



Model-based clustering using non-parametric hidden Markov models

Elisabeth Gassiat, Ibrahim Kaddouri, Zacharie Naulet

► To cite this version:

Elisabeth Gassiat, Ibrahim Kaddouri, Zacharie Naulet. Model-based clustering using non-parametric hidden Markov models. 2024. hal-04535823

HAL Id: hal-04535823

<https://hal.science/hal-04535823>

Preprint submitted on 7 Apr 2024

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.

MODEL-BASED CLUSTERING USING NON-PARAMETRIC HIDDEN MARKOV MODELS

BY ÉLISABETH GASSIAT^{1,a}, IBRAHIM KADDOURI^{1,b} AND ZACHARIE NAULET^{1,c}

¹Université Paris-Saclay, CNRS, Laboratoire de mathématiques d'Orsay, 91405, Orsay, France,

^aelisabeth.gassiat@universite-paris-saclay.fr; ^bibrahim.kaddouri@universite-paris-saclay.fr;

^czacharie.naulet@universite-paris-saclay.fr

Thanks to their dependency structure, non-parametric Hidden Markov Models (HMMs) are able to handle model-based clustering without specifying group distributions. The aim of this work is to study the Bayes risk of clustering when using HMMs and to propose associated clustering procedures. We first give a result linking the Bayes risk of classification and the Bayes risk of clustering, which we use to identify the key quantity determining the difficulty of the clustering task. We also give a proof of this result in the i.i.d. framework, which might be of independent interest. Then we study the excess risk of the plugin classifier. All these results are shown to remain valid in the online setting where observations are clustered sequentially. Simulations illustrate our findings.

1. Introduction.

1.1. *Motivation and context.* Clustering is a statistical approach that seeks to infer interesting hidden structure in the data. Various clustering algorithms have been proposed in the literature. Some clustering algorithms do not aim to learn the distribution of the observations but only exploit their geometric properties to perform clustering. This is mainly the case for the hierarchical clustering algorithm, the k -means algorithm and their variants. Another interesting family of clustering procedures is that of *model-based clustering*, see [35, 22] for review papers. Under this framework, observations are assumed to be derived from a population mixture model, in which the populations define the clusters. Unlike geometric clustering algorithms such as k -means, model-based clustering uses a soft-max assignment: the algorithm assigns the observation to the cluster with the highest probability given the data. For example, if the observations are assumed to be independently distributed from a mixture of Gaussian distributions, one can learn the parameters using maximum likelihood (MLE) with the expectation-maximization (EM) algorithm, and then compute the probability of each class given the observation with these parameters. While similar parametric mixture models can be useful in many applications [43, 14, 22, 25], they fall short in some contexts due to the complex nature of real-world data, which may not always be modelled using a parametric mixture [10, 27]. This is the reason why it is important to study model-based clustering in the case where the observations come from a non-parametric mixture. Unfortunately, it is impossible to retrieve the mixture components without additional assumptions on the model, because otherwise the non-parametric mixture model is not identifiable. Several attempts to get identifiable mixture models with more flexible modeling of the populations distributions are described in [17, 5, 6, 42]. Our work focuses on clustering under the assumption that observations come from a HMM. Indeed, HMMs have been proved to be identifiable without any assumption on the populations distributions [18, 4]. In a HMM, to each observation Y_i is associated a non-observed random variable (thus hidden) X_i which corresponds to the true

MSC2020 subject classifications: Primary 62M05; secondary 62G99.

Keywords and phrases: Clustering, Nonparametric populations, Hidden Markov models, Mixture models.

class to which the observation belongs. These hidden random variables are assumed to form a Markov chain and the observations are independent conditionally to the sequence of hidden states. The HMM is stationary as soon as the hidden Markov chain itself is stationary. Clustering using HMMs is usually done in two steps. First, the parameters of the HMM (initial distribution and transition matrix of the hidden Markov chain, population densities – usually called emission densities) are estimated, then, using these estimated densities, the celebrated forward-backward algorithm is used to determine the plugin estimator of the a posteriori distribution of the hidden states given the observations. The clustering is then performed by soft assignment. Such a procedure is widely used in applications, see [26, 40, 20], to cite a few examples.

For each clustering procedure, the clustering risk is defined as the expected proportion of data that are not correctly clustered, and the Bayes risk of clustering is the minimum of the clustering risk over all clustering procedures. While there is a significant body of research on HMMs and their applications, the theoretical guarantees on their performance in the context of clustering are not as well-developed as some other clustering algorithms, such as k -means and its variants [21, 9, 15, 32]. This can make it difficult to interpret and rely on HMM-based clustering methods. In this work, we are interested in a precise analysis of the Bayes risk of clustering. We also aim to provide procedures with clustering risk comparable to the Bayes risk of clustering: up to bounded multiplicative constants and/or additive constants decaying to zero, and to identify the key quantity driving the difficulty of the clustering task, the so-called signal to noise ratio (SNR) in the Gaussian mixture model-based clustering literature). We focus only on the case where the dimension of the observations is fixed and independent of the sample size. To the best of our knowledge, this work is the first to provide theoretical guarantees for HMM-based clustering in the parametric and non-parametric settings.

1.2. Contribution. In the unsupervised setting, the clusters are identified only up to a permutation of the labels, which is known as the *label-switching* issue. Thus, the clustering risk and classification risk differ, the latter upper bounding the former (see Section 2 for details). The classification risk is minimized by the so-called Bayes classifier and the behavior of the Bayes risk of classification can then be closely analyzed. In contrast, the risk of clustering is difficult to analyze, with no obvious minimizer. Hence, it is desirable to establish an equivalence between the two risks. Our first main result proves that the Bayes risk of clustering is lower bounded by the Bayes risk of classification up to an *additive* constant involving the number of clusters, a forgetting coefficient of the hidden Markov chain and the number of observations n . While this constant is of order $\mathcal{O}(n^{-1/2})$ in n , we can obtain a similar lower-bound with an additive term that decrease exponentially in n , at the cost of a *multiplicative* constant smaller than 1. This shows that as long as the Bayes risk of classification does not decrease exponentially in n , the two risks are equivalent in the sense that the ratio of the Bayes risk of clustering over the Bayes risk of classification is smaller than one and bounded away from 0 uniformly over a large set of model parameters. Surprisingly, when there are only two populations, the two risks are shown to be equivalent for a sufficiently large number of observations regardless of the magnitude of the Bayes risk of clustering. However, we provide a counterexample illustrating that this equivalence does not hold whenever observations are composed of more than two clusters. A similar result is proved when the observations are independent identically distributed and we believe it is of independent interest. All these results remain true when observations are clustered sequentially. See Theorems 1 and 2.

Thanks to Theorem 2, we are able to conduct a precise analysis of the Bayes risk of clustering relying on the simpler Bayes risk of classification. We are interested in understanding the difficulty of the clustering task with respect to the separation between the emission densities only, the transition matrix is seen as fixed. In Theorem 3, we identify the key quantity

driving the difficulty of clustering which is

$$\Lambda := \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y).$$

Here, $\mathbb{X} = \{1, \dots, J\}$ and the $(f_x)_{x \in \mathbb{X}}$ are the population densities of the distributions $(F_x)_{x \in \mathbb{X}}$ with respect to a dominating measure \mathcal{L} over the observation space \mathbb{Y} , that is $dF_x = f_x d\mathcal{L}$, J being the number of clusters. When $J = 2$, $\Lambda = 1 - \|F_1 - F_2\|_{\text{TV}}$, where $\|F_1 - F_2\|_{\text{TV}}$ is the total variation distance between the two distributions. In the particular situation where the observations are independent and modelled through a Gaussian mixture model, we recover in a wide range of regimes the SNR identified in [21, 38], see the discussion at the end of Section 3.3.

We then investigate the usual plugin procedure. Theorem 4 bounds the excess risk of the plugin classifier by the expected errors in the estimation of the model parameters. By combining carefully the kernel density estimator of [1] and the estimator of the transition matrix and initial distribution of [12], we construct an estimator of the model parameters whose excess risk can be shown to decay to zero at a rate that can be specified thanks to Theorem 4. Finally, we present numerical simulations to highlight the fact that HMM modeling allows clustering data for which other usual methods can not recover the hidden structure.

1.3. Related work. Mixture models have the advantage of providing a mathematical definition of a cluster. However, given their non-identifiability in the general case [41], model-based clustering in the non-parametric mixture setting requires additional assumptions making the model identifiable and thus the mixture components learnable. Similar assumptions include dependence of the observations [19, 17], symmetry of the mixture components [24] or separation between the mixture components [5]. In [5], the authors retrieve the identifiability of the non-parametric mixture thanks to assumptions on the proximity between the mixture components. Then they propose an algorithm for clustering i.i.d. observations of this mixture. Such a trend of work is complementary to ours. Our HMM modeling covers a wider setting of cluster distributions but requires a specific dependence in consecutive observations.

Identifiability of non-parametric HMMs has been recently discovered [18, 4]. Given that the transition matrix of the hidden Markov chain is non singular and ergodic, all parameters, including the number J of hidden states of the Markov chain and the emission distributions, can be recovered from the distribution of K consecutive observations (Y_1, \dots, Y_K) as long as the emission distributions are distinct and K is large enough compared to J [4]. If we assume the emission distributions to be linearly independent in the set of all distributions, we can take $K = 3$. It is not needed here that they have densities with respect to some known dominating measure. Several non-parametric density estimation procedures have been proposed in the literature [8, 12, 30, 31, 1, 2]. The way in which filtering and smoothing distributions are affected when plugging-in estimators of the parameters is studied in [12].

Given that model-based clustering procedures learn the model parameters, the quality of their clustering depends strongly on whether the model is identifiable or not. In the case where the dependence structure is lost, the model becomes non-identifiable and learning the model parameters becomes impossible. With a finite number of observations, the performances can become quite bad if the data, still dependent, are not far to be independent. The works [3, 2] quantify precisely how the estimation performances are affected when getting close to independence.

The Gaussian mixture model (GMM) is the object of intensive research since it serves as the standard framework for theoretical analysis of clustering algorithms. In the context of this model, the optimality of spectral clustering was proved in [33]. Constraining the Gaussian mixture to have only two components allowed showing a variant of Lloyd's algorithm

to be minimax optimal and an exact phase transition for perfect recovery of the clusters was characterized in [38]. While these findings allow a precise analysis of the risk of clustering, they are restricted to the parametric setting of Gaussian mixtures and do not apply to more complex modelling of the observations. The SubGaussian Mixture Model, on the other hand, is a more flexible model. Lloyd's algorithm was investigated in [32] within the framework of this model and precise guarantees on the error of clustering were proved under suitable initialization of the algorithm. Under the same modelling, the clustering performance of the relaxed k -means algorithm was studied in [21], revealing the signal-to-noise ratio with respect to which the clustering error decays. These two results, however, require a minimum separation between the means of the clusters, and thus, do not cover all the possible regimes.

Online algorithms play an important role in the context of HMMs. Unlike batch methods, which require all data to be available in advance, online algorithms enable incremental learning of the HMM parameters, which makes these algorithms well-suited for real-time applications. Plugin methods can be used to perform online clustering. See [28, 29] for examples of similar algorithms.

1.4. Organization of the paper. The assumptions of the model and the definitions of the various risks we shall use in this work are described in Section 2. We state our main results in Section 3, the proofs of which are detailed in Section 6. Section 4 is devoted to simulation experiments, and Section 5 to possible further work.

2. Setting and definitions.

2.1. Notations. We note for $i \leq j$ $X_{i:j} := (X_i, \dots, X_j)$ and $Y_{i:j} := (Y_i, \dots, Y_j)$. We denote $[n] := \{1, \dots, n\}$ and $\mathcal{P}[n]$ is the set of partitions of $[n]$. The set of permutations of $[J]$ will be denoted by \mathcal{S}_J , and for any $\tau \in \mathcal{S}_J$, we consider θ^τ , the model parameter when labels are permuted with τ , that is $\theta^\tau = (\nu^\tau, Q^\tau, (f_{\tau(x)})_{x \in \mathbb{X}})$ where $\nu^\tau = (\nu_{\tau(x)})_{x \in \mathbb{X}}$ and $Q^\tau = (Q_{\tau(x), \tau(x')})_{x, x' \in \mathbb{X}}$. For h a measurable function, $\|h\|_\infty$ denotes the essential supremum of h , possibly infinite. Frobenius norm is denoted by $\|\cdot\|_F$ and $\|\cdot\|$ stands for the operator norm. Given two sequences of positive numbers $(a_n)_n$ and $(b_n)_n$, $a_n \gtrsim b_n$ means that there exists an absolute constant $c > 0$ and $n_0 \in \mathbb{N}$ such that $(\forall n \geq n_0) \quad a_n \geq cb_n$. In the HMM modeling, for any parameter θ and any $\tau \in \mathcal{S}_J$, the distribution of the observations under \mathbb{P}_θ is the same as under \mathbb{P}_{θ^τ} . In other words, this distribution is invariant up to permutation of the labels given by the hidden states. This is known as the *label-switching* issue.

2.2. The model. We consider a Hidden Markov Model with J hidden states taking value in a set of labels $\mathbb{X} = \{1, \dots, J\}$, and observations in a Polish space endowed with its Borel σ -field $(\mathbb{Y}, \mathcal{Y})$. We denote $\mathbf{X} = (X_1, X_2, \dots)$ and $\mathbf{Y} = (Y_1, Y_2, \dots)$ respectively the sequence of hidden states forming the Markov chain and the observations. We assume that the emission distributions have densities f_x , $x = 1, \dots, J$, with respect to a dominating measure \mathcal{L} on $(\mathbb{Y}, \mathcal{Y})$. The HMM assumption boils down to:

$$\begin{aligned} \mathbf{X} &\sim \text{Markov}(\nu, Q) \\ \mathbb{P}_\theta(Y_i \in \cdot \mid \mathbf{X}) &= f_{X_i}(\cdot) d\mathcal{L} \end{aligned}$$

where ν is the initial distribution of the chain and Q the transition matrix of the hidden chain. Throughout the paper, it is assumed that only the beginning $Y_{1:n}$ of the sequence \mathbf{Y} is observed, and nothing else. We set $\theta = (\nu, Q, (f_x)_{x \in \mathbb{X}})$ the parameter of the model and Θ denotes the space of all valid parameters. The following assumption is made:

- *Independent case.* In this case we assume that all the lines of Q are identical and equal to the vector ν of weights forming its stationary distribution. This corresponds to the usual mixture model with independent latent variables. The set of these parameters will be denoted Θ^{ind} .
- *Dependent case.* In this case we assume the lines of the transition matrix Q not all equal, so that the $(X_i)_{i \geq 1}$ are not independent. The set of these parameters will be denoted Θ^{dep} .

Throughout this work, we shall consider the transition matrix being fixed, and we will be interested in how the separation of populations, understood as some quantity depending on the populations densities, drives the difficulty of the clustering task. We show that, under the HMM modelling, it is possible to cluster general populations without assuming they belong to some prescribed parametric family. While our primary interest is in the dependent case, in which the emission densities can be identified without any further constraint, we obtain, however, results that have interest also in the widely used independent case, in particular in the traditional Gaussian location-scale mixture model (see Section 3.1).

One main difference between the HMM and the i.i.d. situation in the analysis of the Bayes risk of classification is that in the HMM modeling, the probability of a label X_i given the observations Y_1, \dots, Y_n depends on all the observations, whereas in the i.i.d. case it depends only on the associated observation Y_i . The vector of probabilities of the X_i 's given Y_1, \dots, Y_n are called the smoothing distributions, they depend on i, n and the observations. The vectors of probabilities of the X_i 's given the observations up to time i , Y_1, \dots, Y_i , are called the filtering distributions, they depend on i, n and the involved observations. Recursive formulas verified by the filtering and smoothing distributions make the computations tractable under the HMM modelling. See Section 3 of [7] for details.

2.3. Offline clustering. For any $n \geq 1$, the finite sequence $X_{1:n} = (X_1, \dots, X_n)$ induces a random partition $\Pi_n = \{C_1, C_2, \dots\}$ of $[n]$ whose blocks – the so-called *clusters* – are the equivalence classes for the random equivalence relation $i \sim j \iff X_i = X_j$. The goal of clustering is to uncover this partition Π_n on the sole basis of the observation $Y_{1:n} = (Y_1, \dots, Y_n)$. We define a *clusterer*:

DEFINITION 1 (Clusterer). A n -clusterer is a measurable map $g : \mathbb{Y}^n \rightarrow \mathcal{P}[n]$. We denote by \mathcal{G}_n the set of all n -clusterers.

We measure the loss incurred by guessing $g(Y_{1:n})$ in place of Π_n via the maximum overlap of a *matching* between the two partitions. To define a matching, we build the complete bipartite graph $(g(Y_{1:n}), \Pi_n, \mathcal{E}(g(Y_{1:n}), \Pi_n))$ on vertices $g(Y_{1:n})$ and Π_n with edge set $\mathcal{E}(g(Y_{1:n}), \Pi_n) := \{\{C, C'\} : C \in g(Y_{1:n}), C' \in \Pi_n\}$. Then we recall that a matching M is a set $M \subseteq \mathcal{E}(g(Y_{1:n}), \Pi_n)$ of edges without common vertices (*i.e.* each block of Π_n and $g(Y_{1:n})$ appears in at most one edge of the matching). An example of a matching is depicted in Figure 1. Then, we consider the following loss [37, 36] of $g(Y_{1:n})$ with respect to the true partition Π_n

$$(g(Y_{1:n}), \Pi_n) \mapsto 1 - \frac{1}{n} \sup_{\substack{M \subseteq \mathcal{E}(g(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C').$$

with associated risk function $\mathcal{R}_n^{\text{clust}} : \Theta \times \mathcal{G}_n \rightarrow [0, 1]$

$$(1) \quad \mathcal{R}_n^{\text{clust}}(\theta, g) := \mathbb{E}_\theta \left[1 - \frac{1}{n} \sup_{\substack{M \subseteq \mathcal{E}(g(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C') \right].$$

A closely related notion is that of a *classifier*:

DEFINITION 2 (Classifier). A n -classifier is a measurable map $h : \mathbb{Y}^n \rightarrow \mathbb{X}^n$. We denote by \mathcal{H}_n the set of all n -classifiers.

One may find our Definition 2 different from standard textbook definitions [13]: we explain the reason of this choice later in Remark 1. A classifier differs from a clusterer in that it not only seeks for the hidden partition, but also for the labels of the observations. Hence, the usage of a classifier only makes sense in a supervised framework where access to some labeled data is allowed in some way. In an unsupervised framework such as our model, \mathbf{Y} does not contain any information about the labels and recovering them better than a lucky guess is impossible. It is however true that, to any n -classifier $h \in \mathcal{H}_n$ corresponds a unique n -clusterer $g \in \mathcal{G}_n$ which can be built via the map $\pi_n : \mathbb{X}^n \rightarrow \mathcal{P}[n]$ such that

$$g(Y_{1:n}) = \pi_n \circ h(Y_{1:n}) = \{\{i : h_i(Y_{1:n}) = x\} : x \in \mathbb{X}\} \setminus \{\emptyset\}$$

and any clusterer can be represented that way by choosing a specific labelling of the clusters. For this reason, the notions of clusterer and classifier are very much often amalgamated in the literature. We argue that it would be better to define them separately in order to avoid confusions between the risk of clustering $\mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h)$ and the risk of classification $\mathcal{R}_n^{\text{class}} : \Theta \times \mathcal{H}_n \rightarrow [0, 1]$ (relative to the loss counting number of misclassified observations)

$$(2) \quad \mathcal{R}_n^{\text{class}}(\theta, h) := \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq X_i} \right].$$

Here again we insist that the definition of $\mathcal{R}_n^{\text{class}}$ in Equation (2) – although mathematically correct – has no statistical interest in an unsupervised model where the classifier h is not allowed to see some of the labels. An easy exercise (see Lemma 14 in Section 6.8) shows that the risk of clustering of $\pi_n \circ h$ can be rewritten as

$$(3) \quad \mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h) = \mathbb{E}_\theta \left[\min_{\tau \in \mathcal{S}_J} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \right]$$

and differs from (2). Note that it does not depend on the classifier h chosen to represent the clusterer thanks to the infimum over the permutations of the labels. Note also that the infimum inside the expectation reflects the ease of the clustering problem compared to the classification problem, because unlike classification, clustering does not seek to identify the true labels themselves.

It is customary to compare the performance of a given classifier h to the best performance attainable by an oracle classifier, namely the *Bayes risk of classification*:

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h).$$

In particular, it is well-known that the previous optimization problem is solved by the (albeit non necessarily unique) so-called Bayes classifier $h_\theta^* = (h_{\theta,1}^*, \dots, h_{\theta,n}^*)$ such that

$$\mathbb{P}_\theta(X_i = h_{\theta,i}^*(Y_{1:n}) \mid Y_{1:n}) = \max_{x \in \mathbb{X}} \mathbb{P}_\theta(X_i = x \mid Y_{1:n}), \quad i = 1, \dots, n.$$

In an unsupervised learning context, it makes sense to compare the risk of a clusterer to the *Bayes risk of clustering*:

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g).$$

As for the classification risk, the solution to the previous optimization problem exists and is obtained by the Bayes clusterer g_θ^* such that $g_\theta^*(Y_{1:n})$ is the partition that minimizes

$$(4) \quad g \mapsto \mathbb{E}_\theta \left[1 - \frac{1}{n} \sup_{\substack{M \subseteq \mathcal{E}(g(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C') \middle| Y_{1:n} \right].$$

In contrast with classification, however, the Bayes clusterer has usually no simple expression. It is to be noted that there is no reason that $g_\theta^* = \pi_n \circ h_\theta^*$. Although the inequality $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$ is true and easily proved from (3), it is not guaranteed that $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)$ and $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$ are equivalent (*ie.* have comparable magnitude). In Sections 3.1 and 3.2 we show that the equivalence holds both in the independent and in the dependent scenarios as long as the Bayes risk of classification is not exponentially small in n . Though we prove the two risks to be equivalent in some specific contexts, we provide a counter-example showing that this is not always true.

REMARK 1. It is common [13] to define a classifier as a function $h : \mathbb{Y} \rightarrow \mathbb{X}$ [as opposed to $\mathbb{Y}^n \rightarrow \mathbb{X}^n$ in our Definition 2]. This usual definition is motivated by the fact that in the independent scenario the law of $X_i | Y_{1:n}$ is that of $X_1 | Y_1$. Hence if one is willing to classify the vector $Y_{1:n}$, the Bayes classifier rewrites as $h_\theta^* = (h_{\theta,1}^*(Y_1), \dots, h_{\theta,1}^*(Y_n))$ with $h_{\theta,1}(y)$ maximizing $x \mapsto \mathbb{P}_\theta(X_1 = x | Y_1 = y)$ and the Bayes risk of classification equals $\mathbb{P}_\theta(h_{\theta,1}^*(Y_1) \neq X_1)$. Thus classifying Y_1 or $Y_{1:n}$ is not very different. In the dependent case, however, the situation differs since $\mathbf{X} | \mathbf{Y}$ is a inhomogeneous Markov chain. This implies in particular that classifying Y_1 or $Y_{1:n}$ are different problems when $n \geq 2$, and the optimal solution for classifying $Y_{1:n}$ cannot be obtained from the optimal solution of classifying Y_1 .

REMARK 2. In [34], the authors define the risk of clustering of a n -classifier h in a different way:

$$\mathcal{R}_n^{\text{MRSS}}(\theta, h) := \mathbb{E}_\theta \left[\inf_{\tau \in \mathcal{S}_J} \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau(X_i) \neq h_i(Y_{1:n})} \middle| Y_{1:n} \right] \right]$$

Similarly, in [32], the risk of clustering is defined by:

$$\tilde{\mathcal{R}}_n^{\text{clust}}(\theta, h) := \inf_{\tau \in \mathcal{S}_J} \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau(X_i) \neq h_i(Y_{1:n})} \right]$$

Their definition is mathematically convenient as one can easily show that (see Lemma 5):

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{MRSS}}(\theta, h) = \inf_{h \in \mathcal{H}_n} \tilde{\mathcal{R}}_n^{\text{clust}}(\theta, h) = \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$$

Hence, the Bayes clusterer relative to their risks is $g^* = \pi_n \circ h^*$ with h^* the Bayes classifier relative to the risk $\mathcal{R}_n^{\text{class}}$. However, contrary to our definition, it seems that there is no suitable statistical interpretation of these risks of clustering.

2.4. *Online clustering.* We are also interested in the situation where the observations must be clustered *sequentially*; *ie.* estimating $(\Pi_k)_{k \geq 1}$ sequentially. In this situation, we are aiming for a sequence of clusterers $\mathbf{g} = (g^{(k)})_{k \geq 1}$ with $g^{(k)} \in \mathcal{G}_k$ for every $k \geq 1$ (thus each $g^{(k)}$ is $Y_{1:k}$ -measurable). The major difference with offline clustering is that, the sequence must satisfy the constraint that the partition $g^{(i+1)}(Y_{1:i+1})$ is obtained from $g^{(i)}(Y_{1:i})$ by either assigning Y_{i+1} to an existing block, or a new block, but no other change is allowed.

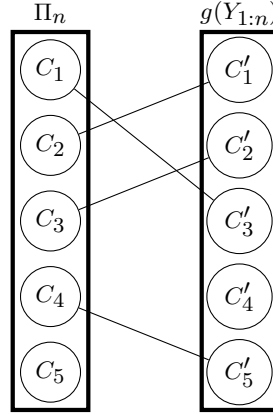


Figure 1: Example of a matching. Nodes on the left represent the clusters induced by the partition of Π_n ; those on the right are the clusters of $g(Y_{1:n})$. Edges form a matching between the two partitions.

Another way to put this is that the sequence \mathbf{g} must be *consistent* in the sense that for all $1 \leq m \leq n$, the partition $g^{(n)}(Y_{1:n})$ restricted to $[m]$ must equal $g^{(m)}(Y_{1:m})$, \mathbb{P}_θ -a.s, *ie*.

$$(5) \quad \{C \cap [m] : C \in g^{(n)}(Y_{1:n})\} \setminus \{\emptyset\} = g^{(m)}(Y_{1:m}) \quad \mathbb{P}_\theta\text{-a.s}$$

Note that the interest of studying online methods lies in the advantages they offer in terms of theoretical and computational properties: online methods are generally easier to implement in the HMM framework and they have plausible theoretical properties making their theoretical study less complicated compared to their offline counterparts.

DEFINITION 3 (Online clusterer). An online clusterer is a sequence $\mathbf{g} \in \prod_{k \geq 1} \mathcal{G}_k$ that satisfies the consistency constraint (5). We denote by \mathcal{G}^{on} the set of all online clusterers.

For each time horizon $n \geq 1$, we consider the risk function $\mathcal{R}_n^{\text{clust, on}} : \Theta \times \mathcal{G}^{\text{on}} \rightarrow [0, 1]$

$$\mathcal{R}_n^{\text{clust, on}}(\theta, \mathbf{g}) := \mathcal{R}_n^{\text{clust}}(\theta, g^{(n)}) = \mathbb{E}_\theta \left[1 - \frac{1}{n} \sup_{\substack{M \subseteq \mathcal{E}(g^{(n)}(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C') \right].$$

Since each g_k in the sequence $\mathbf{g} \in \mathcal{G}^{\text{on}}$ can be represented (in a non-unique way) as $\pi_k \circ h^{(k)}$ for some classifier $h^{(k)} \in \mathcal{H}_k$, the whole sequence can be represented as $\mathbf{g} = (\pi_k \circ h^{(k)})_{k \geq 1} \equiv \pi \circ \mathbf{h}$. The consistency constraint (5) then translates on the sequence $\mathbf{h} = (h^{(k)})_{k \geq 1}$ as

$$(6) \quad \forall 1 \leq m \leq n, \quad h_m^{(m)}(y_{1:m}) = h_m^{(n)}(y_{1:n}) \quad \mathbb{P}_\theta\text{-a.s}$$

Although online classifiers have no admitted definition [23], we adopt the following definition which is consistent with our framework.

DEFINITION 4 (Online classifier). An online classifier is a sequence $\mathbf{h} \in \prod_{k \geq 1} \mathcal{H}_k$ that satisfies the consistency constraint (6). We denote by \mathcal{H}^{on} the set of all online classifiers.

Then, for a classifier $\mathbf{h} \in \mathcal{H}^{\text{on}}$, we define the risk of online classification by:

$$(7) \quad \mathcal{R}_n^{\text{class, on}}(\theta, \mathbf{h}) := \mathcal{R}_n^{\text{class}}(\theta, h^{(n)}) = \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i^{(n)}(Y_{1:n}) \neq X_i} \right].$$

Here again, the same remarks we made in the offline context still apply: the risk of online classification is of no statistical interest in an unsupervised context where labels are not observed. However, its analysis would be crucial in obtaining results concerning the risk of clustering.

In view of the definition of $\mathcal{R}_n^{\text{clust, on}}$, of (3) and (6), we can easily show that

$$(8) \quad \forall \mathfrak{h} \in \mathcal{H}^{\text{on}}, \forall \theta \in \Theta, \quad \mathcal{R}_n^{\text{clust, on}}(\theta, \pi \circ \mathfrak{h}) = \mathbb{E}_\theta \left[\min_{\tau \in \mathcal{S}_J} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{h_i^{(n)}(Y_{1:n}) \neq \tau(X_i)\}} \right].$$

Similarly to the offline framework, the *Bayes risk of online classification* can then be defined as :

$$\inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class, on}}(\theta, \mathfrak{h}) = \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i^{(n)}(Y_{1:n}) \neq X_i} \right]$$

This risk is minimized by the Bayes classifier $h_\theta^{*,(n)} = (h_{\theta,1}^{*,(n)}, \dots, h_{\theta,n}^{*,(n)})$ such that

$$\mathbb{P}_\theta(X_i = h_{\theta,i}^{*,(n)}(Y_{1:n}) \mid Y_{1:i}) = \max_{x \in \mathbb{X}} \mathbb{P}_\theta(X_i = x \mid Y_{1:i}), \quad i = 1, \dots, n.$$

Note that we did not use the usual format of Definition 4 in the expression of the online Bayes classifier, but we rather used $h_\theta^{*,(n)}(y_{1:n}) = (h_{\theta,1}^{*,(n)}(y_{1:n}), \dots, h_{\theta,n}^{*,(n)}(y_{1:n}))$. This is justified by the fact that risk depends solely on the n -th component of the online classifier \mathfrak{h} and this is the format we will adopt in Section 3.4 when studying plugin methods.

3. Main results.

3.1. Comparison of clustering and classification risks in the independent case. The Bayes risk of classification offers an advantage due to its closed formula, stemming from the explicit identification of the Bayes classifier. In contrast, the Bayes risk of clustering lacks this straightforward formulation. However, thanks to Equation (3), the risk of a clustering procedure can be seen as the risk of classification of the associated classifier up to the best random permutation (the one minimizing the sum inside the expectation in Equation (3)). This is why a common idea is that risk of clustering and risk of classification are closely related. Following this intuition, authors of [44, 16] propose an inequality linking the minimax (over a specific class) risk of clustering to the minimax risk of classification, in the context of model-based community detection under the stochastic block model. Their inequality is applied in [38] in the context of the mixture of two Gaussian distributions and in [32] for general subGaussian mixtures. Although their argument is neat, it relies heavily on two key ingredients: (i) in their model the latent partition is deterministic and viewed as a parameter, and (ii) to lower bound the minimax risk, it is enough to consider some well-chosen parameters. Thus their argument is not transposable to the situation where one is interested in bounding the Bayes risk of clustering at every possible value of the parameter with random labels. Here we propose another route. We believe this contribution is of independent interest. The case of independent labels is considered in this section, while the more challenging case of Markovian labels is deferred to the next section.

THEOREM 1. *In the case of independent labels, if $J = 2$, then for all $\theta \in \Theta^{\text{ind}}$ and all $n \geq 1$*

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq \left[1 - 2e \left(\sqrt{\frac{\log(2)}{2n}} \right)^{1 - \frac{2}{n}} \right] \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h).$$

If $J \geq 3$, then for all $\theta \in \Theta^{\text{ind}}$ and all $n \geq 1$ the following bounds are true

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - \sqrt{\frac{\log(J!)}{2n}},$$

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq \left[1 - \frac{4e}{\beta} \left(\sqrt{\frac{\log(J!)}{2n}} \right)^{1-4/(n\beta)} \right] \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - J^2 e^{-n\beta/8}$$

where $\beta = \min_{j \neq k} (\nu_j + \nu_k)$. An exactly similar result holds in the online framework.

The proof of Theorem 1 is given in Section 6.2.

Theorem 1 establishes the remarkable fact that, when $J = 2$, the Bayes risk of clustering and the Bayes risk of classification are uniformly comparable provided $n \geq 17$. When $J = 3$, however, the bound become vacuous if $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \lesssim J^2 e^{-n\beta/8}$. The next proposition shows that this is not an artifact of our bound, but rather that the two risks are genuinely not comparable uniformly over the parameter space.

PROPOSITION 1. Whenever $J \geq 3$ and $n \geq 1$:

$$\inf_{\theta \in \Theta^{\text{ind}}} \frac{\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)}{\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)} = 0.$$

The proof of Proposition 1 is given in Section 6.3. The proof sheds light on why the previously established equivalence when there are only two classes no longer holds when the number of classes exceeds three. Indeed, the absence of equivalence in our specific example is partially attributed to the configuration where observations form two clusters with independent supports, with one cluster being significantly larger and the other scarcely observed. This particular scenario does not present an issue in the context of only two classes, as both the losses of classification and clustering will be null. However, as J increases to three or more, a substantial disparity emerges between the magnitudes of the risk associated with clustering and that of classification. This difference in magnitudes makes the equivalence impossible.

3.2. *Comparison of clustering and classification risks in the dependent case.* We now establish results similar to those of the previous section, in the case of Markovian labels.

ASSUMPTION 1. $\delta = \min_{x, x'} Q_{x, x'} > 0$ and $\min_x \nu_x \geq \delta$.

The positive lower bound δ introduced in Assumption 1 makes the hidden Markov chain irreducible. It will be used in proving deviation inequalities or when forgetting properties of the chain are needed.

THEOREM 2. In the case of dependent labels and under Assumption 1, if $J = 2$, then for all $\theta \in \Theta^{\text{dep}}$ and all $n \geq 1$

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq \left[1 - 2e \left(\frac{1 - \delta}{\delta} \right)^4 \left(\frac{\log(2)}{2n} \right)^{\frac{1}{2} - \frac{1}{n}} \right] \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$$

If $J \geq 3$, then for all $\theta \in \Theta^{\text{dep}}$ and all $n \geq 1$ the following bounds are true

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - \frac{1}{1 - \rho_0} \sqrt{\frac{\log(J!)}{2n}}$$

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq \left[1 - \frac{5}{\beta(1 - \rho_0)} \sqrt{\frac{\log(J!)}{2n}} \right] \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - (J^2 + 1)e^{-2n(1 - \rho_0)^2 \beta^2 / 25}$$

where $\beta = \min_{i,j \neq k} \mathbb{P}_\theta(X_i \in \{j, k\})$, $\rho_0 = \frac{1-J\delta}{1-(J-1)\delta}$, $\rho_1 = 1 - J\delta$. An exactly similar result holds in the online framework.

The proof of Theorem 2 is given in Section 6.2.

Note that when the HMM is stationary, $\beta = \min_{j \neq k}(\nu_j + \nu_k)$ as in the i.i.d. case. Notice also that when $\delta = 1/J$, the observations are i.i.d. with uniform distribution over the set \mathbb{X} , and we recover the first inequality for i.i.d. observations. However, we do not recover the inequality for i.i.d. observations in general from that of HMM observations. We give specific proofs for each case in Section 6.2. Thanks to this clear link between classification and clustering risks, every bound on the Bayes risk of classification can be extrapolated to the Bayes risk of clustering. Given the closed formula of the risk of classification and the recursive formulas ensured by the filtering and smoothing distributions, it is simpler to analyse the risk of classification than that of clustering.

3.3. A key quantity for the Bayes risk of clustering. We now state our main result which proves upper and lower bounds on the Bayes risks for online and offline classification in function of a quantity measuring the separation between the emission densities up to constants depending on the transition matrix. These bounds can be extrapolated to the risk of clustering thanks to Theorem 2. Let $\Lambda := \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y)$.

THEOREM 3. *Under Assumption 1, for all $n \geq 1$, for all $\theta \in \Theta^{\text{dep}}$, the Bayes risks for online and offline classification satisfy the following inequalities:*

$$\frac{\delta}{1 - (J-1)\delta} \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathfrak{h}) \leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \leq \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathfrak{h})$$

$$\delta \Lambda \leq \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathfrak{h}) \leq (1 - (J-1)\delta) \Lambda$$

The proof of Theorem 3 is given in Section 6.4.

Note that the bounds vanish when all the emission densities have disjoint supports. In the case where $\delta = \frac{1}{J}$, all the quantities involved in the inequalities above are equal to Λ/J . This situation corresponds to i.i.d. observations derived from a mixture with J components of equal weights. This proves in particular the tightness of the bounds which cannot be improved by any absolute multiplicative constant without restricting the parameter space.

In the case of a HMM with two hidden states, the theorem simplifies to:

$$\delta \int_{\mathbb{Y}} [f_1 \wedge f_2](y) d\mathcal{L}(y) \leq \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathfrak{h}) \leq (1 - \delta) \int_{\mathbb{Y}} [f_1 \wedge f_2](y) d\mathcal{L}(y)$$

Summarizing Theorems 2 and 3, we first see that:

$$\begin{aligned} \frac{\delta^2}{1 - (J-1)\delta} \Lambda - \frac{1}{1 - \rho_0} \sqrt{\frac{\log(J!)}{2n}} &\leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \\ &\leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \\ &\leq \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathfrak{h}) \\ &\leq \Lambda \end{aligned}$$

Let $C(\nu, Q, J)$ be a positive constant depending only on ν , Q and J . In the regime where $\Lambda \geq C(\nu, Q, J)e^{-cn(1-\rho_0)^2\beta^2}$:

$$\alpha\Lambda \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \Lambda$$

where $\alpha = \left[1 - \frac{5}{\delta\beta} \sqrt{\frac{\log(J!)}{2n}}\right] \frac{\delta^2}{1-(J-1)\delta} - \frac{J^2+1}{C(\nu, Q, J)}$ which is positive provided $n \gtrsim \frac{\log(J!)}{(\delta\beta)^2}$ and $C(\nu, Q, J) \gtrsim \left(\frac{J}{\delta}\right)^2$. In the same regime, one has:

$$\alpha'\Lambda \leq \inf_{\mathfrak{g} \in \mathcal{G}_{\text{on}}} \mathcal{R}_n^{\text{clust, on}}(\theta, \mathfrak{g}) \leq \Lambda$$

where $\alpha' = \left[1 - \frac{5}{\delta\beta} \sqrt{\frac{\log(J!)}{2n}}\right] \delta - \frac{J^2+1}{C(\nu, Q, J)}$ which is positive provided $n \gtrsim \frac{\log(J!)}{(\delta\beta)^2}$ and $C(\nu, Q, J) \gtrsim \frac{J^2}{\delta}$. In other words, in the regime where $\Lambda \geq C(\nu, Q, J)e^{-cn(1-\rho_0)^2\beta^2}$, Λ is the quantity driving the difficulty of clustering.

In the context of Gaussian HMMs with two hidden states, the previous results translate in a clear identification of the signal-to-noise ratio which drives the Bayes risk of clustering. In the case where the two emissions have means μ_0 and μ_1 with the same covariance matrix Σ , the Bayes risk of clustering ensures:

$$\alpha \exp\left(-\frac{\text{SNR}}{2}\right) \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \exp\left(-\frac{\text{SNR}}{8}\right)$$

where $\text{SNR} = (\mu_0 - \mu_1)^\top \Sigma^{-1}(\mu_0 - \mu_1)$. No restriction was made on the dimension of the observation space, the SNR is valid for both low and high dimension. A similar result holds for the online risk.

In many regimes, this signal-to-noise ratio is somewhat similar to that retrieved in [21] when studying the relaxed k -means algorithm and in [38] where the minimax risk of clustering is considered. To be able to compare, consider the simple setting where $\Sigma = \sigma^2 I_p$, with I_p the identity matrix in \mathbb{R}^p , and set $\Delta = \|\mu_1 - \mu_2\|$. Then, [38] gives a lower bound for the minimax clustering risk of the form $C \exp(-cr_n^2)$, with $C > 0$ and $r_n = \frac{\Delta^2/\sigma^2}{\sqrt{\Delta^2/\sigma^2 + p/n}}$, whereas ours is of the form $C \exp(-c\Delta^2/\sigma^2) - \exp(-c'n)$, in the context of non-parametric HMMs. Both bounds have the same magnitude when $a\frac{p}{n} \leq \frac{\Delta^2}{\sigma^2} \leq bn$ for specific positive constants a, b . Moreover, in this regime, the magnitude of our upper-bound ($C \exp(-c\Delta^2/\sigma^2)$) coincides with theirs too. Unlike their result, we obtain our SNR without prior geometric modeling of the distributions of the populations thanks to the HMM assumption. We prove below upper bounds on the excess risk when nothing is known about the populations distributions.

3.4. Excess risk of the plugin procedure. The results in Section 3.3 say that the clustering performance attained by the clusterer $\pi_n \circ h_\theta^*$ with $h_\theta^* \in \arg \min_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$ shall be close enough to the performance of a “true” Bayes clusterer in $\arg \min_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)$. This is good news since the Bayes clusterer has no simple expression, but h_θ^* does. The issue, as usual, is that h_θ^* depends on the unknown parameter $\theta \in \Theta$. A simple yet classical remedy is to first estimate θ by $\hat{\theta}$ and use the *plugin* classifier $h_{\hat{\theta}}^*$ with the hope its performance will mimic that of h_θ^* . In this section, we focus on the clustering rule obtained via the plugin classifier, that is for the offline case

$$\hat{h}^\tau(Y_{1:n}) := \left(\arg \max_{x \in \mathbb{X}} \mathbb{P}_{\hat{\theta}(Y_{1:n})^\tau}(X_i = x \mid Y_{1:n}) \right)_{1 \leq i \leq n}$$

and for the online case

$$\hat{\mathfrak{h}}^\tau(Y) := \left(\left(\arg \max_{x \in \mathbb{X}} \mathbb{P}_{\hat{\theta}(Y_{1:i})^\tau} (X_i = x \mid Y_{1:i}) \right)_{1 \leq i \leq j} \right)_{j \geq 1}$$

where for all $i \geq 1$, $\hat{\theta}(Y_{1:i}) = (\hat{\nu}^{(i)}, \hat{Q}^{(i)}, (\hat{f}_x^{(i)})_{x \in \mathbb{X}})$ is an estimator using the observations $Y_{1:i}$, and where τ is some permutation of \mathbb{X} . When τ is the identity we shall simply write \hat{h} and \mathfrak{h} . When $i = n$, this estimator will be denoted simply by $\hat{\theta}(Y_{1:n}) = (\hat{\nu}, \hat{Q}, (f_x)_{x \in \mathbb{X}})$. Theorem 4 below controls the excess risk of classification of the plugin classifier. The proof is detailed in Section 6.5. Its cornerstone is the analysis of how errors in estimation of the parameters propagate to errors in the filtering and smoothing distributions proved in [12]. To apply the result therein we introduce the following assumptions.

ASSUMPTION 2. The initial distribution ν is the stationary distribution of \mathbf{X} .

Assumption 2 was not needed in Theorem 2 when comparing the Bayes risk of clustering to the Bayes risk of classification or in Theorem 3 when exhibiting the key quantity Λ . Notice however that under Assumption 2, the second part of Assumption 1 follows from the first part.

ASSUMPTION 3. The emission densities $(f_x)_{x \in \mathbb{X}}$ are compactly supported and $C^* = \int \frac{dy}{\sum_{x \in \mathbb{X}} f_x(y)} < \infty$, the integral is over the union of the supports of the emission densities.

THEOREM 4. For all $\theta \in \Theta^{\text{dep}}$ satisfying assumptions 1, 2 and 3 and for all $n \geq 1$,

$$\begin{aligned} & \mathcal{R}_n^{\text{class, on}}(\theta, \hat{\mathfrak{h}}) - \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class, on}}(\theta, \mathfrak{h}) \\ & \leq \frac{C}{n} \sum_{i=1}^n \left(\mathbb{E}_\theta \left[\delta \rho^{i-1} \|\nu - \hat{\nu}^{(i)}\|_2 + \|Q - \hat{Q}^{(i)}\|_F \right] + \sqrt{C^*} \mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x^{(i)}\|_\infty^2 \right]^{1/2} \right), \end{aligned}$$

and,

$$\begin{aligned} \mathcal{R}_n^{\text{class}}(\theta, \hat{h}) - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) & \leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_F \right] \\ & \quad + \frac{\delta C \sqrt{C^*}}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[(\hat{\rho} \vee \rho)^{2|l-i|} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2}, \end{aligned}$$

where $C = \frac{8(1-\delta)}{\delta^3}$, $\rho = \frac{1-2\delta}{1-\delta}$, $\hat{\rho} = \frac{1-2\hat{\delta}}{1-\hat{\delta}}$ and $\hat{\delta} = \min_{x, x'} \hat{Q}_{x, x'}$.

The proof of Theorem 4 is given in Section 6.5.

To end up with a proper rate for the excess risk, we need to specify an estimator. For finite observation space \mathbb{Y} , that is when the emission distributions are multinomial, a sufficient condition for identifiability is that the vectors of multinomial probabilities given the states are linearly independent, and the transition matrix Q is full-rank and aperiodic. Then, estimation can be done at parametric rate, so that clustering using the plugin method in the Bayes classifier achieves the Bayes risk of clustering up to a c/\sqrt{n} term for some $c > 0$. This procedure is then guaranteed to work well in the case where the Bayes risk is at least larger than c/\sqrt{n} . We introduce now sufficient assumptions to obtain consistent non-parametric estimators.

ASSUMPTION 4. Q is full-rank and aperiodic.

Under Assumptions 1, 2 and 4, the hidden Markov chain is stationary ergodic. The following assumption is sufficient to build estimators using the empirical distribution of the distribution of three consecutive observations.

ASSUMPTION 5. $(f_x)_{x \in \mathbb{X}}$ are linearly independent and belong to $C^s(\mathbb{R})$ the space of locally Hölder continuous functions.

COROLLARY 1. *Let $D_n \rightarrow +\infty$ arbitrarily slowly. There exists a sequence of randomized estimators $(\hat{\theta}_n)_{n \geq 1}$ such that for all $\theta \in \Theta^{\text{dep}}$ satisfying Assumption 1 and Assumptions 2 to 5 there is a sequence of random permutations $(\tau_n)_{n \geq 1}$ such that:*

$$\mathbb{E}[\mathcal{R}_n^{\text{class, on}}(\theta, \hat{h}^{\tau_n})] - \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class, on}}(\theta, \mathfrak{h}) = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right)$$

$$\mathbb{E}[\mathcal{R}_n^{\text{class}}(\theta, \hat{h}^{\tau_n})] - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right)$$

where the expectation $\mathbb{E}[\cdot]$ is understood with respect to the randomness of the algorithm.

The proof of Corollary 1 is given in Section 6.6.

Notice that the rate on the excess risk is exactly the rate of estimation of Hölder regular functions in sup-norm. However, since we do not observe realizations of each emission density separately, this rate is not a straightforward consequence of density estimation theory. See [1] where the usual non-parametric rate for the estimation of densities in sup-norm is obtained in the HMM context.

The estimator in the proof of Corollary 1 is built using the spectral method as in [12] but requires to carefully prove the existence of a random permutation achieving the clustering risk. A straightforward consequence of Corollary 1 and Theorem 2 is the following.

COROLLARY 2. *Let $D_n \rightarrow +\infty$ arbitrarily slowly. There exists a sequence of randomized estimators $(\hat{\theta}_n)_{n \geq 1}$ such that for all $\theta \in \Theta^{\text{dep}}$ satisfying Assumption 1 and Assumptions 2 to 5:*

$$\mathbb{E}[\mathcal{R}_n^{\text{clust, on}}(\theta, \pi_n \circ \hat{\mathfrak{h}})] - \inf_{\mathfrak{g} \in \mathcal{G}^{\text{on}}} \mathcal{R}_n^{\text{clust, on}}(\theta, \mathfrak{g}) = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right)$$

and

$$\mathbb{E}[\mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ \hat{h})] - \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right)$$

where the expectation $\mathbb{E}[\cdot]$ is understood with respect to the randomness of the algorithm.

Notice that the aim of the modified spectral algorithm detailed in Section 6.6 is to guarantee that the same random permutation is used for all parameters. Indeed, in the previous works, the aim was more to get upper bounds on the risks up to label-switching, and the analysis was done for Q and the emission distributions separately. Though we analyse precisely the question for the spectral estimator, we believe that results similar to Corollary 2 hold for most estimation procedures previously proposed in the literature, putting in the upper bounds the bounds on the estimation risk up to label-switching obtained in those works. This is why we shall use in Section 4 the least-squares estimation method for which a public code exists and which we found very efficient [11].

REMARK 3. For the problem of clustering i.i.d. observations coming from a non-parametric mixture with two components, one can easily show that as for the HMM, the Bayes risk has the same order of magnitude as $\int f_0 \wedge f_1$ (Cf. [13] section 2.4). However, the associated Bayes classifier can never be learnt nor approximated from the data because the model is not identifiable and thus, the mixture components can not be estimated. The added-value of Hidden Markov modelling relies on the possibility of constructing consistent estimators of the model parameters which makes it possible to approximate the Bayes classifier using the plugin classifier with these estimators.

4. Numerical simulations. We present here the results of numerical simulations which leverage the added value of non-parametric clustering under hidden Markov modelling. We will consider two examples in which the hidden states will be generated through the same transition matrix $Q = \begin{pmatrix} 0.8 & 0.2 \\ 0.3 & 0.7 \end{pmatrix}$.

EXAMPLE 1. A sample of size $n = 5.10^4$ is generated from two Gaussian mixtures: $\frac{1}{2}(\mathcal{N}(1.7, 0.2) + \mathcal{N}(7, 0.15))$ and $\frac{1}{2}(\mathcal{N}(3.5, 0.2) + \mathcal{N}(5, 0.4))$.

EXAMPLE 2. A sample of size $n = 10^5$ is generated from two Gaussian mixtures: $\frac{1}{2}(\mathcal{N}(3, 0.6) + \mathcal{N}(7, 0.4))$ and $\frac{1}{2}(\mathcal{N}(5, 0.3) + \mathcal{N}(9, 0.4))$.

The first possibility to perform clustering on these samples is to:

- Estimate the model parameters.
- Compute the expectation in (4) using the estimated parameters.
- Repeat the second step for each choice of the classifier h .

The Bayes clusterer would correspond then to the one minimizing the expectation appearing in (4). The problem with this method is that it requires a minimization over the set of all classifiers whose cardinal is 2^n . This method is time consuming given the size of our datasets. This is why we will rather use the plugin classifier. Thanks to Corollary 2, the result of this clustering procedure should not be too far from that of the Bayes clusterer. Our procedure follows the following steps:

- First, the emission densities are estimated using the observations. We use the penalized least squares estimator proposed in [8] whose code has been made public in [11]. We use the histogram basis for the estimation.
- Second, the Forward-Backward algorithm is used to compute the a posteriori distributions of the hidden states under the estimated model parameters and given the observations.
- Third, the hidden states are estimated by maximizing the a posteriori distributions.

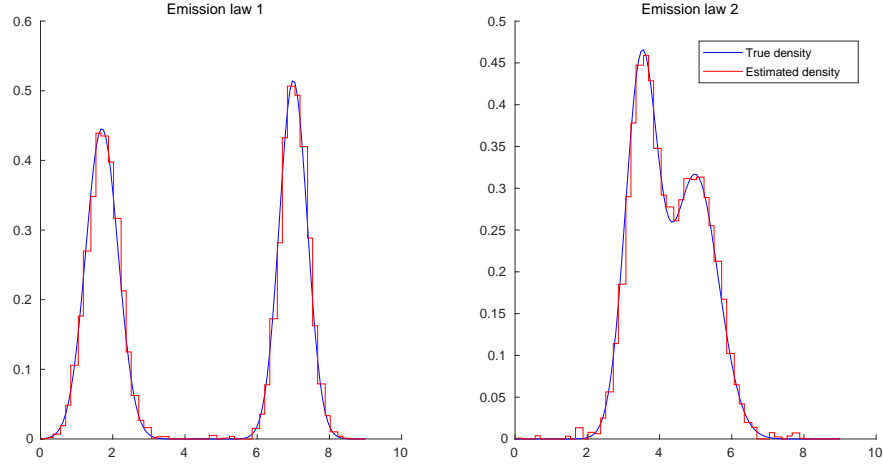
The results of clustering using the Forward-Backward algorithm will be compared to those using the k -means algorithm. Since we have access to the hidden states, the error of clustering can be estimated by choosing the best permutation. Figures 2, 3, and 4 display the results of estimation and clustering. Performance of Bayes classifier, plugin classifier and k -means are reported in Table 1.

If the observations were independent and the emissions modelled non-parametrically, the unique quantity that could have been estimated consistently would be the stationary distribution. However, under the HMM assumption, the estimation of each emission density with the minimax rate is possible. This is due to the identifiability of the model which holds even if the emission densities were not constrained to belong to any specific family. This is not possible in the independent case. Figure 2 shows the estimation results of the emission densities and

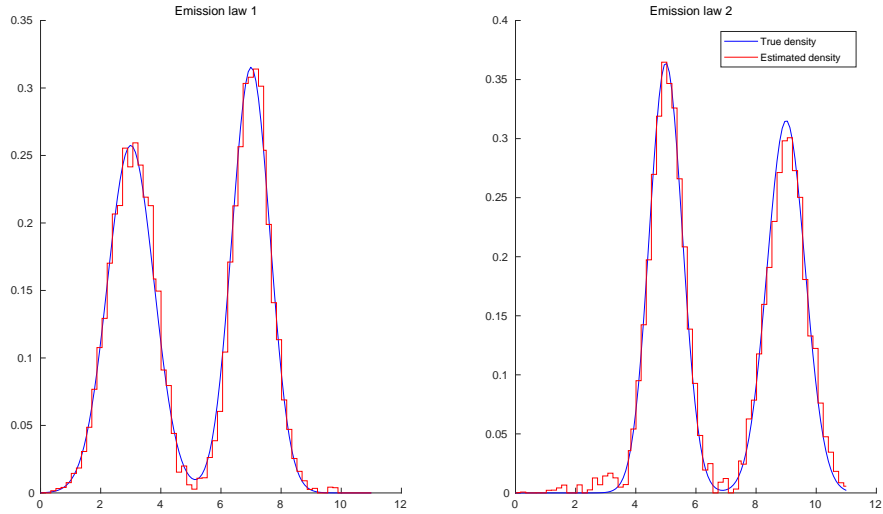
	Bayes classifier	Plugin classifier	k -means algorithm	Λ
Example 1	1.56%	1.61%	46.7%	0.046
Example 2	6.42%	6.51%	47.3%	0.165

TABLE 1

Errors of clustering using three clustering rules: the Bayes classifier (using the true model parameters), the plugin classifier (using the estimated parameters) and the k -means algorithm.



(a) Estimation results for Example 1



(b) Estimation results for Example 2

Figure 2: Non-parametric penalized least squares density estimation using the histogram basis for Example 1 and Example 2

confirms the theoretical properties of the estimator. On the other hand, Theorem 2 proves that in the case of a HMM with two classes, classification errors (and by Theorem 3 clustering error also) could appear only in zones where the two emission densities overlap. In Figures

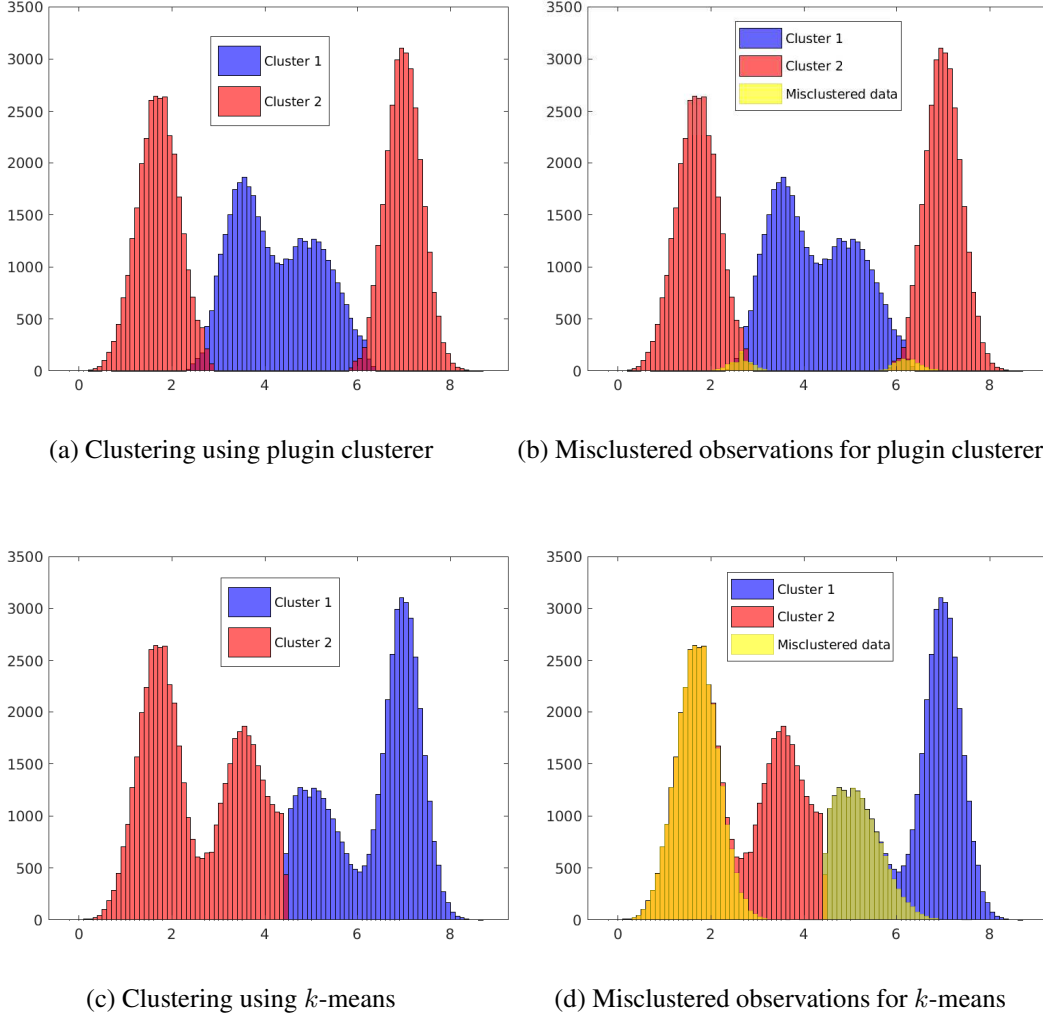


Figure 3: Histograms of clusters and clustering errors for Example 1

3b, and 4b, misclustered observations appear only in the overlaps between the emission densities. Compared to the k -means algorithm which is purely geometric and does not exploit the distribution of the observations, the plugin procedure allows combining together observations even if they are geometrically distant from each other, which is not possible with the k -means algorithm. In this context of Gaussian mixtures, the performance of the k -means algorithm is mediocre as depicted in Table 1 and does not improve significantly when the overlap between the emission densities is small. However, for the plugin procedure, the more separated are the emission densities, the better are the results of clustering. Note that the results of clustering remain the same whether the observations are clustered online or offline.

5. Discussions and Perspectives. This work focuses on an in-depth study of Bayes risk of clustering. This analysis has led us to prove a form of equivalence between the risk of clustering and the risk of classification. After identifying the key quantity which measures the difficulty of the classification task, it was extrapolated to the Bayes risk of clustering in several regimes. Finally, the excess risk of the plugin procedure was studied. Although the analysis is sufficiently detailed to ensure a thorough understanding of risk of clustering, there

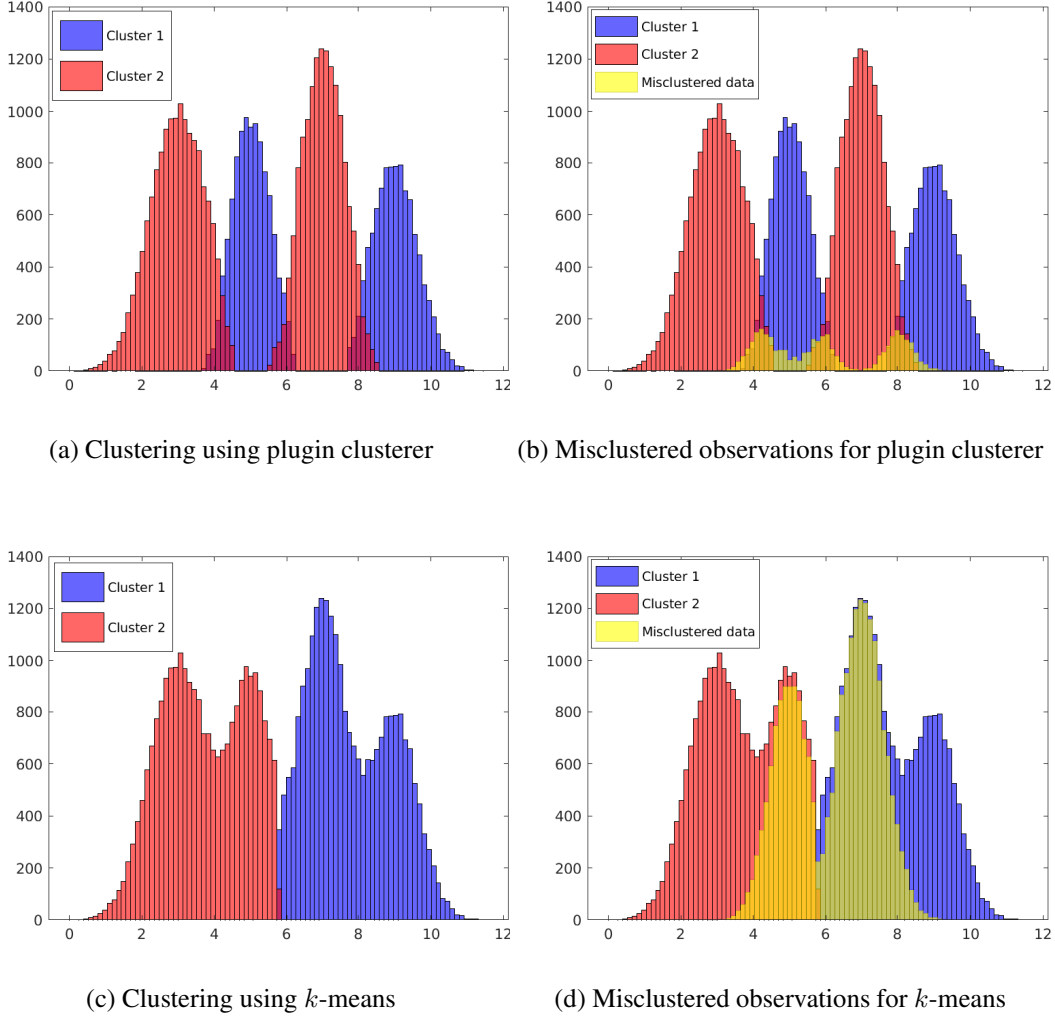


Figure 4: Histograms of clusters and clustering errors for Example 2

are still some interesting questions which were not covered by this work. We give a small overview below.

Lower-bound on entries of Q . Throughout our analysis of the Bayes risk of clustering, we have used Assumption 1 which was crucial in obtaining the lower-bounds. In the absence of such an assumption, the lower-bound of Theorem 3 no more matches the upper-bound and the magnitude of the Bayes risk of classification cannot be precisely understood. The same thing applies to the Bayes risk of clustering. In addition, the optimality of the plugin procedure is no more guaranteed.

Approaching the frontier to independence. This situation happens when the emission distributions are nearly similar or when the transition matrix has almost equal lines. In this case, one can hope to improve the coefficient $\frac{\delta}{1-(J-1)\delta}$ which links the Bayes risk of online and offline classification (Theorem 3). In fact, in this situation, the dependence between the observations is so weak and the effect of future observations on the current classification rule is so negligible that the magnitude of the Bayes risk of classification is the same whether in online or offline frameworks. On the other hand, as shown in [2], estimation of the model

parameters no more becomes possible when approaching the frontier. The plugin procedure should not work as well.

Lower bounds on the Bayes risk of clustering when it is very small. Theorems 1 and 2 establish lower bounds on the Bayes risk of clustering in term of the Bayes risk of classification. When $J = 2$ these bounds are meaningful regardless of how small is the Bayes risk of clustering. When $J = 3$, however, these bounds can be vacuous if the Bayes risk of classification gets too small. This is not an artifact of our bounds since we have shown in Proposition 1 that the two Bayes risks are not uniformly comparable when $J \geq 3$. Whence from the current work we only know that the Bayes risk of clustering is driven by Λ in the region of parameters for which it is not exponentially small in n and that it cannot be driven by Λ otherwise. If it was the case, then it would be equivalent to the Bayes risk of classification in contradiction with the Proposition 1. Understanding the the Bayes risk of clustering in the region of extreme parameters is still an open question.

High dimension. The results concerning the excess risk of the plugin procedure (Corollary 1 and Corollary 2) no more hold in high dimension, that is when the dimension of the observations grows with n . The reason is that the estimators of the HMM parameters are not guaranteed to work in high dimension and the constants appearing in the control of the errors of estimation may blow up. However, the other results still apply in this context. Notably, the relationship between classification and clustering risks remains unaffected by observation space dimensionality (Theorem 1 and Theorem 2). Furthermore, the Bayes risk of classification is driven by the same measure of separation Λ regardless of the dimension of the observations (Theorem 3).

Fast rates. Corollary 2 has interest mainly in the situation where $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \gtrsim (\log(n)/n)^{\frac{s}{2s+1}}$. In this regime, Corollary 2 tells us that the plugin procedure has a risk of the same magnitude. However, in the situation where the magnitude of the Bayes risk of clustering is much smaller, that is when the emission distributions are very separated, one can hope to obtain faster rates. For example, when $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) = \mathcal{O}(e^{-cn})$ for c a positive absolute constant, one can hope to show that the risk of the plugin procedure is exponentially small in n . The following lemma represents a first step for the proof of such a result.

LEMMA 1. *For all $0 < \gamma < 1/2$ and all $\theta \in \Theta$*

$$(9) \quad \mathcal{R}_n^{\text{class}}(\theta, \hat{h}) \leq \frac{\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)}{1/2 - \gamma} + \frac{1}{n} \sum_{i=1}^n \mathbb{P}_\theta \left(\left\| \phi_{\theta, i|n} - \phi_{\hat{\theta}, i|n} \right\|_{\text{TV}} > \gamma \right).$$

where \hat{h} is the plugin classifier defined in Section 3.4.

The proof of Lemma 1 is given in Section 6.7.

Observe that the second term of the rhs of (9) is a large deviation term which may eventually decrease exponentially fast in n . The only price to pay to obtain large deviation type of decay is a constant factor of at least 2 in front of $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$. In situations where $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$ is small, this might be advantageous compared to the bound in Corollary 1. Allegedly, the inequality (16) can be used to control the second term in the rhs of (9). This would require to derive large deviation inequalities for all the terms involved in (16), which turns out to be a rather challenging problem.

Optimal excess risk. Although we obtain upper bounds on the excess risk of the plugin procedure (Corollaries 1 and 2), we do not know the optimal rate of decay of the excess risk. In particular, it is unknown if the plugin procedure achieves optimal excess risk. The Lemma 1 suggests that when the Bayes risk is smaller than $\mathcal{O}(n^{-s/(2s+1)})$ then our upper bounds on the excess risk of the plugin could be improved. Yet without optimality guarantees. Determining the optimal excess risk is an open and interesting question to investigate.

Alternatives to plugin. Under the hidden Markov modeling, the most straightforward way to take advantage of the identifiability of the model is to estimate the model parameters and use them for clustering through the plugin procedure. Unlike the i.i.d. case where algorithms such as k -means can be used as an alternative, we do not know of any alternative to the plugin in the HMM case. It would be nice to find clustering procedures that leverage the nonparametric identifiability of HMM without relying on estimating the parameters first.

6. Proofs.

6.1. *Preliminaries.* In this section we shall recall basic results for HMMs that can be found in [7] about the distribution of the hidden states given a set of observations.

For any parameter θ , any integers $k, i \leq j$, the distribution of X_k given $Y_{i:j}$ under \mathbb{P}_θ will be denoted $\phi_{\theta,k|i:j}(\cdot, Y_{i:j})$. For any integers $i \leq n$, we shall simplify the so-called filtering distribution $\phi_{\theta,n|i:n}(\cdot, Y_{i:n})$ to $\phi_{\theta,n}(\cdot, Y_{i:n})$.

Conditional on observations $Y_{i:n}$, the sequence of the hidden states is an inhomogeneous Markov chain, with transition matrices called *forward kernels*. For each $k \leq n-1$, the forward kernel is denoted $(F_{\theta,k|n}[Y_{k+1:n}])$ to emphasize that it only depends on $Y_{k+1:n}$. When $k \geq n$, the kernel does not depend on the observations and is equal to the transition matrix Q , so that $F_{\theta,k|n}[Y_{k+1:n}] := Q$ for $k \geq n$. In other words, for any $n \in \mathbb{Z}$, for any index $i \leq n$ and $k \geq i$ and any real-valued function f on \mathbb{X} (understood as a vector in \mathbb{R}^J),

$$\mathbb{E}_\theta[f(X_{k+1}) | X_{i:k}, Y_{i:n}] = F_{\theta,k|n}[Y_{k+1:n}]f = \sum_{x \in \mathbb{X}} F_{\theta,k|n}[Y_{k+1:n}](X_k, x)f(x).$$

Conditional on observations $Y_{i:n}$, the reverse time sequence of hidden states is also an inhomogeneous Markov chain with transition matrices $(B_{\theta,k}[Y_{i:k}])_{k \leq n-1}$ called *backward kernels*. In other words, for any $n \in \mathbb{Z}$, $i \leq k \leq n-1$ and any function f on \mathbb{X} :

$$\mathbb{E}_\theta[f(X_k) | X_{k+1:n}, Y_{i:n}] = B_{\theta,k}[Y_{i:k}]f = \sum_{x \in \mathbb{X}} B_{\theta,k}[Y_{i:k}](X_{k+1}, x)f(x).$$

Here, the backward kernel $B_{\theta,k}[Y_{i:k}]$ depends only on the observations up to time k . It is given by:

$$(10) \quad B_{\theta,k}[Y_{i:k}](\tilde{x}, x) = \frac{\phi_{\theta,k}(x, Y_{i:k})Q(x, \tilde{x})}{\sum_{x' \in \mathbb{X}} \phi_{\theta,k}(x', Y_{i:k})Q(x', \tilde{x})}.$$

Note that the denominator is always positive thanks to Assumption 1.

For any transition kernel T , we denote $\delta(T)$ is the Dobrushin coefficient of T defined by:

$$\delta(T) = \sup_{(x, x') \in \mathbb{X} \times \mathbb{X}} \|T(x, \cdot) - T(x', \cdot)\|_{\text{TV}}$$

where $\|\cdot\|_{\text{TV}}$ is the total variation norm. We recall the following two lemmas which can be found in [7].

LEMMA 2. *Let f be a real-valued function over $(\mathbb{X}, \mathcal{P}(\mathbb{X}))$. Let ξ and ξ' be two probability measures over $(\mathbb{X}, \mathcal{P}(\mathbb{X}))$ and let T be a transition kernel from $(\mathbb{X}, \mathcal{P}(\mathbb{X}))$ to $(\mathbb{Y}, \mathcal{Y})$. Then one has:*

- $|\xi f - \xi' f| \leq \|\xi - \xi'\|_{TV} \text{osc}(f)$
- $\|\xi T - \xi' T\|_{TV} \leq \delta(T) \|\xi - \xi'\|_{TV}$

where $\text{osc}(f) = \sup_{(x,x') \in \mathbb{X} \times \mathbb{X}} |f(x) - f(x')|$.

To end with, notice that under Assumption 1, for any subset A of \mathbb{X} ,

$$J\delta\gamma(A) \leq \sum_{x' \in A} Q(x, x') \leq J(1 - (J-1)\delta)\gamma(A)$$

with γ the uniform distribution over \mathbb{X} . Using Lemma 4.3.13 in [7], this leads to

LEMMA 3. *Under Assumption 1, for any (possibly non positive) integers k and n , the Dobrushin coefficient of the forward kernel $F_{\theta,k|n}$ satisfies:*

$$\delta(F_{\theta,k|n}) \leq \begin{cases} \rho_0 & k < n \\ \rho_1 & k \geq n \end{cases}$$

with $\rho_0 = 1 - \frac{J\delta}{J(1-(J-1)\delta)} = \frac{1-J\delta}{1-(J-1)\delta}$ and $\rho_1 = 1 - J\delta$.

Using Equation (10) we get that, under Assumption 1, for any (possibly non positive) integers $i \leq k$,

$$\forall x \in \mathbb{X}, \quad B_{\theta,k}[Y_{i:k}](x, \cdot) \geq \frac{\delta}{1 - (J-1)\delta} \phi_{\theta,k}(\cdot, Y_{i:k}),$$

so that using applying Lemma 4.3.13 in [7],

LEMMA 4. *Under Assumption 1, for any (possibly non positive) integers $i \leq k \leq n-1$, the Dobrushin coefficient of the backward kernel $B_{\theta,k}[Y_{i:k}]$ satisfies:*

$$\delta(B_{\theta,k}[Y_{i:k}]) \leq 1 - \frac{\delta}{1 - (J-1)\delta} = \rho_0.$$

6.2. *Proofs of Theorems 1 and 2.* The proofs of Theorems 1 and 2 in the online framework follow the exact same lines of those in the offline case without any change. In addition, the constants appearing in the bounds remain unchanged. We will therefore only present the proofs of the results in the offline framework.

6.2.1. *Common elements to the proof of Theorems 1 and 2.* Recall the definition of $\mathcal{R}_n^{\text{MRSS}}$ in Remark 2: for all $(\theta, h) \in \Theta \times \mathcal{H}_n$,

$$(11) \quad \mathcal{R}_n^{\text{MRSS}}(\theta, h) = \mathbb{E}_\theta \left[\min_{\tau \in \mathcal{S}_J} \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n}) \right].$$

where $U_{n,\tau}(h) := \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau(X_i) \neq h_i(Y_{1:n})}$.

Instead of comparing $\mathcal{R}_n^{\text{clust}}$ and $\mathcal{R}_n^{\text{class}}$, we compare $\mathcal{R}_n^{\text{clust}}$ and $\mathcal{R}_n^{\text{MRSS}}$, which is enough to obtain the result thanks to the following easy lemma.

LEMMA 5. *For all $\theta \in \Theta$ and all $n \geq 1$*

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) = \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{MRSS}}(\theta, h).$$

PROOF. The optimal permutation $\hat{\tau}$ such that $\hat{p}_{\hat{\tau}}(h) = \min_{\tau \in \mathcal{S}_J} \mathbb{E}_{\theta}(U_{n,\tau}(h) \mid Y_{1:n})$ is a $Y_{1:n}$ -measurable permutation valued random variable. Since any $h \in \mathcal{H}_n$ is also $Y_{1:n}$ -measurable, the result is immediate. \square

In the next we then focus on comparing $\mathcal{R}_n^{\text{clust}}$ and $\mathcal{R}_n^{\text{MRSS}}$. We let $N_j := \sum_{i=1}^n \mathbf{1}_{\{X_i=j\}}$ and $N_{(1)} \leq N_{(2)} \leq \dots \leq N_{(J)}$ denote the order statistics of (N_1, \dots, N_J) . We also make use of the notation $\hat{p}_{\tau}(h) := \mathbb{E}_{\theta}(U_{n,\tau}(h) \mid Y_{1:n})$. We also let $\hat{\tau}_h$ denote a $Y_{1:n}$ -measurable permutation satisfying:

$$\hat{p}_{\hat{\tau}_h}(h) = \mathbb{E}_{\theta}(U_{n,\hat{\tau}_h}(h) \mid Y_{1:n}) = \min_{\tau} \mathbb{E}_{\theta}(U_{n,\tau}(h) \mid Y_{1:n}) = \min_{\tau} \hat{p}_{\tau}(h).$$

PROPOSITION 2. *A generic lower bound that works for any latent model (i.i.d. or HMM or whatever). For all classifiers h , all ε and all η*

$$\begin{aligned} \mathbb{E}_{\theta} \left(\min_{\tau} U_{n,\tau}(h) \right) &\geq \mathbb{E}_{\theta} \left(\min_{\tau} \mathbb{E}_{\theta}(U_{n,\tau}(h) \mid Y_{1:n}) \right) \\ &\quad - \mathbb{E}_{\theta} \left(\mathbb{E}_{\theta} \left(\max_{\tau} [-U_{n,\tau}(h) + \hat{p}_{\tau}(h)] \mid Y_{1:n} \right) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}} \right) \\ &\quad - \mathbb{E}_{\theta} \left(\mathbb{P}_{\theta}(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right) \\ &\quad - \mathbb{P}_{\theta}(N_{(1)} + N_{(2)} < 2n\eta). \end{aligned}$$

PROOF. For $\varepsilon \in [0, 1]$, we decompose

$$(12) \quad \mathbb{E}_{\theta} \left(\min_{\tau} U_{n,\tau}(h) \right) = \mathbb{E}_{\theta} \left(\min_{\tau} U_{n,\tau}(h) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}} \right) + \mathbb{E}_{\theta} \left(\min_{\tau} U_{n,\tau}(h) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right).$$

The first term in the rhs of (12) is handled via a small deviation principle, remarking that on the event $\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}$

$$\begin{aligned} \mathbb{E}_{\theta} \left(\min_{\tau} U_{n,\tau}(h) \mid Y_{1:n} \right) &= \mathbb{E}_{\theta} \left(\min_{\tau} [\hat{p}_{\tau}(h) + U_{n,\tau}(h) - \hat{p}_{\tau}(h)] \mid Y_{1:n} \right) \\ &\geq \min_{\tau} \hat{p}_{\tau}(h) - \mathbb{E}_{\theta} \left(\max_{\tau} [-U_{n,\tau}(h) + \hat{p}_{\tau}(h)] \mid Y_{1:n} \right). \end{aligned}$$

The second term in (12) is handled via a large deviation principle. By Lemma 6, on the event $\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}$

$$\begin{aligned} \mathbb{E}_{\theta} \left(\min_{\tau} U_{n,\tau}(h) \mid Y_{1:n} \right) &\geq \mathbb{E}_{\theta} \left(\min_{\tau} U_{n,\tau}(h) \mathbf{1}_{\{U_{n,\hat{\tau}_h}(h) \leq (N_{(1)} + N_{(2)})/(2n)\}} \mid Y_{1:n} \right) \\ &= \mathbb{E}_{\theta} \left(U_{n,\hat{\tau}_h}(h) \mathbf{1}_{\{U_{n,\hat{\tau}_h}(h) \leq (N_{(1)} + N_{(2)})/(2n)\}} \mid Y_{1:n} \right) \\ &\geq \hat{p}_{\hat{\tau}_h}(h) - \mathbb{P}_{\theta} \left(U_{n,\hat{\tau}_h}(h) > \frac{N_{(1)} + N_{(2)}}{2n} \mid Y_{1:n} \right). \end{aligned}$$

But,

$$\mathbb{P}_{\theta} \left(U_{n,\hat{\tau}_h}(h) > \frac{N_{(1)} + N_{(2)}}{2n} \mid Y_{1:n} \right) \leq \mathbb{P}_{\theta}(U_{n,\hat{\tau}_h} > \eta \mid Y_{1:n}) + \mathbb{P}_{\theta}(N_{(1)} + N_{(2)} < 2n\eta \mid Y_{1:n}).$$

The generic lower bound follows. \square

LEMMA 6. *If $U_{n,\tau'}(h) \leq \frac{N_{(1)} + N_{(2)}}{2n}$ then $\min_{\tau} U_{n,\tau}(h) = U_{n,\tau'}(h)$.*

PROOF. Let $U_{n,\tau'}(h) \leq \frac{N_{(1)}+N_{(2)}}{2n}$ and suppose $\min_{\tau} U_{n,\tau}(h) < U_{n,\tau'}(h)$. Then, there exists a permutation $\tau'' \neq \tau'$ such that $U_{n,\tau''}(h) < U_{n,\tau'}(h)$. But then, letting $I = \{i \in \{1, \dots, n\} : \tau'(X_i) = h_i(Y_{1:n})\}$:

$$\begin{aligned}
n(U_{n,\tau'}(h) - U_{n,\tau''}(h)) &= \sum_{i=1}^n \left(\mathbf{1}_{\{\tau'(X_i) \neq h_i(Y_{1:n})\}} - \mathbf{1}_{\{\tau''(X_i) \neq h_i(Y_{1:n})\}} \right) \\
&= \sum_{i=1}^n \left(\mathbf{1}_{\{\tau''(X_i) = h_i(Y_{1:n})\}} - \mathbf{1}_{\{\tau'(X_i) = h_i(Y_{1:n})\}} \right) \\
&= - \sum_{i \in I} \mathbf{1}_{\{\tau'(X_i) \neq \tau''(X_i)\}} + \sum_{i \in I^c} \mathbf{1}_{\{\tau'(X_i) = h_i(Y_{1:n})\}} \\
&= - \sum_{i=1}^n \mathbf{1}_{\{\tau'(X_i) \neq \tau''(X_i)\}} + \sum_{i \in I^c} \left(\mathbf{1}_{\{\tau'(X_i) = h_i(Y_{1:n})\}} + \mathbf{1}_{\{\tau'(X_i) \neq \tau''(X_i)\}} \right) \\
&\leq -(N_{(1)} + N_{(2)}) + 2|I^c| \\
&= -(N_{(1)} + N_{(2)}) + 2nU_{n,\tau'}(h)
\end{aligned}$$

where we have used that since $\tau' \neq \tau''$, it must be that $\sum_{i=1}^n \mathbf{1}_{\{\tau'(X_i) \neq \tau''(X_i)\}} \geq N_{(1)} + N_{(2)}$. Rearranging the previous:

$$U_{n,\tau''}(h) \geq \frac{N_{(1)} + N_{(2)}}{n} - U_{n,\tau'}(h) \geq U_{n,\tau'}(h).$$

which contradicts that $U_{n,\tau''}(h) < U_{n,\tau'}(h)$. Hence $\min_{\tau} U_{n,\tau}(h) \geq U_{n,\tau'}(h)$. \square

6.2.2. *Proof of Theorem 1 (independent scenario).* Here we apply the result of Proposition 2 to the i.i.d. case. The first trivial bound is obtained by choosing $\varepsilon = \eta = 0$. With this choice, Proposition 2 gives

$$\begin{aligned}
\mathbb{E}_{\theta} \left(\min_{\tau} U_{n,\tau}(h) \right) &\geq \mathbb{E}_{\theta} \left(\min_{\tau} \mathbb{E}_{\theta}(U_{n,\tau}(h) \mid Y_{1:n}) \right) - \mathbb{E}_{\theta} \left(\mathbb{E}_{\theta} \left(\max_{\tau} [-U_{n,\tau}(h) + \hat{p}_{\tau}(h)] \mid Y_{1:n} \right) \right) \\
&\geq \mathbb{E}_{\theta} \left(\min_{\tau} \mathbb{E}_{\theta}(U_{n,\tau}(h) \mid Y_{1:n}) \right) - \sqrt{\frac{\log(J!)}{2n}}
\end{aligned}$$

by Lemma 7 below. The other bounds are obtained using Lemma 7 and Markov's inequality to establish that for any $\varepsilon > 0$

$$\begin{aligned}
&\mathbb{E}_{\theta} \left(\mathbb{E}_{\theta} \left(\max_{\tau} [-U_{n,\tau}(h) + \hat{p}_{\tau}(h)] \mid Y_{1:n} \right) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}} \right) \\
&\leq \sqrt{\frac{\log(J!)}{2n}} \mathbb{P}_{\theta}(\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon) \\
&\leq \frac{1}{\varepsilon} \sqrt{\frac{\log(J!)}{2n}} \mathbb{E}_{\theta} \left(\min_{\tau} \mathbb{E}_{\theta}(U_{n,\tau}(h) \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}} \right)
\end{aligned}$$

and then using Lemma 8 to establish that for all $\varepsilon, \eta > 0$

$$\begin{aligned}
\mathbb{E}_{\theta} \left(\mathbb{P}_{\theta}(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right) &\leq \mathbb{E}_{\theta} \left(\hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)} \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right) \\
&\leq \frac{e}{\eta} \left(\frac{e\varepsilon}{\eta} \right)^{n\eta-1} \mathbb{E}_{\theta} \left(\min_{\tau} \mathbb{E}_{\theta}(U_{n,\tau}(h) \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right).
\end{aligned}$$

From here, observe that when $J = 2$ it must be that $\mathbb{P}_\theta(N_{(1)} + N_{(2)} < n) = 0$. Then, the bound in this case follows by taking $\eta \rightarrow \frac{1}{2}$ [by below] and $\varepsilon = \frac{1}{2e}[\log(J!)/(2n)]^{1/n}$.

When $J \geq 3$, Lemma 9 can be used to find that

$$\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) \leq J^2 e^{-\frac{n(\beta-2\eta)^2}{2\beta}}.$$

Then the remaining bound is obtained by choosing $\eta = \beta/4$ and $\varepsilon = \frac{\beta}{4e}[\log(J!)/(2n)]^{2/(n\beta)}$.

LEMMA 7. *For all θ , \mathbb{P}_θ -almost-surely*

$$\mathbb{E}_\theta \left(\max_{\tau} [-U_{n,\tau}(h) + \hat{p}_\tau(h)] \mid Y_{1:n} \right) \leq \sqrt{\frac{\log(J!)}{2n}}.$$

PROOF. For any $\lambda > 0$,

$$\begin{aligned} \mathbb{E}_\theta \left[\sup_{\tau} \{-U_{n,\tau}(h) + \hat{p}_\tau(h)\} \right] &\leq \frac{1}{\lambda} \log \left(\mathbb{E}_\theta \left[\exp \left(\sup_{\tau} \{-\lambda(U_{n,\tau}(h) - \hat{p}_\tau(h))\} \right) \mid Y_{1:n} \right] \right) \\ &= \frac{1}{\lambda} \log \left(\mathbb{E}_\theta \left[\sup_{\tau} \exp(-\lambda(U_{n,\tau}(h) - \hat{p}_\tau(h))) \mid Y_{1:n} \right] \right) \\ &\leq \frac{1}{\lambda} \log \left(J! \sup_{\tau} \mathbb{E}_\theta \left[\exp(-\lambda(U_{n,\tau}(h) - \hat{p}_\tau(h))) \mid Y_{1:n} \right] \right) \\ &\leq \frac{1}{\lambda} \log \left(J! \exp \left(\frac{\lambda^2}{8} \times n \times \left(\frac{1}{n} \right)^2 \right) \right) \text{ (Hoeffding's lemma)} \\ &\leq \inf_{\lambda > 0} \left\{ \frac{\log(J!)}{\lambda} + \frac{\lambda}{8n} \right\} \\ &\leq \sqrt{\frac{\log(J!)}{2n}}. \end{aligned}$$

Hoeffding's lemma applies because conditionally to the sequence of observations $Y_{1:n}$, the labels $X_{1:n}$ are still independent. \square

LEMMA 8. *For all $\theta \in \Theta$, \mathbb{P}_θ -almost-surely*

$$\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \leq \hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)}.$$

PROOF. By Chernoff's bound [with $q_i(h) = \mathbb{P}_\theta(h_i(Y_{1:n}) \neq \hat{\tau}_h(X_i) \mid Y_{1:n})$]:

$$\begin{aligned} \mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) &= \mathbb{P}_\theta \left(\sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \hat{\tau}_h(X_i)} > n\eta \mid Y_{1:n} \right) \\ &\leq \inf_{\lambda > 0} \exp \left(-\lambda n\eta + \sum_{i=1}^n \log \left(q_i(h) e^\lambda + 1 - q_i(h) \right) \right) \\ &\leq \inf_{\lambda > 0} \exp \left(-\lambda n\eta + n\hat{p}_{\hat{\tau}_h}(h)(e^\lambda - 1) \right) \\ &= \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta} e^{-n\hat{p}_{\hat{\tau}_h}(h)} \\ &\leq \hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)}. \end{aligned}$$

□

LEMMA 9. Let $\beta = \min_{j \neq k} (\nu_j + \nu_k)$. If $J \geq 3$, then

$$\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) \leq J^2 e^{-\frac{n(\beta-2\eta)^2}{2\beta}}.$$

PROOF. If $J = 2$, remark that $N_{(1)} + N_{(2)} = n$ P_θ -almost-surely, hence the result. Now we assume that $J \geq 3$. It holds,

$$\begin{aligned} \mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) &= P_\theta(\exists j \neq k, N_j + N_k \leq 2n\eta) \\ &\leq J^2 \max_{j \neq k} \mathbb{P}_\theta(N_j + N_k \leq 2n\eta). \end{aligned}$$

Then observe that $N_j + N_k = \sum_{i=1}^n (\mathbf{1}_{X_i=j} + \mathbf{1}_{X_i=k}) = \sum_{i=1}^n \mathbf{1}_{X_i \in \{j,k\}}$ whenever $j \neq k$. In other words, when $j \neq k$ the random variables $N_j + N_k$ has a Binomial distribution with parameters $(n, \nu_j + \nu_k)$ under \mathbb{P}_θ . The conclusion follows using Chernoff's bound on the Binomial distribution [recall the Binomial distribution is sub-Gaussian on the left-tail]. □

6.2.3. *Proof of Theorem 2 (dependent scenario).* We apply the result of Proposition 2 to the HMM case. The first trivial bound is obtained by choosing $\varepsilon = \delta = 0$. Proposition 2 yields:

$$\begin{aligned} \mathbb{E}_\theta\left(\min_{\tau} U_{n,\tau}(h)\right) &\geq \mathbb{E}_\theta\left(\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n})\right) - \mathbb{E}_\theta\left(\mathbb{E}_\theta\left(\max_{\tau} [\hat{p}_{\tau}(h) - U_{n,\tau}(h)] \mid Y_{1:n}\right)\right) \\ &\geq \mathbb{E}_\theta\left(\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n})\right) - \frac{1}{1-\rho_0} \sqrt{\frac{\log(J!)}{2n}} \end{aligned}$$

by Lemma 10 below. The first inequality follows by taking the infimum over h on both sides.

The other bounds are obtained using Lemma 10 and Markov's inequality as in the independent case:

$$\begin{aligned} &\mathbb{E}_\theta\left(\mathbb{E}_\theta\left(\max_{\tau} [-U_{n,\tau}(h) + \hat{p}_{\tau}(h)] \mid Y_{1:n}\right) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}}\right) \\ &\leq \frac{1}{1-\rho_0} \sqrt{\frac{\log(J!)}{2n}} \mathbb{P}_\theta(\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon) \\ &\leq \frac{1}{\varepsilon(1-\rho_0)} \sqrt{\frac{\log(J!)}{2n}} \mathbb{E}_\theta\left(\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}}\right) \end{aligned}$$

and then using Lemma 11 to establish that for all $\varepsilon, \eta > 0$

$$\begin{aligned} &\mathbb{E}_\theta\left(\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}}\right) \\ &\leq \mathbb{E}_\theta\left(\hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{1-(J-1)\delta}{\delta}\right)^{2n} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta}\right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)} \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}}\right) \\ &\leq \frac{e}{\eta} \left(\frac{1-(J-1)\delta}{\delta}\right)^{2n} \left(\frac{e\varepsilon}{\eta}\right)^{n\eta-1} \mathbb{E}_\theta\left(\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}}\right). \end{aligned}$$

From here, observe that when $J = 2$ it must be that $\mathbb{P}_\theta(N_{(1)} + N_{(2)} < n) = 0$. Then, the bound in this case follows by taking $\eta \rightarrow \frac{1}{2}$ [by below] and $\varepsilon = \frac{1}{2e} \left(\frac{\delta}{1-\delta}\right)^4 [\log(J!)/(2n)]^{1/n}$.

For the remaining inequality, when $J \geq 3$, Lemma 12 ensures:

$$\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) \leq J^2 e^{-2n(1-\rho_1)^2(\beta-2\eta)^2}.$$

On the other hand, we have by Lemma 10:

$$\mathbb{E}_\theta \left[\max_{\tau} \{ \hat{p}_\tau(h) - U_{n,\tau}(h) \} \mid Y_{1:n} \right] \leq \frac{1}{1-\rho_0} \sqrt{\frac{\log(J!)}{2n}}$$

On the event: $\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}$, for $\varepsilon < \eta$:

$$\begin{aligned} \mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) &= \mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) - \hat{p}_{\hat{\tau}_h}(h) > \eta - \hat{p}_{\hat{\tau}_h}(h) \mid Y_{1:n}) \\ &\leq \mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) - \hat{p}_{\hat{\tau}_h}(h) > \eta - \varepsilon \mid Y_{1:n}) \\ &\leq e^{-\lambda(\eta-\varepsilon)} \mathbb{E}_\theta \left[e^{\lambda\{U_{n,\hat{\tau}_h}(h) - \hat{p}_{\hat{\tau}_h}(h)\}} \mid Y_{1:n} \right] \\ &\leq e^{-\lambda(\eta-\varepsilon)} e^{\frac{1}{8n} \left(\frac{\lambda}{1-\rho_0} \right)^2} \end{aligned}$$

where the last inequality is due to the argument using Marton coupling as shown in the proof of Lemma 10. Taking the minimum over λ yields:

$$\mathbb{E}_\theta \left[\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \mathbf{1}_{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon} \right] \leq e^{-2n(1-\rho_0)^2(\eta-\varepsilon)^2}$$

Using Lemma 12, the final bound reads:

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \right] &\geq \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta[U_{n,\tau}(h) \mid Y_{1:n}] \right] \\ &\quad - \frac{1}{\varepsilon(1-\rho_0)} \sqrt{\frac{\log(J!)}{2n}} \mathbb{E}_\theta \left[\hat{p}_{\hat{\tau}_h}(h) \mathbf{1}_{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon} \right] \\ &\quad - e^{-2n(1-\rho_0)^2(\eta-\varepsilon)^2} - J^2 e^{-2n(1-\rho_1)^2(\beta-2\eta)^2} \end{aligned}$$

Choosing $\varepsilon = \frac{\eta}{2}$ and $\eta = \frac{2}{5}\beta$ and noting that $\rho_1 < \rho_0$ one obtains:

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \right] &\geq \left[1 - \frac{5}{\beta(1-\rho_0)} \sqrt{\frac{\log(J!)}{2n}} \right] \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta[U_{n,\tau}(h) \mid Y_{1:n}] \right] \\ &\quad - (J^2 + 1)e^{-cn(1-\rho_0)^2\beta^2}. \end{aligned}$$

The result follows.

LEMMA 10. Under Assumptions 1 and 2, for all $\theta \in \Theta^{\text{dep}}$, \mathbb{P}_θ -almost-surely

$$\mathbb{E}_\theta \left[\max_{\tau} \{ \hat{p}_\tau(h) - U_{n,\tau}(h) \} \mid Y_{1:n} \right] \leq \frac{1}{1-\rho_0} \sqrt{\frac{\log(J!)}{2n}}$$

where $\rho_0 = \frac{1-J\delta}{1-(J-1)\delta}$.

PROOF. Given that for any $\lambda > 0$,

$$\mathbb{E}_\theta \left[\max_{\tau} \{ \hat{p}_\tau(h) - U_{n,\tau}(h) \} \right] \leq \frac{1}{\lambda} \log \left(J! \max_{\tau} \mathbb{E}_\theta \left[\exp(-\lambda(U_{n,\tau}(h) - \hat{p}_\tau(h))) \mid Y_{1:n} \right] \right)$$

we shall exhibit an upper bound of the rhs term by applying Theorem 2.9 of [39], conditional on $Y_{1:n}$. So that for now we consider $Y = Y_{1:n}$ as fixed. Define $f_Y^{h,\tau}$ for any $x = x_{1:n}$ by:

$$f_Y^{h,\tau}(x) = -\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau(x_i) \neq h_i(Y_{1:n})}.$$

Then, for any $x = x_{1:n}$ and $x' = x'_{1:n}$,

$$f_Y^{h,\tau}(x) - f_Y^{h,\tau}(x') \leq \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{x_i \neq x'_i}.$$

We thus may apply (2.5) in Theorem 2.9 of [39] to get

$$(13) \quad \mathbb{E}_\theta \left[e^{\lambda(f_Y^{h,\tau}(X) - \mathbb{E}_\theta[f_Y^{h,\tau}(X)|Y_{1:n}])} \middle| Y_{1:n} \right] \leq e^{\frac{\lambda^2}{8n^2} \sum_{i=1}^n (\sum_{j=i}^n \Gamma_{i,j})^2},$$

where Γ comes from a Marton coupling (see Definition 2.1 in [39]) and is given by:

$$\Gamma_{j,i} := 0, \quad \Gamma_{i,j} := \sup_{x_1, \dots, x_i, x'_i \in \mathbb{X}} \mathbb{P}_\theta \left[X_{1,j}^{(x_1, \dots, x_i, x'_i, Y_{1:n})} \neq X_{2,j}^{(x_1, \dots, x_i, x'_i, Y_{1:n})} \middle| Y_{1:n} \right]$$

for $1 \leq i < j \leq n$. Now,

$$\begin{aligned} \Gamma_{i,j} \leq \sup_{x_{1:i-1} \in \mathbb{X}^{i-1}} & \left\| \mathbb{P}_\theta \left(X_{1,j}^{(x_{1:i-1}, \tilde{a}, \tilde{b}, Y_{1:n})} \in \cdot \middle| X_{1:i-1} = x_{1:i-1}, Y_{1:n} \right) \right. \\ & \left. - \mathbb{P}_\theta \left(X_{2,j}^{(x_{1:i-1}, \tilde{a}, \tilde{b}, Y_{1:n})} \in \cdot \middle| X_{1:i-1} = x_{1:i-1}, Y_{1:n} \right) \right\|_{\text{TV}}, \end{aligned}$$

ie.,

$$\begin{aligned} \Gamma_{i,j} \leq \sup_{x_{1:i-1} \in \mathbb{X}^{i-1}} & \left\| \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_{1:i-1} = x_{1:i-1}, X_i = \tilde{a}, Y_{1:n} \right) \right. \\ & \left. - \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_{1:i-1} = x_{1:i-1}, X_i = \tilde{b}, Y_{1:n} \right) \right\|_{\text{TV}}, \end{aligned}$$

ie.,

$$\Gamma_{i,j} \leq \left\| \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_i = \tilde{a}, Y_{1:n} \right) - \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_i = \tilde{b}, Y_{1:n} \right) \right\|_{\text{TV}}$$

since conditional on $Y_{1:n}$, the hidden states form a nonhomogeneous Markov chain with transition kernels $(F_{k|n}[Y_{k+1:n}])$. Exponential forgetting of the smoothing distributions in HMMs (Proposition 4.3.26 in [7]) allows to conclude that

$$\left\| \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_i = \tilde{a}, Y_{1:n} \right) - \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_i = \tilde{b}, Y_{1:n} \right) \right\|_{\text{TV}} \leq \rho_0^{j-i}$$

where $\rho_0 = \frac{1-J\delta}{1-(J-1)\delta}$ (see also Lemma 3). By inequality (13):

$$\max_{\tau} \mathbb{E}_\theta \left[e^{\lambda(f_Y^{h,\tau}(X) - \mathbb{E}_\theta[f_Y^{h,\tau}(X)|Y_{1:n}])} \middle| Y_{1:n} \right] \leq e^{\frac{\lambda}{8n} \left(\frac{\lambda}{1-\rho_0} \right)^2}.$$

Thus,

$$\begin{aligned} \mathbb{E}_\theta \left[\max_{\tau} \{ \hat{p}_\tau(h) - U_{n,\tau}(h) \} \middle| Y_{1:n} \right] & \leq \inf_{\lambda > 0} \left\{ \frac{\log(J!)}{\lambda} + \frac{\lambda}{8n(1-\rho_0)^2} \right\} \\ & = \frac{1}{1-\rho_0} \sqrt{\frac{\log(J!)}{2n}} \end{aligned}$$

□

LEMMA 11. For all $\theta \in \Theta$, \mathbb{P}_θ -almost-surely

$$\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \leq \hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)}.$$

PROOF. Let $S = \sum_{i=1}^n \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}$. We consider the following operators defined on $L^\infty(\{0,1\})$, for $i \in [n]$:

$$(M_{\theta,i}.f)(x) := f(0)e^{\lambda \mathbf{1}_{\hat{\tau}_h(0) \neq h_i(Y_{1:n})}} B_{\theta,i}(x,0) + f(1)e^{\lambda \mathbf{1}_{\hat{\tau}_h(1) \neq h_i(Y_{1:n})}} B_{\theta,i}(x,1)$$

where $B_{\theta,i}$ is the Backward kernel defined by:

$$B_{\theta,i}(x,y) = \frac{\phi_{\theta,i}(y)Q(y,x)}{\sum_{y'} \phi_{\theta,i}(y')Q(y',x)}.$$

Then observe that,

$$\begin{aligned} \mathbb{E}_\theta[e^{\lambda S} \mid Y_{1:n}] &= \mathbb{E}_\theta \left[e^{\lambda \sum_{i=2}^n \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} \mathbb{E}_\theta \left[e^{\lambda \mathbf{1}_{\hat{\tau}_h(X_1) \neq h_1(Y_{1:n})}} \mid X_{2:n}, Y_{1:n} \right] \mid Y_{1:n} \right] \\ &= \mathbb{E}_\theta \left[e^{\lambda \sum_{i=2}^n \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} (M_{\theta,1}.\mathbf{1})(X_2) \mid Y_{1:n} \right] \end{aligned}$$

Repeating inductively the same trick leads to

$$\mathbb{E}_\theta[e^{\lambda S} \mid Y_{1:n}] = \mathbb{E}_\theta[(M_{\theta,n} \dots M_{\theta,1}.\mathbf{1})(X_{n+1}) \mid Y_{1:n}]$$

Hence

$$\mathbb{E}_\theta[e^{\lambda S} \mid Y_{1:n}] \leq \|(M_{\theta,n} \dots M_{\theta,1}.\mathbf{1})\|_\infty \leq \prod_{i=1}^n \|M_{\theta,i}\|_\infty$$

where

$$\begin{aligned} \|M_{\theta,i}\|_\infty &:= \sup_{f, \|f\|_\infty=1} \|M_{\theta,i}.f\|_\infty \\ &= \max((M_{\theta,i}.f)(0), (M_{\theta,i}.f)(1)) \\ &\leq \max \left(\sum_{z \in \{0,1\}} e^{\lambda \mathbf{1}_{\hat{\tau}_h(z) \neq h_i(Y_{1:n})}} B_{\theta,i}(0,z), \sum_{z \in \{0,1\}} e^{\lambda \mathbf{1}_{\hat{\tau}_h(z) \neq h_i(Y_{1:n})}} B_{\theta,i}(1,z) \right). \end{aligned}$$

But, given that

$$B_{\theta,i}(x,y) = \frac{\phi_{\theta,i}(y)Q(y,x)}{\sum_{y'} \phi_{\theta,i}(y')Q(y',x)} \leq \frac{1 - (J-1)\delta}{\delta} \phi_{\theta,i}(y)$$

and that

$$\phi_{\theta,i|n}(y) = \phi_{\theta,i+1|n} B_{\theta,i}(y) = \frac{\sum_x \phi_{\theta,i+1|n}(x) \phi_{\theta,i}(y) Q(y,x)}{\sum_{y'} \phi_{\theta,i}(y') Q(y',x)} \geq \frac{\delta \phi_{\theta,i}(y)}{1 - (J-1)\delta}$$

one obtains

$$B_{\theta,i}(x,y) \leq \left(\frac{1 - (J-1)\delta}{\delta} \right)^2 \phi_{\theta,i|n}(y).$$

Thus,

$$\begin{aligned} \|M_{\theta,i}\|_\infty &\leq \left(\frac{1 - (J-1)\delta}{\delta} \right)^2 \left(e^{\lambda \mathbf{1}_{\hat{\tau}_h(0) \neq h_i(Y_{1:n})}} \phi_{i|n}(0) + e^{\lambda \mathbf{1}_{\hat{\tau}_h(1) \neq h_i(Y_{1:n})}} \phi_{i|n}(1) \right) \\ &= \left(\frac{1 - (J-1)\delta}{\delta} \right)^2 \mathbb{E}_\theta \left[e^{\lambda \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} \mid Y_{1:n} \right]. \end{aligned}$$

Finally,

$$\mathbb{E}_\theta[\exp(\lambda S) \mid Y_{1:n}] \leq \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \prod_{i=1}^n \mathbb{E}_\theta \left[e^{\lambda \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} \mid Y_{1:n} \right].$$

One can then use Chernoff's bound [with $q_i(h) = \mathbb{P}_\theta(h_i(Y_{1:n}) \neq \hat{\tau}_h(X_i) \mid Y_{1:n})$]:

$$\begin{aligned} \mathbb{P}_\theta(U_{n, \hat{\tau}_h}(h) > \eta \mid Y_{1:n}) &= \mathbb{P}_\theta \left(\sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \hat{\tau}_h(X_i)} > n\eta \mid Y_{1:n} \right) \\ &\leq \inf_{\lambda > 0} e^{-\lambda n\eta} \mathbb{E}_\theta \left[e^{\lambda S} \mid Y_{1:n} \right] \\ &\leq \inf_{\lambda > 0} e^{-\lambda n\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \prod_{i=1}^n \mathbb{E}_\theta \left[e^{\lambda \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} \mid Y_{1:n} \right] \\ &\leq \inf_{\lambda > 0} e^{-\lambda n\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} e^{\sum_{i=1}^n \log(q_i(h) e^\lambda + 1 - q_i(h))} \\ &\leq \inf_{\lambda > 0} e^{-\lambda n\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} e^{n\hat{p}_{\hat{\tau}_h}(h)(e^\lambda - 1)} \\ &\leq \hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta - 1} e^{-n\hat{p}_{\hat{\tau}_h}(h)} \end{aligned}$$

□

LEMMA 12. Let $\beta = \min_{i,j \neq k} \mathbb{P}_\theta(X_i \in \{j, k\})$ and $\rho_1 = 1 - J\delta$. If $J \geq 3$, then for $\eta < \frac{\beta}{2}$

$$\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) \leq J^2 e^{-2n(1-\rho_1)^2(\beta-2\eta)^2}.$$

PROOF. Let $\lambda > 0$, $\eta < \frac{\beta}{2}$, $j \neq k$, $\beta_i(j, k) = \mathbb{P}_\theta(X_i \in \{j, k\})$.

$$\begin{aligned} \mathbb{P}_\theta(N_j + N_k < 2n\eta) &= \mathbb{P}_\theta \left(\sum_{i=1}^n (\mathbf{1}_{X_i \in \{j, k\}} - \beta_i(j, k)) < 2n\eta - \sum_{i=1}^n \beta_i(j, k) \right) \\ &= \mathbb{P}_\theta \left(e^{\lambda \sum_{i=1}^n (\beta_i(j, k) - \mathbf{1}_{X_i \in \{j, k\}})} > e^{\lambda (\sum_{i=1}^n \beta_i(j, k) - 2n\eta)} \right) \\ &\leq e^{-\lambda (\sum_{i=1}^n \beta_i(j, k) - 2n\eta)} \mathbb{E}_\theta \left[e^{\lambda \sum_{i=1}^n (\beta_i(j, k) - \mathbf{1}_{X_i \in \{j, k\}})} \right] \end{aligned}$$

$\mathbb{E}_\theta \left[e^{\lambda \sum_{i=1}^n (\beta_i(j, k) - \mathbf{1}_{X_i \in \{j, k\}})} \right]$ can be controlled by the same technique using the Marton coupling that was used in the proof of Lemma 10:

$$\mathbb{E}_\theta \left[e^{\lambda \sum_{i=1}^n (\beta_i(j, k) - \mathbf{1}_{X_i \in \{j, k\}})} \right] \leq e^{\frac{(n\lambda)^2}{8n(1-\rho_1)^2}} = e^{\frac{n\lambda^2}{8(1-\rho_1)^2}}$$

It follows that

$$\begin{aligned} \mathbb{P}_\theta(N_j + N_k < 2n\eta) &\leq \inf_{\lambda > 0} e^{-\lambda (\sum_{i=1}^n \beta_i(j, k) - 2n\eta) + \frac{n\lambda^2}{8(1-\rho_1)^2}} \\ &\leq e^{-2n(1-\rho_1)^2 \left(\frac{\sum_{i=1}^n \beta_i(j, k)}{n} - 2\eta \right)^2} \\ &\leq e^{-2n(1-\rho_1)^2(\beta-2\eta)^2} \end{aligned}$$

Thus,

$$\begin{aligned}\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) &\leq J^2 \max_{j \neq k} \mathbb{P}_\theta(N_j + N_k < 2n\eta) \\ &\leq J^2 e^{-2n(1-\rho_1)^2(\beta-2\eta)^2}.\end{aligned}$$

□

6.3. Proof of Proposition 1. The result is straightforward when $n = 1$. We assume in what follows $n \geq 2$. We prove the proposition by showing that when the probability of having small clusters is high, the two risks are not necessarily equivalent; and $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)$ may be much smaller than $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$.

Consider $J = 3$ (similar examples can be constructed for any $J \geq 3$) with $\nu_1 = 1 - 2\eta$, $\nu_2 = \nu_3 = \eta$. We take $F_1 = U(0, 1/2)$, $F_2 = U(3/4, 1)$ and $F_3 = U(3/4 - \varepsilon, 1 - \varepsilon)$ for some $0 < \varepsilon < 1/4$. In this case,

$$\begin{aligned}\mathbb{P}_\theta(X_i \in \cdot \mid Y_i \in (0, 1/2)) &= \delta_1(\cdot) \\ \mathbb{P}_\theta(X_i \in \cdot \mid Y_i \in (3/4 - \varepsilon, 3/4)) &= \delta_3(\cdot) \\ \mathbb{P}_\theta(X_i \in \cdot \mid Y_i \in (1 - \varepsilon, 1)) &= \delta_2(\cdot) \\ \mathbb{P}_\theta(X_i \in \cdot \mid Y_i \in (3/4, 1 - \varepsilon)) &= \frac{1}{2}\delta_2(\cdot) + \frac{1}{2}\delta_3(\cdot).\end{aligned}$$

So in this case

$$\begin{aligned}\mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n}) &= \mathbb{E}_\theta\left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \mid Y_{1:n}\right) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{P}_\theta(h_i(Y_{1:n}) \neq \tau(X_i) \mid Y_{1:n}).\end{aligned}$$

A Bayes “classifier” minimizing $h \mapsto \mathcal{R}_n^{\text{class}}(\theta, h)$ in this case is given by

$$h_i^* = \begin{cases} 1 & \text{if } Y_i \in (0, 1/2) \\ 2 & \text{if } Y_i \in (1 - \varepsilon, \varepsilon) \\ 3 & \text{if } Y_i \in (3/4 - \varepsilon, 1 - \varepsilon). \end{cases}$$

With this choice, the optimal permutation is identity and

$$\begin{aligned}\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h^*) \mid Y_{1:n}) &= \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (0, 1/2)} \mathbb{P}_\theta(1 \neq X_i \mid Y_i) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4)} \mathbb{P}_\theta(3 \neq X_i \mid Y_i) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (1 - \varepsilon, 1)} \mathbb{P}_\theta(2 \neq X_i \mid Y_i) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)} \mathbb{P}_\theta(3 \neq X_i \mid Y_i),\end{aligned}$$

ie.,

$$\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h^*) \mid Y_{1:n}) = \frac{1}{2n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)}.$$

Thus,

$$\begin{aligned}\mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h^*) \mid Y_{1:n}) \right] &= \frac{1}{2} \mathbb{P}_\theta(Y_1 \in (3/4, 1 - \varepsilon)) \\ &= \frac{1}{2} (\eta \cdot (1 - 4\varepsilon) + \eta \cdot (1 - 4\varepsilon)) \\ &= \eta(1 - 4\varepsilon).\end{aligned}$$

We now investigate $\mathbb{E}_\theta(\min_{\tau} U_{n,\tau}(h^*))$, for the previous Bayes classifier h^* [which do not necessarily minimize $h \mapsto \mathbb{E}_\theta(\min_{\tau} U_{n,\tau}(h))$]. We rewrite,

$$\begin{aligned}\mathbb{E}_\theta(\min_{\tau} U_{n,\tau}(h^*)) &= \mathbb{E}_\theta(\min_{\tau} U_{n,\tau}(h^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0}) \\ &\quad + \sum_{m=1}^n \mathbb{E}_\theta(\min_{\tau} U_{n,\tau}(h^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = m})\end{aligned}$$

Let first consider $\mathbb{E}_\theta[\min_{\tau} U_{n,\tau}(h^*) \mid Y_{1:n}]$ on the event that $\{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0\}$. Let define $N = \sum_{i=1}^n \mathbf{1}_{X_i \neq 1}$:

- if $N = 0$ this means that all the Y_i are in $(0, 1/2)$ and our classifier h^* combined with the identity permutation will make zero error, *i.e.* $\min_{\tau} U_{n,\tau}(h^*) = 0$ on this event;
- if $N = 1$ this means that there is only one $j \in \{1, \dots, n\}$ such that $Y_j \in (3/4, 1 - \varepsilon)$ [by assumption it cannot be in $(3/4 - \varepsilon)$ or $(1 - \varepsilon, 1)$]. Our classifier will predict $h_i^* = 1$ for all $i \neq j$ and $h_j = 3$. Now, necessarily $X_i = 1$ for $i \neq j$. If $X_j = 3$ then $h^* \circ \text{Id}$ will have loss zero, and if $X_j = 2$ then $h^* \circ \begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix}$ will have loss zero. So in the event that $\{N = 1\}$ we also have that $\min_{\tau} U_{n,\tau}(h^*) = 0$.
- If $N \geq 2$ our classifier will still classify perfectly all the $Y_i \in (0, 1/2)$ so the loss cannot exceed $\min_{\tau} U_{n,\tau}(h^*) \leq N/n$ in this case.

So on the event $\{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0\}$:

$$\begin{aligned}\mathbb{E}_\theta(\min_{\tau} U_{n,\tau}(h^*) \mid Y_{1:n}) &\leq \mathbb{E}_\theta\left(\frac{N}{n} \mathbf{1}_{N \geq 2} \mid Y_{1:n}\right) \\ &\leq \sum_{k=2}^n \frac{k}{n} \mathbb{P}_\theta(N = k \mid Y_{1:n})\end{aligned}$$

But under the law of $X_{1:n} \mid Y_{1:n}$ in the considered event, we have that N is almost-surely equal to the number of $Y_i \in (3/4, 1 - \varepsilon)$, so

$$\begin{aligned}\mathbb{E}_\theta(\min_{\tau} U_{n,\tau}(h^*) \mid Y_{1:n}) &\leq \sum_{k=2}^n \frac{k}{n} \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1-\varepsilon)} = k} \\ &= \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1-\varepsilon)}\right) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1-\varepsilon)} \geq 2}.\end{aligned}$$

Deduce that,

$$\begin{aligned}\mathbb{E}_\theta(\min_{\tau} U_{n,\tau}(h^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0}) \\ \leq \mathbb{E}_\theta\left(\left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1-\varepsilon)}\right) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1-\varepsilon)} \geq 2} \mid \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0\right)\end{aligned}$$

Conditional on $\{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0\}$, the random variable $\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1-\varepsilon)}$ has a Binomial($n, 2\eta$) distribution. Then,

$$\begin{aligned} \mathbb{E}_\theta(\min_\tau U_{n,\tau}(h^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0}) &\leq \frac{1}{n} \sum_{k=2}^n \binom{n}{k} k \cdot (2\eta)^k (1-2\eta)^{n-k} \\ &= \frac{2\eta(1 - (1-2\eta)^n - 2\eta)}{1-2\eta} \\ &\asymp 4(n-1)\eta^2 \end{aligned}$$

when $\eta \ll 1/n$.

Next [remark that this can be largely improved, but this is indeed for our purpose],

$$\begin{aligned} \sum_{m=1}^n \mathbb{E}_\theta(\min_\tau U_{n,\tau}(h^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = m}) \\ \leq \mathbb{P}_\theta\left(\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} \geq 1\right) \\ = 1 - \mathbb{P}_\theta\left(\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0\right) \\ = 1 - \mathbb{P}_\theta\left(\forall i, Y_i \notin (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)\right) \\ = 1 - \left((1-2\eta) + 2\eta(1-4\varepsilon)\right)^n \\ = 1 - (1-8\eta\varepsilon)^n \\ \asymp 8n\eta\varepsilon \end{aligned}$$

when $\eta \ll n$. So by choosing $\varepsilon \asymp \eta$, we have shown that whenever $\eta \ll 1/n$

$$\inf_h \mathbb{E}_\theta(\min_\tau U_{n,\tau}(h)) \leq \mathbb{E}_\theta(\min_\tau U_{n,\tau}(h^*)) \lesssim n\eta^2$$

but

$$\inf_h \mathbb{E}_\theta[\min_\tau \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n})] = \mathbb{E}_\theta[\min_\tau \mathbb{E}_\theta(U_{n,\tau}(h^*) \mid Y_{1:n})] \sim \eta$$

so that

$$\frac{\inf_h \mathbb{E}_\theta(\min_\tau U_{n,\tau}(h))}{\inf_h \mathbb{E}_\theta[\min_\tau \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n})]} \lesssim n\eta$$

which goes to zero as $\eta \rightarrow 0$.

6.4. Proof of Theorem 3.

6.4.1. *Bounds on the online Bayes risk.* Simple computations lead to the expression of the online Bayes risk:

$$\inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class, on}}(\theta, \mathfrak{h}) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \mathbb{P}_\theta(X_i \neq x_0 \mid Y_{1:i}) \right] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \phi_{\theta,i}(x) \right) \right].$$

First, for $i = 1$,

$$\begin{aligned}\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \mathbb{P}_\theta (X_1 \neq x_0 \mid Y_1) \right] &= \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \frac{\sum_{x \neq x_0} \nu_x f_x(Y_1)}{\sum_x \nu_x f_x(Y_1)} \right] \\ &= \int \min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \nu_x f_x(y) d\mathcal{L}(y),\end{aligned}$$

so that using Assumption 1,

$$\begin{aligned}(14) \quad \delta \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y) &\leq \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \mathbb{P}_\theta (X_1 \neq x_0 \mid Y_1) \right] \\ &\leq (1 - (J - 1)\delta) \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y).\end{aligned}$$

Then, using the Forward recursions (see Equation (3.22) in Proposition 3.2.5 of [7]), for any $i \geq 2$,

$$\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \phi_{\theta,i}(x) \right) \right] = \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\frac{\sum_{x \neq x_0} \sum_{x' \in \mathbb{X}} Q_{x',x} \phi_{\theta,i-1}(x') f_x(Y_i)}{\sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q_{x',x} \phi_{\theta,i-1}(x') f_x(Y_i)} \right) \right].$$

Let $A \in \mathcal{Y}$. One has:

$$\begin{aligned}\mathbb{P}_\theta (Y_i \in A \mid Y_{1:i-1}) &= \sum_{x \in \mathbb{X}} \mathbb{P}_\theta (Y_i \in A, X_i = x \mid Y_{1:i-1}) \\ &= \sum_{x \in \mathbb{X}} \mathbb{P}_\theta (Y_i \in A \mid X_i = x, Y_{1:i-1}) \mathbb{P}_\theta (X_i = x \mid Y_{1:i-1}) \\ &= \sum_{x \in \mathbb{X}} \mathbb{P}_\theta (X_i = x \mid Y_{1:i-1}) \int_A f_x(y) d\mathcal{L}(y) \\ &= \sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} \mathbb{P}_\theta (X_{i-1} = x', X_i = x \mid Y_{1:i-1}) \int_A f_x(y) d\mathcal{L}(y) \\ &= \sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} \mathbb{P}_\theta (X_i = x \mid X_{i-1} = x', Y_{1:i-1}) \phi_{\theta,i-1}(x') \int_A f_x(y) d\mathcal{L}(y) \\ &= \int_A \left(\sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q_{x',x} \phi_{\theta,i-1}(x') f_x(y) \right) d\mathcal{L}(y),\end{aligned}$$

so that, conditionally on $Y_{1:i-1}$, Y_i has density $\sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q_{x',x} \phi_{\theta,i-1}(x') f_x$ with respect to the dominating measure \mathcal{L} . We thus get :

$$\begin{aligned}\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \phi_{\theta,i}(x) \right) \right] &= \mathbb{E}_\theta \left[\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\frac{\sum_{x \neq x_0} \sum_{x' \in \mathbb{X}} Q_{x',x} \phi_{\theta,i-1}(x') f_x(Y_i)}{\sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q_{x',x} \phi_{\theta,i-1}(x') f_x(Y_i)} \right) \mid Y_{1:i-1} \right] \right] \\ &= \mathbb{E}_\theta \left[\int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \sum_{x' \in \mathbb{X}} Q_{x',x} \phi_{\theta,i-1}(x') f_x(y) \right) d\mathcal{L}(y) \right].\end{aligned}$$

Then, under Assumption 1, for any $i \geq 2$,

$$(15) \quad \delta \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y) \leq \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \phi_{\theta,i}(x) \right) \right] \\ \leq (1 - (J-1)\delta) \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y),$$

so that using (14) and (15) we get for all $n \geq 1$,

$$\delta \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y) \leq \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathfrak{h}) \leq (1 - (J-1)\delta) \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y).$$

6.4.2. *Bounds on the offline Bayes risk.* The offline Bayes risk can be written as:

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \mathbb{P}_\theta(X_i \neq x_0 \mid Y_{1:n}) \right] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta,i|1:n}(x) \right].$$

Now, for any $i \leq n-1$, using the Backward recursions (see Proposition 3.3.9 in [7]),

$$\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta,i|1:n}(x) \right] = \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \sum_{x' \in \mathbb{X}} \phi_{\theta,i+1|1:n}(x') B_{\theta,i}[Y_{1:i}](x', x) \right].$$

Now, using Assumption 1 and Equation (10), we get

$$B_{\theta,i}[Y_{1:i}](x', x) = \frac{\phi_{\theta,i}(x) Q(x, x')}{\sum_{\tilde{x} \in \mathbb{X}} \phi_{\theta,i}(\tilde{x}) Q(\tilde{x}, x')} \geq \frac{\phi_{\theta,i}(x) Q(x, x')}{1 - (J-1)\delta}$$

so that

$$\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta,i|1:n}(x) \right] \geq \frac{\delta}{1 - (J-1)\delta} \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x'} \phi_{\theta,i+1}(x') \sum_{x \neq x_0} \phi_{\theta,i}(x) \right] \\ \geq \frac{\delta}{1 - (J-1)\delta} \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta,i}(x) \right] \\ = \frac{\delta}{1 - (J-1)\delta} \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta,i}(x) \right].$$

Now, for $i = n$, the same inequality obviously holds, so that we get

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \geq \frac{\delta}{1 - (J-1)\delta} \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathfrak{h})$$

6.5. Proof of Theorem 4.

Case of offline clustering. Recall that $\phi_{\theta,i|n}(\cdot) = \mathbb{P}_\theta(X_i \in \cdot \mid Y_{1:n})$. Then,

$$\begin{aligned}
\mathcal{R}_n^{\text{class}}(\theta, \hat{h}) - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) &= \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i \neq \hat{h}_i} \right] - \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i \neq h_i^*} \right] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\mathbb{P}_\theta(X_i = h_i^* \mid Y_{1:n}) - \mathbb{P}_\theta(X_i = \hat{h}_i \mid Y_{1:n}) \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\mathbb{P}_\theta(X_i = h_i^* \mid Y_{1:n}) - \mathbb{P}_{\hat{\theta}}(X_i = \hat{h}_i \mid Y_{1:n}) + \|\phi_{\hat{\theta},i|n} - \phi_{\theta,i|n}\|_{\text{TV}} \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\mathbb{P}_\theta(X_i = h_i^* \mid Y_{1:n}) - \mathbb{P}_{\hat{\theta}}(X_i = h_i^* \mid Y_{1:n}) + \|\phi_{\hat{\theta},i|n} - \phi_{\theta,i|n}\|_{\text{TV}} \right] \\
&\leq \frac{2}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\|\phi_{\hat{\theta},i|n} - \phi_{\theta,i|n}\|_{\text{TV}} \right]
\end{aligned}$$

where the penultimate line follows because by definition \hat{h}_i maximizes $x \mapsto \mathbb{P}_{\hat{\theta}}(X_i = x \mid Y_{1:n})$.

Then, by an application of Proposition 2.2 of [12], one has

$$\begin{aligned}
(16) \quad \|\phi_{\theta,i|n}(\cdot, Y_{1:n}) - \phi_{\hat{\theta},i|n}(\cdot, Y_{1:n})\|_{\text{TV}} &\leq \frac{4(1-\delta)}{\delta^2} \left(\rho^{i-1} \|\nu - \hat{\nu}\|_2 \right. \\
&\quad \left. + (1/(1-\rho) + 1/(1-\hat{\rho})) \|Q - \hat{Q}\|_{\text{F}} + \sum_{l=1}^n \delta \frac{(\hat{\rho} \vee \rho)^{|l-i|}}{c^*(Y_l)} \max_{x \in \mathbb{X}} |f_x(Y_l) - \hat{f}_x(Y_l)| \right)
\end{aligned}$$

where $c^*(y) = \min_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q(x, x') f_{x'}(y)$, $\rho = 1 - \delta/(1-\delta)$ and $\hat{\rho} = 1 - \hat{\delta}/(1-\hat{\delta})$ where $\hat{\delta}$ is an estimator of δ .

Thus under Assumption 3,

$$\begin{aligned}
\mathcal{R}_n^{\text{class}}(\theta, \hat{h}) - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) &\leq \frac{8(1-\delta)}{\delta^2} \mathbb{E}_\theta \left[\frac{1}{n\delta} \|\nu - \hat{\nu}\|_2 + \left(\frac{1}{\delta} + \frac{1}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} + \frac{1}{n} \sum_{i=1}^n \sum_{l=1}^n \delta \frac{(\hat{\rho} \vee \rho)^{|l-i|}}{c^*(Y_l)} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty \right] \\
&\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} + \frac{\delta^2}{n} \sum_{i=1}^n \sum_{l=1}^n \frac{(\hat{\rho} \vee \rho)^{|l-i|}}{c^*(Y_l)} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty \right] \\
&\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} \right] + \frac{\delta^2 C}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[\frac{(\hat{\rho} \vee \rho)^{|l-i|}}{c^*(Y_l)} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty \right] \\
&\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} \right] + \frac{\delta^2 C}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[\frac{1}{c^*(Y_l)^2} \right]^{1/2} \mathbb{E}_\theta \left[(\hat{\rho} \vee \rho)^{2|l-i|} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2} \\
&\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} \right] + \frac{\delta C \sqrt{C^*}}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[(\hat{\rho} \vee \rho)^{2|l-i|} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2}
\end{aligned}$$

Case of online clustering. The proof follows the same lines but uses Proposition 2.1 of [12] instead of Proposition 2.2. First, computations similar to those of the offline case yield:

$$\mathcal{R}_n^{\text{class,on}}(\theta, \hat{\mathbf{h}}) - \inf_{\mathbf{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathbf{h}) \leq \frac{2}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\|\phi_{\hat{\theta},i} - \phi_{\theta,i}\|_{\text{TV}} \right].$$

By application of Proposition 2.1 of [12] and Assumption 3 one has:

$$\begin{aligned} \|\phi_{\theta,i}(\cdot, Y_{1:i}) - \phi_{\hat{\theta},i}(\cdot, Y_{1:i})\|_{\text{TV}} &\leq \frac{4(1-\delta)}{\delta^2} \left(\rho^{i-1} \|\nu - \hat{\nu}^{(i)}\|_2 + \frac{\|Q - \hat{Q}^{(i)}\|_{\text{F}}}{1-\rho} \right. \\ &\quad \left. + \delta \sum_{l=1}^i \frac{\rho^{i-l}}{c^\star(Y_l)} \max_{x \in \mathbb{X}} |f_x(Y_l) - \hat{f}_x^{(i)}(Y_l)| \right) \end{aligned}$$

where $(\hat{\nu}^{(i)}, \hat{Q}^{(i)}, \hat{f}_x^{(i)})$ is the estimator of θ using observations $Y_{1:i}$. It follows that:

$$\begin{aligned} \mathcal{R}_n^{\text{class,on}}(\theta, \hat{\mathbf{h}}) - \inf_{\mathbf{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathbf{h}) &\leq \frac{\delta C}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\rho^{i-1} \|\nu - \hat{\nu}^{(i)}\|_2 + \frac{\|Q - \hat{Q}^{(i)}\|_{\text{F}}}{1-\rho} \right. \\ &\quad \left. + \delta \sum_{l=1}^i \frac{\rho^{i-l}}{c^\star(Y_l)} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x^{(i)}\|_\infty \right]. \end{aligned}$$

Given that

$$\begin{aligned} \mathbb{E}_\theta \left[\frac{\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x^{(i)}\|_\infty}{c^\star(Y_l)} \right] &\leq \mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x^{(i)}\|_\infty^2 \right]^{1/2} \mathbb{E}_\theta \left[\frac{1}{c^\star(Y_l)^2} \right]^{1/2} \\ &\leq \frac{\sqrt{C^\star}}{\delta} \mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x^{(i)}\|_\infty^2 \right]^{1/2}, \end{aligned}$$

one gets:

$$\begin{aligned} &\mathcal{R}_n^{\text{class,on}}(\theta, \hat{\mathbf{h}}) - \inf_{\mathbf{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class,on}}(\theta, \mathbf{h}) \\ &\leq \frac{\delta C}{n} \sum_{i=1}^n \left(\mathbb{E}_\theta \left[\rho^{i-1} \|\nu - \hat{\nu}^{(i)}\|_2 + \frac{\|Q - \hat{Q}^{(i)}\|_{\text{F}}}{\delta} \right] + \frac{\sqrt{C^\star}}{\delta} \mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x^{(i)}\|_\infty^2 \right]^{1/2} \right). \end{aligned}$$

6.6. Proof of Corollary 1. The Algorithm 1 will yield the estimates in the statement of the theorem. This algorithm merges the spectral algorithms of [12, 1] with some slight modifications. Note that all the expectations and probabilities of this proof are with respect to the observations and the random unit matrices. Also note that the algorithm outputs estimates of the densities that are not necessarily *bona-fide* densities. This is not problematic for the plugin procedure as one typically uses the Forward-Backward algorithm [7] which works even if the emissions are not correctly normalized.

First, we start by controlling $\mathbb{E}_\theta \left[\|Q - \hat{Q}\|_{\text{F}}^2 \right]$, $\mathbb{E}_\theta \left[\frac{1}{\delta^2} \right]$ and $\mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]$ using the estimates yielded by the algorithm. Thanks to step 7 of the algorithm, one can obtain a slightly different version of Theorem 3.1 of [12] (Note that this version is used in the proof of Corollary 3.2 in [12]). It ensures the existence of positive constants C, x_0, y_0, D_0 and n_1 such that for all $D \geq D_0$ there exist a permutation $\tau_D \in \mathcal{S}_J$ such that for all

Algorithm 1 Non-parametric spectral estimation of the transition matrix and the emission laws**Input**

- Number of states J , integers D and r .
- Data $(Y_i)_{i \leq n+2}$ drawn from a HMM with J states.
- Functions $(\varphi_d)_{d \in \mathbb{N}}$ **uniformly bounded** such that $O = (\mathbb{E}_\theta[\varphi_d(Y_1) \mid X_1 = j])_{1 \leq d \leq D, 1 \leq j \leq J}$ is of rank J with $\sigma_J(O)$ **bounded away from 0 uniformly** in D , at least for D large enough.
- K a Lipschitz-continuous kernel

Output

- Spectral estimators \hat{Q} and $(\hat{f}_j)_{1 \leq j \leq J}$

Estimation

[Step 1] For all $a, b, c \in \llbracket 1, D \rrbracket$, consider the following empirical estimators:

$$\begin{aligned}\hat{L}(a) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) \\ \hat{N}(a, b) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) \varphi_b(Y_{s+1}) \\ \hat{P}(a, c) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) \varphi_c(Y_{s+2}) \\ \hat{M}(a, b, c) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) \varphi_b(Y_{s+1}) \varphi_c(Y_{s+2}) \\ \hat{M}^{x, L}(a, b) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) K_L(x, Y_{s+1}) \varphi_b(Y_{s+2})\end{aligned}$$

where L is chosen such that: $2^L \asymp (n/\log(n))^{1/(2s+1)}$ and s is the smoothness of the emissions.

[Step 2] Let \hat{V} be the $D \times J$ matrix of orthonormal right singular vectors of \hat{P} corresponding to its top J singular values.

[Step 3] For all $d \in \llbracket 1, D \rrbracket$, set $\hat{B}(d) = (\hat{V}^\top \hat{P} \hat{V})^{-1} \hat{V}^\top \hat{M}(\cdot, d, \cdot) \hat{V}$

[Step 4] Generate Ω a $J \times J$ unit matrix uniformly drawn, for all $x \in \llbracket 1, J \rrbracket$, $\hat{C}(x) = \sum_{d=1}^D (\hat{V} \Omega)(d, x) \hat{B}(d)$

[Step 5] Compute \hat{R}_1 a $J \times J$ unit Euclidean norm columns matrix that diagonalizes the matrix $\hat{C}(1)$:

$$\hat{R}_1^{-1} \hat{C}(1) \hat{R}_1 = \text{Diag}[(\hat{\Lambda}(1, 1), \dots, \hat{\Lambda}(1, J))]$$

[Step 6] For all $x, x' \in \llbracket 1, J \rrbracket$, $\hat{\Lambda}(x, x') = (\hat{R}_1^{-1} \hat{C}(x) \hat{R}_1)_{x', x'}$.

[Step 7] Repeat steps 4 to 6 r times and take Ω_r maximizing $i \mapsto \min_{k \leq J} \min_{k_1 \neq k_2} |\hat{\Lambda}_i(k, k_1) - \hat{\Lambda}_i(k, k_2)|$

[Step 8] Set $\hat{O} = \hat{V} \Omega_r \hat{\Lambda}$, $\hat{\nu} = (\hat{V}^\top \hat{O})^{-1} \hat{V}^\top \hat{L}$ and $\hat{Q} = \Pi_{TM} \left((\hat{V}^\top \hat{O} \text{Diag}[\hat{\nu}])^{-1} \hat{V}^\top \hat{N} \hat{V} (\hat{O}^\top \hat{V})^{-1} \right)$

where Π_{TM} denotes the projection (with respect to the scalar product given by the Frobenius norm) onto the convex set of transition matrices.

[Step 9] For $x \in \mathbb{R}$, set $\hat{B}^x = \hat{B}^{x, D, L} = (\hat{V}^\top \hat{P} \hat{V})^{-1} \hat{V}^\top \hat{M}^{x, L} \hat{V}$

[Step 10] Set $\hat{R}_2 = \hat{Q} \hat{O}^\top \hat{V}$ and take $\tilde{f}_j(x) = (\hat{R}_2 \hat{B}^x \hat{R}_2^{-1})_{j, j}$

[Step 11] $\hat{f}_j(x) = \begin{cases} \tilde{f}_j(x) & \text{si } |\tilde{f}_j(x)| \leq n^\beta \\ n^\beta \text{sign}(\tilde{f}_j(x)) & \text{otherwise} \end{cases}$ for $\beta > 0$ fixed (but arbitrary).

$n \geq n_1 \eta_3^2(\varphi_D) x(y + \log(r)) e^{y/r}$, $x \geq x_0$, $y \geq y_0$ and $r \geq 1$, with probability at least $1 - 4e^{-x} - 2e^{-y}$:

$$(17) \quad \begin{aligned} \|Q^{\tau_D} - \hat{Q}\|_F^2 &\leq C \eta_3^2(\varphi_D) x(y + \log(r)) e^{y/r} / n \\ \max_{x \in \mathbb{X}} \|f_{D, \tau_D}(x) - \hat{f}_{D, x}^{(r)}\|_2^2 &\leq C \eta_3^2(\varphi_D) x(y + \log(r)) e^{y/r} / n \end{aligned}$$

where:

- $(\varphi_k)_{k \in \mathbb{N}}$ is the basis used in Algorithm 1
- $\eta_3^2(\varphi_D) = \sup_{y, y' \in \mathcal{Y}^3} \sum_{a, b, c=1}^D (\varphi_a(y_1)\varphi_b(y_2)\varphi_c(y_3) - \varphi_a(y'_1)\varphi_b(y'_2)\varphi_c(y'_3))^2$
- $f_{D,x} = \sum_{d=1}^D \langle f_x, \varphi_d \rangle \varphi_d$ the projection of the density f_x on the subspace spanned by the first D components of the basis.
- $\hat{f}_{D,x}^{(r)} = \sum_{d=1}^D \hat{O}_{d,x} \varphi_d$ where \hat{O} is the matrix constructed at step 8 of Algorithm 1.

As detailed in [12], when using a wavelet basis or trigonometric polynomials basis, $\eta_3(\varphi_D)$ ensures for a constant $C_\eta > 0$:

$$(18) \quad \eta_3(\varphi_D) \leq C_\eta D^{3/2} \text{ and } \max_{l \in \mathbb{N}} \|\varphi_l\|_\infty < \infty$$

We assume a similar basis is used.

It is important to note that the estimator $\hat{f}_{D,x}^{(r)}$ is not the one yielded by the algorithm 1 but it is rather the one used in [12]. We do not use it for the estimation because it does not allow obtaining the appropriate rate in infinite norm (see [12] for more details). However, we will use it in our proof because $\|\hat{O}(\cdot, k) - O(\cdot, \tau_{D_n}(k))\|_2 = \|\hat{f}_{D_n,k}^{(r)} - f_{D_n, \tau_{D_n}(k)}\|_2$.

Assume that the parameters $x = x_n$, $y = y_n$, $D = D_n$ and $r = r_n$ are increasing with respect to n and that $n \geq n_1 C_\eta^2 D_n^3 x_n (y_n + \log(r_n)) e^{y_n/r_n}$.

Control of $\mathbb{E}_\theta [\|Q^{\tau_{D_n}} - \hat{Q}\|_F^2]$. The control of $\|Q^{\tau_{D_n}} - \hat{Q}\|_F^2$ in expectation is already proved in Corollary 3.2 of [12] using Inequality (17). The proof chooses $r_n \propto \log(n)$ and $\eta_3(\varphi_{D_n}) = o(\sqrt{n}/\log(n))$ and yields:

$$\mathbb{E}_\theta [\|Q^{\tau_{D_n}} - \hat{Q}\|_F^2] = \mathcal{O}(\eta_3^2(\varphi_{D_n}) \log(n)/n) = \mathcal{O}(D_n^3 \log(n)/n) \quad (\text{by (18)})$$

for a sequence $(\tau_{D_n})_n$ of permutations. We will keep the same values of r_n and D_n in what follows.

One of the advantages of this algorithm with respect to the previous versions is that it allows obtaining the appropriate rate on the errors in the estimation of all the model parameters. This is done thanks to the use of the kernel estimator of the emission densities for which the error of approximation is tuned (through the parameter L) independently of the error of estimation of the transition matrix. The shortcoming of the algorithm proposed in [12] is that it does not allow controlling the rate on the emission densities without altering that of the transition matrix as it is clear in Corollary 3.3 of that paper.

Control of $\mathbb{E}_\theta [\frac{1}{\delta^2}]$. Let $\tilde{\delta} = \frac{\delta}{2}$. Then,

$$\begin{aligned} \mathbb{E}_\theta \left[\frac{1}{\delta^2} \right] &= \mathbb{E}_\theta \left[\frac{1}{\delta^2} \mathbf{1}_{\delta < \tilde{\delta}} \right] + \mathbb{E}_\theta \left[\frac{1}{\delta^2} \mathbf{1}_{\delta \geq \tilde{\delta}} \right] \\ &\leq \mathbb{E}_\theta \left[\frac{1}{\delta^2} \mathbf{1}_{\delta < \tilde{\delta}} \right] + \frac{1}{\tilde{\delta}^2} \end{aligned}$$

On the other hand:

$$\tilde{\delta} - \hat{\delta} = \delta - \hat{\delta} - \frac{\delta}{2} \leq \left| \delta - \hat{\delta} \right| - \frac{\delta}{2} \leq \max_{i,j} |Q_{i,j}^{\tau_{D_n}} - \hat{Q}_{i,j}| - \frac{\delta}{2} \leq \|Q^{\tau_{D_n}} - \hat{Q}\|_F - \frac{\delta}{2}$$

where we have used the inequality $|\min_{i,j} Q_{i,j} - \min_{i,j} \hat{Q}_{i,j}| \leq \max_{i,j} |Q_{i,j} - \hat{Q}_{i,j}|$.

We assume that all the entries of \hat{Q} are between $n^{-\alpha/2}$ and $1 - n^{-\alpha/2}$ for $\alpha \geq 2$. If it is not the case, modifying the entries of \hat{Q} to obtain a similar property induces an error of order

$n^{-\alpha/2}$ which is negligible with respect to the rates we seek and all the subsequent results remain unchanged. It follows that $\hat{\delta} \geq n^{-\alpha/2}$ and:

$$\mathbb{E}_\theta \left[\frac{1}{\hat{\delta}^2} \right] \leq n^\alpha \mathbb{P}_\theta \left(\|Q^{\tau_{D_n}} - \hat{Q}\|_F > \frac{\delta}{2} \right) + \frac{1}{\hat{\delta}^2}$$

Choosing for example $x_n = y_n = r_n = \alpha \log(n)$, then one obtains for n large enough:

$$\begin{aligned} \mathbb{P}_\theta \left(\|Q^{\tau_{D_n}} - \hat{Q}\|_F > \frac{\delta}{2} \right) &\leq \mathbb{P}_\theta \left(\|Q^{\tau_{D_n}} - \hat{Q}\|_F > C\eta_3^2(\varphi_{D_n})x_n(y_n + \log(r_n))e/n \right) \\ &\leq 4e^{-x_n} + 2e^{-y_n} \end{aligned}$$

It follows that $\mathbb{E}_\theta \left[\frac{1}{\hat{\delta}^2} \right]$ is upper-bounded by an absolute constant.

The values x_n , y_n and r_n will be kept the same in what follows.

Control of $\mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_{\tau_{D_n}(x)} - \hat{f}_x\|_\infty^2 \right]$. The difficulty of the control of this quantity lies in ensuring that the same permutation τ_{D_n} used for the control of Q still works for the control of the emission densities. In the spectral algorithm of [1], the matrix \hat{R} is chosen independently of Q or \hat{Q} (In fact, this algorithm does not even estimate \hat{Q}). Had we used this matrix, there would not be any reason for which the same permutation τ_{D_n} works for the control of the emission densities. To solve this problem, we choose a matrix \hat{R} that depends explicitly on \hat{Q} so that the same permutation that works for the control of Q works also for that of the emission densities. We follow here the steps of the proof of Theorem 5 in [1].

Let $M^{x,L}, P, O$ be the quantities estimated by $\hat{M}^{x,L}, \hat{P}, \hat{O}$ and construct \tilde{f}_j, \tilde{f}_j using Algorithm 1. Let E_n be the event with probability greater than $1 - 4e^{-x_n} - 2e^{-y_n}$ on which the control of (17) holds when $x_n = y_n = r_n = \alpha \log(n)$ and $D = D_n$. For $\gamma > 0$ there exists $c = c(\gamma)$ such that the event:

$$\mathcal{A}_n = \left\{ \|\hat{P} - P\| \leq cD_n \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}}, \quad \sup_{x \in \mathbb{R}} \|\hat{M}^{x,L} - M^{x,L}\| \leq cD_n^2 \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right\}$$

is measurable and has probability $n^{-\gamma}$ (Cf. Lemma 25.a of [1]). Given that E_n has probability greater than $1 - 4e^{-x_n} - 2e^{-y_n} = 1 - 6n^{-\alpha}$, it follows that $\mathcal{A}_n \cap E_n$ has probability greater than $1 - n^{-\gamma} - 6n^{-\alpha}$. Note that the difference with the original proof is that we use the event $\mathcal{A}_n \cap \mathcal{E}_n$ instead of the event \mathcal{A}_n . This is compulsory to control the errors of \hat{Q} and \hat{f} simultaneously.

On the event $\mathcal{A}_n \cap E_n$, and at step 10 of Algorithm 1, instead of using the matrix \hat{R} appearing in the spectral algorithm in [1], we rather use the matrix $\hat{R}_2 = \hat{Q}\hat{O}^\top \hat{V}$ where the components \hat{V} , \hat{Q} and \hat{O} are constructed in the Algorithm 1. On the other hand, since the columns of \hat{R} are not normalized, we choose $\tilde{R}_2 = QO^\top \hat{V}$ on the contrary to what is done in the proof of Theorem 5 in [1]. By denoting $Q^{\tau_{D_n}} = P_{\tau_{D_n}} Q P_{\tau_{D_n}}^{-1}$, one obtains:

$$\hat{R}_2 - P_{\tau_{D_n}} \tilde{R}_2 = \hat{Q}\hat{O}^\top \hat{V} - P_{\tau_{D_n}} QO^\top \hat{V} = \left(\hat{Q}(\hat{O} - OP_{\tau_{D_n}}^\top)^\top + (\hat{Q} - Q^{\tau_{D_n}})P_{\tau_{D_n}} O^\top \right) \hat{V}$$

It follows by using operator norm:

$$(19) \quad \|\hat{R}_2 - P_{\tau_{D_n}} \tilde{R}_2\| \leq \left(\|\hat{Q}\| \|\hat{O} - OP_{\tau_{D_n}}^\top\| + \|P_{\tau_{D_n}} O^\top\| \|\hat{Q} - Q^{\tau_{D_n}}\| \right) \|\hat{V}\|$$

First, note that on the event $\mathcal{A}_n \cap E_n$:

$$\begin{aligned}
\|\hat{Q}\| &\leq 1 \\
\|\hat{V}\| &\leq \|\hat{V}\|_F = J^{1/2} \quad (\text{columns are normalized}) \\
\|\hat{Q} - Q^{\tau_{D_n}}\| &\leq \|\hat{Q} - Q^{\tau_{D_n}}\|_F \\
&\leq \left(CC_\eta^2 D_n^3 x_n (y_n + \log(r_n)) e^{y_n/r_n} / n \right)^{1/2} \quad (\text{by (17) and (18)}) \\
\|P_{\tau_{D_n}} O^\top\| &= \|O\| = \sup_{\|v\|=1} \left(\sum_{j=1}^J \left(\sum_{l=1}^{D_n} v_l \langle f_j, \varphi_l \rangle \right)^2 \right)^{1/2} \\
&\leq (JD_n)^{1/2} \max_{1 \leq l \leq D_n} \|\varphi_l\|_\infty \\
\|\hat{O} - OP_{\tau_{D_n}}^\top\| &\leq \|\hat{O} - OP_{\tau_{D_n}}^\top\|_F = \left(\sum_{k=1}^K \|\hat{O}(\cdot, k) - O(\cdot, \tau_{D_n}(k))\|_2^2 \right)^{1/2} \\
&= \left(\sum_{k=1}^K \|\hat{f}_{D_n, k}^{(r)} - f_{D_n, \tau_{D_n}(k)}\|_2^2 \right)^{1/2} \\
&\leq \left(KCC_\eta^2 D_n^3 x_n (y_n + \log(r_n)) e^{y_n/r_n} / n \right)^{1/2} \quad (\text{by (17) and (18)}).
\end{aligned}$$

By keeping the previous choices of x_n, y_n and r_n then by Inequality (19), there exists a constant C' such that:

$$\mathbf{1}_{\mathcal{A}_n \cap E_n} \|\hat{R}_2 - P_{\tau_{D_n}} \tilde{R}_2\| \leq C' D_n^{1/2} D_n^{3/2} \frac{\log(n)}{\sqrt{n}}.$$

By Lemma 25.b of [1], for n large enough, \hat{P} has rank J , $(\hat{V}^T \hat{P} \hat{V})$ and $(\hat{V}^T P \hat{V})$ are invertible and the matrices $(\hat{B}(d))_{1 \leq d \leq D_n}$ appearing in algorithm 1 are then well-defined. By Lemma 11 and 25.b of [1], $\tilde{B}^x = (\hat{V}^T P \hat{V})^{-1} \hat{V}^T M^x \hat{V}$ satisfies:

$$(20) \quad \tilde{B}^x = (QO^T \hat{V})^{-1} D^x (QO^T \hat{V}) = \tilde{R}_2^{-1} D^x \tilde{R}_2$$

where $D^x = (K_L[f_j](x))_{j \leq J}$. Using the fact that $\sigma_J(\hat{V}) = 1$ (the columns of \hat{V} are orthonormal) and $\sigma_J(QO^T) \geq \sigma_J(Q)\sigma_J(O) > 0$ (because Q in full rank and $\sigma_J(O)$ is bounded from below by an absolute constant by assumption of the algorithm), it follows that $\|\tilde{R}_2^{-1}\|^{-1}$ is bounded from below by an absolute constant because:

$$\begin{aligned}
(21) \quad \|\tilde{R}_2^{-1}\|^{-1} &= \frac{1}{\sigma_1(\tilde{R}_2^{-1})} = \sigma_J(\tilde{R}_2) \\
&= \sigma_J(QO^T \hat{V}) \geq \sigma_J(QO^T) \sigma_J(\hat{V}) = \sigma_J(QO^T) \geq \sigma_J(Q) \sigma_J(O) > 0.
\end{aligned}$$

Thus, for n large enough, the assumption of Lemma 13 (stated below) is verified with $A_t = \tilde{B}^t$, $\hat{A}_t = \hat{B}^t$, $R = \tilde{R}_2 = QO^T \hat{V}$. This ensures that:

$$\begin{aligned}
&\mathbf{1}_{\mathcal{A}_n \cap E_n} \max_{x \in \mathbb{X}} \|\tilde{f}_x - K_L[f_{\tau_{D_n}}(x)]\|_\infty \\
&\leq 4\kappa(\tilde{R}_2) \left[\mathbf{1}_{\mathcal{A}_n \cap E_n} \sup_t \|\tilde{B}^t - \hat{B}^t\| + \lambda_{\max} \kappa(\tilde{R}_2) \|\tilde{R}_2^{-1}\| \mathbf{1}_{\mathcal{A}_n \cap E_n} \|\hat{R}_2 - P_{\tau_{D_n}} \tilde{R}_2\| \right]
\end{aligned}$$

where $\lambda_{\max} = \sup_t \max_j |\lambda_{t,j}| = \sup_t \max_j |K_L[f_j]| < \infty$ and $\lambda_{t,j}$ is the j -th eigenvalue of \tilde{B}^t . By Lemma 35.c of [1], one has for some constant C :

$$\mathbf{1}_{\mathcal{A}_n \cap E_n} \sup_t \|\tilde{B}^t - \hat{B}^t\| \leq C D_n^2 \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}}.$$

By Lemma 35.b in [1], one obtains: $\kappa(\tilde{R}_2) \leq \tilde{C} D_n^{1/2}$ and $\|\tilde{R}_2^{-1}\|$ upper-bounded by an absolute constant by (21). We finally obtain that:

$$\begin{aligned} \mathbf{1}_{\mathcal{A}_n \cap E_n} \max_{x \in \mathbb{X}} \|\tilde{f}_x - K_L[f_{\tau_{D_n}(x)}]\|_\infty &\leq c' D_n^{1/2} \left[D_n^2 \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} + D_n^{1/2} D_n^{1/2} D_n^{3/2} \frac{\log(n)}{\sqrt{n}} \right] \\ &\leq c'' D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \end{aligned}$$

for some constants c', c'' . The choice of L (cf. Algorithm 1) allows obtaining then:

$$\mathbf{1}_{\mathcal{A}_n \cap E_n} \max_{x \in \mathbb{X}} \|\tilde{f}_x - f_{\tau_{D_n}(x)}\|_\infty \leq c'' D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}}.$$

Given that for n large enough $\|f_j\|_\infty \leq n^\beta$, it follows that: $\|\hat{f}_j - f_{\tau_{D_n}(j)}\|_\infty \leq \|\tilde{f}_j^L - f_{\tau_{D_n}(j)}\|_\infty$.

Finally, thanks to the truncation of the emission densities, it is possible to obtain the same rate in expectation:

$$\begin{aligned} \mathbb{E}_\theta \left[\|\hat{f}_j - f_{\tau_{D_n}(j)}\|_\infty^2 \right] &\leq c'' D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{2s}{2s+1}} + 2n^\beta \mathbb{P}((\mathcal{A}_n \cap E_n)^c) \\ &\leq c'' D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{2s}{2s+1}} + 2n^\beta (n^{-\gamma} + 6n^{-\alpha}). \end{aligned}$$

By choosing α and γ sufficiently large, one obtains:

$$\mathbb{E}_\theta \left[\|\hat{f}_j - f_{\tau_{D_n}(j)}\|_\infty^2 \right] = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{2s}{2s+1}} \right).$$

Control of $\frac{1}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[(\rho \vee \hat{\rho})^{2|l-i|} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2}$. Let $\tilde{\rho} = \frac{1-2\tilde{\delta}}{1-\delta}$ where $\tilde{\delta} = \frac{\delta}{2}$.

$$\begin{aligned} \mathbb{E}_\theta \left[(\rho \vee \hat{\rho})^{2|l-i|} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2} &\leq \tilde{\rho}^{|l-i|} \mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2} \\ &\quad + \mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \mathbf{1}_{\hat{\rho} \geq \tilde{\rho}} \right]^{1/2}. \end{aligned}$$

The term $\mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \mathbf{1}_{\hat{\rho} \geq \tilde{\rho}} \right]^{1/2}$ can be made of order $n^{-\alpha}$ for arbitrary $\alpha > 0$ by the same large deviation argument used in the control of $\mathbb{E}_\theta \left[\frac{1}{\delta^2} \right]$. Then, summing up over i and n yields:

$$\frac{1}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[(\rho \vee \hat{\rho})^{2|l-i|} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2} = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right)$$

Finally, The control of the offline excess risk follows directly from Theorem 4 and the previous controls. For the online excess risk, remarking that $\mathbb{E}_\theta \left[\|Q - \hat{Q}^{(i)}\|_F \right] \leq \mathbb{E}_\theta \left[\|Q - \hat{Q}^{(i)}\|_F^2 \right]^{1/2}$, we have:

$$\begin{aligned} \mathcal{R}_n^{\text{class, on}}(\theta, \mathfrak{h}^{\tau_{D_n}^{-1}}) - \inf_{\mathfrak{h} \in \mathcal{H}^{\text{on}}} \mathcal{R}_n^{\text{class, on}}(\theta, \mathfrak{h}) &\leq \mathcal{O} \left(\frac{D_n^{5/2}}{n} \sum_{i=1}^n \left(\frac{\log(i)}{i} \right)^{\frac{s}{2s+1}} \right) \\ &= \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right). \end{aligned}$$

LEMMA 13. Suppose $(A_t, t \in \mathbb{R})$ are $J \times J$ matrices simultaneously diagonalized by a matrix R :

$$RA_tR^{-1} = \text{diag}(\lambda_{t,1}, \dots, \lambda_{t,J}), t \in \mathbb{R}.$$

Let \hat{R} be a matrix such that for some permutation τ of $\{1, \dots, J\}$, we have:

$$\|\hat{R} - P_\tau R\| = \varepsilon_R \leq \frac{\|R^{-1}\|^{-1}}{2}$$

Assume $\lambda_{\max} = \sup_t \max_j |\lambda_{t,j}| < \infty$. For matrices $(\hat{A}_t)_{t \in \mathbb{R}}$, write $\varepsilon_A = \sup_t \|A_t - \hat{A}_t\|$ and define

$$\hat{\lambda}_{t,j} = e_j^T \hat{R} \hat{A}_t \hat{R}^{-1} e_j \quad \text{and} \quad \lambda_{t,\tau(j)} = e_j^T (P_\tau R) A_t (P_\tau R)^{-1} e_j.$$

Then

$$\sup_t \max_j |\hat{\lambda}_{t,j} - \lambda_{t,\tau(j)}| \leq 4\kappa(R) [\varepsilon_A + \lambda_{\max} \kappa(R) \|R^{-1}\| \varepsilon_R].$$

PROOF. Let $\hat{\zeta}_j^T = e_j^T \hat{R}$, let $\hat{\xi}_j = \hat{R}^{-1} e_j$ and define $\zeta_j^T = e_j^T P_\tau R$ and $\xi_j = (P_\tau R)^{-1} e_j$. Then, $\lambda_{t,\tau(j)} = \zeta_j^T A_t \xi_j$, $\hat{\lambda}_{t,j} = \hat{\zeta}_j^T \hat{A}_t \hat{\xi}_j$ and we have:

$$\begin{aligned} |\hat{\lambda}_{t,j} - \lambda_{t,\tau(j)}| &= |\hat{\zeta}_j^T \hat{A}_t \hat{\xi}_j - \zeta_j^T A_t \xi_j| \\ &= |\hat{\zeta}_j^T \hat{A}_t (\hat{\xi}_j - \xi_j) + (\hat{\zeta}_j^T - \zeta_j^T) A_t \xi_j + \hat{\zeta}_j^T (\hat{A}_t - A_t) \xi_j| \\ &\leq \|\hat{\zeta}_j^T\| \|\hat{A}_t\| \|\hat{\xi}_j - \xi_j\| + \|\hat{\zeta}_j^T - \zeta_j^T\| \|A_t \xi_j\| + \|\hat{\zeta}_j^T\| \|\xi_j\| \varepsilon_A \\ \|\zeta_j^T\| &= \|e_j^T P_\tau R\| \leq \|P_\tau R\| = \|R\| \\ \|\hat{\zeta}_j^T - \zeta_j^T\| &= \|e_j^T (\hat{R} - P_\tau R)\| \leq \|\hat{R} - P_\tau R\| = \varepsilon_R \\ \|\xi_j\| &= \|(P_\tau R)^{-1} e_j\| \leq \|(P_\tau R)^{-1}\| = \|R^{-1}\| \\ \|\hat{\xi}_j - \xi_j\| &= \|(\hat{R}^{-1} - (P_\tau R)^{-1}) e_j\| \leq \|\hat{R}^{-1} - (P_\tau R)^{-1}\| \\ &\leq \frac{\|R^{-1}\|^2 \varepsilon_R}{1 - \varepsilon_R \|R^{-1}\|} \quad (\text{by Lemma 37 of [1]}) \\ \|A_t\| &= \|R^{-1} \text{diag}(\lambda_{t,\cdot}) R\| \leq \lambda_{\max} \kappa(R) \\ \|A_t \xi_j\| &= \|\lambda_{t,\tau(j)} \xi_j\| \leq \lambda_{\max} \|R^{-1}\| \\ \|\hat{\zeta}_j^T\| &= \|\hat{\zeta}_j^T - \zeta_j^T + \zeta_j^T\| \leq \|\hat{\zeta}_j^T - \zeta_j^T\| + \|\zeta_j^T\| \leq \varepsilon_R + \|R\| \\ \|\hat{A}_t\| &\leq \varepsilon_A + \|A_t\| \leq \varepsilon_A + \lambda_{\max} \kappa(R). \end{aligned}$$

Thus, by the inequalities above:

$$\begin{aligned}
|\hat{\lambda}_{t,j} - \lambda_{t,\tau(j)}| &\leq (\varepsilon_R + \|R\|)(\varepsilon_A + \lambda_{\max}\kappa(R)) \frac{\|R^{-1}\|^2 \varepsilon_R}{1 - \varepsilon_R \|R^{-1}\|} \\
&\quad + \lambda_{\max} \varepsilon_R \|R^{-1}\| + \|R^{-1}\|(\varepsilon_R + \|R\|)\varepsilon_A \\
&\leq 2(\varepsilon_R + \|R\|)(\varepsilon_A + \lambda_{\max}\kappa(R))\|R^{-1}\|^2 \varepsilon_R \\
&\quad + \lambda_{\max}\|R^{-1}\|\varepsilon_R + \|R^{-1}\|(\varepsilon_R + \|R\|)\varepsilon_A \\
&\leq 2\|R^{-1}\|^2 \varepsilon_R^2 \varepsilon_A + 2\kappa(R)\|R^{-1}\|\varepsilon_A \varepsilon_R + 2\lambda_{\max}\kappa(R)\|R^{-1}\|^2 \varepsilon_R^2 \\
&\quad + 2\lambda_{\max}\kappa(R)^2\|R^{-1}\|\varepsilon_R + \lambda_{\max}\|R^{-1}\|\varepsilon_R + \kappa(R)\varepsilon_A + \|R^{-1}\|\varepsilon_R \varepsilon_A \\
&\leq 2\|R^{-1}\|^2 \varepsilon_R^2 \varepsilon_A + (2\kappa(R) + 1)\|R^{-1}\|\varepsilon_A \varepsilon_R \\
&\quad + 2\lambda_{\max}\kappa(R)((\|R^{-1}\|\varepsilon_R)^2 + \kappa(R)(\|R^{-1}\|\varepsilon_R)) + \lambda_{\max}\|R^{-1}\|\varepsilon_R + \kappa(R)\varepsilon_A \\
&\leq 2\left(\frac{1}{2}\right)^2 \varepsilon_A + \frac{3}{2}\kappa(R)\varepsilon_A + 2\lambda_{\max}\kappa(R)\|R^{-1}\|\varepsilon_R \left(\kappa(R) + \frac{1}{2}\right) \\
&\quad + \lambda_{\max}\|R^{-1}\|\varepsilon_R + \kappa(R)\varepsilon_A \\
&\leq 3\kappa(R)\varepsilon_A + \lambda_{\max}(1 + 3\kappa(R)^2)\|R^{-1}\|\varepsilon_R \\
&\leq 4\kappa(R) [\varepsilon_A + \lambda_{\max}\kappa(R)\|R^{-1}\|\varepsilon_R]
\end{aligned}$$

where we have used $\|R^{-1}\|\varepsilon_R \leq 1/2$ and $\kappa(R) \geq 1$. \square

6.7. Proof of Lemma 1. Let define $p_{n,i}(\theta) := \max_k \mathbb{P}_\theta(X_i = k \mid Y_{1:n}) = \mathbb{P}_\theta(X_i = h_{\theta,i}^* \mid Y_{1:n})$; defining $h_{\theta,i}^*$ realizing the maximum. Suppose $p_{n,i}(\theta) \geq \frac{1}{2} + \gamma$ for some $0 < \gamma \leq 1/2$. Then,

$$\begin{aligned}
p_{n,i}(\theta) - \max_{k \neq h_{\theta,i}^*} \mathbb{P}_\theta(X_i = k \mid Y_{1:n}) &\geq p_{n,i}(\theta) - \sum_{k \neq h_{\theta,i}^*} \mathbb{P}_\theta(X_i = k \mid Y_{1:n}) \\
&= p_{n,i}(\theta) - [1 - p_{n,i}(\theta)] \\
&= 2p_{n,i}(\theta) - 1 \\
&\geq 2\gamma.
\end{aligned}$$

Consequently if $p_{n,i}(\theta) \geq \frac{1}{2} + \gamma$,

$$\begin{aligned}
\mathbb{P}_{\hat{\theta}}(X_i = h_{\theta,i}^* \mid Y_{1:n}) &\geq p_{n,i}(\theta) - \left\| \phi_{\theta,i|n} - \phi_{\hat{\theta},i|n} \right\|_{\text{TV}} \\
&\geq \max_{k \neq h_{\theta,i}^*} \mathbb{P}_\theta(X_i = k \mid Y_{1:n}) + 2\gamma - \left\| \phi_{\theta,i|n} - \phi_{\hat{\theta},i|n} \right\|_{\text{TV}} \\
&\geq \max_{k \neq h_{\theta,i}^*} \mathbb{P}_{\hat{\theta}}(X_i = k \mid Y_{1:n}) + 2\gamma - 2 \left\| \phi_{\theta,i|n} - \phi_{\hat{\theta},i|n} \right\|_{\text{TV}}
\end{aligned}$$

We have shown that on the intersection of the two events

$$E_{n,i} := \left\{ p_{n,i}(\theta) \geq \frac{1}{2} + \gamma \right\}, \quad F_{n,i} := \left\{ \left\| \phi_{\theta,i|n} - \phi_{\hat{\theta},i|n} \right\|_{\text{TV}} < \gamma \right\}$$

the plugin rule $h_{\hat{\theta}}^*$ maximizing $k \mapsto \mathbb{P}_{\hat{\theta}}(X_i = k \mid Y_{1:n})$ is unique and it must be that $h_{\hat{\theta},i}^* = h_{\theta,i}^*$. Then, we bound the risk as follows,

$$\begin{aligned} \mathcal{R}_n^{\text{class}}(\theta, h_{\hat{\theta}}^*) &\leq \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{h_{\hat{\theta},i}^* \neq X_i\}} \mathbf{1}_{E_{n,i}} \cap F_{n,i} \right] + \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{h_{\hat{\theta},i}^* \neq X_i\}} \mathbf{1}_{E_{n,i}^c} \right] \\ &\quad + \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{h_{\hat{\theta},i}^* \neq X_i\}} \mathbf{1}_{F_{n,i}^c} \right] \\ &\leq \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbb{P}_{\theta}(h_{\hat{\theta},i}^* \neq X_i) \mathbf{1}_{E_{n,i}} \right] + \frac{1}{n} \sum_{i=1}^n \mathbb{P}_{\theta}(E_{n,i}^c) + \frac{1}{n} \sum_{i=1}^n \mathbb{P}_{\theta}(F_{n,i}^c) \end{aligned}$$

Finally, notice that,

$$\begin{aligned} \mathbb{P}_{\theta}(E_{n,i}^c) &= \mathbb{P}_{\theta} \left(p_{n,i}(\theta) < \frac{1}{2} + \gamma \right) \\ &= \mathbb{P}_{\theta} \left(\mathbb{P}_{\theta}(X_i = h_{\hat{\theta},i}^* \mid Y_{1:n}) < \frac{1}{2} + \gamma \right) \\ &= \mathbb{P}_{\theta} \left(\mathbb{P}_{\theta}(X_i \neq h_{\hat{\theta},i}^* \mid Y_{1:n}) > \frac{1}{2} - \gamma \right) \\ &\leq \frac{1}{1/2 - \gamma} \mathbb{E}_{\theta} \left(\mathbb{P}_{\theta}(h_{\hat{\theta},i}^* \neq X_i \mid Y_{1:n}) \mathbf{1}_{E_{n,i}^c} \right). \end{aligned}$$

Hence the result.

6.8. Equivalence of the definitions of the risk of clustering.

LEMMA 14. *The risk of clustering of $\pi_n \circ h$ can be rewritten as*

$$(22) \quad \mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h) := \mathbb{E}_{\theta} \left[\min_{\tau \in \mathcal{S}_J} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \right]$$

PROOF. It suffices to show that

$$\sup_{\substack{M \subseteq \mathcal{E}(\pi_n \circ h(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C') = \sup_{\tau \in \mathcal{S}_J} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) = \tau(X_i)}$$

Let $C_k = \{i \in [n] \mid h_i(Y_{1:n}) = k\}$ and $C'_k = \{i \in [n] \mid X_i = k\}$. Since the two partitions Π_n and $\pi_n \circ h(Y_{1:n})$ have the same number of clusters J (with possibly empty clusters), the supremum is reached on matchings with J edges. Using this fact, it follows that the matching reaching the supremum is of the form:

$$M = \left\{ (C_k, C'_{\tau(k)}) \mid 1 \leq k \leq J \right\}$$

where τ is a permutation of $\{1, \dots, J\}$. One obtains:

$$\begin{aligned}
 \sup_{\substack{M \subseteq \mathcal{E}(\pi_n \circ h(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C') &= \sup_{\tau \in \mathcal{S}_J} \sum_{k=1}^J \text{Card}(C_k \cap C'_{\tau(k)}) \\
 &= \sup_{\tau \in \mathcal{S}_J} \sum_{k=1}^J \sum_{i=1}^n \mathbf{1}_{\tau^{-1}(X_i)=h_i(Y_{1:n})=k} \\
 &= \sup_{\tau \in \mathcal{S}_J} \sum_{i=1}^n \mathbf{1}_{\tau(X_i)=h_i(Y_{1:n})}
 \end{aligned}
 \tag{23}$$

□

REFERENCES

- [1] ABRAHAM, K., CASTILLO, I. and GASSIAT, E. (2022). Multiple testing in nonparametric hidden Markov models: An empirical Bayes approach. *J. Mach. Learn. Res.* **23** 4061–4117.
- [2] ABRAHAM, K., GASSIAT, E. and NAULET, Z. (2023). Frontiers to the learning of nonparametric hidden Markov models. *arXiv preprint*. <https://doi.org/10.48550/arXiv.2306.16293>
- [3] ABRAHAM, K., GASSIAT, E. and NAULET, Z. (2023). Fundamental limits for learning hidden Markov model parameters. *IEEE Trans. Inform. Theory* **69** 1777–1794. <https://doi.org/10.1109/tit.2022.3213429> MR4564681
- [4] ALEXANDROVICH, G., HOLZMANN, H. and LEISTER, A. (2016). Nonparametric identification and maximum likelihood estimation for hidden Markov models. *Biometrika* **103** 423–434. <https://doi.org/10.1093/biomet/asw001> MR3509896
- [5] ARAGAM, B., DAN, C., P. XING, E. and RAVIKUMAR, P. (2020). Identifiability of Nonparametric mixture models and Bayes Optimal clustering. *Ann. Statist.* **48** 2277–2302. <https://doi.org/10.1214/19-AOS1887>
- [6] BORDES, L., MOTTELET, S. and VANDEKERKHOVE, P. (2006). Semiparametric estimation of a two-component mixture model. *Ann. Statist.* **34** 1204–1232. <https://doi.org/10.1214/009053606000000353>
- [7] CAPPÉ, O., MOULINES, E. and RYDEN, T. (2005). *Inference in Hidden Markov Models (Springer Series in Statistics)*. Springer-Verlag.
- [8] CASTRO, Y. D., GASSIAT, É. and LACOUR, C. (2016). Minimax Adaptive Estimation of Nonparametric Hidden Markov Models. *J. Mach. Learn. Res.* **17** 1–43.
- [9] CHRÉTIEN, S., DOMBRY, C. and FAIVRE, A. (2016). A Semi-Definite Programming approach to low dimensional embedding for unsupervised clustering. *arXiv preprint*. <https://doi.org/10.48550/arXiv.1606.09190>
- [10] COUVREUR, L. and COUVREUR, C. (2000). Wavelet-based non-parametric HMM’s: theory and applications. In *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.00CH37100)* **1** 604–607 vol.1. <https://doi.org/10.1109/ICASSP.2000.862054>
- [11] DE CASTRO, Y. (2016). Nonparametric HMM. <https://github.com/ydecastro/nonparametric-hmm>.
- [12] DE CASTRO, Y., GASSIAT, É. and LE CORFF, S. (2017). Consistent Estimation of the Filtering and Marginal Smoothing Distributions in Nonparametric Hidden Markov Models. *IEEE Trans. Inform. Theory* **63** 4758–4777. <https://doi.org/10.1109/TIT.2017.2696959>
- [13] DEVROYE, L., GYÖRFI, L. and LUGOSI, G. (1996). *A Probabilistic Theory of Pattern Recognition* **31**. <https://doi.org/10.1007/978-1-4612-0711-5>
- [14] DING, Q., HAN, J., ZHAO, X. and CHEN, Y. (2015). Missing-Data Classification With the Extended Full-Dimensional Gaussian Mixture Model: Applications to EMG-Based Motion Recognition. *IEEE Trans. on Indus. Elec.* **62** 4994–5005. <https://doi.org/10.1109/TIE.2015.2403797>
- [15] G MIXON, D., VILLAR, S. and WARD, R. (2017). Clustering subgaussian mixtures by semidefinite programming. *Information and Inference* **6** 389–415. <https://doi.org/10.1093/imaiai/iax001>
- [16] GAO, C., MA, Z., ZHANG, A. Y. and ZHOU, H. H. (2018). Community detection in degree-corrected block models. *Ann. Statist.* **46** 2153–2185. <https://doi.org/10.1214/17-AOS1615> MR3845014
- [17] GASSIAT, E. (2019). Mixtures of nonparametric components and hidden Markov models. In *Handbook of mixture analysis. Chapman & Hall/CRC Handb. Mod. Stat. Methods* 343–360. CRC Press, Boca Raton, FL. MR3889699

- [18] GASSIAT, E., CLEYNEN, A. and ROBIN, S. (2016). Inference in finite state space non parametric hidden Markov models and applications. *Stat. Comput.* **26** 61–71. <https://doi.org/10.1007/s11222-014-9523-8> MR3439359
- [19] GASSIAT, E. and ROUSSEAU, J. (2016). Non parametric finite translation mixtures with dependent regime. *Bernoulli* **22**. <https://doi.org/10.3150/14-BEJ631>
- [20] GHASSEMPOUR, S., GIROSI, F. and MAEDER, A. (2014). Clustering Multivariate Time Series Using Hidden Markov Models. *International Journal of Environmental Research and Public Health* **3**. <https://doi.org/10.3390/ijerph110302741>
- [21] GIRAUD, C. and VERZELEN, N. (2018). Partial recovery bounds for clustering with the relaxed K-means. *Math. Stat. Learn.* **3** 317–374.
- [22] GRÜN, B. (2019). Model-based clustering. In *Handbook of mixture analysis. Chapman & Hall/CRC Handb. Mod. Stat. Methods* 157–192. CRC Press, Boca Raton, FL. MR3889693
- [23] HOI, S. C., SAHOO, D., LU, J. and ZHAO, P. (2021). Online learning: A comprehensive survey. *Neuro-computing* **459** 249–289.
- [24] HUNTER, D. R., WANG, S. and HETTMANSBERGER, T. P. (2007). Inference for mixtures of symmetric distributions. *Ann. Statist.* **35** 224 – 251. <https://doi.org/10.1214/009053606000001118>
- [25] J. DANG, U., P. B. GALLAUGHER, M., P. BROWNE, R. and D. MCNICHOLAS, P. (2023). Model-Based Clustering and Classification Using Mixtures of Multivariate Skewed Power Exponential Distributions. *Journal of Classification* **40** 145–167.
- [26] KHIATANI, D. and GHOSE, U. (2017). Weather forecasting using Hidden Markov Model. In *2017 International Conference on Computing and Communication Technologies for Smart Nation (IC3TSN)* 220–225. <https://doi.org/10.1109/IC3TSN.2017.8284480>
- [27] LAMBERT, M. F., WHITING, J. P. and METCALFE, A. V. (2003). A non-parametric hidden Markov model for climate state identification. *Hydrology and Earth System Sciences* **7** 652–667. <https://doi.org/10.5194/hess-7-652-2003>
- [28] LE CORFF, S. and FORT, G. (2013). Online Expectation Maximization based algorithms for inference in Hidden Markov Models. *Electron. J. Statist.* **7** 763–792. <https://doi.org/10.1214/13-EJS789>
- [29] LE CORFF, S., FORT, G. and MOULINES, E. (2012). New Online EM Algorithms for General Hidden Markov Models. Application to the SLAM Problem. In *Proceedings of the 10th International Conference on Latent Variable Analysis and Signal Separation* 131–138. https://doi.org/10.1007/978-3-642-28551-6_17
- [30] LEHÉRICY, L. (2018). State-by-state minimax adaptive estimation for nonparametric hidden Markov models. *J. Mach. Learn. Res.* **19** Paper No. 39, 46. MR3862446
- [31] LEHÉRICY, L. (2021). Nonasymptotic control of the MLE for misspecified nonparametric hidden Markov models. *Electron. J. Stat.* **15** 4916–4965. <https://doi.org/10.1214/21-ejs1890> MR4347381
- [32] LU, Y. and H. ZHOU, H. (2016). Statistical and Computational Guarantees of Lloyd’s Algorithm and Its Variants. *arXiv preprint*. <https://doi.org/10.48550/arXiv.1612.02099>
- [33] LÖFFLER, M., Y. ZHANG, A. and ZHOU, H. H. (2021). Optimality of spectral clustering in the Gaussian mixture model. *Ann. Statist.* **49** 2506–2530. <https://doi.org/10.1214/20-AOS2044>
- [34] MARANDON, A., REBAFKA, T. and ROQUAIN, E. (2023). False clustering rate control in mixture models. *arXiv preprint*. <https://doi.org/10.48550/arXiv.2203.02597>
- [35] MCNICHOLAS, P. D. (2016). Model-based clustering. *Journal of Classification* **33** 331–373.
- [36] MEILÄ, M. (2005). Comparing clusterings: an axiomatic view. In *Proceedings of the 22nd international conference on Machine learning* 577–584.
- [37] MEILÄ, M. and HECKERMAN, D. (2001). An experimental comparison of model-based clustering methods. *Machine learning* **42** 9–29.
- [38] NDAOUD, M. (2022). Sharp optimal recovery in the two component Gaussian mixture model. *Ann. Statist.* **50** 2096 – 2126. <https://doi.org/10.1214/22-AOS2178>
- [39] PAULIN, D. (2015). Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electron. J. Probab.* **20** 1 – 32. <https://doi.org/10.1214/EJP.v20-4039>
- [40] SCHLIEP, A., SCHÖNHUTH, A. and STEINHOFF, C. (2003). Using hidden Markov models to analyze gene expression time course data. *Bioinformatics*. <https://doi.org/10.1093/bioinformatics/btg1036>
- [41] TEICHER, H. (1963). Identifiability of Finite Mixtures. *The Annals of Mathematical Statistics* **34** 1265 – 1269. <https://doi.org/10.1214/aoms/1177703862>
- [42] TEICHER, H. (1967). Identifiability of mixtures of product measures. *The Annals of Mathematical Statistics* **38** 1300–1302.
- [43] UEDA, N. and SAITO, K. (2002). Parametric Mixture Models for Multi-Labeled Text. In *Proceedings of the 15th International Conference on Neural Information Processing Systems* 737–744. MIT Press.
- [44] ZHANG, A. Y. and ZHOU, H. H. (2016). Minimax rates of community detection in stochastic block models. *Ann. Statist.* **44** 2252–2280. <https://doi.org/10.1214/15-AOS1428> MR3546450