



HAL
open science

False clustering rate control in mixture models

Ariane Marandon, Tabea Rebafka, Etienne Roquain, Nataliya Sokolovska

► **To cite this version:**

Ariane Marandon, Tabea Rebafka, Etienne Roquain, Nataliya Sokolovska. False clustering rate control in mixture models. 2022. hal-03625469

HAL Id: hal-03625469

<https://hal.science/hal-03625469v1>

Preprint submitted on 30 Mar 2022

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.

False clustering rate control in mixture models

Ariane Marandon¹, Tabea Rebafka¹, Etienne Roquain¹, Nataliya Sokolovska²

¹ *Sorbonne Université, LPSM*

² *Sorbonne Université, NutriOmics*

February 28, 2022

Abstract

The clustering task consists in delivering labels to the members of a sample. For most data sets, some individuals are ambiguous and intrinsically difficult to attribute to one or another cluster. However, in practical applications, misclassifying individuals is potentially disastrous. To overcome this difficulty, the idea followed here is to classify only *a part* of the sample in order to obtain a small misclassification rate. This approach is well known in the supervised setting, and referred to as classification with an abstention option. The purpose of this paper is to revisit this approach in an unsupervised mixture-model framework. The problem is formalized in terms of controlling the false clustering rate (FCR) below a prescribed level α , while maximizing the number of classified items. New procedures are introduced and their behavior is shown to be close to the optimal one by establishing theoretical results and conducting numerical experiments. An application to breast cancer data illustrates the benefits of the new approach from a practical viewpoint.

1 Introduction

1.1 Background

Clustering is a standard statistical task that aims at grouping together individuals with similar features. However, it is common that data sets include ambiguous individuals that are inherently difficult to classify, which makes the clustering result potentially unreliable. To illustrate this point, consider a Gaussian mixture model with overlapping mixture components. Then it is difficult, or even impossible, to assign the correct cluster label to data points that fall in the overlap of those clusters, see Figure 1. Hence, when the overlap is large (Figure 1 panel (b)), the misclassification rate of a standard clustering method is inevitably elevated.

This issue is critical in applications where misclassifications come with a high cost for the user and should be avoided. This is for example the case for medical diagnosis, where an error can have severe consequences on the individual’s health. When there is too much uncertainty, a solution is to avoid classification for such individuals, and to adopt a wiser “abstention decision”, that leaves the door open for further medical exams.

In a supervised setting, classification with a reject (or abstention) option is a long-standing statistical paradigm, that can be traced back to Chow (1970), with more recent works including Herbei and Wegkamp (2006); Bartlett and Wegkamp (2008); Wegkamp and Yuan (2011), among others. In this line of research, rejection is accounted for by adding a term to the risk that penalizes any rejection (i.e., non classification). However, this approach does not provide a prescribed control of the classification error among the classified items (those that are not rejected).

More confidence in the labelling can be brought via conformal prediction, Vovk et al. (1999, 2005), see also Angelopoulos et al. (2021) for a review. This generic technique can be applied to supervised classification for producing confidence sets for the label vector. However, such inference does not incorporate the abstention rule. Very recently, Denis and Hebiri (2020)

proposed a conformal approach with a rejection option. The method consists in minimizing the probability of misclassification conditionally on the fact that it has been classified, while controlling the probability of being classified. This brings a safety guarantee on the label assigned to the classified items.

1.2 Aim and approach

The goal of the present work is to propose a labelling guarantee on the classified items in the vein of Denis and Hebiri (2020), but in the more challenging unsupervised setting, where no training set is available and data are assumed to be generated from a finite mixture model. This is achieved by the possibility to refuse to cluster ambiguous individuals and by using the false clustering rate (FCR), which is defined as the average proportion of misclassifications among the classified objects. Our procedures are devised to keep the FCR below some nominal level α , while classifying a maximum number of items.

It is important to understand the role of the nominal level α in our approach. It is chosen by the user and depends on their acceptance or tolerance for misclassified objects. Since the FCR is the misclassification risk that is allowed on the classified objects, the final interpretation of FCR control at level α is clear: if for instance α is chosen to be 5% and 100 items are finally chosen to be classified by the method, then the number of misclassified items is expected to be at most 5. This high interpretability is similar to the one of the false discovery rate (FDR) in multiple testing, which has known a great success in applications since its introduction by Benjamini and Hochberg (1995). This is a clear advantage of our approach for practical use compared to the methods with a rejection option that are based on a penalized risk.

In our framework, a procedure is composed of two intertwined decisions:

- a selection rule deciding which of the items are labelled;
- a clustering method inferring the labels for the selected items.

Both decisions heavily rely on the appropriate quantification of the uncertainty of the cluster labels. As such, our approach is model-based, and can be viewed as a method that thresholds the posterior probabilities of the cluster labels with a data-driven choice of the threshold. The performance of the method will depend on the quality of the estimates of these posterior probabilities in the mixture model.

The adaptive character of our method is illustrated in Figure 1: when the clusters are well separated (panel (a)), the new procedure only discards few items and provides a clustering close to the correct one. However, when the clusters are overlapping (panel (b)), to avoid a high misclassification error, the procedure discards most of the items and only provides few labels, for which the uncertainty is low. In both cases, the proportion of misclassified items among the selected ones is small and in particular close to the target level α (here 10%). Hence, by adapting the amount of labeled or discarded items, our method always delivers a reliable clustering result, inspite of the varying intrinsic difficulty of the clustering task.

1.3 Presentation of the results

Let us now describe in more details the main contributions of the paper.

- We introduce three new data-driven procedures that perform simultaneously selection and clustering: the plug-in procedure (illustrated in Figure 1) and two bootstrap procedures (parametric and non-parametric), see Section 3.2.

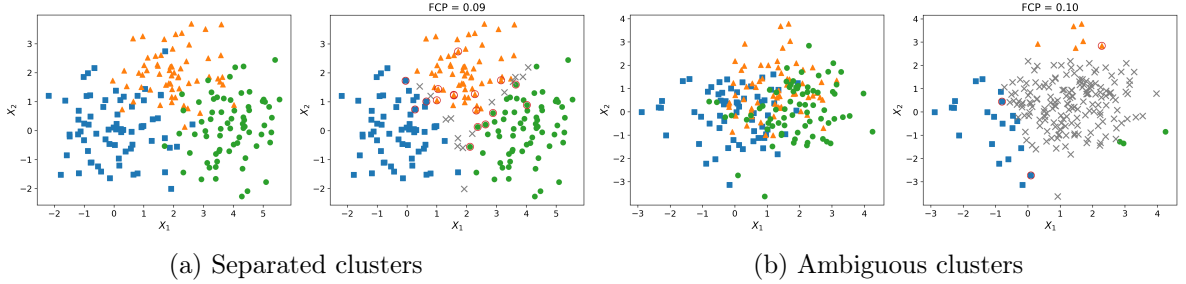


Figure 1: Data from Gaussian mixtures with three components ($n = 200$), in a fairly separated case (panel (a)) and an ambiguous case (panel (b)). In each panel, the left part displays the true clustering, while the right part illustrates the new procedure (plug-in procedure at level $\alpha = 10\%$), that does not cluster all items. The points not classified are depicted by grey crosses. Points circled in red correspond to erroneous labels.

- We prove that the plug-in procedure controls the FCR at the desired level α , up to a reminder term which becomes small when the sample size grows. In addition, this procedure is shown to satisfy the following optimality property: any other procedure that provides an FCR control necessarily classifies equal or less items than the plug-in procedure, up to a small remainder term. All results are non-asymptotic in the sample size (Theorem 6), and it is shown that the reminder terms tend to zero when the sample size tends to infinity (Corollary 7). We also put forward specific convergence rates depending on the parameter estimation accuracy (Corollary 8).
- Numerical experiments establish that the bootstrap procedures improve the plug-in procedure in terms of non-asymptotic FCR control, and thus are more reliable for practical use, where the sample size may be moderate, see Section 5.1.
- Our analysis also shows that the fixed threshold procedure that only labels items with a maximum posterior probability larger than $1 - \alpha$ is generally suboptimal for an FCR control at level α , see Section 5.1. To this extent, our procedures can be seen as more refined algorithms that classify more individuals while maintaining the FCR control.
- The practical impact of our approach is demonstrated on a real data set, see Section 5.2.

1.4 Relation to previous work

Other clustering guarantees in unsupervised learning While we provide a specific FCR control guarantee on the clustering, other criteria, not particularly linked to a rejection option, have been previously proposed in an unsupervised setting. Considering the cluster labels as fixed parameters, minimizing the misclassification risk has been proposed in Lei and Rinaldo (2015); Lu and Zhou (2016), among others. When clusters are rather considered as latent variables generated concomitantly with the observations (as we do here), other criteria have been investigated: the probability to make a different decision than the Bayes rule (Azizyan et al., 2013), the probability of exact recovery (Arora and Kannan, 2005; Bickel and Chen, 2009), or the fact that all clusters are mostly homogeneous with high probability (Najafi et al., 2020). All these works provide a guarantee only if the setting is favorable enough. By contrast, providing a rejection option is the key to obtain a guarantee in any setting (in the worst situation, the procedure will not classify any item).

Comparison to Denis and Hebiri (2020) and Mary-Huard et al. (2021) We describe here two recent studies that are related to ours, because they also use a FCR-like criterion. The first one is the work of Denis and Hebiri (2020), which also relies on a thresholding of the (estimated) posterior probabilities. However, the control is different, because it does not provide an FCR control, but rather a type-II error control concerning the probability of classifying an item. Also, the proposed procedure therein uses an additional labeled sample (semi-supervised setting), which is forbidden in our case.

The work of Mary-Huard et al. (2021) also proposes a control of the FCR. However, the analysis therein is solely based on the case where the model parameters are known (thus corresponding to the oracle case developed in Section 3.1 here). Compared to Mary-Huard et al. (2021), the present work provides number of new contributions, which are all given in Section 1.3. Let us also emphasize that we handle the label switching problem in the FCR, which seems to be overlooked in Mary-Huard et al. (2021).

Relation to false discovery rate The FCR is closely related to the false discovery rate (FDR), defined as the average proportion of errors among the discoveries, in a multiple testing context. In fact, we can roughly view the problem of designing an abstention rule as testing, for each item i , whether the clustering rule correctly classifies item i or not. With this analogy, our selection rule is based on quantities similar to the local FDR values (Efron et al., 2001), a key quantity to build optimal FDR controlling procedure in multiple testing mixture models, see, e.g., Storey (2003); Sun and Cai (2007); Cai et al. (2019); Rebafka et al. (2019). In particular, our final selection procedure shares similarities with the procedure introduced in Sun and Cai (2007), also named cumulative ℓ -value procedure Abraham et al. (2021). In addition, our theoretical analysis is related to the work of Rebafka et al. (2019), although the nature of the algorithm developed therein is different from here: they use the q -value procedure of Storey (2003), while our method rather relies on the cumulative ℓ -value procedure.

1.5 Organization of the paper

The paper is organized as follows: Section 2 introduces the model and relevant notation, namely the FCR criterion, with a particular care of the label switching problem. Section 3 presents the methods: the oracle, plug-in and the bootstrap approaches. Our main theoretical results are provided in Section 4, after introducing the appropriate assumptions. Section 5 presents numerical experiments and an application to a real data set, while a conclusion is given in Section 6. Proofs of the results and other technical details are deferred to appendices.

2 Setting

This section presents the notation, model, procedures and criteria that will be used throughout the manuscript.

2.1 Model

Let $\mathbf{X} = (X_1, \dots, X_n)$ be an observed random sample of size n . Each X_i is an i.i.d. copy of a d -dimensional real random vector, which is assumed to follow the standard mixture model:

$$\begin{aligned} Z &\sim \mathcal{M}(\pi_1, \dots, \pi_Q), \\ X|Z = q &\sim F_{\phi_q}, \quad 1 \leq q \leq Q, \end{aligned}$$

where $\mathcal{M}(\pi_1, \dots, \pi_Q)$ denotes the multinomial distribution of parameter π (equivalently, $\pi_q = \mathbb{P}(Z = q)$ for each q). The model parameters are given by

- the probability distribution π on $\{1, \dots, Q\}$ that is assumed to satisfy $\pi_q > 0$ for all q . Hence, π_q corresponds to the probability of being in class q ;
- the parameter $\phi = (\phi_1, \dots, \phi_Q) \in \mathcal{U}^Q$, where $\{F_u, u \in \mathcal{U}\}$ is a collection of distributions on \mathbb{R}^d . We moreover assume that each distribution F_u has a density with respect to the Lebesgue measure on \mathbb{R}^d , that we denote by f_u . Also, we assume that the ϕ_q 's are all distinct.

The parameter Q , coding for the number of classes, is assumed to be known and fixed throughout the manuscript (see Section 6 for a discussion). The overall parameter is thus $\theta = (\pi, \phi)$, the parameter set is denoted by Θ , and the distribution of (Z, X) is denoted by P_θ . The distribution family $\{P_\theta, \theta \in \Theta\}$ is the considered statistical model. We also assume that Θ is an open subset of \mathbb{R}^K for some $K \geq 1$ with the corresponding topology.

In this mixture model, the latent vector $\mathbf{Z} = (Z_1, \dots, Z_n)$ encodes a partition of the n observations into Q classes given by $\{1 \leq i \leq n : Z_i = q\}$, $1 \leq q \leq Q$. We refer to this model-based, random partition as the true latent clustering in the sequel.

In what follows, the “true” parameter that generates (Z, X) is assumed to be fixed and is denoted by $\theta^* \in \Theta$.

2.2 Procedure and criteria

Our approach starts with a given clustering rule, that aims at recovering the true latent clustering. In general, a clustering rule is defined as a (measurable) function of the observation \mathbf{X} returning a vector label $\hat{\mathbf{Z}} = (\hat{Z}_i)_{1 \leq i \leq n} \in \{1, \dots, Q\}^n$ for which the label q is given to individual i if and only if $\hat{Z}_i = q$.

The classification error of $\hat{\mathbf{Z}}$ is thus given by $\varepsilon(\hat{\mathbf{Z}}, \mathbf{Z}) = \sum_{i=1}^n \mathbb{1}\{Z_i \neq \hat{Z}_i\}$. Obviously, since switching the labels in \mathbf{Z} does not change the true latent partition, we should allow for switching the labels when making the decision $\hat{\mathbf{Z}}$ from the data. The clustering risk of $\hat{\mathbf{Z}}$ is thus defined by

$$R(\hat{\mathbf{Z}}) = \mathbb{E}_{\theta^*} \left(\min_{\sigma \in [Q]} \mathbb{E}_{\theta^*} \left(n^{-1} \varepsilon(\sigma(\hat{\mathbf{Z}}), \mathbf{Z}) \mid \mathbf{X} \right) \right), \quad (1)$$

where $[Q]$ denotes the set of all permutations on $\{1, \dots, Q\}$. Above, the minimum over σ is added to handle the aforementioned label switching problem.

Remark 1. *Note that the position of the minimum w.r.t. σ matters in the risk (1): σ is allowed to depend on X but not on Z . Hence, this risk has to be understood as being computed up to a data dependent label switching. This definition coincides with the usual definition of misclassification risk in the situation where the true clustering is deterministic, see Lei and Rinaldo (2015); Lu and Zhou (2016). Hence, it can be seen as a natural extension of the latter to a mixture model where the true clustering is random.*

Classically, we aim to find a clustering rule $\hat{\mathbf{Z}}$ such that the clustering risk is “small”. However, as mentioned above, whether this is possible or not depends on the intrinsic difficulty of the clustering problem and thus of the true parameter θ^* (see Figure 1). Therefore, the idea is to provide a selection rule, that is, a (measurable) function of the observation \mathbf{X} returning a subset $S \subset \{1, \dots, n\}$, such that the clustering risk *with restriction to S* is small. Throughout the paper, a *procedure* refers to a couple $\mathcal{C} = (\hat{\mathbf{Z}}, S)$, where $\hat{\mathbf{Z}}$ is a clustering rule and S is a selection rule.

Definition 1 (False clustering rate). *The false clustering rate (FCR) of a procedure $\mathcal{C} = (\widehat{\mathbf{Z}}, S)$ is given by*

$$\text{FCR}_{\theta^*}(\mathcal{C}) = \mathbb{E}_{\theta^*} \left(\min_{\sigma \in [Q]} \mathbb{E}_{\theta^*} \left(\frac{\varepsilon_S(\sigma(\widehat{\mathbf{Z}}), \mathbf{Z})}{\max(|S|, 1)} \mid \mathbf{X} \right) \right), \quad (2)$$

where $\varepsilon_S(\widehat{\mathbf{Z}}, \mathbf{Z}) = \sum_{i \in S} \mathbb{1}\{Z_i \neq \widehat{Z}_i\}$ denotes the misclassification error restricted to the subset S .

In this work, the aim is to find a procedure \mathcal{C} such that the false clustering rate is controlled at a nominal level α , that is, $\text{FCR}_{\theta^*}(\mathcal{C}) \leq \alpha$. Obviously, choosing S empty implies $\varepsilon_S(\sigma(\widehat{\mathbf{Z}}), \mathbf{Z}) = 0$ a.s. for any permutation σ and thus satisfies this control. Hence, while maintaining the control $\text{FCR}_{\theta^*}(\mathcal{C}) \leq \alpha$, we aim to classify as much individuals as possible, that is, to make $\mathbb{E}_{\theta^*}|S|$ as large as possible.

The FCR (2) involves an expectation of a ratio, which is more difficult to handle than a ratio of expectations. Hence, the following simpler alternative criterion will also be useful in our analysis.

Definition 2 (Marginal false clustering rate). *The marginal false clustering rate (mFCR) of a procedure $\mathcal{C} = (\widehat{\mathbf{Z}}, S)$ is given by*

$$\text{mFCR}_{\theta^*}(\mathcal{C}) = \frac{\mathbb{E}_{\theta^*} \left(\min_{\sigma \in [Q]} \mathbb{E}_{\theta^*} \left(\varepsilon_S(\sigma(\widehat{\mathbf{Z}}), \mathbf{Z}) \mid \mathbf{X} \right) \right)}{\mathbb{E}_{\theta^*}(|S|)}, \quad (3)$$

with the convention $0/0 = 0$.

Note that the mFCR is similar to the criterion introduced in Denis and Hebiri (2020) in the supervised setting.

2.3 Notation

We will extensively use the following notation: for all $q \in \{1, \dots, Q\}$ and $\theta = (\pi, \phi) \in \Theta$, we let

$$\ell_q(X, \theta) = \mathbb{P}_{\theta}(Z = q | X) = \frac{\pi_q f_{\phi_q}(X)}{\sum_{\ell=1}^Q \pi_{\ell} f_{\phi_{\ell}}(X)}; \quad (4)$$

$$T(X, \theta) = 1 - \max_{q \in \{1, \dots, Q\}} \ell_q(X, \theta) \in [0, 1 - 1/Q]. \quad (5)$$

Hence, $\ell_q(X, \theta)$ is the posterior probability of being in class q given the measurement X under the distribution P_{θ} . The quantity $T(X, \theta)$ is a measure of the risk when classifying X : it is close to 0 when there exists a class q such that $\ell_q(X, \theta)$ is close to 1, that is, when X can be classified with a large confidence.

3 Methods

In this section, we introduce new methods for controlling the FCR. We start by identifying an *oracle* method, in the sense that it uses the true value of the parameter θ^* . Substituting in that oracle the unknown parameter θ^* by an estimator provides our first method, called the *plug-in* procedure. We then define a refined version of the plug-in procedure, that accounts for the variability of the estimator and is based on a *bootstrap* approach.

3.1 Oracle procedures

MAP clustering Here, we proceed as if an oracle had given us the true value of θ^* and we introduce an oracle procedure $\mathcal{C}^* = (\widehat{\mathbf{Z}}^*, S^*)$ based on this value. As the following lemma shows, the best clustering rule is well-known and given by the Bayes clustering $\widehat{\mathbf{Z}}^* = (\widehat{Z}_1^*, \dots, \widehat{Z}_n^*)$ which can be written as

$$\widehat{Z}_i^* \in \operatorname{argmax}_{q \in \{1, \dots, Q\}} \ell_q(X_i, \theta^*), \quad i \in \{1, \dots, n\}, \quad (6)$$

where $\ell_q(\cdot)$ is the posterior probability given by (4).

Lemma 1. *We have $\min_{\widehat{\mathbf{Z}}} R(\widehat{\mathbf{Z}}) = R(\widehat{\mathbf{Z}}^*) = n^{-1} \sum_{i=1}^n \mathbb{E}_{\theta^*}(T_i^*)$, for the Bayes clustering $\widehat{\mathbf{Z}}^*$ defined by (6) and for*

$$T_i^* = T(X_i, \theta^*) = \mathbb{P}_{\theta^*}(Z_i \neq \widehat{Z}_i^* | X_i), \quad i \in \{1, \dots, n\}, \quad (7)$$

where $T(\cdot)$ is given by (5).

In words, Lemma 1 establishes that the oracle statistics T_i^* corresponds to the posterior misclassification probabilities of the Bayes clustering. To decrease the overall misclassification risk, it is natural to avoid classification for points with a high value of the test statistic T_i^* .

Thresholding selection rules In this section, we introduce the selection rule, which assesses which items have to be classified. From the above paragraph, it is natural to consider a thresholding based selection rule $S = \{i \in \{1, \dots, n\} : T_i^* \leq t\}$, for some threshold t to be suitably chosen. The following result gives insights for choosing such a threshold t .

Lemma 2. *For a procedure $\mathcal{C} = (\widehat{\mathbf{Z}}^*, S)$ with Bayes clustering and an arbitrary selection S ,*

$$\operatorname{FCR}_{\theta^*}(\mathcal{C}) = \mathbb{E}_{\theta^*} \left(\frac{\sum_{i \in S} T_i^*}{\max(|S|, 1)} \right). \quad (8)$$

As a consequence, a first way to build an (oracle) selection is to set

$$S = \{i \in \{1, \dots, n\} : T_i^* \leq \alpha\}.$$

Since an average of numbers smaller than α is also smaller than α , the corresponding procedure controls the FCR at level α . This procedure is referred to as the procedure *with fixed threshold* in the sequel. It corresponds to the following naive approach: to get a clustering with a risk of α , we only keep the items that are in their corresponding class with a posterior probability at least $1 - \alpha$. By contrast, the selection rule considered here is rather

$$S = \{i \in \{1, \dots, n\} : T_i^* \leq t(\alpha)\},$$

for a threshold $t(\alpha) \geq \alpha$ maximizing $|S|$ under the constraint $\sum_{i \in S} T_i^* \leq \alpha |S|$. It uniformly improves the procedure with fixed threshold and will in general lead to a (much) broader selection. This gives rise to the *oracle procedure*, that can be easily implemented by ordering the T_i^* 's, see Algorithm 1.

Algorithm 1 Oracle procedure

Input: Parameter θ^* , sample (X_1, \dots, X_n) , level α .

1. Compute the posterior probabilities $\mathbb{P}_{\theta^*}(Z_i = q|X_i)$, $1 \leq i \leq n$, $1 \leq q \leq Q$;
2. Compute the Bayes clustering \widehat{Z}_i^* , $1 \leq i \leq n$, given by (6);
3. Compute the probabilities T_i^* , $1 \leq i \leq n$, given by (7);
4. Order these probabilities in increasing order $T_{(1)}^* \leq \dots \leq T_{(n)}^*$;
5. Choose k^* the maximum of $k \in \{0, \dots, n\}$ such that $\max(k, 1)^{-1} \sum_{j=1}^k T_{(j)}^*(\mathbf{X}) \leq \alpha$;
6. Select S_α^* , the index corresponding to the k^* smallest elements among the T_i^* 's.

Output: Oracle procedure $\mathcal{C}_\alpha = (\widehat{\mathbf{Z}}^*, S_\alpha^*)$.

Algorithm 2 Plug-in procedure

Input: Sample (X_1, \dots, X_n) , level α .

1. Compute an estimator $\hat{\theta}$ of θ ;
2. Run the oracle procedure given in Algorithm 1 with $\hat{\theta}$ in place of θ^* .

Output: Plug-in procedure $\widehat{\mathcal{C}}_\alpha^{\text{PI}} = (\widehat{\mathbf{Z}}_\alpha^{\text{PI}}, \widehat{S}_\alpha^{\text{PI}})$.

3.2 Empirical procedures

Plug-in procedure The oracle procedure cannot be used in practice since θ^* is generally unknown. A natural idea is then to approach θ^* by an estimator $\hat{\theta}$ and then to plug this estimate into the oracle procedure. The resulting procedure, denoted $\widehat{\mathcal{C}}_\alpha^{\text{PI}} = (\widehat{\mathbf{Z}}_\alpha^{\text{PI}}, \widehat{S}_\alpha^{\text{PI}})$, is called the *plug-in procedure* and is implemented in Algorithm 2.

Obviously, the quality of the plug-in procedure relies on the chosen estimator $\hat{\theta}$. In a mixture model, the maximum likelihood estimator (MLE) is not computable in general. A standard way to approach the MLE is to use an Expectation Maximization (EM) algorithm (Dempster et al., 1977). The property of the EM estimator $\hat{\theta}$ in the current framework has been extensively studied both empirically and theoretically in the literature, see Balakrishnan et al. (2017) and references therein.

In Section 4, we establish that the plug-in procedure has suitable properties: when n tends to infinity, provided that the chosen estimator $\hat{\theta}$ behaves well and under mild regularity assumptions on the model, the FCR of the plug-in procedure is close to the level α while it is nearly optimal in terms of average selection number.

Bootstrap procedure Despite the favorable theoretical properties shown in Section 4, the plug-in procedure achieves an FCR that can exceed α in some situations, as we will see in our numerical experiments (Section 5). This is in particular the case when the estimator $\hat{\theta}$ is too rough. Indeed, the uncertainty of $\hat{\theta}$ near θ^* is ignored by the plug-in procedure.

To take into account this effect, we propose to use a bootstrap approach. It is based on the following result.

Lemma 3. For a given level $\alpha \in (0, 1)$, the FCR of the plug-in procedure $\widehat{\mathcal{C}}_\alpha^{\text{PI}}$ is given by

$$\text{FCR}(\widehat{\mathcal{C}}_\alpha^{\text{PI}}) = \mathbb{E}_{\mathbf{X} \sim P_{\theta^*}} \left(\min_{\sigma \in [Q]} \frac{\sum_{i=1}^n \{1 - \ell_{\sigma(\widehat{Z}_i^{\text{PI}}(\mathbf{X}))}(X_i, \theta^*)\} \mathbb{1}\{i \in \widehat{S}_\alpha^{\text{PI}}(\mathbf{X})\}}{\max(|\widehat{S}_\alpha^{\text{PI}}(\mathbf{X})|, 1)} \right). \quad (9)$$

The general idea is as follows: since $\text{FCR}(\widehat{\mathcal{C}}_\alpha^{\text{PI}})$ can exceed α , we choose α' as large as possible such that $\widehat{\text{FCR}}_{\alpha'} \leq \alpha$, for which $\widehat{\text{FCR}}_{\alpha'}$ is a bootstrap approximation of $\text{FCR}(\widehat{\mathcal{C}}_\alpha^{\text{PI}})$ based on (9).

Algorithm 3 Bootstrap procedure

Input: Sample (X_1, \dots, X_n) , level α , number B of bootstrap runs.

1. Choose a grid of increasing level $(\alpha(k))_{1 \leq k \leq K}$;
2. Compute $\widehat{\text{FCR}}_{\alpha(k)}^B$, $1 \leq k \leq K$, according to (10);
3. Choose \tilde{k} according to (11).

Output: Bootstrap procedure $\widehat{\mathcal{C}}_{\alpha}^{\text{boot}} = \widehat{\mathcal{C}}_{\alpha(\tilde{k})}^{\text{PI}}$.

The bootstrap approximation reads as follows: in the RHS of (9), we replace the true parameter θ^* by $\hat{\theta}$ and $\mathbf{X} \sim P_{\theta^*}$ by $\mathbf{X}' \sim \hat{P}$, where \hat{P} is an empirical substitute of P_{θ^*} . This empirical distribution \hat{P} is $P_{\hat{\theta}}$ for the parametric bootstrap and the uniform distribution over the X_i 's for the non-parametric bootstrap. This yields the bootstrap approximation of $\text{FCR}(\widehat{\mathcal{C}}_{\alpha}^{\text{PI}})$ given by

$$\widehat{\text{FCR}}_{\alpha} = \mathbb{E}_{\mathbf{X}' \sim \hat{P}} \left(\min_{\sigma \in [Q]} \frac{\sum_{i=1}^n \{1 - \ell_{\sigma(\hat{Z}_i^{\text{PI}}(\mathbf{X}'))}(X'_i, \hat{\theta}(\mathbf{X}))\} \mathbb{1}\{i \in \hat{S}_{\alpha}^{\text{PI}}(\mathbf{X}')\}}{\max(|\hat{S}_{\alpha}^{\text{PI}}(\mathbf{X}')|, 1)} \mid \mathbf{X} \right).$$

Classically, the latter is itself approximated by a Monte Carlo scheme:

$$\widehat{\text{FCR}}_{\alpha}^B = \frac{1}{B} \sum_{b=1}^B \min_{\sigma \in [Q]} \frac{\sum_{i=1}^n \{1 - \ell_{\sigma(\hat{Z}_i^{\text{PI}}(\mathbf{X}^b))}(X_i^b, \hat{\theta}(\mathbf{X}))\} \mathbb{1}\{i \in \hat{S}_{\alpha}^{\text{PI}}(\mathbf{X}^b)\}}{\max(|\hat{S}_{\alpha}^{\text{PI}}(\mathbf{X}^b)|, 1)}, \quad (10)$$

with $\mathbf{X}^1, \dots, \mathbf{X}^B$ i.i.d. $\sim \hat{P}$ corresponding to the bootstrap samples of \mathbf{X} .

Let $(\alpha(k))_{1 \leq k \leq K} \in (0, 1)^K$ a grid of increasing nominal levels (possibly in restriction to values slightly below the targeted level α). Then, the bootstrap procedure at level α is defined as $\widehat{\mathcal{C}}_{\alpha}^{\text{boot}} = \widehat{\mathcal{C}}_{\alpha(\tilde{k})}^{\text{PI}}$ where

$$\tilde{k} = \max \left\{ k \in \{1, \dots, K\} : \widehat{\text{FCR}}_{\alpha(k)}^B \leq \alpha \right\}. \quad (11)$$

This procedure is implemented in Algorithm 3.

Remark 2 (Parametric versus non parametric bootstrap). *The usual difference between parametric and non parametric bootstrap also holds in our context: the parametric bootstrap is fully based on $P_{\hat{\theta}}$, while the non parametric bootstrap builds an artificial sample (with replacement) from the original sample, which does not come from a P_{θ} -type distribution. This gives rise to different behaviors in practice: when $\hat{\theta}$ is too optimistic (which will be typically the case here when the estimation error is large), the correction brought by the parametric bootstrap (based on $P_{\hat{\theta}}$) is often weaker than that of the non parametric one. By contrast, when $\hat{\theta}$ is close to the true parameter, the parametric bootstrap approximation is more faithful because it uses the model, see Section 5.*

4 Theoretical guarantees for the plug-in procedure

In this section, we derive theoretical properties for the plug-in procedure: we show that its FCR and mFCR is close to α while its expected selection number is close to be optimal under some conditions.

4.1 Additional notation and assumptions

We make use of an optimality theory for mFCR control, that will be developed in detail in Section A.1. This approach extensively relies on the following quantities (recall the definition of $T(X, \theta)$ in (5)):

$$\text{mFCR}_t^* = \mathbb{E}_{\theta^*} (T(X, \theta^*) \mid T(X, \theta^*) < t); \quad (12)$$

$$t^*(\alpha) = \sup \{t \in [0, 1] : \text{mFCR}_t^* \leq \alpha\} \quad (13)$$

$$\alpha_c = \inf \{\text{mFCR}_t^* : t \in (0, 1], \text{mFCR}_t^* > 0\}; \quad (14)$$

$$\bar{\alpha} = \text{mFCR}_1^*. \quad (15)$$

In words, mFCR_t^* is the mFCR of an oracle procedure that selects the T_i^* smaller than a threshold t (Lemma 11). Then, $t^*(\alpha)$ is the optimal threshold such that this procedure has a mFCR controlled at level α . Next, α_c and $\bar{\alpha}$ are the lower and upper bounds for the level α , respectively, for which the optimality theory can be applied.

Now, we introduce our main assumption, which will be ubiquitous in our analysis.

Assumption 1. *For all $\theta \in \Theta$, under P_{θ^*} , the r.v. $T(X, \theta)$ given by (5) is continuous, that is, $t \mapsto \mathbb{P}_{\theta^*}(T(X, \theta) < t)$ is continuous on $[0, 1]$. In addition, it is increasing on $(\alpha_c, \bar{\alpha})$.*

Assumption 1 is useful to several regards: first, it prohibits ties in the $T(X_i, \theta)$'s, $1 \leq i \leq m$, so that the selection rule (see Algorithm 1) can be truly formulated as a thresholding rule. Second, it entails interesting properties for function $t \mapsto \text{mFCR}_t^*$, see Lemma 11 (this in particular ensures that the supremum in (13) is a maximum). Also note that the inequality $0 \leq \alpha_c < \bar{\alpha} < 1 - 1/Q$ holds under Assumption 1.

The next assumption ensures that the density family $\{f_u, u \in \mathcal{U}\}$ is smooth, and will be useful to establish consistency results.

Assumption 2. *For P_{θ^*} -almost all $x \in \mathbb{R}^d$, $u \in \mathcal{U} \mapsto f_u(x)$ is continuous.*

Let us now introduce some additional quantities measuring the regularity of the model and that will appear in our remainder terms (recall definition (4), (5) and (13) of $\ell_q(X, \theta)$, $T(X, \theta)$ and $t^*(\alpha)$, respectively): for $\epsilon, \delta, v > 0$,

$$\mathcal{W}_\ell(\epsilon) = \sup \left\{ \mathbb{E}_{\theta^*} \left[\max_{1 \leq q \leq Q} |\ell_q(X, \theta^*) - \ell_q(X, \theta)| \right], \|\theta - \theta^*\|_2 \leq \epsilon, \theta \in \Theta \right\}; \quad (16)$$

$$\mathcal{W}_T(\delta) = \sup \{ |\mathbb{P}_{\theta^*}(T(X, \theta^*) < t') - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t)|, t, t' \in [0, 1], |t' - t| \leq \delta \}; \quad (17)$$

$$\Psi(\epsilon) = \inf_{\delta \in (0, 1)} \{ \mathcal{W}_T(\delta) + \mathcal{W}_\ell(\epsilon)/\delta \}; \quad (18)$$

$$\mathcal{W}_{t^*, \alpha}(v) = \sup \{ |t^*(\alpha + \beta) - t^*(\alpha)|, |\beta| \leq v \}. \quad (19)$$

Note that $\Psi(\epsilon) \leq \mathcal{W}_T(\mathcal{W}_\ell(\epsilon)^{1/2}) + \mathcal{W}_\ell(\epsilon)^{1/2}$ for ϵ small enough. The following lemma is straightforward.

Lemma 4. *Under Assumption 1, we have $\lim_{\delta \rightarrow 0} \mathcal{W}_T(\delta) = 0$, $\lim_{v \rightarrow 0} \mathcal{W}_{t^*, \alpha}(v) = 0$. Under Assumption 2, we have $\lim_{\epsilon \rightarrow 0} \mathcal{W}_\ell(\epsilon) = 0$. Under both assumptions, we have $\lim_{\epsilon \rightarrow 0} \Psi(\epsilon) = 0$.*

We can derive rates for the convergences in Lemma 4 under the following additional regularity condition.

Assumption 3. *There exists positive constants $r = r(\theta^*)$, $C_1 = C_1(\theta^*)$, $C_2 = C_2(\theta^*, \alpha)$, $C_3 = C_3(\theta^*, \alpha)$ such that*

(i) for \mathbb{P}_{θ^*} -almost all $x, u \in \mathcal{U} \mapsto f_u(x)$ is continuously differentiable, and

$$\sum_{1 \leq q \leq Q} \mathbb{E}_{\theta^*} \left(\sup_{\substack{\theta \in \Theta \\ \|\theta - \theta^*\| \leq r}} \|\nabla_{\theta} \ell_q(X, \theta)\| \right) \leq C_1;$$

(ii) for all $t, t' \in [0, 1]$, $|\mathbb{P}_{\theta^*}(T(X, \theta^*) < t) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t')| \leq C_2|t - t'|$;

(iii) for all $\beta \in [(\alpha_c + \alpha)/2, (\alpha + \bar{\alpha})/2]$, $|t^*(\beta) - t^*(\alpha)| \leq C_3|\beta - \alpha|$.

Lemma 5. *Under Assumption 3, we have $\mathcal{W}_\ell(\epsilon) \leq C_1\epsilon$, $\mathcal{W}_T(\delta) \leq C_2\delta$, $\mathcal{W}_{t^*,\alpha}(v) \leq C_3v$ and $\Psi(\epsilon) \leq (1 + C_2)C_1^{1/2}\epsilon^{1/2}$ for ϵ, δ, v small enough.*

In Appendix D, it is proved that Assumptions 1, 2 and 3 hold true in the homoscedastic two-component multivariate Gaussian model, see Lemma 18.

Let us now discuss conditions on the estimator $\hat{\theta}$ on which the plug-in procedure is based. We start by introducing the following assumption (used in the concentration part of the proof, see Lemma 14).

Assumption 4. *The estimator $\hat{\theta}$ is assumed to be valued in a countable subset \mathcal{D} of Θ .*

This assumption is a minor restriction, because we can always choose $\mathcal{D} \subset \mathbb{Q}^K$ (recall $\Theta \subset \mathbb{R}^K$). Next, we additionally define a quantity measuring the quality of the estimator: for all $\epsilon > 0$,

$$\eta(\epsilon, \theta^*) = \mathbb{P}_{\theta^*} \left(\min_{\sigma \in [Q]} \|\hat{\theta}^\sigma - \theta^*\|_2 \geq \epsilon \right). \quad (20)$$

In the Gaussian case, the quantity $\eta(\epsilon, \theta^*)$ is provided to be small for a constrained MLE, see Proposition 16.

4.2 A general result

The following theorem is our main theoretical result: it provides non-asymptotical bounds on the mFCR/FCR of the plug-in procedure and on its average selection number.

Theorem 6. *Let Assumption 1 be true. For any $\alpha \in (\alpha_c, \bar{\alpha})$, for $s^* = s^*(\alpha, \theta^*) \in (0, 1)$ and $e^* = e(\alpha, \theta^*) > 0$ two constants only depending on α and θ^* , the following holds. Consider the plug-in procedure $\hat{\mathcal{C}}_\alpha^{PI} = (\hat{\mathbf{Z}}^{PI}, \hat{S}_\alpha^{PI})$ introduced in Algorithm 2 and based on an estimator $\hat{\theta}$ satisfying Assumption 4, with $\eta(\epsilon, \theta^*)$ defined by (20). Then for $\epsilon \leq e^*$ and $n \geq 4$, letting*

$$\Delta_n(\epsilon) = 2(\mathcal{W}_T(\mathcal{W}_{t^*,\alpha}(2\delta_n + 8\Psi(\epsilon)/s^*)) + 4\Psi(\epsilon) + 2\delta_n) + 50/n + 10\eta(\epsilon, \theta^*),$$

for $\delta_n = (2/s^*)\sqrt{(\log n)/n}$ and with the quantities \mathcal{W}_T , \mathcal{W}_ℓ , Ψ , $\mathcal{W}_{t^*,\alpha}$ defined by (17), (16), (18), (19), respectively, it holds:

- $\hat{\mathcal{C}}_\alpha^{PI}$ controls both the FCR and mFCR at level close to α in the following sense:

$$\max \left(\text{mFCR}(\hat{\mathcal{C}}_\alpha^{PI}), \text{FCR}(\hat{\mathcal{C}}_\alpha^{PI}) \right) \leq \alpha + \Delta_n(\epsilon)/s^*;$$

- $\widehat{\mathcal{C}}_\alpha^{PI}$ is nearly optimal in the following sense: for any other procedure $\mathcal{C} = (\widehat{\mathbf{Z}}, S)$ that controls the mFCR at level α , that is, with $\text{mFCR}(\mathcal{C}) \leq \alpha$, we have

$$n^{-1} \mathbb{E}_{\theta^*}(|\widehat{S}_\alpha^{PI}|) \geq n^{-1} \mathbb{E}_{\theta^*}(|S|) - \Delta_n(\epsilon).$$

The proof employs techniques that share similarities with the work of Rebafka et al. (2019) developed in a different context, see Appendix A. Here, a difficulty is to handle the new statistics $T(X_i, \hat{\theta})$ which is defined as an extremum, see (5).

4.3 Corollaries

In this section, we present two straightforward consequences of the main result. By Lemma 4, assuming Assumptions 1 and 2, we have that for all ϵ small enough, the limit of the remainder term $\limsup_n \Delta_n(\epsilon)$ can be bounded by a function of ϵ that vanishes when ϵ tends to 0. This yields the following consistency result.

Corollary 7 (Asymptotic optimality of the plug-in procedure). *Consider the setting of Theorem 6 where, in addition, Assumption 2 is supposed to hold. Then, we have*

$$\limsup_n \text{FCR}(\widehat{\mathcal{C}}_\alpha^{PI}) \leq \alpha, \quad \limsup_n \text{mFCR}(\widehat{\mathcal{C}}_\alpha^{PI}) \leq \alpha,$$

and for any procedure $\mathcal{C} = (\widehat{\mathbf{Z}}, S)$ that controls the mFCR at level α , we have

$$\liminf_n \{n^{-1} \mathbb{E}_{\theta^*}(|\widehat{S}_\alpha^{PI}|) - n^{-1} \mathbb{E}_{\theta^*}(|S|)\} \geq 0.$$

As shown by Lemma 5, by strengthening the assumptions, Theorem 6 delivers in addition convergence rates.

Corollary 8 (Optimality of the plug-in procedure with rates). *Consider the setting of Theorem 6, where in addition Assumption 3 holds. Then, with constants $C_1 = C_1(\alpha, \theta^*) > 0$ and $C_2 = C_2(\alpha, \theta^*) > 0$, we have for any sequence $\epsilon_n > 0$ tending to zero, for n larger than a constant only depending on α and θ^* ,*

$$\max(\text{FCR}(\widehat{\mathcal{C}}_\alpha^{PI}), \text{mFCR}(\widehat{\mathcal{C}}_\alpha^{PI})) \leq \alpha + C_1 \left(\epsilon_n^{1/2} + \sqrt{(\log n)/n} + \eta(\epsilon_n, \theta^*) \right) \quad (21)$$

$$n^{-1} \mathbb{E}_{\theta^*}(|\widehat{S}_\alpha^{PI}|) - n^{-1} \mathbb{E}_{\theta^*}(|S|) \geq -C_2 \left(\epsilon_n^{1/2} + \sqrt{(\log n)/n} + \eta(\epsilon_n, \theta^*) \right), \quad (22)$$

for any procedure $\mathcal{C} = (\widehat{\mathbf{Z}}, S)$ that controls (non asymptotically) the mFCR at level α .

The assumptions of Corollary 8 are all satisfied in the homoscedastic two-component multivariate Gaussian model. Applying this result in this case with $\epsilon_n = c\sqrt{(\log n)/n}$ (for some constant $c > 0$ not depending on n) and with $\hat{\theta}$ being a constrained MLE, Proposition 16 ensures that $\eta(\epsilon_n, \theta^*)$ is smaller than $1/n$. This indicates that the convergence rates in (21) and (22) are of order $((\log n)/n)^{1/4}$ in this case.

5 Experiments

In this section, we evaluate the behavior of the new procedures: plug-in (Algorithm 2), parametric bootstrap and non parametric bootstrap (Algorithm 3). For this, we use both synthetic and real data.

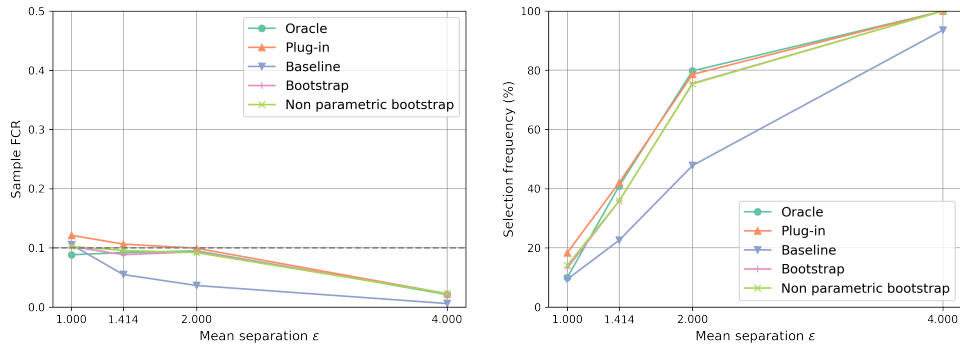


Figure 2: FCR (left panel) and selection frequency (right panel) as a function of the mean separation ϵ . Known mixture proportions and covariances setting with $Q = 2$, $n = 100$, $d = 2$, $\alpha = 0.1$.

5.1 Synthetic data set

The performance of our procedures is studied via simulations in different settings with various difficulties. All of them are gaussian mixture models, with possible restrictions on the parameter space. For parameter estimation, the classical EM algorithm is applied with 100 iterations and 10 starting points chosen with Kmeans++ (Arthur and Vassilvitskii, 2006). In the bootstrap procedures $B = 1000$ bootstrap samples are generated. The performance of all procedures is assessed via the *sample FCR* and the proportion of classified data points, which is referred to as the *selection frequency*. For every setting and every set of parameters, depicted results display the mean over 100 simulated datasets. As a baseline, we consider the fixed threshold procedure in which one selects data points that have a maximum posterior group membership probability that exceeds $1 - \alpha$. The oracle procedure (Algorithm 1) is also considered in our experiments for comparison.

Finally note that some figures are postponed in Appendix E for space reasons.

Known proportions and covariances In the first setting, the true mixture proportions and covariance matrices are known and used in the EM algorithm. We consider the case $Q = 2$, $\pi_1 = \pi_2 = 1/2$ and $\Sigma_1 = \Sigma_2 = I_d$ with I_d the $(d \times d)$ -identity matrix. For the mean vectors, we set $\mu_1 = 0$ and $\mu_2 = (\epsilon/\sqrt{d}, \dots, \epsilon/\sqrt{d})$. The quantity ϵ corresponds to the mean separation, that is, $\|\mu_1 - \mu_2\|_2 = \epsilon$ and accounts for the difficulty of the clustering problem.

Figure 2 displays the FCR for nominal level $\alpha = 0.1$, sample size $n = 100$, dimension $d = 2$ and varying mean separation $\epsilon \in \{1, \sqrt{2}, 2, 4\}$. Globally, our procedures all have an FCR close to the target level α (excepted for the very well separated case $\epsilon = 4$ for which the FCR is much smaller because a large part of the items can be trivially classified). In addition, the selection rate is always close to the one of the oracle procedure. On the other hand, the baseline procedure is too conservative: its FCR can be well below the nominal level and it selects up to 50% less than the other procedures. This is well expected, because unlike our procedures, the baseline has a fixed threshold and thus does not adapt to the difficulty of the problem.

We also note that the FCR of the plug-in approach is slightly inflated for a weak separation ($\epsilon = 1$). This comes from the parameter estimation, which is difficult in that case. This also illustrates the interest of the bootstrap methods, that allow to recover the correct level in that case, by appropriately correcting the plug-in approach.

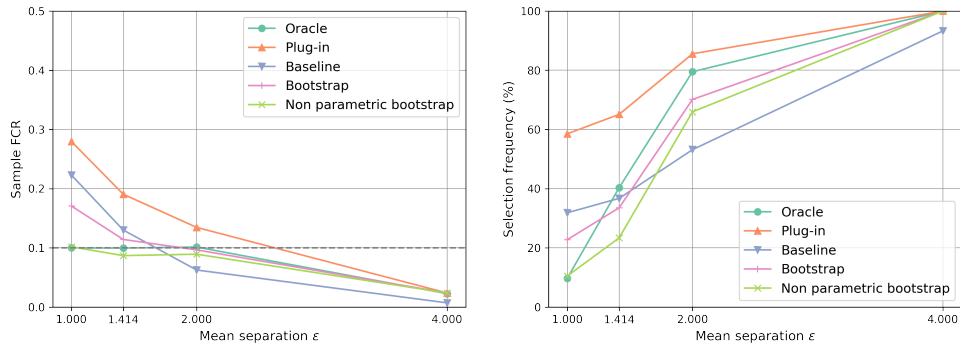


Figure 3: Same as Figure 2 in the diagonal covariance setting.

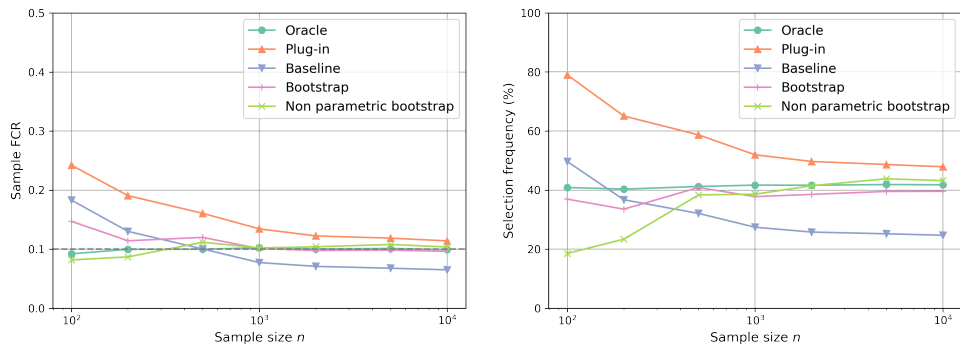
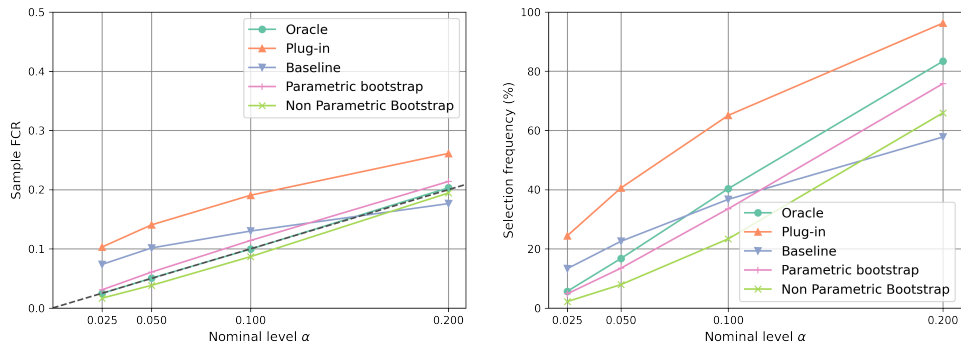


Figure 4: FCR (left panel) and selection frequency (right panel) as a function of the sample size n . Diagonal covariances setting with $Q = 2$, $d = 2$, $\epsilon = \sqrt{2}$, $\alpha = 10\%$.

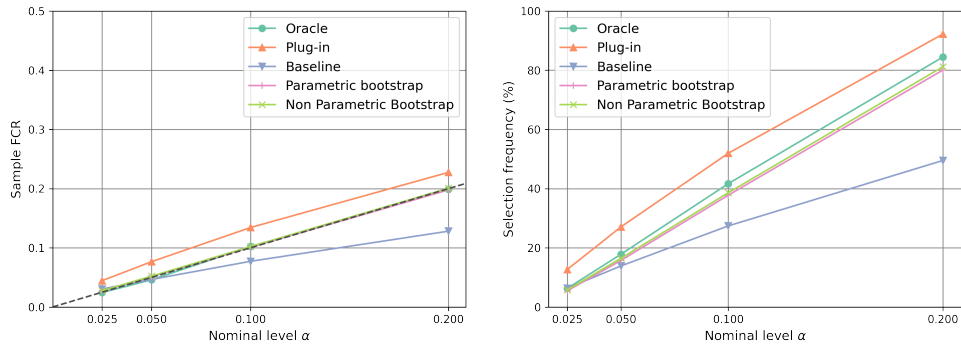
Diagonal covariances In this setting, the true parameters are the same as in previous paragraph, but the true mixture proportions and covariance matrices are unknown and not used in the EM algorithm. However, to help the estimation, the diagonal structure of Σ_1 and Σ_2 is used in this algorithm.

Figure 3 displays the FCR and the selection frequency in function of the separation ϵ for this case. The conclusion is qualitatively the same as in the previous case, but with larger FCR values for a weak separation. Overall, it shows that the plug-in procedure is anti-conservative and that the bootstrap corrections are able to recover an FCR/selection frequency close to the one of the oracle. However, for a weak separation $\epsilon = 1$, note that the parametric bootstrap correction is not enough and that the latter procedure still overshoots the level α . Indeed, in our simulations, it appears that $P_{\hat{\theta}}$ is typically a distribution that is more favorable than P_{θ^*} from a statistical point of view (for instance, with more separated clusters). Note that these conclusions also hold when making the FCR/selection frequency vary with n , see Figure 4.

Figure 5 displays the FCR and the selection frequency for varying nominal level α , with $\epsilon = \sqrt{2}$ and $n \in \{200, 1000\}$. The plug-in is still anti-conservative while the bootstrap procedures have an FCR that is close to α for both sample sizes and uniformly on the considered α range. Moreover, we note that for all our procedures (including the plug-in), the gap between the FCR and the nominal level is roughly constant with α : this illustrates the adaptive aspect of our procedures. This is in contrast with the baseline procedure, for which this gap highly depends on α , and which may be either anti-conservative or sub-optimal depending on the α value.



(a) $n = 200$



(b) $n = 1000$

Figure 5: FCR (left panel) and selection frequency (right panel) as a function of the nominal level α . Diagonal covariances setting with $Q = 2$, $d = 2$, $\epsilon = \sqrt{2}$, $n \in \{200, 1000\}$.

Other settings We also consider the following settings:

- Larger dimension (Figure 9). Same as in the previous paragraph but for $d = 20$.
- Three components (Figure 10): Same as in the previous paragraph but for $Q = 3$, $\pi_1 = \pi_2 = \pi_3 = 1/3$, $\Sigma_1 = \Sigma_2 = \Sigma_3 = I_2$, $\mu_1 = 0$, $\mu_2 = (\epsilon/\sqrt{2}, \epsilon/\sqrt{2})$, $\mu_3 = (0, \sqrt{\epsilon})$ (so that μ_3 is at equal L2-distances from μ_1 and μ_2).
- Unconstrained model (Figure 11): In this case, the EM algorithm is agnostic in the sense that does not use any knowledge, whatsoever, of the structure or values of the parameters. The true parameters are the same as the previous paragraph.
- An additional “typical” setting (Figure 6): an unconstrained model with three components $Q = 3$, dimension $d = 4$ and a separation $\epsilon = 2$.

In the three first settings, the difficulty of the estimation problem is increased with respect to the previous section. Hence, the behaviors the aforementioned procedures is deteriorated. In particular, it puts forward that the non parametric bootstrap can be anti-conservative (resp. over-conservative) in some cases when n is too large (resp. small). Finally, the last setting is intended to be a faithful summary of the procedure behaviors in a “typical” situation. For a fairly large sample size $n = 1000$, the new bootstrap procedures are correctly mimicking the oracle. The more challenging case of a small sample size $n = 100$ is handled differently by these two corrections: the parametric one is anti-conservative while the non-parametric is very conservative.

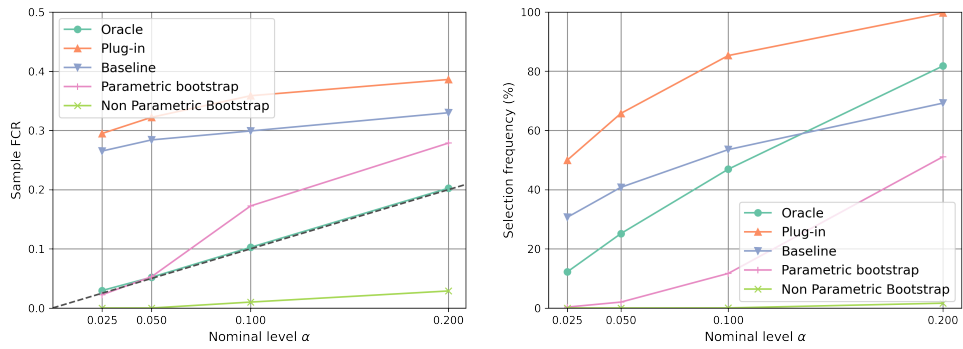
In conclusion, in our experiments, the parametric bootstrap procedure shows overall the more “stable” behavior: it uniformly improves the plug-in procedure across all the explored parameter ranges. In addition, it gets an FCR and a selection frequency close to the one of the oracle when the sample size n is fairly large. For more challenging cases where the sample size is small and we seek for a stricter FCR control, the non-parametric bootstrap is an interesting alternative.

5.2 Real data set

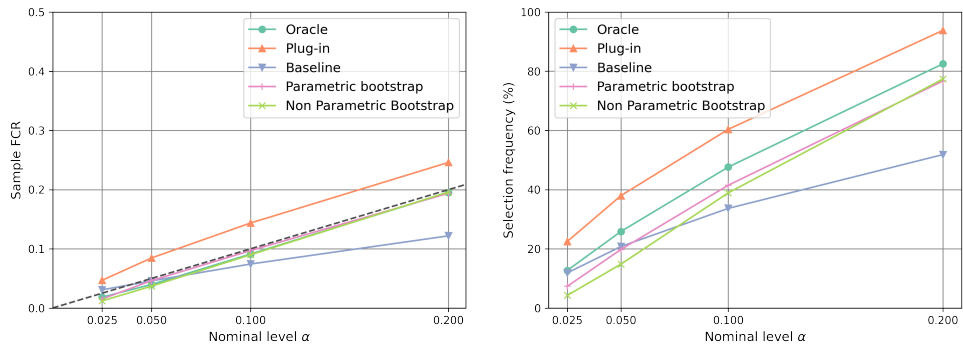
We consider the Wisconsin Breast Cancer Diagnosis (WCBBD) dataset from the UCI ML repository. The data consists of features computed from a digitalized image of a fine needle aspirate (FNA) of a breast mass, on a total of 569 patients (each corresponds to one FNA sample) of which 212 are diagnosed as Benign and 357 as Malignant. Ten real-valued measures were computed for each cell nucleus present in any one image (e.g. radius, perimeter, texture, etc.). Then, the mean, standard error and mean of the three largest values of these measures were computed for each image, resulting in a total of 30 features. Here, we restrict the analysis to the variables that correspond to the mean of these measures.

We choose to model the data as a mixture of Students as proposed in Peel and McLachlan (2000). Student mixtures are appropriate for data containing observations with longer than normal tails or atypical observations leading to overlapping clusters. Compared to Gaussian mixtures, Students are less concentrated and thus produce estimates of the posterior probabilities of class memberships that are less extreme, which is favorable for our selection procedures. In our study, the degree of freedom of each component is set to 4, and no constraints are put on the rest of the parameters. The t-mixture is fit via the EM algorithm provided by the Python package `studenttmixture` (Peel and McLachlan, 2000).

For illustration, Figure 7 (panel (a)) displays the first two variables of the dataset, the mean radius and the mean texture of the images. Different colors indicate the ground truth labels



(a) $n = 200$



(b) $n = 1000$

Figure 6: Same as Figure 5 with dimension $Q = 3$, $d = 4$, $\epsilon = 2$, for an unconstrained model.

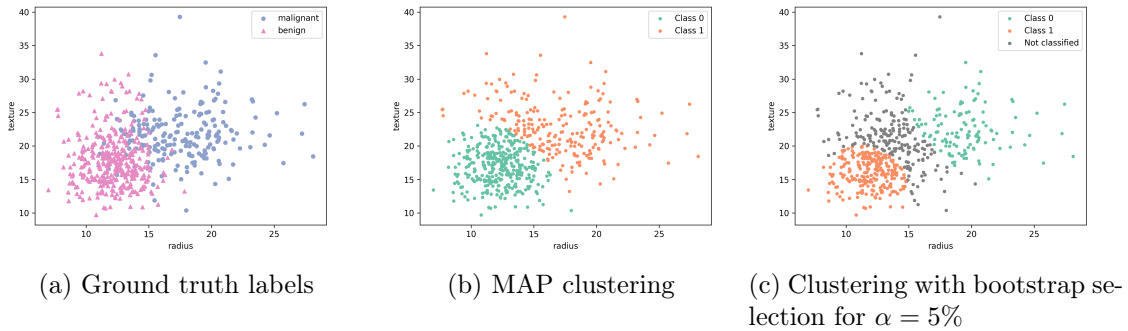


Figure 7: Comparison of the clustering result with t-mixture modelling with ground truth labels on the WCBD dataset, restricted to the variables *radius* and *texture*, with and without selection. With the parametric bootstrap procedure applied at $\alpha = 5\%$, the FCR w.r.t. the ground truth labels is of 3% versus 14% without selection.

(this information is not used in the clustering). One can see that the Student approximation is fairly good for each of the groups, and there is some overlap between them. Figure 7 (panel (b)) displays the MAP clustering result for the t-mixture model without any selection. The FCR is computed with respect to the ground truth labels and amounts to 14 %. Finally, Figure 7 (panel (c)) provides the result of our parametric bootstrap procedure with nominal level $\alpha = 5\%$. The procedure does not classify points that are at the intersection of the clusters and the FCR equals 3%, which is below the targeted level.

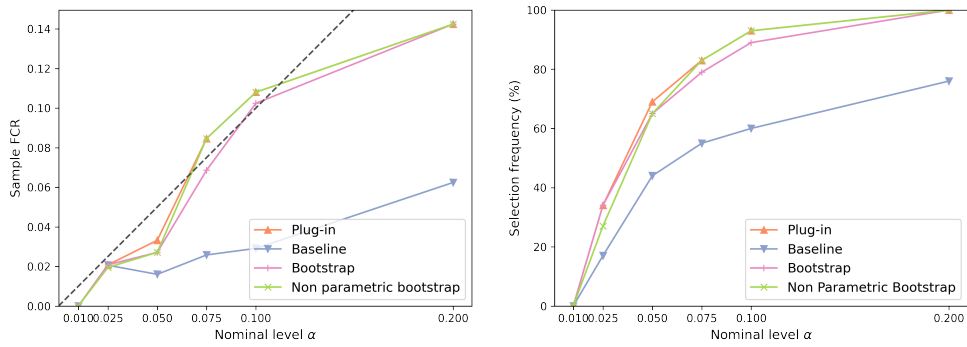
Figure 8 (panel (a)) provides a comparison of all our procedures and the baseline for varying nominal level α . The new procedures yield quite similar results. In particular, we see that the false clustering rate is always close to the nominal level, which illustrates once again the adaptive behavior of our procedures. By contrast, the false clustering rate of the baseline procedure is significantly below the nominal level and selects far less data points than our procedures, especially for large values of α . In addition, Figure 8 (panel (b)) displays results for the dataset where seven variables ($d = 7$) are analyzed and the conclusions are qualitatively the same.

6 Conclusion and discussion

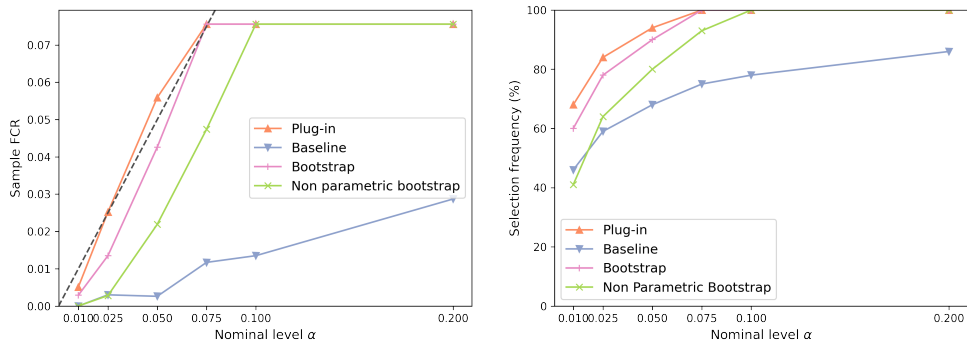
We have presented new data-driven methods providing both clustering and selection that ensure an FCR control guarantee in a mixture model. The plug-in approach was shown to be theoretically valid both when the parameter estimation is accurate and the sample size is large enough. When this is not necessarily the case, we proposed two second-order bootstrap corrections that have been shown to increase the FCR control ability on numerical experiments. Finally, applying our unsupervised methods to a supervised data set, our approach has been qualitatively validated by considering the attached labels as revealing the true clusters: our approach produces a selection/clustering close to what one would have done by observing the labels with some misclassification error range.

We would like to underline that the cluster number Q is assumed to be fixed and known throughout the study. In practice, it can be fitted from the data by using the standard AIC or BIC criteria. In addition, if several values of Q make sense from a practical viewpoint, we recommend to provide to the practitioner the collection of the corresponding outputs.

Concerning the pure task of controlling the FCR in the mixture model, our methods provide a correct FCR control in some area of the parameter space, leaving other less favorable



(a) Variables: radius, texture



(b) Variables: radius, texture, smoothness, compactness, concavity, symmetry, fractal dimension

Figure 8: FCR and selection frequency as a function of the nominal level α .

parameter configurations with a slight inflation in the FCR level. This phenomenon is well known for FDR control in the two group mixture multiple testing model (Sun and Cai, 2007; Roquain and Verzelen, 2019), and facing a similar problem in our framework is well expected. On the one hand, in some cases, this problem can certainly be solved by improving on parameter estimation: here the EM algorithm seems to over-estimate the extreme posterior probabilities, which makes the plug-in procedure too anti-conservative. On the other hand, it could be hopeless to expect a robust FCR control uniformly valid over all configurations, while being optimal in the favorable cases. To illustrate that point, we refer to the work Roquain and Verzelen (2019) that shows that such a procedure does not exist in the FDR controlling case, when the null distribution is Gaussian with an unknown scaling parameter (which is a framework sharing similarities with the one considered here). Investigating such a “lower bound” result in the current setting would provide better guidelines for the practitioner and is therefore an interesting direction for future research. In addition, in these unfavorable cases, adding labeled samples and considering a semi-supervised framework can be an appropriate alternative for practical use. This new sample is likely to considerably improve the inference. Studying the FCR control in that setting is another promising avenue.

Acknowledgments

This work has been supported by ANR-16-CE40-0019 (SansSouci), ANR-17-CE40-0001 (BASICS) and by the GDR ISIS through the “projets exploratoires” program (project TASTY). A. Marandon has been supported by a grant from Région Île-de-France (“DIM Math Innov”). We would like to thank Gilles Blanchard and Stéphane Robin for interesting discussions. We are also grateful to Eddie Aamari and Nhat Ho for their help for proving Lemma 17.

References

- Kweku Abraham, Ismael Castillo, and Etienne Roquain. Empirical bayes cumulative ℓ -value multiple testing procedure for sparse sequences. [arXiv preprint arXiv:2102.00929](https://arxiv.org/abs/2102.00929), 2021.
- Anastasios Nikolas Angelopoulos, Stephen Bates, Michael I. Jordan, and Jitendra Malik. Uncertainty sets for image classifiers using conformal prediction. In [9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021](https://openreview.net/forum?id=eNdiU_DbM9). OpenReview.net, 2021. URL https://openreview.net/forum?id=eNdiU_DbM9.
- Sanjeev Arora and Ravi Kannan. Learning mixtures of separated nonspherical Gaussians. [The Annals of Applied Probability](https://doi.org/10.1214/105051604000000512), 15(1A):69 – 92, 2005. doi: 10.1214/105051604000000512. URL <https://doi.org/10.1214/105051604000000512>.
- David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. Technical report, Stanford, 2006.
- Martin Azizyan, Aarti Singh, and Larry Wasserman. Minimax theory for high-dimensional gaussian mixtures with sparse mean separation. In [Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS’13](https://proceedings.neurips.cc/paper_files/paper/2013/file/Azizyan_Singh_Wasserman_NIPS13.pdf), page 2139–2147, Red Hook, NY, USA, 2013. Curran Associates Inc.
- Sivaraman Balakrishnan, Martin J. Wainwright, and Bin Yu. Statistical guarantees for the EM algorithm: From population to sample-based analysis. [The Annals of Statistics](https://doi.org/10.1214/16-AOS1435), 45(1): 77 – 120, 2017. doi: 10.1214/16-AOS1435. URL <https://doi.org/10.1214/16-AOS1435>.

- Peter L. Bartlett and Marten H. Wegkamp. Classification with a reject option using a hinge loss. Journal of Machine Learning Research, 9(59):1823–1840, 2008. URL <http://jmlr.org/papers/v9/bartlett08a.html>.
- Yoav Benjamini and Yoel Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. J. Roy. Statist. Soc. Ser. B, 57(1):289–300, 1995. ISSN 0035-9246.
- Peter J. Bickel and Aiyu Chen. A nonparametric view of network models and new-man–girvan and other modularities. Proceedings of the National Academy of Sciences, 106(50):21068–21073, 2009. ISSN 0027-8424. doi: 10.1073/pnas.0907096106. URL <https://www.pnas.org/content/106/50/21068>.
- T. Cai, Wenguang Sun, and Weinan Wang. Covariate-assisted ranking and screening for large-scale two-sample inference. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 81(2):187–234, 2019.
- Jiahua Chen. Consistency of the MLE under Mixture Models. Statistical Science, 32(1):47 – 63, 2017. doi: 10.1214/16-STS578. URL <https://doi.org/10.1214/16-STS578>.
- C. Chow. On optimum recognition error and reject tradeoff. IEEE Transactions on Information Theory, 16(1):41–46, 1970. doi: 10.1109/TIT.1970.1054406.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society. Series B (Methodological), 39(1):1–38, 1977.
- Christophe Denis and Mohamed Hebiri. Consistency of plug-in confidence sets for classification in semi-supervised learning. Journal of Nonparametric Statistics, 32(1):42–72, 2020. doi: 10.1080/10485252.2019.1689241. URL <https://doi.org/10.1080/10485252.2019.1689241>.
- Bradley Efron, Robert Tibshirani, John D Storey, and Virginia Tusher. Empirical bayes analysis of a microarray experiment. Journal of the American Statistical Association, 96 (456):1151–1160, 2001. doi: 10.1198/016214501753382129.
- Radu Herbei and Marten H. Wegkamp. Classification with reject option. The Canadian Journal of Statistics / La Revue Canadienne de Statistique, 34(4):709–721, 2006. ISSN 03195724. URL <http://www.jstor.org/stable/20445230>.
- Nhat Ho and XuanLong Nguyen. On strong identifiability and convergence rates of parameter estimation in finite mixtures. Electronic Journal of Statistics, 10(1):271 – 307, 2016. doi: 10.1214/16-EJS1105. URL <https://doi.org/10.1214/16-EJS1105>.
- Jing Lei and Alessandro Rinaldo. Consistency of spectral clustering in stochastic block models. The Annals of Statistics, 43(1):215 – 237, 2015. doi: 10.1214/14-AOS1274. URL <https://doi.org/10.1214/14-AOS1274>.
- Yu Lu and Harrison H Zhou. Statistical and computational guarantees of lloyd’s algorithm and its variants. arXiv preprint arXiv:1612.02099, 2016.
- Tristan Mary-Huard, Vittorio Perduca, Gilles Blanchard, and Martin-Magniette Marie-Laure. Error rate control for classification rules in multiclass mixture models, 2021.

- Pascal. Massart. Concentration Inequalities and Model Selection Ecole d'Été de Probabilités de Saint-Flour XXXIII - 2003. École d'Été de Probabilités de Saint-Flour, 1896. Springer Berlin Heidelberg, Berlin, Heidelberg, 1st ed. 2007. edition, 2007. ISBN 1-280-85333-6.
- Volodymyr Melnykov. On the distribution of posterior probabilities in finite mixture models with application in clustering. Journal of Multivariate Analysis, 122:175–189, 2013.
- Amir Najafi, Seyed Abolfazl Motahari, and Hamid R. Rabiee. Reliable clustering of Bernoulli mixture models. Bernoulli, 26(2):1535 – 1559, 2020. doi: 10.3150/19-BEJ1173. URL <https://doi.org/10.3150/19-BEJ1173>.
- D. Peel and G. J. McLachlan. Robust mixture modelling using the t distribution. Statistics and Computing, 10(4):339–348, 2000. doi: 10.1023/A:1008981510081. URL <https://doi.org/10.1023/A:1008981510081>.
- Tabea Rebařka, Etienne Roquain, and Fanny Villers. Graph inference with clustering and false discovery rate control, 2019.
- Etienne Roquain and Nicolas Verzelen. False discovery rate control with unknown null distribution: is it possible to mimic the oracle? arXiv preprint arXiv:1912.03109, 2019.
- John D Storey. The positive false discovery rate: a bayesian interpretation and the q-value. The Annals of Statistics, 31(6):2013–2035, 2003.
- Wenguang Sun and T. Tony Cai. Oracle and adaptive compound decision rules for false discovery rate control. Journal of the American Statistical Association, 102(479):901–912, 2007.
- Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. Algorithmic learning in a random world. Springer Science & Business Media, 2005.
- Volodya Vovk, Alexander Gammerman, and Craig Saunders. Machine-learning applications of algorithmic randomness. In Proceedings of the Sixteenth International Conference on Machine Learning, ICML '99, page 444–453, 1999.
- Marten Wegkamp and Ming Yuan. Support vector machines with a reject option. Bernoulli, 17(4):1368–1385, 2011.

A Proof of our main result

A.1 An optimal procedure

We consider in this section the procedure that serves as an optimal procedure in our theory. For $t \in [0, 1]$, let $\mathcal{C}_t^* = (\widehat{\mathbf{Z}}^*, S_t^*)$ be the procedure using the Bayes clustering $\widehat{\mathbf{Z}}^*$ (6) and the selection rule $S_t^* = \{i \in \{1, \dots, n\} : T_i^* < t\}$. Let us consider the map $t \in [0, 1] \mapsto \text{mFCR}(\mathcal{C}_t^*)$ and note that $\text{mFCR}(\mathcal{C}_t^*) = \text{mFCR}_t^*$ as defined by (12). Lemma 11 below provides the key properties for this function.

Definition 3. *The optimal procedure at level α is defined by $\mathcal{C}_{t^*(\alpha)}^*$ where $t^*(\alpha)$ is defined by (13).*

Note that the optimal procedure is not the same as the oracle procedure defined in Section 3.1, although these two procedures are expected to behave roughly in the same way (at least for a large n).

Under Assumption 1, Lemma 11 entails that, for $\alpha > \alpha_c$, $\text{mFCR}(\mathcal{C}_{t^*(\alpha)}^*) \leq \alpha$. Hence, $\mathcal{C}_{t^*(\alpha)}^*$ controls the mFCR at level α . In addition, it is optimal in the following sense: any other mFCR controlling procedure should select less items than $\mathcal{C}_{t^*(\alpha)}^*$.

Lemma 9 (Optimality of $\mathcal{C}_{t^*(\alpha)}^*$). *Let Assumption 1 be true and choose $\alpha \in (\alpha_c, \bar{\alpha}]$. Then the oracle procedure $\mathcal{C}_{t^*(\alpha)}^* = (\widehat{\mathbf{Z}}^*, S_{t^*(\alpha)}^*)$ satisfies the following:*

(i) $\text{mFCR}(\mathcal{C}_{t^*(\alpha)}^*) = \alpha$;

(ii) for any procedure $\mathcal{C} = (\widehat{\mathbf{Z}}, S)$ such that $\text{mFCR}(\mathcal{C}) \leq \alpha$, we have $\mathbb{E}_{\theta^*}(|S|) \leq \mathbb{E}_{\theta^*}(|S_{t^*(\alpha)}^*|)$.

A.2 Preliminary steps for proving Theorem 6

To keep the main proof concise, we need to define several additional notation. Let for $t \in [0, 1]$ and $\theta \in \Theta$ (recall (5))

$$\widehat{\mathbf{L}}_0(\theta, t) = \frac{1}{n} \sum_{i=1}^n T(X_i, \theta) \mathbb{1}_{T(X_i, \theta) < t};$$

$$\widehat{\mathbf{L}}_1(\theta, t) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{T(X_i, \theta) < t}.$$

Denote $\widehat{\mathbf{L}} = \widehat{\mathbf{L}}_0 / \widehat{\mathbf{L}}_1$, $\mathbf{L}_0 = \mathbb{E}_{\theta^*} \widehat{\mathbf{L}}_0$, $\mathbf{L}_1 = \mathbb{E}_{\theta^*} \widehat{\mathbf{L}}_1$, $\mathbf{L} = \mathbf{L}_0 / \mathbf{L}_1$ (with the convention $0/0 = 0$). Note that for any $\alpha > \alpha_c$, the mFCR of the optimal procedure $\mathcal{C}_{t^*(\alpha)}^*$ defined in Section A.1 is given by $\text{mFCR}(\mathcal{C}_{t^*(\alpha)}^*) = \mathbf{L}(\theta^*, t^*(\alpha)) = \alpha$.

Also, we denote from now on $\ell_{i,q}^* = \mathbb{P}_{\theta^*}(Z_i = q | X_i)$ for short and introduce for any parameter $\theta \in \Theta$ (recall (4) and (5))

$$\bar{q}(X_i, \theta) \in \operatorname{argmax}_{q \in \{1, \dots, Q\}} \ell_q(X_i, \theta), \quad 1 \leq i \leq n; \quad (23)$$

$$U(X_i, \theta) = 1 - \ell_{i, \bar{q}(X_i, \theta)}^*, \quad 1 \leq i \leq n; \quad (24)$$

$$\widehat{\mathbf{M}}_0(\theta, t) = \frac{1}{n} \sum_{i=1}^n U(X_i, \theta) \mathbb{1}_{T(X_i, \theta) < t}, \quad t \in [0, 1], \quad (25)$$

Note that $\widehat{\mathbf{M}}_0(\theta^*, t) = \widehat{\mathbf{L}}_0(\theta^*, t)$ but in general $\widehat{\mathbf{M}}_0(\theta, t)$ is different from $\widehat{\mathbf{L}}_0(\theta, t)$. We denote $\widehat{\mathbf{M}} = \widehat{\mathbf{M}}_0 / \widehat{\mathbf{L}}_1$, $\mathbf{M}_0 = \mathbb{E}_{\theta^*} \widehat{\mathbf{M}}_0$ and $\mathbf{M} = \mathbf{M}_0 / \mathbf{L}_1$ (with the convention $0/0 = 0$).

When $\alpha \in (\alpha_c, \bar{\alpha}]$ (recall (14) and (15)), we also let

$$s^* = s^*(\alpha, \theta^*) = n^{-1} \mathbb{E}_{\theta^*} \left(|S_{t^*(\frac{\alpha + \alpha_c}{2})}^*| \right) = \mathbf{L}_1(\theta^*, t^*((\alpha + \alpha_c)/2)) > 0. \quad (26)$$

We easily see that the latter is positive: if it was zero then $S_{t^*((\alpha + \alpha_c)/2)}^*$ would be empty which would entails that $\text{mFCR}(\mathcal{C}_{t^*((\alpha + \alpha_c)/2)}^*)$ is zero. This is excluded by definition (14) of α_c because $(\alpha + \alpha_c)/2 > \alpha_c$.

Also, we are going to extensively use the event

$$\Omega_\epsilon = \left\{ \min_{\sigma \in [Q]} \|\hat{\theta}^\sigma - \theta^*\|_2 < \epsilon \right\}.$$

On this event, we fix any permutation $\sigma \in [Q]$ (possibly depending on X) such that $\|\hat{\theta}^\sigma - \theta^*\|_2 < \epsilon$. Now using Lemma 12, the plug-in selection rule can be rewritten as $\widehat{S}_\alpha^{\text{PI}} = \{i \in \{1, \dots, n\} : \widehat{T}_i < \hat{t}(\alpha)\}$ (denoted by \widehat{S} in the sequel for short), where

$$\hat{t}(\alpha) = \sup\{t \in [0, 1] : \widehat{\mathbf{L}}(\hat{\theta}, t) \leq \alpha\}. \quad (27)$$

With the above notation, we can upper bound what is inside the brackets of $\text{FCR}(\widehat{\mathcal{C}}^{\text{PI}})$ and $\text{mFCR}(\widehat{\mathcal{C}}^{\text{PI}})$ as follows.

Lemma 10. *For the permutation σ in Ω_ϵ realizing $\|\hat{\theta}^\sigma - \theta^*\|_2 < \epsilon$, we have on the event Ω_ϵ the following relations:*

$$\begin{aligned} |\widehat{S}| &= \widehat{\mathbf{L}}_1(\hat{\theta}^\sigma, \hat{t}(\alpha)); \\ \min_{\sigma' \in [Q]} \mathbb{E}_{\theta^*} \left(\varepsilon_{\widehat{S}}(\sigma'(\widehat{\mathbf{Z}}), \mathbf{Z}) \mid \mathbf{X} \right) &\leq \widehat{\mathbf{M}}_0(\hat{\theta}^\sigma, \hat{t}(\alpha)); \\ \min_{\sigma' \in [Q]} \mathbb{E}_{\theta^*} \left(\frac{\varepsilon_{\widehat{S}}(\sigma'(\widehat{\mathbf{Z}}), \mathbf{Z})}{\max(|\widehat{S}|, 1)} \mid \mathbf{X} \right) &\leq \widehat{\mathbf{M}}(\hat{\theta}^\sigma, \hat{t}(\alpha)). \end{aligned}$$

Finally, we make use of the concentration of the empirical processes $\widehat{\mathbf{L}}_0(\theta, t)$, $\widehat{\mathbf{L}}_1(\theta, t)$, and $\widehat{\mathbf{M}}_0(\theta, t)$, uniformly with respect to $\theta \in \mathcal{D}$ (where \mathcal{D} is defined in Assumption 4). Thus, we define the following events, for $\delta, \epsilon > 0$ (recall s^* defined by (26)):

$$\begin{aligned} \Gamma_{0,\delta,t} &= \left\{ \sup_{\theta \in \mathcal{D}} \left| \widehat{\mathbf{L}}_0(\theta, t) - \mathbf{L}_0(\theta, t) \right| \leq \delta \right\}; \\ \Gamma_{1,\delta,t} &= \left\{ \sup_{\theta \in \mathcal{D}} \left| \widehat{\mathbf{L}}_1(\theta, t) - \mathbf{L}_1(\theta, t) \right| \leq \delta \right\}; \\ \Gamma_{\delta,t} &= \left\{ \sup_{\substack{\theta \in \mathcal{D}, \\ \mathbf{L}_1(\theta, t) \geq s^*}} \left| \widehat{\mathbf{L}}(\theta, t) - \mathbf{L}(\theta, t) \right| \leq \delta \right\}; \\ \Upsilon_{0,\delta,t} &= \left\{ \sup_{\theta \in \mathcal{D}} \left| \widehat{\mathbf{M}}_0(\theta, t) - \mathbf{M}_0(\theta, t) \right| \leq \delta \right\}. \end{aligned}$$

Note that the following holds:

$$\Gamma_{0,\delta s^*/2,t} \cap \Gamma_{1,\delta s^*/2,t} \subset \Gamma_{\delta,t}. \quad (28)$$

Indeed, on the event $\Gamma_{0,\delta s^*/2,t} \cap \Gamma_{1,\delta s^*/2,t}$, provided that $\mathbf{L}_1(\theta, t) \geq s^*$, we have

$$\begin{aligned} &\left| \frac{\widehat{\mathbf{L}}_0(\theta, t)}{\widehat{\mathbf{L}}_1(\theta, t)} - \frac{\mathbf{L}_0(\theta, t)}{\mathbf{L}_1(\theta, t)} \right| \\ &\leq \left| \frac{\mathbf{L}_0(\theta, t) - \widehat{\mathbf{L}}_0(\theta, t)}{\mathbf{L}_1(\theta, t)} \right| + \widehat{\mathbf{L}}_0(\theta, t) \left| \frac{1}{\widehat{\mathbf{L}}_1(\theta, t)} - \frac{1}{\mathbf{L}_1(\theta, t)} \right| \\ &\leq (\delta s^*/2)/s^* + (\delta s^*/2)/s^* = \delta, \end{aligned}$$

because $\widehat{\mathbf{L}}_0(\theta, t) \leq \widehat{\mathbf{L}}_1(\theta, t)$. This proves the desired inclusion.

A.3 Proof of Theorem 6

Let us now provide a proof for Theorem 6.

Step 1: bounding $\hat{t}(\alpha)$ w.r.t. $t^*(\alpha)$ Recall (13), (27) and (26). In this part, we only consider realizations on the event Ω_ϵ . Let $\beta \in [\frac{2\alpha+\alpha_c}{3}, \frac{\alpha+\bar{\alpha}}{2}]$. By Lemma 13, we have

$$\mathbf{L}_1(\hat{\theta}^\sigma, t^*(\beta)) \geq \mathbf{L}_1(\theta^*, t^*(\beta)) - \Psi(\|\hat{\theta}^\sigma - \theta^*\|_2) \geq \mathbf{L}_1(\theta^*, t^*((2\alpha + \alpha_c)/3)) - \Psi(\epsilon),$$

because $t^*(\beta) \geq t^*(\frac{2\alpha+\alpha_c}{3})$ since $t^*(\cdot)$ is non decreasing by Lemma 11. Hence $\mathbf{L}_1(\hat{\theta}^\sigma, t^*(\beta)) \geq s^*$ for ϵ smaller than a threshold only depending on θ^* and α . Hence, we have on $\Gamma_{\delta, t^*(\beta)}$ that

$$\mathbf{L}(\hat{\theta}^\sigma, t^*(\beta)) - \delta \leq \widehat{\mathbf{L}}(\hat{\theta}^\sigma, t^*(\beta)) \leq \delta + \mathbf{L}(\hat{\theta}^\sigma, t^*(\beta)).$$

By using again Lemma 13, we have

$$\mathbf{L}(\theta^*, t^*(\beta)) - 3\Psi(\epsilon)/s^* \leq \mathbf{L}(\hat{\theta}^\sigma, t^*(\beta)) \leq \mathbf{L}(\theta^*, t^*(\beta)) + 3\Psi(\epsilon)/s^*.$$

Given that $\mathbf{L}(\theta^*, t^*(\beta)) = \text{mFCR}(\mathcal{C}_{t^*(\beta)}^*) = \beta$ (see Lemma 9 (i)), it follows that for $\gamma = \gamma(\epsilon, \delta) = \delta + 4\Psi(\epsilon)/s^*$, on the event $\Gamma_{\delta, t^*(\alpha-\gamma)} \cap \Gamma_{\delta, t^*(\alpha+\gamma)}$,

$$\widehat{\mathbf{L}}(\hat{\theta}^\sigma, t^*(\alpha - \gamma)) \leq \alpha, \quad \widehat{\mathbf{L}}(\hat{\theta}^\sigma, t^*(\alpha + \gamma)) > \alpha,$$

where we indeed check that $\alpha - \gamma \geq \frac{2\alpha+\alpha_c}{3}$ and $\alpha + \gamma \leq \frac{\alpha+\bar{\alpha}}{2}$ for δ and ϵ smaller than some threshold only depending on θ^* and α . In a nutshell, we have established

$$\Gamma_{\delta, t^*(\alpha-\gamma)} \cap \Gamma_{\delta, t^*(\alpha+\gamma)} \cap \Omega_\epsilon \subset \{t^*(\alpha - \gamma) \leq \hat{t}(\alpha) \leq t^*(\alpha + \gamma)\}. \quad (29)$$

Step 2: upper-bounding the FCR Let us consider the event

$$\Lambda_{\alpha, \delta, \epsilon} := \Gamma_{0, \delta s^*/2, t^*(\alpha-\gamma)} \cap \Gamma_{1, \delta s^*/2, t^*(\alpha-\gamma)} \cap \Gamma_{0, \delta s^*/2, t^*(\alpha+\gamma)} \cap \Gamma_{1, \delta s^*/2, t^*(\alpha+\gamma)} \cap \Upsilon_{0, \delta, t^*(\alpha+\gamma)} \cap \Omega_\epsilon,$$

where the different events have been defined in the previous section.

Let us prove (21). By using Lemma 10 and (29),

$$\begin{aligned} \text{FCR}(\hat{\mathcal{C}}) &\leq \mathbb{E}_{\theta^*}[\widehat{\mathbf{M}}(\hat{\theta}^\sigma, \hat{t}(\alpha)) \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}}] + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c) \\ &\leq \mathbb{E}_{\theta^*} \left[\frac{\widehat{\mathbf{M}}_0(\hat{\theta}^\sigma, t^*(\alpha + \gamma))}{\widehat{\mathbf{L}}_1(\hat{\theta}^\sigma, t^*(\alpha - \gamma))} \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}} \right] + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c). \end{aligned}$$

Now using a concentration argument on the event $\Lambda_{\alpha, \delta, \epsilon} \subset \Gamma_{1, \delta, t^*(\alpha-\gamma)} \cap \Upsilon_{0, \delta, t^*(\alpha+\gamma)}$, we have

$$\begin{aligned} \text{FCR}(\hat{\mathcal{C}}) &\leq \mathbb{E}_{\theta^*} \left[\frac{\mathbf{M}_0(\hat{\theta}^\sigma, t^*(\alpha + \gamma)) + \delta}{\mathbf{L}_1(\hat{\theta}^\sigma, t^*(\alpha - \gamma)) - \delta} \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}} \right] + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c) \\ &\leq \frac{\mathbf{M}_0(\theta^*, t^*(\alpha + \gamma)) + 3\Psi(\epsilon) + \delta}{\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta} + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c) \\ &= \frac{\mathbf{L}_0(\theta^*, t^*(\alpha + \gamma)) + 3\Psi(\epsilon) + \delta}{\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta} + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c), \end{aligned} \quad (30)$$

by using Lemma 13 and that $\mathbf{M}_0(\theta^*, t) = \mathbf{L}_0(\theta^*, t)$ for all t by definition. Now, using again Lemma 13, we have

$$\begin{aligned} \mathbf{L}_0(\theta^*, t^*(\alpha + \gamma)) &\leq \mathbf{L}_0(\theta^*, t^*(\alpha - \gamma)) + \mathcal{W}_T(t^*(\alpha + \gamma) - t^*(\alpha - \gamma)) \\ &\leq \mathbf{L}_0(\theta^*, t^*(\alpha - \gamma)) + \mathcal{W}_T(\mathcal{W}_{t^*, \alpha}(2\gamma)) \end{aligned}$$

This entails

$$\begin{aligned} \text{FCR}(\hat{\mathcal{C}}) &\leq \frac{\mathbf{L}_0(\theta^*, t^*(\alpha - \gamma)) + \mathcal{W}_T(\mathcal{W}_{t^*, \alpha}(2\gamma)) + 3\Psi(\epsilon) + \delta}{\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta} + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c) \\ &\leq \frac{\mathbf{L}_0(\theta^*, t^*(\alpha - \gamma))}{\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta} + (s^*/2)^{-1} (\mathcal{W}_T(\mathcal{W}_{t^*, \alpha}(2\gamma)) + 3\Psi(\epsilon) + \delta) + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c), \end{aligned}$$

by choosing ϵ, δ smaller than a threshold (only depending on θ^* and α) so that $\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta \geq s^*/2$. Now using $\mathbf{L}_0(\theta^*, t^*(\alpha - \gamma)) = (\alpha - \gamma)\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma))$, we have

$$\begin{aligned} \frac{\mathbf{L}_0(\theta^*, t^*(\alpha - \gamma))}{\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta} &= (\alpha - \gamma) \left(1 + \frac{\Psi(\epsilon) + \delta}{\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta} \right) \\ &\leq \alpha (1 + (s^*/2)^{-1}(\Psi(\epsilon) + \delta)). \end{aligned}$$

This leads to

$$\text{FCR}(\hat{\mathcal{C}}) \leq \alpha + (2/s^*) (\mathcal{W}_T(\mathcal{W}_{t^*, \alpha}(2\delta + 8\Psi(\epsilon)/s^*)) + 4\Psi(\epsilon) + 2\delta) + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c),$$

which holds true for δ, ϵ smaller than a threshold only depending on θ^* and α .

Step 3: upper-bounding the mFCR We apply a similar technique as for step 2. By using Lemma 10 and (29),

$$\begin{aligned} \text{mFCR}(\hat{\mathcal{C}}) &\leq \frac{\mathbb{E}_{\theta^*}[\widehat{\mathbf{M}}_0(\hat{\theta}^\sigma, \hat{t}(\alpha)) \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}}] + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)}{\mathbb{E}_{\theta^*}[\widehat{\mathbf{L}}_1(\hat{\theta}^\sigma, \hat{t}(\alpha)) \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}}]} \\ &\leq \frac{\mathbb{E}_{\theta^*}[\widehat{\mathbf{M}}_0(\hat{\theta}^\sigma, t^*(\alpha + \gamma)) \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}}] + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)}{\mathbb{E}_{\theta^*}[\widehat{\mathbf{L}}_1(\hat{\theta}^\sigma, t^*(\alpha - \gamma)) \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}}]}. \end{aligned}$$

Now using a concentration argument on $\Lambda_{\alpha, \delta, \epsilon} \subset \Gamma_{1, \delta, t^*(\alpha - \gamma)} \cap \Upsilon_{0, \delta, t^*(\alpha + \gamma)}$, we have

$$\begin{aligned} \text{mFCR}(\hat{\mathcal{C}}) &\leq \frac{\mathbb{E}_{\theta^*}[(\mathbf{M}_0(\hat{\theta}^\sigma, t^*(\alpha + \gamma)) + \delta) \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}}] + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)}{\mathbb{E}_{\theta^*}[(\mathbf{L}_1(\hat{\theta}^\sigma, t^*(\alpha - \gamma)) - \delta) \mathbb{1}_{\Lambda_{\alpha, \delta, \epsilon}}]} \\ &\leq \frac{\mathbf{M}_0(\theta^*, t^*(\alpha + \gamma)) + 3\Psi(\epsilon) + \delta + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)}{\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta - \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)} \\ &= \frac{\mathbf{L}_0(\theta^*, t^*(\alpha + \gamma)) + 3\Psi(\epsilon) + \delta + \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)}{\mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta - \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)}, \end{aligned}$$

by using Lemma 13 and that $\mathbf{M}_0(\theta^*, t) = \mathbf{L}_0(\theta^*, t)$ by definition. Letting $x = \mathbf{L}_0(\theta^*, t^*(\alpha + \gamma)) + 3\Psi(\epsilon) + \delta$, $y = \mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta$ and $u = \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)$, we have obtained the bound $(x + u)/(y - u)$, which has to be compared with the FCR bound (30), which reads $x/y + u$. Now, when $y \in [0, 1]$, $x \geq 0$, $x/y \leq 2$, $u/y \leq 1/2$, $y - u \geq s^*/2$, we have

$$(x + u)/(y - u) \leq \frac{x/y}{1 - u/y} + (2/s^*)u \leq x/y(1 + 2u/y) + (2/s^*)u \leq x/y + (10/s^*)u.$$

As a result, for ϵ, δ small enough, and $\mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c) \leq s^*/4$, we obtain the same bound as for the FCR, with $\mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)$ replaced by $(10/s^*)\mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)$.

Step 4: lower-bounding the selection rate In Step 3, when bounding the mFCR, we derived a lower bound for the denominator of the mFCR, that is, $\mathbb{E}_{\theta^*}(|\hat{S}|)$. It reads

$$\begin{aligned} n^{-1} \mathbb{E}_{\theta^*}(|\hat{S}|) &\geq \mathbf{L}_1(\theta^*, t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta - \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c) \\ &\geq \mathbf{L}_1(\theta^*, t^*(\alpha)) - \mathcal{W}_T(t^*(\alpha) - t^*(\alpha - \gamma)) - \Psi(\epsilon) - \delta - \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c) \\ &\geq n^{-1} \mathbb{E}_{\theta^*}(|S_{t^*}^*(\alpha)|) - \mathcal{W}_T(\mathcal{W}_{t^*, \alpha}(\gamma)) - \Psi(\epsilon) - \delta - \mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c), \end{aligned}$$

by using (17) and (19). Now consider another procedure $\mathcal{C} = (\hat{\mathbf{Z}}, S)$ that controls the mFCR at level α , that is, $\text{mFCR}(\mathcal{C}) \leq \alpha$. By Lemma 9, we then have $\mathbb{E}_{\theta^*}(|S_{t^*}^*(\alpha)|) \geq \mathbb{E}_{\theta^*}(|S|)$.

Step 5: concentration Finally, we bound $\mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c)$ by using Lemma 14 with $u = \sqrt{(\log n)/n}$ and $x = 2\sqrt{(\log n)/n}$. This gives for $\delta = 2x/s^*$, and n such that $\log(n) \geq 1$,

$$\mathbb{P}((\Lambda_{\alpha, \delta, \epsilon})^c) \leq 5/n + \mathbb{P}(\Omega_\epsilon^c).$$

B Proofs of lemmas

Proof of Lemma 1 The clustering risk of $\hat{\mathbf{Z}}$ is given by

$$\begin{aligned} R(\hat{\mathbf{Z}}) &= \mathbb{E}_{\theta^*} \left(\min_{\sigma \in [Q]} \mathbb{E}_{\theta^*} \left(n^{-1} \sum_{i=1}^n \mathbb{1}\{Z_i \neq \sigma(\hat{Z}_i)\} \mid \mathbf{X} \right) \right) \\ &= \mathbb{E}_{\theta^*} \left(\min_{\sigma \in [Q]} n^{-1} \sum_{i=1}^n \mathbb{P}_{\theta^*}(Z_i \neq \sigma(\hat{Z}_i) \mid \mathbf{X}) \right) \\ &\geq \mathbb{E}_{\theta^*} \left(\min_{\hat{\mathbf{Z}}} n^{-1} \sum_{i=1}^n \mathbb{P}_{\theta^*}(Z_i \neq \hat{Z}_i \mid \mathbf{X}) \right), \end{aligned}$$

where, by independence, the minimum in the lower bound is achieved for the Bayes clustering. Thus, $R(\hat{\mathbf{Z}}) \geq n^{-1} \sum_{i=1}^n \mathbb{E}_{\theta^*}(T_i^*)$. Moreover, $n^{-1} \sum_{i=1}^n \mathbb{E}_{\theta^*}(T_i^*) \geq R(\hat{\mathbf{Z}}^*)$, since

$$R(\hat{\mathbf{Z}}^*) = \mathbb{E}_{\theta^*} \left(\min_{\sigma \in [Q]} n^{-1} \sum_{i=1}^n \mathbb{P}_{\theta^*}(Z_i \neq \sigma(\hat{Z}_i^*) \mid \mathbf{X}) \right) \leq \mathbb{E}_{\theta^*} \left(n^{-1} \sum_{i=1}^n \mathbb{P}_{\theta^*}(Z_i \neq \hat{Z}_i^* \mid \mathbf{X}) \right).$$

Thus, $\min_{\hat{\mathbf{Z}}} R(\hat{\mathbf{Z}}) = R(\hat{\mathbf{Z}}^*)$ and the proof is completed.

Proof of Lemma 2 Following the reasoning of the proof of Lemma 1 that, we have

$$\text{FCR}_{\theta^*}(\mathcal{C}) = \mathbb{E}_{\theta^*} \left(\min_{\sigma \in [Q]} \mathbb{E}_{\theta^*} \left(\frac{\sum_{i \in S} \mathbb{1}\{Z_i \neq \sigma(\hat{Z}_i^*)\}}{\max(|S|, 1)} \mid \mathbf{X} \right) \right) = \mathbb{E}_{\theta^*} \left(\frac{\sum_{i \in S} T_i^*}{\max(|S|, 1)} \right).$$

Proof of Lemma 3 By definition, we have

$$\text{FCR}(\hat{\mathcal{C}}_\alpha^{\text{PI}}) = \mathbb{E}_{\theta^*} \left(\min_{\sigma \in [Q]} \mathbb{E}_{\theta^*} \left(\frac{\sum_{i=1}^n \mathbb{1}_{Z_i \neq \sigma(\hat{Z}_i^{\text{PI}}(\mathbf{X}))} \mathbb{1}\{i \in \hat{S}^{\text{PI}}(\mathbf{X})\}}{\max(|\hat{S}^{\text{PI}}(\mathbf{X})|, 1)} \mid \mathbf{X} \right) \right)$$

so that (9) follows by a direct integration w.r.t. the latent variable Z .

Proof of Lemma 4 The only non-trivial fact is for $\mathcal{W}_{t^*,\alpha}(v)$. Assumption 1 and Lemma 11 provide that $t \mapsto \text{mFCR}_t^*$ is a one-to-one continuous increasing map from $(t^*(\alpha_c), t^*(\bar{\alpha}))$ to $(\alpha_c, \bar{\alpha})$. Hence, for $\alpha \in (\alpha_c, \bar{\alpha})$, $\beta \mapsto t^*(\alpha + \beta)$ is continuous in 0 and $\lim_{v \rightarrow 0} \mathcal{W}_{t^*,\alpha}(v) = 0$.

Proof of Lemma 9 By Lemma 11, we have that $\text{mFCR}(\mathcal{C}_t^*)$ is monotonous in t and continuous w.r.t. t on $(t^*(\alpha_c), 1]$, thus for $\alpha \in (\alpha_c, \bar{\alpha}]$, $\text{mFCR}(\mathcal{C}_{t^*(\alpha)}^*) = \alpha$ which gives (i). For (ii), let $\mathcal{C} = (\widehat{\mathbf{Z}}, S)$ be a procedure such that $\text{mFCR}(\mathcal{C}) \leq \alpha$. Let us consider the procedure \mathcal{C}' with the Bayes clustering $\widehat{\mathbf{Z}}^*$ and the same selection rule S . Since \mathcal{C}' is based on a Bayes clustering, by the same reasoning leading to $R(\widehat{\mathbf{Z}}^*) \leq R(\widehat{\mathbf{Z}})$ in Section 3.1, we have that $\text{mFCR}(\mathcal{C}') \leq \text{mFCR}(\mathcal{C}) \leq \alpha$ with

$$\text{mFCR}(\mathcal{C}') = \frac{\mathbb{E}_{\theta^*}(\sum_{i \in S} T_i^*)}{\mathbb{E}_{\theta^*}(|S|)}.$$

Hence,

$$\mathbb{E}_{\theta^*} \left(\sum_{i \in S} T_i^* \right) \leq \alpha \mathbb{E}_{\theta^*}(|S|). \quad (31)$$

Now we use an argument similar to the proof of Theorem 1 in Cai et al. (2019). By definition of $S_{t^*(\alpha)}^*$, we have that

$$\sum_{i=1}^n \left(\mathbb{1}_{i \in S_{t^*(\alpha)}^*}(\mathbf{X}) - \mathbb{1}_{i \in S}(\mathbf{X}) \right) (T_i^* - t^*(\alpha)) \leq 0$$

which we can rewrite as

$$\sum_{i=1}^n \left(\mathbb{1}_{i \in S_{t^*(\alpha)}^*}(\mathbf{X}) - \mathbb{1}_{i \in S}(\mathbf{X}) \right) (T_i^* - t^*(\alpha) + \alpha - \alpha) \leq 0$$

and so

$$\begin{aligned} \mathbb{E}_{\theta^*} \left(\sum_{i=1}^n \left(\mathbb{1}_{i \in S_{t^*(\alpha)}^*}(\mathbf{X}) - \mathbb{1}_{i \in S}(\mathbf{X}) \right) (T_i^* - \alpha) \right) &\leq (t^*(\alpha) - \alpha) \mathbb{E}_{\theta^*} \left(\sum_{i=1}^n \left(\mathbb{1}_{i \in S_{t^*(\alpha)}^*}(\mathbf{X}) - \mathbb{1}_{i \in S}(\mathbf{X}) \right) \right) \\ &= (t^*(\alpha) - \alpha) (\mathbb{E}_{\theta^*}(|S_{t^*(\alpha)}^*|) - \mathbb{E}_{\theta^*}(|S|)). \end{aligned}$$

On the other hand, $\text{mFCR}(\mathcal{C}_{t^*(\alpha)}^*) = \alpha$ together with (31) implies that

$$\begin{aligned} &\mathbb{E}_{\theta^*} \left(\sum_{i=1}^n \left(\mathbb{1}_{i \in S_{t^*(\alpha)}^*}(\mathbf{X}) - \mathbb{1}_{i \in S}(\mathbf{X}) \right) (T_i^* - \alpha) \right) \\ &= \mathbb{E}_{\theta^*} \left(\sum_{i \in S_{t^*(\alpha)}^*} T_i^* - \alpha |S_{t^*(\alpha)}^*| - \sum_{i \in S} T_i^* + \alpha |S| \right) \geq 0. \end{aligned}$$

Combining, the relations above provides

$$(t^*(\alpha) - \alpha) (\mathbb{E}_{\theta^*}(|S_{t^*(\alpha)}^*|) - \mathbb{E}_{\theta^*}(|S|)) \geq 0.$$

Finally, noting that $t^*(\alpha) - \alpha > 0$ since $\alpha = \text{mFCR}(\mathcal{C}_{t^*(\alpha)}^*) < t^*(\alpha)$ by (ii) Lemma 11, this gives $\mathbb{E}_{\theta^*}(|S_{t^*(\alpha)}^*|) - \mathbb{E}_{\theta^*}(|S|) \geq 0$ and concludes the proof.

Proof of Lemma 10 First, we have by definition $\ell_q(X_i, \theta^\sigma) = \ell_{\sigma(q)}(X_i, \theta)$ and thus $T(X_i, \hat{\theta}) = T(X_i, \hat{\theta}^\sigma)$ by taking the maximum over q . This gives $\hat{S}^\sigma = \hat{S}$ and yields the first equality. Next, we have on Ω_ϵ ,

$$\begin{aligned} \min_{\sigma' \in [Q]} \mathbb{E}_{\theta^*} \left(\varepsilon_{\hat{S}}(\sigma'(\hat{\mathbf{Z}}), \mathbf{Z}) \mid \mathbf{X} \right) &\leq \mathbb{E}_{\theta^*} \left(\varepsilon_{\hat{S}}(\sigma(\hat{\mathbf{Z}}), \mathbf{Z}) \mid \mathbf{X} \right) \\ &\leq \mathbb{E}_{\theta^*} \left(\varepsilon_{\hat{S}^\sigma}(\sigma(\hat{\mathbf{Z}}), \mathbf{Z}) \mid \mathbf{X} \right), \end{aligned}$$

still because $\hat{S}^\sigma = \hat{S}$. Now observe that,

$$\begin{aligned} \mathbb{E}_{\theta^*} \left(\varepsilon_{\hat{S}^\sigma}(\sigma(\hat{\mathbf{Z}}), \mathbf{Z}) \mid \mathbf{X} \right) &= \frac{1}{n} \sum_{i=1}^n \mathbb{P}_{\theta^*}(Z_i \neq \sigma(\bar{q}(X_i, \hat{\theta}^\sigma)) \mid \mathbf{X}) \mathbb{1}_{T(X_i, \hat{\theta}^\sigma) < \hat{t}(\alpha)} \\ &= \frac{1}{n} \sum_{i=1}^n (1 - \ell_{i, \sigma(\bar{q}(X_i, \hat{\theta}^\sigma))}^*) \mathbb{1}_{T(X_i, \hat{\theta}^\sigma) < \hat{t}(\alpha)} \\ &= \widehat{\mathbf{M}}_0(\hat{\theta}^\sigma, \hat{t}(\alpha)), \end{aligned}$$

because $\sigma(\bar{q}(X_i, \hat{\theta}^\sigma)) = \bar{q}(X_i, \hat{\theta}^\sigma)$. This proves the result.

C Auxiliary results

Lemma 11. *Let us consider the procedure \mathcal{C}_t^* defined in Section A.1 and the functional mFCR_t^* defined by (12). Then we have*

$$\text{mFCR}(\mathcal{C}_t^*) = \frac{\mathbb{E}_{\theta^*} \left(\sum_{i=1}^n T_i^* \mathbb{1}_{T_i^* < t} \right)}{\mathbb{E}_{\theta^*} \left(\sum_{i=1}^n \mathbb{1}_{T_i^* < t} \right)} = \text{mFCR}_t^*, \quad t \in [0, 1]. \quad (32)$$

Moreover, the following properties for the function $t \in [0, 1] \mapsto \text{mFCR}(\mathcal{C}_t^*)$:

- (i) $\text{mFCR}(\mathcal{C}_t^*)$ is non-decreasing in $t \in [0, 1]$ and, under Assumption 1, it is increasing in $t \in (t^*(\alpha_c), t^*(\bar{\alpha}))$;
- (ii) $\text{mFCR}(\mathcal{C}_t^*) < t$ for $t \in (0, 1]$;
- (iii) Under Assumption 1, $\text{mFCR}(\mathcal{C}_t^*)$ is continuous w.r.t. t on $(t^*(\alpha_c), 1]$, where $t^*(\alpha_c)$ is given by (14).

Proof. First, (32) is obtained similarly than (8). For proving (i), let $t_1, t_2 \in [0, 1]$ such that $t_1 < t_2$. We show that $\text{mFCR}(\mathcal{C}_{t_1}^*) \leq \text{mFCR}(\mathcal{C}_{t_2}^*)$. Remember here the convention $0/0 = 0$ and that $\text{mFCR}(\mathcal{C}_t^*) = \mathbb{E}_{\theta^*} (T(X, \theta^*) \mid T(X, \theta^*) < t)$. First, if $\mathbb{P}_{\theta^*} (T(X, \theta^*) < t_1) = 0$ then the result is immediate. Otherwise, we have that

$$\begin{aligned} &\text{mFCR}(\mathcal{C}_{t_1}^*) - \text{mFCR}(\mathcal{C}_{t_2}^*) \\ &= (\mathbb{P}_{\theta^*} (T(X, \theta^*) < t_1))^{-1} \mathbb{E}_{\theta^*} \left(T(X, \theta^*) \left\{ \mathbb{1}_{T(X, \theta^*) < t_1} - \frac{\mathbb{P}_{\theta^*} (T(X, \theta^*) < t_1)}{\mathbb{P}_{\theta^*} (T(X, \theta^*) < t_2)} \mathbb{1}_{T(X, \theta^*) < t_2} \right\} \right) \end{aligned}$$

where, given that $t_1 < t_2$, the quantity in the brackets is positive when $T(X, \theta^*) < t_1$ and is negative or zero otherwise. Hence,

$$\begin{aligned} &T(X, \theta^*) \left\{ \mathbb{1}_{T(X, \theta^*) < t_1} - \frac{\mathbb{P}_{\theta^*} (T(X, \theta^*) < t_1)}{\mathbb{P}_{\theta^*} (T(X, \theta^*) < t_2)} \mathbb{1}_{T(X, \theta^*) < t_2} \right\} \\ &\leq t_1 \left\{ \mathbb{1}_{T(X, \theta^*) < t_1} - \frac{\mathbb{P}_{\theta^*} (T(X, \theta^*) < t_1)}{\mathbb{P}_{\theta^*} (T(X, \theta^*) < t_2)} \mathbb{1}_{T(X, \theta^*) < t_2} \right\}. \end{aligned}$$

Taking the expectation makes the right-hand-side equal to zero, from which the result follows. Now, to show the increasingness, if $\text{mFCR}(\mathcal{C}_{t_1}^*) = \text{mFCR}(\mathcal{C}_{t_2}^*)$ for $t^*(\alpha_c) < t_1 < t_2 < t^*(\bar{\alpha})$, then the above reasoning shows that

$$(T(X, \theta^*) - t_1) \left\{ \mathbb{1}_{T(X, \theta^*) < t_1} - \frac{\mathbb{P}_{\theta^*}(T(X, \theta^*) < t_1)}{\mathbb{P}_{\theta^*}(T(X, \theta^*) < t_2)} \mathbb{1}_{T(X, \theta^*) < t_2} \right\} \leq 0$$

and has an expectation equal to 0. Hence, given that $T(X, \theta^*)$ is continuous, we derive that almost surely

$$\mathbb{P}_{\theta^*}(T(X, \theta^*) < t_2) \mathbb{1}_{T(X, \theta^*) < t_1} = \mathbb{P}_{\theta^*}(T(X, \theta^*) < t_1) \mathbb{1}_{T(X, \theta^*) < t_2},$$

that is, $\mathbb{P}_{\theta^*}(t_1 \leq T_i^* < t_2) = 0$, which is excluded by Assumption 1. This entails $\text{mFCR}(\mathcal{C}_{t_1}^*) < \text{mFCR}(\mathcal{C}_{t_2}^*)$.

For proving (ii), let $t > 0$. If $\mathbb{P}_{\theta^*}(T(X, \theta^*) < t) = 0$ then the result is immediate. Otherwise, we have that $\text{mFCR}(\mathcal{C}_t^*) - t = (\mathbb{P}_{\theta^*}(T(X, \theta^*) < t))^{-1} \mathbb{E}_{\theta^*}((T(X, \theta^*) - t) \mathbb{1}\{T(X, \theta^*) < t\})$. The latter is clearly not positive, and is moreover negative because $(T(X, \theta^*) - t) \mathbb{1}\{T(X, \theta^*) < t\} \leq 0$ and $\mathbb{P}_{\theta^*}(T(X, \theta^*) = t) = 0$ by Assumption 1.

For proving (iii), let $\psi_0(t) = \mathbb{E}_{\theta^*}(T(X, \theta^*) \mathbb{1}\{T(X, \theta^*) < t\})$ and $\psi_1(t) = \mathbb{P}_{\theta^*}(T(X, \theta^*) < t)$, the numerator and denominator of $\text{mFCR}(\mathcal{C}_t^*) = \text{mFCR}_t^*$, respectively. $\psi_1(t)$ is non-decreasing in t , with $\psi_1(0) = 0$ and $\psi_1(1) > 0$. Moreover, ψ_0 and ψ_1 are both continuous under Assumption 1. Then denote by t_c the largest t s.t. $\psi_1(t) = 0$. ψ_1 is zero on $[0, t_c]$ then strictly positive and non-decreasing on $(t_c, 1]$, and we have that $t_c = t^*(\alpha_c)$. Hence, $\text{mFCR}(\mathcal{C}_t^*)$ is zero on $[0, t_c]$ then strictly positive and continuous on $(t_c, 1]$. \square

Remark 3. *With the notation of the above proof, $t \mapsto \text{mFCR}(\mathcal{C}_t^*)$ may have a discontinuity point at t_c since for $t_n \xrightarrow[t_n > t_c]{} t_c$, as $\psi_1(t_n) \rightarrow 0$, one does not necessarily have that $\text{mFCR}(\mathcal{C}_{t_n}^*) \rightarrow 0$.*

Lemma 12 (Expression of plug-in procedure as a thresholding rule). *For any $\alpha \in (0, 1)$, let us consider the plug-in procedure $\hat{\mathcal{C}}_\alpha^{PI} = (\hat{\mathbf{Z}}_\alpha^{PI}, \hat{S}_\alpha^{PI})$ defined by Algorithm 2 and denote $K = |\hat{S}_\alpha^{PI}|$ the maximum of the $k \in \{0, \dots, n\}$ such that $\max(k, 1)^{-1} \sum_{j=1}^k \hat{T}_{(j)} \leq \alpha$ for $\hat{T}_i = 1 - \max_q \ell_q(X_i, \hat{\theta})$, $1 \leq i \leq n$. Consider also $\hat{t}(\alpha)$ defined by (27). Let Assumption 1 be true and consider an estimator $\hat{\theta}$ satisfying Assumption 4. Then it holds that $\hat{t}(\alpha) = \hat{T}_{(K+1)}$ and*

$$\hat{S}_\alpha^{PI} = \{i \in \{1, \dots, n\} : \hat{T}_i < \hat{t}(\alpha)\}.$$

Proof. If $\hat{T}_{(K)} < \hat{T}_{(K+1)}$ then the result is immediate. Thus it suffices to show that $\hat{T}_{(K)} = \hat{T}_{(K+1)}$ occurs with probability 0. From Assumption 4 (with the countable set \mathcal{D} defined therein), we have

$$\mathbb{P}_{\theta^*}(\hat{T}_{(K)} = \hat{T}_{(K+1)}) \leq \mathbb{P}_{\theta^*} \left(\bigcup_{i \neq j} \{\hat{T}_i = \hat{T}_j\} \right) \leq \mathbb{P}_{\theta^*} \left(\bigcup_{\theta \in \mathcal{D}} \bigcup_{i \neq j} \{T(X_i, \theta) = T(X_j, \theta)\} \right).$$

Now, the right term is a countable union of events which are all of null probability under Assumption 1. The result follows.

Lemma 13. *We have for all $\theta \in \Theta$,*

$$\sup_{t \in [0,1]} |\mathbf{L}_1(\theta, t) - \mathbf{L}_1(\theta^*, t)| \leq \Psi(\|\theta^* - \theta\|); \quad (33)$$

$$\sup_{t \in [0,1]} |\mathbf{L}_0(\theta, t) - \mathbf{L}_0(\theta^*, t)| \leq 2\Psi(\|\theta^* - \theta\|); \quad (34)$$

$$\sup_{t \in [t^*(\frac{\alpha + \alpha_c}{2}), 1]} |\mathbf{L}(\theta, t) - \mathbf{L}(\theta^*, t)| \leq 3\Psi(\|\theta^* - \theta\|)/s^*; \quad (35)$$

$$\sup_{t \in [0,1]} |\mathbf{M}_0(\theta, t) - \mathbf{M}_0(\theta^*, t)| \leq 3\Psi(\|\theta^* - \theta\|); \quad (36)$$

$$\sup_{t \in [t^*(\frac{\alpha + \alpha_c}{2}), 1]} |\mathbf{M}(\theta, t) - \mathbf{M}(\theta^*, t)| \leq 4\Psi(\|\theta^* - \theta\|)/s^*; \quad (37)$$

where $\alpha \in (\alpha_c, \bar{\alpha}]$ and $s^* > 0$ is given by (26). In addition, for all $\theta \in \Theta$ and $t, t' \in [0, 1]$,

$$|\mathbf{L}_0(\theta, t) - \mathbf{L}_0(\theta, t')| \leq 4\Psi(\|\theta^* - \theta\|) + \mathcal{W}_T(|t - t'|). \quad (38)$$

Proof. Fix $\theta \in \Theta$ and $t \in [0, 1]$. We have for any $\delta > 0$,

$$\begin{aligned} & |\mathbb{P}_{\theta^*}(T(X, \theta) < t) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t)| \\ & \leq (\mathbb{P}_{\theta^*}(T(X, \theta^*) < t + \delta) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t)) \vee (\mathbb{P}_{\theta^*}(T(X, \theta^*) < t) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t - \delta)) \\ & \quad + \mathbb{P}_{\theta^*}(|T(X, \theta^*) - T(X, \theta)| > \delta) \\ & \leq \mathcal{W}_T(\delta) + \mathbb{E}_{\theta^*}(|T(X, \theta^*) - T(X, \theta)|)/\delta. \end{aligned}$$

In addition, by definition (5),

$$\begin{aligned} |T(X, \theta^*) - T(X, \theta)| & \leq \left| \max_{1 \leq q \leq Q} \ell_q(X, \theta^*) - \max_{1 \leq q \leq Q} \ell_q(X, \theta) \right| \\ & \leq \max_{1 \leq q \leq Q} |\ell_q(X, \theta^*) - \ell_q(X, \theta)|. \end{aligned}$$

Hence,

$$|\mathbb{P}_{\theta^*}(T(X, \theta) < t) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t)| \leq \inf_{\delta \in (0,1)} \{\mathcal{W}_T(\delta) + \mathcal{W}_\ell(\|\theta^* - \theta\|)/\delta\} = \Psi(\|\theta^* - \theta\|),$$

which establishes (33).

Next, we have

$$\begin{aligned} & \mathbf{L}_0(\theta, t) - \mathbf{L}_0(\theta^*, t) \\ & = \mathbb{E}_{\theta^*}[T(X, \theta)(\mathbb{1}_{T(X, \theta) < t} - \mathbb{1}_{T(X, \theta^*) < t}) + \mathbb{1}_{T(X, \theta^*) < t}(T(X, \theta) - T(X, \theta^*))] \\ & \leq t|\mathbb{P}_{\theta^*}(T(X, \theta) < t) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t)| + |\mathbb{E}_{\theta^*}[\mathbb{1}_{T(X, \theta^*) < t}(T(X, \theta) - T(X, \theta^*))]| \\ & \leq |\mathbb{P}_{\theta^*}(T(X, \theta) < t) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t)| + \mathbb{E}_{\theta^*}|T(X, \theta) - T(X, \theta^*)| \leq 2\Psi(\|\theta^* - \theta\|) \end{aligned}$$

By exchanging the role of θ and θ^* in the above reasoning, the same bound holds for $\mathbf{L}_0(\theta^*, t) - \mathbf{L}_0(\theta, t)$, which gives (34). To prove (35), we use for any $t \in [t^*(\frac{\alpha + \alpha_c}{2}), 1]$,

$$\begin{aligned} & \left| \frac{\mathbf{L}_0(\theta, t) - \mathbf{L}_0(\theta^*, t)}{\mathbf{L}_1(\theta, t) - \mathbf{L}_1(\theta^*, t)} \right| \\ & \leq \left| \frac{\mathbf{L}_0(\theta, t) - \mathbf{L}_0(\theta^*, t)}{\mathbf{L}_1(\theta^*, t)} \right| + \mathbf{L}_0(\theta, t) \left| \frac{1}{\mathbf{L}_1(\theta^*, t)} - \frac{1}{\mathbf{L}_1(\theta, t)} \right| \\ & \leq 2\Psi(\|\theta^* - \theta\|)/s^* + \frac{1}{\mathbf{L}_1(\theta^*, t)} \frac{\mathbf{L}_0(\theta, t)}{\mathbf{L}_1(\theta, t)} |\mathbb{P}_{\theta^*}(T(X, \theta^*) < t) - \mathbb{P}_{\theta^*}(T(X, \theta) < t)| \\ & \leq 3\Psi(\|\theta^* - \theta\|)/s^*, \end{aligned}$$

because $\mathbf{L}_0(\theta, t) \leq \mathbf{L}_1(\theta, t)$ and $\mathbf{L}_1(\theta^*, t) \geq s^*$ by monotonicity. Similarly to the bound on \mathbf{L}_0 , we derive

$$\begin{aligned} & |\mathbf{M}_0(\theta, t) - \mathbf{M}_0(\theta^*, t)| \\ & \leq |\mathbb{P}_{\theta^*}(T(X, \theta) < t) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t)| + \mathbb{E}_{\theta^*} |U(X, \theta) - U(X, \theta^*)|. \end{aligned}$$

Define $\bar{q}(X, \theta) \in \operatorname{argmax}_{q \in \{1, \dots, Q\}} \ell_q(X, \theta)$. Now, since $U(X, \theta^*) \leq U(X, \theta)$ by definition (24), we have

$$\begin{aligned} \mathbb{E}_{\theta^*} |U(X, \theta) - U(X, \theta^*)| &= \mathbb{E}_{\theta^*} [U(X, \theta) - U(X, \theta^*)] \\ &= \mathbb{E}_{\theta^*} [\ell_{\bar{q}(X, \theta)}(X, \theta^*) - \ell_{\bar{q}(X, \theta^*)}(X, \theta^*)] \\ &= \mathbb{E}_{\theta^*} [\ell_{\bar{q}(X, \theta)}(X, \theta^*) - \ell_{\bar{q}(X, \theta)}(X, \theta) + \ell_{\bar{q}(X, \theta)}(X, \theta) - \ell_{\bar{q}(X, \theta^*)}(X, \theta^*)] \\ &\leq \mathbb{E}_{\theta^*} [\max_{1 \leq q \leq Q} |\ell_q(X, \theta^*) - \ell_q(X, \theta)|] + \mathbb{E}_{\theta^*} [\max_{1 \leq q \leq Q} \ell_q(X, \theta) - \max_{1 \leq q \leq Q} \ell_q(X, \theta^*)] \\ &\leq 2 \mathbb{E}_{\theta^*} [\max_{1 \leq q \leq Q} |\ell_q(X, \theta^*) - \ell_q(X, \theta)|] \leq 2\Psi(\|\theta^* - \theta\|). \end{aligned}$$

This proves (36) and leads to (37) by following the reasoning that provided (35). \square

Next, we have for $0 \leq t' \leq t \leq 1$, by (34),

$$|\mathbf{L}_0(\theta, t) - \mathbf{L}_0(\theta, t')| \leq |\mathbf{L}_0(\theta^*, t) - \mathbf{L}_0(\theta^*, t')| + 4\Psi(\|\theta^* - \theta\|).$$

Moreover,

$$\begin{aligned} |\mathbf{L}_0(\theta^*, t) - \mathbf{L}_0(\theta^*, t')| &= \mathbf{L}_0(\theta^*, t) - \mathbf{L}_0(\theta^*, t') = \mathbb{E}_{\theta^*} [T(X, \theta^*) \mathbb{1}_{t' \leq T(X, \theta^*) < t}] \\ &\leq \mathbb{E}_{\theta^*} [\mathbb{1}_{t' \leq T(X, \theta^*) < t}] = \mathbb{P}_{\theta^*}(T(X, \theta^*) < t) - \mathbb{P}_{\theta^*}(T(X, \theta^*) < t'), \end{aligned}$$

which is below $\mathcal{W}_T(t - t')$ by (17). This leads to (38).

Lemma 14 (Concentration of $\widehat{\mathbf{L}}_0$, $\widehat{\mathbf{L}}_1$, and $\widehat{\mathbf{M}}_0$). *For all $n \geq 1$, $x > 0$, $t \in [0, 1]$,*

$$\mathbb{P}_{\theta^*} \left(\sup_{\theta \in \mathcal{D}} \left| \widehat{\mathbf{L}}_0(\theta, t) - \mathbf{L}_0(\theta, t) \right| > x \right) \leq e^{-nu^2}; \quad (39)$$

$$\mathbb{P}_{\theta^*} \left(\sup_{\theta \in \mathcal{D}} \left| \widehat{\mathbf{L}}_1(\theta, t) - \mathbf{L}_1(\theta, t) \right| > x \right) \leq e^{-nu^2}; \quad (40)$$

$$\mathbb{P}_{\theta^*} \left(\sup_{\theta \in \mathcal{D}} \left| \widehat{\mathbf{M}}_0(\theta, t) - \mathbf{M}_0(\theta, t) \right| > x \right) \leq e^{-nu^2}, \quad (41)$$

where $u > 0$ is s.t. $x \geq \frac{2}{\sqrt{n}} + u$.

Proof. The concentration of each is a direct application of Talagrand's inequality. To be more precise, Theorem 5.3. in (Massart, 2007)) states the following. Let ξ_1, \dots, ξ_n independent r.v., \mathcal{F} a countable class of measurable functions s.t. $a \leq f \leq b$ for every $f \in \mathcal{F}$ for some real numbers a, b , and $W = \sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n f(\xi_i) - \mathbb{E}(f(\xi_i)) \right|$. Then, for any $x > 0$,

$$\mathbb{P}(W - \mathbb{E}(W) \geq x) \leq e^{-\frac{2x^2}{n^2(b-a)^2}}.$$

The result then follows by applying Lemma 15, which bounds the expectation term $\mathbb{E}(W)$ in each case by $2\sqrt{n}$. \square

Lemma 15. Let ξ_1, \dots, ξ_n be independent r.v. and \mathcal{F} be a countable set of measurable functions s.t. $\|f\|_\infty \leq 1$ for $f \in \mathcal{F}$. Then, one has

$$\mathbb{E} \left(\sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n f(\xi_i) - \mathbb{E}(f(\xi_i)) \right| \right) \leq 2\sqrt{n}.$$

Proof. We have that

$$\mathbb{E} \left(\sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n f(\xi_i) - \mathbb{E} \left(\sum_{i=1}^n f(\xi_i) \right) \right| \right) \leq \mathbb{E} \left(\sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n f(\xi_i) - f(\xi'_i) \right| \right),$$

where the ξ'_i s are independent r.v. with the same distribution as the ξ_i s. Let $(\epsilon_i)_{1 \leq i \leq n}$ be independent Rademacher variables. Then, we have

$$\sum_{i=1}^n f(\xi_i) - f(\xi'_i) \stackrel{D}{=} \sum_{i=1}^n \epsilon_i (f(\xi_i) - f(\xi'_i))$$

and

$$\sum_{i=1}^n f(\xi_i) - f(\xi'_i) = \mathbb{E} \left(\left| \sum_{i=1}^n \epsilon_i (f(\xi_i) - f(\xi'_i)) \right| \middle| \xi_1, \dots, \xi_n, \xi'_1, \dots, \xi'_n \right).$$

Thus,

$$\begin{aligned} \mathbb{E} \left(\sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n f(\xi_i) - \mathbb{E} \left(\sum_{i=1}^n f(\xi_i) \right) \right| \right) &\leq 2 \sup_{f \in \mathcal{F}} \mathbb{E} \left(\left| \sum_{i=1}^n \epsilon_i f(\xi_i) \right| \middle| \xi_1, \dots, \xi_n \right) \\ &\leq 2 \sup_{f \in \mathcal{F}} \sqrt{\mathbb{V} \left(\sum_{i=1}^n \epsilon_i f(\xi_i) \middle| \xi_1, \dots, \xi_n \right)} \\ &= 2 \sup_{f \in \mathcal{F}} \sqrt{\sum_{i=1}^n (f(\xi_i))^2} \\ &\leq 2\sqrt{n}. \end{aligned}$$

where the second inequality is due to Jensen's inequality. \square

D Auxiliary results for the Gaussian case

D.1 Convergence rate

The likelihood of Gaussian mixture models is unbounded in general (Chen, 2017). Thus, the convergence of MLEs has been studied over a constrained subspace of the parameter space that may expand as the sample size increases, which is often referred to as a sieve MLE.

Proposition 16. Consider the mixture model (Section 2.1) in the d -multivariate Gaussian case where $\phi_q = (\pi_q, \Sigma_q) \in \mathcal{U}$ with constrained parameter space defined as follows: $\mathcal{U} = [-a_n, a_n]^d \times \{\Sigma \in S_d^{++}, \underline{\lambda} \leq \lambda_1(\Sigma) \leq \lambda_2(\Sigma) \leq \bar{\lambda}\}$ where $a_n \leq L(\log n)^\gamma$ for some $L, \gamma > 0$, S_d^{++} denotes the space of positive definite matrices, with $\underline{\lambda}, \bar{\lambda}$ some strictly positive constants. Then the (constrained) MLE is such that the corresponding quantity $\eta(\cdot, \theta^*)$ defined by (20) satisfies $\eta(\delta_n, \theta^*) \leq 1/n$ for n large enough, where δ_n is a sufficiently large multiple of $\sqrt{\log n/n}$.

Proof. In the considered model, any mixture can be defined in terms of $\{f_u, u \in \mathcal{U}\}$ and a discrete mixing measure $G = \sum_{q=1}^Q \pi_q \delta_{\phi_q}$ with Q support points, as $\sum_{q=1}^Q \pi_q f_{\phi_q} = \int f_u(x) dG(u)$. As shown by Ho and Nguyen (2016), the convergence of mixture model parameters can be measured in terms of a Wasserstein distance on the space of mixing measures. Let $G_1 = \sum_{q=1}^Q \pi_q^1 \delta_{\phi_q^1}$ and $G_2 = \sum_{q=1}^Q \pi_q^2 \delta_{\phi_q^2}$ be two discrete probability measures on some parameter space, which is equipped with metric $\|\cdot\|$. The Wasserstein distance of order 1 between G_1 and G_2 is given by

$$W_1(G_1, G_2) = \inf_p \sum_{q,l} p_{q,l} \|\phi_q^1 - \phi_l^2\|$$

where the infimum is over all couplings $(p_{q,l})_{1 \leq q,l \leq Q} \in [0,1]^{Q \times Q}$ such that $\sum_l p_{q,l} = \pi_q^1$ and $\sum_q p_{q,l} = \pi_l^2$. Let G^*, \hat{G}_n denote the true mixing measure and the mixing measure that corresponds to the restricted MLE considered here, respectively. Theorem 4.2. in Ho and Nguyen (2016) implies that, with the notation of Ho and Nguyen (2016), for any $\epsilon_n \geq (\sqrt{C_1}/c)\delta_n$, and $\delta_n \leq C\sqrt{\log n/n}$, we have $\mathbb{P}_{\theta^*}(W_1(\hat{G}_n, G^*) \geq (c/C_1)\epsilon_n) \leq ce^{-n\epsilon_n^2}$. We apply this relation for $\epsilon_n = \max((\sqrt{C_1}/c)\delta_n, \sqrt{\log(cn)/n})$. In that case, we have still ϵ_n of order $\sqrt{\log n/n}$ and the upper-bound is at most $1/n$. On the other hand, if we have a convergence rate in terms of W_1 , then we have convergence of the mixture model parameters in terms of $\|\cdot\|$ at the same rate, see Lemma 17. This concludes the proof. \square

Lemma 17. *Let $G_n = \sum_{q=1}^Q \pi_q^n \delta_{\phi_q^n}$ be a sequence of discrete probability measures on \mathcal{U} , and let G^*, W_1 be defined as in the proof of Proposition 16. There exists a constant C only depending on G^* such that if $W_1(G_n, G^*) \rightarrow 0$, then for sufficiently large n ,*

$$W_1(G_n, G^*) \geq C \min_{\sigma \in [Q]} \|\theta_n^\sigma - \theta^*\|.$$

Proof. In what follows, we let $\{p_{q,l}^n\}$ denote the corresponding probabilities of the optimal coupling for the pair (G_n, G^*) . We start by showing that $(\phi_q^n)_q \rightarrow (\phi_q^*)_q$ in $\|\cdot\|$ up to a permutation of the labels. Let σ^n the permutation of the labels such that $\|\phi_q^n - \phi_l^*\| \geq \|\phi_{\sigma^n(l)}^n - \phi_l^*\|$ for all $q, l \in \{1, \dots, Q\}$. Then, by definition,

$$\begin{aligned} W_1(G_n, G^*) &\geq \sum_{1 \leq q, l \leq Q} p_{q,l}^n \|\phi_{\sigma^n(l)}^n - \phi_l^*\| \\ &= \sum_l \pi_l^* \|\phi_{\sigma^n(l)}^n - \phi_l^*\|. \end{aligned}$$

It follows that each $\|\phi_{\sigma^n(l)}^n - \phi_l^*\|$ must converge to zero. Since $(\phi_q^n)_q \rightarrow (\phi_q^*)_q$ up to a permutation of the labels, without loss of generality we can assume that $\phi_q^n \rightarrow \phi_q^*$ for all q . Let $\Delta\phi_q^n := \phi_q^n - \phi_q^*$ and $\Delta\pi_q^n := \pi_q^n - \pi_q^*$. Write $W_1(G_n, G^*)$ as

$$W_1(G_n, G^*) = \sum_q p_{qq}^n \|\Delta\phi_q^n\| + \sum_{q \neq l} p_{ql}^n \|\phi_q^n - \phi_l^*\|$$

Define $C_{ql} = \|\phi_q^* - \phi_l^*\|$ and $C = \min_{q \neq l} C_{ql} > 0$. It follows from the convergence of ϕ^n that for $q \neq l$, $\|\phi_q^n - \phi_l^*\| \geq C/2$ for sufficiently large n . Thus,

$$W_1(G_n, G^*) \geq \frac{C}{2} \sum_{q \neq l} p_{ql}^n$$

We deduce that $\sum_{q \neq l} p_{ql}^n \rightarrow 0$. As a result, $p_{qq}^n = \pi_q^* - \sum_{l \neq q} p_{lq}^n \rightarrow \pi_q^*$, and so, $p_{qq}^n \geq (1/2)\pi_{\min}^* := \min_l \pi_l^*$ for sufficiently large n . On the other hand, $\sum_{q \neq l} p_{ql}^n = \sum_q \pi_q^n - p_{qq}^n =$

$\sum_q \pi_q^* - p_{qq}^n$ where $p_{qq}^n \leq \min(\pi_q^n, \pi_q^*)$. Thus, $\sum_{q \neq l} p_{ql}^n \geq \sum_q \pi_q^n - \min(\pi_q^n, \pi_q^*) = \sum_{q, \pi_q^n \geq \pi_q^*} \pi_q^n - \pi_q^* = \sum_{q, \pi_q^n \geq \pi_q^*} |\pi_q^n - \pi_q^*|$ and similarly we have that $\sum_{q \neq l} p_{ql}^n \geq \sum_{q, \pi_q^* \geq \pi_q^n} |\pi_q^n - \pi_q^*|$. It follows that $2 \sum_{q \neq l} p_{ql}^n \geq \sum_q |\pi_q^n - \pi_q^*|$. Therefore, for sufficiently large n ,

$$W_1(G_n, G^*) \geq \frac{1}{2} \pi_{\min}^* \sum_q \|\Delta \phi_q^n\| + \frac{C}{4} \sum_q |\Delta \pi_q^n|.$$

This gives the result. \square

D.2 Gaussian computations

The following lemma holds.

Lemma 18. *Let us consider the multivariate Gaussian case where $\phi_q = (\mu_q, \Sigma_q)$, $1 \leq q \leq Q$, with $Q = 2$, $\Sigma_1 = \Sigma_2$ is an invertible covariance matrix and μ_1 and μ_2 are two different vectors of \mathbb{R}^d . Then Assumptions 1, 2 and 3 hold true for $\alpha_c = 0$ and for a level $\alpha \in (0, \bar{\alpha}) \setminus \mathcal{E}$ for \mathcal{E} a set of Lebesgue measure 0.*

Proof. Let us first prove that $T(X, \theta)$ is a continuous random variable under \mathbb{P}_{θ^*} (this is established below without assuming $\Sigma_1 = \Sigma_2$ for the sake of generality). We have

$$\begin{aligned} \mathbb{P}_{\theta^*}(T(X, \theta) = t) &= \mathbb{P}_{\theta^*} \left(1 - \max_{1 \leq q \leq Q} \left\{ \frac{\pi_q f_{\phi_q}(X)}{\sum_{\ell=1}^Q \pi_\ell f_{\phi_\ell}(X)} \right\} = t \right) \\ &= \mathbb{P}_{\theta^*} \left(\max_{1 \leq q \leq Q} \{ \pi_q f_{\phi_q}(X) \} = (1-t) \sum_{\ell=1}^Q \pi_\ell f_{\phi_\ell}(X) \right) \\ &\leq \sum_{q=1}^Q \mathbb{P}_{\theta^*} \left(\pi_q f_{\phi_q}(X) = (1-t) \sum_{\ell=1}^Q \pi_\ell f_{\phi_\ell}(X) \right). \end{aligned}$$

Since $Q = 2$, the latter probability is (for say $q = 1$), equal to

$$\begin{aligned} &\mathbb{P}_{\theta^*}(\pi_1 f_{\phi_1}(X) = (1/t - 1) \pi_2 f_{\phi_2}(X)) \\ &= \mathbb{P}_{\theta^*}(f_{\phi_1}(X)/f_{\phi_2}(X) = (1/t - 1) \pi_2/\pi_1) \\ &= \mathbb{P}_{\theta^*}((X - \mu_1)^t \Sigma_1^{-1} (X - \mu_1) - (X - \mu_2)^t \Sigma_2^{-1} (X - \mu_2) = -2 \log((1/t - 1) \pi_2/\pi_1) - \log(|\Sigma_1|/|\Sigma_2|)). \end{aligned}$$

Now,

$$\begin{aligned} &(X - \mu_1)^t \Sigma_1^{-1} (X - \mu_1) - (X - \mu_2)^t \Sigma_2^{-1} (X - \mu_2) \\ &= (X - \mu_1)^t \Sigma_1^{-1} (X - \mu_1) - (X - \mu_1)^t \Sigma_2^{-1} (X - \mu_2) - (\mu_1 - \mu_2)^t \Sigma_2^{-1} (X - \mu_2) \\ &= (X - \mu_1)^t (\Sigma_1^{-1} - \Sigma_2^{-1}) (X - \mu_1) - (X - \mu_1)^t \Sigma_2^{-1} (\mu_1 - \mu_2) - (\mu_1 - \mu_2)^t \Sigma_2^{-1} (X - \mu_2) \\ &= (X - \mu_1)^t (\Sigma_1^{-1} - \Sigma_2^{-1}) (X - \mu_1) - (\mu_1 - \mu_2)^t \Sigma_2^{-1} (2X - \mu_2 - \mu_1). \end{aligned}$$

Since the real matrix $\Sigma_1^{-1} - \Sigma_2^{-1}$ is symmetric, we can diagonalize it and we end up with a subset of \mathbb{R}^d of the form

$$\left\{ y \in \mathbb{R}^d : \sum_{j=1}^d (\alpha_j y_j^2 + \beta_j y_j) + \gamma = 0 \right\},$$

for some real parameters $\alpha_j, \beta_j, \gamma$. The result follows because this set has a Lebesgue measure equal to 0 in any case.

Now, since $\Sigma_1 = \Sigma_2 = \Sigma$, we have for all $t \in (0, 1)$,

$$\begin{aligned} \{T(X, \theta) > t\} &= \left\{ \forall q \in \{1, \dots, Q\}, \pi_q f_{\phi_q}(X) < (1-t) \sum_{\ell=1}^Q \pi_\ell f_{\phi_\ell}(X) \right\} \\ &= \{\pi_1 f_{\phi_1}(X) < (1/t-1)\pi_2 f_{\phi_2}(X)\} \cap \{\pi_2 f_{\phi_2}(X) < (1/t-1)\pi_1 f_{\phi_1}(X)\} \\ &= \left\{ (1/t-1)^{-1} < \frac{\pi_1 f_{\phi_1}(X)}{\pi_2 f_{\phi_2}(X)} < (1/t-1) \right\}. \end{aligned}$$

Applying $2 \log(\cdot)$ on each part of the relation, we obtain

$$\{T(X, \theta) > t\} = \{-2 \log(1/t-1) < a^t X + b < 2 \log(1/t-1)\},$$

for

$$\begin{aligned} a &= a(\theta) = 2\Sigma^{-1}(\mu_1 - \mu_2) \in \mathbb{R}^d \setminus \{0\} \\ b &= b(\theta) = -(\mu_1 - \mu_2)^t \Sigma^{-1}(\mu_1 + \mu_2) + 2 \log(\pi_1/\pi_2) \in \mathbb{R}^d. \end{aligned}$$

Since under P_{θ^*} we have $X \sim \pi_1^* \mathcal{N}(\mu_1^*, \Sigma^*) + \pi_2^* \mathcal{N}(\mu_2^*, \Sigma^*)$, we have $a^t X + b \sim \pi_1^* \mathcal{N}(a^t \mu_1^* + b, a^t \Sigma^* a) + \pi_2^* \mathcal{N}(a^t \mu_2^* + b, a^t \Sigma^* a)$. This yields for all $t \in (0, 1)$,

$$\begin{aligned} \mathbb{P}_{\theta^*}(T(X, \theta) > t) &= \pi_1 \left[\Phi \left(\frac{2 \log(1/t-1) - a^t \mu_1^* - b}{(a^t \Sigma^* a)^{1/2}} \right) - \Phi \left(\frac{-2 \log(1/t-1) - a^t \mu_1^* - b}{(a^t \Sigma^* a)^{1/2}} \right) \right] \\ &\quad + \pi_2 \left[\Phi \left(\frac{2 \log(1/t-1) - a^t \mu_2^* - b}{(a^t \Sigma^* a)^{1/2}} \right) - \Phi \left(\frac{-2 \log(1/t-1) - a^t \mu_2^* - b}{(a^t \Sigma^* a)^{1/2}} \right) \right]. \end{aligned} \tag{42}$$

A direct consequence is that for all $t \in (0, 1)$, we have $\mathbb{P}_{\theta^*}(T(X, \theta) > t) < 1$, that is, $\mathbb{P}_{\theta^*}(T(X, \theta) \leq t) = \mathbb{P}_{\theta^*}(T(X, \theta) < t) > 0$. Hence, α_c defined in (14) is equal to zero. Moreover, from (42), we clearly have that $t \in (0, 1) \mapsto \mathbb{P}_{\theta^*}(T(X, \theta) > t)$ is decreasing, so that $t \in (0, 1) \mapsto \mathbb{P}_{\theta^*}(T(X, \theta) \leq t)$ is increasing. This proves that Assumption 1 holds in that case.

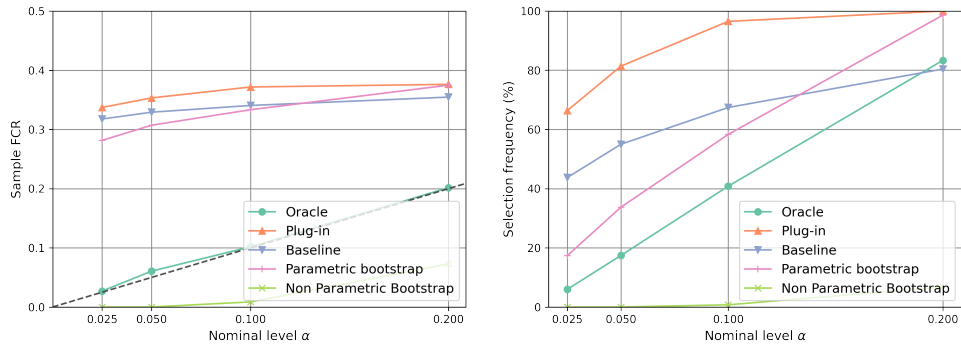
Let us now check Assumptions 2 and 3. Assumptions 2 and 3 (i) follow from Result 2.1 in Melnykov (2013).

As for Assumption 3 (ii), from (42), we only have to show that the function $t \in (0, 1) \mapsto \frac{\partial}{\partial t} \Phi \left(\frac{\log(1/t-1) - \alpha^*}{\beta^*} \right)$ is uniformly bounded by some constant $C = C(\alpha^*, \beta^*)$, for any $\alpha^* \in \mathbb{R}$ and $\beta^* > 0$. A straightforward calculation leads to the following: for all $t \in (0, 1)$,

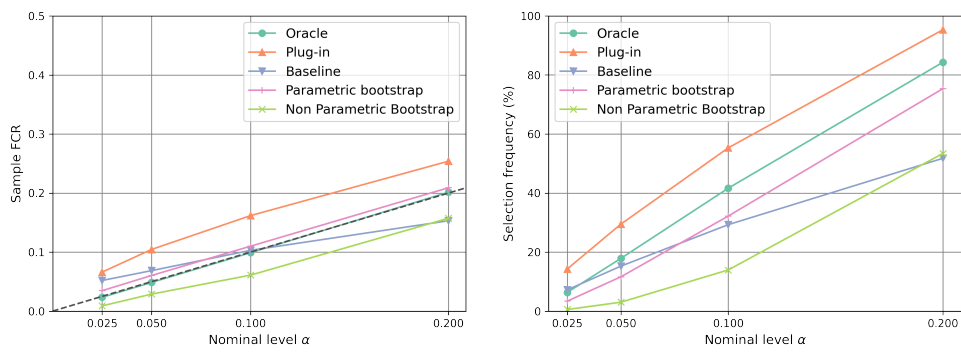
$$\left| \frac{\partial}{\partial t} \Phi \left(\frac{\log(1/t-1) - \alpha^*}{\beta^*} \right) \right| = \frac{e^{-\left(\frac{\log(1/t-1) - \alpha^*}{\beta^*}\right)^2 / 2}}{\beta^* \sqrt{2\pi}} \frac{1}{t(1-t)}. \tag{43}$$

Consider now $t_0 = t_0(\alpha^*, \beta^*) \in (0, 1/2)$ such that $\left(\frac{\log(1/t-1) - \alpha^*}{\beta^*}\right)^2 \geq 2 \log(1/t)$ for all $t \in (0, t_0)$. It is clear that the right-hand-side of (43) is upper-bounded by $\frac{1}{\beta^* \sqrt{2\pi}(1-t_0)}$ on $t \in (0, t_0)$. Similarly, let $t_1 = t_1(\alpha^*, \beta^*) \in (1/2, 1)$ such that $\left(\frac{\log(1/t-1) - \alpha^*}{\beta^*}\right)^2 \geq 2 \log(1/(1-t))$ for all $t \in (t_1, 1)$. It is clear that the right-hand-side of (43) is upper-bounded by $\frac{1}{\beta^* \sqrt{2\pi} t_1}$ on $t \in (t_1, 1)$. Finally, for $t \in [t_0, t_1]$, the upper-bound $\frac{1}{\beta^* \sqrt{2\pi} t_0(1-t_1)}$ is valid. This proves that Assumption 3 (ii) holds.

Let us now finally turn to Assumption 3 (iii). Lemma 11 ensures that $t \in (0, t^*(\bar{\alpha})) \mapsto \text{mFCR}_t^*$ is continuous increasing. Hence, $t^* : \beta \in (0, \bar{\alpha}) \mapsto t^*(\beta)$ defined in (13) is the inverse of this function and is also continuous increasing. It is therefore differentiable almost everywhere in $(0, \bar{\alpha})$, so everywhere in $(0, \bar{\alpha}) \setminus \mathcal{E}$ where \mathcal{E} is a set of Lebesgue measure 0. By taking α in $(0, \bar{\alpha}) \setminus \mathcal{E}$, this ensures that t^* is differentiable in α and thus that Assumption 3 (iii) holds. \square



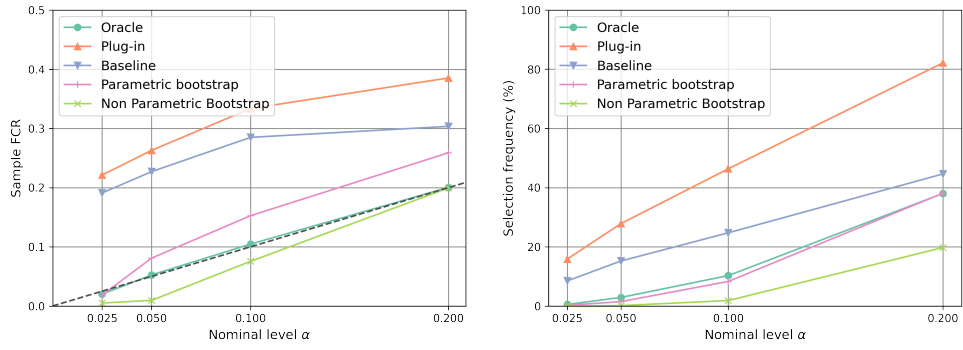
(a) $n = 200$



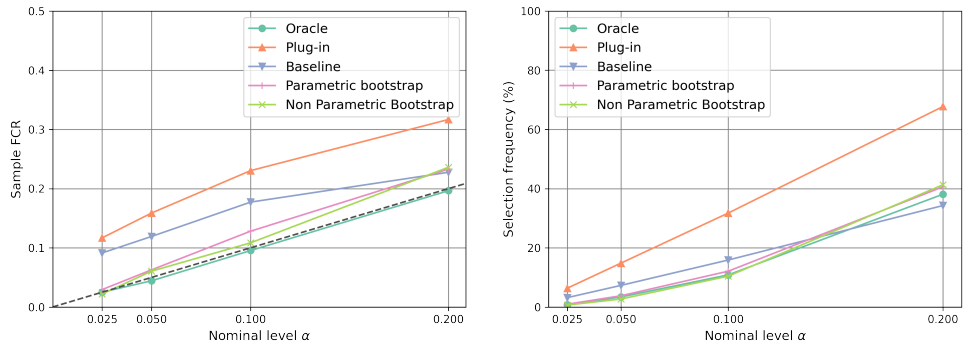
(b) $n = 1000$

Figure 9: Same as Figure 5 with dimension $d = 20$.

E Further experiments

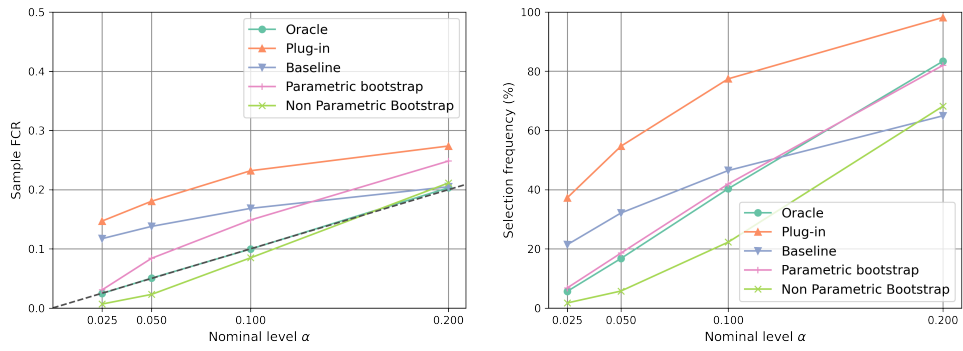


(a) $n = 200$

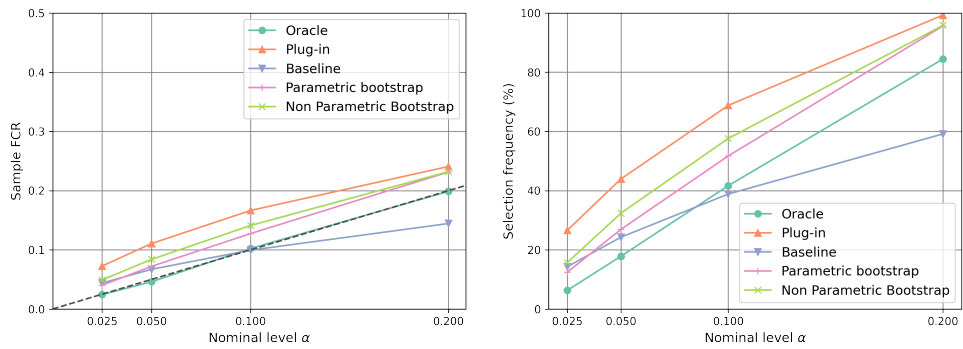


(b) $n = 1000$

Figure 10: FCR and selection frequency as a function of the nominal level α . Three-component model $Q = 3$, $\pi_1 = \pi_2 = \pi_3 = 1/3$, $\Sigma_1 = \Sigma_2 = \Sigma_3 = I_2$, $\mu_1 = 0$, $\mu_2 = (\epsilon/\sqrt{2}, \epsilon/\sqrt{2})$, $\mu_3 = (0, 2\epsilon/\sqrt{2})$.



(a) $n = 200$



(b) $n = 1000$

Figure 11: FCR and selection frequency as a function of the nominal level α . Unconstrained model with $Q = 2$, $d = 2$, $\epsilon = \sqrt{2}$, $n \in \{200, 1000\}$.