvalet, Rakotomamonjy, Keshet, and Canu, 2009, Wegkamp and Yuan, 2011), active learning (Zhang and Chaudhuri, 2014), and fair learning (Schreuder and Chzhen, 2021).

Alternative partial classification rules include the tight clustering algorithm introduced in Tseng and Wong (2005) and extended to deal with large datasets in Karmakar, Das, Bhattacharya, Sarkar, and Mukhopadhyay (2019). In the latter work, when applying the extended tight clustering algorithm to a dataset consisting of more than 50,000 gene expression probes for individuals suffering from psoriasis, more than 30,000 probes were not classified in any of the six clusters that were identified.

All aforementioned approaches (partial rules, rules with different misclassification costs and rule with abstention costs) are attractive, however there exist very few guidelines about how to tune parameters such as α , c_{kl} , or the cost of abstaining, and what will be the impact of this tuning on the final misclassification rate. In this article, we propose a method that addresses both the problem of partial classification and the presence of classes of no interest. In particular, we propose to design the abstention region based on a directly interpretable constraint taking the form of an overall statistical confidence level on the effective classifications. Based on the control of type I error rate in statistical test theory, we introduce optimal classification rules that guarantee that as many observations as possible are classified, under the constraint that the rate of misclassification in the classes of interest is controlled. As such, our method can be understood as an extension of the Bayesian False Discovery Rate (BFDR) developed in Efron and Tibshirani (2002) to general finite mixture models. It is shown that different classification rules should be considered to achieve optimality according to the number K of classes of interest ($K = P$ or $K < P$). Sub-optimality of the thresholded classification rule (2) is also demonstrated.

The principle of controlling a specific error rate related to classical statistical confidence criteria has also been considered in, e.g. Scott and Nowak (2005) and Scott, Bellala, Willett et al. (2009) in the context of binary classification; see also Tong, Feng, and Zhao (2016) for a more recent survey. In that setting, an asymmetry is introduced between the classes: the goal is to maximize correct classification rate in class 1 subject to a fixed control at a prescribed level of the classification error in class 0, either in absolute-value (Neyman-Pearson classification) or in the sense of FDR. The setting we introduce here is

multi-class and includes the possibility of abstention. Furthermore all the classes of interest have the same status (the classes of no interest have a different status, however even if all classes are of interest, our setting makes sense and abstention remains a possibility). The goal we aim at is to guarantee a certain confidence level over classified examples while minimizing the probability of abstention of classes of interest. The statistical criteria we introduce generalize the confidence/coverage tradeoff considered in El-Yaniv and Wiener (2010), Wiener and El-Yaniv (2015), Denis and Hebiri (2020), who only considered the binary classification case where all classes are of interest.

The problem of classification with confidence has also been considered under a related but different angle by Lei (2014). In Lei’s approach, only the binary classification setting is considered and two regions C_0 and C_1 , where the respective classes are predicted, are constructed such that C_i has a prescribed coverage $1 - \alpha_i$ for class i and the ”ambiguous classification region” $C_0 \cap C_1$ has minimal overall probability. We observe that the latter overlap region could be interpreted as similar to an abstention region in our setting. However, the criterion used by Lei (2014) is different from ours, and finding the optimal classification regions amounts to determine appropriate thresholds for two different level sets of the posterior class probability function. Furthermore, it is unclear how to extend the considered criterion to more than two classes. In our approach, we consider a global constraint which in our view is more readily interpretable in terms of classification confidence, is naturally adapted to more than two classes, and only necessitates the determination of a single level set of an appropriate function.

The paper is organized as follows. In Section 2, notions corresponding to type I and type II error rates are defined in the multiclass classification framework along with optimality for a classification rule. Optimal rules are then exhibited in Section 3, and heuristic procedures to estimate these rules are presented. The empirical behavior of the procedures is briefly investigated and an application of the proposed strategy to the analysis of differential methylation is presented in Section 4. Lastly, some discussion is developed in Section 5.

2 Definitions

2.1 Restricted classification rules

Because the classification task may be arbitrarily difficult for some points of $\mathcal{X}$, we consider restricted rules $\psi_R(x)$, that classify points in a subset $R \subseteq \mathcal{X}$ only:

$$\psi_R : \mathcal{X} \rightarrow \{0, 1, \dots, P\},$$

where 0 is the status “not classified” given to observations in $\bar{R}$, the complementary set of R in $\mathcal{X}$:

$$\forall x \in \bar{R}, \psi_R(x) = 0.$$

As an illustration, consider the thresholded classification rule ψ^α provided in equation (2), and define

$$\tau^*(x) = \max_{1 \leq p \leq P} \tau_p(x),$$

which corresponds to the maximal posterior probability for point x . Then classification rule ψ^α corresponds to the restricted classification rule ψ_{R_α} where

$$R_\alpha = \{x | \tau^*(x) \geq 1 - \alpha\}.$$

In some cases, the number K of classes of interest may be lower than the total number of classes P in the mixture. In the following, and without loss of generality, we will assume that the classes of interest are the first K classes of the mixture. In this context where $K < P$, it will be convenient to consider classification rules that are also restricted to predict labels in $1, \dots, K$:

$$\psi_{K,R} : \mathcal{X} \rightarrow \{0, 1, \dots, K\}.$$

In this situation, the maximal posterior probability of interest is

$$\tau_K^*(x) = \max_{1 \leq k \leq K} \tau_k(x).$$

Note that in the case $K = P$ one has $\tau_P^*(x) = \tau^*(x)$. One can extend the definition of the thresholded classification rule to the case $K < P$ by defining

$$\psi_K^\alpha(x) = \begin{cases} k & \text{if } k = \underset{1 \leq k' \leq K}{\operatorname{argmax}} \tau_{k'}(x) \text{ and } \tau_K^*(x) > 1 - \alpha, \\ 0 & \text{(not classified) otherwise,} \end{cases} \quad (3)$$

2.2 Error rates

A “good” classification rule should classify as many observations as possible while controlling the classification error rate. As in the statistical testing setting, one can introduce the type I and type II error rates associated to any given classification rule. We consider two definitions of the type I error rate and one for the type II error rate in the multiclass classification context:

Definition 2.1 Let $\{1, \dots, K\}$ be the classes of interest (possibly with $K = P$) and $\psi_{K,R}$ be a restricted classification rule. Quantities

$$\begin{aligned} MNPR(\psi_{K,R}) &= \mathbb{P}(\psi_{K,R}(X) \neq Z, \psi_{K,R}(X) > 0) = \mathbb{P}(\psi_{K,R}(X) \neq Z, X \in R) \\ \text{and } MFDR(\psi_{K,R}) &= \begin{cases} \mathbb{P}(\psi_{K,R}(X) \neq Z | \psi_{K,R}(X) > 0) = \mathbb{P}(\psi_{K,R}(X) \neq Z | X \in R), \\ \text{if } \mathbb{P}(X \in R) > 0 \\ 0, \text{ if } \mathbb{P}(X \in R) = 0, \end{cases} \end{aligned}$$

are called the multiclass Neyman-Pearson error rate and the multiclass false discovery rate, respectively. The quantity

$$MFNR(\psi_{K,R}) = \mathbb{P}(Z \in \{1, \dots, K\}, \psi_{K,R}(X) = 0) = \mathbb{P}(Z \in \{1, \dots, K\}, X \in \bar{R})$$

is called the multiclass false negative rate.

The MFNR quantifies the proportion of observations that belong to a class of interest and that were not classified by rule $\psi_{K,R}$. This quantity corresponds to the false negative proportion in the testing setting, and should be as small as possible. The MNPR and MFDR correspond to possible classification error rates one may want to control at a given level α . Note that a criterion equivalent to the MFDR has been considered by El-Yaniv and Wiener (2010), Wiener and El-Yaniv (2015), Denis and Hebiri (2020), albeit only in the situation $K = P = 2$.

To exemplify the previous definitions, here again one can consider the thresholded classification rule $\psi^\alpha (= \psi_p^\alpha)$ defined above in the by-default context where all classes are of interest. One can show that the MFDR of ψ^α is always lower than α , whatever the number of classes of interest. Indeed one has

$$\begin{aligned} MFDR(\psi^\alpha) &= \mathbb{P}(\psi^\alpha(X) \neq Z | X \in R_\alpha) \\ &= \frac{\mathbb{P}(\psi^\alpha(X) \neq Z \cap X \in R_\alpha)}{\mathbb{P}(X \in R_\alpha)}. \end{aligned}$$

The numerator in the last expression can be reformulated as follows:

$$\begin{aligned} \mathbb{P}(\psi^\alpha(X) \neq Z \cap X \in R_\alpha) &= \int_{\mathcal{X}} \mathbb{E}_{Z|X=x} [I_{\{\psi^\alpha(x) \neq Z\}}] I_{\{x \in R_\alpha\}} f(x) dx \\ &= \int_{\mathcal{X}} \mathbb{E}_{Z|X=x} [Z \neq p^*] I_{\{x \in R_\alpha\}} f(x) dx \\ &= \int_{\mathcal{X}} (1 - \tau_p^*(x)) I_{\{x \in R_\alpha\}} f(x) dx \end{aligned}$$

where $p^* = \arg \max_p \tau_p(x)$. Plugging this expression into the previous equation leads to

$$\begin{aligned} MFDR(\psi^\alpha) &= \frac{1}{\mathbb{P}(X \in R_\alpha)} \int_{R_\alpha} (1 - \tau_{p^*}^*(x)) f(x) dx \\ &\leq \frac{1}{\mathbb{P}(X \in R_\alpha)} \int_{R_\alpha} \alpha f(x) dx \\ &\leq \alpha. \end{aligned}$$

2.3 Optimal classification rules

We now introduce a formal definition for the *optimality* of a restricted classification rule.

Definition 2.2 *Given a level α ($0 \leq \alpha \leq 1$) and a set of classes of interest $\{1, \dots, K\}$ (possibly with $K = P$), a classification rule $\psi_{K,R}^*$ is MNPR-optimal at level α if*

$$\psi_{K,R}^* = \arg \min_{\psi_{K,R}} MFNR(\psi_{K,R}) \text{ u.c. } MNPR(\psi_{K,R}) \leq \alpha. \quad (4)$$

Alternatively, $\psi_{K,R}^$ is MFDR-optimal at level α if*

$$\psi_{K,R}^* = \arg \min_{\psi_{K,R}} MFNR(\psi_{K,R}) \text{ u.c. } MFDR(\psi_{K,R}) \leq \alpha. \quad (5)$$

Note that the definition implies both a region R as large as possible and an optimal classification at each point x in region R .

3 Optimal classification rules

The goal of the present section is to exhibit the shape of the optimal classification rules for problems (4) and (5). To this end, we will first prove that for a fixed region R it is optimal to apply the restricted MAP rule. As a consequence, looking for an optimal rule ψ_R^* actually boils down to finding the optimal region R^* where to apply the (restricted) MAP rule. The theoretical form of the optimal region will then be derived for problems (4) and (5). Each optimal region requires the tuning of a unknown threshold λ , for which we will provide a heuristic estimation method.

3.1 Optimal rule when R is fixed

Let us first define the restricted MAP classification rule $\psi_{K,R}^{MAP}$ as

$$\psi_{K,R}^{MAP}(x) = \begin{cases} \arg \max_{1 \leq k \leq K} \tau_k(x) & \text{if } x \in R \\ 0 & \text{otherwise.} \end{cases}$$

Let first consider a classification $\psi_{K,R}$ restricted to the region R . At any point $x \in R$ one has

$$\begin{aligned}\mathbb{P}(\psi_{K,R}(x) \neq Z) &= \sum_{k=1}^K (1 - \tau_k(x)) I_{\{\psi_{K,R}(x)=k\}} \\ &\geq (1 - \tau_K^*(x)) \\ &\geq \mathbb{P}(\psi_{K,R}^{MAP}(x) \neq Z)\end{aligned}$$

and therefore

$$\begin{aligned}MNPR(\psi_{K,R}) &\geq MNPR(\psi_{K,R}^{MAP}) \\ \text{and } MFDR(\psi_{K,R}) &\geq MFDR(\psi_{K,R}^{MAP}).\end{aligned}$$

Alternatively it is straightforward to observe that

$$MFNR(\psi_{K,R}) = MFNR(\psi_{K,R}^{MAP}).$$

As a consequence, one can conclude that for a fixed region R the optimal restricted classification rule is $\psi_{K,R}^{MAP}$. Therefore finding the solutions of problems (4) and (5) boils down to finding the region R^* to which the MAP classification rule should be applied.

In the following the upperscript MAP and the lowerscript K in $\psi_{K,R}^{MAP}$ will be skipped for simplicity, and ψ_R should be understood as $\psi_{K,R}^{MAP}$, with $K = P$ or $K < P$ depending on the context.

3.2 Optimal classification region for MNPR control

If all classes are of interest ($K = P$), one looks for a region R^* such that

$$R^* \in \arg \min_R P(X \in \bar{R}) \quad \text{u.c.} \quad \mathbb{P}(\psi_R(X) \neq Z, X \in R) \leq \alpha.$$

If $K < P$, the region R^* should satisfy

$$R^* \in \arg \min_R P(Z \in \{1, \dots, K\}, X \in \bar{R}) \quad \text{u.c.} \quad \mathbb{P}(\psi_R(X) \neq Z, X \in R) \leq \alpha. \quad (6)$$

Proposition 1 *If $MNPR(\psi_{\mathcal{X}}) \leq \alpha$, then $R^* = \mathcal{X}$ is a solution of (6). Otherwise, any R^* satisfying $MNPR(\psi_{R^*}) = \alpha$ and such that*

$$\left\{ x \mid \frac{\sum_{k=1}^K \tau_k(x)}{1 - \tau_K^*(x)} > \lambda \right\} \subseteq R^* \subseteq \left\{ x \mid \frac{\sum_{k=1}^K \tau_k(x)}{1 - \tau_K^*(x)} \geq \lambda \right\}, \quad (7)$$

for some $\lambda \geq 0$, is a solution of (6).

If the marginal distribution of X is without atoms, there exists R^* satisfying the above conditions.

In the particular case where $K = P$, condition (7) can be reduced to:

$$\{x \mid \tau^*(x) > \lambda\} \subseteq R^* \subseteq \{x \mid \tau^*(x) \geq \lambda\}. \quad (8)$$

for some $\lambda \geq 0$.

The proof is based on the original proof of the Neyman-Pearson theorem, Neyman and Pearson (1933), and can be found in Appendix A.

3.3 Optimal classification region for MFDR control

If all classes are of interest, one looks for a region R^* such that

$$R^* \in \arg \min_R P(X \in \bar{R}) \text{ u.c. } \mathbb{P}(\psi_R(X) \neq Z \mid X \in R) \leq \alpha.$$

If $K < P$, R^* should satisfy

$$R^* \in \arg \min_R P(Z \in \{1, \dots, K\}, X \in \bar{R}) \text{ u.c. } \mathbb{P}(\psi_R(X) \neq Z \mid X \in R) \leq \alpha. \quad (9)$$

Proposition 2 If $MFDR(\psi_{\mathcal{X}}) \leq \alpha$, then $R^* = \mathcal{X}$ is a solution of (9). If $\tau_K^*(X) < 1 - \alpha$ almost surely, then $R^* = \emptyset$ is a solution of (9).

Otherwise, any R^* satisfying $MFDR(\psi_{R^*}) = \alpha$ and such that

$$\left\{x \mid \frac{1 - \alpha - \tau_K^*(x)}{\sum_{k=1}^K \tau_k(x)} < \lambda\right\} \subseteq R^* \subseteq \left\{x \mid \frac{1 - \alpha - \tau_K^*(x)}{\sum_{k=1}^K \tau_k(x)} \leq \lambda\right\}, \quad (10)$$

for some $\lambda \geq 0$, is a solution of (9); if the marginal distribution of X is without atoms, such an R^* exists.

In the case $K = P$, condition (10) can be reduced to:

$$\{x \mid \tau^*(x) > 1 - \alpha - \lambda\} \subseteq R^* \subseteq \{x \mid \tau^*(x) \geq 1 - \alpha - \lambda\}, \quad (11)$$

for some $\lambda \geq 0$.

The proof of Proposition 2 can be found in Appendix B. In the case $K = P$, the optimal region takes the form of a thresholded rule (2) for an appropriate threshold. This was observed by Denis and Hebiri (2020) who studied optimal rules for the MFDR/coverage

tradeoff in the case $K = P = 2$. Interestingly, when $K < P$, the optimal rule is different from a simple threshold rule.

The form of the optimal region R^* in (8), (10) is $C_{<\lambda} \subseteq R^* \subseteq C_{\leq\lambda}$, where $C_{<\lambda}, C_{\leq\lambda}$ are sublevel sets (in the sense of strict, resp. non-strict inequality) of a criterion depending on the posterior probabilities at point x . This “sandwiching” relation is theoretically relevant in situations where $\mathbb{P}(C_{<\lambda}) < \mathbb{P}(C_{\leq\lambda})$, for example if the criterion used is piece-wise constant, and the level set at value λ has nonzero probability. In this situation, strictly speaking to achieve exactly the target MNPR or MFDR rate, R^* should only include part of the level set at value λ (and this part can be chosen arbitrarily provided the constraint is satisfied). The assumption of atom-free X -marginal could also be lifted to grant the existence of optimal rules in general, provided randomized rules are allowed, as in classical Neyman-Pearson theory. Since such situations are obviously of little relevance for most applications, in the remainder of this work, to simplify exposition we assume that $\mathbb{P}(C_{<\lambda}) = \mathbb{P}(C_{\leq\lambda})$ and that R^* exactly coincides with a sublevel set.

3.4 Comparison between thresholded rule and the MFDR-optimal rule

The restriction region R_α of the thresholded rule (3) can now be compared to the MFDR-optimal regions R^* found in the previous section. One has

$$\begin{aligned} R_\alpha &= \{x | \tau_K^*(x) \geq 1 - \alpha\} \text{ with } K \leq P, \\ R^* &= \left\{ x | \tau_K^*(x) \geq 1 - \alpha - \lambda \sum_{k=1}^K \tau_k(x) \right\} \text{ with } \lambda \geq 0, \text{ if } K \leq P. \end{aligned}$$

According to the shape of R^* , the threshold for classifying an observation x based on the maximal posterior probability $\tau_K^*(x)$ depends not only on α , but also on $\sum_{k=1}^K \tau_k(x)$, that is the probability of belonging to the overall group of populations of interest. When $K < P$ this can be interpreted as follows: when $\tau_K^*(x)$ is mild, the optimal rule may still classify x if the probability of x to belong to the overall group of the populations of interest is high. This interesting feature is not accounted for in the thresholded rule that classifies an observation only based on $\tau_K^*(x)$. In this sense, the thresholded rule is not based on the appropriate classification criterion when one is interested only in a subset of all possible populations.

More generally, one can observe that whatever the number of populations of interest, both ψ_{R^*} and ψ^α classify observations based on $\tau^*(x)$, but using a threshold that is always lower for ψ_{R^*} . Consequently i) the MFNR will always be higher for ψ^α , and ii) the MFDR of ψ^α will be lower than the requested nominal level, i.e. the thresholded rule is a *conservative*

classification strategy to control the MFDR at level α . In Section 4 we illustrate on simulated and real data the fact that the gap in term of MFNR between ψ_{R^*} and ψ^α may be high.

3.5 Estimation of parameter λ

We present here an heuristic strategy to choose λ from the data at hand. First note that all optimal regions defined in Sections 3.2 and 3.3 are of the form of level sets

$$R^* = \{ x \mid \text{Crit}(x) > \lambda \},$$

where $\text{Crit}(x)$ is a criterion based on the posterior probabilities at point x , that depends on the risk one wants to control and on the number of classes of interest as follows:

★ if one aims at controlling the MNPR then

$$\text{Crit}(x) = \frac{\sum_{k=1}^K \tau_k(x)}{1 - \tau_K^*(x)} \text{ (case } K < P) \quad \text{or} \quad \text{Crit}(x) = \tau^*(x) \text{ (case } K = P),$$

★ if one aims at controlling the MFDR then

$$\text{Crit}(x) = \frac{\tau_K^*(x) + \alpha - 1}{\sum_{k=1}^K \tau_k(x)} \text{ (case } K < P) \quad \text{or} \quad \text{Crit}(x) = \tau^*(x) + \alpha - 1 \text{ (case } K = P)$$

where $\lambda \leq 0$.

Assuming an n -sample $D_n = (X_1, \dots, X_n)$ is available, the MNPR and MFDR can be estimated for any region R by

$$\begin{aligned} \widehat{\mathbb{P}}(\psi(X) \neq Z, X \in R) &= \frac{1}{n} \sum_{X_i \in R} (1 - \tau_K^*(X_i)) \\ \text{and} \quad \widehat{\mathbb{P}}(\psi(X) \neq Z \mid X \in R) &= \frac{1}{n_R} \sum_{X_i \in R} (1 - \tau_K^*(X_i)) \end{aligned}$$

respectively, where n_R is the number of observations in D_n belonging to R (and we recall that $\tau_K^*(x) = \tau^*(x)$ if $K = P$). This leads to the following general heuristic for the evaluation of threshold λ :

Algorithm 1

Require: $\alpha, x_1, \dots, x_n$

1/ Order observations $x_1, \dots, x_n$ according to $Crit(x_i)$:

$$Crit(x_{(1)}) \geq \dots \geq Crit(x_{(N)}),$$

2/ Find the largest index i_{max} such that $\frac{1}{M(i_{max})} \sum_{i=1}^{i_{max}} (1 - \tau_K^*(x_{(i)})) \leq \alpha$,

3/ Set $\hat{\lambda} = Crit(x_{(i_{max})})$,

4/ Output $\hat{\lambda}$.

If the goal is to control MFDR, then in the second step of the heuristic $M(i_{max}) = i_{max}$, otherwise $M(i_{max}) = n$ for a MNPR control.

4 Applications

The aim of this section is to illustrate the performance of the MFDR classification rule derived in the previous section. We first consider different scenarios based on simulated data. In a first scenario the true posterior probabilities are available, which corresponds to the theoretical setting of Sections 2 and 3. In this scenario the only parameter to be estimated is λ , and we empirically evaluate the ability of the previous algorithm to efficiently estimate λ and control the error rate at the required nominal level. The performance of the optimal rule is also compared to the performance of the thresholded rule. In a second scenario the true posterior probabilities are assumed to be unknown - the by-default setting of most application cases. This scenario allows us to evaluate the impact of estimating the posterior probabilities on the error rate control procedure. Lastly, we present an application on real data, where the MFDR control procedure is applied to the differential analysis of methylation profiles.

4.1 Simulation setting

Datasets are simulated from a mixture of 3 bidimensional Gaussian distributions whose mean vectors are $(-1, 0)$, $(0, D)$ and $(1, 0)$, respectively. Here D is a parameter that tunes the distance between populations 1 and 3 and population 2. When $D = 0$ the three populations are highly overlapping, whereas a high value of D makes the populations more distinct. In each class the covariance matrix is diagonal, with an identical variance σ^2 . Weights π_1, π_2, π_3 are all fixed at 1/3. For each dataset, 200 observations per population are simulated. In this context the easiness of the classification task is ruled by parameters

D and σ^2 : the higher D (respectively the lower σ^2), the easier the classification. Several values are considered for D (0, 1, 2, 3) and σ^2 (0.5, 1, 2). An illustration of different configurations ranging from an “easy” classification case ($D = 3$, $\sigma^2 = 0.5$) to a “hard” one ($D = 0$, $\sigma^2 = 2$) can be found in Appendix C. Finally, 100 datasets are generated for each configuration.

In a supplementary simulation study, datasets were generated from a mixture of three bivariate Student distributions; details can be found in Appendix F.

4.2 Case 1: posterior probabilities are known

In what follows the objective is to perform classification with a MFDR controlled at a nominal threshold fixed at $\alpha = 0.05$.

In Section 2.2, a straightforward calculation has shown that the thresholded classification rule ψ_α with a threshold fixed at $1 - \alpha$ guarantees a MFDR control at a level α . However the results of Section 3.3 suggest that fixing the threshold at $1 - \alpha$ may result in a high MFNR. The alternative procedure for the optimal rule described in Section 3.5 should yield better performance thanks to a less conservative choice for the threshold. The potential gain (in terms of MFNR) of the optimal rule compared to the $1 - \alpha$ thresholded rule is investigated in this section.

The analysis is performed as follows: for each dataset the posterior probabilities are computed for each observation using the true parameters of the model. The procedure (10) is then applied to estimate the threshold $\hat{\lambda}$ in order to control the MFDR at nominal level $\alpha = 0.05$. Observations are classified using either the MAP rule, the thresholded rule (with $\alpha = 0.05$) or the estimated optimal rule (i.e. the optimal rule applied with estimated threshold $\hat{\lambda}$). For all rules the actual *MFDR* and *MFNR* are evaluated by comparing the predicted and true labels of the classified observations. Results are displayed in Figure 1 for the case where all classes are of interest. The case where only classes 1 and 3 are of interest is provided in Appendix D.

In the case where $\sigma^2 = 0.5$ one can observe that all classification rules achieve a low MFDR - that may still be much higher than 0.05 in the case of the MAP rule. Both the thresholded and the optimal rule efficiently achieve control of the MFDR at the nominal level, but with significant differences in terms of MFNR. As illustrated in Figure 2, the estimated threshold for the optimal rule may be much lower (below 0.5 in some configurations) than the $1 - \alpha = 0.95$ value used in the thresholded rule. As a consequence, the MFNR is twice higher for the thresholded rule than for the optimal rule, illustrating how conservative the thresholded rule can be.

In the case where $\sigma^2 = 2$ the classification task becomes quite difficult, and the MAP rule yields MFDR that are higher than 50%. In such configurations both the thresholded

Figure 1: Performances of the MAP, optimal and $1 - \alpha$ thresholded classification rules in terms of realized MFDR and MFNR. Columns correspond to different values of the variance σ^2 . The black line on the top graphs corresponds to the nominal level of 5%. This figure appears in color in the electronic version of this article, and color refers to that version.

Figure 2: Boxplots of the estimated threshold for the optimal rule. Colors correspond to levels of the variance σ^2 . The horizontal line indicates the threshold 0.95 used in the thresholded rule. This figure appears in color in the electronic version of this article.

and the optimal rules do not classify any observation in most cases, since the highest posterior probability observed (in a class of interest) is lower than 0.95.

In conclusion, whatever the configuration, the optimal rule $\psi_{\hat{R}^*}$ with estimated threshold $\hat{\lambda}$ controls the misclassification rate as requested, whereas the thresholded classification rule is more conservative. When the classification task is easy or intermediate the MFNR of the optimal rule can be much smaller than the MFNR of the thresholded classification rule. Results are unchanged if the prior proportions of the mixture vary (not shown), and/or if only a subset of the three classes are of interest (see Appendix D).

4.3 Case 2: posterior probabilities are inferred

So far we assumed that the true parameters of the mixture are known, but in practice these parameters are usually estimated from the data at hand. To evaluate the impact of parameter estimation on misclassification rate, we performed the same analysis as in the previous section, except that parameters are now supposed to be unknown and are estimated using the `mclust` R package (Scrucca, Fop, Murphy, and Raftery, 2016), fixing the number of classes to its true value (i.e. 3). Results are displayed in Figure 3.

When the classification problem is tractable (low or moderate values of σ^2 and/or high values of D) both the thresholded and the optimal rules efficiently control the MFDR. As in the previous scenario, a significant gain in terms of MFDR is observed when using the optimal rule rather than the thresholded rule. When the classification problem becomes too difficult (i.e. classes strongly overlap) the estimated posterior probabilities get inaccurate, and any classification rule based on these quantities becomes irrelevant. Although neither the thresholded nor the optimal rule control the MFDR at the nominal level, it is still worth to use one of these classification rules rather than the MAP rule. Considering the case where only a subset of classes is of interest leads to similar conclusions (see Figure 7 in Appendix D).

The performance of the optimal rule depends on the accuracy of the posterior probability estimates. In order to explore the robustness of our method under model misspecification that could possibly lead to biased posterior probabilities estimates, we generated datasets according to a mixture of three Student bidimensional distributions with the same centers and covariance matrices as in the main Gaussian simulation scheme of “intermediate” difficulty ($\sigma = 1$). We then estimated the posterior probabilities using the EM algorithm for Gaussian mixtures implemented in `mclust`. Results are shown in Appendix F. As expected, for higher degrees of freedom, the error rates closely match the ones in the central column ($\sigma = 1$) of Figure 3. For lower degrees of freedom, performance breaks down in the “difficult” configurations given by low values of D . For more tractable configurations, MFDR control is generally achieved and we do not observe notably differences

Figure 3: Same figure as Figure 1, except that posterior probabilities are now estimated. This figure appears in color in the electronic version of this article.

in terms of MFNR with respect to the Gaussian simulations.

4.4 Application to transcriptomic data

We consider the unsupervised classification problem described in Bérard et al. (2011). A methylation experiment was performed to compare the methylation profiles of two organs (leaf and seed) of *Arabidopsis thaliana*. The methylation profiles were measured on a same plant using a tiling array technology. The genomic sequence of *Arabidopsis thaliana* is represented on the array by approximately 5×10^5 probes, covering both genic and intergenic regions. For each probe, the methylation signal is measured in the two organs. From a statistical point of view the sampled population is the population of probes, each of them being described by a bivariate signal (hybridization in leaf and seed), and the goal of the analysis is to identify differentially methylated probes, i.e. probes whose methylation signals in leaf and seed differ. In the following only probes corresponding to chromosome 4 (107,199 probes) are considered. A more thorough description of the data can be found in Bérard et al. (2011).

In Bérard et al. (2011) a constrained Gaussian bivariate mixture model was fitted to the data. Comparing two samples requires distinguishing four different classes of probes that can be biologically interpreted as follows : a class of probes with low methylation signals (1), a class of probes exhibiting a similar intensity level in both organs (2), a class of probes with lower methylation intensities in leaf compared to seed (3), and a symmetric class of probes with higher intensities in leaf compared to seed (4) (see Appendix E for a graphical representation of the model, and the original article for technical details). Note that in the initial article the number of classes of the mixture was directly deduced from the biological comparison to be performed. Based on this four component mixture model, posterior probabilities to belong to each class were computed for each probe. In the original article the MAP rule was applied to infer class memberships, and probes classified into one of the two classes of interest (under and over-methylated classes) were identified and further investigated. Starting from the same posterior probabilities as in the original publication, we performed probe classification into the classes of interest using the optimal classification rule defined in section 3. Here $P = 4$ and $K = 2$, and the MFDR is controlled at nominal level $\alpha = 0.1$. Among the 15,801 probes initially classified as under or over-methylated, 13,065 are classified by the resulting MFDR classification rule, yielding an MFNR estimated at 0.0624.

The different classification rules are illustrated in Figure 4. The left panel represents the MAP rule used in the initial analysis (black = non-methylated, red = identically methylated, blue = under-methylated in leaf, green = over-methylated in leaf). The center panel corresponds to the optimal rule. Here the same colouring of the points is applied, with

Figure 4: **Left:** Probes of chromosome 4 coloured according to the MAP classification rule. **Center:** Probes of chromosome 4 coloured according to the optimal classification rule. **Right:** Histogram of the maximum posterior probabilities for purple probes classified by the MAP rule but not by the optimal rule. This figure appears in color in the electronic version of this article.

grey points corresponding to probes unclassified by the optimal rule and purple points corresponding to probes initially classified as over or under-methylated by the MAP rule but unclassified by the optimal rule.

As expected, all purple points are positioned on the boundaries between the two classes of interest and the two other classes. Although the posterior probability to be either over-methylated or under-methylated is higher than any other one for these probes, the actual value of the maximal posterior probability may be quite low, ranging between 0.25 and 0.6 (see Figure 4, right). In contrast, all probes classified by the optimal rule have posterior probabilities higher than 0.57. One can also notice that the optimal classification rule discards some obvious spurious classifications, such as the ones observed for very low values of the methylation signal, that were probably due to the constrained shapes of the covariance matrices of the adjusted model. Such points are discarded due to their low posterior probabilities.

In order to gain some additional insight regarding the boundary probes, we investigated to which extent the classification of a given probe is consistent with the ones of the adjacent (left and right) probes¹. For a given probe, the classification consistency score counts the number of adjacent probes sharing the same classification status. This score takes value

¹Although not accounted for in the initial publication, the information about the genomic position of the probes is available.

	0	1	2
Ident. methyl.	3236 (0.28)	4374 (0.38)	4019 (0.35)
Non-methyl.	3159 (0.04)	18750 (0.24)	57858 (0.73)
Over methyl.	3469 (0.61)	1304 (0.23)	894 (0.16)
Under methyl.	5850 (0.58)	2544 (0.25)	1740 (0.17)
(a)			
	0	1	2
Not classified	8514 (0.09)	23650 (0.25)	61968 (0.66)
Over methyl.	2650 (0.57)	1150 (0.25)	869 (0.19)
Under methyl.	4550 (0.54)	2172 (0.26)	1674 (0.2)
(b)			

Table 1: **(a)** Distribution of the probe counts in terms of MAP classification (rows) and classification consistency with neighbours (columns). The four classes correspond to “non-methylated” (1), “identically methylated” (2), “under-methylated in leaf” (3) and “over-methylated in leaf” (4). The classification consistency score counts the number of adjacent probes sharing the same classification status as the current probe. Numbers in brackets correspond to per row fractions. **(b)** Same distribution for the optimal classification.

0, 1 or 2, with 2 indicating a high classification consistency between the probe and its two neighbours.

Table 1 displays the distribution of probes in terms of class assignment and classification consistency, for the MAP rule (left) and the optimal rule (right). Focusing on class “over-methylated” one notices that the proportion of probes with score 1 or 2 is higher when using the optimal classification rule rather than the MAP classification rule. This illustrates the fact that most of the probes discarded by the optimal rule (corresponding to the purple points of Figure 4, center) are isolated probes exhibiting weak evidence for methylation. One can conclude that applying the optimal rule yields a more robust set of candidate probes by shaving inconsistent candidates.

5 Discussion

The methodology presented here builds on two central ideas: i) in many unsupervised settings only a subset of the classes is of interest for the practitioner, and ii) one would like to provide some guarantees (in terms of error rates) regarding the classification of observations into these specific classes of interest. To this end, classes of interest must be identified beforehand. In some contexts this identification is straightforward. In the

methylation application of Section 4 both the *a priori* number of classes P and the number of classes of interest K are known, a situation that may arise whenever the application context corresponds to a differential analysis setting where two or more conditions are compared. Such examples correspond to the ideal application cases of MFDR control procedures. However, the procedure may also be directly applied to other contexts where P and/or K is unknown. Two examples of such applications are

- cases where P is unknown but there exists a clear H_0 class with known distribution, and the goal is to identify observations that do not belong to the H_0 class; here both P and K are initially unknown, and K can be set to $P - 1$ once the fitting and selection of the model is performed. Note that in this context and for a fixed nominal level of MFDR, the MFNR could reach different values for different choices of P (hence K). In particular, one could expect the MFNR to increase with higher values of P . Fixing a maximum level for the MFNR could then guide the model selection, by e.g. choosing $\hat{P}$ as the largest P satisfying the constraint - a larger P leading to a finer granularity of the classification.
- cases where both P and K are unknown but prior information about some observations belonging to a class of interest is available. For instance, in a genomics context one may identify classes of genes based on their expression profiles, then i) identify classes where some genes have a known biological function, and ii) apply the MFDR control procedure to classes corresponding to these functions.

In such application cases where either P or K are unknown, the MFDR procedure may be impacted by the model selection, something that was not investigated here. To what extent this impact will affect the procedure is difficult to quantify and will directly depend on the quality of the *a priori* knowledge one has at hand. For instance in the genomics application mentioned above, applying the MFDR procedure to a well-characterized class of genes may be highly efficient even if the number of classes P is poorly estimated or if some (other) classes are poorly fitted by the inferred model.

While the present paper focused on finding optimal classification rules in an unsupervised framework, a same motivation may exist in the context of (semi-) supervised classification. From a practical point of view, since the optimal classification rules derived in this article only depend on posterior probabilities, they can be applied to any statistical method that yields such probabilities, e.g. logistic regression or discriminant analysis, and could also be extended to methods for which pseudo posterior probabilities can be obtained (see Tao, Wu, Wang, and Wang, 2005, Grandvalet, Mariéthoz, and Bengio, 2006 and references therein).

Future work also includes the definition and derivation of optimal restricted classification rules in contexts where the labels are not assumed to be independent, such as hidden Markov models and/or latent variables models for network data (stochastic block model and latent block model, see Matias and Robin, 2014) where existing results are restricted to the binary classification case (Sun and Cai, 2009). Additionally the consistency of the MFDR estimates presented here could be investigated using techniques similar to the ones developed in Denis and Hebiri (2020), such that their theoretical results could be extended to the case where $P > 2$ and only some classes are of interest.

Acknowledgments

G. Blanchard acknowledges support from Agence Nationale de la Recherche (ANR) via the project ANR-19-CHIA-0021-01 (BiSCottE), and the project ANR-16-CE40-0019 (SansSouci); and from the Franco-German University through the binational Doktorandenkolleg CDFA 01-18. GQE and IPS2 benefit from the support of the LabEx Saclay Plant Sciences-SPS (ANR-17-EUR-0007).

References

- Bartlett, P. and M. Wegkamp (2008): “Classification with a reject option using a hinge loss.” *Journal of Machine Learning Research*, 9.
- Bérard, C., M.-L. Martin-Magniette, V. Brunaud, S. Aubourg, and S. Robin (2011): “Un-supervised classification for tiling arrays: Chip-chip and transcriptome,” *Statistical applications in genetics and molecular biology*, 10.
- Chow, C. (1970): “On optimum recognition error and reject tradeoff,” *IEEE Transactions on Information Theory*, 16, 41–46.
- Denis, C. and M. Hebiri (2020): “Consistency of plug-in confidence sets for classification in semi-supervised learning,” *Journal of Nonparametric Statistics*, 32, 42–72.
- Efron, B. and R. Tibshirani (2002): “Empirical bayes methods and false discovery rates for microarrays,” *Genetic Epidemiology*, 23, 70–86.
- El-Yaniv, R. and Y. Wiener (2010): “On the foundations of noise-free selective classification.” *Journal of Machine Learning Research*, 11.

- Friedman, J., T. Hastie, and R. Tibshirani (2009): *The elements of statistical learning: data mining, inference, and prediction*, Springer series in statistics New York.
- Grandvalet, Y., J. Mariéthoz, and S. Bengio (2006): “A probabilistic interpretation of svms with an application to unbalanced classification,” in *Advances in Neural Information Processing Systems*, 467–474.
- Grandvalet, Y., A. Rakotomamonjy, J. Keshet, and S. Canu (2009): “Support vector machines with a reject option,” in *Advances in Neural Information Processing Systems 21 (NIPS 2008)*, MIT press, 537–544.
- Herbei, R. and M. H. Wegkamp (2006): “Classification with reject option,” *Canadian Journal of Statistics*, 34, 709–721.
- Karmakar, B., S. Das, S. Bhattacharya, R. Sarkar, and I. Mukhopadhyay (2019): “Tight clustering for large datasets with an application to gene expression data,” *Scientific Reports*, 9, 3053.
- Lei, J. (2014): “Classification with confidence,” *Biometrika*, 101, 755–769.
- Matias, C. and S. Robin (2014): “Modeling heterogeneity in random graphs through latent space models: a selective review,” *ESAIM: Proceedings and Surveys*, 47, 55–74.
- McLachlan, G. J. and D. Peel (2000): *Finite mixture models*, New York: Wiley.
- Neyman, J. and E. S. Pearson (1933): “On the problem of the most efficient tests of statistical hypotheses,” *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231, 289–337.
- Pillai, I., G. Fumera, and F. Roli (2013): “Multi-label classification with a reject option,” *Pattern Recognition*, 46, 2256–2266.
- Schreuder, N. and E. Chzhen (2021): “Classification with abstention but without disparities,” .
- Scott, C., G. Bellala, R. Willett, et al. (2009): “The false discovery rate for statistical pattern recognition,” *Electronic Journal of Statistics*, 3, 651–677.
- Scott, C. and R. Nowak (2005): “A Neyman-Pearson approach to statistical learning,” *IEEE Transactions on Information Theory*, 51, 3806–3819.

- Scrucca, L., M. Fop, T. Murphy, and A. Raftery (2016): “mclust 5: clustering, classification and density estimation using Gaussian finite mixture models,” *The R Journal*, 8, 289–317.
- Sun, W. and T. T. Cai (2009): “Large-scale multiple testing under dependence,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71, 393–424.
- Tao, Q., G.-W. Wu, F.-Y. Wang, and J. Wang (2005): “Posterior probability support vector machines for unbalanced data,” *IEEE Transactions on Neural Networks*, 16, 1561–1573.
- Tong, X., Y. Feng, and A. Zhao (2016): “A survey on neyman-pearson classification and suggestions for future research,” *Wiley Interdisciplinary Reviews: Computational Statistics*, 8, 64–81.
- Tseng, G. C. and W. H. Wong (2005): “Tight clustering: A resampling-based approach for identifying stable and tight patterns in data,” *Biometrics*, 61, 10–16.
- Wegkamp, M. and M. Yuan (2011): “Support vector machines with a reject option,” *Bernoulli*, 17, 1368–1385.
- Wiener, Y. and R. El-Yaniv (2015): “Agnostic pointwise-competitive selective classification,” *Journal of Artificial Intelligence Research*, 52, 171–201.
- Zhang, C. and K. Chaudhuri (2014): “Beyond disagreement-based agnostic active learning,” in *Advances in Neural Information Processing Systems 27 (NIPS 2014)*, 442–450.

Appendix A: Proof of Proposition 1

We denote ψ the MAP rule amongst class 1 to K , and ψ_R this rule restricted to a region R . If $MNPR(\psi_{\mathcal{X}}) \leq \alpha$, since $MFNR(\psi_R) \geq MFNR(\psi_{\mathcal{X}})$ for any region R , obviously $R^* = \mathcal{X}$ is a solution.

Otherwise, let R^* be a region satisfying the conditions of the theorem, and let R be any region such that the MAP classification rule restricted to R satisfies $MNPR(\psi_R) \leq \alpha = MNPR(\psi_{R^*})$. Let us prove that $MFNR(\psi_R) \geq MFNR(\psi_{R^*})$. Observing that $R = R \cap (R^* \cup \bar{R}^*)$ and $R^* = R^* \cap (R \cup \bar{R})$, one has :

$$\begin{aligned} & \mathbb{P}(\psi(X) \neq Z, X \in R) \leq \mathbb{P}(\psi(X) \neq Z, X \in R^*) \\ \Leftrightarrow & \mathbb{P}(\psi(X) \neq Z, X \in R \cap \bar{R}^*) \leq \mathbb{P}(\psi(X) \neq Z, X \in R^* \cap \bar{R}) \\ \Leftrightarrow & \int_{R \cap \bar{R}^*} (1 - \tau_K^*(x)) f(x) dx \leq \int_{R^* \cap \bar{R}} (1 - \tau_K^*(x)) f(x) dx, \end{aligned}$$

where the last equation follows from the application of the MAP rule on R and R^* respectively.

Similarly,

$$\begin{aligned} & P(Z \in \{1, \dots, K\}, X \in \bar{R}) - P(Z \in \{1, \dots, K\}, X \in \bar{R}^*) \geq 0 \\ \Leftrightarrow & P(Z \in \{1, \dots, K\}, X \in \bar{R} \cap R^*) - P(Z \in \{1, \dots, K\}, X \in \bar{R}^* \cap R) \geq 0. \end{aligned}$$

Therefore

$$\begin{aligned} P(Z \in \{1, \dots, K\}, X \in \bar{R} \cap R^*) &= \int_{\bar{R} \cap R^*} \left(\sum_{k=1}^K \tau_k(x) \right) f(x) dx \\ &\geq \lambda \int_{\bar{R} \cap R^*} (1 - \tau_K^*(x)) f(x) dx \\ &\geq \lambda \int_{\bar{R}^* \cap R} (1 - \tau_K^*(x)) f(x) dx \\ &\geq \int_{\bar{R}^* \cap R} \left(\sum_{k=1}^K \tau_k(x) \right) f(x) dx \\ &= P(Z \in \{1, \dots, K\}, X \in \bar{R}^* \cap R), \end{aligned}$$

where the first and third inequalities above follow from the property of R^* (and thus $\bar{R}^*$) given in equation (7).

Assuming that the marginal distribution of X is without atoms, we now establish that there exists a region R^* satisfying (7) for some $\lambda \geq 0$, and satisfying the MNPR constraint with equality, provided $MNPR(\psi_{\mathcal{X}}) > \alpha$. Define the event $A = \{\psi(X) \neq Z\}$, and Q the measure on the space $\mathcal{X}$ as $Q(R) = P(\{X \in R\} \cap A) = MNPR(\psi_R)$. Since the marginal distribution of X is without atoms, so is Q .

Denote $\xi(x) = (\sum_{k=1}^K \tau_k(x)) / (1 - \tau_K^*(x))$, and define the level sets for any $\lambda \geq 0$:

$$L(\lambda) = \{x \in \mathcal{X} : \xi(x) \geq \lambda\}; \quad L^>(\lambda) = \{x \in \mathcal{X} : \xi(x) > \lambda\}.$$

Observe that it holds $L^>(\lambda) \subseteq L(\lambda)$ for any λ , and $L(\lambda) = \bigcap_{\lambda' < \lambda} L^>(\lambda')$, $L^>(\lambda) = \bigcup_{\lambda' > \lambda} L^>(\lambda')$, so that

$$\limsup_{\lambda' \searrow \lambda} Q(L^>(\lambda')) = Q(L^>(\lambda)) \leq Q(L(\lambda)) = \liminf_{\lambda' \nearrow \lambda} Q(L^>(\lambda')). \quad (12)$$

Define

$$\lambda^* = \sup\{\lambda \geq 0 : Q(L^>(\lambda)) \geq \alpha\} \cup \{0\},$$

By the definition of λ^* , it holds

$$Q(L^>(\lambda^*)) \leq \alpha \leq Q(L(\lambda^*));$$

this is true by (12) if $\lambda^* > 0$, and, if $\lambda^* = 0$, it holds as well since we assumed $Q(L(0)) = Q(\mathcal{X}) = MNPR(\psi_{\mathcal{X}}) > \alpha$. (It can be also checked that it still holds if $\lambda^* = \infty$, which can happen in principle when $\mathbb{P}(\xi(X) = \infty) > 0$.) Therefore, since Q is without atoms, there exists a set $R \subseteq \mathcal{X}$ such that $L^>(\lambda^*) \subseteq R \subseteq L(\lambda^*)$ with $Q(R) = MNPR(\psi_R) = \alpha$, as required.

Appendix B: Proof of Proposition 2

If $MFDR(\psi_{\mathcal{X}}) \leq \alpha$, since $MFNR(\psi_R) \geq MFNR(\psi_{\mathcal{X}})$ for any region R , obviously $R^* = \mathcal{X}$ is a solution. For any region $R \subseteq \mathcal{X}$ with $P(X \in R) > 0$, it holds that $MFDR(\psi_R) = \mathbb{E}[1 - \tau_K^*(X)|X \in R]$. Therefore, if $\tau_K^*(X) < 1 - \alpha$ almost surely, it holds $MFDR(\psi_R) > \alpha$ for any region R with $P(X \in R) > 0$. Thus, in this case only a region R with null probability (in particular $R^* = \emptyset$) satisfies the MFDR constraint since by convention $MFDR(\Psi_{\emptyset}) = 0$.

Otherwise, let R^* be a region satisfying (10) and the MFDR constraint with equality. Consider R any region satisfying

$$\mathbb{P}(\psi(X) \neq Z | X \in R) = \mathbb{P}(\psi(X) \neq Z | X \in R^*) = \alpha.$$

On the one hand, one has :

$$\mathbb{P}(X \in R) - \mathbb{P}(X \in R^*) = \mathbb{P}(X \in R \cap \bar{R}^*) - \mathbb{P}(X \in R^* \cap \bar{R}).$$

On the other hand,

$$\begin{aligned} \mathbb{P}(\psi(X) \neq Z | X \in R) = \alpha &\Leftrightarrow \mathbb{P}(\psi(X) \neq Z, X \in R) = \alpha \mathbb{P}(X \in R) \\ \Rightarrow \alpha [\mathbb{P}(X \in R) - \mathbb{P}(X \in R^*)] &= \mathbb{P}(\psi(X) \neq Z, X \in R) - \mathbb{P}(\psi(X) \neq Z, X \in R^*) \\ &= \mathbb{P}(\psi(X) \neq Z, X \in R \cap \bar{R}^*) - \mathbb{P}(\psi(X) \neq Z, X \in R^* \cap \bar{R}). \end{aligned}$$

Hence

$$\begin{aligned} \mathbb{P}(\psi(X) \neq Z, X \in R \cap \bar{R}^*) - \mathbb{P}(\psi(X) \neq Z, X \in R^* \cap \bar{R}) &= \alpha [\mathbb{P}(X \in R \cap \bar{R}^*) - \mathbb{P}(X \in R^* \cap \bar{R})] \\ \Rightarrow \int_{R \cap \bar{R}^*} (1 - \tau_K^*(x) - \alpha) f(x) dx - \int_{R^* \cap \bar{R}} (1 - \tau_K^*(x) - \alpha) f(x) dx &= 0 \\ \Rightarrow \lambda \left(\int_{R \cap \bar{R}^*} \left(\sum_{k=1}^K \tau_k(x) \right) f(x) dx - \int_{R^* \cap \bar{R}} \left(\sum_{k=1}^K \tau_k(x) \right) f(x) dx \right) &\leq 0 \\ \Rightarrow \mathbb{P}(Z \in \{1, \dots, K\}, X \in R \cap \bar{R}^*) - \mathbb{P}(Z \in \{1, \dots, K\}, X \in R^* \cap \bar{R}) &\leq 0 \\ \Rightarrow \mathbb{P}(Z \in \{1, \dots, K\}, X \in \bar{R}) - \mathbb{P}(Z \in \{1, \dots, K\}, X \in \bar{R}^*) &\geq 0. \end{aligned}$$

Analogously to the proof of Proposition 1, we establish the existence of R^* satisfying the required conditions, provided we exclude the edge cases, i.e. we assume that $MFDR(\psi_{\mathcal{X}}) > \alpha$, and also $\mathbb{P}(\tau_K^*(X) \geq 1 - \alpha) > 0$.

Recall $MFDR(\psi_R) = \mathbb{E}[1 - \tau_K^*(X)|X \in R]$ (provided $\mathbb{P}(X \in R) > 0$), so $MFDR(\psi_R) = \alpha$ is equivalent to $\mathbb{E}[(1 - \tau_K^*(X) - \alpha)I_{\{X \in R\}}] = 0$. Let Q be the signed measure on the space $\mathcal{X}$ defined as $Q(R) = \mathbb{E}[(1 - \tau_K^*(X) - \alpha)I_{\{X \in R\}}]$.

Denote $\xi(x) = (1 - \tau_K^*(x) - \alpha) / (\sum_{k=1}^K \tau_k(x))$, and define the level sets for any $\lambda \geq 0$:

$$L(\lambda) = \{x \in \mathcal{X} : \xi(x) \leq \lambda\}; \quad L^<(\lambda) = \{x \in \mathcal{X} : \xi(x) < \lambda\}.$$

Observe that it holds $L^<(\lambda) \subseteq L(\lambda)$ for any λ , and $L(\lambda) = \bigcap_{\lambda' > \lambda} L^<(\lambda')$, $L^<(\lambda) = \bigcup_{\lambda' < \lambda} L^<(\lambda')$, so that

$$\limsup_{\lambda' \nearrow \lambda} Q(L^<(\lambda')) = Q(L^<(\lambda)) \leq Q(L(\lambda)) = \liminf_{\lambda' \searrow \lambda} Q(L^<(\lambda')). \quad (13)$$

Define

$$\lambda^* = \inf\{\lambda \geq 0 : Q(L^<(\lambda)) \geq 0\};$$

we take $\lambda^* = \infty$ if the above set is empty. By the definition of λ^* , it holds

$$Q(L^<(\lambda^*)) \leq 0 \leq Q(L(\lambda^*));$$

this is true by (13) if $0 < \lambda^* < \infty$; if $\lambda^* = \infty$, it holds as well because $Q(L(\infty)) = Q(\mathcal{X}) > 0$ since $MFDR(\psi_{\mathcal{X}}) > \alpha$; if $\lambda^* = 0$, it still holds because $Q(L^<(0)) \leq 0$ in general. Therefore, since Q is without atoms, there exists a set $R \subseteq \mathcal{X}$ such that $L^>(\lambda^*) \subseteq R \subseteq L(\lambda^*)$ with $Q(R) = 0$. This will imply $MFDR(\psi_R) = \alpha$, as required, provided $\mathbb{P}(X \in R) > 0$, which we still have to check to finish the proof.

Recall we assumed $\mathbb{P}(\tau_K^*(X) \geq 1 - \alpha) > 0$. Since $\{x \in \mathcal{X} : \tau_K^*(X) \geq 1 - \alpha\} = L(0)$, this means $\mathbb{P}(X \in L(0)) > 0$. If $\lambda^* > 0$, it holds $L(0) \subseteq R$ and we are finished. In the special case $\lambda^* = 0$, it holds both $Q(L(0)) \leq 0$ by definition of $L(0)$, and $Q(L(0)) \geq 0$ by definition of λ^* . Hence $Q(L(0)) = 0$ in this case, and we can take $R = L(0)$.

Appendix C: Illustration of different configurations of the simulation study

Figure 5: Three examples of simulated data, with different parameter values: an easy case (left) corresponding to $D = 3$ and $\sigma^2 = 0.5$, and intermediate case (center) corresponding to $D = 2$ and $\sigma^2 = 1$ and a difficult case (right) corresponding to $D = 0$ and $\sigma^2 = 2$. Colors correspond to class labels.

Appendix D: Analysis of simulated data when only some classes are of interest.

Here the same simulation setting is considered as in the main article, the only difference being that it is now assumed that only classes 1 and 3 are of interest. Figure 6 displays the performance of the different classification rules when the true parameters of the model are known. Figure 7 displays the same result when the parameters of the model are inferred from the data.

Figure 6: Performances of the MAP, optimal and $1 - \alpha$ thresholded classification rules in terms of estimated and true MFDR and MFNR. All rules are based on the true posterior probabilities. Only classes 1 and 3 are assumed to be of interest.

Figure 7: Same Figure as 6, except that posterior probabilities are inferred from the data.

Appendix E: Illustration of the dataset and model used in Berard et al. (2011).

Figure 8 (left) provides a graphical representation of the dataset. Each point corresponds to a probe, represented by its methylation signal in the leaf (x-axis) and in the seed (y-axis). One can observe four different clouds: the upper (respectively lower) cloud corresponds to probes that are over-methylated (resp. under methylated) in the seed compared to the leaf. The two other clouds are positioned on the first bisector and correspond to non methylated probes (signals close to 0 in the two organs) or identically methylated probes. The model is graphically represented on Figure 8 (right), where the four ellipses represent the (constrained) shapes of the covariance matrices of the four Gaussian components of the constrained Gaussian bivariate mixture.

Figure 8: **Left:** Methylation data for chromosome 4. Each point corresponds to a probe, represented by its methylation levels in the leaf (x-axis) and seed (y-axis) **Right:** Same graph with the four ellipses corresponding to the four Gaussian bivariate components.

Appendix F: Simulations based on a mixture of Student bivariate distributions

For this supplementary simulation study, datasets were simulated from a mixture of three bidimensional Student distributions. We recall that given a Gaussian vector $\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \Sigma)$, a location vector μ and $u \sim \chi_v^2$, the vector

$$\mathbf{x} = \frac{\mathbf{y}}{\sqrt{\frac{u}{v}}} + \mu$$

follows a multivariate Student distribution with parameters Σ, μ, v . For each of the three components in the mixture, we simulated 200 observations with the same locations $\mu_1 = (-1, 0), \mu_2 = (0, D)$ and $\mu_3 = (1, 0)$ as in our principal Gaussian simulations. We took the scale matrix Σ to be the diagonal matrix $\sigma'^2 I$ where σ' was chosen so that the covariance matrix $\frac{v}{v-2} \Sigma$ is the same as the one in the Gaussian

scenario of “intermediate” difficulty, i.e. $\frac{\nu}{\nu-2}\sigma^2 = 1$. We considered several degrees of freedom ν (5, 10, 20, 50), and the same values of D as in our Gaussian simulations (0, 1, 2, 3). We generated 100 simulations for each configuration.

Figure 9 shows the MFDR and MFNR of the considered rules using exact posterior probabilities in the case when all the classes are of interest, while Figure 10 shows error rates when the posterior probabilities are estimated using the EM algorithm for Gaussian mixtures implemented in `mclust`. Note that these results should be compared to the results shown in the central column of Figures 1 and 3 ($\sigma = 1$) in the main text.

Figure 9: Performances of the MAP, optimal and $1 - \alpha$ thresholded classification rules on mixture of three bivariate Student distributions. All rules are based on the true posterior probabilities. Columns correspond to different values of the degrees of freedom ν . All classes are considered of interest.

Figure 10: Same Figure as Figure 9, except that posterior probabilities are inferred from the data.