

Three Rates of Convergence or Separation via U-Statistics in a Dependent Framework

Quentin Duchemin

*LAMA, Univ Gustave Eiffel, Univ Paris Est Creteil,
CNRS, F-77447 Marne-la-Vallée, France*

QUENTIN.DUCHEMIN@UNIV-EIFFEL.FR

Yohann De Castro

*Institut Camille Jordan
École Centrale de Lyon, Lyon, France*

YOHANN.DE-CASTRO@EC-LYON.FR

Claire Lacour

*LAMA, Univ Gustave Eiffel, Univ Paris Est Creteil,
CNRS, F-77447 Marne-la-Vallée, France*

CLAIRE.LACOUR@UNIV-EIFFEL.FR

Editor: Aryeh Kontorovich

Abstract

Despite the ubiquity of U-statistics in modern Probability and Statistics, their non-asymptotic analysis in a dependent framework may have been overlooked. In a recent work, a new concentration inequality for U-statistics of order two for uniformly ergodic discrete time Markov chains has been proved. In this paper, we put this theoretical breakthrough into action by pushing further the current state of knowledge in three different active fields of research. First, we establish a new exponential inequality for the estimation of spectra of integral operators with MCMC methods. The novelty is that this result holds for kernels with positive and negative eigenvalues, which is new as far as we know.

In addition, we investigate generalization performance of online algorithms working with pairwise loss functions and Markov chain samples. We provide an online-to-batch conversion result by showing how we can extract a low risk hypothesis from the sequence of hypotheses generated by any online learner.

We finally give a non-asymptotic analysis of a goodness-of-fit test on the density of the stationary measure of a Markov chain. We identify some classes of alternatives over which our test based on the L^2 distance has a prescribed power.

Keywords: U-statistics, Markov chains, Concentration inequality, Integral operators, Online learning, Non-parametric hypothesis testing

1. Introduction

For the last twenty years, the phenomenon of the concentration of measure has received much attention. The main interesting feature of concentration inequalities is that, unlike central limit theorems or large deviations inequalities, they are nonasymptotic. Among others, Pascal Massart, Michel Ledoux and Gabor Lugosi produced a series of works that led to a large span of powerful inequalities. Their results have found application in model selection (cf. Massart, 2007; Lerasle et al., 2016), statistical learning (cf. Cléménçon et al., 2020), online learning (cf. Wang et al., 2012) or random graphs (cf. De Castro et al., 2019; Duchemin and De Castro, 2022). Most of the concentration inequalities are formulated for

U-statistics of order m (cf. Van der Vaart, 2000, Chapter 12), which are defined as a sum of the form

$$\sum_{1 \leq i_1 < \dots < i_m \leq n} h_{i_1, \dots, i_m}(X_{i_1}, \dots, X_{i_m}),$$

where X_1, \dots, X_n are random variables taking values in a measurable space (E, Σ) (with E Polish) and where h_{i_1, \dots, i_m} are measurable functions of m variables $h_{i_1, \dots, i_m} : E^m \rightarrow \mathbb{R}$. The pioneering works considered independent random variables $(X_i)_{i \geq 1}$, an assumption that can be prohibitive for practical applications which often involve some dependence structure. To cope with this issue, some researchers left the independent setting by working with Markov chains or by adopting some mixing conditions and we refer for example to Fan et al. (2021); Jiang et al. (2018); Paulin (2015); Adamczak (2008); Cl  men  on et al. (2020). The previous mentioned papers considered U-statistics of order $m = 1$ and the non-asymptotic behaviour of tails of U-statistics of order $m \geq 2$ in a dependent framework remains so far understudied. Recently, the two papers Duchemin et al. (2022) and Shen et al. (2020) made a first step to fill this gap. While Shen et al. (2020) consider U-statistics of arbitrary order with smooth and symmetric kernels and work with mixing conditions, Duchemin et al. (2022) are focused on U-statistics of order two for uniformly ergodic discrete time Markov chains and bounded kernels. Let us highlight that we work with the result of the former paper rather than the one from Shen et al. (2020) since we need a concentration inequality valid for any initial distribution of the chain. We give further details at the beginning of Section 2.2.

Our paper is in the line of work of Massart (2000) where concentration of measure is applied to tackle problems arising from model selection. In this work, we shed light on the large number of potential theoretical breakthroughs allowed by a better understanding of the non-asymptotic tail behavior of U-statistics of order two in a dependent framework. We present new theoretical results in three different branches of Statistics ranging from online learning to goodness-of-fit tests. In Section 1.1 we describe in details our three contributions, highlighting their applications to learning theory and the proof innovations.

1.1 Our contributions

Our new results - that we referred to as applications for brevity - push further the current state of knowledge in three different active areas of research in Probability, Statistics and Machine Learning. Although the recent progress in concentration inequality for U-statistics with dependent random variables is a key element in our proofs, our contributions are not a direct consequence of it. The purpose of this section is threefold: (i) we present concisely our main results, (ii) we highlight the proof innovations of our approach compared to previous works and (iii) we propose relevant connections between our work and other important topics in learning theory and Statistics.

- **Estimation of spectra of signed integral operator with MCMC algorithms** (Section 3)

We study the convergence of sequence of spectra of kernel matrices towards the spectrum of some integral operator. Previous important works may include Adamczak and Bednorz (2015a) and, as far as we know, they all assume that the kernel is of positive-type (i.e. giving an integral operator with non-negative eigenvalues). Getting counterpart of those results for signed integral operators is of great interest since they

arise for example in random graphs with latent space (cf. De Castro et al., 2019) which can be characterized by the so-called graphon (cf. Lovász, 2012). For the first time, this paper proves a non-asymptotic result of convergence of spectra for kernels that are not of positive-type. We further prove that *independent Hastings algorithms* are valid sampling schemes to apply our result.

Proof innovations. In Section 3.2, we propose a comparison between our result and the one from Adamczak and Bednorz (2015a). We explain why working with integral operators of positive-type allows Adamczak and Bednorz to make use of a powerful decoupling technique. Thanks to this elegant argument, they are reduced to prove a concentration inequality for a sum of Banach space valued random variables where the i -th summand depends only on the i -th visited state of the Markov chain. By considering signed integral operators, the approach of the former paper cannot be adapted. Our proof relies on a low rank approximation of the kernel and on a concentration result for U-statistics with dependent random variables.

Application to learning theory. A large number of learning algorithms aim at estimating the eigenvalues and/or the eigenvectors of data-dependent matrices. This is for example the case for Principal Component Analysis (PCA) or some manifold methods (cf. Rosasco et al., 2010). It appears that these matrices can often be interpreted as the empirical versions of continuous objects such as integral operators. As highlighted in Rosasco et al. (2010), the theoretical analysis of the above mentioned learning algorithms requires to quantify the difference between the eigen-structure of the empirical operators and their continuous counterparts. Specific examples coming from the Machine Learning and the Statistics communities where our result may find an echo include the estimation of the entire spectrum of a Markov operator (cf. Chakraborty and Khare, 2019), estimation procedures in random graphs (cf. Duchemin and De Castro, 2022) or the analysis of the generalization properties of neural networks (cf. Zhang et al., 2021).

- **Online learning with pairwise loss functions** (Section 4)

In Machine Learning, several important problems involve a pairwise loss function, i.e. a loss function which depends on a pair of examples. One typical example is the problem of metric learning (cf. Jin et al., 2009) where one aims to learn a metric so that instances with the same labels are close while ones with different labels are far away from each other. Other pairwise learning tasks include preference learning (cf. Xing et al., 2002), ranking (cf. Agarwal and Duchi, 2012), gradient learning (cf. Meir and Zhang, 2003) and AUC maximization (cf. Zhao et al., 2011). Batch learning algorithms with pairwise loss functions have been extensively studied and their generalization properties have been well established. However, batch algorithms have some limitations especially when data becomes available in a sequential order or for large scale learning problems where their computational cost can be prohibitive. Online algorithms have been designed to efficiently solve learning problems in such situations: they deal with data coming on fly and try to improve the learned model along time based on the new observations. The performance of online learning algorithms is typically analyzed through the notion of *regret* which compares the payoff obtained by the algorithm along time with the one that would have been obtained by taking the optimal decision at each time step (cf.

Bubeck and Cesa-Bianchi, 2012). The regret quantifies the number of mistakes made by the algorithm without requiring assumptions on the way the training sequence is generated. When the sequence of observations is the realization of some stochastic process, one can analyze online algorithms through a different lens by wondering how they generalize on future examples. More precisely, we would like to convert a regret bound of an online learner into a control of the excess risk. In the online learning research community, these types of results are called *online-to-batch conversion* and we refer to (Hoi et al., 2021, Section 3.7) for a comprehensive introduction to this topic. Online-to-batch conversion results for online learning with univariate or pairwise loss functions working with i.i.d. samples have been considered for quite a while in both Machine Learning and Statistics literature (cf. Wang et al., 2012; Ying and Zhou, 2017; Guo et al., 2017; Chen and Lei, 2018). For dependent data sequences, generalization bounds for online algorithms have also been proved in the last decades with univariate loss functions (cf. Agarwal and Duchi, 2012). However, theoretical guarantees for the generalization performance of online algorithms with pairwise loss functions with non i.i.d. data have been so far little studied. Inspired by Wang et al. (2012), our work is one of the first to bring results regarding this problem. In Section 4.1.4, we establish clear connections with the existing literature.

Proof innovations. Wang et al. (2012) was a pioneering work for the study of generalization performance of online learning algorithms with pairwise loss functions and worked with i.i.d. observations. In this paper, we extend the result of Wang et al. (2012) by considering a dependent framework that makes the theoretical analysis more challenging. In our proofs, we bypass the additional issues arising from data dependency using properties of uniformly ergodic Markov chains, concentration inequalities for U-statistics (of order one and two) of dependent random variables and reversibility of Markov chains by considering the time-reversed sequence. Using the marker $\hat{\otimes}$, we shed light in Section B on the specific parts of the proof where the arguments used in the i.i.d. framework fail, requiring a specific theoretical work handling a sequence of dependent observations.

- **Adaptive goodness-of-fit tests in a density model** (Section 5)

Several works have already proposed goodness-of-fit tests for the density of the stationary distribution of a sequence of dependent random variables. In Li and Tkacz (2001), a test based on an L^2 -type distance between the nonparametrically estimated conditional density and its model-based parametric counterpart is proposed. In Bai (2003) a Kolmogorov-type test is considered. Chwialkowski et al. (2016) derive a test procedure for τ -mixing sequences using Stein discrepancy computed in a reproducing kernel Hilbert space. In all the above mentioned papers, asymptotic properties of the test statistic are derived but no non-asymptotic analysis of the methods is conducted. As far as we know, this paper is the first to provide a non-asymptotic condition on the classes of alternatives ensuring that the statistical test reaches a prescribed power working in a dependent framework.

1.2 Outline

In Section 2, we introduce useful notations for our paper and we present the concentration inequality for U-statistics that is an important argument of our proofs. The next three sections are dedicated to our main results. We start by providing a convergence result for the estimation of spectra of integral operators with MCMC algorithms (see Section 3). We show that independent Hastings algorithms satisfy under mild conditions the assumptions of Section 2.2 and we illustrate our result with the estimation of the spectra of some Mercer kernels. For the second application of our concentration inequality, we investigate the generalization performance of online algorithms with pairwise loss functions in a Markovian framework (see Section 4). We motivate the study of such problems and we provide an online-to-batch conversion result. In a third and final application, we propose a goodness-of-fit test for the density of the stationary measure of a Markov chain (see Section 5). We give an explicit condition on the set of alternatives to ensure that the statistical test proposed reaches a prescribed power. The proofs related to the three applications are given in Section A, Section B and Sections C.1-C.3 respectively.

2. Notations and Concentration inequality for U-statistics with dependent random variables

2.1 Notations

Let us consider an arbitrary measurable space (F, \mathcal{F}) . For any measure ω on (F, \mathcal{F}) , the total variation norm of ω is defined by $\|\omega\|_{\text{TV}} := \sup_{A \in \mathcal{F}} |\omega(A)|$. The space of square summable functions on F with respect to the measure ω defined by

$$L^2(\omega) := \{f : F \rightarrow \mathbb{R} \text{ measurable} \mid \int_F f(x)^2 d\omega(x) < \infty\},$$

endowed with the inner product

$$(f, g) \in L^2(\omega) \times L^2(\omega) \mapsto \langle f, g \rangle := \int_F f(x)g(x)d\omega(x),$$

is a Hilbert space and we denote by $\|\cdot\|_2$ the norm induced by $\langle \cdot, \cdot \rangle$. For any function $h : F \rightarrow \mathbb{R}$, we define the supremum norm of h by $\|h\|_\infty := \sup_{x \in F} |h(x)|$. We denote by $\mathcal{B}(\mathbb{R})$ the Borel algebra on \mathbb{R} and we set $\mathbb{N}^* := \mathbb{N} \setminus \{0\}$. For any $x \in \mathbb{R}_+$, we denote by $\lfloor x \rfloor$ (resp. $\lceil x \rceil$) the largest integer that is less than or equal to x (resp. the smallest integer greater than or equal to x). For any $x, y \in \mathbb{R}$, we set $x \vee y := \max(x, y)$ and $x \wedge y := \min(x, y)$. Given a sequence of real valued random variables $(X_n)_{n \in \mathbb{N}}$ and a sequence of positive reals $(a_n)_{n \in \mathbb{N}}$, the notation $X_n = \mathcal{O}_{\mathbb{P}}(a_n)$ means that $(X_n/a_n)_{n \in \mathbb{N}}$ converges to zero in probability as $n \rightarrow \infty$.

2.2 Concentration inequality for U-statistics of uniformly ergodic Markov chains

In this section, we present the concentration result from Duchemin et al. (2022) for U-statistics of uniformly ergodic discrete time Markov chains that will be an essential tool in our proofs. Let us mention that we do not work with the concentration inequality from Shen et al. (2020)

since it only holds for stationary chains if the kernel h is π -canonical (see Assumption 3). Stationarity may be seen as a strong assumption which would make our main results from Section 3 of little interest since MCMC methods are used when we are not able to directly sample from the distribution π . Regarding Sections 4 and 5, the concentration inequality for U-statistics used in the proofs of our results needs to hold for any initial distribution of the chain.

Let (E, Σ) be a measurable space. We consider a Markov chain $(X_i)_{i \geq 1}$ on (E, Σ) with transition kernel $P : E \times E \rightarrow [0, 1]$ and with a unique stationary distribution π . We consider a measurable function $h : (E \times E, \Sigma \otimes \Sigma) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ and we are interested in the following U-statistic

$$U_{\text{stat}}(n) := \sum_{1 \leq i \neq j \leq n} (h(X_i, X_j) - \mathbb{E}_{(X,Y) \sim \pi \otimes \pi} [h(X, Y)]).$$

We will work under the following set of assumptions.

Assumption 1 *The Markov chain $(X_i)_{i \geq 1}$ is ψ -irreducible (cf. Meyn and Tweedie, 1993, Section 4.2) for some maximal irreducibility measure ψ on Σ . Moreover, there exist some natural number m and a constant $\delta_m > 0$ such that*

$$\forall x \in E, \forall A \in \Sigma, \quad \delta_m \mu(A) \leq P^m(x, A). \tag{1}$$

for some probability measure μ .

A Markov chain satisfying Assumption 1 is called uniformly ergodic (cf. Meyn and Tweedie, 1993, Chapter 16) and admits a unique stationary distribution denoted by π . Assumption 1 also implies that the regeneration times associated to the split chain are exponentially integrable, meaning that their Orlicz norm with respect to the function $\psi_1(x) = \exp(x) - 1$ are bounded by some constant $\tau > 0$. We refer to (Duchemin et al., 2022, Section 2.3) for details.

Assumption 2 can be read as a reverse Doeblin's condition and is used in Duchemin et al. (2022) as a decoupling tool. In their paper, the authors give several natural examples for which this condition holds.

Assumption 2 *There exist $\delta_M > 0$ and some probability measure ν such that*

$$\forall x \in E, \forall A \in \Sigma, \quad P(x, A) \leq \delta_M \nu(A).$$

The last assumption introduces the notion of π -canonical kernel, which is the counterpart in the Markovian setting of the canonical (or degenerate) property of the independent framework.

Assumption 3 *Denoting by π the stationary distribution of the Markov chain $(X_i)_{i \geq 1}$, we assume that $h : (E \times E, \Sigma \otimes \Sigma) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ is measurable, bounded and is π -canonical, namely*

$$\forall x, y \in E, \quad \mathbb{E}_\pi[h(X, x)] = \mathbb{E}_\pi[h(X, y)] = \mathbb{E}_\pi[h(x, X)] = \mathbb{E}_\pi[h(y, X)].$$

This common expectation will be denoted $\mathbb{E}_\pi[h]$.

Let us mention that several important kernels are π -canonical. This is the case of translation-invariant kernels which have been widely studied in the Machine Learning community (cf. Lerasle et al. (2016)). Another example of π -canonical kernel is a rotation invariant kernel when $E = \mathbb{S}^{d-1} := \{x \in \mathbb{R}^d : \|x\|_2 = 1\}$ with π also rotation invariant (cf. De Castro et al., 2019; Duchemin and De Castro, 2022). Note also that if the kernel h is not π -canonical, the U-statistic decomposes into a linear term and a π -canonical U-statistic. This is called the *Hoeffding decomposition* (cf. Giné and Nickl, 2016, p.176) and takes the following form

$$\begin{aligned} \sum_{i \neq j} (h(X_i, X_j) - \mathbb{E}_{(X,Y) \sim \pi \otimes \pi} [h(X, Y)]) &= \sum_{i \neq j} \left(\tilde{h}(X_i, X_j) - \mathbb{E}_\pi [\tilde{h}(X, \cdot)] \right) \\ &+ \sum_{i \neq j} (\mathbb{E}_{X \sim \pi} [h(X, X_j)] - \mathbb{E}_{(X,Y) \sim \pi \otimes \pi} [h(X, Y)]) \\ &+ \sum_{i \neq j} (\mathbb{E}_{X \sim \pi} [h(X_i, X)] - \mathbb{E}_{(X,Y) \sim \pi \otimes \pi} [h(X, Y)]), \end{aligned}$$

where the kernel \tilde{h} is π -canonical with

$$\forall x, y \in E, \quad \tilde{h}(x, y) = h(x, y) - \mathbb{E}_{X \sim \pi} [h(x, X)] - \mathbb{E}_{X \sim \pi} [h(X, y)].$$

We will use this method several times in our proofs (for example in Eq.(19)).

We are now ready to state the result from Duchemin et al. (2022) that is one key theoretical tool to derive our three contributions presented in the next section.

Theorem 1 *Suppose that Assumptions 1, 2 and 3 are satisfied. Then there exist constants $\beta, \kappa > 0$ (depending on the Markov chain $(X_i)_{i \geq 1}$) such that for any $u \geq 1$ and any $n \geq 2$, with probability at least $1 - \beta e^{-u} \log n$,*

$$\frac{2}{n(n-1)} U_{\text{stat}}(n) \leq \kappa \|h\|_\infty \log n \left\{ \frac{u}{n} + \left(\frac{u}{n} \right)^2 \right\}.$$

3. Estimation of spectra of signed integral operator with MCMC algorithms

3.1 MCMC estimation of spectra of signed integral operators

Let us consider a Markov chain $(X_n)_{n \geq 1}$ on E satisfying the assumptions of Theorem 1 with stationary distribution π , and some symmetric kernel $h : E \times E \rightarrow \mathbb{R}$ such that $h \in L^2(\pi \otimes \pi)$. We can associate to h the kernel of a linear operator \mathbf{H} defined by

$$\mathbf{H}f(x) := \int_E h(x, y) f(y) d\pi(y). \quad (2)$$

This is a Hilbert-Schmidt operator on $L^2(\pi)$ and thus it has a real spectrum consisting of a square summable sequence of eigenvalues (cf. Conway, 2019, p.267). In the following, we will denote the eigenvalues of \mathbf{H} by $\lambda(\mathbf{H}) := (\lambda_1, \lambda_2, \dots)$. For some $n \in \mathbb{N}^*$, we consider

$$\tilde{\mathbf{H}}_n := \frac{1}{n} (h(X_i, X_j))_{1 \leq i, j \leq n} \quad \text{and} \quad \mathbf{H}_n := \frac{1}{n} ((1 - \delta_{i,j})h(X_i, X_j))_{1 \leq i, j \leq n}, \quad (3)$$

with respective eigenvalues $\lambda(\tilde{\mathbf{H}}_n)$ and $\lambda(\mathbf{H}_n)$. Following (Koltchinskii and Giné, 2000, Section 2), we introduce in Definition 1 the rearrangement distance δ_2 which measures closeness of spectra.

Definition 1 *Given two sequences x, y of reals – completing finite sequences by zeros – such that*

$$\sum_i x_i^2 + y_i^2 < \infty,$$

we define the ℓ_2 rearrangement distance $\delta_2(x, y)$ as

$$\delta_2^2(x, y) := \inf_{\sigma \in \mathfrak{S}} \sum_i (x_i - y_{\sigma(i)})^2,$$

where \mathfrak{S} is the set of permutations of natural numbers. δ_2 is a pseudometric on ℓ_2 , where ℓ_2 is the Hilbert space of all square summable sequences.

Theorem 2 gives conditions ensuring that both the spectrum of \mathbf{H}_n and the one of $\tilde{\mathbf{H}}_n$ converge towards the spectrum of the integral operator \mathbf{H} as $n \rightarrow \infty$. Theorem 2 holds under Assumption 4 that we discuss in details in Section 3.2. The proof of Theorem 2 is postponed to Section A.

Assumption 4 *$h : E \times E \rightarrow \mathbb{R}$ is a bounded and symmetric function square integrable with respect to $\pi \otimes \pi$. Moreover there exist continuous functions $\varphi_r : E \rightarrow \mathbb{R}$, $r \in I$ (where $I = \mathbb{N}$ or $I = 1, \dots, N$) that form an orthonormal basis of $L^2(\pi)$ and a sequence of real numbers $(\lambda_r)_{r \in I} \in \ell_2$ such that we have pointwise*

$$h(x, y) = \sum_{r \in I} \lambda_r \varphi_r(x) \varphi_r(y),$$

with $\Upsilon := \sup_{r \in I} \|\varphi_r\|_\infty^2 < \infty$ and $S := \sup_{x \in E} \sum_{r \in I} |\lambda_r| \varphi_r(x)^2 < \infty$.

We further denote $\Lambda := \sup_{r \in I} |\lambda_r|$.

Theorem 2 *Let $(X_i)_{i \geq 1}$ be a Markov chain on E satisfying Assumptions 1 and 2 described in Section 2.2 with stationary distribution π . Suppose that Assumption 4 is satisfied. Then for any $t > 0$,*

$$\begin{aligned} \mathbb{P} \left(\frac{1}{4} \delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))^2 \geq \frac{S^2(1 + \kappa) \log n}{n} + 2 \sum_{i > \lceil n^{1/4} \rceil, i \in I} \lambda_i^2 + t \right) \\ \leq 32\sqrt{n} \exp(-\mathcal{C} \min(nt^2, \sqrt{nt})) + \beta \log(n) \exp\left(-\frac{n}{\log n} \min(\mathcal{B}t, (\mathcal{B}t)^{1/2})\right), \end{aligned}$$

where for some universal constant $K > 0$, we have $\mathcal{B} = (K\kappa S)^{-1}$, $\mathcal{C} = (K^{1/2}m\tau(S + \Lambda\Upsilon))^{-2}$. $\kappa > 0$ and $\beta > 0$ are the constants from Theorem 1 and depend on the Markov chain. We refer to Assumption 1 and the following remark for the definitions of the constants m and τ .

Remark The same bound holds for $\delta_2(\lambda(\mathbf{H}), \lambda(\tilde{\mathbf{H}}_n))^2$.

3.2 Comparison with the existing literature

Previous works. In Adamczak and Bednorz (2015a), the authors studied the convergence properties of MCMC methods to estimate the spectrum of integral operators with bounded *positive* kernels (i.e. such that \mathbf{H} has non-negative eigenvalues). They show a sub-exponential tail behavior for the δ_2 distance between the spectrum of \mathbf{H} and the one of the random matrix \mathbf{H}_n . Their result has the merit to hold for geometrically ergodic Markov chains, but they work with the restrictive assumption that the eigenvalues of \mathbf{H} are non-negative. This makes their kernel of positive-type (cf. Eq.(14) Adamczak and Bednorz, 2015a), allowing them to use of powerful decoupling argument. They are reduced to study a sum of rank one operators on $L^2(\pi)$ and the proof is concluded by using the splitting technique (cf. Meyn and Tweedie, 1993, Section 5.1) with a Bernstein-type inequality for sums of independent Banach space valued random variables.

Comparing Theorem 2 with previous works. In this paper we consider signed integral operators and we cannot adapt the proof proposed by Adamczak and Bednorz (2015a). Working with stronger conditions on the Markov chain $(X_i)_{i \geq 1}$ compared to the former paper, Theorem 2 proves a new concentration inequality for the δ_2 distance between $\lambda(\mathbf{H})$ and $\lambda(\mathbf{H}_n)$ that holds for **arbitrary signs of the eigenvalues** of \mathbf{H} . We provide a synthetic description of our proof at the beginning of Section A . Note that the set of operators handled by Theorem 2 is not a superset of the ones handled by (Adamczak and Bednorz, 2015a, Theorem 2.2). The difference lies in the fact that we ask the family of functions $(\varphi_r)_{r \in I}$ to be uniformly bounded (cf. Assumption 4). Let us mention that the set of assumptions considered in Adamczak and Bednorz (2015a) implies that $S < \infty$. In the following, we comment our extra assumption $\Upsilon < \infty$ with more details.

1. The basis functions $(\varphi_r)_{r \in I}$ are continuous and Assumption 2 typically holds for a compact space E . Hence, by considering that E is compact and that the sequence $(\lambda_r)_{r \in I}$ has finite support (i.e. $I = [N]$ for some natural number N), it holds $\Upsilon < \infty$.
2. Asking for $\Upsilon < \infty$ is only useful to apply a concentration inequality for Markov chains at one specific step of our proof (cf. Eq.(15)). Hence this assumption might be weakened by applying other exponential tail control for empirical processes of Markov chains. Nevertheless we point out that Theorem 2 holds for uniformly ergodic Markov chains which is equivalent to the standard *drift condition* where the drift function V is bounded (cf. Meyn and Tweedie, 1993, Chap.16). Hence, the assumptions needed for the exponential inequality from (Adamczak and Bednorz, 2015b, Theorem 1.1) or the one from (Durmus et al., 2021, Theorem 5) imply that $\Upsilon < \infty$. Hence, weakening the condition $\Upsilon < \infty$ seems challenging but we believe that it might be done in some specific settings using for instance the work from (Ciolek and Bertail, 2019, Section 3.2).

3.3 Admissible sampling schemes: Independent Hastings algorithm

One can use the previous result to estimate the spectrum of the integral operator \mathbf{H} using MCMC methods. To do so, we need to make sure that the Markov chain used for the MCMC method satisfies the conditions of Theorem 1. It is for example well known that Metropolis random walks on \mathbb{R} are not uniformly ergodic (cf. Meyn and Tweedie, 1993). In the following,

we show that an independent Hastings algorithm can be used on bounded state space to generate a uniformly ergodic chain with the desired stationary distribution.

3.3.1 INDEPENDENT HASTINGS ALGORITHM ON BOUNDED STATE SPACE.

Let $E \subset \mathbb{R}^k$ be a bounded subset of \mathbb{R}^k equipped with the Borel σ -algebra $\mathcal{B}(E)$. We consider a density π which is only known up to a factor and a probability density q with respect to the Lebesgue measure λ_{Leb} on E , satisfying $\pi(y), q(y) > 0$ for all $y \in E$. In the independent Hastings algorithm, a candidate transition generated according to the law $q\lambda_{Leb}$ is then accepted with probability $\alpha(x, y)$ given by

$$\alpha(x, y) = \min \left(1, \frac{\pi(y)q(x)}{\pi(x)q(y)} \right).$$

With an approach similar to Theorem 2.1 from Mengersen and Tweedie (1996), Proposition 2 shows that under some conditions on the densities π and q , the independent Hastings algorithm satisfies the Assumptions 1 and 2.

Proposition 2 *Let us assume that $\sup_{x \in E} q(x) < \infty$ and that there exists $\beta > 0$ such that*

$$\frac{q(y)}{\pi(y)} > \beta, \quad \forall y \in E.$$

Then, the independent Hastings algorithm satisfies the Assumptions 1 and 2.

Proof We denote P the transition kernel of the Markov chain generated with the independent Hastings algorithm. For any $x \in E$, the density with respect to λ_{Leb} of the absolutely continuous part of $P(x, dy)$ is $p(x, \cdot) = q(\cdot)\alpha(x, \cdot)$, while the singular part is given by $\mathbb{1}_x(\cdot) \left(\int_{z \in E} q(z)\alpha(x, z)d\lambda_{Leb}(z) \right)$. For fixed $x \in E$, we have either $\alpha(x, y) = 1$ in which case $p(x, y) = q(y) \geq \beta\pi(y)$, or else

$$p(x, y) = q(y) \frac{\pi(y)q(x)}{\pi(x)q(y)} = q(x) \frac{\pi(y)}{\pi(x)} \geq \beta\pi(y).$$

We deduce that for any $x \in E$, it holds

$$P(x, A) \geq \beta \int_{y \in A} \pi(y)d\lambda_{Leb}(y),$$

which proves that the chain is uniformly ergodic (cf. Eq.(1)). Hence Assumption 1 is satisfied. Assumption 2 trivially holds since E is bounded and $\sup_{y \in E} q(y) < \infty$. ■

From Proposition 2 and Theorem 2, we deduce that one can use a MCMC approach to estimate the spectrum of a signed integral operator \mathbf{H} (that satisfies Assumption 4) as defined in Eq.(2) where E is a bounded subset of \mathbb{R}^k . More precisely, if the density π of Eq.(2) is known up to a factor and if there exists some probability density q with respect to λ_{Leb} satisfying the assumptions of Proposition 2, the Independent Hastings algorithm provides a Markov chain that satisfies Assumptions 1 and 2. Hence the non-asymptotic bound from Theorem 2 holds. We put this methodology into action in the new section by estimating the spectrum of some Mercer kernels on the d -dimensional sphere.

3.4 Estimation of the spectrum of Mercer kernels

In this example, we illustrate Theorem 2 by computing the eigenvalues of an integral operator naturally associated with a Mercer kernel using a MCMC algorithm. A function $h : E \times E \rightarrow \mathbb{R}$ is called a Mercer kernel if E is a compact metric space and $h : E \times E \rightarrow \mathbb{R}$ is a continuous symmetric and positive definite function. It is well known that if h is a Mercer kernel, then the integral operator L_h associated with h is a compact and bounded linear operator, self-adjoint and semi-definite positive. The spectral theorem implies that if h is a Mercer kernel, then there is a complete orthonormal system $(\varphi_1, \varphi_2, \dots)$ of eigenvectors of L_h . The eigenvalues $(\lambda_1, \lambda_2, \dots)$ are real and non-negative. The Mercer Theorem (see e.g. Christmann and Steinwart, 2008, Theorem 4.49) shows that the eigen-structure of L_h can be used to get a representation of the Mercer kernel h as a sum of a convergent sequence of product functions for the uniform norm. In this context, Theorem 2 allows to derive the convergence rate in the δ_2 metric of the estimated spectrum towards the one of the integral operator \mathbf{H} as presented in Proposition 3.

Proposition 3 *We keep the notations and the assumptions of Theorem 2. We assume further that there exists $s > 0$, a (Sobolev) regularity parameter, such that for some constant $C(s) > 0$,*

$$\forall R > 1, \quad \sum_{i>R} \lambda_i^2 \leq C(s)R^{-2s}.$$

Then it holds

$$\delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))^2 = \begin{cases} \mathcal{O}_{\mathbb{P}}\left(\sqrt{\frac{\log n}{n}}\right) & \text{if } s \geq 1 \\ \mathcal{O}_{\mathbb{P}}\left(\frac{1}{n^{s/2}}\right) & \text{if } s \in (0, 1) \end{cases}.$$

Proof Proposition 3 directly follows from Theorem 2 by choosing $t = \sqrt{\frac{\log n}{n}}$. ■

To illustrate our purpose, we consider the d -dimensional sphere $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\|_2 = 1\}$. We consider a positive definite kernel on \mathbb{S}^{d-1} defined by $\forall x, y \in \mathbb{S}^{d-1}, \quad h(x, y) = \psi(x^\top y)$ where $\psi : [-1, 1] \rightarrow \mathbb{R}$ is continuous. From the Funk-Hecke Theorem (see e.g. Müller, 2012, p.30), we know that the eigenvalues of the Mercer kernel h are

$$\lambda_k = |\mathbb{S}^{d-2}| \int_{-1}^1 \psi(t) P_k(d; t) (1 - t^2)^{\frac{d-3}{2}} dt, \tag{4}$$

where $|\mathbb{S}^{d-2}|$ is the surface area of \mathbb{S}^{d-2} and $P_k(d; t)$ is the Legendre polynomial of degree k in dimension d . For any $k \in \mathbb{N}$, the multiplicity of the eigenvalue λ_k is the dimension of the space of spherical harmonics of degree k . To build the Markov chain $(X_i)_{i \geq 1}$, we start by sampling randomly X_1 on \mathbb{S}^{d-1} . Then, for any $i \in \{2, \dots, n\}$, we sample


- a unit vector $Y_i \in \mathbb{S}^{d-1}$ uniformly, orthogonal to X_{i-1} .
- a real $r_i \in [-1, 1]$ encoding the distance between X_{i-1} and X_i . r_i is sampled from a distribution $f_{\mathcal{L}} : [-1, 1] \rightarrow [0, 1]$.

then X_i is defined by

$$X_i = r_i \times X_{i-1} + \sqrt{1 - r_i^2} \times Y_i.$$

By assuming that $\min_{r \in [-1,1]} f_{\mathcal{L}}(r) > 0$ and $\|f_{\mathcal{L}}\|_{\infty} < \infty$, Assumptions 1 and 2 hold and the stationary distribution of the chain $(X_i)_{i \geq 1}$ is the Haar measure on \mathbb{S}^{d-1} (cf. Duchemin and De Castro, 2022).

In Figure 1, we plot the non-zero eigenvalues using function $\psi : t \mapsto (1+t)^2$ and taking $f_{\mathcal{L}}$ proportional to $r \mapsto f_{(5,1)}(\frac{r+2}{4})$ where $f_{(5,1)}$ is the pdf of the Beta distribution with parameter $(5, 1)$. We plot both the true eigenvalues and the ones computed using a MCMC approach. Figure 1: Consider function $\psi : t \mapsto (1+t)^2$, $d = 2$ and $n = 1000$. The true eigenvalues can be computed using (4), but in this case, we know the exact values of the three non-zero eigenvalues namely $\lambda_0 = 3\pi$, $\lambda_1 = 2\pi$ and $\lambda_2 = \pi/2$. Their respective multiplicities are 1, 2 and 2. The estimated eigenvalues are the eigenvalues of the matrix $\mathbf{H}_n = \frac{1}{n} \left((1 - \delta_{i,j}) \psi(X_i^{\top} X_j) \right)_{1 \leq i, j \leq n}$ where the n points X_1, X_2, \dots, X_n are sampled on the Euclidean sphere \mathbb{S}^{d-1} using a Markovian dynamic.



4. Online Learning with Pairwise Loss Functions

4.1 Brief introduction to online learning and motivations

4.1.1 PRESENTATION OF THE TRADITIONAL ONLINE LEARNING SETTING

Online learning is an active field of research in Machine Learning in which data becomes available in a sequential order and is used to update the best predictor for future data at each step. This method aims at learning some function $f : E \rightarrow \mathcal{Y}$ where E is the space of inputs and \mathcal{Y} is the space of outputs. At each time step t , we observe a new example $(x_t, y_t) \in E \times \mathcal{Y}$. Traditionally, the random variables (x_t, y_t) are supposed i.i.d. with common joint probability distribution $(x, y) \mapsto p(x, y)$ on $E \times \mathcal{Y}$. In this setting, the loss function is given as $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, such that $\ell(f(x), y)$ measures the difference between the predicted value $f(x)$ and true value y . The goal is to select at each time step t a function $h_t : E \rightarrow \mathcal{Y}$ in a fixed set \mathcal{H} based on the observed examples until time t (namely $(x_i, y_i)_{1 \leq i \leq t}$) such that h_t has “small” risk \mathcal{R} defined by

$$\mathcal{R}(h) = \mathbb{E}_{(X,Y) \sim p} [\ell(h(X), Y)],$$

where h is any measurable mapping from E to \mathcal{Y} .

Online learning is used when data is coming *on the fly* and we do not want to wait for the acquisition of the complete dataset to take a decision. In such cases, online learning algorithms allow to dynamically adapt to new patterns in the data.

4.1.2 ONLINE LEARNING WITH PAIRWISE LOSS FUNCTIONS

In some cases, the framework provided in the previous paragraph is not appropriated to solve the task at stake. Consider the example of ranking problems. The state space is E and there exists a function $f : E \rightarrow \mathbb{R}$ which assigns to each state $x \in E$ a label $f(x) \in \mathbb{R}$. f naturally defines a partial order on E . At each time step t , we observe an example $x_t \in E$ together with its label $f(x_t)$ and we suppose that the random variables $(x_t)_t$ are i.i.d. with common distribution p . Our goal is to learn the partial order of the items in E induced by the function f . More precisely, we consider a space $\mathcal{H} \subset \{h : E \times E \rightarrow \mathbb{R}\}$, called the set of hypotheses. An *ideal* hypothesis $h \in \mathcal{H}$ would satisfy

$$\forall x, u \in E, \quad f(x) \geq f(u) \Leftrightarrow (h(x, u) \geq 0 \text{ and } h(u, x) \leq 0).$$

We consider a loss function $\ell : \mathcal{H} \times E \times E \rightarrow \mathbb{R}$ such that $\ell(h, x, u)$ measures the ranking error induced by h and a typical choice is the 0-1 loss

$$\ell(h, x, u) = \mathbb{1}_{\{(f(x)-f(u))h(x,u)<0\}}.$$

U-statistics naturally arise in such settings as for example in Cléménçon et al. (2008) where Cléménçon and al. study the consistency of the empirical risk minimizer of ranking problems using the theory of U-processes in an i.i.d. framework.

Example: Bipartite ranking problems

*We describe the concrete problem of bipartite ranking. We consider that we have as input a training set of examples. Each example is described by some feature vector and is associated with a binary label. Typically one can consider that we have access to health data of an individual along time. We know at each time step her/his health status x_t and her/his label which is 0 if the individual is healthy and 1 if she/he is sick. In the bipartite ranking problem, we want to learn a **scorer** which maps any feature vector describing the health status of the individual to a real number such that sick states have a higher score than healthy ones. Following the health status of individuals is time-consuming and we cannot afford to wait for the end of the data acquisition process to understand the relationship between the feature vector describing the health status of the individual and her/his sickness. In such settings where data is coming on the fly, online algorithms are common tools that allow to learn a scorer function along time. At each time step the scorer function is updated based on the new measurement provided.*

4.1.3 GENERALIZATION BOUNDS FOR ONLINE LEARNING

The performance of online learning algorithms is often analyzed with the notion of *regret* which compares the payoff obtained by the algorithm along time with the one that would have been obtained by taking the optimal decision at each time step (cf. Hoi et al., 2021; Bubeck and Cesa-Bianchi, 2012). It is natural to wonder if stronger theoretical guarantees can be obtained when some probabilistic structure underlies the sequence of examples, or loss functions, presented to the online algorithm. As asked in Agarwal and Duchi (2012), ‘if the sequence of examples are generated by a stochastic process, can the online learning algorithm output a good predictor for future samples from the same process?’ In other words,

we want to study the generalization ability of some online learner that generates a sequence of hypothesis $(h_t)_{t \geq 1}$ by bounding with high probability the excess risk defined as

$$\frac{1}{n} \sum_{t=1}^n \mathcal{R}(h_t) - \min_{h \in \mathcal{H}} \mathcal{R}(h).$$

Generalization bounds for online learning with pairwise loss functions working with i.i.d. samples have been considered for quite a while in both Machine Learning and Statistics literature (cf. Kar et al., 2013; Ying and Zhou, 2017; Guo et al., 2017; Chen and Lei, 2018). For dependent data sequences, generalization bounds for online algorithms have also been proved in the last decades with univariate loss functions (cf. Zhang, 2005; Xu et al., 2014; Agarwal and Duchi, 2012). However, theoretical guarantees for the generalization performance of online algorithms with pairwise loss functions with non i.i.d. data have been so far little studied. A quick and incomplete review of the literature is presented in Table 1.

	Univariate loss function	Pairwise loss function
i.i.d. data	Hoi et al., Section 3.7 and references therein	Kar et al.; Ying and Zhou; Guo et al.; Chen and Lei
Dependent data	Zhang; Xu et al.; Agarwal and Duchi	Our work

Table 1: Overview of the literature providing generalization bounds for online learning algorithms.

4.1.4 GENERALIZATION BOUNDS FOR PAIRWISE ONLINE LEARNING WITH DEPENDENT DATA

Connection with the existing literature. As far as we know, the few papers that investigate the generalization performance of pairwise online learning algorithms with non i.i.d. data have studied specific algorithms and/or specific learning tasks (cf. Qin et al., 2021; Zeng et al., 2021). In Zeng et al. (2021), the authors analyze online pairwise support vector machine while the work Qin et al. (2021) is focused on online regularized pairwise learning algorithm with least squares loss function. One possible reason explaining this gap in the literature is that *‘for pairwise learning [where] pairs of training examples are not i.i.d., [...] standard techniques can not be directly applied.’* (cf. Zeng et al., 2021).

With the upcoming application, we are the first - as far as we know - to provide a generalization bounds for online algorithms with pairwise loss functions and Markov chain samples that hold for an arbitrary online learner, covering a large span of settings.

Online learning with a Markovian dynamic. The theoretical analysis of Machine Learning algorithms with an underlying Markovian distribution of the data has become a very active field of research. The first papers to study online learning with samples drawn from non-identical distributions were Smale and Zhou (2009) and Steinwart et al. (2009) where online learning for least square regression and off-line support vector machines are investigated. In Zou et al. (2009), the generalization performance of the empirical risk minimization algorithm is studied with uniformly ergodic Markov chain samples. Hence

the analysis of online algorithms with dependent samples is recent and several works make the assumption that the sequence is a uniformly ergodic Markov chain. We motivate the Markovian assumption on the example of the previous paragraph.

Example (continued): Interest in online algorithms with Markovian dynamic

The health status of the individual at time $n + 1$ is not independent from the past and a simple way to model this time evolution would be to consider that it only depends on the last measured health status namely the feature vector x_n . This is a Markovian assumption on the sequence of observed health status of the individual.

We have explained why pairwise loss functions capture ranking problems and naturally arise in several Machine Learning problems such as metric learning or bipartite ranking (cf. Cléménçon et al. (2008)). We have shown the interest to provide a generalization bounds for online learning with pairwise loss functions with a Markovian assumption on the distribution of the sequence of examples and this is the goal of the next section.

4.2 Online-to-batch conversion for pairwise loss functions with Markov chains

We consider a reversible Markov chain $(X_i)_{i \geq 1}$ with state space E satisfying Assumption 1 with stationary distribution π . Using (Meyn and Tweedie, 1993, Theorem 16.0.2), we deduce that there exist constants $0 < \rho < 1$ and $L > 0$ such that

$$\|P^n(x, \cdot) - \pi\|_{\text{TV}} \leq L\rho^n, \quad \forall n \geq 0, \pi\text{-a.e } x \in E. \tag{5}$$

We assume that we have a function $f : E \rightarrow \mathbb{R}$ which defines the ordering of the objects in E . We aim at finding a relevant approximation of the ordering of the objects in E by selecting a function h (called a *hypothesis* function) in a space \mathcal{H} based on the observation of the random sequence $(X_i, f(X_i))_{1 \leq i \leq n}$. To measure the performance of a given hypothesis $h : E \times E \rightarrow \mathbb{R}$, we use a pairwise loss function of the form $\ell(h, X, U)$. Typically, one could use the *misranking loss* defined by

$$\ell(h, x, u) = \mathbb{1}_{\{(f(x) - f(u))h(x, u) < 0\}},$$

which is 1 if the examples are ranked in the wrong order and 0 otherwise. The goal of the learning problem is to find a hypothesis h which minimizes the *expected misranking risk*

$$\mathcal{R}(h) := \mathbb{E}_{(X, X') \sim \pi \otimes \pi} [\ell(h, X, X')].$$

We show that the investigation of the generalization performance of online algorithms with pairwise loss functions provided by Wang et al. (2012) can be extended to a Markovian framework. Our contribution is two-fold.

- Firstly, we prove that with high probability, the average risk of the sequence of hypotheses generated by an arbitrary online learner is bounded by some easily computable statistic.
- This first technical result is then used to show how we can extract a low risk hypothesis from a given sequence of hypotheses selected by an online learner. This is an *online-to-batch* conversion for pairwise loss functions with a Markovian assumption on the distribution of the observed states.

Given a sequence of hypotheses $(h_i)_{1 \leq i \leq n} \in \mathcal{H}^n$ generated by any online algorithm, we define the *average paired empirical risk* $\mathcal{M}^n(h_1, \dots, h_{n-1-b_n})$ (see Eq.(6)) averaging the *paired empirical risks* M_t (see Eq.(7)) of hypotheses h_{t-b_n} when paired with X_t as follows

$$\mathcal{M}^n(h_1, \dots, h_{n-1-b_n}) := \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} M_t, \quad (6)$$

$$\text{and } M_t := \frac{1}{t - b_n} \sum_{i=1}^{t-b_n} \ell(h_{t-b_n}, X_t, X_i), \quad (7)$$

where

$$c_n = \lceil c \times n \rceil \text{ for some } c \in (0, 1) \quad \text{and} \quad b_n = \lfloor q \log(n) \rfloor, \quad (8)$$

for an arbitrarily chosen $q > \frac{1}{\log(1/\rho)}$ where ρ is a constant related to the uniform ergodicity of the Markov chain, see Eq.(5). In the following, we will simply denote $\mathcal{M}^n(h_1, \dots, h_{n-1-b_n})$ by \mathcal{M}^n when the sequence of considered hypotheses is clear from the context.

M_t is the *paired empirical risk* of hypothesis h_{t-b_n} with X_t . It measures the performance of the hypothesis h_{t-b_n} on the example X_t when paired with examples seen before time $t - b_n$. \mathcal{M}^n is the mean value of a proportion $1 - c$ of these paired empirical risks. Hence the parameter $c \in (0, 1)$ controls the proportion of hypotheses h_{t-b_n} whose paired empirical risk M_t does not appear in the average paired empirical risk value \mathcal{M}^n . The parameter b_n controls the time gaps between elements of pairs (X_t, X_i) appearing in Eq.(7) in such way that their joint law is close to the product law $\pi \otimes \pi$ (mixing of the chain is met). The use of the burning parameter b_n is the main difference with the work Wang et al. (2012) when defining \mathcal{M}^n and M_t in Eq.(6) and Eq.(7). From a pragmatic point of view,

- we discard the first hypotheses that are not reliable, namely we do not consider hypothesis h_i for $i \leq c_n - b_n$. These first hypotheses are considered as not reliable since the online learner selected them based on a too small number of observed examples.
- since h_{t-b_n} is learned from X_1, \dots, X_{t-b_n} , we test the performance of h_{t-b_n} on X_t (and not on some X_i with $t - b_n + 1 \leq i < t$) to ensure that the distribution of X_t conditionally on $\sigma(X_1, \dots, X_{t-b_n})$ is approximately the stationary distribution of the chain π (see Assumption 1 and Equation (5)). Stated otherwise, this ensures that sufficient mixing has occurred.

Note that we assume n large enough to ensure that $c_n - b_n \geq 1$. For any $\eta > 0$, we denote $\mathcal{N}(\mathcal{H}, \eta)$ the L^∞ η -covering number for the hypothesis class \mathcal{H} (see Definition 4).

Definition 4 (cf. Wainwright, 2019, Chapter 5.1) *Let us consider some $\eta > 0$. A L^∞ η -cover of a set \mathcal{H} is a set $\{g_1, \dots, g_N\} \subset \mathcal{H}$ such that for any $h \in \mathcal{H}$, there exists some $i \in \{1, \dots, N\}$ such that $\|g_i - h\|_\infty \leq \eta$. The L^∞ η -covering number $\mathcal{N}(\mathcal{H}, \eta)$ is the cardinality of the smallest L^∞ η -cover of the set \mathcal{H} .*

Theorem 3 bounds the average risk of the sequence of hypotheses in terms of its empirical counterpart \mathcal{M}^n and is proved in Section B.1.

Theorem 3 *Assume that the Markov chain $(X_i)_{i \geq 1}$ is reversible and satisfies Assumption 1. Assume the hypothesis space $(\mathcal{H}, \|\cdot\|_\infty)$ is compact. Let $h_0, h_1, \dots, h_n \in \mathcal{H}$ be the ensemble of hypotheses generated by an arbitrary online algorithm working with a pairwise loss function ℓ such that,*

$$\ell(h, x_1, x_2) = \varphi(f(x_1) - f(x_2), h(x_1, x_2)),$$

where $\varphi : \mathbb{R} \times \mathbb{R} \rightarrow [0, 1]$ is a Lipschitz function w.r.t. the second variable with a finite Lipschitz constant $Lip(\varphi)$. Let $\xi > 0$ be an arbitrary positive number and let us consider $q = \frac{\xi+1}{\log(1/\rho)}$ for the definition of b_n (see Eq.(8)). Then for all $c > 0$ and for all $\varepsilon > 0$ such that $\varepsilon \underset{n \rightarrow \infty}{=} o(n^\xi)$, we have for sufficiently large n

$$\begin{aligned} \mathbb{P} \left(\left| \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \mathcal{M}^n \right| \geq \varepsilon \right) &\leq 2 \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8Lip(\varphi)} \right) + 1 \right] b_n \\ &\times \exp \left(- \frac{(c_n - b_n)C(m, \tau)\varepsilon^2}{16b_n^2} \right), \end{aligned}$$

where $C(m, \tau)^{-1} = 7 \times 10^3 \times m^2 \tau^2$. We refer to Assumption 1 and the following remark (or to Duchemin et al., 2022, Section 2) for the definitions of the constants m and τ that depend on the Markov chain $(X_i)_{i \geq 1}$.

Theorem 3 shows that average paired empirical risk \mathcal{M}^n (see Eq.(6)) is close to average risk given by

$$\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}).$$

Quantitative errors bounds can be given assuming that the L_∞ -metric entropy (l.h.s of the next equation) satisfies

$$\log \mathcal{N}(\mathcal{H}, \eta) = \mathcal{O}(\eta^{-\theta}), \tag{9}$$

where $\theta > 0$ is an exponent, depending on the dimension of state space E and the regularity of hypotheses of \mathcal{H} , that can be computed in some situations (Lipschitz function, higher order smoothness classes, see (Wainwright, 2019, Chapter 5.1) for instance). Theorem 4 made this statement rigorous (cf. Eq.(10)).

As previously mentioned, online learning algorithms are often studied through the lens of *regret*. The definition of a regret bound in our context is provided in Definition 5.

Definition 5 *An online learning algorithm will be said to have a regret bound \mathfrak{R}_n if it presents an ensemble h_1, \dots, h_{n-1} such that*

$$\mathcal{M}^n \leq \min_{h \in \mathcal{H}} \{ \mathcal{M}^n(h, \dots, h) \} + \mathfrak{R}_n.$$

In the literature of learning theory Cucker and Zhou (2007), we are often interested in the averaged excess generalization error

$$\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \mathcal{R}(h^*),$$

where h^* is the population risk minimizer and is given by $h^* \in \arg \min_{h \in \mathcal{H}} \mathcal{R}(h)$.

As a consequence, most of works focused on online-to-batch conversion are interested in the overall convergence rate of the excess generalization error for online learners that achieve a given regret bound. Examples can be found with (Guo et al., 2017, Corollary 4) or with (Kar et al., 2013, Theorem 5) where both papers work with pairwise loss functions with i.i.d. observations. In Theorem 4 (cf. Eq.(11)) we provide the overall rate for the averaged excess generalization error for an online learning satisfying a given regret bound. Theorem 4 is proved in Section B.2 and should be understood as an extension of the above mentioned results from Kar et al. (2013) and Guo et al. (2017).

Theorem 4 *We keep the notations and assumptions of Theorem 3. Assume further that \mathcal{H} satisfies Eq.(9). Then it holds*

$$\left| \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \mathcal{M}^n \right| = \mathcal{O}_{\mathbb{P}} \left[\frac{\log(n) \log(\log n)}{n^{\frac{1}{2+\theta}}} \right]. \quad (10)$$

Moreover, if the online learner has a regret bound \mathfrak{R}_n (cf. Definition 5), it holds

$$\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \mathcal{R}(h^*) = \mathcal{O}_{\mathbb{P}} \left[\frac{\log(n) \log(\log n)}{n^{\frac{1}{2+\theta}}} + \mathfrak{R}_n \right]. \quad (11)$$

4.3 Batch hypothesis selection

Theorems 3 and 4 are results on the performance of online learning algorithms. We will use these results to study the generalization performance of such online algorithms in the batch setting (see Theorem 5). Hence we are now interested in *selecting a good hypothesis from the ensemble of hypotheses generated by the online learner* namely that has a small empirical risk.

We measure the risk for h_{t-b_n} on the last $n - t$ examples of the sequence X_1, \dots, X_n , and penalize each h_{t-b_n} based on the number of examples on which it is evaluated. More precisely, let us define the empirical risk of hypothesis h_{t-b_n} on $\{X_{t+1}, \dots, X_n\}$ as

$$\widehat{\mathcal{R}}(h_{t-b_n}, t+1) := \binom{n-t}{2}^{-1} \sum_{k>i, i \geq t+1}^n \ell(h_{t-b_n}, X_i, X_k).$$

For a confidence parameter $\gamma \in (0, 1)$ that will be specified in Theorem 5, the hypothesis \widehat{h} is chosen to minimize the following penalized empirical risk,

$$\widehat{h} = h_{\widehat{t}-b_n} \quad \text{and} \quad \widehat{t} \in \arg \min_{c_n \leq t \leq n-1} \left(\widehat{\mathcal{R}}(h_{t-b_n}, t+1) + c_\gamma(n-t) \right), \quad (12)$$

where

$$c_\gamma(x) = \sqrt{\frac{C(m, \tau)^{-1}}{x} \log \frac{64(n-c_n)(n-c_n+1)}{\gamma}},$$

with $C(m, \tau)^{-1} = 7 \times 10^3 \times m^2 \tau^2$.

Theorem 5 proves that the model selection mechanism previously described select a hypothesis \widehat{h} from the hypotheses of an arbitrary online learner whose risk is bounded relative to \mathcal{M}^n . The proof of Theorem 5 is postponed to Section B.3.

Theorem 5 *Assume that the Markov chain $(X_i)_{i \geq 1}$ is reversible and satisfies Assumptions 1 and 2. Let h_0, \dots, h_n be the set of hypotheses generated by an arbitrary online algorithm \mathcal{A} working with a pairwise loss ℓ which satisfies the conditions given in Theorem 3. Let $\xi > 0$ be an arbitrary positive number and let us consider $q = \frac{\xi+1}{\log(1/\rho)}$ for the definition of b_n (see Eq.(8)). For all $\varepsilon > 0$ such that $\varepsilon \underset{n \rightarrow \infty}{=} o(n^\xi)$, if the hypothesis is selected via Eq.(12) with the confidence γ chosen as*

$$\gamma = 64(n - c_n + 1) \exp\left(- (n - c_n)\varepsilon^2 C(m, \tau)/128\right),$$

then, when n is sufficiently large, we have

$$\mathbb{P}\left(\mathcal{R}(\hat{h}) \geq \mathcal{M}^n + \varepsilon\right) \leq 32 \left[\mathcal{N}\left(\mathcal{H}, \frac{\varepsilon}{16\text{Lip}(\varphi)}\right) + 1 \right] \exp\left(-\frac{(c_n - b_n)C(m, \tau)\varepsilon^2}{(16b_n)^2} + 2\log n\right).$$

Analogously to the previous section, we can derive from Theorem 5 a bound for the excess risk of the selected hypothesis \hat{h} .

Corollary 1 *We keep the notations and assumptions of Theorem 5. Assume further that \mathcal{H} satisfies Eq.(9). Then it holds*

$$\left| \mathcal{R}(\hat{h}) - \mathcal{M}^n \right| = \mathcal{O}_{\mathbb{P}}\left[\frac{\log^2 n}{n^{\frac{1}{2+\theta}}}\right].$$

Moreover, if the online learner has a regret bound \mathfrak{R}_n (cf. Definition 5), it holds

$$\mathcal{R}(\hat{h}) - \mathcal{R}(h^*) = \mathcal{O}_{\mathbb{P}}\left[\frac{\log^2 n}{n^{\frac{1}{2+\theta}}} + \mathfrak{R}_n\right].$$

5. Adaptive goodness-of-fit tests in a density model

5.1 Goodness-of-fit tests and review of the literature

In its original formulation, the goodness-of-fit test aims at determining if a given distribution q matches some unknown distribution p from samples $(X_i)_{i \geq 1}$ drawn independently from p . Classical approaches to solve the goodness-of-fit problem use the empirical process theory. Most of the popular tests such as the Kolmogorov-Smirnov, Cramer-von Mises, and Anderson-Darling statistics are based on the empirical distribution function of the samples. Other traditional approaches may require space partitioning or closed-form integrals Baringhaus and Henze (1988), Beirlant et al. (2008). In Rudzkiš and BakshaeV (2013), a non-parametric method is proposed with a test based on a kernel density estimator. In the last decade, a lot of effort has been put into finding more efficient goodness-of-fit tests. The motivation was mainly coming from graphical models where the distributions are known up to a normalization factor that is often computationally intractable. To address this problem, several tests have been proposed based on Reproducing Kernel Hilbert Space (RKHS) embedding. A large span of them use classes of Stein transformed RKHS functions (Liu et al., 2016; Gorham and Mackey, 2017). For example in Chwialkowski et al. (2016), a goodness-of-fit test is proposed for both i.i.d or non i.i.d samples. The test statistic uses the squared Stein discrepancy,

which is naturally estimated by a V-statistic. One drawback of such approach is that the theoretical results provided are only asymptotic. This paper is part of a large list of works that proposed a goodness-of-fit test and where the use of U-statistics naturally emerge (cf. Liu et al., 2016; Fan, 1997; Butucea et al., 2007; Fan and Ullah, 1999; Fernández and Gretton, 2019; Fromont and Laurent, 2006). To conduct a non-asymptotic analysis of the goodness-of-fit tests proposed for non i.i.d samples, a concentration result for U-statistics with dependent random variables is much needed.

5.2 Goodness-of-fit test for the density of the stationary measure of a Markov chain

In this section, we provide a goodness-of-fit test for Markov chains whose stationary distribution has density with respect to the Lebesgue measure λ_{Leb} on \mathbb{R} . Our work is inspired from Fromont and Laurent (2006) where Fromont and Laurent tackled the goodness-of-fit test with i.i.d samples. Conducting a non-asymptotic theoretical study of our test, we are able to identify the classes of alternatives over which our method has a prescribed power.

Let X_1, \dots, X_n be a Markov chain with stationary distribution π with density f with respect to the Lebesgue measure on \mathbb{R} . Let f_0 be some given density in $L^2(\lambda_{Leb})$ and let α be in $]0, 1[$. Assuming that f belongs to $L^2(\lambda_{Leb})$, we construct a level α test of the null hypothesis " $f = f_0$ " against the alternative " $f \neq f_0$ " from the observation (X_1, \dots, X_n) . The test is based on the estimation of $\|f - f_0\|_2^2$ that is $\|f\|_2^2 + \|f_0\|_2^2 - 2\langle f, f_0 \rangle$. $\langle f, f_0 \rangle$ is usually estimated by the empirical estimator $\sum_{i=1}^n f_0(X_i)/n$ and the cornerstone of our approach is to find a way to estimate $\|f\|_2^2$. We follow the work of Fromont and Laurent (2006) and we introduce a set $\{S_m, m \in \mathcal{M}\}$ of linear subspaces of $L^2(\lambda_{Leb})$. For all m in \mathcal{M} , let $\{p_l, l \in \mathcal{L}_m\}$ be some orthonormal basis of S_m . The variable

$$\hat{\theta}_m = \frac{1}{n(n-1)} \sum_{l \in \mathcal{L}_m} \sum_{i \neq j=1}^n p_l(X_i)p_l(X_j)$$

estimates $\|\Pi_{S_m}(f)\|_2^2$ where Π_{S_m} denotes the orthogonal projection onto S_m . Then $\|f - f_0\|_2^2$ can be approximated by

$$\hat{T}_m = \hat{\theta}_m + \|f_0\|_2^2 - \frac{2}{n} \sum_{i=1}^n f_0(X_i),$$

for any m in \mathcal{M} . Denoting by $t_m(u)$ the $(1-u)$ quantile of the law of \hat{T}_m under the hypothesis " $f = f_0$ " and considering

$$u_\alpha = \sup_{u \in]0, 1[} \mathbb{P}_{f_0} \left(\sup_{m \in \mathcal{M}} (\hat{T}_m - t_m(u)) > 0 \right) \leq \alpha,$$

we introduce the test statistic T_α defined by

$$T_\alpha = \sup_{m \in \mathcal{M}} (\hat{T}_m - t_m(u_\alpha)). \tag{13}$$

The test consists in rejecting the null hypothesis if T_α is positive. This approach can be read as a multiple testing procedure. Indeed, for each m in \mathcal{M} , we construct a level u_α test

of the null hypothesis " $f = f_0$ " by rejecting this hypothesis if \widehat{T}_m is larger than its $(1 - u_\alpha)$ quantile under the hypothesis " $f = f_0$ ". We thus obtain a collection of tests and we decide to reject the null hypothesis if for some of the tests of the collection this hypothesis is rejected.

Now we define the different collection of linear subspaces $\{S_m, m \in \mathcal{M}\}$ that we will use in the following. We will focus on constant piecewise functions, scaling functions and, in the case of compactly supported densities, trigonometric polynomials.

- For all D in \mathbb{N}^* and $k \in \mathbb{Z}$, let

$$I_{D,k} = \sqrt{D} \mathbb{1}_{[k/D, (k+1)/D]}.$$

For all $D \in \mathbb{N}^*$, we define $S_{(1,D)}$ as the space generated by the functions $\{I_{D,k}, k \in \mathbb{Z}\}$ and

$$\widehat{\theta}_{(1,D)} = \frac{1}{n(n-1)} \sum_{k \in \mathbb{Z}} \sum_{i \neq j=1}^n I_{D,k}(X_i) I_{D,k}(X_j).$$

- Let us consider a pair of compactly supported orthonormal wavelets (φ, ψ) such that for all $J \in \mathbb{N}$, $\{\varphi_{J,k} = 2^{J/2} \varphi(2^J \cdot -k), k \in \mathbb{Z}\} \cup \{\psi_{j,k} = 2^{j/2} \psi(2^j \cdot -k), j \in \mathbb{N}, j \geq J, k \in \mathbb{Z}\}$ is an orthonormal basis of $L^2(\lambda_{Leb})$. For all $J \in \mathbb{N}$ and $D = 2^J$, we define $S_{(2,D)}$ as the space generated by the scaling functions $\{\varphi_{J,k}, k \in \mathbb{Z}\}$ and

$$\widehat{\theta}_{(2,D)} = \frac{1}{n(n-1)} \sum_{k \in \mathbb{Z}} \sum_{i \neq j=1}^n \varphi_{J,k}(X_i) \varphi_{J,k}(X_j).$$

- Let us consider the Fourier basis of $L^2([0,1])$ given by

$$\begin{aligned} g_0(x) &= \mathbb{1}_{[0,1]}(x), \\ g_{2p-1}(x) &= \sqrt{2} \cos(2\pi p x) \mathbb{1}_{[0,1]}(x) \quad \forall p \geq 1, \\ g_{2p}(x) &= \sqrt{2} \sin(2\pi p x) \mathbb{1}_{[0,1]}(x) \quad \forall p \geq 1. \end{aligned}$$

For all $D \in \mathbb{N}^*$, we define $S_{(3,D)}$ as the space generated by the functions $\{g_l, l = 0, \dots, D\}$ and

$$\widehat{\theta}_{(3,D)} = \frac{1}{n(n-1)} \sum_{l=0}^D \sum_{i \neq j=1}^n g_l(X_i) g_l(X_j).$$

We denote $\mathbb{D}_1 = \mathbb{D}_3 = \mathbb{N}^*$ and $\mathbb{D}_2 = \{2^J, J \in \mathbb{N}\}$. For l in $\{1, 2, 3\}$, D in \mathbb{D}_l , $\Pi_{S_{(l,D)}}$ denotes the orthogonal projection onto $S_{(l,D)}$ in $L^2(\lambda_{Leb})$. For all l in $\{1, 2, 3\}$, we take $\mathcal{D}_l \subset \mathbb{D}_l$ with $\cup_{l \in \{1,2,3\}} \mathcal{D}_l \neq \emptyset$ and $\mathcal{D}_3 = \emptyset$ if the X_i 's are not included in $[0,1]$. Let $\mathcal{M} = \{(l, D), l \in \{1, 2, 3\}, D \in \mathcal{D}_l\}$.

Theorem 6 describes classes of alternatives over which the corresponding test has a prescribed power. We work under the additional Assumption 5. We refer to Section C.1 for the proof of Theorem 6.

Assumption 5 *The initial distribution of the Markov chain $(X_i)_{i \geq 1}$, denoted χ , is absolutely continuous with respect to the stationary measure π and its density, denoted by $\frac{d\chi}{d\pi}$, has finite p -moment for some $p \in (1, \infty]$, i.e*

$$\infty > \left\| \frac{d\chi}{d\pi} \right\|_{\pi, p} := \begin{cases} \left[\int \left| \frac{d\chi}{d\pi} \right|^p d\pi \right]^{1/p} & \text{if } p < \infty, \\ \text{ess sup} \left| \frac{d\chi}{d\pi} \right| & \text{if } p = \infty. \end{cases}$$

In the following, denote $q = \frac{p}{p-1} \in [1, \infty)$ (with $q = 1$ if $p = +\infty$) which satisfies $\frac{1}{p} + \frac{1}{q} = 1$.

Theorem 6 *Let X_1, \dots, X_n a Markov chain on \mathbb{R} satisfying the Assumptions 1, 2 and 5 with stationary measure π . We assume that π has density f with respect the Lebesgue measure on \mathbb{R} and let f_0 be some given density. Let T_α be the test statistic defined by Eq.(13). Assume that f_0 and f belong to $L^\infty(\mathbb{R})$ (the space of essentially bounded measurable functions on \mathbb{R}) and that there exist $p_1, p_2 \in (1, +\infty]$ such that*

$$C_\chi := \left\| \frac{1}{f} \frac{d\chi}{d\lambda_{Leb}} \right\|_{f\lambda_{Leb}, p_1} \vee \left\| \frac{1}{f_0} \frac{d\chi}{d\lambda_{Leb}} \right\|_{f_0\lambda_{Leb}, p_2} < \infty,$$

where we used the notations of Assumption 5. We fix some γ in $]0, 1[$. For any $\varepsilon \in]0, 2[$, there exist some positive constants C_1, C_2, C_3 such that, setting for all $m = (l, D)$ in \mathcal{M} ,

$$\begin{aligned} V_m(\gamma) &= C_1 \|f\|_\infty \frac{\log(3C_\chi/\gamma)}{\varepsilon n} + C_2 (\|f\|_\infty \log(D+1) + \|f_0\|_\infty) \frac{\log(3C_\chi/\gamma)}{n} \\ &\quad + C_3 (\|f\|_\infty + 1) DR \left(n, \log \left\{ \frac{3\beta \log n}{\gamma} \right\} \right), \end{aligned}$$

with

$$R(n, u) = \log n \left\{ \frac{u}{n} + \left(\frac{u}{n} \right)^2 \right\},$$

if f satisfies

$$\|f - f_0\|_2^2 > (1 + \varepsilon) \inf_{m \in \mathcal{M}} \{ \|f - \Pi_{S_m}(f)\|_2^2 + t_m(u_\alpha) + V_m(\gamma) \}, \quad (14)$$

then

$$\mathbb{P}_f(T_\alpha \leq 0) \leq \gamma.$$

In order to make the condition (14) more explicit and to study its sharpness, we define the uniform separation rate which provides for any $\gamma \in (0, 1)$ the smallest distance between the set of null hypotheses and the set of alternatives to ensure that the power of our statistic test with level α is at least $1 - \gamma$.

Definition 6 *Given $\gamma \in]0, 1[$ and a class of functions $\mathcal{B} \subset L^2(\lambda_{Leb})$, we define the uniform separation rate $\rho(\Phi_\alpha, \mathcal{B}, \gamma)$ of a level α test Φ_α of the null hypothesis " $f \in \mathcal{F}$ " over the class \mathcal{B} as the smallest number ρ such that the test guarantees a power at least equal to $(1 - \gamma)$ for all alternatives $f \in \mathcal{B}$ at a distance ρ from \mathcal{F} . Stated otherwise, denoting by $d_2(f, \mathcal{F})$ the L^2 -distance between f and \mathcal{F} and by \mathbb{P}_f the distribution of the observation (X_1, \dots, X_n) ,*

$$\rho(\Phi_\alpha, \mathcal{B}, \gamma) = \inf \{ \rho > 0, \forall f \in \mathcal{B}, d_2(f, \mathcal{F}) \geq \rho \implies \mathbb{P}_f(\Phi_\alpha \text{ rejects}) \geq 1 - \gamma \}.$$

In the following, we derive an explicit upper bound on the uniform separation rates of the test proposed above over several classes of alternatives. For $s > 0, P > 0, M > 0$ and $l \in \{1, 2, 3\}$, we introduce

$$\mathcal{B}_s^{(l)}(P, M) = \left\{ f \in L^2(\lambda_{Leb}) \mid \forall D \in \mathcal{D}_l, \quad \|f - \Pi_{S_{(l,D)}}(f)\|_2^2 \leq P^2 D^{-2s}, \|f\|_\infty \leq M \right\}.$$

These sets of functions include some Hölder balls or Besov bodies with smoothness s , as highlighted in Fromont and Laurent (2006, Section 2.3). Corollary 2 gives an upper bound for the uniform separation rate of our testing procedure over the classes $\mathcal{B}_s^{(l)}(P, M)$ and is proved in Section C.3.

Corollary 2 *Let T_α be the test statistic defined by (13). Assume that for $l \in \{1, 2, 3\}$, \mathcal{D}_l is $\{2^J, 0 \leq J \leq \log_2(n/(\log(n) \log \log n)^2)\}$ or \emptyset . For all $s > 0, M > 0, P > 0$ and $l \in \{1, 2, 3\}$ such that $\mathcal{D}_l \neq \emptyset$, there exists some positive constant $C = C(s, \alpha, \gamma, M, \|f_0\|_\infty)$ such that the uniform separation rate of the test $\mathbf{1}_{T_\alpha > 0}$ over $\mathcal{B}_s^{(l)}(P, M)$ satisfies for n large enough*

$$\rho\left(\mathbf{1}_{T_\alpha > 0}, \mathcal{B}_s^{(l)}(P, M), \gamma\right) \leq C' P^{\frac{1}{2s+1}} \left(\frac{\log(n) \log \log n}{n}\right)^{\frac{s}{2s+1}}.$$

Remark In Corollary 2, the condition *n large enough* corresponds to

$$\left(\log(n) \frac{\log \log n}{n}\right)^{1/2} \leq P \leq \frac{n^s}{(\log(n) \log \log n)^{2s+1/2}}.$$

For the problem of testing the null hypothesis " $f = \mathbb{1}_{[0,1]}$ " against the alternative $f = \mathbb{1}_{[0,1]} + g$ with $g \neq 0$ and $g \in B_s(P)$ where $B_s(P)$ is a class of smooth functions (like some Hölder, Sobolev or Besov ball in $L^2([0, 1])$) with unknown smoothness parameter s , Ingster in Ingster (1993) established in the case where the random variables $(X_i)_{i \geq 1}$ are i.i.d. that the adaptive minimax rate of testing is of order $(\sqrt{\log \log n}/n)^{2s/(4s+1)}$. From Corollary 2, we see that our procedure leads to a rate which is close (at least for sufficiently large smoothness parameter s) to the one derived by Ingster in the i.i.d. framework since the upper bound on the uniform separation rate from Corollary 2 can be read (up to a log factor) as $([\log \log n]/n)^{\frac{2s}{4s+2}}$.

5.3 Simulations

We propose to test our method on three practical examples.¹ In all our simulations, we use Markov chains of length $n = 100$. We choose different alternatives to test our method and we use i.i.d. samples from these distributions. We chose a level $\alpha = 5\%$ for all our experiments. All tests are conducted as follows.

1. We start by the estimation of the $(1 - u)$ quantiles $t_m(u)$ of the variables $\widehat{T}_m = \widehat{\theta}_m + \|f_0\|_2^2 - \frac{2}{n} \sum_{i=1}^n f_0(X_i)$ under the hypothesis " $f = f_0$ " for u varying on a regular grid of $]0, \alpha[$. We sample 5,000 sequences of length $n = 100$ with i.i.d. random variables with distribution f_0 . We end up with an estimation $\widehat{t}_m(u)$ of $t_m(u)$ for any u in the grid and any $m \in \mathcal{M}$.

1. The code is available at <https://github.com/quentin-duchemin/goodness-of-fit-MC>.

2. Then, we estimate the value of u_α . We sample again 5,000 sequences of length $n = 100$ with i.i.d. random variables with distribution f_0 . We use them to estimate the probabilities $\mathbb{P}_{f_0}(\sup_{m \in \mathcal{M}}(\widehat{T}_m - \widehat{t}_m(u)) > 0)$ for any u in the grid and we keep the larger value of u such that the corresponding probability is still larger than α . The selected value of the grid is called u_α . Thanks to the first step, we have the estimates $\widehat{t}_m(u_\alpha)$ of $t_m(u_\alpha)$ for any $m \in \mathcal{M}$.
3. Finally, we sample 5,000 Markov chains with length $n = 100$ with stationary distribution f . For each sequence, we can compute \widehat{T}_m . Dividing by 5,000 the number of sequences for which $\sup_{m \in \mathcal{M}}(\widehat{T}_m - \widehat{t}_m(u_\alpha)) > 0$, we get an estimation of the power of the test.

To define comparison points, we compare the power of our test with the classical Kolmogorov-Smirnov test (KS test) and the Chi-squared test (χ^2 test). The rejection region associated with a test of level 5% is set by *sampling under the null* for both the KS test and the χ^2 test. With Figure 2, we provide a visualization of the density of the stationary distribution of the Markov chain and of the density of the alternative that gives the smaller power on our experiments.

5.3.1 EXAMPLE 1: AR(1) PROCESS

Let us consider some $\theta \in (0, 1)$. Then, we define the AR(1) process $(X_i)_{i \geq 1}$ starting from $X_1 = 0$ with for any $n \geq 1$,

$$X_{n+1} = \theta X_n + \xi_{n+1},$$

where $(\xi_n)_n$ are i.i.d. random variables with distribution $\mathcal{N}(0, \tau^2)$ with $\tau > 0$. From Example 1 of Duchemin et al. (2022, Section 2.6), we know that Assumptions 1 and 2 hold. The stationary measure π of the Markov chain $(X_i)_{i \geq 1}$ is $\mathcal{N}\left(0, \frac{\tau^2}{1-\theta^2}\right)$, i.e. π has density f with respect to the Lebesgue measure on \mathbb{R} with

$$\forall y \in \mathbb{R}, \quad f(y) = \frac{\sqrt{1-\theta^2}}{\sqrt{2\pi\tau^2}} \exp\left(-\frac{(1-\theta^2)y^2}{2\tau^2}\right).$$

We focus on the following alternatives

$$f_{\mu, \sigma^2}(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

Table 2 shows the estimated powers for our test, the KS test and the χ^2 test.

5.3.2 EXAMPLE 2: MARKOV CHAIN GENERATED FROM INDEPENDENT METROPOLIS HASTING ALGORITHM

Let us consider the probability measure π with density f with respect to the Lebesgue measure on $[-3, 3]$ where

$$\forall x \in [-3, 3], \quad f(x) = \frac{1}{Z} e^{-x^2} (3 + \sin(5x) + \sin(2x)),$$

(μ, σ^2)	Our test	χ^2 test	KS test	$\ \mathbf{f} - \mathbf{f}_{\mu, \sigma^2}\ _2$
(2, 1.5)	0.99	0.85	0.98	0.39
(0, 1)	0.97	0.9	0.8	0.2
(-0.2, 1.2)	0.86	0.63	0.84	0.17
(0, 1.2)	0.81	0.64	0.82	0.16
(0, 2)	0.1	0.03	0.29	0.06

Table 2: Estimated powers of the tests for Markov chains with size $n = 100$. We worked with $\tau = 1$, $\theta = 0.8$ and $\mathcal{M} = \{(1, i) : i \in \{1, \dots, 10\}\}$. Hence, the stationary distribution of the chain is approximately $\mathcal{N}(0, 2.8)$. For the χ^2 test, we work on the interval $[-5, 5]$ that we split into 20 regular parts.

with Z a normalization constant such that $\int_{-3}^3 f(x)dx = 1$. To construct a Markov chain with stationary measure π , we use an independent Metropolis-Hasting algorithm with proposal density $q(x) \propto \exp(-x^2/6)$. Using Proposition 2, we get that the above built Markov chain $(X_i)_{i \geq 1}$ satisfies Assumptions 1 and 2. We focus on the following alternatives

$$g_{\mu, \sigma^2}(x) = \frac{1}{Z(\mu, \sigma^2)} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right) \mathbf{1}_{[-3, 3]}(x),$$

where $Z(\mu, \sigma^2)$ is a normalization constant such that $\int g_{\mu, \sigma^2}(x)dx = 1$. Table 3 shows the estimated powers for our test, the KS test and the χ^2 test.

(μ, σ^2)	Our test	χ^2 test	KS test	$\ \mathbf{f} - \mathbf{g}_{\mu, \sigma^2}\ _2$
(0, 1)	0.96	0.91	0.9	0.29
(0, 0.7 ²)	0.95	0.84	0.93	0.23
(0.3, 0.7 ²)	0.92	0.87	0.93	0.19

Table 3: Estimated powers of the tests for Markov chains with size $n = 100$. We used $\mathcal{M} = \{(1, i) : i \in \{1, \dots, 10\}\}$. For the χ^2 test, we work on the interval $[-3, 3]$ that we split into 20 regular parts.

5.3.3 EXAMPLE 3: ARCH PROCESS

Let us consider some $\theta \in (-1, 1)$. We are interested in the simple threshold auto-regressive model $(X_n)_{n \geq 1}$ defined by $X_1 = 0$ and for any $n \geq 1$,

$$X_{n+1} = \theta|X_n| + (1 - \theta^2)^{1/2}\xi_{n+1},$$

where the random variables $(\xi_n)_{n \geq 2}$ are i.i.d. with standard Gaussian distribution. From Example 3 of Duchemin et al. (2022, Section 2.6), we know that Assumptions 1 and 2 hold. The transition kernel of the Markov chain $(X_i)_{i \geq 1}$ is

$$\forall x, y \in \mathbb{R}, \quad P(x, y) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(y - \theta|x|)^2}{2(1 - \theta^2)}\right).$$

The stationary distribution π of the Markov chain has density f with respect to the Lebesgue measure on \mathbb{R} with

$$\forall y \in \mathbb{R}, \quad f(y) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{y^2}{2}\right) \Phi\left(\frac{\theta y}{(1-\theta^2)^{1/2}}\right),$$

where Φ is the standard normal cumulative distribution function. We focus on the following alternatives

$$f_{\mu,\sigma^2}(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

Table 4 shows the estimated powers for our test, the KS test and the χ^2 test.

(μ, σ^2)	Our test	χ^2 test	KS test	$\ f - f_{\mu,\sigma^2}\ _2$
(0, 1)	0.98	0.85	0.95	0.3
(1, 0.8 ²)	0.95	0.79	0.88	0.22
(0.5, 1)	0.3	0.07	0.5	0.14
(0.6, 0.8 ²)	0.35	0.16	0.4	0.036

Table 4: Estimated powers of the tests for Markov chains with size $n = 100$. We used $\theta = 0.8$ and $\mathcal{M} = \{(1, i) : i \in \{1, \dots, 10\}\}$. For the χ^2 test, we work on the interval $[-20, 20]$ that we split into 20 regular parts.




Figure 2: In solid line, we plot the density of the stationary measure of the Markov chain for the three examples of our simulations. In dotted line, we plot the density of the alternative that gives the smaller power on our experiments.

5.3.4 COMMENTS ON OUR NUMERICAL EXPERIMENTS

Our experiments show that the χ^2 goodness-of-fit test give in general the smaller power compared to our method and to the KS test. The χ^2 test is better suited to deal with discrete probability distributions and it seems to suffer to small power in our continuous setting. Note that using the χ^2 test with continuous densities on \mathbb{R} require to specify some hyperparameters (such as a compact interval and the number of bins to discretize it). In practice, the test results can change drastically by modifying these hyperparameters, making the test unreliable. Our experiments also show that when the L^2 norm between the true density f and the alternative one f_0 is large enough, our method reaches higher power

compared to the two other procedures considered. Nevertheless, our approach seems less powerful compared to the KS test when the L^2 norm $\|f - f_0\|_2$ is getting smaller. This is not surprising since our testing procedure is based on the L^2 norm while the KS test relies on the sup norm between cumulative distribution functions (CDFs). We conduct a final experiment to better stress this distinction between our procedure and the KS test. We consider the notations and the framework of the example from Section 5.3.1 with the following alternatives

$$f^{(L,\delta)}(x) = \begin{cases} f_{0, \frac{\tau^2}{1-\theta^2}}(x) & \text{if } |x| \geq \delta \\ f_{0, \frac{\tau^2}{1-\theta^2}}(x) - L & \text{if } -\delta < x \leq 0 \\ f_{0, \frac{\tau^2}{1-\theta^2}}(x) + L & \text{if } 0 < x < \delta \end{cases},$$

where $L, \delta > 0$ are chosen so that $f^{(L,\delta)}(x) \geq 0$ for any $x \in \mathbb{R}$. We work with $\tau = 1, \theta = 0.8$ and $\mathcal{M} = \{(1, i) : i \in \{1, \dots, 10\}\}$. Figure 3 shows the alternatives considered. The sup norm between the CDFs of f and $f^{(L,\delta)}$ is equal to $L\delta$ while the squared L^2 norm between f and $f^{(L,\delta)}$ is $2L^2\delta^3/3$. Hence, we expect that powers will increase for both tests when L and/or δ are increasing. Moreover, we expect the power of our method to be more sensitive to the parameters L and δ . Those intuitions are confirmed with the numerical experiments presented in Table 5.




Figure 3: Alternative considered.

(\mathbf{L}, δ)	0.25	0.5	0.75	1
0.05	0.06	0.12	0.2	0.21
	0.1	0.15	0.2	0.22
0.05	0.16	0.33	0.36	0.4
	0.23	0.26	0.33	0.37
0.1	0.33	0.66	0.8	0.83
	0.26	0.35	0.46	0.48
0.15	0.82	0.87	0.9	0.95
	0.35	0.45	0.55	0.67
0.2	0.9	0.93	0.95	0.98
	0.46	0.54	0.72	0.87

Table 5: Estimated powers of the tests for Markov chains with size $n = 100$. Gray cells are the powers of our method while blank cells are the ones obtained with the KS test.

Acknowledgments

We would like to acknowledge support for this project from the Région Ile-de-France.

Appendices

Guidelines for the supplementary material.

- Sections A, B and C: Proofs
Sections A, B and C provide respectively the proofs of our main results from Sections 3, 4 and 5.
- Section D: Technical Lemmas
This section contains some Lemmas useful for our proofs.

A. Proofs for Section 3

Let us explain in a nutshell the structure of our proof. For any natural number R , we denote \mathbf{H}^R the integral operator with kernel function h_R at resolution R , namely

$$h_R(x, y) := \sum_{r \in I, r \leq R} \lambda_r \varphi_r(x) \varphi_r(y), \quad \mathbf{H}^R f(x) := \int_E h_R(x, y) f(y) d\pi(y).$$

We define $\tilde{\mathbf{H}}_n^R$ and \mathbf{H}_n^R analogously by using the kernel h_R in Eq.(3). Using the triangle inequality, we split the distance $\delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))$ into four terms.

1. $\delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}^R))$ is a bias term induced by working at resolution R .
2. A non-trivial preliminary work allows to prove that $\delta_2(\lambda(\mathbf{H}^R), \lambda(\tilde{\mathbf{H}}_n^R))$ can be written as an empirical process of the Markov chain $(X_i)_{i \geq 1}$ whose tail can be controlled by applying concentration inequalities for sums of functions of uniformly ergodic Markov chains (this is where we use the assumption that Υ is finite). We refer to Eq.(15).
3. Since the matrices \mathbf{H}_n^R and $\tilde{\mathbf{H}}_n^R$ only differ at diagonal elements, $\delta_2(\lambda(\tilde{\mathbf{H}}_n^R), \lambda(\mathbf{H}_n^R))$ can be coarsely bounded by $n^{-1/2} \|h_R\|_\infty$ (cf. Eq.(16)).
4. Applying the Hoffman-Wielandt inequality, one can notice that $\delta_2(\lambda(\mathbf{H}_n^R), \lambda(\mathbf{H}_n))$ can be upper-bounded by a U-statistic of order two of the Markov chain $(X_i)_{i \geq 1}$ (cf. Eq.(17)). We control the tail behaviour of this U-statistic by applying Theorem 1.

The proof is then concluded by choosing the resolution level R so that $R^2 = \lceil \sqrt{n} \rceil$.

A.1 Deviation inequality for the spectrum of signed integral operators

As shown in Section A.2, Theorem 2 is a direct consequence of the concentration result provided by Theorem 7.

Theorem 7 *We keep notations of Section 3. Assume that $(X_n)_{n \geq 1}$ is a Markov chain on E satisfying Assumptions 1 and 2 described in Section 2.2 with stationary distribution π . Let us consider some symmetric kernel $h : E \times E \rightarrow \mathbb{R}$, square integrable with respect to $\pi \otimes \pi$. Let us consider some $R \in \mathbb{N}^*$. We assume that there exist continuous functions $\varphi_r : E \rightarrow \mathbb{R}$, $r \in I$*

(where $I = \mathbb{N}$ or $I = 1, \dots, N$) that form an orthonormal basis of $L^2(\pi)$ such that it holds pointwise

$$h(x, y) = \sum_{r \in I} \lambda_r \varphi_r(x) \varphi_r(y),$$

with

$$\Lambda_R := \sup_{r \in I, r \leq R} |\lambda_r| \text{ and } \Upsilon_R := \sup_{r \in I, r \leq R} \|\varphi_r\|_\infty^2.$$

We define $h_R(x, y) = \sum_{r \in I, r \leq R} \lambda_r \varphi_r(x) \varphi_r(y)$ and we assume that $\|h_R\|_\infty, \|h - h_R\|_\infty < \infty$. Then there exists a universal constant $K > 0$ such that for any $t > 0$, it holds

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{4} \delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))^2 \geq (\|h_R\|_\infty^2 + \kappa \|h - h_R\|_\infty^2) \frac{\log n}{n} + 2 \sum_{i > R, i \in I} \lambda_i^2 + t \right) \\ & \leq 16 \exp \left(-n \frac{t^2}{K m^2 \tau^2 \|h - h_R\|_\infty^2} \right) + \beta \log(n) \exp \left(-\frac{n}{16 \log n} \left\{ \left[\frac{t}{c} \right] \wedge \left[\frac{t}{c} \right]^{1/2} \right\} \right) \\ & + 16 R^2 \exp \left(-\frac{nt}{K m^2 \tau^2 R^2 \Lambda_R^2 \Upsilon_R^2} \right). \end{aligned}$$

where $c = \kappa \|h - h_R\|_\infty$ with $\kappa > 0$ depending on δ_M, τ, L, m and ρ . β depends only on ρ .

Proof of Theorem 7. For any integer $R \geq 1$, we denote

$$\begin{aligned} X_{n,R} &:= \frac{1}{\sqrt{n}} (\varphi_r(X_i))_{1 \leq i \leq n, 1 \leq r \leq R} \in \mathbb{R}^{n \times R} \\ A_{n,R} &:= \left(X_{n,R}^\top X_{n,R} \right)^{1/2} \in \mathbb{R}^{R \times R} \\ K_R &:= \text{Diag}(\lambda_1, \dots, \lambda_R) \\ \tilde{\mathbf{H}}_n^R &:= X_{n,R} K_R X_{n,R}^\top \\ \mathbf{H}_n^R &:= \left((1 - \delta_{i,j}) \left(\tilde{\mathbf{H}}_n^R \right)_{i,j} \right)_{1 \leq i, j \leq n}. \end{aligned}$$

We remark that $A_{n,R}^2 = I_R + E_{R,n}$ where $(E_{R,n})_{r,s} = (1/n) \sum_{i=1}^n (\varphi_r(X_i) \varphi_s(X_i) - \delta_{r,s})$ for all $r, s \in [R]$. Denoting $\lambda(\mathbf{H}^R) = (\lambda_1, \dots, \lambda_R)$, we have

$$\delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))^2 \leq 4 \left[\delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}^R))^2 + \delta_2(\lambda(\mathbf{H}^R), \lambda(\tilde{\mathbf{H}}_n^R))^2 + \delta_2(\lambda(\tilde{\mathbf{H}}_n^R), \lambda(\mathbf{H}_n^R))^2 + \delta_2(\lambda(\mathbf{H}_n^R), \lambda(\mathbf{H}_n))^2 \right].$$

Bounding $\delta_2(\lambda(\mathbf{H}^R), \lambda(\tilde{\mathbf{H}}_n^R))^2$.

Let us consider some $\varepsilon > 0$.

Using a singular value decomposition of $X_{n,R}$, one can show that $\lambda(X_{n,R} K_R X_{n,R}^\top) = \lambda(A_{n,R} K_R A_{n,R})$ which leads to

$$\begin{aligned} \delta_2 \left(\lambda(\mathbf{H}^R), \lambda(\tilde{\mathbf{H}}_n^R) \right) &= \delta_2 \left(\lambda(K_R), \lambda(X_{n,R} K_R X_{n,R}^\top) \right) \\ &= \delta_2 \left(\lambda(K_R), \lambda(A_{n,R} K_R A_{n,R}) \right) \\ &\leq \|K_R - A_{n,R} K_R A_{n,R}\|_F, \end{aligned}$$

Using Equation (4.8) from Koltchinskii and Giné (2000, page 127), we get

$$\delta_2 \left(\lambda(\mathbf{H}^R), \lambda(\tilde{\mathbf{H}}_n^R) \right)^2 \leq 2 \|K_R E_{R,n}\|_F^2 = 2 \sum_{1 \leq r, s \leq R} \lambda_s^2 \left(\frac{1}{n} \sum_{i=1}^n \varphi_r(X_i) \varphi_s(X_i) - \delta_{r,s} \right)^2. \quad (15)$$

Hence,

$$\begin{aligned} & \mathbb{P} \left(\delta_2 \left(\lambda(\mathbf{H}^R), \lambda(\tilde{\mathbf{H}}_n^R) \right)^2 \geq t \right) \\ & \leq \sum_{1 \leq s, r \leq R} \mathbb{P} \left(\sqrt{2} |\lambda_s| \left| \frac{1}{n} \sum_{i=1}^n \varphi_r(X_i) \varphi_s(X_i) - \delta_{r,s} \right| \geq \sqrt{t}/R \right) \\ & \leq \sum_{1 \leq s, r \leq R, \lambda_s \neq 0} \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \varphi_r(X_i) \varphi_s(X_i) - \delta_{r,s} \right| \geq \sqrt{t}/(\sqrt{2}R|\lambda_s|) \right) \\ & \leq \sum_{1 \leq s, r \leq R, \lambda_s \neq 0} 16 \exp \left(- (Km^2\tau^2)^{-1} \frac{nt}{R^2 |\lambda_s|^2 \Upsilon_R^4} \right) \\ & = 16R^2 \exp \left(- (Km^2\tau^2)^{-1} \frac{nt}{R^2 \Lambda_R^2 \Upsilon_R^4} \right), \end{aligned}$$

where the last inequality follows from Proposition 7 and where $K > 0$ is a universal constant.

Bounding $\delta_2(\lambda(\tilde{\mathbf{H}}_n^R), \lambda(\mathbf{H}_n^R))^2$.

$$\delta_2(\lambda(\tilde{\mathbf{H}}_n^R), \lambda(\mathbf{H}_n^R))^2 \leq \|\tilde{\mathbf{H}}_n^R - \mathbf{H}_n^R\|_F^2 = \frac{1}{n^2} \left(\sum_{i=1}^n h_R^2(X_i, X_i) \right) \leq \frac{\|h_R\|_\infty^2}{n}. \quad (16)$$

Bounding $\delta_2(\lambda(\mathbf{H}_n^R), \lambda(\mathbf{H}_n))^2$.

$$\delta_2(\lambda(\mathbf{H}_n^R), \lambda(\mathbf{H}_n))^2 \leq \|\tilde{\mathbf{H}}_n^R - \tilde{\mathbf{H}}_n\|_F^2 = \frac{1}{n^2} \left(\sum_{1 \leq i, j \leq n, i \neq j} (h - h_R)(X_i, X_j)^2 \right). \quad (17)$$

Let us consider,

$$\forall x, y \in E, \quad m_R(x, y) := (h - h_R)^2(x, y) - s_R(x) - s_R(y) - \mathbb{E}_{\pi \otimes \pi}[(h - h_R)^2(X, Y)],$$

where $s_R(x) = \mathbb{E}_\pi[(h - h_R)^2(x, X)] - \mathbb{E}_{\pi \otimes \pi}[(h - h_R)^2(X, Y)]$. One can check that for any $x \in E$, $\mathbb{E}_\pi[m_R(x, X)] = \mathbb{E}_\pi[m_R(X, x)] = 0$. Hence, m_R is π -canonical.

$$\frac{1}{n(n-1)} \left(\sum_{1 \leq i, j \leq n, i \neq j} (h - h_R)(X_i, X_j)^2 \right) \quad (18)$$

$$= \frac{1}{n(n-1)} \sum_{1 \leq i, j \leq n, i \neq j} m_R(X_i, X_j) + \frac{2}{n} \sum_{i=1}^n s_R(X_i) + \mathbb{E}_{\pi \otimes \pi}[(h - h_R)^2(X, Y)]. \quad (19)$$

Using Theorem 1, we get that there exist two constants $\beta, \kappa > 0$ such that for any $u \geq 1$, it holds with probability at least $1 - \beta e^{-u} \log(n)$,

$$\frac{1}{n(n-1)} \sum_{1 \leq i, j \leq n, i \neq j} m_R(X_i, X_j) \leq \kappa \|h - h_R\|_\infty \log n \left\{ \frac{u}{n} \vee \left(\frac{u}{n} \right)^2 \right\}.$$

Let us now consider some $t > 0$ such that

$$\kappa \|h - h_R\|_\infty \log n \left\{ \frac{u}{n} \vee \left(\frac{u}{n} \right)^2 \right\} \leq t. \quad (20)$$

The condition (20) is equivalent to

$$u \leq n \left\{ \frac{t}{\kappa \|h - h_R\|_\infty \log n} \wedge \left(\frac{t}{\kappa \|h - h_R\|_\infty \log n} \right)^{1/2} \right\},$$

which is satisfied in particular if t and u are such that

$$u = \frac{n}{\log n} \left\{ \left[\frac{t}{c} \right] \wedge \left[\frac{t}{c} \right]^{1/2} \right\},$$

where $c = \kappa \|h - h_R\|_\infty$. One can finally notice that for this choice of u , the condition $u \geq 1$ holds in particular for n large enough in order to have $n/\log n \geq \kappa \|h - h_R\|_\infty t^{-1}$.

We deduce from this analysis that for any $t > 0$, we have for n large enough to satisfy $n/\log n \geq \kappa \|h - h_R\|_\infty t^{-1}$,

$$\mathbb{P} \left(\frac{1}{n(n-1)} \sum_{1 \leq i, j \leq n, i \neq j} m_R(X_i, X_j) \geq t \right) \leq \beta \log(n) \exp \left(-\frac{n}{\log n} \left\{ \left[\frac{t}{c} \right] \wedge \left[\frac{t}{c} \right]^{1/2} \right\} \right).$$

Using Proposition 7, we get that for some universal constant $K > 0$,

$$\mathbb{P} \left(\frac{2}{n} \left| \sum_{i=1}^n s_R(X_i) \right| \geq t \right) \leq 16 \exp \left(-n \frac{t^2}{K m^2 \tau^2 \|h - h_R\|_\infty^2} \right).$$

We deduce that for some universal constant $K > 0$ it holds

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{n^2} \left(\sum_{1 \leq i, j \leq n, i \neq j} (h - h_R)(X_i, X_j)^2 \right) - \mathbb{E}_{\pi \otimes \pi} [(h - h_R)^2] \geq t \right) \\ & \leq 16 \exp \left(-n \frac{t^2}{K m^2 \tau^2 \|h - h_R\|_\infty^2} \right) + \beta \log(n) \exp \left(-\frac{n}{4 \log n} \left\{ \left[\frac{t}{c} \right] \wedge \left[\frac{t}{c} \right]^{1/2} \right\} \right). \end{aligned}$$

Since $\mathbb{E}_{\pi \otimes \pi} [(h - h_R)^2(X, Y)] = \sum_{i > R, i \in I} \lambda_i^2$, we deduce that

$$\begin{aligned} & \mathbb{P} \left(\delta_2(\lambda(\mathbf{H}_n^R), \lambda(\mathbf{H}_n))^2 - \sum_{i > R, i \in I} \lambda_i^2 \geq t \right) \\ & \leq 16 \exp \left(-n \frac{t^2}{K m^2 \tau^2 \|h - h_R\|_\infty^2} \right) + \beta \log(n) \exp \left(-\frac{n}{4 \log n} \left\{ \left[\frac{t}{c} \right] \wedge \left[\frac{t}{c} \right]^{1/2} \right\} \right). \end{aligned}$$

Hence we proved that for any $u > 0$ such that $n/\log n \geq \kappa \|h - h_R\|_\infty u^{-1}$,

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{4} \delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))^2 \geq \frac{\|h_R\|_\infty^2}{n} + 2 \sum_{i>R, i \in I} \lambda_i^2 + u \right) \\ & \leq 16 \exp \left(-n \frac{u^2}{K m^2 \tau^2 \|h - h_R\|_\infty^2} \right) + \beta \log(n) \exp \left(-\frac{n}{16 \log n} \left\{ \left[\frac{u}{c} \right] \wedge \left[\frac{u}{c} \right]^{1/2} \right\} \right) \\ & + 16R^2 \exp \left(-\frac{nu}{K m^2 \tau^2 R^2 \Lambda_R^2 \Upsilon_R^2} \right). \end{aligned}$$

Considering $t > 0$ and applying the previous inequality with $u = t + \frac{\kappa \|h - h_R\|_\infty \log n}{n}$, we get

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{4} \delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))^2 \geq (\|h_R\|_\infty^2 + \kappa \|h - h_R\|_\infty^2) \frac{\log n}{n} + 2 \sum_{i>R, i \in I} \lambda_i^2 + t \right) \\ & \leq 16 \exp \left(-n \frac{t^2}{K m^2 \tau^2 \|h - h_R\|_\infty^2} \right) + \beta \log(n) \exp \left(-\frac{n}{16 \log n} \left\{ \left[\frac{t}{c} \right] \wedge \left[\frac{t}{c} \right]^{1/2} \right\} \right) \\ & + 16R^2 \exp \left(-\frac{nt}{K m^2 \tau^2 R^2 \Lambda_R^2 \Upsilon_R^2} \right). \end{aligned}$$

This concludes the proof of Theorem 7.

A.2 Proof of Theorem 2.

We consider any $R \in \mathbb{N}^*$. We remark that for any $x, y \in E$,

$$\begin{aligned} |h_R(x, y)| &= \left| \sum_{r=1}^R \lambda_r \varphi_r(x) \varphi_r(y) \right| \\ &\leq \left(\sum_{r=1}^R |\lambda_r| \varphi_r(x)^2 \right)^{1/2} \times \left(\sum_{r=1}^R |\lambda_r| \varphi_r(y)^2 \right)^{1/2} \quad (\text{Using Cauchy-Schwarz inequality}) \\ &\leq S, \end{aligned}$$

which proves that $\|h_R\|_\infty \leq S$. Similar computations lead to $\|h - h_R\|_\infty \leq S$.

Using Theorem 7 we get for any $t > 0$,

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{4} \delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))^2 \geq \frac{S^2(1 + \kappa) \log n}{n} + 2 \sum_{i>R, i \in I} \lambda_i^2 + t \right) \\ & \leq 16 \exp \left(-n \frac{t^2}{K m^2 \tau^2 S^2} \right) + \beta \log(n) \exp \left(-\frac{n}{16 \log n} \left\{ \left[\frac{t}{\kappa S} \right] \wedge \left[\frac{t}{\kappa S} \right]^{1/2} \right\} \right) \\ & + 16R^2 \exp \left(-\frac{nt}{K m^2 \tau^2 R^2 \Lambda^2 \Upsilon^2} \right), \end{aligned}$$

where $\Lambda := \sup_{r \geq 1} |\lambda_r| < \infty$. Choosing $R^2 = \lceil \sqrt{n} \rceil$, we get

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{4} \delta_2(\lambda(\mathbf{H}), \lambda(\mathbf{H}_n))^2 \geq \frac{S^2(1+\kappa) \log n}{n} + 2 \sum_{i > \lceil n^{1/4} \rceil, i \in I} \lambda_i^2 + t \right) \\ & \leq 32\sqrt{n} \exp(-\mathcal{C} \min(nt^2, \sqrt{nt})) + \beta \log(n) \exp\left(-\frac{n}{\log n} \min(\mathcal{B}t, (\mathcal{B}t)^{1/2})\right), \end{aligned}$$

where $\mathcal{B} = (K\kappa S)^{-1}$ and $\mathcal{C} = K^{-1} (m^2 \tau^2 (S + \Lambda \Upsilon))^{-2}$.

B. Proofs for Section 4

In this section, for any $k \geq 0$ we denote \mathbb{E}_k the conditional expectation with respect to the σ -algebra $\sigma(X_1, \dots, X_k)$.

B.1 Proof of Theorem 3

By definition of \mathcal{M}^n , we want to bound

$$\mathbb{P} \left(\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} M_t \geq \varepsilon \right),$$

which takes the form

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} [\mathcal{R}(h_{t-b_n}) - \mathbb{E}_{t-b_n}[M_t]] - \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} [M_t - \mathbb{E}_{t-b_n}[M_t]] \geq \varepsilon \right) \\ & \leq \mathbb{P} \left(\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} [\mathcal{R}(h_{t-b_n}) - \mathbb{E}_{t-b_n}[M_t]] \geq \varepsilon/2 \right) + \mathbb{P} \left(\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} [\mathbb{E}_{t-b_n}[M_t] - M_t] \geq \varepsilon/2 \right). \end{aligned} \tag{21}$$

B.1.1 STEP 1: MARTINGALE DIFFERENCE

We first deal with the second term of Eq.(21). Note that we can write

$$\sum_{t=c_n}^{n-1} [\mathbb{E}_{t-b_n}[M_t] - M_t] = \sum_{t=c_n}^{n-1} \sum_{k=1}^{b_n} [\mathbb{E}_{t-k}[M_t] - \mathbb{E}_{t-k+1}[M_t]] = \sum_{k=1}^{b_n} \sum_{t=c_n}^{n-1} [\mathbb{E}_{t-k}[M_t] - \mathbb{E}_{t-k+1}[M_t]].$$

Let us consider some $k \in \{1, \dots, b_n\}$, then we have that $V_t^{(k)} = (\mathbb{E}_{t-k}[M_t] - \mathbb{E}_{t-k+1}[M_t]) / (n - c_n)$ is a martingale difference sequence, i.e. $\mathbb{E}_{t-k}[V_t^{(k)}] = 0$. Since the loss function is bounded in $[0, 1]$, we have $|V_t^{(k)}| \leq 2/(n - c_n)$, $t = 1, \dots, n$. Therefore by the Hoeffding-Azuma inequality, $\sum_t V_t^{(k)}$ can be bounded such that

$$\mathbb{P} \left(\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} [\mathbb{E}_{t-k}[M_t] - \mathbb{E}_{t-k+1}[M_t]] \geq \frac{\varepsilon}{2b_n} \right) \leq \exp \left(-\frac{(1-c)n\varepsilon^2}{8b_n^2} \right).$$

We deduce that

$$\mathbb{P}\left(\frac{1}{n-c_n}\sum_{t=c_n}^{n-1}[\mathbb{E}_{t-b_n}[M_t] - M_t] \geq \varepsilon/2\right) \leq b_n \exp\left(-\frac{(1-c)n\varepsilon^2}{8b_n^2}\right). \quad (22)$$

B.1.2 STEP 2: SYMMETRIZATION BY A GHOST SAMPLE

In this step we bound the first term in Eq.(21). Let us start by introducing a ghost sample $\{\xi_j\}_{1 \leq j \leq n}$, where the random variables ξ_j i.i.d with distribution π . Recall the definition of M_t and define \widetilde{M}_t as

$$M_t = \frac{1}{t-b_n} \sum_{i=1}^{t-b_n} \ell(h_{t-b_n}, X_t, X_i), \quad \widetilde{M}_t = \frac{1}{t-b_n} \sum_{i=1}^{t-b_n} \ell(h_{t-b_n}, X_t, \xi_i).$$

The difference between \widetilde{M}_t and M_t is that M_t is the sum of the loss incurred by h_{t-b_n} on the current instance X_t and all the previous examples X_j , $j = 1, \dots, t-b_n$ on which h_{t-b_n} is trained, while \widetilde{M}_t is the loss incurred by the same hypothesis h_{t-b_n} on the current instance X_t and an independent set of examples ξ_j , $j = 1, \dots, t-b_n$.

First remark that we have

$$\begin{aligned} & \frac{1}{n-c_n} \sum_{t=c_n}^{n-1} [\mathcal{R}(h_{t-b_n}) - \mathbb{E}_{t-b_n}[M_t]] \\ = & \frac{1}{n-c_n} \sum_{t=c_n}^{n-1} [\mathcal{R}(h_{t-b_n}) - \mathbb{E}_{t-b_n}[\widetilde{M}_t]] + \frac{1}{n-c_n} \sum_{t=c_n}^{n-1} [\mathbb{E}_{t-b_n}[\widetilde{M}_t] - \mathbb{E}_{t-b_n}[M_t]]. \end{aligned} \quad (23)$$

The first term of Eq.(23) is handled in Wang et al. (2012) by relying heavily on the assumption that samples are i.i.d (see Wang et al., 2012, Claim 1). Hence, the approach of Wang and al. cannot be adapted in our framework. To overcome this difficulty, we use the uniform ergodicity of the Markov chain. This is where the use of the burning parameter b_n is essential.

Since ℓ is in $[0, 1]$, the first term can be bounded directly using the uniform ergodicity of the Markov chain $(X_i)_i$ as follows

$$\begin{aligned} & \frac{1}{n-c_n} \sum_{t=c_n}^{n-1} [\mathcal{R}(h_{t-b_n}) - \mathbb{E}_{t-b_n}[\widetilde{M}_t]] \\ = & \frac{1}{n-c_n} \sum_{t=c_n}^{n-1} \int_{x \in E} \left(d\pi(x) \mathbb{E}_{X \sim \pi}[\ell(h_{t-b_n}, x, X)] - P^{b_n}(X_{t-b_n}, dx) \mathbb{E}_{X \sim \pi}[\ell(h_{t-b_n}, x, X)] \right) \\ = & \frac{1}{n-c_n} \sum_{t=c_n}^{n-1} \int_{x \in E} \mathbb{E}_{X \sim \pi}[\ell(h_{t-b_n}, x, X)] \left(d\pi(x) - P^{b_n}(X_{t-b_n}, dx) \right) \\ \leq & \frac{1}{n-c_n} \sum_{t=c_n}^{n-1} \int_{x \in E} \left| d\pi(x) - P^{b_n}(X_{t-b_n}, dx) \right| \\ \leq & L\rho^{b_n}, \end{aligned}$$

where we used Eq.(5).

It remains to control

$$\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \left[\mathbb{E}_{t-b_n}[\widetilde{M}_t] - \mathbb{E}_{t-b_n}[M_t] \right],$$

and we follow an approach similar to Wang et al. (2012). Let us remind that M_t and \widetilde{M}_t depend on the hypothesis h_{t-b_n} and let us define $L_t(h_{t-b_n}) = \left[\mathbb{E}_{t-b_n}[\widetilde{M}_t] - \mathbb{E}_{t-b_n}[M_t] \right]$. We have

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} L_t(h_{t-b_n}) \geq \varepsilon \right) \\ & \leq \mathbb{P} \left(\sup_{\widehat{h}_{c_n-b_n}, \dots, \widehat{h}_{n-1-b_n}} \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} L_t(\widehat{h}_{t-b_n}) \geq \varepsilon \right) \\ & \leq \sum_{t=c_n}^{n-1} \mathbb{P} \left(\sup_{\widehat{h} \in \mathcal{H}} L_t(\widehat{h}) \geq \varepsilon \right). \end{aligned} \quad (24)$$

To bound the right hand side of Eq.(24) we give first the following Lemma.

Lemma 1 *Given any function $f \in \mathcal{H}$ and any $t \geq c_n$,*

$$\forall \varepsilon > 0, \quad \mathbb{P}(L_t(f) \geq \varepsilon) \leq 16 \exp(-(t - b_n)C(m, \tau)\varepsilon^2).$$

In the i.i.d. framework, the counterpart of Lemma 1 follows from a straightforward application of McDiarmid's inequality (see Wang et al., 2012, Lemma 5). In our work, we consider uniformly ergodic Markov chains and the proof of Lemma 1 requires extra work. We apply a concentration inequality for Markov chains (see Proposition 7) which needs to hold for any initial distribution. We apply Proposition 7 by considering the time-reversed sequence and this is where we use the reversibility of the chain.



Proof of Lemma 1.

Note that

$$\begin{aligned} L_t(f) &= \mathbb{E}_{t-b_n}[\widetilde{M}_t] - \mathbb{E}_{t-b_n}[M_t] \\ &= \frac{1}{t - b_n} \sum_{i=1}^{t-b_n} (\mathbb{E}_{t-b_n}[\ell(f, X_t, \xi_i)] - \mathbb{E}_{t-b_n}[\ell(f, X_t, X_i)]) \\ &= \frac{1}{t - b_n} \sum_{i=1}^{t-b_n} \mathbb{E}_{\xi \sim \pi} \left[\mathbb{E}_{X_t \sim P^{b_n}(X_{t-b_n}, \cdot)} \{\ell(f, X_t, \xi)\} \right] - \mathbb{E}_{X_t \sim P^{b_n}(X_{t-b_n}, \cdot)} \{\ell(f, X_t, X_i)\}. \end{aligned}$$

Hence, denoting $m(f, X_{t-b_n}, x) = \mathbb{E}_{X_t \sim P^{b_n}(X_{t-b_n}, \cdot)} \{\ell(f, X_t, x)\}$, we get

$$L_t(f) \leq \frac{1}{t - b_n} \sum_{i=1}^{t-b_n} \{ \mathbb{E}_{\xi \sim \pi} [m(f, X_{t-b_n}, \xi)] - m(f, X_{t-b_n}, X_i) \}.$$

By the reversibility of the chain $(X_i)_{i \geq 1}$, we know that the sequence $(X_{t-b_n}, X_{t-b_n-1}, \dots, X_1)$ conditionally on X_{t-b_n} is a Markov chain with stationary distribution π . Applying Proposition 7 (see Section D) we get that

$$\begin{aligned} & \mathbb{P}(L_t(f) \geq \varepsilon \mid X_{t-b_n}) \\ & \leq \mathbb{P}\left(\frac{1}{t-b_n} \sum_{i=1}^{t-b_n} \{\mathbb{E}_{\xi_i \sim \pi} [m(f, X_{t-b_n}, \xi_i)] - m(f, X_{t-b_n}, X_i)\} \geq \varepsilon \mid X_{t-b_n}\right) \\ & \leq 16 \exp(-(t-b_n)C(m, \tau)\varepsilon^2), \end{aligned}$$

for some constant $C(m, \tau) > 0$ depending only on m and τ . Then we deduce that

$$\begin{aligned} \mathbb{P}(L_t(f) \geq \varepsilon) &= \mathbb{E}[\mathbb{E}\{\mathbf{1}_{L_t(f) \geq \varepsilon} \mid X_{t-b_n}\}] \\ &= \mathbb{E}[\mathbb{P}\{L_t(f) \geq \varepsilon \mid X_{t-b_n}\}] \\ &\leq 16 \exp(-(t-b_n)C(m, \tau)\varepsilon^2), \end{aligned}$$

which concludes the proof of Lemma 1. ■

The following two Lemmas are key elements to prove Lemma 4. Their proofs are strictly analogous to the proofs of Lemmas 6, 7 and 8 from Wang et al. (2012).

Lemma 2 (cf. Wang et al., 2012, Lemma 6) For any two functions $h_1, h_2 \in \mathcal{H}$, the following equation holds

$$|L_t(h_1) - L_t(h_2)| \leq 2\text{Lip}(\varphi)\|h_1 - h_2\|_\infty.$$

Lemma 3 Let $\mathcal{H} = S_1 \cup \dots \cup S_l$ and $\varepsilon > 0$. Then

$$\mathbb{P}\left(\sup_{h \in \mathcal{H}} L_t(h) \geq \varepsilon\right) \leq \sum_{j=1}^l \mathbb{P}\left(\sup_{h \in S_j} L_t(h) \geq \varepsilon\right).$$

Lemma 4 (cf. Wang et al., 2012, Lemma 6) For any $c_n \leq t \leq n$, it holds

$$\mathbb{P}\left(\sup_{h \in \mathcal{H}} L_t(h) \geq \varepsilon\right) \leq 16\mathcal{N}\left(\mathcal{H}, \frac{\varepsilon}{4\text{Lip}(\varphi)}\right) \exp\left(-\frac{(t-b_n)C(m, \tau)\varepsilon^2}{4}\right).$$

Combining Lemma 4 and Eq.(24), we have

$$\mathbb{P}\left(\frac{1}{n-c_n} \sum_{t=c_n}^{n-1} L_t(h_{t-b_n}) \geq \varepsilon\right) \leq 16\mathcal{N}\left(\mathcal{H}, \frac{\varepsilon}{4\text{Lip}(\varphi)}\right) n \exp\left(-\frac{(c_n-b_n)C(m, \tau)\varepsilon^2}{4}\right).$$

We deduce that

$$\begin{aligned} & \mathbb{P}\left(\frac{1}{n-c_n} \sum_{t=c_n}^{n-1} [\mathcal{R}(h_{t-b_n}) - \mathbb{E}_{t-b_n}[M_t]] \geq \varepsilon/2\right) \\ & \leq \mathbb{P}\left(L\rho^{b_n} + \frac{1}{n-c_n} \sum_{t=c_n}^{n-1} [\mathbb{E}_{t-b_n}[\widetilde{M}_t] - \mathbb{E}_{t-b_n}[M_t]] \geq \varepsilon/2\right) \\ & \leq 16\mathcal{N}\left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)}\right) n \exp\left(-\frac{(c_n-b_n)C(m, \tau)(\varepsilon/2 - L\rho^{b_n})^2}{4}\right). \end{aligned}$$

B.1.3 STEP 3: CONCLUSION OF THE PROOF



By considering dependent random variables, we needed the introduction of the burning parameter b_n (see Eq.(23)). This situation brings extra technicalities to conclude the proof.

From the previous inequality and (22), we get

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} M_t \geq \varepsilon \right) \\ & \leq b_n \exp \left(-\frac{(1-c)n\varepsilon^2}{8b_n^2} \right) + 16\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)} \right) n \exp \left(-\frac{(c_n - b_n)C(m, \tau) (\varepsilon/2 - L\rho^{b_n})^2}{4} \right). \end{aligned}$$

Note that $(c_n - b_n)\varepsilon\rho^{b_n} \underset{n \rightarrow \infty}{=} o(n\varepsilon n^{q \log(\rho)}) \underset{n \rightarrow \infty}{=} o(n^{1+\xi+q \log(\rho)})$ because by assumption $\varepsilon \underset{n \rightarrow \infty}{=} o(n^\xi)$. However, by choice of q we have

$$1 + \xi + q \log(\rho) = 1 + \xi + \frac{1 + \xi}{\log(1/\rho)} \log(\rho) = 0,$$

and we finally get that $(c_n - b_n)\varepsilon\rho^{b_n} \underset{n \rightarrow \infty}{=} o(1)$. We deduce that for n large enough it holds

$$\exp \left(-\frac{(c_n - b_n)C(m, \tau) (\varepsilon/2 - L\rho^{b_n})^2}{4} \right) \leq 2 \exp \left(-\frac{(c_n - b_n)C(m, \tau)\varepsilon^2}{16} \right).$$

Then, noticing that

$$\exp \left(-\frac{(1-c)n\varepsilon^2}{8b_n^2} \right) \underset{n \rightarrow \infty}{=} \mathcal{O} \left(\exp \left(-\frac{(c_n - b_n)C(m, \tau)\varepsilon^2}{16b_n^2} \right) \right),$$

we finally get for n large enough

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} M_t \geq \varepsilon \right) \\ & \leq \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)} \right) + 1 \right] b_n \exp \left(-\frac{(c_n - b_n)C(m, \tau)\varepsilon^2}{16b_n^2} \right). \end{aligned}$$

B.2 Proof of Theorem 4

Theorem 3 shows that

$$\begin{aligned} & \mathbb{P} \left(\left| \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \mathcal{M}^n \right| \geq \varepsilon \right) \\ & \leq \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)} \right) + 1 \right] b_n \exp \left(-\frac{(c_n - b_n)C(m, \tau)\varepsilon^2}{16b_n^2} \right), \end{aligned} \quad (25)$$

and the assumption on the space \mathcal{H} gives that for some $\theta > 0$, it holds for any $\eta > 0$, $\log \mathcal{N}(\mathcal{H}, \eta) = \mathcal{O}(\eta^{-\theta})$. By taking $\varepsilon = \frac{\log(n) \log(\log n)}{n^{\frac{1}{2+\theta}}}$ it is straightforward to prove that the logarithm of the right hand side of Eq.(25) goes to $-\infty$ as $n \rightarrow +\infty$. This concludes the proof of the first part of Theorem 4.

Since the result from Theorem 3 trivially holds by considering $h_1 = \dots = h_{n-1} = h^*$, the previous computations show that for any $\delta > 0$ there exists some $N \in \mathbb{N}$ such that for any $n \geq N$ it holds with probability at least $1 - \delta$,

$$\left| \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \mathcal{M}^n \right| \vee |\mathcal{M}^n(h^*, \dots, h^*) - \mathcal{M}^n| \leq \frac{\log(n) \log(\log n)}{n^{\frac{1}{2+\theta}}}.$$

Hence, by considering that the online learner has a regret bound \mathfrak{R}_n (cf. Definition 5), we get that for any $\delta > 0$ there exists some $N \in \mathbb{N}$ such that for any $n \geq N$ it holds with probability at least $1 - \delta$,

$$\begin{aligned} & \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \mathcal{R}(h^*) \\ \leq & \frac{1}{n - c_n} \sum_{t=c_n}^{n-1} \mathcal{R}(h_{t-b_n}) - \mathcal{M}^n + \mathcal{M}^n - \mathcal{M}^n(h^*, \dots, h^*) + \mathcal{M}^n(h^*, \dots, h^*) - \mathcal{R}(h^*) \\ \leq & 2 \frac{\log(n) \log(\log n)}{n^{\frac{1}{2+\theta}}} + \mathcal{M}^n - \inf_{h \in \mathcal{H}} \mathcal{M}^n(h, \dots, h) \leq 2 \frac{\log(n) \log(\log n)}{n^{\frac{1}{2+\theta}}} + \mathfrak{R}_n, \end{aligned}$$

which concludes the proof of Theorem 4.

B.3 Proof of Theorem 5

The proof of Theorem 5 has two main steps. First, we show that $\mathcal{R}(\widehat{h})$ is close to $\min_{c_n \leq t \leq n-1} \mathcal{R}(h_{t-b_n}) + 2c_\gamma(n-t)$ with high probability. Then we show that $\min_{c_n \leq t \leq n-1} \mathcal{R}(h_{t-b_n}) + 2c_\gamma(n-t)$ is close to \mathcal{M}^n with high probability. The second step is similar to the proof of Wang et al. (2012). For the first step, we need a concentration inequality for U-statistics of order two for uniformly ergodic Markov chains. This is where we use the Hoeffding decomposition and Theorem 1 (see Section 2.2).

Let us recall that for any $1 \leq t \leq n-2$, $\widehat{\mathcal{R}}(h_{t-b_n}, t+1) = \binom{n-t}{2}^{-1} \sum_{k>i, i \geq t+1}^n \ell(h_{t-b_n}, X_i, X_k)$. We define

$$\ell(h, x) := \mathbb{E}_\pi[\ell(h, X, x)] - \mathcal{R}(h), \text{ and } \widetilde{\ell}(h, x, y) = \ell(h, x, y) - \ell(h, x) - \ell(h, y) - \mathcal{R}(h).$$

Then for any $t \in \{b_n + 1, \dots, n-2\}$ we have the following decomposition

$$\widehat{\mathcal{R}}(h_{t-b_n}, t+1) - \mathcal{R}(h_{t-b_n}) = \binom{n-t}{2}^{-1} \sum_{k>i, i \geq t+1}^n \widetilde{\ell}(h_{t-b_n}, X_i, X_k) + \frac{2}{n-t} \sum_{i=t+1}^n \ell(h_{t-b_n}, X_i). \quad (26)$$

One can check that for any $x \in E$, $\mathbb{E}_\pi [\tilde{\ell}(h, X, x)] = \mathbb{E}_\pi [\tilde{\ell}(h, x, X)] = 0$. Moreover, for any hypothesis $h \in \mathcal{H}$, $\|\tilde{\ell}(h, \cdot, \cdot)\|_\infty \leq 4$ (because the loss function ℓ takes its value in $[0, 1]$). Hence, for any fixed hypothesis $h \in \mathcal{H}$, the kernel $\tilde{\ell}(h, \cdot, \cdot)$ satisfies Assumption 3. Applying Theorem 1, we know that there exist constants $\beta, \kappa > 0$ such that for any $t \in \{b_n + 1, \dots, n - 2\}$ and for any $\gamma \in (0, 1)$, it holds with probability at least $1 - \gamma$,

$$\left| \binom{n-t}{2}^{-1} \sum_{k>i, i \geq t+1}^n \tilde{\ell}(h_{t-b_n}, X_i, X_k) \right| \leq \kappa \frac{\log(n-t-1)}{n-t-1} \log((\beta \vee e^1) \log(n-t+1)/\gamma)^2.$$

Note that we used that for $u = \log((\beta \vee e^1) \log(n-t+1)/\gamma) \geq 1$ it holds

$$\log n \left\{ \frac{u}{n} \vee \left(\frac{u}{n} \right)^2 \right\} \leq \frac{\log n}{n} u^2.$$

Using Proposition 7, we also have that for any $t \in \{b_n + 1, \dots, n - 2\}$ and any $\varepsilon > 0$,

$$\mathbb{P} \left(\left| \frac{2}{n-t} \sum_{i=t+1}^n \ell(h_{t-b_n}, X_i) \right| > \varepsilon \right) \leq 32 \exp(-C(m, \tau)(n-t)\varepsilon^2),$$

where $C(m, \tau) = (Km^2\tau^2)^{-1} > 0$ for some universal constant K (one can check from the proof of Proposition 7 that $K = 7 \times 10^3$ fits). We get that for any $t \in \{b_n + 1, \dots, n - 2\}$ and any $\gamma \in (0, 1)$, it holds with probability at least $1 - \gamma$,

$$\left| \frac{2}{n-t} \sum_{i=t+1}^n \ell(h_{t-b_n}, X_i) \right| \leq \frac{\log(32/\gamma)^{1/2} C(m, \tau)^{-1/2}}{\sqrt{n-t}}.$$

We deduce that for any $t \in \{b_n + 1, \dots, n - 2\}$ and any fixed $\gamma \in (0, 1)$, it holds with probability at least $1 - \gamma$,

$$\left| \widehat{\mathcal{R}}(h_{t-b_n}, t+1) - \mathcal{R}(h_{t-b_n}) \right| \leq C(m, \tau)^{-1/2} \sqrt{\frac{\log(64/\gamma)}{n-t}},$$

i.e.

$$\mathbb{P} \left(\left| \widehat{\mathcal{R}}(h_{t-b_n}, t+1) - \mathcal{R}(h_{t-b_n}) \right| \geq c_\gamma(n-t) \right) \leq \frac{\gamma}{(n-c_n)(n-c_n+1)}. \quad (27)$$

Based on the selection procedure of the hypothesis \widehat{h} defined in Eq.(12), the concentration result Eq.(27) allows us to show that $\mathcal{R}(\widehat{h})$ is close to $\min_{c_n \leq t \leq n-1} \mathcal{R}(h_{t-b_n}) + 2c_\gamma(n-t)$ with high probability. This is stated by Lemma 5 which is proved in Section B.4.

Lemma 5 *Let h_0, \dots, h_{n-1} be the set of hypotheses generated by an arbitrary online algorithm \mathcal{A} working with a pairwise loss ℓ which satisfies the conditions given in Theorem 3. Then for any $\gamma \in (0, 1)$, we have*

$$\mathbb{P} \left(\mathcal{R}(\widehat{h}) > \min_{c_n \leq t < n-1} (\mathcal{R}(h_{t-b_n}) + 2c_\gamma(n-t)) \right) \leq \gamma.$$

To conclude the proof, we need to show that $\min_{c_n \leq t \leq n-1} \mathcal{R}(h_{t-b_n}) + 2c_\gamma(n-t)$ is close to \mathcal{M}^n .

First we remark that

$$\begin{aligned}
 & \min_{c_n \leq t \leq n-1} \mathcal{R}(h_{t-b_n}) + 2c_\gamma(n-t) \\
 = & \min_{c_n \leq t \leq n-1} \min_{t \leq i \leq n-1} \mathcal{R}(h_{i-b_n}) + 2c_\gamma(n-i) \\
 \leq & \min_{c_n \leq t \leq n-1} \frac{1}{n-t} \sum_{i=t}^{n-1} (\mathcal{R}(h_{i-b_n}) + 2c_\gamma(n-i)) \\
 \leq & \min_{c_n \leq t \leq n-1} \left(\frac{1}{n-t} \sum_{i=t}^{n-1} \mathcal{R}(h_{i-b_n}) + \frac{2}{n-t} \sum_{i=t}^{n-1} \sqrt{\frac{C(m, \tau)^{-1} \log \frac{64(n-c_n)(n-c_n+1)}{\gamma}}{n-i}} \right) \\
 \leq & \min_{c_n \leq t \leq n-1} \left(\frac{1}{n-t} \sum_{i=t}^{n-1} \mathcal{R}(h_{i-b_n}) + \frac{2}{n-t} \sum_{i=t}^{n-1} \sqrt{\frac{2C(m, \tau)^{-1} \log \frac{64(n-c_n+1)}{\gamma}}{n-i}} \right) \\
 \leq & \min_{c_n \leq t \leq n-1} \left(\frac{1}{n-t} \sum_{i=t}^{n-1} \mathcal{R}(h_{i-b_n}) + 4\sqrt{\frac{2C(m, \tau)^{-1} \log \frac{64(n-c_n+1)}{\gamma}}{n-t}} \right),
 \end{aligned}$$

where the last inequality holds because $\sum_{i=1}^{n-t} \sqrt{1/i} \leq 2\sqrt{n-t}$. Indeed, $x \mapsto 1/\sqrt{x}$ is a decreasing and continuous function and a classical serie/integral approach leads to

$$\sum_{i=1}^{n-t} \sqrt{1/i} \leq 1 + \int_1^{n-t} \frac{1}{\sqrt{x}} dx = 1 + [2\sqrt{x}]_1^{n-t} \leq 2\sqrt{n-t}.$$

We define $\mathcal{M}_t^n := \frac{1}{n-t} \sum_{m=t}^{n-1} M_m$. From Theorem 3, one can see that for each $t = c_n, \dots, n-1$,

$$\mathbb{P} \left(\frac{1}{n-t} \sum_{i=t}^{n-1} \mathcal{R}(h_{i-b_n}) \geq \mathcal{M}_t^n + \varepsilon \right) \leq \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)} \right) + 1 \right] b_n \exp \left(-\frac{(t-b_n)C(m, \tau)\varepsilon^2}{16b_n^2} \right).$$

Let us set

$$K_t = \mathcal{M}_t^n + 4\sqrt{\frac{2C(m, \tau)^{-1} \log \frac{64(n-c_n+1)}{\gamma}}{n-t}} + \varepsilon.$$

Using the fact that if $\min(a_1, a_2) \leq \min(b_1, b_2)$ then either $a_1 \leq b_1$ or $a_2 \leq b_2$, we can write

$$\begin{aligned}
 & \mathbb{P} \left(\min_{c_n \leq t \leq n-1} \mathcal{R}(h_{t-b_n}) + 2c_\gamma(n-t) \geq \min_{c_n \leq t \leq n-1} K_t \right) \\
 \leq & \mathbb{P} \left(\min_{c_n \leq t \leq n-1} \left(\frac{1}{n-t} \sum_{i=t}^{n-1} \mathcal{R}(h_{i-b_n}) + 4\sqrt{\frac{2C(m, \tau)^{-1}}{n-t} \log \frac{64(n-c_n+1)}{\gamma}} \right) \geq \min_{c_n \leq t \leq n-1} K_t \right) \\
 \leq & \sum_{t=c_n}^{n-1} \mathbb{P} \left(\frac{1}{n-t} \sum_{i=t}^{n-1} \mathcal{R}(h_{i-b_n}) + 4\sqrt{\frac{2C(m, \tau)^{-1}}{n-t} \log \frac{64(n-c_n+1)}{\gamma}} \geq K_t \right) \\
 = & \sum_{t=c_n}^{n-1} \mathbb{P} \left(\frac{1}{n-t} \sum_{i=t}^{n-1} \mathcal{R}(h_{i-b_n}) \geq \mathcal{M}_t^n + \varepsilon \right) \\
 \leq & (n-c_n) \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)} \right) + 1 \right] b_n \exp \left(-\frac{(c_n-b_n)C(m, \tau)\varepsilon^2}{16b_n^2} \right) \\
 \leq & \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)} \right) + 1 \right] \exp \left(-\frac{(c_n-b_n)C(m, \tau)\varepsilon^2}{16b_n^2} + 2 \log n \right).
 \end{aligned}$$

Using Lemma 5, we get

$$\begin{aligned}
 & \mathbb{P} \left(\mathcal{R}(\hat{h}) \geq \min_{c_n \leq t \leq n-1} \mathcal{M}_t^n + 4\sqrt{\frac{2C(m, \tau)^{-1}}{n-t} \log \frac{64(n-c_n+1)}{\gamma}} + \varepsilon \right) \\
 \leq & \gamma + \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)} \right) + 1 \right] \exp \left(-\frac{(c_n-b_n)C(m, \tau)\varepsilon^2}{16b_n^2} + 2 \log n \right),
 \end{aligned}$$

which gives in particular

$$\begin{aligned}
 & \mathbb{P} \left(\mathcal{R}(\hat{h}) \geq \mathcal{M}^n + 4\sqrt{\frac{2C(m, \tau)^{-1}}{n-c_n} \log \frac{64(n-c_n+1)}{\gamma}} + \varepsilon \right) \\
 \leq & \gamma + \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{8\text{Lip}(\varphi)} \right) + 1 \right] \exp \left(-\frac{(c_n-b_n)C(m, \tau)\varepsilon^2}{16b_n^2} + 2 \log n \right).
 \end{aligned}$$

We substitute ε with $\varepsilon/2$ and we choose γ such that $4\sqrt{\frac{2C(m, \tau)^{-1}}{n-c_n} \log \frac{64(n-c_n+1)}{\gamma}} = \varepsilon/2$ with n large enough to ensure that $\gamma < 1$. We have for any $c > 0$,

$$\begin{aligned}
 & \mathbb{P} \left(\mathcal{R}(\hat{h}) \geq \mathcal{M}^n + 4\sqrt{\frac{2C(m, \tau)^{-1}}{n-c_n} \log \frac{64(n-c_n+1)}{\gamma}} + \frac{\varepsilon}{2} \right) \\
 \leq & 64(n-c_n+1) \exp \left(-\frac{(n-c_n)C(m, \tau)\varepsilon^2}{128} \right) + \left[32\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{16\text{Lip}(\varphi)} \right) + 1 \right] \times \\
 & \exp \left(-\frac{(c_n-b_n)C(m, \tau)\varepsilon^2}{(16b_n)^2} + 2 \log n \right) \\
 \leq & 32 \left[\mathcal{N} \left(\mathcal{H}, \frac{\varepsilon}{16\text{Lip}(\varphi)} \right) + 1 \right] \exp \left(-\frac{(c_n-b_n)C(m, \tau)\varepsilon^2}{(16b_n)^2} + 2 \log n \right),
 \end{aligned}$$

where these inequalities hold for n large enough.

B.4 Proof of Lemma 5

Let

$$T^* := \arg \min_{c_n \leq t < n-1} (\mathcal{R}(h_{t-b_n}) + 2c_\gamma(n-t)),$$

and $h^* = h_{T^*-b_n}$ is the corresponding hypothesis that minimizes the penalized true risk and let $\widehat{\mathcal{R}}^* = \widehat{\mathcal{R}}(h^*, T^* + 1)$ to be the penalized empirical risk of $h_{T^*-b_n}$. Set, for brevity

$$\widehat{\mathcal{R}}_{t-b_n} = \widehat{\mathcal{R}}(h_{t-b_n}, t+1),$$

and let

$$\widehat{T} := \arg \min_{c_n \leq t < n-1} (\widehat{\mathcal{R}}_{t-b_n} + c_\gamma(n-t)),$$

where \widehat{h} coincides with $h_{\widehat{T}-b_n}$. Using this notation and since

$$\widehat{\mathcal{R}}_{\widehat{T}-b_n} + c_\gamma(n-\widehat{T}) \leq \widehat{\mathcal{R}}^* + c_\gamma(n-T^*),$$

holds with certainty, we have

$$\begin{aligned} & \mathbb{P} \left(\mathcal{R}(\widehat{h}) > \mathcal{R}(h^*) + \mathcal{E} \right) \\ &= \mathbb{P} \left(\mathcal{R}(\widehat{h}) > \mathcal{R}(h^*) + \mathcal{E}, \widehat{\mathcal{R}}_{\widehat{T}-b_n} + c_\gamma(n-\widehat{T}) \leq \widehat{\mathcal{R}}^* + c_\gamma(n-T^*) \right) \\ &\leq \mathbb{P} \left(\bigcup_{c_n \leq t \leq n-1} \left\{ \mathcal{R}(h_{t-b_n}) > \mathcal{R}(h^*) + \mathcal{E}, \widehat{\mathcal{R}}_{t-b_n} + c_\gamma(n-t) \leq \widehat{\mathcal{R}}^* + c_\gamma(n-T^*) \right\} \right) \\ &\leq \sum_{t=c_n}^{n-1} \mathbb{P} \left(\mathcal{R}(h_{t-b_n}) > \mathcal{R}(h^*) + \mathcal{E}, \widehat{\mathcal{R}}_{t-b_n} + c_\gamma(n-t) \leq \widehat{\mathcal{R}}^* + c_\gamma(n-T^*) \right), \end{aligned}$$

where \mathcal{E} is a positive-valued random variable to be specified. Now we remark that if

$$\widehat{\mathcal{R}}_{t-b_n} + c_\gamma(n-t) \leq \widehat{\mathcal{R}}^* + c_\gamma(n-T^*), \tag{28}$$

holds, then at least one of the following three conditions must hold

$$\begin{aligned} (i) \quad & \widehat{\mathcal{R}}_{t-b_n} && \leq \mathcal{R}(h_{t-b_n}) - c_\gamma(n-t) \\ (ii) \quad & \widehat{\mathcal{R}}^* && > \mathcal{R}(h^*) + c_\gamma(n-T^*) \\ (iii) \quad & \mathcal{R}(h_{t-b_n}) - \mathcal{R}(h^*) && \leq 2c_\gamma(n-T^*). \end{aligned}$$

Stated otherwise, if Eq.(28) holds for some $t \in \{c_n, \dots, n-1\}$ then

- either $t = T^*$ and (iii) holds trivially.
- or $t \neq T^*$ which can occur because
 - $\widehat{\mathcal{R}}_{t-b_n}$ underestimates $\mathcal{R}(h_{t-b_n})$ and (i) holds.
 - $\widehat{\mathcal{R}}^*$ overestimates $\mathcal{R}(h^*)$ and (ii) holds.
 - n is too small to statistically distinguish $\mathcal{R}(h_{t-b_n})$ and $\mathcal{R}(h^*)$, and (iii) holds.

Therefore, for any fixed t , we have

$$\begin{aligned} & \mathbb{P} \left(\mathcal{R}(h_{t-b_n}) > \mathcal{R}(h^*) + \mathcal{E}, \widehat{\mathcal{R}}_{t-b_n} + c_\gamma(n-t) \leq \widehat{\mathcal{R}}^* + c_\gamma(n-T^*) \right) \\ & \leq \mathbb{P} \left(\widehat{\mathcal{R}}_{t-b_n} \leq \mathcal{R}(h_{t-b_n}) - c_\gamma(n-t) \right) + \mathbb{P} \left(\widehat{\mathcal{R}}^* > \mathcal{R}(h^*) + c_\gamma(n-T^*) \right) \\ & \quad + \mathbb{P} \left(\mathcal{R}(h_{t-b_n}) - \mathcal{R}(h^*) \leq 2c_\gamma(n-T^*), \mathcal{R}(h_{t-b_n}) > \mathcal{R}(h^*) + \mathcal{E} \right). \end{aligned}$$

By choosing $\mathcal{E} = 2c_\gamma(n-T^*)$, the last term in the previous inequality is zero and we can write

$$\begin{aligned} & \mathbb{P} \left(\mathcal{R}(\widehat{h}) > \mathcal{R}(h^*) + 2c_\gamma(n-T^*) \right) \\ & \leq \sum_{t=c_n}^{n-1} \mathbb{P} \left(\widehat{\mathcal{R}}_{t-b_n} \leq \mathcal{R}(h_{t-b_n}) - c_\gamma(n-t) \right) + (n-c_n) \mathbb{P} \left(\widehat{\mathcal{R}}^* > \mathcal{R}(h^*) + c_\gamma(n-T^*) \right) \\ & \leq (n-c_n) \frac{\gamma}{(n-c_n)(n-c_n+1)} + (n-c_n) \left\{ \sum_{t=c_n}^{n-1} \mathbb{P} \left(\widehat{\mathcal{R}}_{t-b_n} > \mathcal{R}(h_{t-b_n}) + c_\gamma(n-t) \right) \right\} \quad (\text{Using (27)}) \\ & \leq \frac{\gamma}{n-c_n+1} + (n-c_n)^2 \frac{\gamma}{(n-c_n)(n-c_n+1)} \quad (\text{Using Eq.(27)}) \\ & \leq \frac{\gamma}{n-c_n+1} + (n-c_n) \frac{\gamma}{n-c_n+1} = \gamma. \end{aligned}$$

B.5 Proof of Corollary 1

The proof of Corollary 1 is analogous to the proof of Theorem 4 by applying Theorem 5 (instead of Theorem 3) and by choosing $\varepsilon = \frac{\log^2 n}{n^{\frac{1}{2+\theta}}}$.

C. Proofs for Section 5

C.1 Proof of Theorem 6

In the following, \mathbb{P}_g will denote the distribution of the Markov chain if the stationary distribution of the chain is assumed to have a density g with respect to the Lebesgue measure on \mathbb{R} . We consider $q = q_1 \vee q_2$ where $q_1, q_2 \in [1, \infty)$ are such that $\frac{1}{p_1} + \frac{1}{q_1} = 1$ and $\frac{1}{p_2} + \frac{1}{q_2} = 1$. The main tool of the proof is the Hoeffding (also called canonical) decomposition of the U-statistics $\widehat{\theta}_m$. We introduce the processes U_n and P_n defined by

$$U_n(h) = \frac{1}{n(n-1)} \sum_{i \neq j=1}^n h(X_i, X_j), \quad P_n(h) = \frac{1}{n} \sum_{i=1}^n h(X_i).$$

We also define $P(h) = \langle h, f \rangle$. By setting, for all $m \in \mathcal{M}$,

$$H_m(x, y) = \sum_{l \in \mathcal{L}_m} (p_l(x) - a_l)(p_l(y) - a_l),$$

with $a_l = \langle f, p_l \rangle$, we obtain the decomposition

$$\widehat{\theta}_m = U_n(H_m) + (P_n - P)(2\Pi_{S_m}(f)) + \|\Pi_{S_m}(f)\|_2^2.$$

Let us consider β in $]0, 1[$. Since

$$\mathbb{P}_f(T_\alpha \leq 0) = \mathbb{P}_f \left(\sup_{m \in \mathcal{M}} (\widehat{\theta}_m + \|f_0\|_2^2 - \frac{2}{n} \sum_{i=1}^n f_0(X_i) - t_m(u_\alpha)) \leq 0 \right),$$

we have

$$\mathbb{P}_f(T_\alpha \leq 0) \leq \inf_{m \in \mathcal{M}} \mathbb{P}_f \left(\widehat{\theta}_m + \|f_0\|_2^2 - \frac{2}{n} \sum_{i=1}^n f_0(X_i) - t_m(u_\alpha) \leq 0 \right).$$

Since $\|f - \Pi_{S_m}(f)\|_2^2 = \|f\|_2^2 - \|\Pi_{S_m}(f)\|_2^2$, it holds

$$\begin{aligned} & \widehat{\theta}_m + \|f_0\|_2^2 - \frac{2}{n} \sum_{i=1}^n f_0(X_i) \\ &= U_n(H_m) + (P_n - P)(2\Pi_{S_m}(f)) - \|f - \Pi_{S_m}(f)\|_2^2 + \|f\|_2^2 + \|f_0\|_2^2 - 2P_n(f_0) \\ &= U_n(H_m) + (P_n - P)(2\Pi_{S_m}(f)) - \|f - \Pi_{S_m}(f)\|_2^2 + \|f - f_0\|_2^2 + 2P(f_0) - 2P_n(f_0), \end{aligned}$$

which leads to

$$\begin{aligned} \mathbb{P}_f(T_\alpha \leq 0) &\leq \inf_{m \in \mathcal{M}} \mathbb{P}_f \left(U_n(H_m) + (P_n - P)(2\Pi_{S_m}(f) - 2f) + (P_n - P)(2f - 2f_0) + \|f - f_0\|_2^2 \right. \\ &\quad \left. \leq \|f - \Pi_{S_m}(f)\|_2^2 + t_m(u_\alpha) \right). \end{aligned} \quad (29)$$

We then need to control $U_n(H_m)$, $(P_n - P)(2\Pi_{S_m}(f) - 2f)$, $(P_n - P)(2f - 2f_0)$ for every $m \in \mathcal{M}$. **Control of $U_n(H_m)$.**

H_m is π -canonical and a direct application of Theorem 1 leads to the following Lemma (the proof of Lemma 6 is postponed to Section C.2).

Lemma 6 *Let us assume that the stationary distribution of the Markov chain $(X_i)_{i \geq 1}$ has density f with respect to the Lebesgue measure on \mathbb{R} . For all $m = (l, D)$ with $l \in \{1, 2, 3\}$ and $D \in \mathbb{D}_l$, introduce $\{p_l, l \in \mathcal{L}_m\}$ defined as in page 21 and $Z_m = \frac{1}{n(n-1)} \sum_{i \neq j=1}^n H_m(X_i, X_j)$, with $H_m(x, y) = \sum_{l \in \mathcal{L}_m} (p_l(x) - \langle f, p_l \rangle)(p_l(y) - \langle f, p_l \rangle)$. There exist some constants $C, \beta > 0$ (both depending on the Markov chain $(X_i)_{i \geq 1}$ while C also depends on φ) such that, for all $l \in \{1, 2, 3\}$, $D \in \mathbb{D}_l$ and $u \geq 1$, it holds with probability at least $1 - \beta e^{-u} \log n$,*

$$|Z_{(l,D)}| \leq C (\|f\|_\infty + 1) DR(n, u),$$

where $R(n, u) = \log n \left\{ \frac{u}{n} + \left(\frac{u}{n} \right)^2 \right\}$.

We deduce that there exist $C, \beta > 0$ such that for any $\gamma \in (0, 1 \wedge (e^{-1} 3\beta \log n))$ and any $m = (l, D) \in \mathcal{M}$,

$$\mathbb{P}_f \left(U_n(H_m) \leq -C (\|f\|_\infty + 1) DR \left(n, \log \left\{ \frac{3\beta \log n}{\gamma} \right\} \right) \right) \leq \gamma/3. \quad (30)$$

From Eq.(29) and Eq.(30) we get that

$$\begin{aligned} \mathbb{P}_f(T_\alpha \leq 0) &\leq \frac{\gamma}{3} + \inf_{m \in \mathcal{M}} \mathbb{P}_f \left((P_n - P)(2\Pi_{S_m}(f) - 2f) + (P_n - P)(2f - 2f_0) + \|f - f_0\|_2^2 \right. \\ &\quad \left. \leq \|f - \Pi_{S_m}(f)\|_2^2 + t_m(u_\alpha) + C(\|f\|_\infty + 1) DR \left(n, \log \left\{ \frac{3\beta \log n}{\gamma} \right\} \right) \right). \end{aligned} \quad (31)$$

Control of $(P_n - P)(2\Pi_{S_m}(f) - 2f)$.

It is easy to check that there exists some constant $C' > 0$ such that for all l in $\{1, 2\}$, D in \mathbb{D}_l ,

$$\left| 2\Pi_{S_{(l,D)}}(f)(X_i) - 2f(X_i) \right| \leq C'\|f\|_\infty.$$

Indeed,

- when $l = 1$, for any $k \in \mathbb{Z}$,

$$\langle \sqrt{D}\mathbf{1}_{[k/D, (k+1)/D]}, f \rangle = \int \sqrt{D}\mathbf{1}_{[k/D, (k+1)/D]}(x)f(x)dx \leq D^{-1/2}\|f\|_\infty.$$

Hence,

$$\begin{aligned} \sup_x |\Pi_{S_{(1,D)}}(f)(x)| &\leq \sup_x \sum_{k \in \mathbb{Z}} \left| \langle \sqrt{D}\mathbf{1}_{[k/D, (k+1)/D]}, f \rangle \right| \sqrt{D}\mathbf{1}_{[k/D, (k+1)/D]}(x) \\ &\leq D^{-1/2}\|f\|_\infty \sup_x \sum_{k \in \mathbb{Z}} \sqrt{D}\mathbf{1}_{[k/D, (k+1)/D]}(x) = \|f\|_\infty. \end{aligned}$$

- when $l = 2$, $D = 2^J$ for some $J \in \mathbb{N}$ and we have for any $k \in \mathbb{Z}$,

$$\langle \varphi_{J,k}, f \rangle = \int 2^{J/2}\varphi(2^Jx - k)f(x)dx \leq \|f\|_\infty \int 2^{J/2}|\varphi(2^Jx)|dx \leq 2^{-J/2}\|f\|_\infty\|\varphi\|_1.$$

Hence,

$$\begin{aligned} \sup_x |\Pi_{S_{(2,D)}}(f)(x)| &\leq \sup_x \sum_{k \in \mathbb{Z}} |\langle \varphi_{J,k}, f \rangle| \times |\varphi_{J,k}(x)| \\ &\leq 2^{-J/2}\|f\|_\infty\|\varphi\|_1 \sup_x \sum_{k \in \mathbb{Z}} |2^{J/2}\varphi(2^Jx - k)| \leq c\|f\|_\infty\|\varphi\|_1, \end{aligned}$$

where $c > 0$ is a constant depending only on φ since φ is bounded and compactly supported. Stated otherwise, there is only a finite number of integers $k \in \mathbb{Z}$ (which is independent of x and J) such that for any $x \in \mathbb{R}$ and any $J \in \mathbb{Z}$, $2^Jx - k$ falls into the support of φ .

Moreover, it is proved in DeVore and Lorentz (1993, Page 269), that one can take C' such that for all D in \mathbb{D}_3 ,

$$\left| 2\Pi_{S_{(3,D)}}(f)(X_i) - 2f(X_i) \right| \leq C'\|f\|_\infty \log(D + 1).$$

Since

$$\mathbb{E}_{X \sim \pi} (2\Pi_{S_m}(f)(X) - 2f(X))^2 \leq 4\|f\|_\infty \|\Pi_{S_m}(f) - f\|_2^2,$$

we can deduce using Proposition 11 (see Section D.2) that for all $m = (l, D) \in \mathcal{M}$,

$$\mathbb{P}_f \left((P_n - P)(2\Pi_{S_m}(f) - 2f) < -\frac{2C' \log(3C_\chi/\gamma)qA_1\|f\|_\infty \log(D+1)}{n} \right. \\ \left. - 2\sqrt{\frac{2\log(3C_\chi/\gamma)qA_2\|f\|_\infty}{n}} \|\Pi_{S_m}(f) - f\|_2 \right) \leq \frac{\gamma}{3}.$$

Considering some $\varepsilon \in]0, 2[$, we use the inequality $\forall a, b \in \mathbb{R}, 2ab \leq 4a^2/\varepsilon + \varepsilon b^2/4$ and we obtain that for any $m = (l, D) \in \mathcal{M}$,

$$\mathbb{P}_f \left((P_n - P)(2\Pi_{S_m}(f) - 2f) + \frac{\varepsilon}{4} \|\Pi_{S_m}(f) - f\|_2^2 < -\frac{2C' \log(3C_\chi/\gamma)qA_1\|f\|_\infty \log(D+1)}{n} \right. \\ \left. - \frac{8\log(3C_\chi/\gamma)qA_2\|f\|_\infty}{\varepsilon n} \right) \leq \frac{\gamma}{3}. \quad (32)$$

The control of $(P_n - P)(2f - 2f_0)$ is computed in the same way and we get

$$\mathbb{P}_f \left((P_n - P)(2f - 2f_0) + \frac{\varepsilon}{4} \|f - f_0\|_2^2 < -\frac{4\log(3C_\chi/\gamma)qA_1(\|f\|_\infty + \|f_0\|_\infty)}{n} \right. \\ \left. - \frac{8\log(3C_\chi/\gamma)qA_2\|f\|_\infty}{\varepsilon n} \right) \leq \frac{\gamma}{3}. \quad (33)$$

Finally, we deduce from Eq.(31), Eq.(32) and Eq.(33) that if there exists some $m = (l, D)$ in \mathcal{M} such that

$$\left(1 - \frac{\varepsilon}{4}\right) \|f - f_0\|_2^2 > \left(1 + \frac{\varepsilon}{4}\right) \|f - \Pi_{S_m}(f)\|_2^2 + \frac{8\log(3C_\chi/\gamma)qA_2\|f\|_\infty}{\varepsilon n} \\ + \frac{4\log(3C_\chi/\gamma)qA_1(\|f\|_\infty + \|f_0\|_\infty)}{n} \\ + \frac{8\log(3C_\chi/\gamma)qA_2\|f\|_\infty}{\varepsilon n} + \frac{2C' \log(3C_\chi/\gamma)qA_1\|f\|_\infty \log(D+1)}{n} \\ + t_m(u_\alpha) + C(\|f\|_\infty + 1) DR \left(n, \log \left\{ \frac{3\beta \log n}{\gamma} \right\} \right),$$

i.e. such that

$$\left(1 - \frac{\varepsilon}{4}\right) \|f - f_0\|_2^2 > \left(1 + \frac{\varepsilon}{4}\right) \|f - \Pi_{S_m}(f)\|_2^2 + \frac{16\log(3C_\chi/\gamma)qA_2\|f\|_\infty}{\varepsilon n} \\ + 4(\|f\|_\infty(C' \log(D+1) + 1) + \|f_0\|_\infty) \frac{\log(3C_\chi/\gamma)qA_1}{n} \\ + t_m(u_\alpha) + C(\|f\|_\infty + 1) DR \left(n, \log \left\{ \frac{3\beta \log n}{\gamma} \right\} \right),$$

then

$$\mathbb{P}_f(T_\alpha \leq 0) \leq \gamma.$$

To conclude the proof of Theorem 6, it suffices to notice that for any $\varepsilon \in]0, 2[$, choosing $\eta > 0$ such that $1 + \eta = \frac{1 + \frac{\varepsilon}{4}}{1 - \frac{\varepsilon}{4}}$ leads to $\varepsilon = \frac{4\eta}{2 + \eta}$. One can immediately check that the condition $\varepsilon \in]0, 2[$ is equivalent to $\eta \in]0, 2[$. Noticing further that $\frac{1}{\varepsilon} = \frac{2 + \eta}{4\eta} < \frac{2 + 2}{4\eta} = \frac{1}{\eta}$, we deduce that for any $\eta \in]0, 2[$, if

$$\begin{aligned} \|f - f_0\|_2^2 &> (1 + \eta) \left\{ \|f - \Pi_{S_m}(f)\|_2^2 + \frac{16 \log(3C_\chi/\gamma)qA_2\|f\|_\infty}{\eta n} \right. \\ &\quad \left. + 4(\|f\|_\infty(C' \log(D + 1) + 1) + \|f_0\|_\infty) \frac{\log(3C_\chi/\gamma)qA_1}{n} \right. \\ &\quad \left. + t_m(u_\alpha) + C(\|f\|_\infty + 1)DR \left(n, \log \left\{ \frac{3\beta \log n}{\gamma} \right\} \right) \right\}, \end{aligned}$$

then

$$\mathbb{P}_f(T_\alpha \leq 0) \leq \gamma.$$

C.2 Proof of Lemma 6

Lemma 6 will follow from Theorem 1 if we can show that the function H_m is bounded. Let us denote $m = (l, D)$ for some $l \in \{1, 2, 3\}$ and $D \in \mathbb{D}_l$. Let us first remark that the Bessel's inequality states that

$$\sum_{k \in \mathcal{L}_m} |\langle p_k, f \rangle|^2 \leq \|f\|_2^2 = \int f(x)f(x)dx \leq \|f\|_\infty, \quad (34)$$

since $\int f(x)dx = 1$ and $f(x) \geq 0$, $\forall x$.

- If $l = 1$, then we notice that for any $k \in \mathbb{Z}$,

$$\begin{aligned} |\langle \sqrt{D}\mathbb{1}_{]k/D, (k+1)/D[}, f \rangle| &= \left| \int \sqrt{D}\mathbb{1}_{]k/D, (k+1)/D[}(x)f(x)dx \right| \\ &\leq \|f\|_\infty \sqrt{D} \int \mathbb{1}_{]k/D, (k+1)/D[}(x)dx \\ &\leq D^{-1/2} \|f\|_\infty. \end{aligned}$$

Then for any $x, y \in \mathbb{R}$ it holds

$$\begin{aligned} |H_m(x, y)| &\leq \sum_{k \in \mathcal{L}_m} |p_k(x)p_k(y)| + \sum_{k \in \mathcal{L}_m} |p_k(x)\langle p_k, f \rangle| + \sum_{k \in \mathcal{L}_m} |p_k(y)\langle p_k, f \rangle| + \sum_{k \in \mathcal{L}_m} |\langle p_k, f \rangle|^2 \\ &\leq \sum_{k \in \mathbb{Z}} D\mathbb{1}_{]k/D, (k+1)/D[}(x)\mathbb{1}_{]k/D, (k+1)/D[}(y) \\ &\quad + 2 \sup_z \sum_{k \in \mathbb{Z}} \sqrt{D}|\mathbb{1}_{]k/D, (k+1)/D[}(z)| \times |\langle \sqrt{D}\mathbb{1}_{]k/D, (k+1)/D[}, f \rangle| + \sum_{k \in \mathcal{L}_m} |\langle p_k, f \rangle|^2 \\ &\leq D + 2\|f\|_\infty + \|f\|_\infty, \end{aligned}$$

where in the last inequality we used Eq.(34).

- If $l = 2$ then $D = 2^J$ for some $J \in \mathbb{N}$ and we have for any $k \in \mathbb{Z}$,

$$\langle \varphi_{J,k}, f \rangle = \int 2^{J/2} \varphi(2^J x - k) f(x) dx \leq \|f\|_\infty \int 2^{J/2} |\varphi(2^J x)| dx \leq 2^{-J/2} \|f\|_\infty \|\varphi\|_1.$$

We get that for any $x, y \in \mathbb{R}$,

$$\begin{aligned} |H_m(x, y)| &\leq \sum_{k \in \mathcal{L}_m} |p_k(x) p_k(y)| + \sum_{k \in \mathcal{L}_m} |p_k(x) \langle p_k, f \rangle| + \sum_{k \in \mathcal{L}_m} |p_k(y) \langle p_k, f \rangle| + \sum_{k \in \mathcal{L}_m} |\langle p_k, f \rangle|^2 \\ &\leq \sum_{k \in \mathbb{Z}} 2^J \varphi(2^J x - k) \varphi(2^J y - k) + 2 \sup_z \sum_{k \in \mathbb{Z}} 2^{-J/2} \|f\|_\infty \|\varphi\|_1 2^{J/2} |\varphi(2^{J/2} z - k)| + \sum_{k \in \mathcal{L}_m} |\langle p_k, f \rangle|^2 \\ &\leq c 2^J + c' \|\varphi\|_1 \|f\|_\infty + \|f\|_\infty \\ &= cD + c' \|\varphi\|_1 \|f\|_\infty + \|f\|_\infty, \end{aligned}$$

for some constants $c, c' > 0$. In the last inequality we used Eq.(34) and the fact φ is bounded and compactly supported. Indeed, this implies that there is only a finite number of integers $k \in \mathbb{Z}$ (which is independent of x and J) such that for any $x \in \mathbb{R}$ and any $J \in \mathbb{Z}$, $2^J x - k$ falls into the support of φ .

- If $l = 3$ then we easily get for any $x, y \in [0, 1]$,

$$\begin{aligned} |H_m(x, y)| &\leq \sum_{k \in \mathcal{L}_m} |p_k(x) p_k(y)| + \sum_{k \in \mathcal{L}_m} |p_k(x) \langle p_k, f \rangle| + \sum_{k \in \mathcal{L}_m} |p_k(y) \langle p_k, f \rangle| + \sum_{k \in \mathcal{L}_m} |\langle p_k, f \rangle|^2 \\ &\leq 2D + 4D \|f\|_\infty + \|f\|_\infty. \end{aligned}$$

We deduce that in any case, H_m is bounded $c(1 + \|f\|_\infty)D$ for some constant $c > 0$ (depending only on φ) which concludes the proof of Lemma 6.

C.3 Proof of Corollary 2

Step 1: We start by providing an upper bound on $t_m(u_\alpha)$ with Lemma 7.

Lemma 7 *There exists a constant $C(\alpha) > 0$ such that for any $m = (l, D) \in \mathcal{M}$ it holds,*

$$t_m(u_\alpha) \leq W_m(\alpha),$$

where

$$W_m(\alpha) = C(\alpha) (\|f_0\|_\infty + 1) \left[DR(n, \log \log n) + \frac{\log \log n}{n} \right].$$

Proof of Lemma 7.

Let us recall that $t_m(u)$ denotes the $(1 - u)$ quantile of the distribution of \widehat{T}_m under the null hypothesis. One can easily see that $|\mathcal{M}| \leq 3(1 + \log_2 n)$. So, setting $\alpha_n = \alpha / (3(1 + \log_2 n))$,

$$\begin{aligned} \mathbb{P}_{f_0} \left(\sup_{m \in \mathcal{M}} (\widehat{T}_m - t_m(\alpha_n)) > 0 \right) &\leq \sum_{m \in \mathcal{M}} \mathbb{P}_{f_0} (\widehat{T}_m - t_m(\alpha_n) > 0) \\ &\leq \sum_{m \in \mathcal{M}} \alpha / (3(1 + \log_2 n)) \\ &\leq \alpha. \end{aligned}$$

By definition of u_α , this implies that $\alpha_n \leq u_\alpha$ and for all $m \in \mathcal{M}$,

$$t_m(u_\alpha) \leq t_m(\alpha_n).$$

Hence it suffices to upper bound $t_m(\alpha_n)$. Let $m = (l, D) \in \mathcal{M}$. We use the same notation as in the proof of Theorem 6 to obtain that

$$\widehat{T}_m = U_n(H_m) + (P_n - P)(2\Pi_{S_m}(f)) - 2P_n(f_0) + \|f_0\|_2^2 + \|\Pi_{S_m}(f)\|_2^2.$$

Under the null hypothesis, this reads as

$$\begin{aligned} \widehat{T}_m &= U_n(H_m) + (P_n - P)(2\Pi_{S_m}(f_0) - 2f_0) - \|f_0\|_2^2 + \|\Pi_{S_m}(f_0)\|_2^2 \\ &= U_n(H_m) + (P_n - P)(2\Pi_{S_m}(f_0) - 2f_0) - \|f_0 - \Pi_{S_m}(f_0)\|_2^2. \end{aligned}$$

We control $U_n(H_m)$ and $(P_n - P)(2\Pi_{S_m}(f_0) - 2f_0)$ exactly like in the proof of Theorem 6. From Lemma 6, there exist $C, \beta > 0$ such that for any $m = (l, D) \in \mathcal{M}$, it holds

$$\mathbb{P}_{f_0} \left(U_n(H_m) \leq C (\|f_0\|_\infty + 1) DR \left(n, \log \left\{ \frac{2\beta \log n}{\alpha_n} \right\} \right) \right) \leq \alpha_n/2. \quad (35)$$

Moreover, since

$$|2\Pi_{S_{(l,D)}}(f_0)(X_i) - 2f_0(X_i)| \leq C' \|f_0\|_\infty \log(D+1),$$

and

$$\mathbb{E}_{X \sim \pi} (2\Pi_{S_m}(f_0)(X) - 2f_0(X))^2 \leq 4\|f_0\|_\infty \|\Pi_{S_m}(f_0) - f_0\|_2^2,$$

we get using Proposition 11 (see Section D.2) that for all $m = (l, D) \in \mathcal{M}$,

$$\begin{aligned} \mathbb{P}_{f_0} \left((P_n - P)(2\Pi_{S_m}(f_0) - 2f_0) > \frac{2C' \log(2C_\chi/\alpha_n) q A_1 \|f_0\|_\infty \log(D+1)}{n} \right. \\ \left. + 2\sqrt{\frac{2\log(2C_\chi/\alpha_n) q A_2 \|f_0\|_\infty}{n}} \|\Pi_{S_m}(f_0) - f_0\|_2 \right) \leq \frac{\alpha_n}{2}. \end{aligned}$$

Using the inequality $\forall a, b \in \mathbb{R}, 2ab \leq a^2 + b^2$, and the fact that for $n \geq 16$, $\log(D+1) \leq \log(n^2+1)$, we obtain that there exists $C'' > 0$ such that

$$\mathbb{P}_{f_0} \left((P_n - P)(2\Pi_{S_m}(f_0) - 2f_0) - \|\Pi_{S_m}(f_0) - f_0\|_2^2 > \frac{C'' \|f_0\|_\infty \log(2C_\chi/\alpha_n) \log(n)}{n} \right) \leq \frac{\alpha_n}{2}.$$

We deduce that it holds

$$\mathbb{P}_{f_0} \left(\widehat{T}_m > C (\|f_0\|_\infty + 1) DR \left(n, \log \left\{ \frac{2\beta \log n}{\alpha_n} \right\} \right) + \frac{C'' \|f_0\|_\infty \log(2C_\chi/\alpha_n) \log(n)}{n} \right) \leq \alpha_n.$$

Noticing that there exists some constant $c(\alpha) > 0$ such that

$$\log \left\{ \frac{2\beta \log n}{\alpha_n} \right\} \vee \log(2C_\chi/\alpha_n) \leq c(\alpha) \log \log n,$$

we deduce by definition of $t_m(\alpha_n)$ that for some $c(\alpha) > 0$,

$$t_m(\alpha_n) \leq c(\alpha)C(\|f_0\|_\infty + 1)DR(n, \log \log n) + c(\alpha)\frac{C''\|f_0\|_\infty \log \log n}{n}.$$

■

Step 2: Proof of Corollary 2.

Let us fix $\gamma \in]0, 1[$ and $l \in \{1, 2, 3\}$. From Theorem 6 and Lemma 7, we deduce that if f satisfies

$$\|f - f_0\|_2^2 > (1 + \varepsilon) \inf_{D \in \mathcal{D}_l} \|f - \Pi_{S(l,D)}(f)\|_2^2 + W_{(l,D)}(\alpha) + V_{(l,D)}(\gamma),$$

then

$$\mathbb{P}_f(T_\alpha \leq 0) \leq \gamma.$$

It is thus a matter of giving an upper bound for

$$\inf_{D \in \mathcal{D}_l} \left\{ \|f - \Pi_{S(l,D)}(f)\|_2^2 + W_{(l,D)}(\alpha) + V_{(l,D)}(\gamma) \right\},$$

when f belongs to some specified classes of functions. Recall that

$$\mathcal{B}_s^{(l)}(P, M) = \{f \in L_2(\mathbb{R}) \mid \forall D \in \mathcal{D}_l, \|f - \Pi_{S(l,D)}(f)\|_2^2 \leq P^2 D^{-2s}, \|f\|_\infty \leq M\}.$$

We now assume that f belongs to $\mathcal{B}_s^{(l)}(P, M)$. Since $\|f - \Pi_{S(l,D)}(f)\|_2^2 \leq P^2 D^{-2s}$, we only need an upper bound for

$$\begin{aligned} & \inf_{D \in \mathcal{D}_l} \left\{ P^2 D^{-2s} + C(\alpha)(\|f_0\|_\infty + 1) \left[DR(n, \log \log n) + \frac{\log \log n}{n} \right] + C_1 \|f\|_\infty \frac{\log(3C_\chi/\gamma)}{\varepsilon n} \right. \\ & \left. + C_2 (\|f\|_\infty \log(D+1) + \|f_0\|_\infty) \frac{\log(3C_\chi/\gamma)}{n} + C_3 (\|f\|_\infty + 1) DR\left(n, \log\left\{\frac{3\beta \log n}{\gamma}\right\}\right) \right\}. \end{aligned}$$

Using that f belongs to $\mathcal{B}_s^{(l)}(P, M)$ and the fact that

$$R(n, \log \log n) \vee R\left(n, \log\left\{\frac{3\beta \log n}{\gamma}\right\}\right) \lesssim \log(n) \frac{\log \log n}{n},$$

where \lesssim states that the inequality holds up to some multiplicative constant independent of n , D and P , we deduce that we want to upper bound

$$\inf_{D \in \mathcal{D}_l} \left\{ P^2 D^{-2s} + D \log(n) \frac{\log \log n}{n} + \frac{\log \log n}{n} + \frac{\log(D+1)}{n} \right\}.$$

Since $\log(D+1) \leq D$ for all $D \in \mathcal{D}_l$, we only need to focus on

$$\inf_{D \in \mathcal{D}_l} \left\{ P^2 D^{-2s} + D \log(n) \frac{\log \log n}{n} \right\}.$$

$P^2 D^{-2s} < D \log(n) \frac{\log \log n}{n}$ if and only if $D > \left(\frac{P^4 n^2}{\log^2(n) (\log \log n)^2} \right)^{\frac{1}{4s+2}}$. Hence we define D_* by

$$\log_2(D_*) := \lfloor \log_2 \left(\left(\frac{P^4 n^2}{\log^2(n) (\log \log n)^2} \right)^{\frac{1}{4s+2}} \right) \rfloor + 1.$$

We consider three cases.

- If $D_* < 1$, then $P^2 D^{-2s} < D \log(n) \frac{\log \log n}{n}$ for any $D \in \mathcal{D}_l$ and by choosing $D_0 = 1$ to upper bound the infimum we get

$$\inf_{D \in \mathcal{D}_l} \left\{ \|f - \Pi_{S(l,D)}(f)\|_2^2 + W_{(l,D)}(\alpha) + V_{(l,D)}(\gamma) \right\} \leq \log(n) \frac{\log \log n}{n}.$$

- If $D_* > 2 \lfloor \log_2(n / (\log(n) \log \log n)^2) \rfloor$, then $P^2 D^{-2s} > D \log(n) \frac{\log \log n}{n}$ for any $D \in \mathcal{D}_l$ and by choosing $D_0 = 2^{\lfloor \log_2(\lfloor n / (\log(n) \log \log n)^2 \rfloor) \rfloor}$ to upper bound the infimum we get

$$\inf_{D \in \mathcal{D}_l} \left\{ \|f - \Pi_{S(l,D)}(f)\|_2^2 + W_{(l,D)}(\alpha) + V_{(l,D)}(\gamma) \right\} \lesssim 2P^2 D_0^{-2s} \leq 2^{2s+1} P^2 \left(\frac{(\log(n) \log \log n)^2}{n} \right)^{2s}.$$

- Otherwise D_* belongs to \mathcal{D}_l and we upper bound the infimum by choosing $D_0 = D_*$ and we get

$$\inf_{D \in \mathcal{D}_l} \left\{ \|f - \Pi_{S(l,D)}(f)\|_2^2 + W_{(l,D)}(\alpha) + V_{(l,D)}(\gamma) \right\} \lesssim 4P^{\frac{2}{2s+1}} \left(\frac{\log(n) \log \log n}{n} \right)^{\frac{2s}{2s+1}}.$$

The proof of Corollary 2 ends with simple computations that we provide below for the sake of completeness. Since

$$\begin{aligned} \log(n) \frac{\log \log n}{n} &\leq P^{\frac{2}{2s+1}} \left(\frac{\log(n) \log \log n}{n} \right)^{\frac{2s}{2s+1}} \\ &\Leftrightarrow \left(\log(n) \frac{\log \log n}{n} \right)^{1/2} \leq P. \end{aligned}$$

and since

$$\begin{aligned} &P^2 \left(\frac{(\log(n) \log \log n)^2}{n} \right)^{2s} \leq P^{\frac{2}{2s+1}} \left(\frac{\log(n) \log \log n}{n} \right)^{\frac{2s}{2s+1}} \\ \Leftrightarrow &P \left(\frac{(\log(n) \log \log n)^2}{n} \right)^s \leq P^{\frac{1}{2s+1}} \left(\frac{\log(n) \log \log n}{n} \right)^{\frac{s}{2s+1}} \\ \Leftrightarrow &P^{2s} \left(\frac{(\log(n) \log \log n)^2}{n} \right)^{s(2s+1)} \leq \left(\frac{\log(n) \log \log n}{n} \right)^s \\ \Leftrightarrow &P \left(\frac{(\log(n) \log \log n)^2}{n} \right)^{s+1/2} \leq \left(\frac{\log(n) \log \log n}{n} \right)^{1/2} \\ \Leftrightarrow &P \leq \frac{n^s}{(\log(n) \log \log n)^{2s+1/2}}, \end{aligned}$$

we deduce that if P is chosen such that

$$\left(\log(n) \frac{\log \log n}{n}\right)^{1/2} \leq P \leq \frac{n^s}{(\log(n) \log \log n)^{2s+1/2}}, \quad (36)$$

then the uniform separation rate of the test $\mathbb{1}_{T_\alpha > 0}$ over $\mathcal{B}_s^{(l)}(P, M)$ satisfies

$$\rho\left(\mathbb{1}_{T_\alpha > 0}, \mathcal{B}_s^{(l)}(P, M), \gamma\right) \leq C' P^{\frac{1}{2s+1}} \left(\frac{\log(n) \log \log n}{n}\right)^{\frac{s}{2s+1}}. \quad (37)$$

Remark This final statement can allow the reader to understand our choice for the size of the model $|\mathcal{M}|$ that we considered. Indeed, we chose for any $l \in \{1, 2, 3\}$, $\mathcal{D}_l = \{2^J, 0 \leq J \leq \log_2(n/(\log(n) \log \log n)^2)\}$ in order to ensure that for values of P saturating the right inequality in (36) (i.e. for $P \approx \frac{n^s}{(\log(n) \log \log n)^{2s+1/2}}$), the upper-bound in Eq.(37) still tends to zero as n goes to $+\infty$ for any possible values of the smoothness parameter s .

D. Concentration Lemmas for Markov chains

D.1 Hoeffding inequality for uniformly ergodic Markov chains

Proposition 7 is an Hoeffding bound for uniformly ergodic Markov chains.

Proposition 7 *Let $(X_i)_{i \geq 1}$ be a Markov chain on E uniformly ergodic (namely satisfying Assumption 1) with stationary distribution π and let us consider some function $f : E \rightarrow \mathbb{R}$ such that $\mathbb{E}_{X \sim \pi}[f(X)] = 0$. Then it holds for any $t \geq 0$*

$$\mathbb{P}\left(\left|\sum_{i=1}^n f(X_i)\right| \geq t\right) \leq 16 \exp\left(-\frac{1}{K(m, \tau)} \frac{t^2}{n \|f\|_\infty^2}\right),$$

where $K(m, \tau) = 2Km^2\tau^2$ for some universal constant $K > 0$. We refer to Assumption 1 and the following remark (or to (Duchemin et al., 2022, Section 2)) for the definitions of the constants m and τ .

Proof of Proposition 7. Let us first recall that under Assumption 1, the 1-Orlicz norm of the regeneration times of the split chain are bounded by some finite constant $\tau > 0$ (see the remark after Assumption 1). In this proof, we will use the notations introduced in (Duchemin et al., 2022, Section 2.3). Since the chain $(\tilde{X}_n)_n$ is distributed as $(X_i)_{i \geq 1}$, we will identify $(\tilde{X}_i)_{i \geq 1}$ and $(X_i)_{i \geq 1}$ in the proof.

Let us consider $N = \sup\{i \in \mathbb{N} : mS_{i+1} + m - 1 \leq n\}$. Then,

$$\left|\sum_{i=1}^n f(X_i)\right| = \left|\sum_{l=0}^N Z_l + \sum_{i=m(S_N+1)}^n f(X_i)\right| \leq \left|\sum_{l=0}^{\lfloor N/2 \rfloor} Z_{2l}\right| + \left|\sum_{l=0}^{\lfloor (N-1)/2 \rfloor} Z_{2l+1}\right| + \left|\sum_{i=m(S_N+1)}^n f(X_i)\right|. \quad (38)$$

We have $|\sum_{i=m(S_{N+1})}^n f(X_i)| \leq AmT_{N+1}$. So using the definition of the Orlicz norm and the fact that the random variables $(T_i)_{i \geq 2}$ are i.i.d., it holds for any $t \geq 0$,

$$\mathbb{P}\left(\left|\sum_{i=m(S_{N+1})}^n f(X_i)\right| \geq t\right) \leq \mathbb{P}(T_{N+1} \geq \frac{t}{Am}) \leq \mathbb{P}(\max(T_1, T_2) \geq \frac{t}{Am}) \leq 4 \exp\left(-\frac{t}{Am\tau}\right).$$

In order to control the first two terms in (38), we need to describe the tail behaviour of the random variable N with Lemma 8.

Lemma 8 (cf. Adamczak, 2008, Lemma 5)

We denote $R = \lfloor 3n/(\mathbb{E}T_2) \rfloor$. If $\|T_1\|_{\psi_1}, \|T_2\|_{\psi_1} \leq \tau$, then $\mathbb{P}(N > R) \leq 2 \exp\left(-\frac{n\mathbb{E}T_2}{8\tau^2}\right)$.

The random variable Z_{2l} is $\sigma(X_{m(S_{2l+1})}, \dots, X_{m(S_{2l+1}+1)-1})$ -measurable. Hence the random variables $(Z_{2l})_l$ are independent (see Duchemin et al., 2022, Section 2.3). Moreover, one has that for any l , $\mathbb{E}[Z_{2l}] = 0$. This is due to (Meyn and Tweedie, 1993, Eq.(17.23) Theorem 17.3.1) together with the assumption that $\mathbb{E}_{X \sim \pi}[f(X)] = 0$. Let us finally notice for any $l \geq 0$, $|Z_{2l}| \leq AmT_{2l+1}$, so $\|Z_{2l}\|_{\psi_1} \leq Am \max(\|T_1\|_{\psi_1}, \|T_2\|_{\psi_1}) \leq Am\tau$. One can similarly get that $(Z_{2l+1})_l$ are independent with $\mathbb{E}[Z_{2l+1}] = 0$ and $\|Z_{2l+1}\|_{\psi_1} \leq Am\tau$ for all $l \in \mathbb{N}$. Using these facts we have for any $t \geq 0$,

$$\begin{aligned} & \mathbb{P}\left(\left|\sum_{l=0}^{\lfloor N/2 \rfloor} Z_{2l}\right| + \left|\sum_{l=0}^{\lfloor (N-1)/2 \rfloor} Z_{2l+1}\right| \geq t\right) \\ & \leq \mathbb{P}\left(\left|\sum_{l=0}^{\lfloor N/2 \rfloor} Z_{2l}\right| + \left|\sum_{l=0}^{\lfloor (N-1)/2 \rfloor} Z_{2l+1}\right| \geq t, N \leq R\right) + 2 \exp\left(-\frac{n\mathbb{E}T_2}{8\tau^2}\right) \\ & \leq \mathbb{P}\left(\max_{0 \leq s \leq \lfloor R/2 \rfloor} \left|\sum_{l=0}^s Z_{2l}\right| \geq t/2\right) + \mathbb{P}\left(\max_{0 \leq s \leq \lfloor (R-1)/2 \rfloor} \left|\sum_{l=0}^s Z_{2l+1}\right| \geq t/2\right) + 2 \exp\left(-\frac{n\mathbb{E}T_2}{8\tau^2}\right) \\ & \leq 3\mathbb{P}\left(\left|\sum_{l=0}^{\lfloor R/2 \rfloor} Z_{2l}\right| \geq t/6\right) + 3\mathbb{P}\left(\left|\sum_{l=0}^{\lfloor (R-1)/2 \rfloor} Z_{2l+1}\right| \geq t/6\right) + 2 \exp\left(-\frac{n\mathbb{E}T_2}{8\tau^2}\right) \quad (\text{Using Lemma 10}) \\ & \leq 12 \exp\left(-\frac{1}{8} \min\left(\frac{t^2}{36RA^2m^2\tau^2}, \frac{t}{6Am\tau}\right)\right) + 2 \exp\left(-\frac{n\mathbb{E}T_2}{8\tau^2}\right), \end{aligned}$$

where we used Lemma 9 in the last inequality.

Lemma 9 (Bernstein's ψ_1 inequality, (Van Der Vaart and Wellner, 2013, Lemma 2.2.11) and the subsequent remark).

If Y_1, \dots, Y_n are independent random variables such that $\mathbb{E}Y_i = 0$ and $\|Y_i\|_{\psi_1} \leq \tau$, then for every $t > 0$,

$$\mathbb{P}\left(\left|\sum_{i=1}^n Y_i\right| > t\right) \leq 2 \exp\left(-\frac{1}{K} \min\left(\frac{t^2}{n\tau^2}, \frac{t}{\tau}\right)\right),$$

for some universal constant $K > 0$ ($K = 8$ fits).

Lemma 10 (cf. Kwapien and Woyczyński, 1992, Proposition 1.1.1) *If X_1, X_2, \dots are independent Banach space valued random variables (not necessarily identically distributed), and if $S_k = \sum_{i=1}^k X_i$, then*

$$\mathbb{P} \left(\max_{1 \leq j \leq k} \|S_j\| > t \right) \leq 3 \max_{1 \leq j \leq k} \mathbb{P} (\|S_j\| > t/3).$$

Gathering the previous results, we obtain that for any $t \geq 0$

$$\begin{aligned} \mathbb{P} \left(\left| \sum_{i=1}^n f(X_i) \right| \geq t \right) &\leq 12 \exp \left(-\frac{1}{8} \min \left(\frac{t^2 (\mathbb{E}T_2)}{36 \times 12 \times nA^2 m^2 \tau^2}, \frac{t}{12Am\tau} \right) \right) \\ &\quad + 2 \exp \left(-\frac{n\mathbb{E}T_2}{8\tau^2} \right) + 4 \exp \left(-\frac{t}{2Am\tau} \right). \end{aligned}$$

Since the left hand side of the previous inequality is zero for $t \geq nA$, and since $m \geq 1$, we obtain Proposition 7.

D.2 Bernstein's inequality for non-stationary Markov chains

Proposition 11 is an extension of the Bernstein type concentration inequality from Jiang et al. (2018) to non-stationary Markov chains. A proof can be found in the Appendix of Duchemin et al. (2022). Let us highlight that Proposition 11 is only used in the proofs of the main results from Section 5. One could have used other concentration results such as the one from Paulin (2015) (by using jointly Theorem 3.4 and Proposition 3.10) which would give strictly analogous results.

Proposition 11 *Suppose that the sequence $(X_i)_{i \geq 1}$ is a Markov chain satisfying Assumptions 1 and 5 with stationary distribution π and with an absolute spectral gap $1 - \lambda > 0$. Let us consider some $n \in \mathbb{N} \setminus \{0\}$ and bounded real valued functions $(f_i)_{1 \leq i \leq n}$ such that for any $i \in \{1, \dots, n\}$, $\int f_i(x) d\pi(x) = 0$ and $\|f_i\|_\infty \leq c$ for some $c > 0$. Let $\sigma^2 = \sum_{i=1}^n \int f_i^2(x) d\pi(x)/n$. Then for any $\varepsilon \geq 0$ it holds*

$$\mathbb{P} \left(\sum_{i=1}^n f_i(X_i) \geq \varepsilon \right) \leq \left\| \frac{d\chi}{d\pi} \right\|_{\pi,p} \exp \left(-\frac{\varepsilon^2/(2q)}{A_2\sigma^2 + A_1c\varepsilon} \right),$$

where $A_2 := \frac{1+\lambda}{1-\lambda}$ and $A_1 := \frac{1}{3}\mathbf{1}_{\lambda=0} + \frac{5}{1-\lambda}\mathbf{1}_{\lambda>0}$. q is the constant introduced in Assumption 5. Stated otherwise, for any $u > 0$ it holds

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n f_i(X_i) > \frac{2quA_1c}{n} + \sqrt{\frac{2quA_2\sigma^2}{n}} \right) \leq \left\| \frac{d\chi}{d\pi} \right\|_{\pi,p} e^{-u}.$$

References

- Radosław Adamczak. A tail inequality for suprema of unbounded empirical processes with applications to Markov chains. *Electron. J. Probab.*, 13(34):1000–1034, 2008. doi: 10.1214/EJP.v13-521. URL <https://doi.org/10.1214/EJP.v13-521>.
- Radosław Adamczak and Witold Bednorz. Some remarks on MCMC estimation of spectra of integral operators. *Bernoulli*, 21(4):2073—2092, Nov 2015a. ISSN 1350-7265. doi: 10.3150/14-bej635. URL <http://dx.doi.org/10.3150/14-BEJ635>.
- Radosław Adamczak and Witold Bednorz. Exponential concentration inequalities for additive functionals of Markov chains. *ESAIM: Probability and Statistics*, 19:440–481, 2015b. ISSN 1262-3318. doi: 10.1051/ps/2014032. URL <http://dx.doi.org/10.1051/PS/2014032>.
- Alekh Agarwal and John C Duchi. The generalization ability of online algorithms for dependent data. *IEEE Transactions on Information Theory*, 59(1):573–587, 2012.
- Jushan Bai. Testing parametric conditional distributions of dynamic models. *The Review of Economics and Statistics*, 85(3):531–549, 2003.
- Ludwig Baringhaus and Norbert Henze. A consistent test for multivariate normality based on the empirical characteristic function. *Metrika*, 35(1):339–348, 1988.
- Jan Beirlant, László Györfi, and Gábor Lugosi. On the asymptotic normality of the L1- and L2-errors in histogram density estimation. *Canadian Journal of Statistics*, 22:309 – 318, 12 2008.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Found. Trends Mach. Learn.*, 5(1):1–122, 2012. doi: 10.1561/22000000024. URL <https://doi.org/10.1561/22000000024>.
- Cristina Butucea et al. Goodness-of-fit testing and quadratic functional estimation from indirect observations. *The Annals of Statistics*, 35(5):1907–1930, 2007.
- Saptarshi Chakraborty and Kshitij Khare. Consistent estimation of the spectrum of trace class data augmentation algorithms. *Bernoulli*, 25(4B):3832–3863, 2019.
- Xiaming Chen and Yunwen Lei. Refined bounds for online pairwise learning algorithms. *Neurocomputing*, 275:2656–2665, 2018.
- Andreas Christmann and Ingo Steinwart. Support Vector Machines. *Support Vector Machines: Information Science and Statistics.*, 01 2008. doi: 10.1007/978-0-387-77242-4.
- Kacper Chwiałkowski, Heiko Strathmann, and Arthur Gretton. A kernel test of goodness of fit. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML’16, page 2606–2615. JMLR.org, 2016.
- Gabriela Ciolek and Patrice Bertail. New Bernstein and Hoeffding type inequalities for regenerative Markov chains. *Latin American journal of probability and mathematical statistics*, 16:1–19, 02 2019. doi: 10.30757/ALEA.v16-09.

- Stephan Cléménçon, Patrice Bertail, and Gabriela Cioł ek. Statistical learning based on Markovian data maximal deviation inequalities and learning rates. *Ann. Math. Artif. Intell.*, 88(7):735–757, 2020. ISSN 1012-2443. doi: 10.1007/s10472-019-09670-6. URL <https://doi.org/10.1007/s10472-019-09670-6>.
- Stéphan Cléménçon, Gábor Lugosi, and Nicolas Vayatis. Ranking and Empirical Minimization of U-statistics. *Ann. Statist.*, 36(2):844–874, 04 2008. doi: 10.1214/009052607000000910. URL <https://doi.org/10.1214/009052607000000910>.
- John B Conway. *A course in functional analysis*, volume 96. Springer, 2019.
- Felipe Cucker and Ding Xuan Zhou. *Learning theory: an approximation theory viewpoint*, volume 24. Cambridge University Press, 2007.
- Yohann De Castro, Claire Lacour, and Thanh Mai Pham Ngoc. Adaptive estimation of nonparametric geometric graphs. *Math. Stat. Learn.*, 2(3):217–274, 2019. ISSN 2520-2316.
- R.A. DeVore and G.G. Lorentz. *Constructive Approximation*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 1993. ISBN 9783540506270. URL https://books.google.fr/books?id=cDqNW6k7_ZwC.
- Quentin Duchemin and Yohann De Castro. Markov random geometric graph, MRGG: A growth model for temporal dynamic networks. *Electron. J. Stat.*, 16(1):671–699, 2022. doi: 10.1214/21-ejs1969. URL <https://doi.org/10.1214/21-ejs1969>.
- Quentin Duchemin, Yohann De Castro, and Claire Lacour. Concentration inequality for U-statistics of order two for uniformly ergodic Markov chains. *Bernoulli*, 2022. URL <https://hal.archives-ouvertes.fr/hal-03014763>.
- Alain Durmus, Eric Moulines, Alexey Naumov, and Sergey Samsonov. Probability and moment inequalities for additive functionals of geometrically ergodic Markov chains. *arXiv preprint arXiv:2109.00331*, 2021.
- Jianqing Fan, Bai Jiang, and Qiang Sun. Hoeffding’s inequality for general Markov chains and its applications to statistical learning. *J. Mach. Learn. Res.*, 22(139):1–35, 2021. ISSN 1532-4435.
- Yanqin Fan. Goodness-of-fit tests for a multivariate distribution by the empirical characteristic function. *Journal of Multivariate Analysis*, 62(1):36 – 63, 1997. ISSN 0047-259X. doi: <https://doi.org/10.1006/jmva.1997.1672>. URL <http://www.sciencedirect.com/science/article/pii/S0047259X97916729>.
- Yanqin Fan and Aman Ullah. On goodness-of-fit tests for weakly dependent processes using kernel method. *Journal of Nonparametric Statistics*, 11(1-3):337–360, 1999.
- Tamara Fernández and Arthur Gretton. A maximum-mean-discrepancy goodness-of-fit test for censored data. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2966–2975, 2019.

- Magalie Fromont and Béatrice Laurent. Adaptive goodness-of-fit tests in a density model. *Ann. Statist.*, 34(2):680–720, 2006. ISSN 0090-5364. doi: 10.1214/009053606000000119. URL <https://doi.org/10.1214/009053606000000119>.
- Evarist Giné and Richard Nickl. *Mathematical foundations of infinite-dimensional statistical models*. Cambridge Series in Statistical and Probabilistic Mathematics, [40]. Cambridge University Press, New York, 2016. ISBN 978-1-107-04316-9. doi: 10.1017/CBO9781107337862. URL <https://doi.org/10.1017/CBO9781107337862>.
- Jackson Gorham and Lester Mackey. Measuring sample quality with kernels. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, page 1292–1301. JMLR.org, 2017.
- Zheng-Chu Guo, Yiming Ying, and Ding-Xuan Zhou. Online regularized learning with pairwise loss functions. *Advances in Computational Mathematics*, 43(1):127–150, 2017.
- Steven C. H. Hoi, Doyen Sahoo, Jing Lu, and Peilin Zhao. Online learning: A comprehensive survey. *Neurocomputing*, 459:249–289, 2021.
- Yu. I. Ingster. Asymptotically minimax hypothesis testing for nonparametric alternatives. I. *Math. Methods Statist.*, 2(2):85–114, 1993. ISSN 1066-5307.
- Bai Jiang, Qiang Sun, and Jianqing Fan. Bernstein’s inequality for general Markov chains. *arXiv preprint arXiv:1805.10721*, 2018.
- Rong Jin, Shijun Wang, and Yang Zhou. Regularized Distance Metric Learning: Theory and Algorithm. In *NIPS*, volume 22, pages 862–870. Citeseer, 2009.
- Purushottam Kar, Bharath K. Sriperumbudur, Prateek Jain, and Harish C. Karnick. On the Generalization Ability of Online Learning Algorithms for Pairwise Loss Functions. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML’13, page III–441–III–449. JMLR.org, 2013.
- Vladimir Koltchinskii and Evarist Giné. Random Matrix Approximation of Spectra of Integral Operators. *Bernoulli*, 6, 02 2000. doi: 10.2307/3318636.
- Stanisław Kwapien and Wojbor A. Woyczyński. *Random series and stochastic integrals: single and multiple*. Probability and its Applications. Birkhäuser Boston, Inc., Boston, MA, 1992. ISBN 0-8176-3572-6. doi: 10.1007/978-1-4612-0425-1. URL <https://doi.org/10.1007/978-1-4612-0425-1>.
- Matthieu Lerasle, Nelo Molter Magalhães, and Patricia Reynaud-Bouret. Optimal kernel selection for density estimation. In *High dimensional probability VII*, volume 71 of *Progr. Probab.*, pages 425–460. Springer, [Cham], 2016. doi: 10.1007/978-3-319-40519-3_19. URL https://doi.org/10.1007/978-3-319-40519-3_19.
- Fuchun Li and Greg Tkacz. A Consistent Bootstrap Test for Conditional Density Functions with Time-Dependent Data. Staff working papers, Bank of Canada, 2001.

- Qiang Liu, Jason Lee, and Michael Jordan. A kernelized stein discrepancy for goodness-of-fit tests. In *International conference on machine learning*, pages 276–284, 2016.
- László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Soc., 2012.
- Pascal Massart. Some applications of concentration inequalities to statistics. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 9, pages 245–303, 2000.
- Pascal Massart. *Concentration inequalities and model selection*, volume 1896 of *Lecture Notes in Mathematics*. Springer, 2007. ISBN 978-3-540-48497-4; 3-540-48497-3.
- Ron Meir and Tong Zhang. Generalization error bounds for Bayesian mixture algorithms. *Journal of Machine Learning Research*, 4(Oct):839–860, 2003.
- K. L. Mengersen and R. L. Tweedie. Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Statist.*, 24(1):101–121, 02 1996. doi: 10.1214/aos/1033066201. URL <https://doi.org/10.1214/aos/1033066201>.
- S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Communications and Control Engineering Series. Springer-Verlag London, Ltd., London, 1993. ISBN 3-540-19832-6. doi: 10.1007/978-1-4471-3267-7. URL <https://doi.org/10.1007/978-1-4471-3267-7>.
- Claus Müller. *Analysis of spherical symmetries in Euclidean spaces*, volume 129. Springer Science & Business Media, 2012.
- Daniel Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electron. J. Probab.*, 20:no. 79, 32, 2015. doi: 10.1214/EJP.v20-4039. URL <https://doi.org/10.1214/EJP.v20-4039>.
- Yimo Qin, Bin Zou, Jingjing Zeng, Zhifei Sheng, and Lei Yin. Online regularized pairwise learning with non-iid observations. *International Journal of Wavelets, Multiresolution and Information Processing*, page 2150041, 2021.
- Lorenzo Rosasco, Mikhail Belkin, and Ernesto De Vito. On learning with integral operators. *Journal of Machine Learning Research*, 11(2), 2010.
- Rimantas Rudzakis and Aleksej Bakshaev. Goodness of fit tests based on kernel density estimators. *Informatika*, 24(3):447–460, 2013.
- Yandi Shen, Fang Han, and Daniela Witten. Exponential inequalities for dependent V-statistics via random Fourier features. *Electron. J. Probab.*, 25:1–18, 2020. doi: 10.1214/20-ejp411. URL <https://doi.org/10.1214/20-ejp411>.
- Steve Smale and Ding-Xuan Zhou. Online learning with Markov sampling. *Analysis and Applications*, 7(01):87–113, 2009.
- Ingo Steinwart, Don Hush, and Clint Scovel. Learning from dependent observations. *Journal of Multivariate Analysis*, 100(1):175–194, 2009.

- A. Van Der Vaart and J. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series in Statistics. Springer New York, 2013. ISBN 9781475725452.
- Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Yuyang Wang, Roni Khardon, Dmitry Pechyony, and Rosie Jones. Generalization Bounds for Online Learning Algorithms with Pairwise Loss Functions. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, pages 13.1–13.22, Edinburgh, Scotland, 2012. PMLR. URL <https://proceedings.mlr.press/v23/wang12.html>.
- Eric Xing, Michael Jordan, Stuart J Russell, and Andrew Ng. Distance metric learning with application to clustering with side-information. *Advances in neural information processing systems*, 15:521–528, 2002.
- Jie Xu, Yuan Yan Tang, Bin Zou, Zongben Xu, Luoqing Li, and Yang Lu. The generalization ability of online SVM classification based on Markov sampling. *IEEE transactions on neural networks and learning systems*, 26(3):628–639, 2014.
- Yiming Ying and Ding-Xuan Zhou. Unregularized online learning algorithms with general loss functions. *Applied and Computational Harmonic Analysis*, 42(2):224–244, 2017.
- Jingjing Zeng, Bin Zou, Yimo Qin, Qian Chen, Jie Xu, Lei Yin, and Hongwei Jiang. Generalization ability of online pairwise support vector machine. *Journal of Mathematical Analysis and Applications*, 497(2):124914, 2021.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.
- Tong Zhang. Data dependent concentration bounds for sequential prediction algorithms. In *International Conference on Computational Learning Theory*, pages 173–187. Springer, 2005.
- Peilin Zhao, Steven C. H. Hoi, Rong Jin, and Tianbao Yang. Online AUC maximization. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML’11, page 233–240, Madison, WI, USA, 2011. Omnipress. ISBN 9781450306195.
- Bin Zou, Hai Zhang, and Zongben Xu. Learning from uniformly ergodic Markov chains. *Journal of Complexity*, 25(2):188–200, 2009.