



HAL
open science

General oracle inequalities for a penalized log-likelihood criterion based on non-stationary data

Julien Aubert, Luc Lehéricy, Patricia Reynaud-Bouret

► **To cite this version:**

Julien Aubert, Luc Lehéricy, Patricia Reynaud-Bouret. General oracle inequalities for a penalized log-likelihood criterion based on non-stationary data. 2024. hal-04578260

HAL Id: hal-04578260

<https://hal.science/hal-04578260>

Preprint submitted on 16 May 2024

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

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



Distributed under a Creative Commons Attribution 4.0 International License

General oracle inequalities for a penalized log-likelihood criterion based on non-stationary data

Julien Aubert*, Luc Lehéricy* and Patricia Reynaud-Bouret*

*Université Côte d'Azur, CNRS, LJAD, France

Abstract

We prove oracle inequalities for a penalized log-likelihood criterion that hold even if the data are not independent and not stationary, based on a martingale approach. The assumptions are checked for various contexts: density estimation with independent and identically distributed (i.i.d) data, hidden Markov models, spiking neural networks, adversarial bandits. In each case, we compare our results to the literature, showing that, although we lose some logarithmic factors in the most classical case (i.i.d.), these results are comparable or more general than the existing results in the more dependent cases.

Keywords : log-likelihood, penalization, oracle inequality, dependent data.

AMS Classification : 62G05, 62M05, 62M99

1 Introduction

Maximum likelihood estimator (MLE) (see [Bickel and Doksum, 2015] and the references therein) is often considered a graal, even if it has been often debated w.r.t. robustness in particular [Baraud and Birgé, 2016]. In particular, in the i.i.d. setting, it is known to be asymptotically optimal (variance which is asymptotically minimal w.r.t. the Cramer-Rao bound) under mild conditions (typically differentiability of the density) [Lehmann, 1999]. In 1973, [Akaike, 1973] proposed its famous penalized log-likelihood criterion (AIC) stating that to select in a finite set of models, one should penalize the log-likelihood of the model m by D_m , the number of parameters describing m . There is a large variety of variants of AIC, for which the asymptotic properties are more or less precise (see [Burnham and Anderson, 2004] and the references therein). Under additional differentiability assumptions, the Wilks phenomenon [Wilks, 1938] quantifies more precisely (in an asymptotic way) how the recentered maximal log-likelihood behaves as a chi-square distribution with D_m degrees of freedom. This phenomenon allows to construct asymptotic likelihood ratio test and therefore to perform model selection by multiple testing [Zheng et al., 2019]. This idea has been used in many settings, even quite recently in combination with asymptotic ℓ_1 model selection [Tang and Leng, 2010, Sur et al., 2019, Sur and Candès, 2019]. In short, with AIC-like penalties and the Wilks phenomenon, we have a quite clear asymptotic picture in the i.i.d. setting on how the maximum likelihood behaves in terms of dimension and how it has to be penalized to find the correct model in a finite fixed set of smooth enough models.

Moving on to the non-asymptotic setting, things are more difficult to study. It is possible to get concentration inequalities that mimic the Wilks phenomenon [Boucheron and Massart, 2011]. With this kind of non asymptotic results, it is possible to understand non asymptotically what the MLE does in an exponential family [Spokoiny, 2012], or how to penalize it to make it work even in infinite dimensional settings with finite effective dimension [Spokoiny, 2017]. It is also possible to penalize the log-likelihood to perform model selection [Castellan, 2003, Massart, 2007] for particular family of models in the i.i.d. setting. However, in these works, model selection "à la Birgé-Massart" [Birgé and Massart, 2001] is based on a form of linearity between the contrast and the family of models under consideration, which leads to a very precise tuning of the penalty constants [Arlot, 2019] but which at the same time, prevents the results to be applied in more general settings.

If we drop the independence assumption but keep the stationarity of the data, there have been a vast variety of model selection results by minimization of ℓ_0 or ℓ_1 penalized contrasts with or without log-likelihood. Let us cite but a few: Markov chains [Lacour, 2008b], hidden Markov models [Lacour, 2008a, De Castro et al., 2016, Lehéricy, 2021], spiking neuronal networks with unobserved components [Ost and Reynaud-Bouret, 2020], point processes [Hansen et al., 2015]. Each time, the arguments are a combination of martingale approach and non asymptotic exponential inequalities, that derive from the ergodicity of the process and mixing properties.

However, the Akaike criterion (minus log-likelihood plus penalization proportional to D_m) is routinely used even beyond this setting [Daw, 2011, Wilson and Collins, 2019]. In particular, in learning experiments where the individual has to learn how to perform a task, and learn it only once, data are neither independent nor stationary. Some authors [Ramponi et al., 2020] have tried to assume that individuals are identically distributed, but this is a very strong and quite unlikely setting in practice. In a first work [Aubert et al., 2023], we proved meaningful bounds for the maximum likelihood estimator on an individual learning trajectory. However, up to our knowledge, there is no theoretical work on model selection in this setting or more generally for dependent and non stationary data.

The purpose of this work is to derive a non-asymptotic oracle inequality for AIC-like model selection that is general enough to cover all the setups listed before: i.i.d. samples, Markov chains, hidden Markov models, partially observed neuronal networks, learning models and more.

The proof relies on a martingale approach and an exponential inequality that was recently proved in [Aubert and Lehéricy, 2024] for a supremum of empirical centered processes that are stochastically normalized. This exponential inequality generalizes the works of [Baraud, 2010], [Talagrand, 1996], and is inspired by the works on renormalized martingales due to [Bercu et al., 2015, Bercu and Touati, 2019] and [de la Peña et al., 2004]. This is the key non asymptotic exponential inequality in the present work.

The rest of the paper is organised as follows. In Section 2, we introduce the notations and the general framework. In Section 3, we state our assumptions and oracle inequalities in probability and in expectation, in the bounded and unbounded frameworks. We also discuss precisely how the penalty in D_m is obtained with respect to more classical proof techniques. In Section 4, we discuss various cases that are covered by these general results in light of the existing results of the literature. The appendices are dedicated to the proofs.

2 Framework and notation

Given two integers $a \leq b$ and a sequence $(x_s)_{s \in \mathbb{Z}}$, write $x_a^b = (x_a, \dots, x_b)$ (with x_a^b being the empty sequence when $a > b$). For any two real numbers x, y , write $x \vee y$ their maximum and $x \wedge y$ their minimum. Let \mathbb{N}^* be the set of positive integers, and for any $n \in \mathbb{N}^*$, write $[n] = \{1, \dots, n\}$. Finally, \log denotes the natural logarithm.

Let $n \geq 2$. We observe a process $(X_t)_{1 \leq t \leq n}$ defined on a polish measure space $(\mathcal{X}, \mathcal{F}, \mu)$ and adapted to a filtration $(\mathcal{F}_t)_{1 \leq t \leq n}$. We write \mathbb{P} the corresponding probability and \mathbb{E} the corresponding expectation.

We are interested in the successive conditional distributions of X_t . If \mathcal{X} is discrete and μ is the counting measure, we are therefore interested by the sequence

$$p_t^*(\cdot) = \mathbb{P}(X_t = \cdot | \mathcal{F}_{t-1}), \quad \forall t \in [n].$$

More generally, in the sequel, for general measured spaces $(\mathcal{X}, \mathcal{F}, \mu)$, we always assume that the conditional density of X_t given \mathcal{F}_{t-1} with respect to μ exists and we write it $p_t^*(\cdot)$. Note that for all $x \in \mathcal{X}$, $(p_t^*(x))_{1 \leq t \leq n}$ is therefore predictable with respect to the filtration—we write that p_t^* is predictable for short. We write $p^* = (p_t^*)_{t \in [n]}$ the vectors of all the successive conditional densities.

2.1 Some examples

The emblematic case is the case where we take for filtration $\mathcal{F}_t = \sigma(X_1^t)$ for $t \geq 1$ and \mathcal{F}_0 is the trivial σ -algebra. In this case, for $t > 1$, we can rewrite p_t^* as the conditional density of X_t given X_1^{t-1} and p_1^* as the density of X_1 (in this case p_1^* is deterministic). To emphasize this fact, in this case we note $p_t^*(\cdot | X_1^{t-1})$ instead of just p_t^* . Note that the density of the vector X_1^n with respect to $\mu^{\otimes n}$ is therefore

$$x_1^n \mapsto \prod_{i=1}^n p_i^*(x_i | X_1^{i-1} = x_1^{i-1}). \quad (1)$$

In the even simpler case that the coordinates of X_1^n are independent, for all $t \geq 1$, p_t^* is the (deterministic) density of X_t w.r.t. to μ .

We might be interested by an enlargement of the filtration. For instance, the conditional distribution of X_t can be a function of not only the past realizations of X_t , but also of additional covariates: imagine that each step t , the distribution of X_t depends on the past but also on an observed variable C_t , so that the conditional densities can be written

$$p_t^*(x_t) = p_t^*(x_t | X_1^{t-1}, C_1^t). \quad (2)$$

Every decent model of evolution for X_t depends on X_1^{t-1} but also on C_1^t . The natural filtration in this setting is $\mathcal{F}_{t-1} = \sigma(X_1^{t-1}, C_1^t)$.

Another setup, where the filtration might not be completely observed, is chains with infinite memory that can model potentially infinite neuronal networks. Even if we do not observe the chain on the whole network with infinite past, we might be interested by approximating p_t^* , which exists as a function of the infinite network, with functions that only involve a smaller subset of observed neurons (see Section 4.3 for more details).

2.2 Models and penalized (partial) log-likelihood

Consider a sequence of models $(\{p_\theta^m : \theta \in \Theta_m \subset \mathbb{R}^{D_m}\})_{m \in \mathcal{M}}$ for some countable set \mathcal{M} . Each p_θ^m is a sequence $p_\theta^m = (p_{\theta,t}^m)_{t \in [n]}$, with $p_{\theta,t}^m$ being a candidate at being p_t^* . In particular, the candidate $p_{\theta,t}^m$ must be predictable.

For any $m \in \mathcal{M}$, define the (partial) log-likelihood of parameter $\theta \in \Theta_m$ given the observations by

$$\ell_n(\theta) = \sum_{t=1}^n \log p_{\theta,t}^m(X_t).$$

Note that in the case mentioned above where $\mathcal{F}_t = \sigma(X_1^t)$ for $t \geq 1$ and \mathcal{F}_0 is the trivial σ -algebra, because of (1), $\ell_n(\theta)$ is exactly the log-likelihood $\log p_\theta^m(X_1^n)$. For models as in (2), this partial log-likelihood still has the convenient form of a conditional log-likelihood: $\ell_n(\theta) = \log p_\theta^m(X_1^n | C_1^n)$, and it matches Cox's partial likelihood [Cox, 1975]. This might no longer be the case for larger filtrations.

For each $m \in \mathcal{M}$, define the maximum likelihood estimator of model m by

$$\hat{\theta}^m \in \arg \max_{\theta \in \Theta_m} \frac{1}{n} \ell_n(\theta).$$

Finally, take a penalty $\text{pen} : \mathcal{M} \rightarrow \mathbb{R}_+$ and select a model \hat{m} that minimizes the penalized log-likelihood:

$$\hat{m} \in \arg \min_{m \in \mathcal{M}} \left(-\frac{1}{n} \ell_n(\hat{\theta}^m) + \text{pen}(m) \right).$$

The penalized likelihood estimator of p^* is therefore $\tilde{p} = p_{\hat{\theta}^{\hat{m}}}^{\hat{m}}$. We aim to understand the properties of approximation of \tilde{p} with respect to p^* .

2.3 Stochastic loss function

Classical approaches [Massart, 2007, Spokoiny, 2012, Spokoiny, 2017], generally use an expectation of the contrast to define the loss function. For instance, in i.i.d. examples, the log-likelihood is naturally linked to the Kullback-Leibler divergence between the distributions.

Here, we want to keep the inherent martingale structure that comes with the filtration and with our object of interest p^* . This is why we are using the stochastic loss function \mathbf{K}_n defined as follows: for any sequence of conditional densities $p = (p_t)_{t \in [n]}$, let

$$\mathbf{K}_n(p) = \frac{1}{n} \sum_{t=1}^n \mathbb{E} \left[\log \frac{p_t^*(X_t)}{p_t(X_t)} \middle| \mathcal{F}_{t-1} \right].$$

This can be seen as the mean of the conditional Kullback-Leibler divergence in the sense that

$$\mathbb{E} \left[\log \frac{p_t^*(X_t)}{p_t(X_t)} \middle| \mathcal{F}_{t-1} \right]$$

is a predictable quantity which corresponds to the Kullback-Leibler divergence between the distributions with densities p_t^* and p_t w.r.t μ , conditionally to \mathcal{F}_{t-1} .

Note that in the case where $\mathcal{F}_t = \sigma(X_1^t)$ for $t \geq 1$ and \mathcal{F}_0 is the trivial σ -algebra, because of (1), $n\mathbb{E}[\mathbf{K}_n(p)]$ is exactly the Kullback-Leibler divergence between the distributions defined respectively by p^* and p .

Our oracle inequalities presented below are bounds on the stochastic loss $\mathbf{K}_n(\tilde{p})$ in probability and in expectation respectively.

3 Main results

3.1 Assumptions

Let us first discuss the main assumptions.

Because we are using a Kullback-Leibler-like divergence as a loss, we need to ensure that it does not diverge. A natural assumption in this respect is to assume that p^* and the candidates p_θ^m stay far from 0 and do not explode. This is done almost surely in Assumption 1 and with high probability only in Assumption 1bis.

Assumption 1. *There exists $\varepsilon > 0$ such that a.s., for all $t \in [n]$, $p_t^*(X_t) \in [\varepsilon, \varepsilon^{-1}]$ and for all $m \in \mathcal{M}$ and all $\theta \in \Theta_m$, $p_{\theta,t}^m(X_t) \in [\varepsilon, \varepsilon^{-1}]$. Without loss of generality, we assume $\log \varepsilon < -1$.*

A weaker version of this assumption is the following.

Assumption 1bis. *For all $m \in \mathcal{M}$, there exists a finite constant B_m such that for all $y \geq 1$,*

$$\mathbb{P} \left(\exists t \in [n], \exists m \in \mathcal{M}, \sup_{\delta, \theta \in \Theta_m \cup \{\star\}} \left| \log \frac{p_{\delta,t}^m(X_t)}{p_{\theta,t}^m(X_t)} \right| > B_m y \mid \mathcal{F}_{t-1} \right) \leq e^{-y} \quad a.s.,$$

with the convention $p_\star^m = p^*$. Without loss of generality, we assume $B_m \geq 1$ for all $m \in \mathcal{M}$.

The second category of assumptions replaces the exponential family assumption of the classic asymptotic results on MLE. It states that the parameterization of the models is Lipschitz w.r.t. some norm that is bounded on Θ_m .

Assumption 2. *For all $m \in \mathcal{M}$, there exist a norm $\|\cdot\|_m$ on \mathbb{R}^{D_m} and finite, positive constants L_m and M_m such that a.s., for all $t \in [n]$, for all $m \in \mathcal{M}$ and all $\delta, \theta \in \Theta_m$,*

$$\left| \log \frac{p_{\delta,t}^m(X_t)}{p_{\theta,t}^m(X_t)} \right| \leq L_m \|\delta - \theta\|_m$$

and

$$\|\delta - \theta\|_m \leq M_m.$$

Without loss of generality, we assume $L_m M_m \geq 1$ for all $m \in \mathcal{M}$.

Likewise, a weaker version of this assumption is

Assumption 2bis. For all $m \in \mathcal{M}$, there exist a norm $\|\cdot\|_m$ on \mathbb{R}^{D_m} and finite, positive constants L_m and M_m such that

$$\mathbb{P} \left(\exists t \in [n], \exists m \in \mathcal{M}, \sup_{\delta, \theta \in \Theta_m} \left| \log \frac{p_{\delta, t}^m(X_t)}{p_{\theta, t}^m(X_t)} \right| > L_m \|\delta - \theta\|_m \log n \right) \leq n^{-1}$$

and

$$\|\delta - \theta\|_m \leq M_m.$$

Without loss of generality, we assume $L_m M_m \geq 1$ for all $m \in \mathcal{M}$.

Finally we will need in the proof to define a supremum of $x \mapsto p_t^*(x)$ and $p_{\theta, t}^m(x)$ over \mathcal{X} . The following assumption ensures that we can define it in a way that is measurable w.r.t. the filtration \mathcal{F}_{t-1} .

Assumption 3. There exists a countable dense subset \mathcal{Q} of \mathcal{X} such that almost surely, for all $t \in [n]$, $m \in \mathcal{M}$ and $\theta \in \Theta_m$,

$$\sup_{x \in \mathcal{Q}} p_{\theta, t}^m(x) = \sup_{x \in \mathcal{X}} p_{\theta, t}^m(x)$$

and

$$\sup_{x \in \mathcal{Q}} p_t^*(x) = \sup_{x \in \mathcal{X}} p_t^*(x).$$

Note that if \mathcal{X} is discrete and countable, this assumption is automatically verified. Likewise, since \mathcal{X} is separable (because it is assumed to be Polish), if almost surely, for all $m \in \mathcal{M}$ and $\theta \in \Theta_m$, the (random) functions $x \mapsto p_t^*(x)$ and $x \mapsto p_{\theta, t}^m(x)$ are continuous, this assumption holds.

With this small set of assumptions, we can now state our oracle inequalities.

3.2 Oracle inequality, bounded case

Theorem 1. Assume that $n \geq 2$ and that Assumptions 1, 2 and 3 hold. There exist positive numerical constants C and C' such that the following holds. For each $m \in \mathcal{M}$, let $A_m = L_m M_m + 2 \log(\varepsilon^{-1})$, and assume that

$$\Sigma = \sum_{m \in \mathcal{M}} \log(A_m) e^{-D_m} < +\infty.$$

Let $\kappa \in (0, 1]$. If for all $m \in \mathcal{M}$,

$$pen(m) \geq \frac{C}{\kappa} A_m^2 \log(\varepsilon^{-1})^{3/2} \log(n A_m)^2 \frac{D_m}{n},$$

then for all $x \geq 0$, with probability at least $1 - 18 \log(n) \Sigma e^{-x}$,

$$(1 - \kappa) \mathbf{K}_n(\tilde{p}) \leq \inf_{m \in \mathcal{M}} \left((1 + \kappa) \inf_{\theta \in \Theta^{D_m}} \mathbf{K}_n(p_\theta^m) + 2 pen(m) \right) + \frac{C'}{\kappa} \left(A_m \log(\varepsilon^{-1})^{3/2} \log(n A_m)^2 + A_m \log(\varepsilon^{-1})^{3/2} \log(n A_m)^2 \right) \frac{x}{n}.$$

Proof. See Section A. □

This result is an oracle inequality in probability. The penalty term is proportional to D_m/n , as in classical oracle inequalities for nested or not too complex families of models. This is ensured by the summability condition on Σ (see [Birgé and Massart, 2001, Massart, 2007] for instance for a discussion about the complexity of a family of models). Hence, it can be read as a usual oracle inequality: the loss is—up to a constant factor and a residual term—smaller than the best bias-variance trade-off in the family of models, with a variance which is of order D_m/n .

Due to the generality of the result, the penalty is actually a bit larger than in the original AIC criterion, with additional logarithmic factors, and depends on the lower bound ε and the Lipschitz constants L_m and M_m . Note the presence of a term depending on \hat{m} in the residual term: due to the fact that \mathcal{M} might be infinite, additional assumptions are required to get rid of it, such as the uniform bound on $(A_m)_{m \in \mathcal{M}}$ introduced in the following corollary.

The next corollary gives a result in expectation under a slightly more restrictive assumption (when $\sup_m L_m M_m$ is bounded), and is proved in Section B.

Corollary 2. *Under the assumptions and with the same constants and notations of Theorem 1, if there exists $A(n)$ such that*

$$\sup_{m \in \mathcal{M}} A_m \leq A(n),$$

then

$$(1 - \kappa) \mathbb{E} [\mathbf{K}_n(\tilde{p})] \leq \mathbb{E} \left[\inf_{m \in \mathcal{M}} \left((1 + \kappa) \inf_{\theta \in \Theta^{D_m}} \mathbf{K}_n(p_\theta^m) + 2 \text{pen}(m) \right) \right] + \frac{36C'}{\kappa} \Sigma A(n) \log(\varepsilon^{-1})^{3/2} \log(nA(n))^2 \frac{\log n}{n}.$$

3.3 Oracle inequality, unbounded case

The bounded case might be a bit restrictive for some applications, so we can relax the assumptions, up to additional logarithmic factors.

Theorem 3. *Assume that $n \geq 2$ and that Assumptions 1bis, 2bis and 3 hold. There exist positive numerical constants C and C' such that the following holds. For each $m \in \mathcal{M}$, let $A_m = L_m M_m + B_m$, and assume that*

$$\Sigma = \sum_{m \in \mathcal{M}} \log(A_m) e^{-D_m} < +\infty$$

Let $\kappa \in (0, 1]$. If for all $m \in \mathcal{M}$,

$$\text{pen}(m) \geq \frac{C}{\kappa} A_m^2 B_m^{3/2} \log(n)^{7/2} \log(nA_m)^2 \frac{D_m}{n},$$

then for all $x \geq 0$, with probability at least $1 - 2n^{-1} - 18 \log(n) \Sigma e^{-x}$,

$$(1 - \kappa) \mathbf{K}_n(\tilde{p}) \leq \inf_{m \in \mathcal{M}} \left((1 + \kappa) \inf_{\theta \in \Theta^{D_m}} \mathbf{K}_n(p_\theta^m) + 2 \text{pen}(m) \right) + \frac{C'}{\kappa} \left(A_m B_m^{3/2} \log(nA_m)^2 + A_{\hat{m}} B_{\hat{m}}^{3/2} \log(nA_{\hat{m}})^2 \right) \frac{(\log n)^{5/2} x}{n}.$$

Proof. See Section A. □

Note that we may replace the term $2n^{-1}$ in the probability by $2n^{-\alpha}$ for any $\alpha \in [1, n]$, provided Assumption 2bis holds with probability $n^{-\alpha}$ for a bound $\alpha L_m \|\delta - \theta\|_m \log n$. This changes the constants C and C' of Theorem 3 into $C\alpha^{7/2}$ and $C'\alpha^{5/2}$ respectively.

Except for these extra logarithmic factors, the previous result is essentially the same as Theorem 1. An expectation version of this result holds under similar assumptions on the family of models. Its proof can be found in Section B.

Corollary 4. *Under the assumptions and with the same constants and notations of Theorem 3, if there exist $A(n)$ and $B(n)$ such that*

$$\sup_{m \in \mathcal{M}} A_m \leq A(n) \quad \text{and} \quad \sup_{m \in \mathcal{M}} B_m \leq B(n),$$

then

$$(1 - \kappa) \mathbb{E}[\mathbf{K}_n(\tilde{p})] \leq \mathbb{E} \left[\inf_{m \in \mathcal{M}} \left((1 + \kappa) \inf_{\theta \in \Theta^{D_m}} \mathbf{K}_n(p_\theta^m) + 2 \text{pen}(m) \right) \right] + \frac{40C'}{\kappa} \Sigma A(n) B(n)^{3/2} \log(nA(n))^2 \frac{(\log n)^{7/2}}{n}.$$

3.4 Where does the penalty in D_m/n comes from ?

The common crucial point to all non-asymptotic controls of log-likelihood estimators [Massart, 2007, Castellan, 2003, Spokoiny, 2012, Spokoiny, 2017] lies in the control of the recentered contrast at the estimation point. In our framework, with the above notation, it means controlling

$$\nu(H) = \frac{1}{n} \sum_{t=1}^n [H_t(X_t) - \mathbb{E}(H_t(X_t) | \mathcal{F}_{t-1})],$$

where H_t is equal to

$$H_t = H_{\theta, t}^m = -\log \left(\frac{p_{\theta, t}^m(X_t)}{p_t^*(X_t)} \right).$$

As long as θ is fixed and deterministic, this is a martingale and various exponential tail bounds are applicable. Without going into details right now, let us just say that with high probability,

$$\nu(H) = \mathcal{O} \left(\sqrt{V(H)/n} \right) \tag{3}$$

with

$$V(H) = \frac{1}{n} \sum_{t=1}^n \mathbb{E}(H_t(X_t)^2 | \mathcal{F}_{t-1}).$$

In particular in the i.i.d. framework, where $V(H)$ is deterministic, of the order of the variance of H_1 , we recover what the central limit theorem entails, that is, the fluctuation of the empirical process are of order $n^{-1/2}$ multiplied by the standard deviation. In more general settings, $V(H)$ is random: it is to the bracket of the martingale. It is usually necessary to restrict oneself to an event where $V(H)$ is bounded to obtain such non asymptotic control of $\nu(H)$ (see for instance [Bercu and Touati, 2008]).

This is still not sufficient to conclude: to understand what happens for the maximum likelihood estimator, we need to control $\nu(H_{\hat{\theta}_m, t}^m)$. The fact that $\hat{\theta}_m$ depends on the whole trajectory prevents us from using classical inequalities for centered processes.

Talagrand inequality and Wilks phenomenon in the i.i.d. setting In the classical i.i.d. setting, and as a first approach, Talagrand's inequality (see [Massart, 2007] for various model selection contexts) leads to this kind of control:

$$\sup_{\theta \in \Theta_m} \frac{\nu(H_\theta^m)}{\|\theta\|} = \mathcal{O} \left(\mathbb{E} \left(\sup_{\theta \in \Theta_m, \|\theta\| \leq 1} \nu(H_\theta^m) \right) \right) + \mathcal{O} \left(\sqrt{\frac{\sup_{\theta \in \Theta_m, \|\theta\| \leq 1} V(H_\theta^m)}{n}} \right).$$

It turns out that in many models used in [Massart, 2007], $\nu(H_\theta^m)$ is in fact linear or close to linear with respect to θ , so that the first term is of order $\sqrt{D_m/n}$, the dimension of model m . In particular, this tells us non asymptotically that

$$\nu(H_{\hat{\theta}_m}^m) = \mathcal{O} \left(\|\hat{\theta}_m\| \sqrt{D_m/n} \right).$$

Beyond the i.i.d. setting, similar results can be obtained by replacing Talagrand's inequality with Baraud's inequality [Baraud, 2010].

In comparison, the Wilks phenomenon predicts that the order of magnitude of $\nu(H_{\hat{\theta}_m}^m)$ is in D_m/n , and not $\sqrt{D_m/n}$ (at least if the model m is well specified, that is there exists θ_m^* such that $p^* = p_{\theta_m^*}^m$): this non asymptotic Talagrand-like bound is pessimistic. But the trick above with the linearity with respect to θ shows us that this pessimistic bound can hopefully be multiplied by the norm of $\hat{\theta}_m$. Per se this is not much, but with a careful choice of the set over which the supremum is taken, we can replace $\|\hat{\theta}_m\|$ by the distance to a pivot θ_m^* such that, approximately,

$$\nu(H_{\hat{\theta}_m}^m) = \mathcal{O} \left(\|\hat{\theta}_m - \theta_m^*\| \sqrt{D_m/n} \right) \leq \mathcal{O} \left(\|\hat{\theta}_m - \theta_m^*\|^2 \right) + \mathcal{O} \left(\frac{D_m}{n} \right).$$

It remains to be in a nice enough setting (such as Lipschitz parameterization of $\theta \mapsto p_\theta^m$) that $\mathcal{O} \left(\|\hat{\theta}_m - \theta_m^*\|^2 \right)$ is a small fraction of the loss function we are using (here $\mathbf{K}_n(p_{\hat{\theta}_m}^m)$). All in all, up to constant, the remainder that the penalty is meant to account for is indeed in D_m/n up to multiplicative constants.

Before going further, let us make a small remark about the pivot θ_m^* . If the model m is well specified, it is easy to use the θ_m^* such that $p^* = p_{\theta_m^*}^m$. If this is not the case, usually people use the best approximation of p^* by a p_θ^m for a certain loss function (say the Kullback-Leibler divergence). So in general, and even in non i.i.d. cases, the classical way to choose a pivot is deterministic, because the loss that is used is deterministic as well.

Non asymptotic Wilks-like results In various models, the approach given above cannot work as nicely as it mainly relies on linearity, so the brute force concentration inequality "à la Talagrand" is in fact too pessimistic. What we would like, in an ideal and over-simplified world, is that (3) holds with a variance term that would be directly $V(H_{\hat{\theta}_m}^m)$, without having to take (and pay for) the supremum over all H_θ^m .

In [Boucheron and Massart, 2011], the authors restrict the previous supremum to nice small balls, so that one gets the correct behavior in D_m/n for $\nu(H_{\hat{\theta}_m}^m)$ directly. This can be especially useful when one wants to produce very sharp constants in the penalty [Arlot, 2019].

Another way to think about this, is to provide an upper function of the process, that is to prove that

$$\sup_{\theta \in \Theta_m} \left[\nu(H_\theta^m) - \mathcal{O} \left(\sqrt{V(H_\theta^m)/n} \right) \right]$$

is negative with high probability. This is Spokoiny’s approach [Spokoiny, 2012, Spokoiny, 2017], whose statistical results are the closest approach to ours for their generality, even if the author does not perform model selection per se.

Self-normalized martingales Note that it is far from obvious to obtain an inequality such as (3) when $V(H)$ is random. It is possible for martingales, in which case $V(H)$ is usually a random quantity called the bracket of the martingale. Exponential inequalities for martingales have been developed, especially for point processes where one can go from a control of the martingale with deterministic upper bound on the bracket, which already tells a lot on the properties of the maximum likelihood estimator [van de Geer, 1995], to a control where the bracket of the martingale can replace the deterministic upper bound in the deviation [Hansen et al., 2015] (up to some corrections).

In the same line, many works on self-normalization of martingales try to directly control the ratio of the martingale by its bracket [Bercu and Touati, 2008, de la Peña et al., 2009]. To our knowledge, nothing exists in this direction for a supremum of the ratios (with random renormalization) except our recent result [Aubert and Lehericy, 2024].

Deterministically renormalized empirical process From a more deterministic point of view, several works aim to choose the correct deterministic renormalization of the empirical process in an i.i.d. setting. There are two main ways to use it for model selection.

On the one hand, if the empirical process itself, once renormalized by a deterministic quantity of the form $d^2(\theta, \theta^*)$ ¹, satisfies a nice exponential inequality, then this can be chained to either directly obtain a Talagrand-like concentration on the supremum [Baraud, 2010], or used to get nice upper-functions [Spokoiny, 2012].

On the other hand, one can refine the renormalization inside the supremum and replace

$$\sup_{\theta \in \Theta_m} \frac{\nu(H_\theta^m)}{\|\theta\|} \quad \text{by} \quad \sup_{\theta \in \Theta_m} \frac{\nu(H_\theta^m)}{d^2(\theta, \theta^*) + x^2},$$

for a positive x to be chosen later. It is with this approach that [Massart, 2007] proposed a fairly general approach to model selection with penalty proportional to D_m even in non linear settings. In this setting, we still do not obtain the order of magnitude of the Wilks phenomenon directly, but because, as previously, $d^2(\theta, \theta^*)$ is close to the loss, we can still get a penalty in D_m/n . The huge advantage is that we can go further away from the linearity assumption by adding more flexibility inside the family of models thanks to the use of $d^2(., .)$ instead of the Euclidean norm.

Our method and the difference with existing works In our case, even the loss is stochastic in the most general case. Only the martingale structure remains. As a consequence, we do not have access to a deterministic $d^2(., .)$ and we cannot even properly define a nice deterministic pivot θ_m^* if the model is misspecified. However, we still have access to the bracket of the martingale $V(H_\theta)$

¹for a nice deterministic d distance on Θ_m , that can be linked to $V(H_\theta^m)$, up to possible additional terms

and it turns out that $V(H_{\hat{\theta}_m}^m)$ is nicely comparable to $\mathbf{K}_n(p_{\hat{\theta}_m}^m)$. We recently proved a concentration inequality for the supremum of stochastically normalized processes [Aubert and Lehericy, 2024], so that we can use more or less Massart’s above argument after replacing $d^2(.,.)$ by the bracket of the martingales. The problem of the pivot is solved by working directly in the space of probability and using p^* as pivot, thus bypassing the misspecification issue.

Finally, let us compare our set of assumptions to [Spokoiny, 2012, Spokoiny, 2017] which constitutes, to our knowledge, the most general non asymptotic results on maximum likelihood estimation. Spokoiny is using the Kullback-Leibler divergence as a reference. It allows him to use a deterministic pivot defined as the closest point to the truth inside the model for the deterministic distance [Spokoiny, 2012] or with additional quadratic corrections [Spokoiny, 2017]. His main assumptions rely on renormalized exponential inequalities on the gradient of the likelihood in a neighborhood of the pivot, where the normalization is quadratic. It allows him to prove a quadratic-like behavior for $\nu(H_{\hat{\theta}_m}^m)$, in this sense proving the non asymptotic Wilks phenomenon rather sharply, which we cannot do with our method.

In contrast, our method is applicable in more general settings: we do not need to use a deterministic loss, nor deterministic parametric pivots, or to assume that the log-likelihood is differentiable or that its gradient satisfies exponential inequalities. We only assume that the parameterization is Lipschitz. By martingales properties, the log-likelihood (and not its gradient) automatically satisfies (3), which is the key to prove the results on the supremum. In the end, we obtain a generalized AIC criterion (that is, a penalty proportional to D_m/n) and prove non asymptotic oracle inequalities in settings where none of the existing work applies.

4 Applications

The previous oracle inequalities are very general. This section aims to explicit how they compare to existing results. Let us look at their applications in various settings.

4.1 The i.i.d. sample case

Let us begin with the original AIC setting in its i.i.d. format. Let X_1, \dots, X_n be i.i.d. real valued random variables with density p_1^* . We consider the simple case where the filtration is generated by the observations, so that all the p_t^* are deterministic and equal to p_1^* .

We also assume that under each model, the variables X_1, \dots, X_n are i.i.d., so that $p_{\theta,t}^m = p_{\theta,1}^m$ for all t and these functions are deterministic.

In this case, $\mathbf{K}_n(p_{\theta}^m)$ is directly $\text{KL}(p_1^*d\mu, p_{\theta,1}^m d\mu)$, the Kullback-Leibler divergence between the distribution of X_1 and the distribution with density $p_{\theta,1}^m$.

Validation of the assumptions Assumption 1 is classical in this setting, see e.g. [Massart, 2007], at least for the lower bound. In [Massart, 2007] or [Castellan, 2003], there is no upper bound assumption but there is a twist: it is not the Kullback-Leibler divergence that is controlled but the Hellinger distance. Here we can also relax the lower bound assumption by using Assumption 1bis.

Assumption 2—the L_m part—can be a consequence of the fact that $p_{\theta,1}^m$ is Lipschitz with constant $L_m \varepsilon^{-1}$, or of directly assuming that $\log(p_{\theta,1}^m)$ is Lipschitz with constant L_m . It can be relaxed with Assumption 2bis. Classical asymptotic results (Wilks phenomenon or even just the consistency of the MLE) require very strong differentiability assumptions, that are not needed here. Note that the bound M_m in Assumptions 2 and 2bis entails that the models Θ_m are compact, which is assumed in general to obtain consistency in M-estimation, whether explicitly or implicitly by assuming that the estimator converges to some limit, see e.g. [van der Vaart et al., 1996].

The classical model selection "à la Birgé-Massart" for densities assumes that the models are close to linear to work, so it seems that this is incompatible with our boundedness assumption [Massart, 2007]. However this is not the case because we are forcing the parametrization to be a density that satisfies Assumption 1. To illustrate this, if we want to compare our assumptions with theirs, let us restrict ourselves to a well known case: the histogram selection [Massart, 2007, Castellan, 2003]. There, X_1, \dots, X_n are i.i.d. with density p_1^* with respect to the Lebesgue measure on $[0, 1]$ and the model m is based on a partition with D_m intervals of $[0, 1]$ of equal length. Then, for $\theta \in \Theta_m \subset \mathbb{R}^{D_m}$,

$$p_{\theta,1}^m = \sum_{I \in m} \theta_I \mathbf{1}_I.$$

Since this must be a density, the model is in fact

$$\Theta_m = \left\{ \theta = (\theta_I)_{I \in m} \text{ such that for all } I, \varepsilon \leq \theta_I \leq \varepsilon^{-1} \text{ and } \sum_{I \in m} \frac{\theta_I}{D_m} = 1 \right\}.$$

In this particular case, $p_{\theta,1}^m$ is Lipschitz with constant $L_m = 1$ w.r.t. $\|\theta\|_\infty$ and $M_m = \varepsilon^{-1} - \varepsilon$ is an upper bound of the diameter of Θ_m for this norm.

Assumption 3 is automatically fulfilled with $\mathcal{Q} = \mathbb{Q} \cap [0, 1]$, the set of rational numbers.

Result Following Theorem 1 and its Corollary, A_m does not depend on m anymore. Therefore, our oracle inequality holds as soon as

$$\text{pen}(m) = \mathcal{O} \left((\log n)^2 \frac{D_m}{n} \right).$$

Our penalty is larger than the one in [Massart, 2007] due to extra logarithmic factors and looks more like a BIC criterion. At this price we are able to prove an oracle inequality directly on the Kullback-Leibler divergence, instead of a mixed oracle inequality involving both the Kullback-Leibler divergence and the Hellinger distance.

4.2 Hidden Markov models

A hidden Markov model is a stochastic process $(H_t, X_t)_t$ where only the observations $(X_t)_t$ are observed, such that the hidden process $(H_t)_t$ is a Markov chain and such that the X_s are independent conditionally to $(H_t)_t$ with a distribution depending only on the corresponding H_s . These models have been widely used since their formalization by [Baum and Petrie, 1966], as they are able to account for complex dependencies in time processes while keeping a very simple and easily interpretable structure.

In this section, we consider finite state space hidden Markov model, in which the process $(H_t)_t$ takes values in a finite space \mathcal{H} . The parameters of a hidden Markov model are the initial distribution π and the transition kernel Q of the hidden process $(H_t)_t$ as well as the emission densities, that is the family $\nu = (\nu_h)_{h \in \mathcal{H}}$, where ν_h is the density of the distribution of X_1 conditionally to $H_1 = h$ w.r.t. the Lebesgue measure. The initial distribution π typically cannot be exactly recovered, so, in general, the parameters we wish to estimate are Q and ν .

The closest result to ours in this setting is the one from [Lehéricy, 2021], who proves an oracle inequality for a maximum likelihood estimator on misspecified hidden Markov models. Their assumptions and proofs are similar to ours, although they rely on tools that are specific to hidden Markov models to obtain their oracle inequality.

Validation of the assumptions In what follows, we only assume that the models are HMM, while the true distribution of $(X_t)_t$ may not be one. As such, we treat it separately, before introducing the models.

Concerning Assumption 1bis for the true distribution, the lower tails of $\log p_t^*(X_t)$ are always automatically sub-exponential by direct application of Markov's inequality. The control of the upper tails follows from the assumption that the conditional densities of $(X_t)_t$ admit a finite moment, that is, there exists constants $\delta > 0$ and $M_\delta > 0$ such that almost surely,

$$\sup_{t \in [n]} \mathbb{E}[p_t^*(X_t)^\delta | \mathcal{F}_{t-1}] \leq M_\delta < +\infty. \quad (4)$$

Under this assumption, there exists $B^* > 0$ depending on δ and M_δ such that for all $y \geq 1$ and $t \in [n]$,

$$\mathbb{P}(|\log p_t^*(X_t)| > B^* y | \mathcal{F}_{t-1}) \leq e^{-y}.$$

Let us now introduce the models. Let $C_Q > 0$ and $\alpha \geq 1$. For all m , let $h_m \in \mathbb{N}^*$ be the number of hidden states of model m . Let $S_m = \{g_\eta, \eta \in E_m \subset \mathbb{R}^{e_m}\}$ be a parametric set of probability densities on \mathcal{X} such that $g_\eta \geq n^{-\alpha}$ for all η . The model Θ_m is defined as the set

$$\begin{aligned} \Theta_m = \left\{ (a, q, \eta) \in \mathbb{R}^{h_m-1} \times \mathbb{R}^{h_m(h_m-1)} \times (E_m)^{h_m} \text{ s.t.} \right. \\ (C_Q \log n)^{-1} \leq h_m a(i) \leq C_Q \log n \quad \text{for all } i \in [h_m - 1], \\ (C_Q \log n)^{-1} \leq h_m q(j, i) \leq C_Q \log n \quad \text{for all } i \in [h_m - 1] \text{ and } j \in [h_m], \\ (C_Q \log n)^{-1} \leq h_m \left(1 - \sum_{j=1}^{h_m-1} a(j) \right) \leq C_Q \log n \\ \left. \text{and } (C_Q \log n)^{-1} \leq h_m \left(1 - \sum_{j=1}^{h_m-1} q(i, j) \right) \leq C_Q \log n \quad \text{for all } i \in [h_m] \right\}. \end{aligned}$$

This model is of dimension $D_m = h_m e_m + h_m^2 - 1$, and the parameter $\theta = (a, q, \eta) \in \Theta_m$ is associated

to the HMM parameters

$$\begin{cases} \pi_\theta^m(i) = a(i) & \text{for all } i \in [h_m - 1], \\ \pi_\theta^m(h_m) = 1 - \sum_{i=1}^{h_m-1} a(i), \\ Q_\theta^m(i, j) = q(i, j) & \text{for all } j \in [h_m - 1] \text{ and } i \in [h_m], \\ Q_\theta^m(i, h_m) = 1 - \sum_{j=1}^{h_m-1} q(i, j) & \text{for all } i \in [h_m], \\ \nu_{\theta,i}^m = g_{\eta_i} & \text{for all } i \in [h_m]. \end{cases}$$

The inequalities on the parameters a and q ensure that all entries of the initial distribution and transition matrices are between $(C_Q \log n)^{-1} h_m^{-1}$ and $(C_Q \log n) h_m^{-1}$. Finally, the likelihood of the observations X_1^n under the parameter $\theta \in \Theta_m$ is

$$p_\theta^m(X_1^n) = \sum_{(i_1, \dots, i_n) \in [h_m]^n} \pi_\theta^m(i_1) \nu_{\theta, i_1}^m(X_1) \prod_{t=2}^n Q_\theta^m(i_{t-1}, i_t) \nu_{\theta, i_t}^m(X_t).$$

Given the upper and lower bounds on the initial distribution and transition matrix, Assumptions 1 and 1bis can be replaced by an assumption on the average $\bar{\nu}_\theta^m = \frac{1}{h_m} \sum_{i \in [h_m]} \nu_{\theta, i}^m$ of the emission densities $(\nu_{\theta, i}^m)_{i \in [h_m]}$, since

$$\begin{aligned} p_{\theta, t}^m(X_t) &= \sum_{i, i' \in \mathcal{H}} p_\theta^m(H_{t-1} = i | X_1^{t-1}) Q_\theta^m(i, i') \nu_{\theta, i'}^m(X_t) \\ &\in [(C_Q \log n)^{-1} \bar{\nu}(X_t), (C_Q \log n) \bar{\nu}_\theta^m(X_t)]. \end{aligned}$$

Thus, to check Assumption 1bis, we will control the tails of $\bar{\nu}_\theta^m(X_t)$. Since all the emission densities are lower bounded by $n^{-\alpha}$, it is enough to assume that for all $m \in \mathcal{M}$, there exists $B'_m \geq 1$ such that for all $t \in [n]$,

$$\forall y \geq 1 \quad \mathbb{P} \left(\sup_{\theta \in \Theta_m} \log \bar{\nu}_\theta^m(X_t) > B'_m y \mid \mathcal{F}_{t-1} \right) \leq e^{-y}, \quad (5)$$

and then, if (4) holds, Assumption 1bis is satisfied for $B_m = 2(B^* \vee B'_m \vee \log(C_Q \log n) \vee (\alpha \log n))$.

Assumption 3 holds as soon as the emission densities are continuous.

Assumptions 2 or 2bis require some technical work to verify. They are used in [Aubert and Lehericy, 2024] to ensure that the entropy of the class of log-densities is controlled. A weaker sufficient assumption for hidden Markov models is introduced in [Lehericy, 2021] (Assumptions **[Aentropy]** and **[Agrowth]**) as well as their Section B.2 to see how it relates to the entropy of the class of log-likelihoods. These two assumptions, together with assuming that the mapping $\theta \in \Theta_m \mapsto (\pi_\theta^m, Q_\theta^m)$ is Lipschitz, are enough to obtain an oracle inequality. Note that the approach they use in Section B.2 also provides a way to control the regularity of the mappings $\theta \mapsto \log p_\theta^m$, in the sense that if the parameters of the HMM are Lipschitz and (4) and (5) hold, then $\theta \mapsto \log p_\theta^m$ is α -Hölder, with $\alpha \gg (\log n)^{-2-\epsilon}$ for any $\epsilon > 0$. A careful reading of our proofs and those of [Aubert and Lehericy, 2024] shows that the oracle inequalities remain valid when assuming

the mapping α -Hölder instead of Lipschitz in Assumptions 2bis and 2, up to replacing D_m by D_m/α .

Finally, the criterion \mathbf{K} used by [Lehéricy, 2021] is actually the limit of $\mathbb{E}[\mathbf{K}_n]$ when $n \rightarrow +\infty$. Due to the upper and lower bounds on the entries of the transition matrices, the models forget the past exponentially fast, in the sense that $|\log p_\theta^m(X_t|X_{t-k}^{t-1}) - \log p_\theta^m(X_t|X_{t-k'}^{t-1})| \leq \rho^{(k \wedge k')-1}/(1-\rho)$ where $\rho = 1 - (C_Q \log n)^{-2}$, see for instance [Lehéricy, 2021, Lemma 15] or [Douc et al., 2004]. If a similar property holds for the true process (such as **[A★forget]** in [Lehéricy, 2021]), then

$$|\mathbf{K} - \mathbb{E}[\mathbf{K}_n]| = \mathcal{O}((\log n)^4/n),$$

which is negligible compared to the remainder term of our oracle inequalities. Thus, we may use one or the other loss interchangeably.

Result Under the same assumptions as [Lehéricy, 2021], our results provide an oracle inequality with a penalty that is also of order $(\log n)^\alpha \frac{D_m}{n}$ for some $\alpha > 0$. The main difference is that Assumptions **[A★mixing]** and **[A★forgetting]** of [Lehéricy, 2021], a ρ -mixing assumption used to obtain concentration inequalities and a forgetting assumption used to truncate the dependencies in the past respectively, are not required to apply our results, provided the loss \mathbf{K}_n is used instead of \mathbf{K} .

4.3 Models of neuronal networks in discrete-time

Neurons are electrical cells that communicate via the emission of action potentials, also called spikes [Galves et al., 2024]. The shape of the action potential is essentially constant and the important point is the time at which the spikes are emitted. As such, the network is represented by a process $(X_t^i)_{t \in \mathbb{Z}, i \in I}$, where I is the set of all neurons constituting the network and $X_t^i = 1$ if the neuron i spikes at time t and $X_t^i = 0$ otherwise. We consider the filtration $\mathcal{F}_t = \sigma((X_s^i)_{s \leq t, i \in I})$.

Since communication between neurons is not instantaneous, most authors [Galves et al., 2024, Ost and Reynaud-Bouret, 2020] usually assume that conditionally to \mathcal{F}_{t-1} , the $(X_t^i)_{i \in I}$ are independent, so that the whole activity can be described by just giving the $p_t^{*,i}$'s with

$$\forall i \in I, \forall t \in \mathbb{Z}, \quad p_t^{*,i} = \mathbb{P}(X_t^i = 1 | \mathcal{F}_{t-1}).$$

We assume the process to be stationary.

One of the main neuronal model in discrete-time is the discrete Hawkes process [Ost and Reynaud-Bouret, 2020], which can be modeled by

$$p_t^{i,H} = \phi_i \left(\sum_{j \in I} \sum_{s < t} h_{j \rightarrow i}(t-s) X_s^j \right),$$

where ϕ_i is a rate function that is usually assumed to be Lipschitz, increasing and taking values in $[0, 1]$, and where $h_{j \rightarrow i}$ are interaction functions: if it is positive at delay δ , neuron j excites neuron i after a delay δ ; if it is negative, neuron j inhibits neuron i after a delay δ . For instance, the linear case is the situation where $\phi_i(x) = \mu_i + x$. In the sequel, to simplify, $\phi_i(\cdot)$ is supposed to be fixed and known. Only the functions $h_{j \rightarrow i}$ are unknown.

The Galves-Löcherbach neuronal model [Galves et al., 2024, Galves and Löcherbach, 2013] is slightly different, here

$$p_t^{i,GL} = \phi_i \left(\sum_{j \in I} \sum_{s=L_t^i}^{t-1} h_{j \rightarrow i}(t-s) X_s^j \right),$$

where L_t^i is the time of the last spike of neuron i . In contrast to the Hawkes process, the neurons of this model essentially reset their memory each time they spike.

In practice, only a small finite subset F of I is observed on a finite time duration, say $t = -A, \dots, n$, for some positive A . For a fixed $i \in F$, we are interested by estimating $(p_t^{*,i})_{t \in [n]}$ based on the observations of $(X_t^j)_{-A \leq t \leq n, j \in F}$.

We are interested in a specific neuron of interest $i \in F$, so the process $(X_t)_{t \in [n]}$ from our oracle inequalities is taken to be $(X_t^i)_{t=1, \dots, n}$. The filtration \mathcal{F}_t is the one defined above and generated by the whole network. Finally, in order to define the models, we have access to more information than $(X_t^i)_{t=1, \dots, n}$ but less than the whole network: we may only use the observations $(X_t^j)_{j \in F, t=-A, \dots, n}$, which are indeed \mathcal{F}_t adapted.

Whatever the neuronal model (p^H for Hawkes or p^{GL} for Galves-Löcherbach) that we choose, we need to parameterize it. We define model m by choosing a finite subset of F , called V_m , which is the proposed neighborhood for neuron i in model m , and by choosing $A_m \leq A$ a maximal lag of interaction. In model m , all the $h_{j \rightarrow i}(u)$ are null if $j \notin V_m$ or $u > A_m$, so the model is parameterized by $(\theta_{j,u} := h_{j \rightarrow i}(u))_{j \in V_m, u=1, \dots, A_m} \in \mathbb{R}^{A_m |V_m|}$. Given this parameterization, the conditional distributions are defined by

$$p_t^{i,m,H} = \phi_i \left(\sum_{j \in V_m} \sum_{u=1}^{A_m} \theta_{j,u} X_{t-u}^j \right)$$

and

$$p_t^{i,m,GL} = \phi_i \left(\sum_{j \in V_m} \sum_{u=1}^{\min(A_m, t-L_t^i)} \theta_{j,u} X_{t-u}^j \right).$$

Validation of the assumptions Assumption 1 means that $p_t^{*,i}$ as well as $p_t^{i,m,H}$ or $p_t^{i,m,GL}$ are in $[\varepsilon, 1 - \varepsilon]$. This is a very common assumption in these settings (see [Duarte et al., 2019, Ost and Reynaud-Bouret, 2020]). In this sense, Assumption 1bis can be seen as a relaxation with respect to previous works.

The assumption that ϕ_i is Lipschitz with constant L is a very classical one [Duarte et al., 2019]. Together with Assumption 1, it implies that $\log(\phi_i)$ (probability of a spike) and $\log(1 - \phi_i)$ (probability of no spike) are Lipschitz with constant $2\varepsilon^{-1}L$. Thus, the first part of Assumption 2 is satisfied with $L_m = 2\varepsilon^{-1}L$ and the ℓ_1 norm $\|\theta\|_1$.

For the second part of Assumption 2, it depends on ϕ_i . Indeed, since ϕ_i is increasing, we can define

$$\Theta_m = \left\{ \theta \in \mathbb{R}^{A_m |V_m|} \text{ such that } \varepsilon \leq \phi_i \left(\sum_{j,u} \theta_{j,u} \mathbf{1}_{\theta_{j,u} < 0} \right) \text{ and } \phi_i \left(\sum_{j,u} \theta_{j,u} \mathbf{1}_{\theta_{j,u} > 0} \right) \leq 1 - \varepsilon \right\},$$

to ensure that Assumption 1 is satisfied. Since ϕ_i is increasing and Lipschitz, we can define its inverse, so that if ε and $1 - \varepsilon$ are possible values for ϕ_i (as is typically the case for linear or sigmoid functions) then it automatically follows that for all $\theta \in \Theta_m$

$$\|\theta\|_1 \leq |\phi_i^{-1}(\varepsilon)| + |\phi_i^{-1}(1 - \varepsilon)|,$$

and so the second part of Assumption 2 is satisfied with $M_m = |\phi_i^{-1}(\varepsilon)| + |\phi_i^{-1}(1 - \varepsilon)|$.

Assumption 3 is automatically fulfilled in each of the models because $p_t^{i,m,H}$ and $p_t^{i,m,GL}$ only depend on a finite set of X_s^j . Assumption 3 for p^* can be solved by assuming the following continuity assumption, which is standard in this setting (see [Duarte et al., 2019, Galves et al., 2024]). Let x be a past configuration, i.e. a possible value for $(X_s^j)_{j \in I, s < t}$, and let us remark that by stationarity, $p_t^{*,i}$ can be seen as a function of x and not of t :

$$p^{*,i}(x) = \mathbb{P}(X_t^i = 1 | (X_s^j)_{j \in I, s < t} = x).$$

The continuity assumption of the neuronal model assumes that there exists a nested sequence $(S_k)_{k \geq 1}$ of finite subsets of $I \times \mathbb{Z}_-$ such that $S_k \xrightarrow[k \rightarrow \infty]{} I \times \mathbb{Z}$ and such that

$$\sup\{|p_t^{*,i}(x) - p_t^{*,i}(y)| \text{ such that } x|_{S_k} = y|_{S_k}\} \xrightarrow[k \rightarrow \infty]{} 0,$$

where $x|_{S_k}$ is the configuration restricted to the indices in S_k . Informally, this continuity assumption states that one can approximate $p^{*,i}(x)$ by what happens on a finite number of x_s^j 's.

Result Following Theorem 1 or its corollary, and since A_m does not depend on m , one can take

$$\text{pen}(m) = \mathcal{O}\left(\log(n)^2 \frac{D_m}{n}\right).$$

Note that under mild conditions (see for instance [Galves et al., 2024] or [Ost and Reynaud-Bouret, 2020]), these processes are stationary and $\mathbb{E}(\mathbf{K}_n(p_\theta^m))$ does not depend on n .

Furthermore, since $-\log(p_{\theta,t}^m/p_t^*)$ is upper and lower bounded, one can easily show that

$$\mathbb{E}\left(\log^2 \frac{p_t^*(X_t)}{p_t(X_t)} \middle| \mathcal{F}_{t-1}\right) \lesssim \mathbb{E}\left[\log \frac{p_t^*(X_t)}{p_t(X_t)} \middle| \mathcal{F}_{t-1}\right] \lesssim \mathbb{E}\left(\log^2 \frac{p_t^*(X_t)}{p_t(X_t)} \middle| \mathcal{F}_{t-1}\right), \quad (6)$$

where \lesssim means that the inequality holds up to positive multiplicative constant.

Moreover, whatever the neuronal model (Hawkes or Galves-Löcherbach), we can expand $\theta^m \in \Theta^m$ with zeroes so that it is defined on $I \times \mathbb{N}^*$. If ϕ_i has a derivative that is upper and lower bounded by some positive constant, so do $\log(\phi_i)$ and $\log(1 - \phi_i)$, and therefore, whatever the value of X_t ,

$$\left| \sum_{j \in I} \sum_{s < t} (h_{j \rightarrow i}(t - s) - \theta_{j,t-s}^m) X_s^j \right| \lesssim \left| \log \left(\frac{p_{\theta^m,t}^m(X_t)}{p_t^*(X_t)} \right) \right| \lesssim \left| \sum_{j \in I} \sum_{s < t} (h_{j \rightarrow i}(t - s) - \theta_{j,t-s}^m) X_s^j \right|$$

(for the Hawkes case and with a restricted sum in the lower bound in the Galves-Löcherbach case). Note that both upper and lower bounds are \mathcal{F}_{t-1} measurable. Hence, going back to the oracle

inequality, we can express both the upper bound and the lower bound in terms of the average square distance

$$\frac{1}{n} \sum_{t=1}^n \left| \sum_{j \in I} \sum_{s < t} (h_{j \rightarrow i}(t-s) - \theta_{j,t-s}^m) X_s^j \right|^2$$

and the upper bound of the oracle inequality is a trade-off between the bias measured by the average square distance above and the penalty in D_m/n up to logarithmic terms.

Let us compare this result to the ones in [Ost and Reynaud-Bouret, 2020] for Hawkes and Galves-Löcherbach process. In [Ost and Reynaud-Bouret, 2020], the authors could only envision linear models (i.e. ϕ_i is linear) and were using least-square loss on the p . They can perform variable selection thanks to an ℓ_1 penalization, whereas in our case the summability condition on Σ makes it impossible to perform variable selection by considering the full set of subsets of variables. Despite this difference, their oracle inequality is for the exact same average square distance as mentioned above with, in the upper bound, a trade-off between the bias and a term in D_m/n up to logarithmic terms. However, the dimension D_m in their case, had to be smaller than a given a priori level of sparsity s . Moreover, their constant in front of D_m was given by an RE inequality on the Gram matrix. In the most general case considered by [Ost and Reynaud-Bouret, 2020], this constant depend on the size of the observed network F and explodes with the size of F ; in the Hawkes case, this bound depends on s , and explodes for moderate s , whereas in our case the penalty can handle large D_m thanks to the summability condition in Σ which ensures that the number of models remains reasonable.

For the Galves-Löcherbach model, we can also compare this result with the ones of [Duarte et al., 2019]. In [Duarte et al., 2019] (see also [De Santis et al., 2022] for recent improvements), the authors envision the estimation of the interaction neighborhood of a neuron i . In these two articles, the authors assume that the set of observed neurons F contains the interaction neighborhood of i . In other words, with their assumptions on the process, at least one model is well specified. In this case, our result shows that we can estimate the conditional probability of spiking for i with parametric rate in D_m/n up to logarithmic terms, as the Wilks phenomenon would predict if it was applicable in this case, but without even knowing which model is the true one up front. However, it is not clear if this means that the neighborhood of interaction is correctly estimated by $V_{\hat{m}}$. We cannot manage variable selection per se because the complexity of the family of models would be too large for our general result. But we can at least hope that $V_{\hat{m}}$ would contain the true interaction neighborhood of neuron i with high probability ². On the other hand, our procedure is much less computationally intensive than theirs and does not require the precise examination of patterns of 0 and 1's as they do, which in practice usually prevents them to use it with more than 4 observed neurons [Brochini et al., 2016]. As such, our method could at least be used to restrict the set of observed neurons before using the methods developed in [Duarte et al., 2019, De Santis et al., 2022].

4.4 Adversarial multi-armed bandits as a cognitive learning model

All the previous applications were in the stationary case, so that we were able to compare our results with existing ones, even if per se our oracle inequalities do apply even if we do not assume stationarity. Let us now look at a setup which cannot be stationary: learning data, in

²This would probably require additional assumptions, such as minimal strength of the interaction inside the neighborhood as in [Duarte et al., 2019] and [De Santis et al., 2022].

which we aim to estimate how an individual or system learns to perform a task by observing their training as it takes place. In such a setting, the data that are produced cannot, by essence, be independent or stationary. If in practice many authors have advertised for the use of MLE [Daw, 2011, Wilson and Collins, 2019], this problem was for the first time studied theoretically in [Aubert et al., 2023]. In this previous work, only the property of the MLE on one model is studied, without the model selection part.

The model used in [Aubert et al., 2023] is the **Exp3** algorithm [Auer et al., 2002]. In the Machine Learning literature, this algorithm was originally proposed to solve an the adversarial multi-armed bandit problem [Bubeck and Cesa-Bianchi, 2012], [Lattimore and Szepesvári, 2020], that is a game played sequentially between a learner and an adversary (the environment), where at each round the learner must choose an action k from a set of actions $[K]$ for some integer $K \geq 1$ and the adversary returns a loss $g_{k,t}$ for this action.

Algorithm 1 Exp3 with learning rate η

Initialization: $p_{\eta,1} = (\frac{1}{K}, \dots, \frac{1}{K})$.

for $t = 1, 2, \dots$ **do**

 Draw an action $X_t \sim p_{\eta,t}$ and if $X_t = k$ receive a loss $g_{t,k} \in [0, 1]$.

 Update for all $k \in [K]$:

$$p_{\eta,t+1}(k) = \frac{\exp\left(-\eta \sum_{s=1}^t \hat{L}_{\eta,s}(k)\right)}{\sum_{j=1}^K \exp\left(-\eta \sum_{s=1}^t \hat{L}_{\eta,s}(j)\right)} \quad \text{where } \hat{L}_{\eta,s}(j) = \frac{g_{j,s}}{p_{\eta,s}(j)} \mathbf{1}_{X_s=j}$$

The fact that the learner plays as if the environment is adversarial makes it a realistic model for cognitive processes where humans or animals need to adapt to a changing environment. However in a cognitive experiment, most of the time, the environment is not adversarial even if the learner does not know that and still uses an adversarial strategy.

Comparison with [Aubert et al., 2023] [Aubert et al., 2023] consider only cases where $g_{k,t} = g_k$ depends only on k , which is common in cognitive experiments. It has been shown that if the learning rate θ is fixed, no estimator can achieve polynomial rates of convergence. This is mainly due to the fact that $p_{\eta,t}(k)$ can rapidly become absurdly small, and as a consequence, only an extremely small number of X_t 's are actually relevant to perform the estimation. For instance, if an individual has finished learning and no longer makes mistakes, then it is impossible to improve our estimation of their learning behaviour even with continued observation. Therefore, in [Aubert et al., 2023], another asymptotic in T is proposed. The goal is to estimate $\theta \in [r, R]$ such that $\eta = \theta/\sqrt{T}$ based on the first

$$T_\varepsilon = \left\lfloor \left(\frac{1}{K} - \varepsilon \right) \frac{\sqrt{T}}{R} \right\rfloor,$$

observations. Indeed up to T_ε , all the $p_{\frac{\theta}{\sqrt{T}},t}(k)$ are larger than ε , a fixed constant in $(0, 1)$ meant to prevent the vanishing of the various probabilities. The choice of the rate \sqrt{T} can be relaxed to other powers of T but for the present illustrative purpose, we only consider this case, which also

corresponds to the renormalization of the learning rate for which **Exp3** satisfies sublinear regret bounds [Bubeck and Cesa-Bianchi, 2012].

With this renormalization, [Aubert et al., 2023] proved a convergence in $T_\varepsilon^{-1/2}$ for

$$\frac{1}{T_\varepsilon} \sum_{t=1}^{T_\varepsilon} \sum_{k=1}^K (p_{\frac{\theta}{\sqrt{T}}, t}(k) - p_{\frac{\hat{\theta}}{\sqrt{T}}, t}(k))^2,$$

with θ the true parameter and $\hat{\theta}$ the MLE.

Let us compare this with Theorem 1 or its corollary in the case of a single, well specified, model. Take $\mathcal{F}_t = \sigma(X_1^t)$ and $n = T_\varepsilon$, so that Assumption 1 is directly satisfied. Since log is Lipschitz on $[\varepsilon, 1 - \varepsilon]$, Assumption 2 follows as a direct consequence of Lemma 4.3 of [Aubert et al., 2023], with respect to $\|\theta\|_\infty$, with a constant L that only depends on R and ε . Moreover we can always take $M = R$. Finally, Assumption 3 is trivial since there is only a finite set of values for X_t .

Thus, Theorem1 shows that $\mathbf{K}_n(\tilde{p})$ is of order $\log(T_\varepsilon)^2/T_\varepsilon$, which is faster than [Aubert et al., 2023]. Indeed $\mathbf{K}_n(p)$ is comparable to the square norm used in [Aubert et al., 2023], which can be proved as follows:

$$\mathbf{K}_n(p) = \frac{1}{T_\varepsilon} \sum_{t=1}^{\varepsilon} \sum_{k=1}^K p_t^*(k) \log \frac{p_t^*(k)}{p_t(k)} = \frac{1}{T_\varepsilon} \sum_{t=1}^{\varepsilon} \sum_{k=1}^K p_t^*(k) \phi(\log \frac{p_t(k)}{p_t^*(k)}),$$

with $\phi(u) = e^u - u - 1$. Since $|\log \frac{p_t^*(k)}{p_t(k)}|$ is bounded, $\phi(u) \gtrsim u^2$. This leads to

$$\mathbf{K}_n(p) \gtrsim \frac{1}{T_\varepsilon} \sum_{t=1}^{\varepsilon} \sum_{k=1}^K p_t^*(k) \log^2 \frac{p_t^*(k)}{p_t(k)} \gtrsim \frac{1}{T_\varepsilon} \sum_{t=1}^{\varepsilon} \sum_{k=1}^K (p_t^*(k) - p_t(k))^2,$$

since the derivative of the logarithm is positive and lower bounded on $[\varepsilon, 1 - \varepsilon]$.

A model selection framework Theorem 1 or its corollary goes further than [Aubert et al., 2023] by also allowing for model selection and bias. In particular, this can allow for reward estimation as well as identifying the granularity or precision the learner is able to have in its actions. This is of particular importance when performing a cognitive learning experiment, since it is rarely clear what numerical value to assign to a real-life reward to reflect how much of an impact it has on the learner (for instance, how much is crab meat worth to an octopus in comparison to shrimp?), nor how precise the learner is able to be when picking an action, in the case where the number of possible actions is continuous or very large (since it is impossible to perfectly replicate a given action; near identical actions should provide near identical rewards, so how dissimilar do actions have to be to be considered distinct?).

One possibility is to use **Exp3** with a different parameterization. For a given model m , consider a partition \mathcal{I}^m of the set of possible actions (which may even be continuous). Let D_m be the number of disjoint sets in \mathcal{I}^m . The parameters of model m are the rewards of each possible set I of the partition, that we call g_I^m . Since the role of the learning parameter η introduced in Algorithm 1 is redundant with the one of the (unknown) rewards g 's, we remove it from the model. Thus, under model m , the learner proceeds as follows:

- Initialize $p_1^m = (1/D_m, \dots, 1/D_m)$

- for $t \geq 1$,
 - pick the set I_t according to p_t^m and pick X_t inside I_t uniformly
 - update for all $I \in \mathcal{I}^m$

$$p_{t+1}^m(I) = \frac{\exp\left(-\sum_{s=1}^t \hat{L}_s^m(I)\right)}{\sum_{J=1}^{D_m} \exp\left(-\sum_{s=1}^t \hat{L}_s^m(J)\right)} \quad \text{where } \hat{L}_s^m(J) = \frac{g_J^m}{p_{\eta,s}(j)} \mathbf{1}_{X_s \in J}$$

To obtain a meaningful renormalization, as before, we call $p_{\theta^m,t}^m$ the distribution of X_t given $\mathcal{F}_{t-1} = \sigma(X_1^{t-1})$, when $g_J^m = \frac{\theta_J^m}{\sqrt{T}}$.

As in the previous single model case, Assumption 1 is straightforward by choosing n of the order of \sqrt{T} up to multiplicative constants. We can also adapt Lemma 4.3 of [Aubert et al., 2023] to show that Assumption 2 is satisfied as long as all the θ_J^m 's are in $[r, R]$, with a Lipschitz constant L with respect to $\|\theta\|_\infty$ which does not depend on m . One can also take $M_m = R$. Assumption 3 is again easily satisfied even if the set of actions is continuous.

Therefore, one can choose a penalty of order

$$\text{pen}(m) = \mathcal{O}\left(\log(T_\varepsilon)^2 \frac{D_m}{T_\varepsilon}\right),$$

and obtain an oracle inequality on $\mathbf{K}_n(\tilde{p})$, which as seen above can be implies an oracle inequality on the square distance. In the end, we obtain a way to estimate at the same time

- with the partition $\mathcal{I}_{\hat{m}}$: the precision in the perception and execution of the actions, that is, which sets of actions are conflated and which are considered distinct, and how precise the learner is able to be when choosing its actions,
- with the resulting $\hat{\theta}_J^{\hat{m}}/\sqrt{T_\varepsilon}$: the estimated reward, by which we mean the numerical value that quantifies the average impact of the outcome of the action on the learner's behaviour, modeled by a piecewise constant function on the partition $\mathcal{I}_{\hat{m}}$.

Acknowledgments

This work was supported by the French government, through the UCA^{Jedi} and 3IA Côte d'Azur Investissements d'Avenir managed by the National Research Agency (ANR-15- IDEX-01 and ANR-19-P3IA-0002), by the interdisciplinary Institute for Modeling in Neuroscience and Cognition (NeuroMod) of the Université Côte d'Azur and directly by the National Research Agency ANR-08-JCJC-0125-01, ANR-19-CE40-0024 with the ChaMaNe project, MITI AAP "défi modélisation du vivant" with DYNAMO project and AAP 80PRIME with eXpLAIIn. It is part of the ComputaBrain project.

References

- [Akaike, 1973] Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. *2nd international symposium on information theory, Akademia Kiado, Budapest*, page 267–281.

- [Arlot, 2019] Arlot, S. (2019). Minimal penalties and the slope heuristics: a survey. *J. SFdS*, 160(3):158–168.
- [Aubert and Lehericy, 2024] Aubert, J. and Lehericy, L. (2024). Exponential inequalities for suprema of processes with stochastic normalization. preprint.
- [Aubert et al., 2023] Aubert, J., Lehericy, L., and Reynaud-Bouret, P. (2023). On the convergence of the MLE as an estimator of the learning rate in the exp3 algorithm. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 1244–1275. PMLR.
- [Auer et al., 2002] Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The non-stochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77.
- [Baraud, 2010] Baraud, Y. (2010). A Bernstein-type inequality for suprema of random processes with applications to model selection in non-Gaussian regression. *Bernoulli*, 16(4):1064 – 1085.
- [Baraud and Birgé, 2016] Baraud, Y. and Birgé, L. (2016). Rho-estimators for shape restricted density estimation. *Stochastic Process. Appl.*, 126(12):3888–3912.
- [Baum and Petrie, 1966] Baum, L. E. and Petrie, T. (1966). Statistical inference for probabilistic functions of finite state markov chains. *The annals of mathematical statistics*, 37(6):1554–1563.
- [Bercu et al., 2015] Bercu, B., Delyon, B., Rio, E., et al. (2015). *Concentration inequalities for sums and martingales*. Springer.
- [Bercu and Touati, 2008] Bercu, B. and Touati, A. (2008). Exponential inequalities for self-normalized martingales with applications. *Ann. Appl. Probab.*, 18(5):1848–1869.
- [Bercu and Touati, 2019] Bercu, B. and Touati, T. (2019). New insights on concentration inequalities for self-normalized martingales. *Electronic Communications in Probability*, 24(none):1 – 12.
- [Bickel and Doksum, 2015] Bickel, P. J. and Doksum, K. A. (2015). *Mathematical statistics—basic ideas and selected topics. Vol. 1*. Texts in Statistical Science Series. CRC Press, Boca Raton, FL, second edition.
- [Birgé and Massart, 2001] Birgé, L. and Massart, P. (2001). Gaussian model selection. *J. Eur. Math. Soc. (JEMS)*, 3(3):203–268.
- [Boucheron and Massart, 2011] Boucheron, S. and Massart, P. (2011). A high-dimensional Wilks phenomenon. *Probab. Theory Related Fields*, 150(3-4):405–433.
- [Brochini et al., 2016] Brochini, L., Galves, A., Hodara, P., Ost, G., and Pouzat, C. (2016). Estimation of neuronal interaction graph from spike train data. *arXiv preprint arXiv:1612.05226*.
- [Bubeck and Cesa-Bianchi, 2012] Bubeck, S. and Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *CoRR*, abs/1204.5721.
- [Burnham and Anderson, 2004] Burnham, K. P. and Anderson, D. R. (2004). Multimodel inference: understanding AIC and BIC in model selection. *Sociol. Methods Res.*, 33(2):261–304.
- [Castellan, 2003] Castellan, G. (2003). Density estimation via exponential model selection. *IEEE Trans. Inform. Theory*, 49(8):2052–2060.

- [Cox, 1975] Cox, D. R. (1975). Partial likelihood. *Biometrika*, 62(2):269–276.
- [Daw, 2011] Daw, N. D. (2011). Trial-by-trial data analysis using computational models. *Decision making, affect, and learning: Attention and performance XXIII*, 23(1).
- [De Castro et al., 2016] De Castro, Y., Gassiat, É., and Lacour, C. (2016). Minimax adaptive estimation of nonparametric hidden Markov models. *Journal of Machine Learning Research*, 17:1–43.
- [de la Peña et al., 2004] de la Peña, V. H., Klass, M. J., and Leung Lai, T. (2004). Self-normalized processes: exponential inequalities, moment bounds and iterated logarithm laws. *The Annals of Probability*, 32(3).
- [de la Peña et al., 2009] de la Peña, V. H., Lai, T. L., and Shao, Q.-M. (2009). *Self-normalized processes*. Probability and its Applications (New York). Springer-Verlag, Berlin. Limit theory and statistical applications.
- [De Santis et al., 2022] De Santis, E., Galves, A., Nappo, G., and Piccioni, M. (2022). Estimating the interaction graph of stochastic neuronal dynamics by observing only pairs of neurons. *Stochastic Processes and their Applications*, 149:224–247.
- [Douc et al., 2004] Douc, R., Moulines, É., and Rydén, T. (2004). Asymptotic properties of the maximum likelihood estimator in autoregressive models with Markov regime. *Annals of Statistics*, pages 2254–2304.
- [Duarte et al., 2019] Duarte, A., Galves, A., Löcherbach, E., and Ost, G. (2019). Estimating the interaction graph of stochastic neural dynamics. *Bernoulli*, 25(1):771–792.
- [Galves and Löcherbach, 2013] Galves, A. and Löcherbach, E. (2013). Infinite systems of interacting chains with memory of variable length—a stochastic model for biological neural nets. *Journal of Statistical Physics*, 151(5).
- [Galves et al., 2024] Galves, A., Löcherbach, E., and Pouzat, C. (2024). *Probabilistic Spiking Neuronal Nets – Data, Models and Theorems*. Springer.
- [Ghosal and van der Vaart, 2007] Ghosal, S. and van der Vaart, A. (2007). Posterior convergence rates of dirichlet mixtures at smooth densities. *Ann. Statist.*, 35(1):697–723.
- [Hansen et al., 2015] Hansen, N. R., Reynaud-Bouret, P., and Rivoirard, V. (2015). Lasso and probabilistic inequalities for multivariate point processes. *Bernoulli*, 21(1).
- [Lacour, 2008a] Lacour, C. (2008a). Adaptive estimation of the transition density of a particular hidden Markov chain. *J. Multivariate Anal.*, 99(5):787–814.
- [Lacour, 2008b] Lacour, C. (2008b). Nonparametric estimation of the stationary density and the transition density of a Markov chain. *Stochastic Process. Appl.*, 118(2):232–260.
- [Lattimore and Szepesvári, 2020] Lattimore, T. and Szepesvári, C. (2020). *Bandit Algorithms*. Cambridge University Press.
- [Lehéricy, 2021] Lehéricy, L. (2021). Nonasymptotic control of the mle for misspecified nonparametric hidden markov models. *Electronic Journal of Statistics*, 15(2):4916–4965.

- [Lehmann, 1999] Lehmann, E. L. (1999). *Elements of large-sample theory*. Springer Texts in Statistics. Springer-Verlag, New York.
- [Massart, 2007] Massart, P. (2007). *Concentration inequalities and model selection: Ecole d'Eté de Probabilités de Saint-Flour XXXIII-2003*. Springer.
- [Ost and Reynaud-Bouret, 2020] Ost, G. and Reynaud-Bouret, P. (2020). Sparse space-time models: concentration inequalities and Lasso. *Ann. Inst. Henri Poincaré Probab. Stat.*, 56(4):2377–2405.
- [Ramponi et al., 2020] Ramponi, G., Drappo, G., and Restelli, M. (2020). Inverse reinforcement learning from a gradient-based learner. *Advances in Neural Information Processing Systems*, 33:2458–2468.
- [Shen et al., 2013] Shen, W., Tokdar, S. T., and Ghosal, S. (2013). Adaptive bayesian multivariate density estimation with dirichlet mixtures. *Biometrika*, 100(3):623–640.
- [Spokoiny, 2012] Spokoiny, V. (2012). Parametric estimation. Finite sample theory. *Ann. Statist.*, 40(6):2877–2909.
- [Spokoiny, 2017] Spokoiny, V. (2017). Penalized maximum likelihood estimation and effective dimension. *Ann. Inst. Henri Poincaré Probab. Stat.*, 53(1):389–429.
- [Sur and Candès, 2019] Sur, P. and Candès, E. J. (2019). A modern maximum-likelihood theory for high-dimensional logistic regression. *Proc. Natl. Acad. Sci. USA*, 116(29):14516–14525.
- [Sur et al., 2019] Sur, P., Chen, Y., and Candès, E. J. (2019). The likelihood ratio test in high-dimensional logistic regression is asymptotically a rescaled chi-square. *Probab. Theory Related Fields*, 175(1-2):487–558.
- [Talagrand, 1996] Talagrand, M. (1996). New concentration inequalities in product spaces. *Invent. Math.*, 126(3):505–563.
- [Tang and Leng, 2010] Tang, C. Y. and Leng, C. (2010). Penalized high-dimensional empirical likelihood. *Biometrika*, 97(4):905–919. With supplementary material available online.
- [van de Geer, 1995] van de Geer, S. (1995). Exponential inequalities for martingales, with application to maximum likelihood estimation for counting processes. *Ann. Statist.*, 23(5):1779–1801.
- [van der Vaart et al., 1996] van der Vaart, A. W., Wellner, J. A., van der Vaart, A. W., and Wellner, J. A. (1996). M-estimators. *Weak Convergence and Empirical Processes: With Applications to Statistics*, pages 284–308.
- [Wilks, 1938] Wilks, S. (1938). The large-sample distribution of the likelihood ratio for composite testing hypotheses. *Ann. Math. Stat.*, pages 60–62.
- [Wilson and Collins, 2019] Wilson, R. C. and Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *Elife*, 8:e49547.
- [Zheng et al., 2019] Zheng, C., Ferrari, D., and Yang, Y. (2019). Model selection confidence sets by likelihood ratio testing. *Statist. Sinica*, 29(2):827–851.

A Proof of Theorems 1 and 3

Both theorems share the same proof, up to differences in notations due to picking Assumptions 1 and 2 or Assumptions 1bis and 2bis, which will be introduced when required.

In the sequel to avoid losing track of important dependance, we denote: for any $t \in [n]$, $m \in \mathcal{M}$, $\theta \in \Theta_m$, $x_t \mapsto p_{\theta,t}^m(x_t | \mathcal{F}_{t-1})$ the density of X_t conditionally to \mathcal{F}_{t-1} under the parameter θ , and likewise for $p_t^*(x_t | \mathcal{F}_{t-1})$.

In what follows, fix $m \in \mathcal{M}$ and $\bar{\theta}^m \in \Theta_m$. Consider the following functions, defined for all $t \in [n]$, $x_1^t \in \mathcal{X}^t$, $m' \in \mathcal{M}$ and $\delta \in \Theta_{m'}$ by

$$g_{\delta,t}^{m'}(x_t, \mathcal{F}_{t-1}) = -\log \left(\frac{p_{\delta,t}^{m'}(x_t | \mathcal{F}_{t-1})}{p_t^*(x_t | \mathcal{F}_{t-1})} \right)$$

and write $g_\delta^{m'} = (g_{\delta,t}^{m'})_{t \in [n]}$.

For all $m' \in \mathcal{M}$, let $\hat{\theta}^{m'}$ be a maximizer of $\theta \in \Theta_{m'} \mapsto \frac{1}{n} \ell_n(\theta)$, and let $\text{crit}(m')$ be

$$\text{crit}(m') = -\frac{1}{n} \ell_n(\hat{\theta}^{m'}) + \text{pen}(m'),$$

and define the set \mathcal{M}' as

$$\mathcal{M}' = \{m' \in \mathcal{M}, \quad \text{crit}(m') \leq \text{crit}(m)\}.$$

For any family $h = (h_t)_{t \in [n]}$ of functions X_t that may depend on the past, that is $h_t(X_t, \mathcal{F}_{t-1})$, we write

$$\begin{cases} P(h) = \frac{1}{n} \sum_{t=1}^n h_t(X_t, \mathcal{F}_{t-1}), \\ C(h) = \frac{1}{n} \sum_{t=1}^n \mathbb{E}[h_t(X_t, \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}] \quad \text{the compensator of } P(h), \\ \nu(h) = P(h) - C(h) = \frac{1}{n} \sum_{t=1}^n (h_t(X_t, \mathcal{F}_{t-1}) - \mathbb{E}[h_t(X_t, \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}]). \end{cases}$$

By definition, for every $m' \in \mathcal{M}'$,

$$P(g_{\hat{\theta}^{m'}}^{m'}) + \text{pen}(m') \leq P(g_{\hat{\theta}^m}^m) + \text{pen}(m) \leq P(g_{\hat{\theta}^m}^m) + \text{pen}(m).$$

Therefore, since $P = C + \nu$, for every $m' \in \mathcal{M}'$,

$$C(g_{\hat{\theta}^{m'}}^{m'}) + \nu(g_{\hat{\theta}^{m'}}^{m'}) \leq C(g_{\hat{\theta}^m}^m) + \text{pen}(m) + \nu(g_{\hat{\theta}^m}^m) - \text{pen}(m').$$

Plugging the definition of $g_{\hat{\theta}^m}^m$ in the above leads to

$$\mathbf{K}_n(p_{\hat{\theta}^{m'}}^{m'}) \leq \mathbf{K}_n(p_{\hat{\theta}^m}^m) + \text{pen}(m) - \nu(g_{\hat{\theta}^{m'}}^{m'}) - \text{pen}(m') + \nu(g_{\hat{\theta}^m}^m). \quad (7)$$

So far everything is similar to [Massart, 2007]. The goal is now to control $-\nu(g_{\hat{\theta}^{m'}}^{m'})$.

Let us distinguish between the two sets of assumptions. Since the core of the proof is identical in both cases, we introduce notations to encompass both cases in a single setting.

- Under Assumptions 1 and 2, for each $t \in [n]$ and $m'' \in \mathcal{M}$, let $q_n = 0$, $F_{m''}^\infty = 2 \log(\varepsilon^{-1})$ and $F_{m''}^{\text{lip}} = L_{m''}$, and for any $\delta \in \Theta_{m''}$, let

$$\mathbf{V}_n(p_\delta^{m''}) = \frac{1}{n} \sum_{t=1}^n \mathbb{E} \left[\left(\log \frac{p_t^*(X_t | \mathcal{F}_{t-1})}{p_{\delta,t}^{m''}(X_t | \mathcal{F}_{t-1})} \right)^2 \middle| \mathcal{F}_{t-1} \right].$$

- Under Assumptions 1bis and 2bis, for each $t \in [n]$ and $m'' \in \mathcal{M}$, let $q_n = 2n^{-1}$, $F_{m''}^\infty = B_{m''} \log n$ and $F_{m''}^{\text{lip}} = L_{m''} \log n$, and for any $\delta \in \Theta_{m''}$, let

$$\mathbf{V}_n(p_\delta^{m''}) = \frac{1}{n} \sum_{t=1}^n \mathbb{E} \left[\left(\log \frac{p_t^*(X_t | \mathcal{F}_{t-1})}{p_{\delta,t}^{m''}(X_t | \mathcal{F}_{t-1})} \right)^2 \mathbf{1}_{\left| \log \frac{p_t^*(X_t | \mathcal{F}_{t-1})}{p_{\delta,t}^{m''}(X_t | \mathcal{F}_{t-1})} \right| \leq F_{m''}^\infty} \middle| \mathcal{F}_{t-1} \right].$$

Let

$$A_{m'} = F_{m'}^{\text{lip}} M_{m'} + F_{m'}^\infty \quad \text{and} \quad v_{m'} = A_{m'} \sqrt{2n}. \quad (8)$$

For any $m' \in \mathcal{M}$, let $\sigma_{m'}$ be the solution of the equation

$$\sigma = \left(1 \wedge \frac{v_{m'}}{\sigma} \right) \sqrt{(D_{m'} + 1) \log \left(\frac{v_{m'}}{\sigma} \vee e \right) + \frac{A_{m'}}{\sigma} (D_{m'} + 1) \log \left(\frac{v_{m'}}{\sigma} \vee e \right)} \quad (9)$$

Lemma 5. *Assume Assumption 3 holds, as well as either Assumptions 1 and 2 or Assumptions 1bis and 2bis. For any family $(\eta_{m'})_{m' \in \mathcal{M}}$ taking values in $(0, 1)$, letting $y_{m'} = \eta_{m'}^{-1} \sqrt{\sigma_{m'}^2 + x + D_{m'}}$ for each $m' \in \mathcal{M}$, it holds with probability at least $1 - q_n - 6e^{-x} \sum_{m' \in \mathcal{M}} \log(v_{m'}) e^{-D_{m'}}$ that for all $m' \in \mathcal{M}$,*

$$\sup_{\delta \in \Theta_{m'}} \left(\frac{|\nu(g_\delta^{m'})|}{2\mathbf{V}_n(p_\delta^{m'}) + \frac{1}{n} y_{m'}^2} \right) \leq 80(2\eta_{m'} + \eta_{m'}^2 A_{m'}). \quad (10)$$

Proof. See Section C.1. □

Lemma 6. *Under either Assumptions 1 and 2 or Assumptions 1bis and 2bis, almost surely, for all $m' \in \mathcal{M}$ and $\delta \in \Theta_{m'}$,*

$$\mathbf{V}_n(p_\delta^{m'}) \leq 16(F_{m'}^\infty)^2 \mathbf{K}_n(p_\delta^{m'}).$$

Proof. See Section C.2. □

Fix some sequence $(\eta_{m'})_{m' \in \mathcal{M}}$ in $(0, 1)$ to be determined later and let $y_{m'} = \eta_{m'}^{-1} \sqrt{\sigma_{m'}^2 + x + D_{m'}}$ for all $m' \in \mathcal{M}$. By Lemma 5, and using the definition of Σ and the fact that since $A_{m'} \geq 2$, $n \geq 2$ and $x \geq 0$,

$$\log(v_{m'}) = \log A_{m'} + \frac{1}{2} \log(2n) \leq 3 \log(n) \log A_{m'},$$

it holds with probability at least $1 - q_n - 18 \log(n) \Sigma e^{-x}$, for all $m' \in \mathcal{M}'$,

$$-\nu(g_{\hat{\theta}_{m'}}^{m'}) \leq 80(2\eta_{m'} + \eta_{m'}^2 A_{m'}) \left(32(F_{m'}^\infty)^2 \mathbf{K}_n(p_{\hat{\theta}_{m'}}^{m'}) + \frac{1}{n} y_{m'}^2 \right)$$

and

$$\nu(g_{\hat{\theta}^m}^m) \leq 80(2\eta_m + \eta_m^2 A_m) \left(32(F_m^\infty)^2 \mathbf{K}_n(p_{\hat{\theta}^m}^m) + \frac{1}{n} y_m^2 \right).$$

Injecting this result in (7), with probability at least $1 - q_n - 18 \log(n) \Sigma e^{-x}$, for all $m' \in \mathcal{M}'$,

$$\begin{aligned} & (1 - 2560(F_{m'}^\infty)^2(2\eta_{m'} + \eta_{m'}^2 A_{m'})) \mathbf{K}_n(p_{\hat{\theta}^{m'}}^{m'}) \\ & \leq (1 + 2560(F_m^\infty)^2(2\eta_m + \eta_m^2 A_m)) \mathbf{K}_n(p_{\hat{\theta}^m}^m) \\ & + 80 \left(\frac{2}{\eta_m} + A_m \right) \frac{1}{n} (\sigma_m^2 + x + D_m) + \text{pen}(m) \\ & + 80 \left(\frac{2}{\eta_{m'}} + A_{m'} \right) \frac{1}{n} (\sigma_{m'}^2 + x + D_{m'}) - \text{pen}(m'). \end{aligned} \quad (11)$$

By the definition of $\sigma_{m'}$ in (9), it is smaller than the solution $\sigma'_{m'}$ of

$$\sigma = \sqrt{(D_{m'} + 1) \log \left(\frac{v_{m'}}{\sigma} \vee e \right)} + \frac{A_{m'}}{\sigma} (D_{m'} + 1) \log \left(\frac{v_{m'}}{\sigma} \vee e \right), \quad (12)$$

which satisfies

$$\begin{aligned} \sqrt{A_{m'}(D_{m'} + 1)} & \leq \sigma'_{m'} \leq \sqrt{(D_{m'} + 1) \log \left(\frac{\sqrt{2nA_{m'}}}{\sqrt{D_{m'} + 1}} \vee e \right)} \\ & + \sqrt{A_{m'}(D_{m'} + 1) \log \left(\frac{\sqrt{2nA_{m'}}}{\sqrt{D_{m'} + 1}} \vee e \right)} \\ & \leq 2\sqrt{A_{m'}(D_{m'} + 1) \log((nA_{m'}) \vee e)}. \end{aligned}$$

Note that $nA_{m'} \geq e$. Let $\kappa \in (0, 1]$ and for all $m' \in \mathcal{M}$, let $\eta_{m'} = \frac{c\kappa}{A_{m'}(F_{m'}^\infty)^{3/2}}$ for some small enough numerical constant $c > 0$, and recall that $A_{m'} \geq F_{m'}^\infty$, then there exist numerical constants C_{pen} and C such that if for all $m' \in \mathcal{M}'$,

$$\text{pen}(m') \geq \frac{C_{\text{pen}}}{\kappa} A_{m'}^2 (F_{m'}^\infty)^{3/2} \log(nA_{m'})^2 \frac{D_{m'}}{n}, \quad (13)$$

it holds for all $x \geq 0$, with probability at least $1 - q_n - 18 \log(n) \Sigma e^{-x}$, for all $m' \in \mathcal{M}'$,

$$\begin{aligned} (1 - \kappa) \mathbf{K}_n(p_{\hat{\theta}^{m'}}^{m'}) & \leq (1 + \kappa) \mathbf{K}_n(p_{\hat{\theta}^m}^m) + 2 \text{pen}(m) \\ & + \frac{C}{\kappa} (A_{m'}(F_{m'}^\infty)^{3/2} \log(nA_{m'})^2 + A_m (F_m^\infty)^{3/2} \log(nA_m)^2) \frac{x}{n}. \end{aligned} \quad (14)$$

B Proof of Corollaries 2 and 4

Corollary 2 follows directly from the fact that $\mathbb{E}[Z] \leq \int_{t \geq 0} \mathbb{P}(Z \geq t) dt$ for any random variable Z .

Under Assumptions 1bis, let us first show that $\mathbf{K}_n(p_\delta^{m'})$ is bounded for all δ and m' almost surely. By Assumption 1bis, almost surely, for any $m' \in \mathcal{M}$ and $\delta \in \Theta_{m'}$,

$$\begin{aligned} \mathbf{K}_n(p_\delta^{m'}) &\leq \frac{1}{n} \sum_{t=1}^n \mathbb{E} \left[\left| \log \frac{p_t^*(X_t | \mathcal{F}_{t-1})}{p_{\delta,t}^{m'}(X_t | \mathcal{F}_{t-1})} \right| \middle| \mathcal{F}_{t-1} \right] \\ &\leq \frac{1}{n} \sum_{t=1}^n \left(B_{m'} + B_{m'} \int_1^{+\infty} \mathbb{P} \left(\left| \log \frac{p_t^*(X_t | \mathcal{F}_{t-1})}{p_{\delta,t}^{m'}(X_t | \mathcal{F}_{t-1})} \right| \geq B_{m'} y \middle| \mathcal{F}_{t-1} \right) dy \right) \\ &\leq 2B_{m'}. \end{aligned}$$

To conclude, assume that there exist $A, B > 0$ such that $\sup_{m \in \mathcal{M}} B_m \leq B(n)$ and $\sup_{m \in \mathcal{M}} A_m \leq A(n)$, so that by Theorem 3, with probability at least $1 - 2n^{-1} - 18 \log(n) \Sigma e^{-x}$,

$$\begin{aligned} (1 - \kappa) \mathbf{K}_n(\tilde{p}) &\leq \inf_{m \in \mathcal{M}} \left((1 + \kappa) \inf_{\theta \in \Theta^{D_m}} \mathbf{K}_n(p_\theta^m) + 2 \text{pen}(m) \right) \\ &\quad + \frac{2C'}{\kappa} A(n) B(n)^{3/2} \log(n A(n))^2 \frac{(\log n)^{5/2} x}{n}, \end{aligned}$$

and use that for any random variable Z such that $Z \leq M$ a.s. for some constant $M > 0$, $\mathbb{E}[Z] \leq \int_{t=0}^M \mathbb{P}(Z \geq t) dt$, so that for all $\kappa \in (0, 1]$,

$$\begin{aligned} (1 - \kappa) \mathbb{E} \left[\mathbf{K}_n(p_{\hat{\theta}^m}^m) \right] &\leq \mathbb{E} \left[\inf_{m \in \mathcal{M}} \left((1 + \kappa) \inf_{\theta \in \Theta^{D_m}} \mathbf{K}_n(p_\theta^m) + 2 \text{pen}(m) \right) \right] \\ &\quad + \frac{4B(n)}{n} + \frac{36C'}{\kappa} \Sigma A(n) B(n)^{3/2} \log(n A(n))^2 \frac{(\log n)^{7/2}}{n}, \end{aligned}$$

and the last term dominates the second to last one.

C Proof of the Lemmas

C.1 Proof of Lemma 5

Fix $m' \in \mathcal{M}$. Recall that for any $\delta, \eta \in \Theta_{m'}$,

$$\nu(g_\delta^{m'}) - \nu(g_\eta^{m'}) = \frac{1}{n} \sum_{t=1}^n \int \log \left(\frac{p_{\eta,t}^{m'}(x_t | \mathcal{F}_{t-1})}{p_{\delta,t}^{m'}(x_t | \mathcal{F}_{t-1})} \right) (d\delta_{X_t}(x_t) - p_t^*(x_t | \mathcal{F}_{t-1}) d\mu(x_t)),$$

where δ_a is the Dirac measure in a .

We extend $\Theta_{m'}$ into $\Theta_{m'} \cup \{\star\}$ by defining $p_\star^{m'} = p^*$, so that when $\eta = \star$, $\nu(g_\eta^{m'}) = 0$ and the formula above becomes $\nu(g_\delta^{m'})$. We want to control this uniformly over $\delta, \eta \in \Theta_{m'} \cup \{\star\}$.

Fix $\delta, \eta \in \Theta_{m'} \cup \{\star\}$. For any $t \in [n]$, let

$$\Delta_t = \int \log \left(\frac{p_{\eta,t}^{m'}(x_t | \mathcal{F}_{t-1})}{p_{\delta,t}^{m'}(x_t | \mathcal{F}_{t-1})} \right) (d\delta_{X_t}(x_t) - p_t^*(x_t | \mathcal{F}_{t-1}) d\mu(x_t)).$$

For any $t \in [n + 1]$, let $M_t = \sum_{s=1}^{t-1} \Delta_s$ (in particular, $M_1 = 0$), so that $\frac{1}{n}M_{n+1} = \nu(g_\delta^{m'}) - \nu(g_\eta^{m'})$.

$(M_t)_{t \geq 1}$ is a $(\sigma(X_1^{t-1}))_{t \geq 1}$ -martingale. For $\ell \geq 2$, let $C_1^\ell = 0$ and for $t \geq 2$, let

$$C_t^\ell = \sum_{s=1}^{t-1} \mathbb{E}[\Delta_s^\ell | \mathcal{F}_{t-1}].$$

Note that for all $s \in [n]$,

$$|\Delta_s| \leq 2 \int \left| \log \frac{p_{\eta,s}^{m'}(x_s | \mathcal{F}_{s-1})}{p_{\delta,s}^{m'}(x_s | \mathcal{F}_{s-1})} \right| \frac{d\delta_{X_s}(x_s) + p_s^*(x_s | \mathcal{F}_{s-1})d\mu(x_s)}{2},$$

so that by convexity of $x \mapsto x^\ell$,

$$\begin{aligned} |C_t^\ell| &\leq \sum_{s=1}^{t-1} \mathbb{E} \left[2^\ell \int \left| \log \frac{p_{\eta,s}^{m'}(x_s | \mathcal{F}_{s-1})}{p_{\delta,s}^{m'}(x_s | \mathcal{F}_{s-1})} \right|^\ell \frac{d\delta_{X_s}(x_s) + p_s^*(x_s | \mathcal{F}_{s-1})d\mu(x_s)}{2} \Big| \mathcal{F}_{s-1} \right] \\ &= 2^\ell \sum_{s=1}^{t-1} \int \left| \log \frac{p_{\eta,s}^{m'}(x_s | \mathcal{F}_{s-1})}{p_{\delta,s}^{m'}(x_s | \mathcal{F}_{s-1})} \right|^\ell p_s^*(x_s | \mathcal{F}_{s-1})d\mu(x_s). \end{aligned} \quad (15)$$

Let \mathcal{Q} be a countable dense subset of \mathcal{X} .

Now, let us distinguish two cases:

- Under Assumptions 1 and 2, for each $t \in [n]$, let A_n^t be the event defined by

$$\left\{ \begin{aligned} &\forall s \in [t], \forall m'' \in \mathcal{M}, \forall \delta, \theta \in \Theta_{m''} \cup \{\star\}, \left| \log \frac{p_\delta^{m''}(X_s | \mathcal{F}_{s-1})}{p_\theta^{m''}(X_s | \mathcal{F}_{s-1})} \right| \leq 2 \log \frac{1}{\varepsilon} \\ &\text{and } \forall s \in [t], \forall m'' \in \mathcal{M}, \forall \delta, \theta \in \Theta_{m''}, \left| \log \frac{p_\delta^{m''}(X_s | \mathcal{F}_{s-1})}{p_\theta^{m''}(X_s | \mathcal{F}_{s-1})} \right| \leq L_{m''} \|\delta - \theta\|_{m''} \end{aligned} \right\}.$$

Thanks to Assumptions 1 and 2, $\mathbb{P}(A_n^t) = 1$. Write $F_{m'}^\infty = 2 \log(\varepsilon^{-1})$ and $F_{m'}^{\text{lip}} = L_{m''}$.

- Under Assumptions 1bis and 2bis, for each $t \in [n]$, let A_n^t be the event defined by

$$\left\{ \begin{aligned} &\forall s \in [t], \forall m'' \in \mathcal{M}, \forall \delta, \theta \in \Theta_{m''} \cup \{\star\}, \left| \log \frac{p_\delta^{m''}(X_s | \mathcal{F}_{s-1})}{p_\theta^{m''}(X_s | \mathcal{F}_{s-1})} \right| \leq 2B_{m''} \log n \\ &\text{and } \forall s \in [t], \forall m'' \in \mathcal{M}, \forall \delta, \theta \in \Theta_{m''}, \left| \log \frac{p_\delta^{m''}(X_s | \mathcal{F}_{s-1})}{p_\theta^{m''}(X_s | \mathcal{F}_{s-1})} \right| \leq L_{m''} \|\delta - \theta\|_{m''} \log n \end{aligned} \right\}$$

Thanks to Assumptions 1bis and 2bis, $\mathbb{P}(A_n^t) \geq 1 - 2n^{-1}$. Write $F_{m'}^\infty = 2B_{m''} \log n$ and $F_{m'}^{\text{lip}} = L_{m''} \log n$.

Note that in both cases, A_n^t depends on X_t and \mathcal{F}_{t-1} . So given \mathcal{F}_{t-1} , we can ask ourselves what are the possible values x of X_t , so that A_n^t is true. To simplify notation, we write $x \in A_n^t$ when this is true. The event $\{x \in A_n^t\}$ is therefore \mathcal{F}_{t-1} measurable.

Now, under Assumption 3, we may define

$$R_{\infty,n}(\delta, \eta) = \max_{1 \leq s \leq n} \sup_{x \in \mathcal{Q}} \left(\left| \log \frac{p_{\eta,s}^{m'}(x|\mathcal{F}_{s-1})}{p_{\delta,s}^{m'}(x|\mathcal{F}_{s-1})} \right| \mathbf{1}_{x \in A_n^s} \right)$$

and

$$R_{2,n}(\delta, \eta)^2 = 2 \sum_{s=1}^n \mathbb{E} \left[\left(\log \frac{p_{\eta,s}^{m'}(X_s|\mathcal{F}_{s-1})}{p_{\delta,s}^{m'}(X_s|\mathcal{F}_{s-1})} \right)^2 \mathbf{1}_{A_n^s} \middle| \mathcal{F}_{s-1} \right].$$

Note that because \mathcal{Q} is countable, both $R_{\infty,n}(\delta, \eta)$ and $R_{2,n}(\delta, \eta)$ are well defined and random. So that on the event A_n^n ,

$$|C_t^\ell| \leq 2^{\ell-1} R_{\infty,n}(\delta, \eta)^{\ell-2} R_{2,n}(\delta, \eta)^2,$$

and since $2^{\ell-1} \leq \ell!$ for all $\ell \geq 2$, on the event A_n^n ,

$$|C_{n+1}^\ell| \leq \frac{\ell!}{2} R_{2,n}(\delta, \eta)^2 R_{\infty,n}(\delta, \eta)^{\ell-2}. \quad (16)$$

By definition, on the event A_n^n ,

$$\begin{cases} \forall \delta, \eta \in \Theta_{m'} \cup \{\star\}, & R_{2,n}(\delta, \eta) \leq F_{m'}^\infty \sqrt{2n}, \\ \forall \delta, \eta \in \Theta_{m'}, & R_{2,n}(\delta, \eta) \leq F_{m'}^{\text{lip}} \|\delta - \eta\|_{m'} \sqrt{2n}. \end{cases} \quad (17)$$

Finally, since $R_{2,n}(\delta, \eta)$ is the Euclidean norm of the vector whose coordinate $s \in [n]$ is the $\mathbf{L}^2(\mathcal{X}, p^*(\cdot|\mathcal{F}_{s-1})d\mu)$ distance between $\log p_{\eta,s}^{m'}(\cdot|\mathcal{F}_{s-1})\mathbf{1}_{A_n^s}$ and $\log p_{\delta,s}^{m'}(\cdot|\mathcal{F}_{s-1})\mathbf{1}_{A_n^s}$, it satisfies the triangular inequality: for all $\eta, \delta, \theta \in \Theta_{m'} \cup \{\star\}$,

$$R_{2,n}(\delta, \eta) \leq R_{2,n}(\delta, \theta) + R_{2,n}(\theta, \eta). \quad (18)$$

Likewise, on the event A_n^n ,

$$\begin{cases} \forall \delta, \eta \in \Theta_{m'} \cup \{\star\}, & R_{\infty,n}(\delta, \eta) \leq F_{m'}^\infty, \\ \forall \delta, \eta \in \Theta_{m'}, & R_{\infty,n}(\delta, \eta) \leq F_{m'}^{\text{lip}} \|\delta - \eta\|_{m'}. \end{cases} \quad (19)$$

and for all $\eta, \delta, \theta \in \Theta_{m'} \cup \{\star\}$

$$R_{\infty,n}(\delta, \eta) \leq R_{\infty,n}(\delta, \theta) + R_{\infty,n}(\theta, \eta). \quad (20)$$

Identify $\Theta_{m'} \cup \{\star\}$ with the subset $\tilde{\Theta}_{m'}$ of the vector space $\mathbb{R}^{D_{m'}+1}$ of generic element (δ, u) with $\delta \in \mathbb{R}^{D_{m'}}$ and $u \in \mathbb{R}$, defined as

$$\tilde{\Theta}_{m'} = \{(\theta, 0) : \theta \in \Theta_{m'}\} \cup \{(\bar{\theta}^{m'}, 1)\},$$

for some fixed $\bar{\theta}^{m'} \in \Theta_{m'}$. Endow the vector space $\mathbb{R}^{D_{m'}+1}$ with the norms

$$\tilde{N}_\infty((\delta, u)) = F_{m'}^{\text{lip}} \|\delta\|_{m'} + F_{m'}^\infty |u| \quad \text{and} \quad \tilde{N}_2((\delta, u)) = \sqrt{2n} \tilde{N}_\infty((\delta, u)).$$

Under Assumption 2bis, for any $\delta, \eta \in \Theta_{m'}$, by (17) and (19), on the event A_n^n ,

$$\begin{cases} R_{2,n}(\delta, \eta) \leq F_{m'}^{\text{lip}} \|\delta - \eta\|_{m'} \sqrt{2n} = \tilde{N}_2((\delta, 0) - (\eta, 0)) \\ R_{\infty,n}(\delta, \eta) \leq F_{m'}^{\text{lip}} \|\delta - \eta\|_{m'} = \tilde{N}_\infty((\delta, 0) - (\eta, 0)), \end{cases}$$

and both inequalities extend to $\eta = \star$ since, for $R_{2,n}$,

$$\begin{aligned} R_{2,n}(\delta, \star) &\leq F_{m'}^\infty \sqrt{2n} \\ &\leq F_{m'}^{\text{lip}} \|\delta - \bar{\theta}^{m'}\|_{m'} \sqrt{2n} + F_{m'}^\infty \sqrt{2n} \\ &= \tilde{N}_2((\delta, 0) - (\bar{\theta}^{m'}, 1)) \end{aligned}$$

and likewise for $\delta = \star$ and for $R_{\infty,n}$. Let

$$A_{m'} = F_{m'}^{\text{lip}} M_{m'} + F_{m'}^\infty \quad \text{and} \quad v_{m'} = A_{m'} \sqrt{2n},$$

so that $\tilde{N}_2(\delta - \eta) \leq v_{m'}$ and $\tilde{N}_\infty(\delta - \eta) \leq A_{m'}$ for all $\delta, \eta \in \tilde{\Theta}_{m'}$.

We may now apply Theorem 5 of [Aubert and Lehericy, 2024] to the process $Y_\delta = n\nu(g_\delta^{m'})$ indexed by $\delta \in \Theta_{m'} \cup \{\star\}$, with $c = 0$ and the event $A = A_n^n$: for all $\sigma > 0$ and $x \geq 0$, let

$$\Psi(\sigma, x) = 20 \left[(\sigma \wedge v_{m'}) \sqrt{x + (D_{m'} + 1) \log \left(\frac{v_{m'}}{\sigma} \vee e \right)} + A_{m'} \left(x + (D_{m'} + 1) \log \left(\frac{v_{m'}}{\sigma} \vee e \right) \right) \right].$$

Then, for all $\theta \in \Theta_{m'} \cup \{\star\}$, $\sigma > 0$ and $x \geq 0$,

$$\begin{aligned} \mathbb{P} \left(\left\{ \sup_{\delta \in \Theta_{m'} \cup \{\star\}} \frac{Y_\delta - Y_\theta}{R_{2,n}(\delta, \theta)^2 + \sigma^2} \geq 4\sigma^{-2} \Psi(\sigma, x + D_{m'}) \right\} \cap A_n^n \right) \\ \leq \left(2 \log \left(\frac{v_{m'}}{\sigma} \vee 0 + 1 \right) e^{-(x + D_{m'})} \right). \end{aligned}$$

In particular, by taking the union bound over $m' \in \mathcal{M}$ for $\theta = \star$, for any family of positive numbers $(y_{m'})_{m' \in \mathcal{M}}$, with probability at least $\mathbb{P}(A_n^n) - e^{-x} \sum_{m' \in \mathcal{M}} \left(2 \log \left(\frac{v_{m'}}{y_{m'}} \vee 0 + 1 \right) e^{-D_{m'}} \right)$, for all $m' \in \mathcal{M}$,

$$\begin{aligned} \sup_{\delta \in \Theta_{m'}} \left(\frac{n\nu(g_\delta^{m'})}{R_{2,n}(\delta, \star)^2 + y_{m'}^2} \right) \\ \leq \frac{80}{y_{m'}^2} \left(y_{m'} \sqrt{x} + (y_{m'} \wedge v_{m'}) \sqrt{(D_{m'} + 1) \log \left(\frac{v_{m'}}{y_{m'}} \vee e \right)} \right. \\ \left. + A_{m'} x + A_{m'} (D_{m'} + 1) \log \left(\frac{v_{m'}}{y_{m'}} \vee e \right) \right). \end{aligned}$$

For each $m' \in \mathcal{M}$, let $\sigma_{m'}$ be the solution of the equation

$$\sigma = \left(1 \wedge \frac{v_{m'}}{\sigma}\right) \sqrt{(D_{m'} + 1) \log\left(\frac{v_{m'}}{\sigma} \vee e\right)} + \frac{A_{m'}}{\sigma} (D_{m'} + 1) \log\left(\frac{v_{m'}}{\sigma} \vee e\right),$$

which exists since the right hand side is positive and non-increasing on $(0, +\infty)$. For any family $(y_{m'})_{m' \in \mathcal{M}}$ such that $y_{m'} \geq \sigma_{m'}$ for all $m' \in \mathcal{M}$, for any $x \geq 0$, it holds with probability at least $\mathbb{P}(A_n^n) - e^{-x} \sum_{m' \in \mathcal{M}} \left(2 \log\left(\frac{v_{m'}}{y_{m'}}\right) \vee 0 + 1\right) e^{-D_{m'}}$ that for all $m' \in \mathcal{M}$,

$$\sup_{\delta \in \Theta_{m'}} \left(\frac{n\nu(g_\delta^{m'})}{R_{2,n}(\delta, \star)^2 + y_{m'}^2} \right) \leq \frac{80}{y_{m'}} \left(\sigma_{m'} + \sqrt{x + D_{m'}} + \frac{A_{m'}}{y_{m'}} (x + D_{m'}) \right).$$

Let $\eta \in (0, 1)$ and fix for each $m' \in \mathcal{M}$

$$y_{m'} = \eta^{-1} \sqrt{\sigma_{m'}^2 + x + D_{m'}},$$

for all $x \geq 0$, with probability at least $\mathbb{P}(A_n^n) - e^{-x} \sum_{m' \in \mathcal{M}} \left(2 \log\left(\frac{v_{m'}}{y_{m'}}\right) \vee 0 + 1\right) e^{-D_{m'}}$, for all $m' \in \mathcal{M}$,

$$\sup_{\delta \in \Theta_{m'}} \left(\frac{\nu(g_\delta^{m'})}{\frac{1}{n} R_{2,n}(\delta, \star)^2 + \frac{1}{n} y_{m'}^2} \right) \leq 80(2\eta + \eta^2 A_{m'}),$$

where

- the first term on the right hand side is due to the concavity of $x \in (0, +\infty) \mapsto \sqrt{x}$,
- the second term on the right hand side holds because $x + D_{m'} \leq x + D_{m'} + \sigma_{m'}^2 = \eta^2 y_{m'}^2$.

By definition $y_{m'} \geq \eta^{-1} \sqrt{D_{m'}} \geq 1$, and $v_{m'} = A_{m'} \sqrt{2n} \geq e$. Therefore,

$$2 \log\left(\frac{v_{m'}}{y_{m'}}\right) \vee 0 + 1 \leq 3 \log v_{m'}.$$

The control of $-\nu(g_\delta^{m'})$ is identical, hence we may control all $|\nu(g_\delta^{m'})|$ with probability at least $\mathbb{P}(A_n^n) - 6e^{-x} \sum_{m' \in \mathcal{M}} \log(v_{m'}) e^{-D_{m'}}$ by union bound. To conclude the proof of the Lemma, note that $\frac{1}{n} R_{2,n}(\delta, \star)^2 \leq 2\mathbf{V}_n(p_\delta)$.

C.2 Proof of Lemma 6

Let's begin with a result proved in [Shen et al., 2013] which we slightly adapt to our situation.

Lemma 7 (Adaptation of Lemma 4 in [Shen et al., 2013]). *For any probability measures P and Q with densities p and q , and any $\lambda \in (0, 1/2]$,*

$$P \left\{ \left(\log \frac{p}{q} \right)^2 \mathbf{1}_{|\log(\frac{p}{q})| \leq \log(\frac{1}{\lambda})} \right\} \leq 8 \left(1 + \left(\log \frac{1}{\lambda} \right)^2 \right) P \left\{ \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \mathbf{1}_{|\log(\frac{p}{q})| \leq \log(\frac{1}{\lambda})} \right\}.$$

Proof. See Section C.3. □

Let us apply Lemma 7 to $p = p_s^*$ and $q = p_{\delta,s}^{m'}$ with $\lambda = \exp(-F_{m'}^\infty)$ for all $s \in [n]$ and $\delta \in \Theta_{m'}$. By definition of $F_{m'}^\infty$, $\lambda \leq 1/2$, so that for all $s \in [n]$ and $\delta \in \Theta_{m'}$,

$$\begin{aligned} \int \left| \log \frac{p_s^*(x_s | \mathcal{F}_{s-1})}{p_{\delta,s}^{m'}(x_s | \mathcal{F}_{s-1})} \right|^2 \mathbf{1}_{\left| \log \frac{p_s^*(x_s | \mathcal{F}_{s-1})}{p_{\delta,s}^{m'}(x_s | \mathcal{F}_{s-1})} \right| \leq F_{m'}^\infty} p_s^*(x_s | \mathcal{F}_{s-1}) d\mu(x_s) \\ \leq 16h^2 \left(p_s^*(\cdot | \mathcal{F}_{s-1}), p_{\delta,s}^{m'}(\cdot | \mathcal{F}_{s-1}) \Big| \mathcal{F}_{s-1} \right) (1 + (F_{m'}^\infty)^2), \end{aligned}$$

where

$$h^2 \left(p_t(\cdot | \mathcal{F}_{t-1}), q_t(\cdot | \mathcal{F}_{t-1}) \Big| \mathcal{F}_{t-1} \right) = \frac{1}{2} \int \left(\sqrt{p_t(x_t | \mathcal{F}_{t-1})} - \sqrt{q_t(x_t | \mathcal{F}_{t-1})} \right)^2 d\mu(x_t).$$

Let us recall a classical relation between the Hellinger distance and the Kullback-Leibler divergence, see for instance in [Massart, 2007, Lemma 7.23]: for any probability measures P and Q ,

$$2h^2(P, Q) \leq \text{KL}(P, Q).$$

Applying this inequality to the probability measures $P = p_t^*(\cdot | X_1^{t-1})$ and $Q = p_{t,\hat{\theta}_{m'}}^{m'}(\cdot | X_1^{t-1})$ conditionally to X_1^{t-1} for all $t \in [n]$ shows that

$$\mathbf{V}_n(p_\delta^{m'}) \leq 8(1 + (F_{m'}^\infty)^2) \mathbf{K}_n(p_\delta^{m'}) \leq 16(F_{m'}^\infty)^2 \mathbf{K}_n(p_\delta^{m'})$$

since $F_{m'}^\infty \geq 1$.

C.3 Proof of Lemma 7

The proof follows exactly the same steps as [Shen et al., 2013]. Let $r : (0, +\infty) \rightarrow \mathbb{R}$ be the function defined implicitly by

$$\log(x) = 2(x^{1/2} - 1) - r(x)(x^{1/2} - 1)^2.$$

The function r is non-negative, decreasing, and $r(x) \leq 2 \log(1/x)$ for all $x \in (0, 1/2]$ (see e.g. [Ghosal and van der Vaart, 2001]). Let $\lambda \in (0, 1/2]$. Since for any $x \geq 1$, $|\log(x)| \leq 2|x^{1/2} - 1|$,

$$P \left\{ \left(\log \frac{p}{q} \right)^2 \mathbf{1}_{1 \leq \frac{q}{p} \leq \frac{1}{\lambda}} \right\} \leq 4P \left\{ \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \mathbf{1}_{1 \leq \frac{q}{p} \leq \frac{1}{\lambda}} \right\}.$$

Moreover, by definition of r ,

$$\begin{aligned}
& P \left\{ \left(\log \frac{p}{q} \right)^2 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq 1} \right\} \\
& \leq 8P \left\{ \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq 1} \right\} + 2P \left\{ r^2 \left(\frac{q}{p} \right) \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^4 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq 1} \right\} \\
& \leq 8P \left\{ \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq 1} \right\} + 2r^2(\lambda)P \left\{ \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq 1} \right\} \\
& \leq 8P \left\{ \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq 1} \right\} + 8 \left(\log \frac{1}{\lambda} \right)^2 P \left\{ \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq 1} \right\},
\end{aligned}$$

where

- the first inequality holds because for any $a, b \in \mathbb{R}$, $(a + b)^2 \leq 2a^2 + 2b^2$,
- the second inequality holds because r is decreasing and since $0 \leq \frac{q}{p} \leq 1$, $\left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \leq 1$,
- the third inequality holds because $r(x) \leq 2 \log(1/x)$ for $x \in (0, 1/2]$.

All in all,

$$P \left\{ \left(\log \frac{p}{q} \right)^2 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq \frac{1}{\lambda}} \right\} \leq 8 \left(1 + \left(\log \frac{1}{\lambda} \right)^2 \right) P \left\{ \left(\frac{q^{1/2}}{p^{1/2}} - 1 \right)^2 \mathbf{1}_{\lambda \leq \frac{q}{p} \leq \frac{1}{\lambda}} \right\}.$$