



Order estimation for non-parametric Hidden Markov Models

Luc Lehéricy

► To cite this version:

| Luc Lehéricy. Order estimation for non-parametric Hidden Markov Models. 2016. hal-01321039v2

HAL Id: hal-01321039

<https://hal.science/hal-01321039v2>

Preprint submitted on 21 Nov 2016 (v2), last revised 6 Apr 2017 (v3)

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.



Distributed under a Creative Commons Attribution - ShareAlike 4.0 International License

Order estimation for non-parametric Hidden Markov Models

Luc Lehericy*

luc.lehericy@math.u-psud.fr

Abstract

We propose and study two practically tractable methods to estimate the order of non-parametric hidden Markov models in addition to their parameters. The first one relies on estimating the rank of a matrix derived from the law of two consecutive observations while the second one selects the order by minimizing a penalized least squares criterion. We show strong consistency of both methods and prove an oracle inequality on the least squares estimators of the model parameters. We numerically compare their ability to select the right order in several situations and discuss their algorithmic complexity.

1. Introduction

1.1. Context and motivation

Hidden Markov models (HMM in short) are natural tools to study time-evolving processes on heterogeneous populations. Non-parametric HMMs have already been considered in several papers like Couvreur and Couvreur (2000) for voice activity detection, Lambert et al. (2003) for climate state identification, Lefèvre (2003) for automatic speech recognition, Shang and Chan (2009) for facial expression recognition, Volant et al. (2014) for methylation comparison of proteins, Yau et al. (2011) for copy number variants identification in DNA analysis.

Formally, a hidden Markov model is a markovian process $(X_t, Y_t)_{t \geq 1}$ taking value in $\mathcal{X} \times \mathcal{Y}$ such that $(X_t)_{t \geq 1}$ is a Markov chain and the $(Y_t)_{t \geq 1}$ are such that they are independent given $(X_t)_{t \geq 1}$ and that the law of Y_t only depends on X_t . The $(X_t)_{t \geq 1}$ are assumed to be hidden, so that one only has access to the observations $(Y_t)_{t \geq 1}$.

*Laboratoire de Mathématiques d'Orsay, Univ. Paris-Sud, CNRS, Université Paris-Saclay, 91405 Orsay, France.

In this article, we make three additional assumptions: the Markov chain $(X_t)_{t \geq 1}$ takes its values in a finite set $\mathcal{X}$ of cardinal K^* , it is ergodic with initial law its stationary law, and the law of Y_t conditionally to $\{X_t = k\}$ has density f_k^* with respect to some sigma-finite measure μ . $\mathcal{Y}$ can be any space as long as such a measure exists.

Our goal is to estimate the parameters of the model, that is the number of hidden states K^* (which we call the *order* of the HMM), the transition matrix and the emission densities, by using only the observations $(Y_t)_{t \geq 1}$.

The estimation of the order is a key issue: even if some papers have proved theoretical results when estimating HMM parameters in the non-parametric setting, for instance de Castro et al. (2015a), the range of their work is limited by the assumption that the order of the HMM is known. However, in general, this number is unknown, so that their results can't be applied. Moreover, as far as we know, no consistency result has been proved about order selection for non-parametric HMMs. Even for parametric HMMs, there is still no general proof that a penalized maximum likelihood approach can lead to consistent order estimators without an *a priori* upper bound.

1.2. Contribution

In this paper, we introduce two different estimators of the order of a HMM. We first conduct a theoretical study and prove their consistency in a very general setting. Then, we assess their performances numerically and compare their behaviours.

Both methods start by projecting the non-parametric densities on a sequence of nested parametric subspaces, each equipped with an orthonormal basis, and estimating the corresponding parameters.

The first method uses the matrix containing the coordinates of the density of two consecutive observations in the orthonormal basis. When the emission densities are linearly independent, estimating the order of the HMM is equivalent to estimating the rank of this matrix. In practice, we only have access to an empirical estimator of this matrix, whose rank might differ from the order. We offer both a theoretical procedure to estimate the order by thresholding the singular values of the empirical matrix and a practical method based on a heuristics on the behaviour of its spectrum.

The second method looks for the smallest number of states such that the L^2 distance between the density of L observations and its projection on a given parametric subspace is minimal. If the subspace dimension is large enough, this is exactly the order of the HMM. In practice, we use an empirical equivalent of the L^2 distance to the real density

which is decreasing with the number of states, so that we have to penalize this criterion to obtain a consistent estimator.

In both cases, we prove the strong consistency of the estimator of the order. This step is crucial when one wants to estimate the other parameters of the HMM. For instance, spectral methods (see Anandkumar et al. (2012); Hsu et al. (2012); de Castro et al. (2015b)) need the right order to work properly. An overestimation of the order can lead to aberrant results.

Both of our estimators are derived from generic methods which also allow to recover the other parameters of the HMM. The first one comes from spectral methods, where previous results can be applied once the order has been estimated, see for instance de Castro et al. (2015b). The second one comes from a model selection procedure with a penalized least squares criterion. We prove an oracle inequality for the corresponding least squares estimators using a sharp control of the bracketing entropy of the models.

Our least squares method performs extremely well in almost any situation. It offers several advantages compared to previous methods: it does not need a preliminary estimation of the transition matrix or of the order, unlike de Castro et al. (2015a) who used spectral estimators, while still allowing to recover their adaptative minimax convergence rate for the estimation of the emission densities when we assume the order to be known. This is especially useful to avoid the cases where the spectral estimator fails, for instance the cases where the order is overestimated or where the states are almost linearly dependent, see for instance Lehericy (2015) or Figure 1. Then, it is very resistant to an overestimation of the order, both theoretically and numerically, thanks to our iterative initialization procedure. This initialization method consists in using estimators from smaller models as initial point for the minimization algorithm in order to avoid getting stuck in suboptimal local extrema. We believe it is of great practical interest since it produces very robust estimators and can also be used in other settings, for instance as initialization for expectation maximization algorithm for maximum likelihood estimators.

The spectral method performs slightly not as well as the least squares method, but its main advantage is its speed. It is easy to put in practice and runs extremely fast while giving very satisfying estimators when the number of observations is large enough.

1.3. Related works

Recently, several order estimation methods using penalized likelihood criterion were studied numerically, see for instance Volant et al. (2014) when emission laws are a mixture of parametric densities, Langrock et al. (2015) in a non-parametric case, or Celeux and Durand (2008) for parametric HMMs. The latter also introduced cross-validation procedures.

However, from a theoretical point of view, the order estimation problem remains widely open in the HMM framework. Using tools from coding theory, Gassiat and Boucheron (2003) introduced a penalized maximum likelihood order estimator for which they prove strong consistency without *a priori* upper bound on the order of the HMM. However, their result is restricted to a finite observation space. Even in the parametric HMM case, one usually needs an upper bound of the order to get weak consistency of order estimators for a penalized likelihood criterion (Gassiat (2002)) or a Bayesian approach (Gassiat et al. (2014); van Havre et al. (2016)). This result can be improved when restricting ourselves to Gaussian or Poisson emission laws: Chambaz et al. (2009) show that the penalized maximum likelihood estimator is strongly consistent without any *a priori* upper bound on the order.

Several methods to estimate the parameters of the HMM have been proposed and studied in the case where its order is known. The first problem was to understand when such models are identifiable. This was done in Gassiat et al. (2015) and Alexandrovich and Holzmänn (2014), who proved that they are identifiable as soon as one takes a sufficient number of observations. Hsu et al. (2012) and Anandkumar et al. (2012) proposed a spectral method to estimate the parameters of parametric HMMs and studied its convergency rate. de Castro et al. (2015b) extended this method to a non-parametric setting while de Castro et al. (2015a) used it to obtain an estimator of the transition matrix of the hidden chain and as initial point for a minimax adaptative least squares estimator of the emission densities.

1.4. Outline of the paper

We first introduce the notations, the model, the least squares estimation procedure and the assumptions we will consider in section 2.. In section 3., we introduce the spectral algorithm and propose a strongly consistent estimation procedure for the order, see Theorem 2.

Section 4. deals with the least squares estimation procedure. We first state an identifiability proposition which we use to prove strong consistency of the estimator of the order. This is done in two step. Firstly, we control the probability to underestimate the order. This is done thanks to Proposition 5, and gives an exponential bound on the probability of error, see Theorem 7. Secondly, we control the probability to overestimate the order, see Theorem 8. For this, we introduce a general condition which we use to prove polynomial decrease rate, and illustrate how to easily satisfy this condition. Finally, we state oracle inequalities on the estimators of the density of L consecutive observations and - under a delicate assumption - the parameters of the hidden Markov model, see Theorem 13.

In section 5., we propose practical algorithms to apply both methods. Firstly, we set the parameters on which we will test both procedures and introduce our iterative initialization algorithm for the least squares method. Secondly, we discuss how to practically implement the spectral thresholding, and introduce a heuristics to identify *significant* singular values in the empirical spectrum of the tensors. Thirdly, we discuss how to calibrate the penalty constant in the least squares method with the dimension jump heuristics or the slope heuristics. Lastly, we compare both methods and discuss their performance and complexity.

Our main technical result can be found at the beginning of section 6.. It is used extensively for both the consistency of the estimator of the order and the oracle inequalities on the HMM parameters. Its proof can be found in appendix C. The rest of this section is dedicated to the proofs of the results.

2. Definitions and assumptions

We will use the following notations.

- $\mathbb{N}^* = \{1, 2, \dots\}$ is the set of positive integers.
- If f_1 and f_2 are two functions, we denote by $f_1 \otimes f_2$ their tensor product, defined by $f_1 \otimes f_2(x_1, x_2) = f_1(x_1)f_2(x_2)$.
- $\text{Lin}(a)$ is the linear space spanned by the family a .
- If E_1 and E_2 are two linear spaces, we denote by $E_1 \otimes E_2$ their tensor product, that is the linear space spanned by the tensor products of their elements: $E_1 \otimes E_2 =$

$\text{Lin}(f_1 \otimes f_2 | f_1 \in E_1, f_2 \in E_2).$

- $\Delta_K = \{\pi \in [0, 1]^K | \sum_{k=1}^K \pi_k = 1\}$ is the simplex in dimension K . It will be seen as the set of probability measures on a finite set of size K .
- $\mathcal{Q}_K \subset \mathbb{R}^{K \times K}$ is the set of irreducible transition matrices of size K .
- $\mathbf{L}^2(A, \nu)$ is the Hilbert space of square integrable functions on A with respect to the measure ν .

In the following, L is a positive integer which will denote the number of consecutive observations used for the estimation procedure.

2.1. Hidden Markov models

Let $(X_j)_{j \geq 1}$ be a Markov chain with finite state space $\mathcal{X}$ of size K^* with transition matrix $\mathbf{Q}^*$ and initial distribution π^* . Without loss of generality, we can set $\mathcal{X} = \{1, \dots, K^*\}$.

Let $(Y_j)_{j \geq 1}$ be random variables on a measured space $(\mathcal{Y}, \mu)$ with μ σ -finite such that conditionally on $(X_j)_{j \geq 1}$ the Y_j 's are independent with a distribution depending only on X_j . Let ν_k^* be the law of Y_j conditionally to $\{X_j = k\}$. Assume that ν_k^* has density $f_k^* \in \mathbf{L}^2(\mathcal{Y}, \mu)$ with respect to μ . We call $(\nu_k^*)_{k \in \mathcal{X}}$ the emission laws and $\mathbf{f}^* = (f_1^*, \dots, f_{K^*}^*)$ the emission densities.

Then $(X_j, Y_j)_{j \geq 1}$ is a hidden Markov model (HMM in short) with parameters $(\pi^*, \mathbf{Q}^*, \mathbf{f}^*)$. The hidden chain $(X_j)_{j \geq 1}$ is assumed to be unknown, so that the estimator only has access to the observations $(Y_j)_{j \geq 1}$. Let g^* be the density of $(Y_1, \dots, Y_L)$ with respect to $\mu^{\otimes L}$.

For $K \in \mathbb{N}^*$, $\pi \in \mathbb{R}^K$, $\mathbf{Q} \in \mathbb{R}^{K \times K}$ and $\mathbf{f} \in (\mathbf{L}^2(\mathcal{Y}, \mu))^K$, let

$$g^{\pi, \mathbf{Q}, \mathbf{f}} = \sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \bigotimes_{i=1}^L f_{k_i}.$$

When π is a probability distribution, $\mathbf{Q}$ a transition matrix and $\mathbf{f}$ a K -uple of probability densities, $g^{\pi, \mathbf{Q}, \mathbf{f}}$ is the density of the first L observations of a HMM with parameters $(\pi, \mathbf{Q}, \mathbf{f})$. In particular, $g^* = g^{\pi^*, \mathbf{Q}^*, \mathbf{f}^*}$.

If $\mathbf{Q}$ is irreducible with stationary law π , we simply write $g^{\mathbf{Q}, \mathbf{f}}$.

2.2. Assumptions

We will need the following assumptions.

[HX] $(X_k)_{k \geq 1}$ is a stationary ergodic Markov chain with parameters $(\pi^*, \mathbf{Q}^*)$;

[HidA] $\mathbf{Q}^*$ is invertible, $L \geq 3$ and the family $\mathbf{f}^*$ is linearly independent;

[HidB] $\mathbf{Q}^*$ is invertible, $L \geq (2K^* + 1)((K^*)^2 - 2K^* + 2) + 1$ and the emission densities $(f_k^*)_{k \in \mathcal{X}}$ are all distinct;

[HF] $\mathbf{f}^* \in \mathcal{F}^{K^*}$, $\mathcal{F}$ is closed under projection on $\mathfrak{P}_M$ for all M , $\bigcup_M \mathfrak{P}_M$ is dense in $\mathcal{F}$ and

$$\forall f \in \mathcal{F}, \quad \begin{cases} \|f\|_\infty \leq C_{\mathcal{F},\infty} \\ \|f\|_2 \leq C_{\mathcal{F},2} \end{cases}$$

with $C_{\mathcal{F},\infty}$ and $C_{\mathcal{F},2}$ larger than 1.

[HidA] appears in spectral methods, with the hypothesis that $\pi^* > 0$ elementwise, see for instance Hsu et al. (2012). **[HidA]** and **[HidB]** also appear in identifiability issues, possibly combined with the stationarity hypothesis, see Alexandrovich and Holzmänn (2014) and Gassiat et al. (2015). Note that **[HidB]** only concerns the real order K^* . Lastly, **[HF]** will be used to control the deviations of the empirical loss.

Even though **[HidB]** appears less restrictive than **[HidA]** about the emission densities, two aspects make it very delicate to use in practice. The first one lies in the condition on the number of consecutive observations L . For **[HidB]**, we have to take L larger than an increasing function of the order, so it requires to have an *a priori* bound on the order to choose L . This is less interesting than **[HidA]**, which can work without *a priori* bound. The second problem is a technical one: all our algorithms need to compute L -dimensional tensors, which means they have an exponential complexity in L . Increasing the order leads to very time-consuming algorithms.

3. Spectral estimation

In this section, we assume **[HX]**, **[HidA]** and **[HF]**.

The idea of this method is to use the matrix containing the coordinates of the density of two consecutive observations in an orthonormal basis. Let $\Phi_M = (\varphi_1^{(M)}, \dots, \varphi_M^{(M)})$ be an orthonormal basis of $\mathfrak{P}_M$. For ease of notation, we will drop the dependency in M and write φ_a instead of $\varphi_a^{(M)}$. Let us introduce the matrix $\mathbf{N}_M$ and its empirical

estimator, defined by

$$\begin{aligned} \forall 1 \leq a, b \leq M, \quad \mathbf{N}_M(a, b) &:= \mathbb{E}[\varphi_a(Y_1)\varphi_b(Y_2)], \\ \forall 1 \leq a, b \leq M, \quad \hat{\mathbf{N}}_M(a, b) &:= \frac{1}{N} \sum_{s=1}^N \varphi_a(Y_1^{(s)})\varphi_b(Y_2^{(s)}). \end{aligned}$$

$\mathbf{N}_M$ contains the coordinates of the density of (Y_1, Y_2) with respect to $\mu^{\otimes 2}$ on the basis Φ_M . It holds that

$$\mathbf{N}_M = \mathbf{O}_M \text{Diag}(\pi^*) \mathbf{Q}^* \mathbf{O}_M^\top, \quad (1)$$

with $\mathbf{O}_M$ the coordinates of the emission densities on the orthonormal basis:

$$\forall 1 \leq m \leq M, \forall k \in \mathcal{X}, \quad \mathbf{O}_M(m, k) := \mathbb{E}[\varphi_m(Y_1) | X_1 = k] = \int \varphi_m f_k^* d\mu.$$

When the emission densities are linearly independent, $\mathbf{O}_M$ has full rank for M large enough.

The key remark for our method is that $\mathbf{N}_M$ contains explicit information about the order of the HMM, as stated in the following lemma:

Lemma 1. *There exists M_0 such that for all $M \geq M_0$, $\mathbf{N}_M$ has rank K^* .*

In practice, one only has access to the matrix $\hat{\mathbf{N}}_M$, which can be seen as a noisy version of $\mathbf{N}_M$. In particular, there is no reason for it to have only K^* nonzero singular values. On the contrary, the spectrum becomes noisy, and when some singular values of $\mathbf{N}_M$ are too small, they can be masked by this noise. As seen in equation (1), this can occur when $\mathbf{Q}^*$ or $\mathbf{O}_M$ are close to not having full rank, which means for $\mathbf{O}_M$ that the emission densities are almost linearly dependent.

Denote by $\sigma_1(A) \geq \sigma_2(A) \geq \dots$ the singular values of the matrix A . We can now state the theorem proving the consistency of the spectral order estimator:

Theorem 2. *Let $\hat{K}_{sp}(C) = \#\{i \mid \sigma_i(\hat{\mathbf{N}}_M) > C\sqrt{\log(N)/N}\}$.*

There exists C_0 and N_0 such that for all $C \geq C_0$ and $N \geq N_0$,

$$\mathbb{P}(\hat{K}_{sp}(C) \neq K^*) \leq N^{-2}$$

so that $\hat{K}_{sp}(C) \rightarrow K^$ almost surely.*

Proof. The following result from appendix E of de Castro et al. (2015b) allows to control the difference between the spectra of $\mathbf{N}_M$ and $\hat{\mathbf{N}}_M$.

Lemma 3. *There exists some constant C_* depending only on $\mathbf{Q}^*$ such that for any positive u , M and N ,*

$$\mathbb{P} \left[\|\mathbf{N}_M - \hat{\mathbf{N}}_M\|_F \geq \frac{\eta_2(\Phi_M)C_*}{\sqrt{N}}(1+u) \right] \leq e^{-u^2}$$

In particular, taking $u = \sqrt{2 \log(N)}$ and assuming $u > 1$, one has with probability $1 - N^{-2}$ that

$$\sigma_1(\mathbf{N}_M - \hat{\mathbf{N}}_M) \leq C \sqrt{\frac{\log(N)}{N}}$$

for all $C \geq C_0 := 2\sqrt{2}\eta_2(\Phi_M)C_*$, using that for any matrix A , one has $\sigma_1(A) \leq \|A\|_F$.

Let $C \geq C_0$. We will need Weyl's inequality (a proof may be found in Stewart and Sun (1990) for instance):

Lemma 4 (Weyl's inequality). *Let A, B be $p \times q$ matrices with $p \geq q$, then for all $i = 1, \dots, q$,*

$$|\sigma_i(A + B) - \sigma_i(A)| \leq \sigma_1(B).$$

Using this inequality, one gets that with probability at least $1 - N^{-2}$, for all $1 \leq i \leq K^*$, $\sigma_i(\hat{\mathbf{N}}_M) > \sigma_{K^*}(\mathbf{N}_M) - C\sqrt{\log(N)/N}$ and for all $i > K^*$, $\sigma_i(\hat{\mathbf{N}}_M) < C\sqrt{\log(N)/N}$.

In particular, if $2C\sqrt{\log(N)/N} < \sigma_{K^*}(\mathbf{N}_M)$, then with probability at least $1 - N^{-2}$, the order is exactly the number of singular values of $\hat{\mathbf{N}}_M$ which are larger than $C\sqrt{\log(N)/N}$. This concludes the proof. $\square$

In practice, C_0 might vary a lot depending on the parameters of the HMM, and we have no reliable way to estimate it. This is a downside of the spectral method when compared to the least squares method, for which we can use well-known calibration procedures like dimension jump or slope heuristics.

4. Least squares estimation

4.1. Approximation spaces and estimators

We want to estimate the density of L consecutive observations g^* by minimizing the quadratic loss $t \mapsto \|t - g^*\|_2^2 - \|g^*\|_2^2$. We thus take the corresponding empirical loss

$$\gamma_N(t) = \|t\|_2^2 - \frac{2}{N} \sum_{s=1}^N t(Z_s)$$

where $Z_s = (Y_s, \dots, Y_{s+L-1})$ for a single HMM (X, Y) of length $N + L - 1$.

Let $\mathcal{F}$ be some subset of $\mathbf{L}^2(\mathcal{Y}, \mu)$. Let $(\mathfrak{P}_M)_{M \in \mathcal{M}}$ be a sequence of nested subspaces of $\mathbf{L}^2(\mathcal{Y}, \mu)$ such that $\mathfrak{P}_M$ has dimension M and define for all $K \in \mathbb{N}^*$, $M \in \mathcal{M}$:

$$\begin{aligned} S_{K,M} &:= \{g^{\mathbf{Q}, \mathbf{f}}, \mathbf{Q} \in \mathcal{Q}_K, \mathbf{f} \in (\mathcal{F} \cap \mathfrak{P}_M)^K\} \\ S_K &:= \{g^{\mathbf{Q}, \mathbf{f}}, \mathbf{Q} \in \mathcal{Q}_K, \mathbf{f} \in \mathcal{F}^K\} \end{aligned}$$

In the following, we will always implicitly consider $M \in \mathcal{M}$.

For all K and M , we define the corresponding estimators

$$\hat{g}_{K,M} = g^{\hat{\mathbf{Q}}_{K,M}, \hat{\mathbf{f}}_{K,M}} \in \arg \min_{t \in S_{K,M}} \gamma_N(t)$$

where we dropped the dependency in N for ease of notation. Then, we select the parameters using the penalized empirical loss:

$$(\hat{K}_{\text{l.s.}}, \hat{M}) \in \arg \min_{K \leq N, M \leq N} \{\gamma_N(\hat{g}_{K,M}) + \text{pen}(N, M, K)\}$$

which leads to the estimators

$$\begin{aligned} \hat{g} &:= \hat{g}_{\hat{K}_{\text{l.s.}}, \hat{M}} \\ \hat{\mathbf{Q}} &:= \hat{\mathbf{Q}}_{\hat{K}_{\text{l.s.}}, \hat{M}} \\ \hat{\mathbf{f}} &:= \hat{\mathbf{f}}_{\hat{K}_{\text{l.s.}}, \hat{M}} \end{aligned}$$

4.2. Underestimation of the order

Note that the law of the HMM remains unchanged under permutation of the hidden states. We will therefore use the following pseudo-distance on the set of parameters:

Let $K \geq 1$, $\pi_1, \pi_2 \in \Delta_K$, $\mathbf{Q}_1$ et $\mathbf{Q}_2$ transition matrices of size K , $\mathbf{f}_1, \mathbf{f}_2 \in (\mathbf{L}^2(\mathcal{Y}, \mu))^K$. Let $\mathfrak{S}(\mathcal{X})$ be the set of permutations of $\mathcal{X}$. For all $\tau \in \mathfrak{S}(\mathcal{X})$, define the swapped parameters $\tau\pi_1$, $\tau\mathbf{Q}_1$ and $\tau\mathbf{f}_1$ by

$$\begin{aligned} (\tau\pi_1)(k) &:= \pi_1(\tau(k)) \\ (\tau\mathbf{Q}_1)(k, l) &:= \mathbf{Q}_1(\tau(k), \tau(l)) \\ (\tau\mathbf{f}_1)_k &:= \mathbf{f}_{1, \tau(k)} \end{aligned}$$

and finally

$$d_{\text{perm}}((\pi_1, \mathbf{Q}_1, \mathbf{f}_1), (\pi_2, \mathbf{Q}_2, \mathbf{f}_2)) := \inf_{\tau \in \mathfrak{S}(\mathcal{X})} \left(\|\tau\pi_1 - \pi_2\|_2^2 + \|\tau\mathbf{Q}_1 - \mathbf{Q}_2\|_F^2 + \sum_{k=1}^K \|(\tau\mathbf{f}_1)_k - f_{2,k}\|_2^2 \right)^{1/2}.$$

The following properties will be of use to prove the consistency of the order estimator, but we think it can also be of independent interest. The first one is a generalization of previous identifiability results from Alexandrovich and Holzmman (2014); Gassiat et al. (2015); de Castro et al. (2015a).

Proposition 5. *Let $K \geq 1$, $\pi \in \Delta_K$ such that $\pi_k > 0$ for all $k \in \mathcal{X}$, $\mathbf{Q}$ transition matrix of size K and $\mathbf{f} \in (\mathbf{L}^2(\mathcal{Y}, \mu))^K$ such that **[HidA]** or **[HidB]** hold. Then, for all $K' \geq 1$, for all $\pi' \in \Delta_{K'}$, for all transition matrix $\mathbf{Q}'$ of size K' and all $\mathbf{f}' \in (\mathbf{L}^2(\mathcal{Y}, \mu))^{K'}$, the following holds:*

$$(g^{\pi, \mathbf{Q}, \mathbf{f}} = g^{\pi', \mathbf{Q}', \mathbf{f}'} \text{ and } K' \leq K) \quad \Rightarrow \quad (K = K' \text{ and } d_{\text{perm}}((\pi, \mathbf{Q}, \mathbf{f}), (\pi', \mathbf{Q}', \mathbf{f}')) = 0).$$

Comment. *This property does not require two assumptions that appear in Alexandrovich and Holzmman (2014) and Gassiat et al. (2015): that $\mathbf{f}$ is a family of probability densities and that the Markov chain is stationary.*

In particular, the fact that $\mathbf{f}$ may not be a family of probability densities is crucial in the proof of Corollary 6, which is necessary to prove the strong consistency of the estimator of the order.

Proof. Assume **[HidA]**. The spectral algorithm from de Castro et al. (2015b) applied on the linear space spanned by both sets of densities allows to retrieve the order from two consecutive observations and the parameters from three consecutive observations. Their proof works when the emission densities are not probability densities and when the chain is not stationary.

Assume **[HidB]**. A careful reading of the proofs of Alexandrovich and Holzmman (2014) shows that their result can be extended to general observation spaces and do not require the measures to be probabilities. $\square$

The following corollary states that the $\mathbf{L}^2$ distance between the actual model and the

models where the order is underestimated is strictly positive. Note that we do not need $\mathcal{F}$ to be compact.

Corollary 6. *Assume [HX], ([HidA] or [HidB]) and [HF] hold. Then, for all $K < K^*$:*

$$d_K := \inf_{t \in S_K} \|t - g^*\|_2 > 0$$

Proof. Proof in section 6.2.1.. □

Our first theorem is an expected result: the probability to underestimate the order decreases exponentially with the number of observations. This comes from Corollary 6: since the empirical criterion converges to the L^2 distance, the penalized error will eventually become larger for orders under K^* than for orders over K^* , which means that we won't underestimate the real order. The exponential decrease rate is very similar to the one studied in Gassiat and Boucheron (2003): in both cases, the exponents involve the distance between the actual model and models with underestimated orders, as can be seen in our proof.

Theorem 7. *Assume [HX], ([HidA] or [HidB]) and [HF] hold. There exists positive constants ρ and β depending on $C_{\mathcal{F},2}$, $C_{\mathcal{F},\infty}$, $\mathbf{Q}^*$ and L such that the following holds.*

Assume that

$$\forall N, \forall M, \forall K, \quad \text{pen}(N, M, K) \geq \rho(MK + K^2 - 1) \frac{\log(N)}{N},$$

and

$$\forall M, \forall K, \quad \text{pen}(N, M, K) \xrightarrow{N \rightarrow \infty} 0$$

then there exists N_0 such that for all $N \geq N_0$,

$$\mathbb{P}(\hat{K}_{l.s.} < K^*) \leq e^{-\beta N}.$$

Proof. Proof in section 6.3.1.. □

4.3. Overestimation of the order

Our second theorem controls the probability to overestimate the order. It consists in overpenalizing large models so that the estimated order remains small.

We will need the following technical condition on the penalty:

Condition ([Hpen](α, ρ)). *The penalty function pen satisfies*

$$\exists N_1, \forall N \geq N_1, \forall M \leq N, \forall K \leq N \text{ s.t. } K > K^*,$$

$$\text{pen}(N, M, K) - \text{pen}(N, M, K^*) \geq \rho(MK + K^2 - 1) \frac{\log(N)}{N} + \alpha \frac{\log(N)}{N},$$

We can now state the theorem and its corollary proving the strong consistency of our estimator of the order.

Theorem 8. *Assume [HX] and [HF] hold. There exists positive constants ρ and β depending on $C_{\mathcal{F},2}, C_{\mathcal{F},\infty}, \mathbf{Q}^*$ and L such that the following holds.*

Assume [Hpen](α, ρ) holds for some $\alpha \geq 0$, then there exists N_0 such that for all $N \geq N_0$,

$$\mathbb{P}(\hat{K}_{l.s.} > K^*) \leq N^{-\beta\alpha}.$$

Proof. Proof in section 6.3.1.. □

Corollary 9. *Assume [HX], [HF] and ([HidA] or [HidB]) hold. There exists positive constants ρ and β depending on $C_{\mathcal{F},2}, C_{\mathcal{F},\infty}, \mathbf{Q}^*$ and L such that the following holds.*

Assume that the penalty function satisfies

$$\begin{cases} \forall N, \forall M \leq N, \forall K \leq N, & \text{pen}(N, M, K) \geq \rho(MK + K^2 - 1) \frac{\log(N)}{N} \\ \forall M, \forall K, & \text{pen}(N, M, K) \xrightarrow{N \rightarrow +\infty} 0 \end{cases}$$

and [Hpen]($\alpha/\beta, \rho$) holds for some $\alpha > 1$, then

$$\mathbb{P}(\hat{K}_{l.s.} \neq K^*) = O(N^{-\alpha}).$$

In particular, $\hat{K}_{l.s.} \rightarrow K^$ almost surely.*

Let us comment on this condition when using a penalty of the form $\text{pen}(N, M, K) = C(MK + K^2 - 1) \log(N)/N$ where C only depends on N .

- If one has an *a priori* bound on the order, i.e. if $K^* \leq K_0$ for some known K_0 , then direct computations show that for all α, ρ , there exists $C \geq 0$ (depending on K_0) such that [Hpen](α, ρ) holds for all $K^* \leq K_0$. This means that if one has an *a priori* bound on the order, then by taking C large enough, the probability that $\hat{K}_{l.s.} > K^*$ decreases with polynomial rate.

- If one does not have an *a priori* bound on K^* , taking a constant C does not allow to get $[\mathbf{Hpen}](\alpha, \rho)$ for all possible K^* , which means we can't apply Corollary 9. However, by taking C as a sequence indexed by N that tends to infinity, we get that for all K^* and α, ρ , $[\mathbf{Hpen}](\alpha, \rho)$ holds. This implies consistency with polynomial decrease of the probability of error, at the cost of overpenalizing.

Overpenalizing is actually necessary if one wants to satisfy $[\mathbf{Hpen}]$ for all K^* . This is stated in the following proposition:

Proposition 10. *Let $\rho > 0$ and pen be a positive penalty such that for all K^* , $[\mathbf{Hpen}](0, \rho)$ holds, then there exists a sequence $(u_N)_{N \geq 1} \rightarrow \infty$ such that for all $N \geq 1$, $M \leq N$ and $K \leq N$, $pen(N, M, K) \geq u_N(MK + K^2 - 1) \log(N)/N$.*

Proof. Proof in section 6.3.2.. □

4.4. Oracle inequalities

Our first result for this section is an oracle inequality on the density of L consecutive observations for the least squares estimator.

Theorem 11. *Assume $[\mathbf{HX}]$ and $[\mathbf{HF}]$ hold. Then there exist positive constants N_0 , ρ and A depending on $C_{\mathcal{F},2}$, $C_{\mathcal{F},\infty}$, $\mathbf{Q}^*$ and L such that if the penalty satisfies*

$$\forall N, \forall M \leq N, \forall K \leq N, \quad pen(N, M, K) \geq \rho(MK + K^2 - 1) \frac{\log(N)}{N}$$

then for all $N \geq N_0$, for all $x > 0$, it holds with probability larger than $1 - e^{-x}$ that

$$\|\hat{g} - g^*\|_2^2 \leq 4 \inf_{K \leq N, M \leq N} \{ \|g_{K,M}^* - g^*\|_2^2 + pen(N, M, K) \} + 4A \frac{x}{N}$$

Proof. Proof in section B1.. □

Comment. *We can replace the constant 4 before the infimum by any constant $\kappa > 1$, at the cost of changing the constants N_0 , ρ and A .*

We would like to deduce an oracle inequality on the parameters of the HMM from this result. It is easy to bound the error on the density g^* by the error on the parameters:

Lemma 12. *Assume $[\mathbf{HF}]$ holds. Let $K \geq 1$, $\pi_1, \pi_2 \in \Delta_K$, $\mathbf{Q}_1$ and $\mathbf{Q}_2$ transition matrices of size K and $\mathbf{f}_1, \mathbf{f}_2 \in \mathcal{F}^K$. Then*

$$\|g^{\pi_1, \mathbf{Q}_1, \mathbf{f}_1} - g^{\pi_2, \mathbf{Q}_2, \mathbf{f}_2}\|_2 \leq C_{\mathcal{F},2}^L \sqrt{LK} d_{perm}((\pi_1, \mathbf{Q}_1, \mathbf{f}_1), (\pi_2, \mathbf{Q}_2, \mathbf{f}_2)).$$

Proof. Proof in section B2.. □

Thus, all we need to deduce an oracle inequality on the parameters is to lower bound the error on g^* by the error on the parameters.

To this end, we introduce the following condition **[Hlow](c)**

Condition ([Hlow](c)). For all $\pi \in \Delta_{K^*}$, for all transition matrix $\mathbf{Q}$ of size K^* and for all $\mathbf{f} \in \mathcal{F}^{K^*}$,

$$\|g^{\pi, \mathbf{Q}, \mathbf{f}} - g^{\pi^*, \mathbf{Q}^*, \mathbf{f}^*}\|_2 \geq c \cdot d_{\text{perm}}((\pi, \mathbf{Q}, \mathbf{f}), (\pi^*, \mathbf{Q}^*, \mathbf{f}^*)).$$

The authors of de Castro et al. (2015a) study a similar inequality and prove that it holds under a generic hypothesis on the parameters: it is enough that a known polynomial in the coefficients of π^* , $\mathbf{Q}^*$ and the Gram matrix $(\langle f_k^*, f_{k'}^* \rangle)_{k, k' \in \mathcal{X}}$ isn't equal to zero (see Theorem 6 of de Castro et al. (2015a)). This condition can be generalized to the present setting.

The following theorem is a direct consequence of the above results. It provides an oracle inequality on the parameters conditionally to the fact that the order has been correctly estimated.

Theorem 13. Assume **[HX]**, **[HF]** and **[Hlow](c)** hold for some constant $c > 0$. Then there exist positive constants N_0 , ρ and A depending on $C_{\mathcal{F}, 2}$, $C_{\mathcal{F}, \infty}$, $\mathbf{Q}^*$ and L such that if the penalty satisfies

$$\forall N, \forall M \leq N, \forall K \leq N, \quad \text{pen}(N, M, K) \geq \rho(MK + K^2 - 1) \frac{\log(N)}{N}$$

then for all $N \geq N_0$, for all $x > 0$, conditionally to $\{\hat{K}_{\text{l.s.}} = K^*\}$, with probability larger than $1 - e^{-x}$:

$$d_{\text{perm}}((\hat{\pi}, \hat{\mathbf{Q}}, \hat{\mathbf{f}}), (\pi^*, \mathbf{Q}^*, \mathbf{f}^*)) \leq \frac{4C_{\mathcal{F}, 2}^L \sqrt{LK^*}}{c} \times \left[\inf_{M \leq N} \left\{ \sum_{k=1}^{K^*} \|f_{M,k}^* - f_k^*\|_2^2 + \text{pen}(N, M, K^*) \right\} + A \frac{x}{N} \right],$$

where $f_{M,k}^*$ is the projection of f_k^* on $\mathfrak{P}_M$.

5. Numerical experiments

5.1. Framework and description

We will consider $\mathcal{Y} = [0, 1]$ with μ being the Lebesgue measure. We will use the trigonometric basis on $L^2([0, 1])$ to generate the approximation spaces $(\mathfrak{P}_M)_M$. More precisely, define

$$\begin{aligned}\varphi_0(t) &= 1 \\ \varphi_a(t) &= \sqrt{2} \cos(\pi a t)\end{aligned}$$

for all $t \in [0, 1]$ and $a \in \mathbb{N}^*$. We take $\mathfrak{P}_M = \text{Lin}(\{\varphi_a \mid 0 \leq a < M\})$ the spaces induced by the trigonometric basis.

Comment. – *Taking the same vectors in all bases is not mandatory to obtain theoretical consistency, but in practice it allows us to take an additional initial point for the minimization step and improves the stability of the algorithm (see step 1 below).*

– *Our theoretical study dealt with nested sequences of spaces, but in practice it is possible to take non-nested sequences such as piecewise constant functions on the regular partition of size M .*

We will assume $\mathbf{f}$ to be linearly independent, so that we only need to take 3 observations to estimate the parameters of the HMM.

In order to assess the performances of the different procedures, we generate N observations of a HMM of order 3 for several values of N , using the following parameters:

- Emission laws: Beta laws with parameters (2; 5), (4; 2) and (4; 4) (case **A**) or (1.5; 5), (7; 2) and (6; 6) (case **B**);
- Markov chain parameters:

$$\begin{aligned}\mathbf{Q}^* &= \begin{pmatrix} 0.8 & 0.1 & 0.1 \\ 0.2 & 0.7 & 0.1 \\ 0.07 & 0.13 & 0.8 \end{pmatrix}, \\ \pi^* &= \left(\frac{47}{120} \quad \frac{11}{40} \quad \frac{1}{3} \right) \\ &\approx (0.3917 \quad 0.2750 \quad 0.3333).\end{aligned}$$

Finally, we take $M_{\max} = 50$ and $K_{\max} = 5$.

The simulation codes are available in MATLAB at https://www.normalesup.org/~lehericy/HMM_order_simfiles/.

Spectral method In order to estimate the order, we use the method from section 3.. However, since the calibration of the constant is a delicate task, we use an empirical heuristics to threshold the singular values.

In order to estimate the HMM parameters, we will use the algorithm presented in appendix A. It is the spectral algorithm studied by de Castro et al. (2015b), who also prove theoretical guarantees on the error of the estimators. As shown by de Castro et al. (2015a), it can be used as initial point for a least squares minimization procedure, leading to minimax adaptative estimators as soon as the order of the HMM is known.

Least squares method Since γ_N is not convex, we use CMAES (for Covariance Matrix Adaptation Estimation Strategy, see Hansen (2006)) in order to find a minimizer. This algorithm requires a good initial point since it might otherwise remain stuck in local minima. One part of our method consists in using previous estimates as initial points for further steps since it is likely that this way the estimators stay near the real minimizer.

Here are the main steps of the algorithm:

1. Minimize γ_N on each model, for $M \leq M_{\max}$ and $K \leq K_{\max}$. We take several initial points for model (K, M) according to the following cases:
 - $K = 1$. Use a HMM with a single state and a uniform emission law.
 - $K > 1$. Take the estimator from model $(K - 1, M)$. For each hidden state of the corresponding HMM, use the model where this state is duplicated. More precisely, the Markov chain $\tilde{X}$ where state I is duplicated is obtained by replacing the state I from chain X with two states I_1 and I_2 such that for each state $S \neq I_1, I_2$,

$$\begin{aligned}
\mathbb{P}(\tilde{X}_{t+1} = I_1 \mid \tilde{X}_t = S) &= \frac{1}{2} \mathbb{P}(X_{t+1} = I \mid X_t = S) \\
&= \mathbb{P}(\tilde{X}_{t+1} = I_2 \mid \tilde{X}_t = S) \\
\mathbb{P}(\tilde{X}_{t+1} = S \mid \tilde{X}_t = I_1) &= \mathbb{P}(X_{t+1} = S \mid X_t = I) \\
&= \mathbb{P}(\tilde{X}_{t+1} = S \mid \tilde{X}_t = I_2)
\end{aligned}$$

and

$$\begin{aligned}\mathbb{P}(\tilde{X}_{t+1} = I_2 \mid \tilde{X}_t = I_1) &= \frac{1}{2}\mathbb{P}(X_{t+1} = I \mid X_t = I) \\ &= \mathbb{P}(\tilde{X}_{t+1} = I_1 \mid \tilde{X}_t = I_2)\end{aligned}$$

- $M > 1$. Use estimator from model $(K, M-1)$ with the M -th coordinate of each emission density set to zero. This is only interesting if all $\mathfrak{P}_M$ are induced by the first M vectors of a given orthonormal basis, like for trigonometric spaces.


Then, after minimization from each one of these initial points, take the estimator that minimizes γ_N .

2. Tune the parameter ρ of the penalty with the slope heuristics or the dimension jump method and select $\hat{M}$ and $\hat{K}$ (see section 5.2.2.)
3. Return the estimator for $M = \hat{M}$ and $K = \hat{K}$.


This iterative initialization procedure relies on the heuristics that when the order is underestimated, then several states are "merged" together. Duplicating a merged state will allow to separate them effectively while still taking advantage of the computations done up to now. It is meant to avoid having to recalculate all states at the same time (which could get us stuck in sub-optimal local minima) when the best solution is likely to be a small modification of the previous estimator. In addition, when the order is overestimated, it allows to make sure the empirical criterion is indeed decreasing with the dimension of the model by giving an estimator that performs at least as well as those from smaller models.

5.2. Order estimation

Figure 1 shows the estimators of the emission densities of the least squares and spectral algorithms when considering $N = 19,998$ consecutive observations of a HMM with the above parameters. The constant of the penalty is calibrated using dimension jump heuristics as described in section 5.2.2.. The least squares algorithm chooses the correct order for the HMM in both cases, while the spectral algorithm chooses 2 states in case A.



(a) Case **A**. We took $K = \hat{K}_{\text{l.s.}} = 3$ and $M = \hat{M} = 13$. The bad behaviour of the spectral algorithm when the emission densities are poorly separated is clearly visible on the third emission law.



(b) Case **B**. We took $K = \hat{K}_{\text{l.s.}} = 3$ and $M = \hat{M} = 22$.

Figure 1: Estimators of the emission densities for $N = 19,998$.

5.2.1. Spectral estimation

Let $\hat{K}_{\text{sp.}}(C) = \#\{i \mid \sigma_i(\hat{\mathbf{N}}_M) > C\sqrt{\log(N)/N}\}$. We showed in section 3. that if one takes C larger than some constant C_0 , then this estimator is strongly consistent. However, two problems arise. The first is that the constant C_0 depends on the parameters of the HMM. The second is that the minimum number of observations after which the order is correctly estimated depends on the K^* -th singular value of $\mathbf{N}_M$, which is also a function

of the parameters of the HMM. If one singular value of N_M is smaller than the threshold (which is the case when O_M is close from not being invertible, i.e. that the emission densities are close from being linearly dependent), then this method will not be able to "see" the corresponding hidden state. Figure 2 (in particular 2a and 2d) illustrates this problem, and Figure 1 shows the result when trying to estimate the densities: when the singular value is drowned by the noise, the estimated density is aberrant.




Figure 2: Spectrum of the empirical matrix $\hat{N}_M$ and the theoretical matrix N_M for $M = 40$ and 10 simulations. The first one (in case **A**) or two (in case **B**) singular values are too large to appear here.

Our heuristics for thresholding will not try to estimate the right constant. Instead, it relies on the fact that when one sorts the singular values in decreasing order, then the smallest ones approximately follow an affine relation with respect to their index. This appears in Figure 3.

We proceed as follows. Let M and M_{reg} be two positive integers such that $M_{\text{reg}} \leq M \leq$

$M_{\max}$. We estimate the affine dependence of the singular values of $\hat{\mathbf{N}}_M$ with respect to their index using its M_{reg} smallest singular values. We take $M_{\text{reg}} = M - K_{\max}$ in order to take the largest possible number of singular values that are not associated to an actual state. Then, we set a thresholding parameter $\tau > 1$. We say a singular value is *significant* if it is above τ times the value that the regression predicts for it. Lastly, we choose K as the number of consecutive significant singular values starting from the largest one.

We do not have any theoretical justification for this heuristics, even if it seems to work as soon as τ is large enough, e.g. $\tau = 1.5$.




Figure 3: Spectrum of $\mathbf{N}_M$ for $M = 40$ and $N = 49,998$ in case **A**. The regression (green line) has been performed on the 35 smallest singular values. The two largest singular values are too large to appear here.

5.2.2. Penalty calibration

Dimension jump heuristics In this paragraph, we study the selected parameters

$$\rho \mapsto (\hat{M}(\rho), \hat{K}(\rho)) \in \arg \min \{ \gamma_N(\hat{g}_{K,M}) + \rho \text{pen}_{\text{shape}}(N, M, K) \}$$

and the selected complexity

$$\rho \mapsto \text{Comp}(\rho) = \hat{M}(\rho)\hat{K}(\rho) + \hat{K}(\rho)[\hat{K}(\rho) - 1]$$

with $\text{pen}_{\text{shape}}(N, M, K) = (MK + K^2 - 1) \log(N)/N$.

Assume that there exists κ such that $\kappa \text{pen}_{\text{shape}}$ is a *minimal penalty*, that is a penalty such that as N tends to infinity, for all $\rho > \kappa$, the size of the model chosen for penalty $\rho \text{pen}_{\text{shape}}$ remains small in some sense and for all $\rho < \kappa$, the size of the model becomes huge. Then, for N large enough, this will appear on the graph of the selected model complexity as a "dimension jump": around some constant κ , the complexity will drop from large models to small models. This is clearly the case in Figure 4.

Figure 5 shows the behaviour of $\hat{M}$ and $\hat{K}$ with ρ . A dimension jump also occurs with these functions. It is most visible for $\hat{M}$.

Once the dimension jump location ρ_{jump} has been estimated, we take $\hat{\rho} = 2\rho_{\text{jump}}$ to select the final parameters (this is usual when calibrating penalties, see for instance Baudry et al. (2012)).




Figure 4: Graph of $\rho \mapsto \text{Comp}(\rho)$ for 10 sets of N consecutive observations, case A.

Slope heuristics This heuristics relies on the fact that when $\text{pen}_{\text{shape}}$ is a minimal penalty, then the empirical contrast function is expected to behave like $\rho_{\min} \text{pen}_{\text{shape}}$ for large models and for some constant $\rho_{\min}$. This gives both a way to calibrate the constant of the penalty and to check if the chosen penalty has the right shape (see Baudry et al. (2012)). The final penalty is then taken as $2\hat{\rho}_{\min} \text{pen}_{\text{shape}}$.

Figure 6 shows the graph of the empirical contrast depending on $\text{pen}_{\text{shape}}$. The slope heuristics works very well in this situation, suggesting that our penalty has the right shape.




Figure 5: Graph of $\rho \mapsto \hat{M}(\rho)$ and $\rho \mapsto \hat{K}(\rho)$ for 10 sets of N consecutive observations, case **A**.




Figure 6: Empirical contrast and calibrated penalty for $N = 49,998$, case **A**.

5.2.3. Complexity

The complexity of the spectral method for a single (K, M) is linear in the number of observations. It is bounded by $O(NM^3 + M^2K^2 + MK^3)$. The NM^3 comes from computing the spectral tensors, and the rest comes from standard matrix operations. Thus, computing it on all $1 \leq M \leq M_{\max}$ has complexity $O(NM_{\max}^4 + M_{\max}^3K^2 + M_{\max}^2K^3)$.

The complexity of the least squares method run on all $1 \leq K \leq K_{\max}$ and $1 \leq M \leq M_{\max}$ is a little more delicate to handle. It also starts by computing the empirical tensors $\hat{\mathbf{M}}_M$ with $O(NM^3)$ operations, since this allows to compute the empirical

contrast γ_N in $O(M^3)$ operations. The difficulty comes from the minimization algorithm. Let's assume it is an iterative procedure that stops after N_{stop} steps, like CMAES. Then, since one has to perform $O(M_{\text{max}}K_{\text{max}}^2)$ minimizations, the total complexity will be $O(NM_{\text{max}}^4 + N_{\text{stop}}M_{\text{max}}^4K_{\text{max}}^2)$. When one wants to estimate the density of L consecutive observations instead of 3, the complexity is $O(NM_{\text{max}}^{L+1} + N_{\text{stop}}M_{\text{max}}^{L+1}K_{\text{max}}^2)$. This exponential complexity in L makes taking more consecutive observations into account very costly and makes using **[HidB]** nearly impossible since it requires L to be larger than $2K_{\text{max}}^3 + O(K_{\text{max}}^2)$.

An important aspect is how the complexity varies with the number of observations for both methods. In the spectral method, the largest term is usually NM_{max}^4 and comes from the computation of the empirical tensors. Once these tensors are computed, the cost of the algorithm does not depend on N at all. For the least squares method on the other hand, it is important to notice that N_{stop} might and should depend on N : it is crucial to take it sufficiently large in order for the minimization algorithm to reach an acceptable precision. In our examples, this made the minimization step the most time-consuming part of the algorithm, and it required much more time than the spectral methods to compute. This is the main limitation of the least squares method.

5.2.4. Discussion

N	$\mathbb{P}(\hat{K}_{\text{l.s.}} = K^*)$	$\mathbb{P}(\hat{K}_{\text{sp.}} = K^*)$
7 500	0.3	0
9 999	0.3	0
19 998	0.9	0
30 000	1	0
49 998	1	0.1

(a) Case **A**.

N	$\mathbb{P}(\hat{K}_{\text{l.s.}} = K^*)$	$\mathbb{P}(\hat{K}_{\text{sp.}} = K^*)$
999	0.2	0
3 000	1	0
9 999	1	1
19 998	1	1

(b) Case **B**.

Figure 7: Probability to select the right order for the two methods ($\hat{K}_{\text{l.s.}}$ for the least squares method and $\hat{K}_{\text{sp.}}$ for the spectral method). 10 simulations have been done for each N . Parameters for spectral selection are $M = 40$, $M_{\text{reg}} = 35$ and $\tau = 1.5$.

Figure 7 summarizes the results of both procedures. The least squares method can identify the order correctly with notably less observations than the spectral method.

The spectral method is easily put in practice and runs extremely fast. It doesn't need a time-consuming contrast minimization step or an initial point. Theoretical results (see de Castro et al. (2015b)) control its error on the HMM parameters nonasymptotically and it allows to efficiently get an estimator of the order. However, the thresholding of the singular values is a delicate problem and we have no guarantee on the practical estimator of the order. Moreover, if the order is incorrect, then the theoretical results about the spectral estimators of the parameters don't hold and this method may behave poorly.

The performances of the least squares method are much better (see for instance de Castro et al. (2015a) for a comparison of the two methods for estimating the emission densities, and Figure 7 for comparing the order estimators). In addition, the model selection step is very easy to handle and gives an estimator of the order that we proved to be consistent. However, this comes at a cost. The minimization of the empirical contrast is a very time-consuming step, especially for large samples and large models. The choice of the initial point is also a well-known issue in this setting since the minimization algorithm might get stuck in a local minima, but our iterative initialization procedure should solve it. Another advantage of this method is that it is very robust in case of an overestimation of the order, since it will just duplicate one of the real states. When compared to maximum likelihood methods, this criterion has the advantage to be very simple to compute: it does not need a time-consuming forward-backward pass and can be computed in constant time after one linear-time preparation. However, there is no algorithm as convenient as EM for the optimization step, so that one has to resort to non-convex minimization methods such as CMAES.

6. Proofs

6.1. Main technical result

The following lemma is the main technical result of this paper. It is the key for both the oracle inequalities and the order estimator consistency. It allows to control the difference between the empirical and theoretical losses.

Define $\nu : t \mapsto \frac{1}{N} \sum_{s=1}^N t(Z_s) - \int t g^*$, so that

$$\forall t \in \mathbf{L}^2(\mathcal{Y}^L, \mu^{\otimes L}), \quad \gamma_N(t) + \|g^*\|_2^2 = \|t - g^*\|_2^2 - 2\nu(t) \quad (2)$$

Let

$$\begin{aligned}
s = (s_{K,M})_{K,M} \in \mathbf{S} &:= \prod_{K \in \mathbb{N}^*, M \in \mathcal{M}} \left(\bigcup_K S_K \right) \\
&\longmapsto (Z_{K,M}(s))_{K,M} \\
&:= \left(\sup_{t \in S_{K,M}} \left[\frac{|\nu(t - s_{K,M})|}{\|t - s_{K,M}\|_2^2 + x_{K,M}^2} \right] \right)_{K,M} \quad (3)
\end{aligned}$$

Comment. It is not necessary to assume that $s_{K,M} \in S_{K,M}$. In particular, one can take $s_{K,M} = g^*$ for all K, M . In that case, we will simply write $Z_{K,M}(g^*)$.

Lemma 14. Assume [HX] and [HF] hold. Then there exists a sequence $(x_{K,M})_{K,M}$ and positive constants N_0, ρ and A depending on $C_{\mathcal{F},2}, C_{\mathcal{F},\infty}, \mathbf{Q}^*$ and L such that if the penalty $\widetilde{\text{pen}}$ satisfies

$$\forall N, \forall M \leq N, \forall K \leq N \quad \widetilde{\text{pen}}(N, M, K) \geq \rho(MK + K^2 - 1) \frac{\log(N)}{N}$$

then for all $s \in \mathbf{S}$, $N \geq N_0$ and $x > 0$, one has with probability larger than $1 - e^{-x}$:

$$\begin{cases} \sup_{K' \leq N, M' \leq N} Z_{K',M'}(s) \leq \frac{1}{4} \\ \sup_{K' \leq N, M' \leq N} (2Z_{K',M'}(s)x_{K',M'}^2 - \widetilde{\text{pen}}(N, M', K')) \leq A \frac{x}{N} \end{cases}$$

Comment. We can replace the constant $1/4$ in the first upper bound by any $\epsilon > 0$, at the cost of changing the constants N_0, ρ and A .

The first step is to obtain a Bernstein-like inequality on the densities, which allows to control the deviations of $|\nu(t)|$ for all t in the models. This is done using results on ergodic Markov chains from Paulin (2013). The crucial step of the proof, and our main contribution, is to obtain a sharp control of the bracketing entropy of the model, which is necessary to control the deviations of the supremum of $|\nu(t)|$ on the whole model. We had to introduce a much finer control of the models than in de Castro et al. (2015a) in order to take the dependency in the order of the HMM into account. The penalty from Lemma 14 comes directly from this step. Finally, it is possible using a peeling lemma from Massart (2007) to transpose this result on the $Z_{K,M}$.

The details of the proof can be found in appendix C.

6.2. Identifiability proofs

6.2.1. Proof of Corollary 6

Denote by Proj_A the orthogonal projection on a linear space A .

Since the union of $(\mathfrak{P}_M)_{M \in \mathcal{M}}$ is dense in $\mathcal{F}$, we can take M such that **[HidA]** or **[HidB]** holds for $\mathbf{f}_M^* = (f_{M,k}^*)_{k \in \mathcal{X}} := (\text{Proj}_{\mathfrak{P}_M} f_k^*)_{k \in \mathcal{X}}$.

We will need the following lemma.

Lemma 15.

$$\forall \pi \in \mathbb{R}^K, \quad \forall \mathbf{Q} \in \mathbb{R}^{K \times K}, \quad \forall \mathbf{f} \in \mathcal{F}^K, \quad \forall M, \quad \text{Proj}_{\mathfrak{P}_M^{\otimes L}}(g^{\pi, \mathbf{Q}, \mathbf{f}}) = g^{\pi, \mathbf{Q}, \text{Proj}_{\mathfrak{P}_M}(\mathbf{f})}$$

Proof. By linearity of the projection operator, it is enough to prove that for all $(t_1, \dots, t_L) \in (\mathbf{L}^2(\mathcal{Y}, \mu))^L$,

$$\text{Proj}_{\mathfrak{P}_M^{\otimes L}}(t_1 \otimes \dots \otimes t_L) = \text{Proj}_{\mathfrak{P}_M}(t_1) \otimes \dots \otimes \text{Proj}_{\mathfrak{P}_M}(t_L)$$

which is easy to check. $\square$

We will make a proof by contradiction. Assume that $\inf_{t \in S_K} \|t - g^*\|_2 = 0$ for some $K < K^*$. Then there exists a sequence $(g_n)_{n \geq 1} = (g^{\pi_n, \mathbf{Q}_n, \mathbf{f}_n})_{n \geq 1}$ such that $g_n \rightarrow g^*$ in $\mathbf{L}^2(\mathcal{Y}^L, \mu^{\otimes L})$, with $\pi_n \in \Delta_K$, $\mathbf{Q}_n$ a transition matrix of size K and $\mathbf{f}_n \in \mathcal{F}^K$.

The orthogonal projection on $\mathfrak{P}_M^{\otimes L}$ is continuous, so by using Lemma 15, one gets that

$$g^{\pi_n, \mathbf{Q}_n, \text{Proj}_{\mathfrak{P}_M}(\mathbf{f}_n)} \rightarrow g^{\pi^*, \mathbf{Q}^*, \mathbf{f}_M^*}$$

Then, using the compacity of Δ_K and of the set of transition matrices of size K and the relative compacity of $(\mathcal{F} \cap \mathfrak{P}_M)^K$ (which is a bounded subset of a finite dimension linear space), one gets (up to extraction of a subsequence) that there exists $\pi_\infty \in \Delta_K$, $\mathbf{Q}_\infty$ a transition matrix of size K and $\mathbf{f}_\infty \in (\mathfrak{P}_M)^K$ such that $\pi_n \rightarrow \pi_\infty$, $\mathbf{Q}_n \rightarrow \mathbf{Q}_\infty$ and $\text{Proj}_{\mathfrak{P}_M}(\mathbf{f}_n) \rightarrow \mathbf{f}_\infty$.

Finally, using the continuity of the function $(\pi, \mathbf{Q}, \mathbf{f}) \mapsto g^{\pi, \mathbf{Q}, \mathbf{f}}$ and the unicity of the limit, one gets

$$g^{\pi_\infty, \mathbf{Q}_\infty, \mathbf{f}_\infty} = g^{\pi^*, \mathbf{Q}^*, \mathbf{f}_M^*}.$$

Then Proposition 5 contradicts the assumption $K < K^*$, which is enough to conclude.

6.3. Consistency proofs

6.3.1. Proof of Theorems 7 and 8

The definition of $\hat{K}_{\text{l.s.}}$ is equivalent to the following one:

$$\hat{K}_{\text{l.s.}} \in \arg \min_{K \leq N} \{ \gamma_N(\hat{g}_{K, \hat{M}_K}) + \text{pen}(N, \hat{M}_K, K) \}$$

where

$$\hat{M}_K \in \arg \min_{M \leq N} \{ \gamma_N(\hat{g}_{K, M}) + \text{pen}(N, M, K) \}$$

Choosing K rather than K^* means that K is better than K^* , i.e.

$$\begin{aligned} \{ \hat{K}_{\text{l.s.}} = K \} \subset \left\{ 0 \geq \inf_{M \leq N} \{ \gamma_N(\hat{g}_{K, M}) + \text{pen}(N, M, K) \} \right. \\ \left. - \inf_{M \leq N} \{ \gamma_N(\hat{g}_{K^*, M}) + \text{pen}(N, M, K^*) \} \right\}. \end{aligned}$$

Let

$$\begin{aligned} D_{N, K} &:= \inf_{M \leq N} \{ \gamma_N(\hat{g}_{K, M}) + \text{pen}(N, M, K) \} \\ &\quad - \inf_{M \leq N} \{ \gamma_N(\hat{g}_{K^*, M}) + \text{pen}(N, M, K^*) \} \\ &= \gamma_N(\hat{g}_{K, \hat{M}_K}) + \text{pen}(N, \hat{M}_K, K) \\ &\quad - \inf_{M \leq N} \left\{ \inf_{t \in S_{K^*, M}} \gamma_N(t) + \text{pen}(N, M, K^*) \right\}. \end{aligned}$$

Then

$$\{ \hat{K}_{\text{l.s.}} = K \} \subset \{ D_{N, K} \leq 0 \}.$$

We will thus control the probability of the latter event for all $K < K^*$ in the first case and $K > K^*$ in the second case.

Proof of Theorem 7 Let $M_0 \in \mathcal{M}$. We will choose a suitable value for this integer later in the proof. Assume $N \geq M_0$. Then by definition of $D_{N,K}$ and of ν (equation (2)),

$$\begin{aligned} D_{N,K} &\geq \gamma_N(\hat{g}_{K,\hat{M}_K}) + \text{pen}(N, \hat{M}_K, K) - \gamma_N(g_{K^*,M_0}^*) - \text{pen}(N, M_0, K^*) \\ &\geq \|g^* - \hat{g}_{K,\hat{M}_K}\|_2^2 - \|g^* - g_{K^*,M_0}^*\|_2^2 - 2\nu(\hat{g}_{K,\hat{M}_K} - g_{K^*,M_0}^*) \\ &\quad + \text{pen}(N, \hat{M}_K, K) - \text{pen}(N, M_0, K^*). \end{aligned}$$

Using the definition of $Z_{K,M}$ (equation (3)), one gets that

$$\begin{aligned} |\nu(\hat{g}_{K,\hat{M}_K} - g_{K^*,M_0}^*)| &\leq |\nu(\hat{g}_{K,\hat{M}_K} - g^*)| + |\nu(g^* - g_{K^*,M_0}^*)| \\ &\leq Z_{K,\hat{M}_K}(g^*)(\|g^* - \hat{g}_{K,\hat{M}_K}\|_2^2 + x_{K,\hat{M}_K}^2) \\ &\quad + Z_{K^*,M_0}(g^*)(\|g^* - g_{K^*,M_0}^*\|_2^2 + x_{K^*,M_0}^2). \end{aligned}$$

Let N_0 , ρ and A be as in Lemma 14. We can assume that $N_0 \geq K^*$ so that $K^* \leq N$. Let us introduce the function $\widetilde{\text{pen}}(N, M, K) = \rho(MK + K^2 - 1)^{\frac{\log(N)}{N}}$. Let $N \geq N_0$ and $x > 0$ and assume we are in the event of probability $1 - e^{-x}$ of Lemma 14. Then, for all $K \leq N$:

$$\begin{aligned} |\nu(\hat{g}_{K,\hat{M}_K} - g_{K^*,M_0}^*)| &\leq \frac{1}{4}\|g^* - \hat{g}_{K,\hat{M}_K}\|_2^2 + \frac{1}{2}A\frac{x}{N} + \frac{1}{2}\widetilde{\text{pen}}(N, \hat{M}_K, K) \\ &\quad + \frac{1}{4}\|g^* - g_{K^*,M_0}^*\|_2^2 + \frac{1}{2}A\frac{x}{N} + \frac{1}{2}\widetilde{\text{pen}}(N, M_0, K^*) \end{aligned}$$

and

$$\begin{aligned} D_{N,K} &\geq \frac{1}{2}\|g^* - \hat{g}_{K,\hat{M}_K}\|_2^2 - \frac{3}{2}\|g^* - g_{K^*,M_0}^*\|_2^2 - 2A\frac{x}{N} + \text{pen}(N, \hat{M}_K, K) \\ &\quad - \text{pen}(N, M_0, K^*) - \widetilde{\text{pen}}(N, \hat{M}_K, K) - \widetilde{\text{pen}}(N, M_0, K^*). \end{aligned}$$

We assumed $\text{pen} \geq \widetilde{\text{pen}}$, so that

$$D_{N,K} \geq \frac{1}{2}\|g^* - \hat{g}_{K,\hat{M}_K}\|_2^2 - \frac{3}{2}\|g^* - g_{K^*,M_0}^*\|_2^2 - 2A\frac{x}{N} - 2\text{pen}(N, M_0, K^*)$$

Corollary 6 ensures that

$$d := \inf_{K < K^*} \inf_{t \in S_K} \|t - g^*\|_2 > 0,$$

so that for all $K < K^*$,

$$D_{N,K} \geq \frac{d^2}{2} - \frac{3}{2} \|g^* - g_{K^*,M_0}^*\|_2^2 - 2A \frac{x}{N} - 2\text{pen}(N, M_0, K^*).$$

By density of $(\mathfrak{P}_M)_{M \in \mathcal{M}}$ in $\mathcal{F}$, one gets that

$$\inf_M \|g_{K^*,M}^* - g^*\|_2 = 0$$

so that there exists M_0 such that $\|g^* - g_{K^*,M_0}^*\|_2^2 \leq d^2/6$. If we choose this M_0 , we get that

$$D_{N,K} \geq \frac{d^2}{4} - 2A \frac{x}{N} - 2\text{pen}(N, M_0, K^*).$$

Which implies that $D_{N,K} > 0$ as soon as $2Ax/N < d^2/4 - 2\text{pen}(N, M_0, K^*)$, i.e.

$$x < \left(\frac{d^2}{8} - \text{pen}(N, M_0, K^*) \right) \frac{N}{A}.$$

To conclude, note that there exists $\tilde{N}_0 \geq \max(N_0, M_0)$ such that for all $N \geq \tilde{N}_0$, $\text{pen}(N, M_0, K^*) \leq \frac{d^2}{16}$. Then, letting $\beta = \frac{d^2}{16A}$, one has for all $N \geq \tilde{N}_0$, with probability $1 - e^{-\beta N}$, for all $K < K^*$, $D_{N,K} > 0$, which implies that $\hat{K}_{\text{l.s.}} \neq K$.

Proof of Theorem 8 For all $K \geq K^*$,

$$D_{N,K} \geq \gamma_N(\hat{g}_{K,\hat{M}_K}) + \text{pen}(N, \hat{M}_K, K) - \gamma_N(g_{K^*,\hat{M}_K}^*) - \text{pen}(N, \hat{M}_K, K^*)$$

and

$$\gamma_N(\hat{g}_{K,\hat{M}_K}) - \gamma_N(g_{K^*,\hat{M}_K}^*) = \|\hat{g}_{K,\hat{M}_K} - g^*\|_2^2 - \|g_{K^*,\hat{M}_K}^* - g^*\|_2^2 - 2\nu(\hat{g}_{K,\hat{M}_K} - g_{K^*,\hat{M}_K}^*).$$

Note that $g_{K^*,\hat{M}_K}^* = g_{K,\hat{M}_K}^*$ is the orthogonal projection of g^* on $\mathfrak{P}_{\hat{M}_K}^{\otimes L}$ and $\hat{g}_{K,\hat{M}_K} \in S_{K,\hat{M}_K} \subset \mathfrak{P}_{\hat{M}_K}^{\otimes L}$, so that, using the Pythagorean Theorem,

$$\|\hat{g}_{K,\hat{M}_K} - g^*\|_2^2 - \|g_{K^*,\hat{M}_K}^* - g^*\|_2^2 = \|\hat{g}_{K,\hat{M}_K} - g_{K^*,\hat{M}_K}^*\|_2^2.$$

Let N_0, ρ and A be as in Lemma 14. We can assume that $N_0 \geq K^*$ so that $K^* \leq N$. Let us introduce the function $\widetilde{\text{pen}}(N, M, K) = \rho(MK + K^2 - 1) \frac{\log(N)}{N}$. Let $N \geq N_0$ and $x > 0$ and assume we are in the event of probability $1 - e^{-x}$ of Lemma 14. Then, for all $K \leq N$

such that $K \geq K^*$:

$$\begin{aligned}
|\nu(\hat{g}_{K, \hat{M}_K} - g_{K^*, \hat{M}_K}^*)| &= |\nu(\hat{g}_{K, \hat{M}_K} - g_{K, \hat{M}_K}^*)| \\
&\leq Z_{K, \hat{M}_K}((g_{K', M'}^*)_{K', M'}) \|\hat{g}_{K, \hat{M}_K} - g_{K, \hat{M}_K}^*\|_2^2 \\
&\quad + Z_{K, \hat{M}_K}((g_{K', M'}^*)_{K', M'}) x_{K, \hat{M}_K}^2 \\
&\leq \frac{1}{4} \|\hat{g}_{K, \hat{M}_K} - g_{K, \hat{M}_K}^*\|_2^2 + \frac{1}{2} A \frac{x}{N} + \frac{1}{2} \widetilde{\text{pen}}(N, \hat{M}_K, K),
\end{aligned}$$

which implies

$$\begin{aligned}
\gamma_N(\hat{g}_{K, \hat{M}_K}) - \gamma_N(g_{K^*, \hat{M}_K}^*) &\geq \frac{1}{2} \|\hat{g}_{K, \hat{M}_K} - g_{K, \hat{M}_K}^*\|_2^2 - A \frac{x}{N} - \widetilde{\text{pen}}(N, \hat{M}_K, K) \\
&\geq -A \frac{x}{N} - \widetilde{\text{pen}}(N, \hat{M}_K, K)
\end{aligned}$$

so that for all $K \leq N$ such that $K \geq K^*$:

$$D_{N, K} \geq \text{pen}(N, \hat{M}_K, K) - \text{pen}(N, \hat{M}_K, K^*) - \widetilde{\text{pen}}(N, \hat{M}_K, K) - A \frac{x}{N}.$$

Now, assume that **[Hpen]** (α, ρ) holds for some $\alpha > 0$ and the above constant ρ . Then there exists N_1 such that for all $N \geq N_1$ and for all $K \leq N$ such that $K \geq K^*$,

$$D_{N, K} \geq \alpha \frac{\log(N)}{N} - A \frac{x}{N},$$

which is strictly positive as soon as $x < \alpha \log(N)/A$. Thus, letting $\beta = 1/(2A)$, one has for all $N \geq \max(N_0, N_1, K^*)$, with probability $1 - N^{-\beta\alpha}$, for all $K \leq N$ such that $K > K^*$, $D_{N, K} > 0$, which implies that $\hat{K}_{\text{l.s.}} \neq K$. This concludes the proof.

6.3.2. Proof of Proposition 10

Let $n \geq 3$. Note $r = \frac{n}{n-1}$ and $K_0 = (n-1)^n$. One can check that $K = K_0 r^n \geq 2K_0$ and $K_0 r^k \in \mathbb{N}^*$ for all $k \in \{0, \dots, n\}$.

Denote by $N(K)$ the integer N_1 in the hypothesis **[Hpen]** $(0, \rho)$ corresponding to $K^* = K$. Then for all $N \geq \sup_{k \in \{0, \dots, n-1\}} N(K_0 r^k)$, for all M and for all $k \in \{1, \dots, n\}$,

$$\text{pen}(N, M, K_0 r^k) - \text{pen}(N, M, K_0 r^{k-1}) \geq \rho(M K_0 r^k + K_0^2 (r^2)^k - 1) \frac{\log(N)}{N}.$$

Taking the sum over $k \in \{1, \dots, n\}$, one gets that

$$\begin{aligned} \text{pen}(N, M, K) &\geq \rho \left(M \frac{r}{r-1} (K - K_0) + \frac{r^2}{r^2-1} (K^2 - K_0^2) - n \right) \frac{\log(N)}{N} \\ &\geq \rho \left(M \frac{r}{r-1} (K - K_0) + \frac{r^2}{r^2-1} (K^2 - 2K_0^2) \right) \frac{\log(N)}{N} \end{aligned}$$

since $n \leq K_0^2 = (n-1)^{2n}$. Using that $K \geq 2K_0$,

$$\text{pen}(N, M, K) \geq \frac{\rho}{2} \left(\frac{r}{r-1} MK + \frac{r^2}{r^2-1} K^2 \right) \frac{\log(N)}{N}.$$

Let $v_n = \frac{\rho}{2} \min \left(\frac{r}{r-1}, \frac{r^2}{r^2-1} \right)$. One gets

$$\text{pen}(N, M, K) \geq v_n (MK + K^2) \frac{\log(N)}{N}.$$

Therefore, there exists a non-decreasing sequence $(u_N)_{N \geq 1}$ such that

$$\begin{cases} \forall N, \forall M \leq N, \forall K \leq N, \text{pen}(N, M, K) \geq u_N (MK + K^2 - 1) \frac{\log(N)}{N} \\ \forall n, u_{\max(n, \sup_{k \in \{0, \dots, n-1\}} N(K_0 r^k))} \geq v_n \end{cases}.$$

and since $v_n \rightarrow \infty$, we get that $u_N \rightarrow \infty$, which concludes the proof.

We could for instance take

$$u_N = \max \left(0, \sup \left\{ v_i \mid i \leq N \text{ s.t. } \sup_{k \in \{0, \dots, i-1\}} N(i^k (i-1)^{i-k}) \leq N \right\} \right).$$

7. Acknowledgment

We would like to thank Elisabeth Gassiat for her precious advice and Yohann de Castro for his codes which were at the root of our numerical experiments.

A Spectral algorithm

Here is the spectral algorithm from de Castro et al. (2015a), de Castro et al. (2015b) that we use in our simulations for the estimation of the HMM parameters.

Data: An observed chain $(Y_1, \dots, Y_N)$ and an order K .

Result: Spectral estimators $\hat{\pi}$, $\hat{\mathbf{Q}}$ and the estimators $(\hat{f}_{M,k})_{k \in \mathcal{X}}$ of $(f_k^*)_{k \in \mathcal{X}}$ in $\mathfrak{P}_M$ (equipped with an orthonormal basis $\Phi_M = (\varphi_1, \dots, \varphi_M)$).

[Step 1] Consider the following empirical estimators: for any a, b, c in $\{1, \dots, M\}$,

- $\hat{\mathbf{L}}_M(a) := \frac{1}{N} \sum_{s=1}^N \varphi_a(Y_1^{(s)})$,
- $\hat{\mathbf{M}}_M(a, b, c) := \frac{1}{N} \sum_{s=1}^N \varphi_a(Y_1^{(s)}) \varphi_b(Y_2^{(s)}) \varphi_c(Y_3^{(s)})$,
- $\hat{\mathbf{N}}_M(a, b) := \frac{1}{N} \sum_{s=1}^N \varphi_a(Y_1^{(s)}) \varphi_b(Y_2^{(s)})$,
- $\hat{\mathbf{P}}_M(a, c) := \frac{1}{N} \sum_{s=1}^N \varphi_a(Y_1^{(s)}) \varphi_c(Y_3^{(s)})$.

[Step 2] Let $\hat{\mathbf{U}}$ be the $M \times K$ matrix of orthonormal right singular vectors of $\hat{\mathbf{P}}_M$ corresponding to its top K singular values.

[Step 3] Form the matrices $\hat{\mathbf{B}}(b) := (\hat{\mathbf{U}}^\top \hat{\mathbf{P}}_M \hat{\mathbf{U}})^{-1} \hat{\mathbf{U}}^\top \hat{\mathbf{M}}_M(\cdot, b, \cdot) \hat{\mathbf{U}}$ for all $b \in \{1, \dots, M\}$.

[Step 4] Set Θ a $(K \times K)$ uniformly drawn random unitary matrix and form the matrices $\hat{\mathbf{C}}(k) := \sum_{b=1}^M (\hat{\mathbf{U}} \Theta)(b, k) \hat{\mathbf{B}}(b)$ for all $k \in \{1, \dots, K\}$.

[Step 5] Compute $\hat{\mathbf{R}}$ a $(K \times K)$ unit Euclidean norm columns matrix that diagonalizes the matrix $\hat{\mathbf{C}}(1)$: $\hat{\mathbf{R}}^{-1} \hat{\mathbf{C}}(1) \hat{\mathbf{R}} = \text{Diag}(\hat{\Lambda}(1, 1), \dots, \hat{\Lambda}(1, K))$.

[Step 6] Set $\hat{\Lambda}(k, k') := (\hat{\mathbf{R}}^{-1} \hat{\mathbf{C}}(k) \hat{\mathbf{R}})(k', k')$ for all $k, k' \in \mathcal{X}$ and $\hat{\mathbf{O}}_M := \hat{\mathbf{U}} \Theta \hat{\Lambda}$.

[Step 7] Consider the emission laws estimator $\hat{\mathbf{f}} := (\hat{f}_{M,k})_{k \in \mathcal{X}}$ defined by $\hat{f}_{M,k} := \sum_{m=1}^M \hat{\mathbf{O}}_M(m, k) \varphi_m$ for all $k \in \mathcal{X}$.

[Step 8] Set $\tilde{\pi} := \Pi_{\Delta_K} \left((\hat{\mathbf{U}}^\top \hat{\mathbf{O}}_M)^{-1} \hat{\mathbf{U}}^\top \hat{\mathbf{L}}_M \right)$ where Π_{Δ_K} denotes the projection onto the simplex in dimension K .

[Step 9] Consider the transition matrix estimator:

$$\hat{\mathbf{Q}} := \Pi_{\text{TM}} \left((\hat{\mathbf{U}}^\top \hat{\mathbf{O}}_M \text{Diag} \tilde{\pi})^{-1} \hat{\mathbf{U}}^\top \hat{\mathbf{N}}_M \hat{\mathbf{U}} (\hat{\mathbf{O}}_M^\top \hat{\mathbf{U}})^{-1} \right),$$

where Π_{TM} denotes the projection onto the convex set of transition matrices, and define $\hat{\pi}$ as the stationary distribution of $\hat{\mathbf{Q}}$.

B Proofs of the oracle inequalities

B1. Proof of Theorem 11

Let $K \leq N$ and $M \leq N$. Then

$$\begin{aligned}\gamma_N(\hat{g}) + \text{pen}(N, \hat{M}, \hat{K}_{\text{l.s.}}) &\leq \gamma_N(\hat{g}_{K,M}) + \text{pen}(N, M, K) \\ &\leq \gamma_N(g_{K,M}^*) + \text{pen}(N, M, K)\end{aligned}$$

where the first inequality comes from the definition of $(\hat{K}_{\text{l.s.}}, \hat{M})$ and the second from the definition of $\hat{g}_{K,M}$. Therefore,

$$\gamma_N(\hat{g}) - \gamma_N(g_{K,M}^*) \leq \text{pen}(N, M, K) - \text{pen}(N, \hat{M}, \hat{K}_{\text{l.s.}}).$$

By definition of ν (equation 2),

$$\gamma_N(t_1) - \gamma_N(t_2) = \|t_1 - g^*\|_2^2 - \|t_2 - g^*\|_2^2 - 2\nu(t_1 - t_2)$$

so that

$$\|\hat{g} - g^*\|_2^2 \leq \|g_{K,M}^* - g^*\|_2^2 + \text{pen}(N, M, K) - \text{pen}(N, \hat{M}, \hat{K}_{\text{l.s.}}) + 2\nu(\hat{g}_{\hat{M}, \hat{K}_{\text{l.s.}}} - g_{K,M}^*)$$

Now we want to control the ν term. By linearity,

$$\nu(\hat{g}_{\hat{K}_{\text{l.s.}}, \hat{M}} - g_{K,M}^*) = \nu(\hat{g}_{\hat{K}_{\text{l.s.}}, \hat{M}} - g^*) + \nu(g^* - g_{K,M}^*)$$

Using the definition of $Z_{K,M}$ (equation 3), we get that

$$\begin{cases} |\nu(\hat{g}_{\hat{K}_{\text{l.s.}}, \hat{M}} - g^*)| \leq Z_{\hat{K}_{\text{l.s.}}, \hat{M}}(g^*)(\|\hat{g}_{\hat{K}_{\text{l.s.}}, \hat{M}} - g^*\|_2^2 + x_{\hat{K}_{\text{l.s.}}, \hat{M}}^2) \\ |\nu(g_{K,M}^* - g^*)| \leq Z_{K,M}(g^*)(\|g_{K,M}^* - g^*\|_2^2 + x_{K,M}^2) \end{cases}$$

so that, using Lemma 14, for all $N \geq N_0$ and $x > 0$, with probability larger than $1 - e^{-x}$, for all $M \leq N$ and $K \leq N$,

$$\begin{aligned} |\nu(\hat{g}_{\hat{K}_{\text{l.s.}}, \hat{M}} - g_{K, M}^*)| &\leq \frac{1}{4} \|\hat{g} - g^*\|_2^2 + \frac{1}{4} \|g_{K, M}^* - g^*\|_2^2 + A \frac{x}{N} \\ &\quad + \frac{1}{2} \text{pen}(N, \hat{M}, \hat{K}_{\text{l.s.}}) + \frac{1}{2} \text{pen}(N, M, K) \end{aligned}$$

so that

$$\begin{aligned} \|\hat{g} - g^*\|_2^2 &\leq \|g_{K, M}^* - g^*\|_2^2 + 2 \text{pen}(N, M, K) \\ &\quad + \frac{1}{2} \|\hat{g} - g^*\|_2^2 + \frac{1}{2} \|g_{K, M}^* - g^*\|_2^2 + 2A \frac{x}{N}, \end{aligned}$$

which means that

$$\frac{1}{2} \|\hat{g} - g^*\|_2^2 \leq \frac{3}{2} \|g_{K, M}^* - g^*\|_2^2 + 2 \text{pen}(N, M, K) + 2A \frac{x}{N}$$

and finally

$$\|\hat{g} - g^*\|_2^2 \leq 4 \inf_{K \leq N, M \leq N} \{\|g_{K, M}^* - g^*\|_2^2 + \text{pen}(N, M, K)\} + 4A \frac{x}{N}$$

which is the expected inequality.

B2. Proof of Lemma 12

First, decompose the difference in three terms.

$$\begin{aligned} \|g^{\pi_1, \mathbf{Q}_1, \mathbf{f}_1} - g^{\pi_2, \mathbf{Q}_2, \mathbf{f}_2}\|_2 &\leq \|g^{\pi_1, \mathbf{Q}_1, \mathbf{f}_1} - g^{\pi_2, \mathbf{Q}_1, \mathbf{f}_1}\|_2 + \|g^{\pi_2, \mathbf{Q}_1, \mathbf{f}_1} - g^{\pi_2, \mathbf{Q}_2, \mathbf{f}_1}\|_2 \\ &\quad + \|g^{\pi_2, \mathbf{Q}_2, \mathbf{f}_1} - g^{\pi_2, \mathbf{Q}_2, \mathbf{f}_2}\|_2 \end{aligned}$$

Then we can control each term separately. Let $(\varphi_m)_{m \in \mathbb{N}^*}$ be an orthonormal basis of $\cup_M \mathfrak{P}_M$.

$$\begin{aligned}
\|g^{\pi_1, \mathbf{Q}, \mathbf{f}} - g^{\pi_2, \mathbf{Q}, \mathbf{f}}\|_2^2 &= \left\| \sum_{\mathbf{k} \in \mathcal{X}^L} (\pi_1 - \pi_2)_{k_1} \mathbf{Q}_{k_1, k_2} \cdots \mathbf{Q}_{k_{L-1}, k_L} \bigotimes_{i=1}^L f_{k_i} \right\|_2^2 \\
&= \sum_{\mathbf{m} \in (\mathbb{N}^*)^L} \left(\sum_{\mathbf{k} \in \mathcal{X}^L} (\pi_1 - \pi_2)_{k_1} \mathbf{Q}_{k_1, k_2} \cdots \mathbf{Q}_{k_{L-1}, k_L} \prod_{i=1}^L \langle f_{k_i}, \varphi_{m_i} \rangle \right)^2 \\
&\leq \sum_{\mathbf{k} \in \mathcal{X}^L} (\pi_1 - \pi_2)_{k_1}^2 \mathbf{Q}_{k_1, k_2} \cdots \mathbf{Q}_{k_{L-1}, k_L} \\
&\quad \times \sum_{\mathbf{k}' \in \mathcal{X}^L} \mathbf{Q}_{k'_1, k'_2} \cdots \mathbf{Q}_{k'_{L-1}, k'_L} \prod_{i=1}^L \sum_{m_i \in \mathbb{N}^*} \langle f_{k'_i}, \varphi_{m_i} \rangle^2
\end{aligned}$$

using Cauchy-Schwarz inequality. Then, since $\sum_{m_i \in \mathbb{N}^*} \langle f_{k'_i}, \varphi_{m_i} \rangle^2 = \|f_{k'_i}\|_2^2 \leq C_{\mathcal{F}, 2}^2$ by **[HF]** and $\mathbf{Q}$ is a transition matrix, we get that

$$\|g^{\pi_1, \mathbf{Q}, \mathbf{f}} - g^{\pi_2, \mathbf{Q}, \mathbf{f}}\|_2^2 \leq K C_{\mathcal{F}, 2}^{2L} \|\pi_1 - \pi_2\|^2$$

A similar decomposition leads to

$$\|g^{\pi, \mathbf{Q}_1, \mathbf{f}} - g^{\pi, \mathbf{Q}_2, \mathbf{f}}\|_2^2 \leq (L-1) K C_{\mathcal{F}, 2}^{2L} \|\mathbf{Q}_1 - \mathbf{Q}_2\|_F^2$$

and

$$\|g^{\pi, \mathbf{Q}, \mathbf{f}_1} - g^{\pi, \mathbf{Q}, \mathbf{f}_2}\|_2^2 \leq L K C_{\mathcal{F}, 2}^{2(L-1)} \sum_{k \in \mathcal{X}} \|(\mathbf{f}_1)_k - (\mathbf{f}_2)_k\|_2^2$$

These inequalities remain true if the states of the second set of parameters are swapped. Then, we use that $C_{\mathcal{F}, 2} \geq 1$ by **[HF]** to conclude.

C Proof of the control of $Z_{K, M}$

This section contains the proof of Lemma 14.

C1. Concentration inequality on $Z_{K,M}(s)$

Define for all $\sigma > 0$ the sets

$$B_\sigma = \{t \in S_{K,M}, C_{\mathcal{F},\infty}^{L/2} \|t - s_{K,M}\|_2 \leq \sigma\}$$

Let d_{g^*} be the semi-distance defined by $d_{g^*}^2(t_1, t_2) = \mathbb{E}[(t_1 - t_2)^2(Z_1)] = \int g^*(t_1 - t_2)^2 d\mu^{\otimes L}$, and d_2 the distance induced by the norm on $\mathbf{L}^2(\mathcal{Y}^L, \mu^{\otimes L})$.

Let $N(\epsilon, A, d) = e^{H(\epsilon, A, d)}$ denote the minimal cardinality of a covering of A by brackets of size ϵ for the semi-distance d , that is by sets $[t_1, t_2] = \{t : \mathcal{Y}^L \mapsto \mathbb{R}, t_1(\cdot) \leq t(\cdot) \leq t_2(\cdot)\}$ such that $d(t_1, t_2) \leq \epsilon$. $H(\cdot, A, d)$ is called the *bracketing entropy* of A for the semi-distance d .

The following lemma is a Bernstein-like inequality that follows from Paulin (2013), Theorem 2.4:

Lemma 16. *Let t be a real valued and measurable bounded function on $\mathcal{Y}^L$. Let $V = \mathbb{E}[t^2(Z_1)]$. There exists a positive constant c^* depending only on $\mathbf{Q}^*$ and L such that for all $0 \leq \lambda \leq 1/(2\sqrt{2}c^*\|t\|_\infty)$ and for all $N \in \mathbb{N}$:*

$$\log \mathbb{E} \exp \left[\lambda \sum_{s=1}^N (t(Z_s) - \mathbb{E}t(Z_s)) \right] \leq \frac{2Nc^*V\lambda^2}{1 - 2\sqrt{2}c^*\|t\|_\infty\lambda}$$

The following lemma is an extension of Theorem 6.8 from Massart (2007) and allows to obtain a concentration inequality on the supremum on all functions of a class when one can control its bracketing entropy.

Lemma 17. *Let Ξ be some measurable space, $(\xi_i)_{1 \leq i \leq N}$ a sequence of random variables on Ξ , $\mathcal{T}$ some countable class of real valued and measurable functions on Ξ . Assume that there exists some positive numbers a and b such that for all $t \in \mathcal{T}$, $\|t\|_\infty \leq b$ and $\sup_i \mathbb{E}[t^2(\xi_i)] \leq a^2$.*

Assume also that there exists some constant $C_\xi \geq 1/4$ such that for all $0 \leq \lambda \leq 1/(2\sqrt{2}C_\xi b)$ and for all $t \in \mathcal{T}$:

$$\log \mathbb{E} \exp \left[\lambda \sum_{s=1}^N (t(\xi_s) - \mathbb{E}t(\xi_s)) \right] \leq \frac{2NC_\xi a^2 \lambda^2}{1 - 2\sqrt{2}C_\xi b \lambda} \quad (4)$$

Assume furthermore that for any positive number δ , there exists some finite set $\mathcal{B}_\delta$ of brackets covering $\mathcal{T}$ such that for any bracket $[t_1, t_2] \in \mathcal{B}_\delta$, $\|t_1 - t_2\|_\infty \leq b$ and $\sup_i \mathbb{E}[(t_1 -$

$t_2)^2(\xi_i)] \leq \delta^2$. Let $e^{H(\delta)}$ denote the minimal cardinality of such a covering. Then, there exists a numerical constant $\kappa > 0$ such that for any measurable set A such that $\mathbb{P}(A) > 0$,

$$\mathbb{E}^A \left(\sup_{t \in \mathcal{T}} \sum_{s=1}^N (t(\xi_s) - \mathbb{E}t(\xi_s)) \right) \leq \kappa C_\xi \left[E + a \sqrt{N \log \left(\frac{1}{\mathbb{P}(A)} \right)} + b \log \left(\frac{1}{\mathbb{P}(A)} \right) \right]$$

where

$$E = \sqrt{N} \int_0^a \sqrt{H(u) \wedge N} du + (b + a)H(a)$$

and for any measurable random variable W , $\mathbb{E}^A[W] = \mathbb{E}[W \mathbf{1}_A] / \mathbb{P}(A)$.

Comment. The assumption $C_\xi \geq 1/4$ is only used to factorise the upper bound by C_ξ . Without it, the upper bound would be

$$\kappa' \left[E + a \sqrt{C_\xi N \log \left(\frac{1}{\mathbb{P}(A)} \right)} + b C_\xi \log \left(\frac{1}{\mathbb{P}(A)} \right) \right]$$

In practice, this assumption doesn't cost anything: if equation (4) holds for some constant C_ξ , then it holds for any constant $C' \geq C_\xi$.

We will apply this lemma to $\Xi = \mathcal{Y}^L$ and $\xi_i = Z_i$. Using Lemma 16, equation (4) holds with $C_\xi = \max(c^*, 1/4)$.

Take $\mathcal{T} = (B_\sigma - s_{K,M}) \cup (-B_\sigma + s_{K,M})$, so that $\sup_{t \in \mathcal{T}} \nu(t) = \sup_{t \in B_\sigma} |\nu(t - s_{K,M})|$. Then, one can check using Lemma 24 that the assumptions $\|t\|_\infty \leq b$ and $\sup_i \mathbb{E}[t^2(\xi_i)] \leq a^2$ hold with

- $b = 2C_{\mathcal{F},\infty}^L$
- $a = 2 \min(\sigma, C_{\mathcal{F},\infty}^{L/2} C_{\mathcal{F},2}^L)$

and $H(u) \leq \log(2) + H(u, B_\sigma - s_{K,M}, d_{g^*})$.

We can do without assuming $\mathcal{T}$ to be countable. Indeed, ν is continuous on $\mathcal{T}$ equipped with the infinity norm. This entails that the supremum of ν over $\mathcal{T}$ is equal to the supremum of ν over any dense subset of $(\mathcal{T}, \|\cdot\|_\infty)$. Since $\mathcal{T} \subset (\mathfrak{P}_M)^{\otimes L}$, which is a finite dimensional metric linear space for the infinity norm, it is separable. Therefore, without loss of generality, we can get rid of the countability assumption on $\mathcal{T}$.

Rewriting these results, we get the following lemma:

Lemma 18. *There exists a constant C^* depending only on Q^* and L such that for all $\sigma > 0$, for all measurable A such that $\mathbb{P}(A) > 0$:*

$$\mathbb{E}^A \left(\sup_{t \in B_\sigma} |\nu(t - s_{K,M})| \right) \leq C^* \left[\frac{E}{N} + \sigma \sqrt{\frac{1}{N} \log \left(\frac{1}{\mathbb{P}(A)} \right)} + \frac{2C_{\mathcal{F},\infty}^L}{N} \log \left(\frac{1}{\mathbb{P}(A)} \right) \right]$$

where

$$\begin{aligned} E = & \sqrt{N} \int_0^\sigma \sqrt{H(u, B_\sigma - s_{K,M}, d_{g^*}) \wedge N} du + \log(2) \sigma \sqrt{N} \\ & + 2(C_{\mathcal{F},\infty}^L + C_{\mathcal{F},\infty}^{L/2} C_{\mathcal{F},2}^L) H(\sigma, B_\sigma - s_{K,M}, d_{g^*}) \end{aligned}$$

The core of the proof consists in controlling the bracketing entropy in order to find a "good" function φ and constants C and $\sigma_{K,M}$ depending on $C_{\mathcal{F},2}$, $C_{\mathcal{F},\infty}$ and L such that $x \mapsto \frac{\varphi(x)}{x}$ is nonincreasing and

$$\forall \sigma \geq \sigma_{K,M} \quad E \leq C \varphi(\sigma) \sqrt{N}. \quad (5)$$

For ease of notation, we did not write the dependency of C and φ on K and M .

Let us see how to conclude with such an inequality. We shall use the following result (lemma 4.23 from Massart (2007)).

Lemma 19. *Let S be some countable set, $u \in S$ and $a : S \mapsto \mathbb{R}_+$ such that $a(u) = \inf_{t \in S} a(t)$. Let Z be some process indexed by S and assume that $\sup_{t \in \mathbf{B}(\lambda)} Z(t) - Z(u)$ has finite expectation for any positive number $\lambda \geq 0$, where*

$$\mathbf{B}(\lambda) = \{t \in S, a(t) \leq \lambda\}$$

Then, for any function ϕ on $\mathbb{R}_+$ such that $x \mapsto \phi(x)/x$ is nonincreasing on $\mathbb{R}_+$ and satisfies for some $\lambda_ \geq 0$ to*

$$\forall \lambda \geq \lambda_* \geq 0, \quad \mathbb{E} \left[\sup_{t \in \mathbf{B}(\lambda)} Z(t) - Z(u) \right] \leq \phi(\lambda)$$

one has for any $x \geq \lambda_$:*

$$\mathbb{E} \left[\sup_{t \in S} \left(\frac{Z(t) - Z(u)}{a(t)^2 + x^2} \right) \right] \leq 4 \frac{\phi(x)}{x^2}$$

In our case,

$$\begin{cases} S = S_{K,M} - s_{K,M} \\ u = s_{K,M} \\ a(t) = C_{\mathcal{F},\infty}^{L/2} \|t - s_{K,M}\|_2 \\ Z(t) = |\nu(t - s_{K,M})| \\ \lambda_* = \sigma_{K,M} \\ \phi(x) = C^* \left[C \frac{\varphi(x)}{\sqrt{N}} + x \sqrt{\frac{1}{N} \log \left(\frac{1}{\mathbb{P}(A)} \right)} + \frac{2C_{\mathcal{F},\infty}^L}{N} \log \left(\frac{1}{\mathbb{P}(A)} \right) \right] \end{cases}$$

With this choice of S , a and Z , this proposition holds even if S is not countable for the same reason as in Lemma 18.

It follows that for all $x \geq \sigma_{K,M}$:

$$\mathbb{E}^A \left[\sup_{t \in S_{K,M}} \frac{|\nu(t - s_{K,M})|}{C_{\mathcal{F},\infty}^L \|t - s_{K,M}\|_2^2 + x^2} \right] \leq 4 \frac{\phi(x)}{x^2},$$

so that if $x_{K,M} \geq \frac{\sigma_{K,M}}{C_{\mathcal{F},\infty}^{L/2}}$:

$$\mathbb{E}^A[Z_{K,M}(s)] \leq 4 \frac{\phi(x_{K,M} C_{\mathcal{F},\infty}^{L/2})}{x_{K,M}^2}$$

and then

$$\begin{aligned} \mathbb{E}^A[Z_{K,M}(s)] &\leq 4 \frac{C^*}{x_{K,M}^2} \left[C \frac{\varphi(x_{K,M} C_{\mathcal{F},\infty}^{L/2})}{\sqrt{N}} + x_{K,M} C_{\mathcal{F},\infty}^{L/2} \sqrt{\frac{1}{N} \log \left(\frac{1}{\mathbb{P}(A)} \right)} \right. \\ &\quad \left. + \frac{2C_{\mathcal{F},\infty}^L}{N} \log \left(\frac{1}{\mathbb{P}(A)} \right) \right] \\ &=: \psi \left(\log \left(\frac{1}{\mathbb{P}(A)} \right) \right) \end{aligned}$$

Note that the function ψ is nondecreasing. On the event $A = \{Z_{K,M}(s) \geq \psi(x)\}$,

$$\psi(x) \leq \mathbb{E}^A[Z_{K,M}(s)] \leq \psi \left(\log \left(\frac{1}{\mathbb{P}(A)} \right) \right)$$

so that $x \leq \log \left(\frac{1}{\mathbb{P}(A)} \right)$ and finally $\mathbb{P}(A) \leq e^{-x}$.

It follows that with probability $1 - e^{-z_{K,M}-z}$:

$$Z_{K,M}(s) \leq 4C^* \left[C \frac{\varphi(x_{K,M} C_{\mathcal{F},\infty}^{L/2})}{x_{K,M}^2 \sqrt{N}} + C_{\mathcal{F},\infty}^{L/2} \sqrt{\frac{z_{K,M} + z}{x_{K,M}^2 N}} + 2C_{\mathcal{F},\infty}^L \frac{z_{K,M} + z}{x_{K,M}^2 N} \right] \quad (6)$$

and the last step of the proof will be to choose the right $x_{K,M}$ and $z_{K,M}$ (see section C3.).

C2. Control of the bracketing entropy

The goal of this section is to prove equation (5), that is to find φ , C and $\sigma_{K,M}$ such that

$$\forall \sigma \geq \sigma_{K,M} \quad E \leq C\varphi(\sigma)\sqrt{N}.$$

The bracketing entropy is invariant under translation and increasing with respect to the inclusion relation, so

$$H(u, B_\sigma - s_{K,M}, d_{g^*}) = H(u, B_\sigma, d_{g^*}) \leq H(u, S_{K,M}, d_{g^*})$$

Using Lemma 24, we get that for all $t \in \mathbf{L}^2(\mathcal{Y}^L, \mathbb{R})$, $\int t^2 g^* d\mu \leq C_{\mathcal{F},\infty}^L \|t\|_2^2$. Therefore, a bracket of size $u/C_{\mathcal{F},\infty}^{L/2}$ for d_2 is also a bracket of size u for d_{g^*} , which implies that

$$H(u, B_\sigma - s_{K,M}, d_{g^*}) \leq H\left(\frac{u}{C_{\mathcal{F},\infty}^{L/2}}, S_{K,M}, d_2\right) \quad (7)$$

Let us now rewrite the definition of $S_{K,M}$:

$$\begin{aligned} S_{K,M} &= \left\{ \sum_{\mathbf{k} \in \{1, \dots, K\}^L} \pi_{k_1} \prod_{i=2}^L \mathbf{Q}_{k_{i-1}, k_i} \bigotimes_{i=1}^L f_{k_i}, \mathbf{Q} \in \mathcal{Q}_K, \pi \mathbf{Q} = \pi, \mathbf{f} \in (\mathcal{F} \cap \mathfrak{P}_M)^K \right\} \\ &\subset \left\{ \sum_{\mathbf{k} \in \{1, \dots, K\}^L} \mu_{\mathbf{k}} \phi_{\mathbf{k}}, \mu \in \mathcal{U}, \phi \in \Phi \right\} \end{aligned}$$

where

$$\left\{ \begin{array}{l} \mathcal{U} = \left\{ \left(\pi_{k_1} \prod_{i=2}^L \mathbf{Q}_{k_{i-1}, k_i} \right)_{k_1, \dots, k_L}, \mathbf{Q} \text{ transition matrix } K \times K, \pi \geq 0, \pi \in \mathbb{S}_{K-1} \right\} \\ \Phi = \left\{ \left(\bigotimes_{i=1}^L f_{k_i} \right)_{k_1, \dots, k_L}, \mathbf{f} \in (\mathcal{F} \cap \mathfrak{P}_M)^K \right\} \end{array} \right\}$$

$\mathcal{U}$ is equipped with the distance $d_2(a, b) = (\sum_{\mathbf{k}} (b_{\mathbf{k}}^i - a_{\mathbf{k}}^i)^2)^{1/2}$. A bracket for $\mathcal{U}$ will be a set $[a, b] = \{c \mid \forall \mathbf{k} \in \{1, \dots, K\}^L, a_{\mathbf{k}} \leq c_{\mathbf{k}} \leq b_{\mathbf{k}}\}$.

Φ is equipped with the distance $d_{\infty, 2}(u, v) = \max_{\mathbf{k}} \|v_{\mathbf{k}}^i - u_{\mathbf{k}}^i\|_2$. A bracket Φ will be a set $[u, v] = \{t \mid \forall \mathbf{k} \in \{1, \dots, K\}^L, u_{\mathbf{k}}(\cdot) \leq t_{\mathbf{k}}(\cdot) \leq v_{\mathbf{k}}(\cdot)\}$.

Controlling the bracketing entropy on each of these sets will allow to control the bracketing entropy of $S_{K, M}$. Let us start with them:

Lemma 20. *There exists a bracket covering $\{[a^i, b^i]\}_{1 \leq i \leq N_{\mathcal{U}}(\epsilon)}$ of size ϵ of $\mathcal{U}$ for the distance d_2 with cardinality*

$$N_{\mathcal{U}}(\epsilon) \leq \max \left(\frac{2LK^{L/2}}{\epsilon}, 1 \right)^{K^2-1} \quad (8)$$

such that for all i and $\mathbf{k}$, $0 \leq a_{\mathbf{k}}^i \leq b_{\mathbf{k}}^i \leq 1$.

Lemma 21. *There exists a bracket covering $\{[u^i, v^i]\}_{1 \leq i \leq N_{\Phi}(\epsilon)}$ of size ϵ of Φ for the distance $d_{\infty, 2}$ of cardinality*

$$N_{\Phi}(\epsilon) \leq \max \left(\frac{L(4C_{\mathcal{F}, 2}M)^L}{\epsilon}, 1 \right)^{MK} \quad (9)$$

such that

$$\max_i \max_{\mathbf{k}} \|v_{\mathbf{k}}^i\|_2^2 \leq (8M^2 C_{\mathcal{F}, 2}^2)^L.$$

Let us take such bracketings and consider the following set of brackets:

$$\left\{ \left[\sum_{\mathbf{k}} \mathcal{A}_{\mathbf{k}}^{i,j}, \sum_{\mathbf{k}} \mathcal{B}_{\mathbf{k}}^{i,j} \right] \right\}_{1 \leq i \leq N_{\mathcal{U}}(\epsilon), 1 \leq j \leq N_{\Phi}(\epsilon)}$$

where

$$\forall y \in \mathcal{Y}^L, \quad \begin{cases} \mathcal{A}_{\mathbf{k}}^{i,j}(y) = \min\{a_{\mathbf{k}}^i u_{\mathbf{k}}^j(y), b_{\mathbf{k}}^i u_{\mathbf{k}(y)}^j\} \\ \mathcal{B}_{\mathbf{k}}^{i,j}(y) = \max\{a_{\mathbf{k}}^i v_{\mathbf{k}}^j(y), b_{\mathbf{k}}^i v_{\mathbf{k}(y)}^j\} \end{cases}.$$

This set covers $S_{K, M}$: for all $\mu \in \mathcal{U}$, $\phi \in \Phi$, there exists $i \in \{1, \dots, N_{\mathcal{U}}(\epsilon)\}$ and $j \in \{1, \dots, N_{\Phi}(\epsilon)\}$ such that $\mu \in [a^i, b^i]$ and $\phi \in [u^j, v^j]$, and then by construction $\sum_{\mathbf{k}} \mu_{\mathbf{k}} \phi_{\mathbf{k}} \in$

$$[\sum_{\mathbf{k}} \mathcal{A}_{\mathbf{k}}^{i,j}, \sum_{\mathbf{k}} \mathcal{B}_{\mathbf{k}}^{i,j}].$$

Let us now bound the size of these brackets. Let $[a, b] \in \{[a^i, b^i]\}_{1 \leq i \leq N_{\mathcal{U}}(\epsilon)}$ and $[u, v] \in \{[u^i, v^i]\}_{1 \leq i \leq N_{\Phi}(\epsilon)}$, then if one denotes by $[\mathcal{A}, \mathcal{B}]$ the corresponding bracket,

$$\begin{aligned} \left\| \sum_{\mathbf{k}} \mathcal{A}_{\mathbf{k}} - \sum_{\mathbf{k}} \mathcal{B}_{\mathbf{k}} \right\|_2^2 &\leq K^L \sum_{\mathbf{k}} \|\mathcal{A}_{\mathbf{k}} - \mathcal{B}_{\mathbf{k}}\|_2^2 \\ &\leq 2K^L \sum_{\mathbf{k}} \left((a_{\mathbf{k}} - b_{\mathbf{k}})^2 \|v_{\mathbf{k}}\|_2^2 + (a_{\mathbf{k}}^2 \vee b_{\mathbf{k}}^2) \|u_{\mathbf{k}} - v_{\mathbf{k}}\|_2^2 \right). \end{aligned}$$

Then, by definition of the brackets, $\|u_{\mathbf{k}} - v_{\mathbf{k}}\|_2^2 \leq \epsilon^2$ and $\sum_{\mathbf{k}} (a_{\mathbf{k}} - b_{\mathbf{k}})^2 \leq \epsilon^2$. In addition, we assumed $|a_{\mathbf{k}}| \leq 1, |b_{\mathbf{k}}| \leq 1$ and $\|v_{\mathbf{k}}\|_2^2 \leq (8M^2 C_{\mathcal{F},2}^2)^L$ for all $\mathbf{k}$, so that

$$\left\| \sum_{\mathbf{k}} \mathcal{A}_{\mathbf{k}} - \sum_{\mathbf{k}} \mathcal{B}_{\mathbf{k}} \right\|_2^2 \leq 2K^L \epsilon^2 ((8M^2 C_{\mathcal{F},2}^2)^L + K^L),$$

which implies

$$\begin{aligned} N(\epsilon, S_{K,M}, d_2) &\leq N_{\mathcal{U}} \left(\frac{\epsilon}{\sqrt{2K^L(K^L + (8M^2 C_{\mathcal{F},2}^2)^L)}} \right) \\ &\quad \times N_{\Phi} \left(\frac{\epsilon}{\sqrt{2K^L(K^L + (8M^2 C_{\mathcal{F},2}^2)^L)}} \right), \quad (10) \end{aligned}$$

and finally by combining 7, 8, 9 and 10:

$$\begin{aligned}
H(u, B_\sigma - s_{K,M}, d_{g^*}) &\leq (K^2 - 1) \log \max \left(\frac{C_{\mathcal{F},\infty}^{L/2} \sqrt{2K^L(K^L + (8M^2C_{\mathcal{F},2}^2)^L)} 2LK^{L/2}}{u}, 1 \right) \\
&\quad + MK \log \max \left(\frac{C_{\mathcal{F},\infty}^{L/2} \sqrt{2K^L(K^L + (8M^2C_{\mathcal{F},2}^2)^L)} L(4C_{\mathcal{F},2}M)^L}{u}, 1 \right) \\
&\leq (MK + K^2 - 1) \log \max \left(\frac{C_{\mathcal{F},\infty}^{L/2} \sqrt{2(K^L + (8M^2C_{\mathcal{F},2}^2)^L)} LK^L(4C_{\mathcal{F},2}M)^L}{u}, 1 \right) \\
&\leq (MK + K^2 - 1) \log \max \left(\frac{C_{\mathcal{F},\infty}^{L/2} \sqrt{2(N^L + (8C_{\mathcal{F},2}^2)^L N^{2L})} NN^L(4C_{\mathcal{F},2})^L N^L}{u}, 1 \right) \\
&\leq (MK + K^2 - 1) \log \max \left(\frac{C_{\mathcal{F},\infty}^{L/2} 2(8C_{\mathcal{F},2}^2)^{L/2} N^L NN^L(4C_{\mathcal{F},2})^L N^L}{u}, 1 \right) \\
&\leq (MK + K^2 - 1) \log \max \left(\frac{2(16C_{\mathcal{F},\infty}^{1/2} C_{\mathcal{F},2}^2)^L N^{3L+1}}{u}, 1 \right) \\
&\leq (MK + K^2 - 1) \log \max \left(\frac{N^{6L}}{u}, 1 \right)
\end{aligned}$$

for N large enough ($N \geq N_0 := 16C_{\mathcal{F},\infty}^{1/2} C_{\mathcal{F},2}^2$) because we assumed $M \leq N$, $K \leq N$, $L \leq N$, $C_{\mathcal{F},2} \geq 1$ and $C_{\mathcal{F},\infty} \geq 1$. Thus, Lemma 25 implies that if we write $C_0 = \sqrt{\pi}$ and

$$\varphi(\sigma) = C_0 \sigma \sqrt{MK + K^2 - 1} \left(1 + \sqrt{\log \left(\max \left\{ \frac{N^{6L}}{\sigma}, 1 \right\} \right)} \right)$$

then for all $N \geq N_0$ and $\sigma > 0$,

$$\begin{cases} \sigma^2 H(\sigma, S_{K,M}, d_2) \leq \varphi(\sigma)^2 \\ \int_0^\sigma \sqrt{H(u, S_{K,M}, d_2)} du \leq \varphi(\sigma) \\ \log(2)\sigma \leq \varphi(\sigma) \end{cases}$$

Let us now check that this function φ satisfies equation (5). First, note that $x \mapsto \frac{\varphi(x)}{x}$ is nonincreasing, so that $x \mapsto \frac{\varphi(x)}{x^2}$ is also nonincreasing. Thus, we may define $\sigma_{K,M}$ as the

unique solution of the equation $\varphi(x) = \sqrt{N}x^2$, and then for all $\sigma \geq \sigma_{K,M}$:

$$H(\sigma, B_\sigma - s_{K,M}, d_{g^*}) \leq \frac{\varphi(\sigma)^2}{\sigma^2} \leq \frac{\varphi(\sigma)}{\sigma} \sigma \sqrt{N} = \varphi(\sigma) \sqrt{N}$$

Equation (5) follows immediately with $C = 2(1 + C_{\mathcal{F},\infty}^L + C_{\mathcal{F},\infty}^{L/2} C_{\mathcal{F},2}^L)$.

Proof of Lemma 20 Let $\epsilon \in (0, 2)$.

We start with the family $\{[k/n, (k+1)/n], k \in \{0, \dots, n-1\}\}$ with n an integer between $1/\epsilon$ and $2/\epsilon$, which gives a bracket covering of size ϵ of $[0, 1]$ with cardinality smaller than $2/\epsilon$. These brackets will be used to control each free component of $\mathbf{Q}$ and π , that is $K^2 - 1$ components.

More precisely, we define the following bracket set:

$$\begin{aligned} & \left\{ [A, B] \mid A_{\mathbf{k}} = \frac{1}{n^L} p_{k_1} \prod_{i=2}^L a_{k_{i-1}, k_i}, B_{\mathbf{k}} = \frac{1}{n^L} (p_{k_1} + 1) \prod_{i=2}^L (a_{k_{i-1}, k_i} + 1), \right. \\ & p \in \{0, \dots, n-1\}^{K-1}, \sum_{k=1}^{K-1} p_k < n, p_K = n - \sum_{k=1}^{K-1} (p_k + 1), \\ & a \in \{0, \dots, n-1\}^{K \times (K-1)}, \forall i \in \{1, \dots, K\}, \sum_{k=1}^{K-1} a_{i,k} < n \text{ and} \\ & \left. a_{i,K} = n - \sum_{k=1}^{K-1} (a_{i,k} + 1) \right\}. \end{aligned}$$

This set covers $\mathcal{U}$ and its cardinality is smaller than $\left(\frac{2}{\epsilon}\right)^{K^2-1}$. To get the bracket's size, note that

$$\begin{aligned} & \sum_{\mathbf{k} \in \{1, \dots, K\}^L} \left(\frac{1}{n^L} p_{k_1} \prod_{i=2}^L a_{k_{i-1}, k_i} - \frac{1}{n^L} (p_{k_1} + 1) \prod_{i=2}^L (a_{k_{i-1}, k_i} + 1) \right)^2 \\ &= \frac{1}{n^{2L}} \sum_{\mathbf{k} \in \{1, \dots, K\}^L} \left(\prod_{i=2}^L a_{k_{i-1}, k_i} + \sum_{j=2}^L p_{k_1} \prod_{i \neq j, i \geq 2} a_{k_{i-1}, k_i} \right)^2 \\ &\leq \frac{L^2 n^{2L-2} K^L}{n^{2L}} \\ &\leq L^2 K^L \epsilon^2, \end{aligned}$$

and in the end

$$N(u, \mathcal{U}, d_2) \leq \max \left(\frac{LK^{L/2}}{u}, 1 \right)^{K^2-1}.$$

Proof of Lemma 21 All $f \in \mathcal{F} \cap \mathfrak{P}_M$ can be written as $\sum_{m=1}^M \lambda_m \varphi_m$ where $(\varphi_m)_{m \in \{1, \dots, M\}}$ is an orthonormal basis of $\mathfrak{P}_M$. Then, assumption [HF] implies that $|\lambda_m| \leq C_{\mathcal{F},2}$ for all $m \in \{1, \dots, M\}$.

We will therefore start from a bracket covering of the euclidian ball of radius $C_{\mathcal{F},2}$ of $\mathbb{R}^M$, from which we will construct a covering of $\mathcal{F} \cap \mathfrak{P}_M$ and of Φ .

Lemma 22. *Let $\epsilon \in (0, 4)$. There exists a bracket covering $\{[a^i, b^i]\}_{1 \leq i \leq N_M}$ of size ϵ of the euclidian ball of radius $C_{\mathcal{F},2}$ of $\mathbb{R}^M$ with cardinality*

$$N_M \leq \max \left(\frac{4C_{\mathcal{F},2}\sqrt{M}}{\epsilon}, 1 \right)^M$$

such that for all $m \in \{1, \dots, M\}$, $i \in \{1, \dots, N_M\}$, $-C_{\mathcal{F},2} \leq a_m^i \leq b_m^i \leq C_{\mathcal{F},2}$.

Proof. We start with a bracket covering of size $\epsilon/\sqrt{M}$ of the infinity ball of radius $C_{\mathcal{F},2}$ of $\mathbb{R}^M$. This can be done by a regular partition with $\max(\lceil 2C_{\mathcal{F},2}/\epsilon \rceil, 1)$ pieces along each coordinate. One can easily check that such a covering is also a covering of size ϵ of the euclidian ball of radius $C_{\mathcal{F},2}$ of $\mathbb{R}^M$. To conclude, it is enough to notice that $\lceil x \rceil \leq 2x$ as soon as $x > 1/2$, and that $2C_{\mathcal{F},2}/\epsilon > 1/2$ because $C_{\mathcal{F},2} \geq 1$ and $\epsilon < 4$. $\square$

Let $\{[a^i, b^i]\}_{1 \leq i \leq N_M}$ be such a covering. For all $m \in \{1, \dots, M\}$, $i \in \{1, \dots, N_M\}$ and $y \in \mathcal{Y}$, let

$$u_m^i(y) = \begin{cases} a_m^i & \text{if } \varphi_m(y) \leq 0 \\ b_m^i & \text{otherwise} \end{cases}$$

$$v_m^i(y) = a_m^i + b_m^i - u_m^i(y)$$

and for all $i \in \{1, \dots, N_M\}$ and $y \in \mathcal{Y}$,

$$\begin{cases} U_1^i(y) = \sum_{m=1}^M u_m^i(y) \varphi_m(y) \\ U_2^i(y) = \sum_{m=1}^M v_m^i(y) \varphi_m(y) \end{cases}$$

and finally for all $\mathbf{i} = (i_1, \dots, i_K) \in \{1, \dots, N_M\}^K$ and $\mathbf{k} = (k_1, \dots, k_L) \in \{1, \dots, K\}^L$:

$$\begin{cases} (V^{\mathbf{i}})_{\mathbf{k}} = \min \left\{ \bigotimes_{\beta=1}^L U_{\sigma_{\beta}}^{i_{k_{\beta}}}; \quad \sigma \in \{1, 2\}^L \right\} \\ (W^{\mathbf{i}})_{\mathbf{k}} = \max \left\{ \bigotimes_{\beta=1}^L U_{\sigma_{\beta}}^{i_{k_{\beta}}}; \quad \sigma \in \{1, 2\}^L \right\} \end{cases}.$$

It is enough to show that $\{[V^{\mathbf{i}}, W^{\mathbf{i}}], \mathbf{i} \in \{1, \dots, N_M\}^K\}$ is a bracket covering of size $L(4C_{\mathcal{F},2}M)^{L-1}\sqrt{M}\epsilon$ of Φ that satisfies

$$\max_i \max_{\mathbf{k}} \int (W_{\mathbf{k}}^i)^2 d\mu^{\otimes L} \leq (8M^2 C_{\mathcal{F},2}^2)^L.$$

Applying the Cauchy-Schwarz inequality, one gets that for all $i \in \{1, \dots, N_M\}$,

$$\begin{aligned} \|U_2^i - U_1^i\|_2^2 &= \left\| \sum_{m=1}^M |b_m^i - a_m^i| \cdot |\varphi_m| \right\|_2^2 \\ &\leq M \|b^i - a^i\|_2^2 \\ &\leq M\epsilon^2. \end{aligned}$$

Moreover, for all $i \in \{1, \dots, N_M\}$ and $\sigma \in \{1, 2\}$,

$$\begin{aligned} \|U_{\sigma}^i\|_2^2 &\leq \left\| \sum_{m_i=1}^M |b_m^i + a_m^i| \cdot |\varphi_m| \right\|_2^2 \\ &\leq 2M(\|a^i\|_2^2 + \|b^i\|_2^2) \\ &\leq 4M^2 C_{\mathcal{F},2}^2. \end{aligned}$$

We then use that for all $\mathbf{i} \in \{1, \dots, N_M\}^K$ and $\mathbf{k} \in \{1, \dots, K\}^L$,

$$\begin{aligned} |W_{\mathbf{k}}^{\mathbf{i}} - V_{\mathbf{k}}^{\mathbf{i}}|(y) &\leq \sum_{\gamma=1}^L \left| U_2^{i_{k_\gamma}} - U_1^{i_{k_\gamma}} \right| (y_\gamma) \max_{\mathbf{j} \in \{1,2\}^L} \prod_{\beta \neq \gamma, \beta=1}^L \left| U_{j_\beta}^{i_{k_\beta}} \right| (y_\beta) \\ &\leq \sum_{\gamma=1}^L \left| U_2^{i_{k_\gamma}} - U_1^{i_{k_\gamma}} \right| (y_\gamma) \prod_{\beta \neq \gamma, \beta=1}^L \left(\left| U_1^{i_{k_\beta}} \right| + \left| U_2^{i_{k_\beta}} \right| \right) (y_\beta) \end{aligned}$$

so that

$$\begin{aligned} \|W_{\mathbf{k}}^{\mathbf{i}} - V_{\mathbf{k}}^{\mathbf{i}}\|_2^2 &\leq L \sum_{\gamma=1}^L \left\| U_2^{i_{k_\gamma}} - U_1^{i_{k_\gamma}} \right\|_2^2 \prod_{\beta \neq \gamma, \beta=1}^L 2 \left(\left\| U_1^{i_{k_\beta}} \right\|_2^2 + \left\| U_2^{i_{k_\beta}} \right\|_2^2 \right) \\ &\leq L 2^{L-1} \sum_{\gamma=1}^L M \epsilon^2 \prod_{\beta \neq \gamma, \beta=1}^L (2 \times 4M^2 C_{\mathcal{F},2}^2) \\ &= L^2 (16M^2 C_{\mathcal{F},2}^2)^{L-1} M \epsilon^2 \\ &= (L(4C_{\mathcal{F},2}M))^{L-1} \sqrt{M} \epsilon^2 \end{aligned}$$

and finally $d_{\infty,2}(W^{\mathbf{i}}, V^{\mathbf{i}}) \leq L(4C_{\mathcal{F},2}M)^{L-1} \sqrt{M} \epsilon$ for all $\mathbf{i} \in \{1, \dots, N_M\}^K$.

The last part of the lemma is proved by noting that for all $\mathbf{i}$ and $\mathbf{k}$,

$$\begin{aligned} (W^{\mathbf{i}})_{\mathbf{k}}^2 &= \max_{\sigma \in \{1,2\}^L} \left\{ \bigotimes_{\beta=1}^L (U_{\sigma_\beta}^{i_{k_\beta}})^2 \right\} \\ &\leq \sum_{\sigma \in \{1,2\}^L} \bigotimes_{\beta=1}^L (U_{\sigma_\beta}^{i_{k_\beta}})^2, \end{aligned}$$

so that

$$\begin{aligned} \int (W^{\mathbf{i}})_{\mathbf{k}}^2 d\mu^{\otimes L} &\leq \sum_{\sigma \in \{1,2\}^L} \prod_{\beta=1}^L \|U_{\sigma_\beta}^{i_{k_\beta}}\|_2^2 \\ &\leq \sum_{\sigma \in \{1,2\}^L} (4M^2 C_{\mathcal{F},2}^2)^L \\ &\leq (8M^2 C_{\mathcal{F},2}^2)^L. \end{aligned}$$

C3. Choice of parameters

Let us come back to equation (6). Since $x \mapsto \frac{\varphi(x)}{x}$ is nonincreasing, one has $\frac{\varphi(x_{K,M} C_{\mathcal{F},\infty}^{L/2})}{x_{K,M} \sqrt{N}} \leq \sigma_{K,M} C_{\mathcal{F},\infty}^{L/2}$ as soon as $x_{K,M} \geq \frac{\sigma_{K,M}}{C_{\mathcal{F},\infty}^{L/2}}$, so with probability $1 - e^{-z_{K,M}-z}$:

$$Z_{K,M}(s) \leq 4C^* \left[C C_{\mathcal{F},\infty}^{L/2} \frac{\sigma_{K,M}}{x_{K,M}} + C_{\mathcal{F},\infty}^{L/2} \sqrt{\frac{z_{K,M} + z}{x_{K,M}^2 N}} + 2C_{\mathcal{F},\infty}^L \frac{z_{K,M} + z}{x_{K,M}^2 N} \right].$$

Let $C' = C^* \max(C, 1) C_{\mathcal{F},\infty}^L$. One gets

$$Z_{K,M}(s) \leq 4C' \left[\frac{\sigma_{K,M}}{x_{K,M}} + \sqrt{\frac{z_{K,M} + z}{x_{K,M}^2 N}} + \frac{z_{K,M} + z}{x_{K,M}^2 N} \right].$$

Let $x_{K,M} = \theta^{-1} \sqrt{\sigma_{K,M}^2 + \frac{z_{K,M} + z}{N}}$ with θ such that $2\theta + \theta^2 \leq 1/(16C')$. Then, with probability $1 - e^{-z_{K,M}-z}$:

$$Z_{K,M}(s) \leq 4C'(\theta + \theta + \theta^2) \leq \frac{1}{4}$$

Now choose $z_{K,M} = M + K$, it follows that $\sum_{K \in \mathbb{N}^*, M \in \mathcal{M}} e^{-z_{K,M}} \leq (e - 1)^{-2} \leq 1$ and the first point of the lemma is proved.

Moreover, one has with probability $1 - e^{-z}$, for all K, M :

$$\begin{aligned} Z_{K,M}(s) x_{K,M}^2 &\leq 4C' \left[\sigma_{K,M} x_{K,M} + x_{K,M} \sqrt{\frac{z_{K,M} + z}{N}} + \frac{z_{K,M} + z}{N} \right] \\ &\leq 4C' \left[2\theta^{-1} x_{K,M}^2 + \frac{z_{K,M} + z}{N} \right] \\ &= 4C' \left[2\theta^{-1} \sigma_{K,M}^2 + (2\theta^{-1} + 1) \frac{M + K}{N} + (2\theta^{-1} + 1) \frac{z}{N} \right] \end{aligned}$$

Let $A = 4C'(2\theta^{-1} + 1)$. We get that with probability $1 - e^{-z}$, for all K, M :

$$Z_{K,M}(s) x_{K,M}^2 \leq A \left[\sigma_{K,M}^2 + \frac{M + K}{N} + \frac{z}{N} \right]$$

Therefore the lemma holds as soon as

$$\forall K \leq N, \forall M \leq N, \quad \widetilde{\text{pen}}(N, M, K) \geq A \left[\sigma_{K,M}^2 + \frac{M + K}{N} \right] \quad (11)$$

Lemma 23. *There exists constants C_1 and N_1 such that for all $N \geq N_1$:*

$$\sigma_{K,M} \leq C_1 \sqrt{\frac{MK + K^2 - 1}{N}} (1 + \sqrt{\log(N)})$$

Proof. Let $x(C) = C \sqrt{\frac{MK + K^2 - 1}{N}} (1 + \sqrt{\log(N)})$.

$\sigma_{K,M}$ is defined by the equation $\frac{\varphi(x)}{x^2 \sqrt{N}} = 1$. The function $x \mapsto \frac{\varphi(x)}{x^2}$ is nondecreasing, so it is enough to show that $\frac{\varphi(x(C))}{x(C)^2 \sqrt{N}} \leq 1$ for some constant C that we can assume to be greater than 1.

It is easy to check that there exists a constant N_1 such that for all $N \geq N_1$, $\frac{\varphi(N^{6L})}{(N^{6L})^2 \sqrt{N}} \leq 1$, so that $\sigma_{K,M} \leq N^{6L}$, which makes it possible to assume $x(C) \leq N^{6L}$. Then

$$\begin{aligned} \frac{\varphi(x(C))}{x(C)^2 \sqrt{N}} &= \frac{C_0}{C} \frac{1 + \sqrt{\log\left(\frac{N^{6L+1/2}}{C \sqrt{(MK+K^2-1)}(1+\sqrt{\log(N)})}\right)}}{1 + \sqrt{\log(N)}} \\ &\leq \frac{C_0}{C} \frac{1 + \sqrt{\log(N^{7L})}}{1 + \sqrt{\log(N)}} \\ &= \frac{C_0}{C} \frac{1 + \sqrt{7L} \sqrt{\log(N)}}{1 + \sqrt{\log(N)}} \end{aligned}$$

and by taking $C_1 = \max(C_0 \sqrt{7L}, 1)$, one gets that

$$\frac{\varphi(x(C_1))}{x(C_1)^2 \sqrt{N}} \leq 1$$

which means that $\sigma_{K,M} \leq x(C_1)$. □

The condition of equation (11) becomes

$$\widetilde{\text{pen}}(N, M, K) \geq A \left[\frac{C_1^2 (MK + K^2 - 1) (1 + \sqrt{\log(N)})^2 + M + K}{N} \right]$$

which is implied by

$$\widetilde{\text{pen}}(N, M, K) \geq \rho (MK + K^2 - 1) \frac{\log(N)}{N}$$

for some constant ρ depending only on $C_{\mathcal{F},2}$, $C_{\mathcal{F},\infty}$, $\mathbf{Q}^*$ and L . This concludes the proof.

C4. Auxiliary lemmas

Lemma 24.

$$\forall t \in \bigcup_K S_K, \quad \begin{cases} \|t\|_\infty \leq C_{\mathcal{F},\infty}^L \\ \|t\|_2 \leq C_{\mathcal{F},2}^L \\ \mathbb{E}[t^2] \leq C_{\mathcal{F},\infty}^L \|t\|_2^2 \end{cases}$$

Proof. t can be written as $t = g^{\pi, \mathbf{Q}, \mathbf{f}}$ with π a probability K -uple, $\mathbf{Q}$ a transition matrix of size K and $\mathbf{f} \in \mathcal{F}^K$ for some $K \geq 1$.

The first point follows from

$$\begin{aligned} \|t\|_\infty &= \left\| \sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \bigotimes_{i=1}^L f_{k_i} \right\|_\infty \\ &\leq \sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \left\| \bigotimes_{i=1}^L f_{k_i} \right\|_\infty \\ &\leq \sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \prod_{i=1}^L \|f_{k_i}\|_\infty \\ &\leq C_{\mathcal{F},\infty}^L \sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \\ &= C_{\mathcal{F},\infty}^L \end{aligned}$$

For the second point, we use the Cauchy-Schwarz inequality:

$$\begin{aligned}
\|t\|_2^2 &= \int \left(\sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \prod_{i=1}^L f_{k_i}(y_i) \right)^2 \mathrm{d}\mu(y_1) \dots \mathrm{d}\mu(y_L) \\
&= \int \left(\sum_{k_1, \dots, k_L=1}^K \sqrt{\pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i)} \right. \\
&\quad \left. \left(\sqrt{\pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i)} \prod_{i=1}^L f_{k_i}(y_i) \right) \right)^2 \mathrm{d}\mu(y_1) \dots \mathrm{d}\mu(y_L) \\
&\leq \int \left(\sum_{k'_1, \dots, k'_L=1}^K \pi(k'_1) \prod_{i=2}^L \mathbf{Q}(k'_{i-1}, k'_i) \right) \\
&\quad \left(\sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \prod_{i=1}^L f_{k_i}^2(y_i) \right) \mathrm{d}\mu(y_1) \dots \mathrm{d}\mu(y_L) \\
&= \sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \int \prod_{i=1}^L f_{k_i}^2(y_i) \mathrm{d}\mu(y_1) \dots \mathrm{d}\mu(y_L) \\
&= \sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) \prod_{i=1}^L \|f_{k_i}\|_2^2 \\
&\leq \sum_{k_1, \dots, k_L=1}^K \pi(k_1) \prod_{i=2}^L \mathbf{Q}(k_{i-1}, k_i) C_{\mathcal{F},2}^{2L} \\
&= C_{\mathcal{F},2}^{2L}
\end{aligned}$$

The last point comes from

$$\begin{aligned}
\mathbb{E}[t^2] &= \int g^* t^2 \mathrm{d}\mu^{\otimes L} \\
&\leq \int \|g^*\|_\infty t^2 \mathrm{d}\mu^{\otimes L} \\
&\leq C_{\mathcal{F},\infty}^L \int t^2 \mathrm{d}\mu^{\otimes L} \quad \text{par le premier point} \\
&= C_{\mathcal{F},\infty}^L \|t\|_2^2
\end{aligned}$$

□

Lemma 25. Let $A, B \in \mathbb{R}_+^*$. Let $H : x \in \mathbb{R}_+^* \mapsto A \log \max(\frac{B}{x}, 1)$, and $\varphi(x) : x \in \mathbb{R}_+^* \mapsto x\sqrt{\pi A}(1 + \sqrt{\log \max(\frac{B}{x}, 1)})$. Then:

$$\begin{cases} x^2 H(x) \leq \varphi(x)^2 \\ \int_0^x \sqrt{H(u)} du \leq \varphi(x) \end{cases}$$

Proof. The first point is straightforward.

For the second point, we have two cases.

Case 1: $x \leq B$. Then $H(x) = \log(\frac{B}{x})$. Therefore, we can use that $\int_0^\sigma \sqrt{\log(\frac{B}{x})} dx \leq \sigma(\sqrt{\pi} + \sqrt{\log(\frac{B}{\sigma})})$, which is enough to conclude.

Case 2: $x \geq B$. Then $H(x) = 0$ and $\varphi(x) = x\sqrt{\pi A} \geq B\sqrt{\pi A} = \varphi(B)$. Thus,

$$\begin{aligned} \int_0^x \sqrt{H(u)} du &= \int_0^B \sqrt{H(u)} du \\ &\leq \varphi(B) \\ &\leq \varphi(x) \end{aligned}$$

□

References

- Alexandrovich, Grigory and Hajo Holzmann, “Nonparametric identification of hidden Markov models,” *arXiv preprint arXiv:1404.4210*, 2014.
- Anandkumar, Animashree, Daniel Hsu, and Sham M Kakade, “A method of moments for mixture models and hidden Markov models,” *arXiv preprint arXiv:1203.0683*, 2012.
- Baudry, Jean-Patrick, Cathy Maugis, and Bertrand Michel, “Slope heuristics: overview and implementation,” *Statistics and Computing*, 2012, 22 (2), 455–470.
- Celeux, Gilles and Jean-Baptiste Durand, “Selecting hidden Markov model state number with cross-validated likelihood,” *Computational Statistics*, 2008, 23 (4), 541–564.
- Chambaz, Antoine, Aurelien Garivier, and Elisabeth Gassiat, “A minimum description length approach to hidden Markov models with Poisson and Gaussian emissions. Application to order identification,” *Journal of Statistical Planning and Inference*, 2009, 139 (3), 962–977.

- Couvreur, Laurent and Christophe Couvreur, “Wavelet-based non-parametric HMM’s: theory and applications,” in “Acoustics, Speech, and Signal Processing, 2000. ICASSP’00. Proceedings. 2000 IEEE International Conference on,” Vol. 1 IEEE 2000, pp. 604–607.
- de Castro, Yohann, Elisabeth Gassiat, and Claire Lacour, “Minimax adaptative estimation of non-parametric hidden Markov models,” *JMLR*, 2015. To appear.
- , —, and Sylvain Le Corff, “Consistent estimation of the filtering and marginal smoothing distributions in nonparametric hidden Markov models,” *arXiv preprint arXiv:1507.06510*, 2015.
- Gassiat, Elisabeth, “Likelihood ratio inequalities with applications to various mixtures,” in “Annales de l’IHP Probabilités et statistiques,” Vol. 38 2002, pp. 897–906.
- , Alice Cleynen, and Stéphane Robin, “Finite state space non parametric hidden Markov models are in general identifiable,” *Stat. Comp.*, 2015, pp. 1–11.
- and Stéphane Boucheron, “Optimal error exponents in hidden Markov models order estimation,” *Information Theory, IEEE Transactions on*, 2003, 49 (4), 964–980.
- , Judith Rousseau et al., “About the posterior distribution in hidden Markov models with unknown number of states,” *Bernoulli*, 2014, 20 (4), 2039–2075.
- Hansen, Nikolaus, “The CMA evolution strategy: a comparing review,” in “Towards a new evolutionary computation,” Springer, 2006, pp. 75–102.
- Hsu, Daniel, Sham M Kakade, and Tong Zhang, “A spectral algorithm for learning hidden Markov models,” *Journal of Computer and System Sciences*, 2012, 78 (5), 1460–1480.
- Lambert, Martin F, Julian P Whiting, and Andrew V Metcalfe, “A non-parametric hidden Markov model for climate state identification,” *Hydrology and Earth System Sciences Discussions*, 2003, 7 (5), 652–667.
- Langrock, Roland, Thomas Kneib, Alexander Sohn, and Stacy L DeRuiter, “Nonparametric inference in hidden Markov models using P-splines,” *Biometrics*, 2015, 71 (2), 520–528.
- Lefèvre, Fabrice, “Non-parametric probability estimation for HMM-based automatic speech recognition,” *Computer Speech & Language*, 2003, 17 (2), 113–136.
- Lehéricy, Luc, “Estimation adaptative non paramétrique pour les modèles à chaîne de Markov cachée,” 2015. Mémoire de M2, Orsay.
- Massart, Pascal, “Concentration inequalities and model selection,” in “Lecture Notes in Mathematics,” Vol. 1896, Berlin: Springer, 2007.

- Paulin, Daniel, “Concentration inequalities for Markov chains by Marton couplings,” *arXiv preprint arxiv:1212.2015v2*, 2013.
- Shang, Lifeng and Kwok-Ping Chan, “Nonparametric discriminant HMM and application to facial expression recognition,” in “Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on” IEEE 2009, pp. 2090–2096.
- Stewart, GW and Ji-Guang Sun, “Matrix Perturbation Theory (Computer Science and Scientific Computing),” 1990.
- van Havre, Zoé, Judith Rousseau, Nicole White, and Kerrie Mengersen, “Overfitting hidden Markov models with an unknown number of states,” *arXiv preprint arXiv:1602.02466*, 2016.
- Volant, Stevann, Caroline Bérard, Marie-Laure Martin-Magniette, and Stéphane Robin, “Hidden Markov models with mixtures as emission distributions,” *Statistics and Computing*, 2014, 24 (4), 493–504.
- Yau, C, Omiros Papaspiliopoulos, Gareth O Roberts, and Christopher Holmes, “Bayesian non-parametric hidden Markov models with applications in genomics,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2011, 73 (1), 37–57.