# Variance-Aware Estimation of Kernel Mean Embedding

### **Geoffrey Wolfer**

Center for Data Science Waseda University 1-6-1 Nishiwaseda, Shinjuku-ku Tokyo 169-8050, Japan

#### Pierre Alquier

GEO.WOLFER@AONI.WASEDA.JP

ALQUIER@ESSEC.EDU

ESSEC Business School Asia-Pacific campus 5 Nepal Park 575749 Singapore

Editor: Bharath Sriperumbudur

### Abstract

An important feature of kernel mean embeddings (KME) is that the rate of convergence of the empirical KME to the true distribution KME can be bounded independently of the dimension of the space, properties of the distribution and smoothness features of the kernel. We show how to speed-up convergence by leveraging variance information in the reproducing kernel Hilbert space. Furthermore, we show that even when such information is a priori unknown, we can efficiently estimate it from the data, recovering the desiderata of a distribution agnostic bound that enjoys acceleration in fortuitous settings. We further extend our results from independent data to stationary mixing sequences and illustrate our methods in the context of hypothesis testing and robust parametric estimation.

**Keywords:** kernel mean embedding, maximum mean discrepancy, distribution learning, empirical Bernstein

## 1. Introduction

Estimating a probability distribution  $\mathbb{P}$  over a space  $\mathcal{X}$  from n iid samples  $X_1, \ldots, X_n \sim \mathbb{P}$  is a central problem in computer science and statistics (Kearns et al., 1994; Tsybakov, 2008). To formalize the question, one selects a distance (or at least a contrast function) between distributions and oftentimes introduces assumptions on the underlying probability space (for instance finitely supported, probability function is absolutely continuous with respect to the Lebesgue measure, Hölder continuous, ...). Increasing the stringency of assumptions generally leads to substantially faster minimax rates. For instance, in the finite support  $|\mathcal{X}| < \infty$  case and with respect to total variation, whilst the expectation risk evolves roughly as  $\sqrt{|\mathcal{X}|/n}$ , it is known that this rate can be sharpened by replacing  $|\mathcal{X}|$  with a bound on the "half-norm"  $^1 \|\mathbb{P}\|_{1/2} \doteq (\sum_{x \in \mathcal{X}} \sqrt{\mathbb{P}(x)})^2$  (Berend and Kontorovich, 2013), when defined, and which corresponds to some measure of entropy of the underlying distribution. Furthermore, even without prior knowledge of  $\|\mathbb{P}\|_{1/2}$ , one can construct confidence intervals around  $\mathbb{P}$  that

<sup>1.</sup> The half-norm is not a norm in the proper sense as it violates the triangle inequality.

<sup>©2025</sup> Geoffrey Wolfer and Pierre Alquier.

License: CC-BY 4.0, see https://creativecommons.org/licenses/by/4.0/. Attribution requirements are provided at http://jmlr.org/papers/v26/23-0161.html.

depend on  $\|\widehat{\mathbb{P}}_n\|_{1/2}$ —the half-norm of the empirical distribution  $\widehat{\mathbb{P}}_n(x) = \frac{1}{n} \sum_{t=1}^n \delta[X_t = x]$ (Cohen et al., 2020, Theorem 2.1). Namely, when  $\mathbb{P}$  is supported on  $\mathbb{N}$ , with probability at least  $1 - \delta$ , it holds that

$$\left\|\widehat{\mathbb{P}}_{n} - \mathbb{P}\right\|_{\mathsf{TV}} \leq \frac{1}{\sqrt{n}} \left( \left\|\widehat{\mathbb{P}}_{n}\right\|_{1/2}^{1/2} + 3\sqrt{\frac{1}{2}\log(2/\delta)} \right).$$
(1)

The advantage of the above expression is twofold (a) we do not need make assumptions on  $\|\mathbb{P}\|_{1/2}$ , which could be non-convergent<sup>2</sup>, and (b) in favorable cases where  $\|\widehat{\mathbb{P}}_n\|_{1/2}$  is small, the intervals will be narrower. In this paper, we set out to explore the question of whether analogues of (1) are possible for general probability spaces and with respect to maximum mean discrepancy.

### 1.1 Notation and Background

The set of integers up to  $n \in \mathbb{N}$  is denoted by  $[n] \doteq \{1, 2, \ldots\}$ . Let  $\mathcal{X}$  be a separable topological space, and  $\mathcal{P}(\mathcal{X})$  the set of all Borel probability measures over  $\mathcal{X}$ . For a bounded function  $f: \mathcal{X} \to \mathbb{R}$ , we define for convenience

$$\overline{f} \doteq \sup_{x \in \mathcal{X}} f(x) \qquad \underline{f} \doteq \inf_{x \in \mathcal{X}} f(x), \qquad \Delta f \doteq \overline{f} - \underline{f}.$$
(2)

Let  $k: \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  be a continuous positive definite kernel and  $\mathcal{H}_k$  be the associated reproducing kernel Hilbert space (RKHS) (Berlinet and Thomas-Agnan, 2011). We assume that the kernel is bounded<sup>3</sup> in the sense that  $\sup_{x \in \mathcal{X}} k(x, x) < \infty$ . Letting  $\mathbb{P} \in \mathcal{P}(\mathcal{X})$ , the kernel mean embedding (KME) of  $\mathbb{P}$  is defined as

$$\mu_{\mathbb{P}} \doteq \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X, \cdot) \right] = \int_{\mathcal{X}} k(x, \cdot) d\mathbb{P}(x) \in \mathcal{H}_k,$$

which is interpreted as a Bochner integral (Diestel and Uhl, 1977, Chapter 2). Let  $n \in \mathbb{N}$ and  $X_1, \ldots, X_n$  be a sequence of observations sampled independently<sup>4</sup> from  $\mathbb{P}$ . We write

$$\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) = \widehat{\mu}_{\mathbb{P}} \doteq \frac{1}{n} \sum_{t=1}^n k(X_t,\cdot) \in \mathcal{H}_k,$$
(3)

for the KME of the empirical distribution  $\widehat{\mathbb{P}}_n \doteq \frac{1}{n} \sum_{t=1}^n \delta_{X_t}$ , and call the distance  $\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k}$ the maximum mean discrepancy (MMD) between the true mean embedding and its empirical estimator. A kernel is called characteristic if the map  $\mu \colon \mathbb{P} \to \mu_{\mathbb{P}}$  is injective. This property ensures that  $\|\mu_{\mathbb{P}} - \mu_{\mathbb{P}'}\|_{\mathcal{H}_k} = 0$  if and only if  $\mathbb{P} = \mathbb{P}'$ . A kernel is called translation invariant<sup>5</sup> (TI) when there exists a positive definite function  $\psi \colon \mathcal{X} \to \mathbb{R}$  such that for any  $x, x' \in \mathcal{X},$ 

$$k(x, x') = \psi(x' - x).$$
 (4)

<sup>2.</sup> As opposed to the empirical proxy  $\|\widehat{\mathbb{P}}_n\|_{1/2}$ , which is always a finite quantity. 3. Note that the supremum of the kernel is always reached on the diagonal, that is,  $\sup_{(x,x')\in\mathcal{X}^2} k(x,x') =$  $\sup_{x \in \mathcal{X}} k(x, x).$ 

<sup>4.</sup> Unless otherwise specified, the data will be considered iid. We will address the case of dependent data in Section 5.

<sup>5.</sup> Also sometimes referred to as an *anisotropic stationary* kernel.

In particular,  $\overline{\psi} = \psi(0) = \overline{k}$ . When  $\mathcal{X} = \mathbb{R}^d$ , a kernel k is said to be a radial basis function (RBF) when for any  $x, x' \in \mathcal{X}, k(x, x') = \phi(||x' - x||_2)$  for some function  $\phi \colon \mathbb{R}_+ \to \mathbb{R}$ . Noticeably, k being positive definite does not preclude it from taking negative values. However, when k is RBF, the following lower bound on  $\phi$  holds (see for instance Genton 2001; Stein 1999),

$$\underline{\phi} \ge \overline{\phi} \inf_{t \ge 0} \left\{ \left(\frac{2}{t}\right)^{(d-2)/2} \Gamma(d/2) J_{(d-2)/2}(t) \right\},$$

where  $\Gamma$  is the Gamma function and  $J_{\beta}$  is the Bessel function of the first kind of order  $\beta$ , showing that  $|\phi|$  becomes evanescent as the dimension increases.

#### 1.2 Related Work

From an asymptotic standpoint, the weak law of large numbers asserts that  $\hat{\mu}_{\mathbb{P}}$  converges to the true  $\mu_{\mathbb{P}}$ . Furthermore,

$$\sqrt{n} \left(\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\right)$$

converges in distribution to a zero mean Gaussian process on  $\mathcal{H}_k$  (Berlinet and Thomas-Agnan, 2011, Section 9.1). This work, however, is more concerned with the finite sample theory, and more specifically with the rate of convergence of  $\hat{\mu}_{\mathbb{P}}$  towards  $\mu_{\mathbb{P}}$  with respect to the RKHS norm. Conveniently and perhaps surprisingly, it is possible to derive a rate that depends neither on the smoothness of the considered kernel k, nor the properties of the true distribution. To obtain a distribution independent rate at  $\mathcal{O}_P(n^{-1/2})$ , a typical strategy (see for example Lopez-Paz et al. 2015, Section B.1) consists in first expressing the dual relationship between the norm in the RKHS and the uniform norm of an empirical process,

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{k}} = \sup_{\substack{f \in \mathcal{H}_{k} \\ \|f\|_{\mathcal{H}_{k}} \leq 1}} \langle f, \widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}} = \sup_{\substack{f \in \mathcal{H}_{k} \\ \|f\|_{\mathcal{H}_{k}} \leq 1}} \left(\frac{1}{n} \sum_{t=1}^{n} f(X_{t}) - \mathbb{E}f\right).$$
(5)

A classical symmetrization argument (Mohri et al., 2018) followed by an application of McDiarmid's inequality (McDiarmid et al., 1989) yield that with probability at least  $1 - \delta$ ,

$$\sup_{\substack{f \in \mathcal{H}_k \\ \|f\|_{\mathcal{H}_k} \le 1}} \left( \frac{1}{n} \sum_{t=1}^n f(X_t) - \mathbb{E}f \right) \le 2\mathfrak{R}_n + \sqrt{\frac{2\overline{k}\log(1/\delta)}{n}},\tag{6}$$

where  $\Re_n$  is the Rademacher complexity (Mohri et al., 2018, Definition 3.1) of the class of unit functions in the RKHS. The bound  $\Re_n^2 \leq \overline{k}/n$  (Bartlett and Mendelson, 2002), and an application of Jensen's inequality conclude the claim. With a more careful analysis, Tolstikhin et al. (2017, Proposition A.1) halved the constant of the first term in (6), hence showed that with probability at least  $1 - \delta$ ,

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{k}} \le \sqrt{\frac{\overline{k}}{n}} + \sqrt{\frac{2\overline{k}\log(1/\delta)}{n}}.$$
(7)

What is more, Tolstikhin et al. (2017, Theorems 1,6,8) provide corresponding lower bounds in  $\Omega_P(n^{-1/2})$ , showing that the embedding of the empirical measure achieves an unimprovable rate of convergence. Results similar to (7) (with worse constants) can be derived when the observations are not independent, see Chérief-Abdellatif and Alquier (2022, Lemma 7.2). We note that other estimators have been proposed in the literature for the case where more information is available about the underlying distribution  $\mathbb{P}$ . See for instance Muandet et al. (2014, 2016) who propose a shrinkage estimator inspired by the James-Stein estimator. In this work, we stress that our estimator will be the KME of the empirical distribution, defined in (3).

### **1.3 Main Contributions**

In this paragraph, we briefly summarize our findings. The tilde notation suppresses constants and logarithmic factors in  $1/\delta$ , and the reader is referred to Theorem 1 and Theorem 12 for the full expression of the second order terms. Our first contribution (Theorem 1) consists in deriving a bound on  $\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k}$  that can leverage additional information about the underlying distribution  $\mathbb{P}$  and the selected kernel k. Namely, we show that with probability at least  $1 - \delta$ ,

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \le \sqrt{2v_k(\mathbb{P})\frac{\log(2/\delta)}{n}} + \widetilde{\mathcal{O}}\left(\frac{1}{n}\right),$$

where  $v_k(\mathbb{P})$ —defined in (8)—corresponds to some notion of variance in the RKHS. Notably, it holds that  $v_k(\mathbb{P}) \leq \overline{k}$ , hence our upper bound is superior to and able to recover known bounds in most settings (Remark 2).

Our chief technical contribution (Theorem 12) is in establishing that, at least when k is translation invariant, the dependence of the above confidence interval can be made independent of the underlying distribution by replacing the variance by an empirical proxy. Specifically, with probability at least  $1 - \delta$ ,

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \le \sqrt{2\widehat{v}_k(X_1, \dots, X_n) \frac{\log(4/\delta)}{n}} + \widetilde{\mathcal{O}}\left(\frac{1}{n}\right)$$

where  $\hat{v}_k$ —defined in (12), (13)—is a quantity that only depends on the chosen kernel and the sample. We also extend the latter result to kernels without the translation invariance property in Theorem 29.

Expanding beyond the iid setting, we obtain convergence rates for stationary mixing sequences. For  $\phi$ -mixing processes, Theorem 17 establishes the following rate of convergence.

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{k}} \leq \sqrt{\frac{v_{k} + \Sigma_{n}}{n}} + 4\sqrt{\frac{2v_{k} \|\Gamma\|_{2} \log(1/\delta)}{n}} + \widetilde{\mathcal{O}}\left(\frac{1}{n}\right),$$

where  $\Sigma_n$  is a total measure of covariance between observations in the RKHS (refer to Lemma 16) and  $\|\Gamma\|_2$  is the spectral norm of a coupling matrix defined in (17). We additionally analyze the broader class of  $\beta$ -mixing processes in Theorem 18.

### 1.4 Outline

In Section 2, for the problem of estimating a distribution with respect to maximum mean discrepancy, we give a convergence rate with a dominant term in  $\mathcal{O}_P(n^{-1/2})$  involving a variance term v which depends on both the chosen kernel k and the underlying distribution  $\mathbb{P}$ . We then illustrate how the rate can subdue minimax lower bounds when v is favorable. In particular, we show that for large collections of kernels, and when  $\mathcal{X} \subset \mathbb{R}^d$ , the variance vcan be controlled by a quantity that decouples the influence of the kernel and a measure of total variance of  $\mathbb{P}$ . In Section 3, we proceed to show that even if v is unknown, it is possible to efficiently estimate it from the data with an "empirical Bernstein" approach whenever the kernel is translation invariant, or at least enjoys a mildly varying diagonal. In Section 4, we illustrate the benefits of the variance aware bounds on the classical two-sample test problem. In Section 5, we establish convergence rates for stationary  $\phi$  and  $\beta$  mixing sequences. In Section 6.2, we put our methods into practice, first in the context of hypothesis testing, and second by improving the results of Briol et al. (2019) and Chérief-Abdellatif and Alquier (2022) in the context of robust parametric maximum mean discrepancy estimation. All proofs are deferred to Section 7 for clarity of the exposition.

#### 2. Variance-Aware Convergence Rates

The central quantity we propose to consider is the following variance term in the RKHS,

$$v_k(\mathbb{P}) \doteq \mathbb{E}_{X \sim \mathbb{P}} \| k(X, \cdot) - \mu_{\mathbb{P}} \|_{\mathcal{H}_h}^2.$$
(8)

It is clear that  $v_k(\mathbb{P})$  depends both on the choice of kernel k and on the underlying distribution, and we will simply write  $v = v_k(\mathbb{P})$  to avoid encumbering notation. Simple calculations (refer to Lemma 16) show that

$$\mathbb{E} \|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k}^2 = \frac{v}{n},$$

where the expectation is taken with respect to an independent sample  $X_1, \ldots, X_n \sim \mathbb{P}$ , thus by applications of Jensen's and Chebyshev's inequalities, we readily obtain a rate on the expected risk in terms of v,

$$\mathbb{E} \|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \le \sqrt{\frac{v}{n}},$$

and a deviation bound

$$\mathbb{P}\left(\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} > (1+\tau)\sqrt{v/n}\right) \le 1/(1+\tau)^2, \tau \in (0,\infty).$$

One basic feature of RKHS is that new kernels can be constructed by linearly combining existing kernels—see for instance Gretton (2013). It is therefore noteworthy that v is linear with respect reproducing kernels. Reformulate

$$v = \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X, X) \right] - \mathbb{E}_{X, X' \sim \mathbb{P}} \left[ k(X, X') \right],$$

and suppose that  $k = \sum_{k_i \in \mathcal{K}} \alpha_i k_i$  where  $\mathcal{K} = \{k_1, \ldots, k_{|\mathcal{K}|}\}$  is a collection of reproducing kernels with  $\alpha \in \mathbb{R}^{|\mathcal{K}|}$ . It then holds that

$$v_k = \sum_{k_i \in \mathcal{K}} \alpha_i v_{k_i}.$$

Moving on to high-probability confidence bounds, an application of Bernstein's inequality in Hilbert spaces (Pinelis and Sakhanenko, 1985; Yurinsky, 1995) not only recovers the rate of convergence, as alluded to by Tolstikhin et al. (2017), but in fact yields a maximal inequality.

**Theorem 1 (Variance-aware confidence interval)** Let  $X_1, \ldots, X_n \stackrel{iid}{\sim} \mathbb{P}, k \colon \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  be a reproducing kernel, and let

$$v \doteq \mathbb{E}_{X \sim \mathbb{P}} \left\| k(X, \cdot) - \mu_{\mathbb{P}} \right\|_{\mathcal{H}_{k}}^{2}$$

With probability at least  $1 - \delta$ , it holds that

$$\max_{1 \le t \le n} \left\{ t \left\| \widehat{\mu}_{\mathbb{P}}(X_1, \dots, X_t) - \mu_{\mathbb{P}} \right\|_{\mathcal{H}_k} \right\} \le \sqrt{2vn \log(2/\delta)} + \frac{4}{3} \sqrt{\overline{k}} \log(2/\delta),$$

where  $\overline{k} = \sup_{x \in \mathcal{X}} k(x, x)$ . In particular, with probability at least  $1 - \delta$ , it holds that

$$\|\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n)-\mu_{\mathbb{P}}\|_{\mathcal{H}_k}\leq \mathcal{B}_{k,\delta}(\mathbb{P},n),$$

with

$$\mathcal{B}_{k,\delta}(\mathbb{P},n) = \mathcal{B}_{\delta} \doteq \sqrt{2v \frac{\log(2/\delta)}{n}} + \frac{4}{3}\sqrt{\overline{k}} \frac{\log(2/\delta)}{n}$$

**Remark 2** Since the reproducing kernel is bounded, it is always the case that

$$v \leq \mathbb{E}_{X \sim \mathbb{P}} \|k(X, \cdot)\|^2_{\mathcal{H}_k} \leq \mathbb{E}_{X \sim \mathbb{P}} k(X, X) \leq \overline{k},$$

where the first inequality can be found for example in (Chérief-Abdellatif and Alquier, 2022, Lemma 7.1, Proof). As a result, Theorem 1 strictly supersedes (7) at least when

$$n > \frac{16}{9} \left( \frac{\log(2/\delta)}{1 + \sqrt{2\log(1/\delta)} - \sqrt{2\log(2/\delta)}} \right)^2.$$

For instance, for  $\delta = 0.05$ , it suffices that  $n \ge 46$ .

**Remark 3** The Bernstein approach, in contrast to bounded differences, has the additional advantage of yielding a maximal inequality, further opening the door for early stopping methods.

**Remark 4 (Sharpness of the constants)** We note but do not pursue that there exist other families of concentration inequalities which are known to dominate Bennett–Bernstein-type inequalities. For instance, the class of inequalities pioneered by Bentkus et al. (2006), with empirical counterparts derived by Kuchibhotla and Zheng (2021), which is known to be

nearly optimal for sample averages of independent observations from a log-concave probability distribution. We also mention the family of inequalities obtained by Waudby-Smith and Ramdas (2024) based on martingale methods. However, our problem requires a bound for norms of sums of vectors in Hilbert spaces, or alternatively for the supremum of empirical processes, which we could not locate in the aforementioned references. Additionally, even in the simpler case of sample averages, computation of the bound requires effort, thus a concentration bound in the more complicated Hilbert space setting could be challenging to compute. Finally, Bernstein-type bounds have known extensions for the case of time-dependent data, enabling us to analyze the estimation problem from stationary mixing sequences of observations (refer to Section 5).

We immediately observe that:

- (O1) While the bound in (7) depends solely on chosen quantities and is computable without any knowledge of  $\mathbb{P}$ , this is not the case for  $\mathcal{B}_{\delta}$  in Theorem 1.
- (O2) Perhaps even more concerning, it is a priori unclear how v depends on properties of k and  $\mathbb{P}$ , thus making it difficult to convert assumptions on  $\mathbb{P}$  into a bound for v.

We defer (O1) to Section 3 and first address (O2) by pointing out that when  $\mathcal{X} = \mathbb{R}^d$ , and for numerous hyper-parametrized families of kernels, it is possible to promptly obtain upper bounds on v that decouple the influence of the hyper-parameter and some measure of spread of the underlying distribution.

#### 2.1 Gaussian Kernel

For a TI kernel—see (4)—we can rewrite v as

$$v = \overline{\psi} - \mathbb{E}_{X, X' \sim \mathbb{P}} \left[ \psi(X' - X) \right].$$

The Gaussian kernel with lengthscale parameter  $\gamma > 0$ , defined by

$$k_{\gamma}(x, x') = \exp\left(-\frac{\|x' - x\|_2^2}{2\gamma^2}\right),$$
(9)

is the prototypical example of a characteristic translation invariant kernel, and satisfies  $\overline{\psi} = 1$ . For  $x \in \mathbb{R}^d$ , let  $x^i$  denote the *i*th component of x. The function  $z \mapsto e^{-z}$  is convex, thus from Jensen's inequality

$$v_{\gamma} = \mathbb{E}_{X \sim \mathbb{P}} \left\| k_{\gamma}(X, \cdot) - \mu_{\mathbb{P}} \right\|_{\mathcal{H}_{k_{\gamma}}}^{2} \leq 1 - \exp\left(-\frac{\mathbb{E}_{X, X' \sim \mathbb{P}} \left\| X' - X \right\|_{2}^{2}}{2\gamma^{2}}\right)$$

We can further rewrite the expectation on the right side of this inequality as

$$\mathbb{E}_{X,X'\sim\mathbb{P}}\left\|X'-X\right\|_{2}^{2} = \sum_{i=1}^{d} \mathbb{E}_{X,X'\sim\mathbb{P}}\left[(X'^{i}-X^{i})^{2}\right] = \sum_{i=1}^{d} 2\mathbb{V}_{X\sim\mathbb{P}}X^{i} = 2\operatorname{Tr}\Sigma_{\mathbb{P}},$$

where  $\Sigma_{\mathbb{P}} \doteq \mathbb{V}_{X \sim \mathbb{P}} X = \mathbb{E}_{X \sim \mathbb{P}} \left[ (X - \mathbb{E}X)(X - \mathbb{E}X)^* \right]$  is the covariance matrix of  $\mathbb{P}$  and its trace is interpreted as a measure a total variance, agnostic to correlations between distinct components. As a result, for any fixed  $\gamma > 0$ , we obtain from Theorem 1 that with probability at least  $1 - \delta$ ,

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{k\gamma}} \le \sqrt{2\left(1 - e^{-\operatorname{Tr}\Sigma_{\mathbb{P}}/\gamma^2}\right)\frac{\log(2/\delta)}{n}} + \frac{4}{3}\frac{\log(2/\delta)}{n}.$$
 (10)

While  $\gamma$  has no influence over the right-hand side in (7), it is clear that for  $\gamma \to \infty$ , the rate of convergence will be accelerated in (10).

**Example 1 (Gaussian location model with known variance)** Assume that  $\mathbb{P} = \mathbb{P}_{\theta} = \mathcal{N}(\theta, \sigma^2 I_d)$ , the Gaussian distribution with unknown location parameter  $\theta \in \mathbb{R}^d$ , but with known covariance matrix  $\sigma^2 I_d$ . It holds that  $\operatorname{Tr} \Sigma_{\mathbb{P}_{\theta}} = \sigma^2 d$ .

- (i) Fixing  $\gamma$  and taking  $\sigma \to 0$ , the bound in (10) vanishes, unlike (7).
- (ii) Setting  $\gamma^2 = \lambda \sigma^2 d$  with  $\lambda > 0$ , we readily obtain the variance upper bound  $v \leq (1 e^{-1/\lambda}) \leq 1/\lambda$ , enabling uncomplicated tuning of the convergence rate by  $\lambda$ .

**Remark 5** Example 1 highlights that if we allow the lengthscale parameter to vary with the sample size n, Theorem 1 speeds up the convergence rate dramatically. For instance, setting  $\lambda_n = -1/\log(1-1/n)$ , we achieve a rate in  $\mathcal{O}_P(n^{-1})$ , subverting the lower bounds of Tolstikhin et al. (2017). In fact, we can reach any prescribed rate of convergence between  $\mathcal{O}_P(n^{-1/2})$  and  $\mathcal{O}_P(n^{-1})$  for a suitable choice of  $\lambda_n$ . This stands in sharp contrast with bounds obtained through a bounded differences approach such as (7). However, it is important to remember that a larger lengthscale parameter will flatten the kernel, hence make the left hand side in (10) less informative. Depending on the application, it is possible to achieve the optimal balance between these two effects; some examples are discussed in Section 6.

### 2.2 Convex Radial Square Basis Functions

In fact, a technique similar to that used for obtaining (10) yields a bound on the variance for any (convex) radial kernel.

**Lemma 6** Assume that for any  $x, x' \in \mathcal{X}$ ,  $k(x, x') = r(-\|x' - x\|_2^2)$  for some convex function  $r: \mathbb{R}_+ \to \mathbb{R}$ . Then

$$v \le \overline{r} - r \left( -\operatorname{Tr} \Sigma_{\mathbb{P}} \right),$$

where  $\Sigma_{\mathbb{P}}$  is the covariance matrix that pertains to  $\mathbb{P}$ .

## 2.3 Positive Definitive Matrix on the Finite Space

Consider  $|\mathcal{X}| < \infty$  and a symmetric positive definite matrix K of size  $|\mathcal{X}|$ . Write  $K = \sum_{x \in \mathcal{X}} \lambda_x v_x^{\mathsf{T}} v_x$ , where for any  $x \in \mathcal{X}$ ,  $K v_x = \lambda_x v_x$ , and by positive definiteness,  $\lambda_x > 0$ . Then the feature map can be expressed as

$$K(x,\cdot) = \left(\sqrt{\lambda_{x'}}v_x(x')\right)_{x'\in\mathcal{X}},$$

and by direct calculation,

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{K}}^{2} = \sum_{x \in \mathcal{X}} \lambda_{x} \left(\frac{1}{n} \sum_{t=1}^{n} v_{x}(X_{t}) - \mathbb{E}[v_{x}(X)]\right)^{2},$$
$$v = \sum_{x \in \mathcal{X}} \lambda_{x} \mathbb{V}[v_{x}(X)].$$

In particular, for K = I,

$$\left\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\right\|_{\mathcal{H}_{K}} = \left\|\widehat{\mathbb{P}}_{n} - \mathbb{P}\right\|_{2}, \qquad v = 1 - \left\|\mathbb{P}\right\|_{2}^{2},$$

recovering that with probability at least  $1 - \delta$ ,

$$\left\|\widehat{\mathbb{P}}_n - \mathbb{P}\right\|_2 \le \sqrt{2(1 - \left\|\mathbb{P}\right\|_2^2)\frac{\log(2/\delta)}{n}} + \frac{4}{3}\frac{\log(2/\delta)}{n}$$

So far, we have shown how to obtain for a natural class of kernels an upper bound on v that decouples the choice of k from some measure of total variance of  $\mathbb{P}$ . Provided an upper bound on the latter is available (see Example 1), we obtain an improved convergence rate. However, we understandably may not have a bound on the variance of  $\mathbb{P}$ , or could have insight about its variance, but only have access to contaminated data (see later Section 6.2.1). Recovering a bound that still enjoys the discussed speed-up without a priori knowledge on  $\mathbb{P}$  is the problem we set out to solve in the next section.

### 3. Convergence Rates with Empirical Variance Proxy

We solve the issue (O1) of not knowing v by estimating it from the data, simultaneously to  $\mu_{\mathbb{P}}$ . A short heuristic analysis of the Epanechnikov function (which is not typically a kernel according to our definition) that we conduct in Section 3.1 hints at an "empirical Bernstein" (Audibert et al., 2007; Maurer and Pontil, 2009) approach—replacing the variance term by some empirical proxy—which we explore formally in Section 3.2. The argument is structured around the pivotal definition of a weakly self-bounding function that we first recall.

**Definition 7 (Maurer 2006; Boucheron et al. 2009)** Let  $n \in \mathbb{N}$ ,  $t_0 \in [n]$ ,

$$x = (x_1, \dots, x_{t_0}, \dots, x_n) \in \mathcal{X}^n$$

 $x'_{t_0} \in \mathcal{X}$ , and write

 $x^{(t_0)} = (x_1, \dots, x_{t_0-1}, x'_{t_0}, x_{t_0+1}, \dots, x_n),$ 

for the vector where  $x_{t_0}$  has been replaced with  $x_{t'_0}$ . Let  $f: \mathcal{X}^n \to \mathbb{R}$ .

(i) The function f is called weakly  $(\alpha, \beta)$ -self-bounding when for all  $x \in \mathcal{X}^n$ ,

$$\sum_{t_0=1}^n \left( f(x) - \inf_{x'_{t_0} \in \mathcal{X}} f\left(x^{(t_0)}\right) \right)^2 \le \alpha f(x) + \beta.$$

(ii) The function f is said to have the bounded differences property when for all  $x \in \mathcal{X}^n$ and all  $t_0 \in [n]$ ,

$$f(x) - \inf_{x'_{t_0} \in \mathcal{X}} f\left(x^{(t_0)}\right) \le 1.$$

# 3.1 Intuition in the Hypercube

Let us take  $\mathcal{X} = [0,1]^d$  the binary hypercube for  $d \gg 1$ , and consider the Epanechnikov (parabolic) function

$$q(x, x') = 1 - ||x' - x||_2^2 / d.$$

Note that q does not define a proper kernel (Cuturi, 2009, p.9). Nevertheless, the function q is TI, with  $\psi(t) = 1 - ||t||_2^2/d$ ,  $\overline{\psi} = 1$ , and we can extend the definition of the variance term

$$v = \overline{\psi} - \mathbb{E}_{X, X' \sim \mathbb{P}} \left[ \psi(X' - X) \right] = 2 \operatorname{Tr} \Sigma_{\mathbb{P}} / d.$$

It is natural to define for  $i \in [d]$ ,

$$\widehat{v}^{i}(X_{1},\ldots,X_{n}) \doteq \frac{1}{2n(n-1)} \sum_{s=1}^{n} \sum_{t=1}^{n} \left(X_{t}^{i}-X_{s}^{i}\right)^{2},$$

that will act as an unbiased empirical proxy for  $v^i \doteq \mathbb{V}_{X \sim \mathbb{P}} [X^i]$ , and introduce  $\operatorname{Tr} \widehat{\Sigma} \doteq \sum_{i=1}^d \widehat{v}^i$ , as an estimator for the trace of the covariance matrix  $\Sigma_{\mathbb{P}}$ .

**Lemma 8** For  $x \in [0,1]^d$ , the function  $\mathcal{X}^m \to \mathbb{R}_+, x \mapsto \operatorname{Tr} \widehat{\Sigma}(x)/d$  is weakly (n/(n-1), 0)-self-bounding and has the bounded differences property in the sense of Definition 7.

As a consequence of Lemma 8, the technique of Maurer and Pontil (2009, Theorem 10) provides the following deviation bounds on the square root of the total variance

$$\mathbb{P}\left(b\left[\sqrt{\operatorname{Tr}\Sigma_{\mathbb{P}}/d} - \sqrt{\operatorname{Tr}\widehat{\Sigma}(X)/d}\right] > \sqrt{\frac{2\log(1/\delta)}{n-1}}\right) \le \delta, \qquad b \in \{-1,1\}.$$
(11)

In other words, with high confidence, we can replace the trace of the covariance matrix with its empirical proxy in the convergence bounds.

#### 3.2 Systematic Approach

Our goal is to rigorously develop the approach intuited in the previous section to include a large class of reproducing kernels. We propose the following empirical proxy for the variance term v,

$$\widehat{v}_k(X_1, \dots, X_n) \doteq \frac{1}{n-1} \sum_{t=1}^n \left( k(X_t, X_t) - \frac{1}{n} \sum_{s=1}^n k(X_t, X_s) \right), \tag{12}$$

write more simply  $\hat{v}$  for  $\hat{v}_k$  when k is clear from the context, and promptly verify that  $\hat{v}$  is an unbiased estimator for v.

**Lemma 9 (Unbiasedness)** It holds that  $\mathbb{E}\hat{v} = v$ , where the expectation is taken over the sample  $X_1, \ldots, X_n \stackrel{iid}{\sim} \mathbb{P}$ .

The remainder of this section is devoted to analyzing the self-boundedness properties of  $\hat{v}$  under mild conditions on k and deriving corresponding empirical confidence intervals around  $\mu_{\mathbb{P}}$ . We henceforth let k be a characteristic TI kernel defined by the positive definite function  $\psi$ . We also obtain partial results for the more general case of kernels which are not translation invariant. This analysis is deferred to the appendix Section A (refer to Theorem 29). In the TI setting, the expression of  $\hat{v}$  can be simplified as

$$\widehat{v}(X_1,\ldots,X_n) \doteq \overline{\psi} - \frac{1}{(n-1)n} \sum_{t \neq s} \psi(X_t - X_s).$$
(13)

Since the kernel is characteristic,  $\psi$  cannot be a constant function, that is  $\Delta \psi > 0$ . Under our assumptions, we introduce a function involving  $\hat{v}$  and  $\Delta \psi$  that is weakly self-bounded and has the bounded differences property.

**Lemma 10** Let  $\psi$  define a characteristic TI kernel. The function

$$\mathcal{X}^n \to \mathbb{R}, \, x \mapsto \frac{n}{2\Delta\psi} \widehat{v}(x),$$

is weakly (2,0)-self-bounding and has bounded differences in the sense of Definition 7. This property leads to concentration of  $\sqrt{\hat{v}}$  around  $\sqrt{v}$ .

**Lemma 11** For  $b \in \{-1, 1\}$ , with probability at least  $1 - \delta$ ,

$$b\left[\sqrt{\hat{v}} - \sqrt{v}\right] \le 2\sqrt{\frac{2\Delta\psi\log(1/\delta)}{n}}.$$
(14)

**Theorem 12 (Confidence interval with empirical variance for TI kernel)** Let  $n \in \mathbb{N}$ ,  $X_1, \ldots, X_n \stackrel{iid}{\sim} \mathbb{P}$ , let k be a characteristic translation invariant kernel defined from a positive definite function  $\psi$  [see (4)]. Then with probability at least  $1 - \delta$ , it holds that

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \leq \widehat{\mathcal{B}}_{k,\delta}(X_1, \dots, X_n),$$

with

$$\widehat{\mathcal{B}}_{k,\delta}(X_1,\ldots,X_n) = \widehat{\mathcal{B}}_{\delta} \doteq \sqrt{2\widehat{v}(X_1,\ldots,X_n)\frac{\log(4/\delta)}{n}} + \frac{16}{3}\sqrt{\Delta\psi}\frac{\log(4/\delta)}{n}$$

and where  $\hat{v}$  is the empirical variance proxy defined in (13).

**Example 2** Theorem 12 immediately holds for Gaussian kernels, regardless of the length-scale parameter.

**Remark 13** We make the following observations.

- (i) It still almost surely holds that  $\hat{v} \leq \Delta k \leq 2\overline{k}$ , thus the fully empirical bound can never be more than a constant factor away from (7) or Theorem 1.
- (ii) Crucially, self-boundedness of  $\hat{v}$  and being able to apply Theorem 12 (or Theorem 29 for non TI kernels) only depends on the choice of kernel, and not on the properties of the underlying distribution.
- (iii) Theorem 12 does not depend on smoothness properties of the kernel (also true for Theorem 29 for non TI kernels).

### 3.3 Computability of the Empirical Variance Proxy

The proxy  $\hat{v}$  is computable from data with  $\mathcal{O}(n^2)$  calls to the kernel function. In Section 6, we will discuss how to use variance-aware bounds to improve confidence bounds and test procedures based on the minimum-MMD estimator. In general, this estimator is computed by stochastic gradient descent (SGD) (Dziugaite et al., 2015; Li et al., 2015; Chérief-Abdellatif and Alquier, 2022) or variants like stochastic natural gradient descent (Briol et al., 2019). A single step of SGD requires to sample m iid random variables from  $\mathbb{P}_{\theta}$  where  $\theta$  is the current estimate, and to compute the unbiased estimate of the gradient which requires (among others)  $\mathcal{O}(mn)$  calls to the kernel function. In the natural version, one must also compute a Jacobian at each step, which requires  $\mathcal{O}(n^2)$  calls to the kernel. Morever, in the discussion following Theorem 2 of Briol et al. (2019), it is argued that we should take  $m \gtrsim n$  in these algorithms. Thus, both SGD and its natural variant will require at least  $\mathcal{O}(n^2)$  calls to the kernel at each iteration. In this light, the computation of  $\hat{v}$  does not significantly affect the computational burden. Note however that, in a few situations where the gradient of the MMD is available in close form, it is possible to use a non-stochastic gradient descent. Such examples include Gaussian mean estimation (Chérief-Abdellatif and Alquier, 2022) and Gaussian copula estimation (Alquier et al., 2023). Each step of the gradient descent requires only  $\mathcal{O}(n)$  calls to the kernel and convergence is typically very fast. In these situations, the computation of  $\hat{v}$  can increase the computational cost when n is large.

### 4. Convergence Rates for the Difference of Two Means

We now extend the results of the previous section to the estimation of the norm of the difference of two means  $\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k}$ . One of the applications of these results is two-sample tests, where we test the hypothesis  $\mathbb{P} = \mathbb{Q}$ , as in Gretton et al. (2012). We will cover this application in Section 6. However, they also have an interest on their own, to illustrate the benefits of variance-aware bounds.

In all this section, we assume that two samples are observed:  $X_1, \ldots, X_n$  iid from  $\mathbb{P}$ , and  $Y_1, \ldots, Y_m$  iid from  $\mathbb{Q}$ . We want to estimate the norm of the difference  $\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k}$ .

#### 4.1 First Approach: Estimation by a U-Statistic

A classical approach is to estimate the squared norm  $\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k}^2$  using a U-statistic, as proposed by Gretton et al. (2012), and then to use a Hoeffding or Bernstein-type inequality for U-statistics.

While Gretton et al. (2012) used the U-statistics version of Hoeffding's inequality to control the risk of this estimator, using a Bernstein inequality instead might allow to derive variance-aware bounds. The first version of such inequalities were proven by Hoeffding (1963) and Arcones (1995). More recently, the results of these two papers are included in Theorem 2 of Peel et al. (2010) while Theorem 3 gives an empirical version of each.

We now provide more details on this approach. First, note that

$$\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_{k}}^{2} = \mathbb{E}_{X, X' \sim \mathbb{P}} k(X, X') + \mathbb{E}_{Y, Y' \sim \mathbb{Q}} k(Y, Y') - 2\mathbb{E}_{X \sim \mathbb{P}, Y \sim \mathbb{Q}} k(X, Y)$$

can be directly estimated by the U-statistic:

$$\tilde{U}_{n,m} = \frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j \neq i} k(X_i, X_j) + \frac{1}{m(m-1)} \sum_{i=1}^{m} \sum_{j \neq i} k(Y_i, Y_j) - \frac{2}{mn} \sum_{i=1}^{n} \sum_{j=1}^{m} k(X_i, Y_j)$$

as proposed by Gretton et al. (2012). This is an unbiased estimator:  $\mathbb{E}\tilde{U}_{n,m} = \|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k}^2$ . In the special case m = n, we can replace  $\tilde{U}_{n,n}$  by the simpler

$$\hat{U}_n = \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i} \left[ k(X_i, X_j) + k(Y_i, Y_j) - k(X_i, Y_j) - k(X_j, Y_i) \right].$$

We will illustrate the U-statistics approach in this simpler setting.

Using Theorem 2 in Peel et al. (2010), with their  $q((X_i, Y_i), (X_j, Y_j)) = k(X_i, X_j) + k(Y_i, Y_j) - k(X_i, Y_j) - k(X_j, Y_i)$  and their m = 2, we obtain the first part of the following statement. Using Theorem 3, we obtain the second one.

### Theorem 14 Put

$$\sigma_{\mathbb{P},\mathbb{Q}}^2 = \mathbb{V}_{X \sim \mathbb{P}, Y \sim \mathbb{Q}} [\mathbb{E}_{X' \sim \mathbb{P}} k(X, X') + \mathbb{E}_{Y' \sim \mathbb{Q}} k(Y, Y') - \mathbb{E}_{Y' \sim \mathbb{Q}} k(X, Y') - \mathbb{E}_{X' \sim \mathbb{P}} k(X', Y)]$$

and

$$\hat{\sigma}_{\mathbb{P},\mathbb{Q}}^{2} = \frac{1}{n(n-1)(n-2)} \sum_{i=1}^{n} \sum_{j \neq i} \sum_{k \neq i,j} \left[ k(X_{i}, X_{j}) + k(Y_{i}, Y_{j}) - k(X_{i}, Y_{j}) - k(X_{j}, Y_{i}) \right] \\ \times \left[ k(X_{i}, X_{k}) + k(Y_{i}, Y_{k}) - k(X_{i}, Y_{k}) - k(X_{k}, Y_{i}) \right].$$

Note that  $\mathbb{E}\hat{\sigma}_{\mathbb{P},\mathbb{Q}}^2 = \sigma_{\mathbb{P},\mathbb{Q}}^2 + \|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k}^4$ . For any  $\delta \in (0,1]$ , with probability at least  $1 - \delta$ ,

$$\left| \left\| \mu_{\mathbb{P}} - \mu_{\mathbb{Q}} \right\|_{\mathcal{H}_{k}}^{2} - \hat{U}_{n} \right| \leq 2\Delta k \sqrt{\frac{8\sigma_{\mathbb{P},\mathbb{Q}}^{2}}{n} \log \frac{4}{\delta} + 2\Delta k \frac{64 + \frac{1}{6}}{n} \log \frac{4}{\delta}}.$$

Moreover, with probability at least  $1 - \delta$ ,

$$\left| \left\| \mu_{\mathbb{P}} - \mu_{\mathbb{Q}} \right\|_{\mathcal{H}_{k}}^{2} - \hat{U}_{n} \right| \leq 2\Delta k \sqrt{\frac{8\hat{\sigma}_{\mathbb{P},\mathbb{Q}}^{2}}{n} \log \frac{8}{\delta} + 2\Delta k \frac{64 + \frac{1}{6} + 5\sqrt{2}}{n} \log \frac{8}{\delta}}.$$

Observe that, when  $\mathbb{P} \neq \mathbb{Q}$ , the bound states that  $\sqrt{\hat{U}_n}$  will be of the order of  $\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k}$ . On the other hand, when  $\mathbb{P} = \mathbb{Q}$ , we have both  $\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k} = 0$  and  $\sigma_{\mathbb{P},\mathbb{Q}}^2 = 0$ , and thus the first inequality in the theorem gives:

$$\sqrt{\hat{U}_n} \le \sqrt{2\Delta k} \frac{64 + \frac{1}{6} + 5\sqrt{2}}{n} \log \frac{8}{\delta},\tag{15}$$

which does not depend on the variances of  $\mathbb{P}$  and  $\mathbb{Q}$ .

### 4.2 Variance-Aware Control of the Fluctuations for Each Sample

An alternative approach is to apply the triangle inequality to upper bound separately the fluctuations for each sample:

$$\left| \left\| \mu_{\mathbb{P}} - \mu_{\mathbb{Q}} \right\|_{\mathcal{H}_{k}} - \left\| \hat{\mu}_{\mathbb{P}}(X_{1}, \dots, X_{n}) - \hat{\mu}_{\mathbb{Q}}(Y_{1}, \dots, Y_{n}) \right\|_{\mathcal{H}_{k}} \right|$$

$$\leq \left\| \mu_{\mathbb{P}} - \hat{\mu}_{\mathbb{P}}(X_{1}, \dots, X_{n}) \right\|_{\mathcal{H}_{k}} + \left\| \mu_{\mathbb{Q}} - \hat{\mu}_{\mathbb{Q}}(Y_{1}, \dots, Y_{n}) \right\|_{\mathcal{H}_{k}}.$$

For example, a direct application of Theorem 1 with a union bound gives the following corollary.

**Corollary 15** Let  $X_1, \ldots, X_n \sim \mathbb{P}, Y_1, \ldots, Y_m \sim \mathbb{Q}, k \colon \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  be a reproducing kernel, and let

$$v_{\mathbb{P}} \doteq \mathbb{E}_{X \sim \mathbb{P}} \left\| k(X, \cdot) - \mu_{\mathbb{P}} \right\|_{\mathcal{H}_{k}}^{2} \text{ and } v_{\mathbb{Q}} \doteq \mathbb{E}_{Y \sim \mathbb{Q}} \left\| k(Y, \cdot) - \mu_{\mathbb{Q}} \right\|_{\mathcal{H}_{k}}^{2}$$

and  $\overline{k} = \sup_{x \in \mathcal{X}} k(x, x)$ . With probability at least  $1 - \delta$ , it holds that

$$\left| \|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_{k}} - \|\hat{\mu}_{\mathbb{P}}(X_{1}, \dots, X_{n}) - \hat{\mu}_{\mathbb{Q}}(Y_{1}, \dots, Y_{m})\|_{\mathcal{H}_{k}} \right|$$

$$\leq \sqrt{2v_{\mathbb{P}} \frac{\log(4/\delta)}{n}} + \sqrt{2v_{\mathbb{Q}} \frac{\log(4/\delta)}{m}} + \left(\frac{1}{n} + \frac{1}{m}\right) \frac{4}{3}\sqrt{k} \log(4/\delta).$$

Of course, we can also state results with empirical variance instead of the true variance, by using Theorems 12 and 29.

In order to compare this result to the U-statistics approach, consider the case n = m. Corollary 15 gives the following upper bound:

$$\begin{aligned} \|\hat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) - \hat{\mu}_{\mathbb{Q}}(Y_1,\ldots,Y_n)\|_{\mathcal{H}_k} \\ &\leq \|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k} + \sqrt{2v_{\mathbb{P}}\frac{\log(4/\delta)}{n}} + \sqrt{2v_{\mathbb{Q}}\frac{\log(4/\delta)}{n}} + \frac{8}{3n}\sqrt{\overline{k}}\log(4/\delta). \end{aligned}$$

In particular, when  $\mathbb{P} = \mathbb{Q}$ , the bound becomes:

$$\|\hat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) - \hat{\mu}_{\mathbb{Q}}(Y_1,\ldots,Y_n)\|_{\mathcal{H}_k} \le 2\sqrt{2v_{\mathbb{P}}\frac{\log(4/\delta)}{n}} + \frac{8}{3n}\sqrt{\overline{k}}\log(2/\delta).$$
(16)

Even though  $\|\hat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) - \hat{\mu}_{\mathbb{Q}}(Y_1,\ldots,Y_n)\|_{\mathcal{H}_k}^2$  is not an unbiased estimator of  $\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}_k}^2$ as  $\hat{U}$ , (15) and (16) show the fluctuations of  $\|\hat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) - \hat{\mu}_{\mathbb{Q}}(Y_1,\ldots,Y_n)\|_{\mathcal{H}_k}$  when  $\mathbb{P} = \mathbb{Q}$  are upper bounded by  $\sqrt{v_{\mathbb{P}}/n} + 1/n$  while the ones of  $\sqrt{\hat{U}}$  by  $\sqrt{1/n}$ . This can be a serious improvement if  $v_{\mathbb{P}}$  is small. Note that we do not claim superiority of the plug-in estimator, but rather that the currently available non-asymptotic bounds for this estimator are tighter. This is an illustration of the power of the variance-aware bounds. A way to compare both estimators would be through an accurate study of their asymptotic fluctuations when  $\mathbb{P} = \mathbb{Q}$ . So far, this analysis is only available for the U-statistic, see Theorem 12 in Gretton et al. (2012).

### 5. Convergence Rates with Time-Dependent Data

In this section, we establish convergence rates for cases where the data  $X_1, \ldots, X_n$  is not independent. Namely, we will assume the data to be a stationary mixing sequence (Bradley, 2005; Doukhan, 2012). In this setting, the observations are identically distributed with marginal distribution  $\mathbb{P}$ , but exhibit time dependencies that diminish as the time interval increases.

**Lemma 16** Suppose that  $X_1, \ldots, X_n$  is a stationary sequence of possibly dependent random variables. Then

$$\mathbb{E} \|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k}^2 = \frac{1}{n} \left( v + \Sigma_n \right),$$

with

$$\Sigma_n \doteq \frac{2}{n} \sum_{t=2}^n (n-t+1)\rho_t, \qquad \text{where} \qquad \rho_t \doteq \mathbb{E} \langle k(X_t, \cdot) - \mu_{\mathbb{P}}, k(X_1, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_k},$$

are the covariance coefficients in the RKHS introduced in Chérief-Abdellatif and Alquier (2022). In particular, when the process is iid, the expression above simplifies to

$$\mathbb{E} \|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k}^2 = \frac{v}{n}.$$

#### 5.1 For $\phi$

The first flavor of mixing we consider is  $\phi$ -mixing, introduced by Ibragimov (1962). Recall that the  $\phi$ -mixing coefficient (Bradley, 2005; Doukhan, 2012) is defined for two  $\sigma$ -fields  $\mathcal{A}$  and  $\mathcal{B}$  by

$$\phi(\mathcal{A}, \mathcal{B}) \doteq \sup_{\substack{A \in \mathcal{A}, B \in \mathcal{B} \\ \mathbb{P}(A) > 0}} \left| \mathbb{P}\left(B|A\right) - \mathbb{P}\left(B\right) \right|.$$

For  $s \in \mathbb{N}$ , we further define

$$\phi(s) \doteq \sup_{r \in \mathbb{N}} \phi\left(\sigma\left(\{X_t \colon t \le r\}\right), \sigma\left(\{X_t \colon t \ge r + s\}\right)\right),$$

where for  $T \subset \mathbb{N}$ ,  $\sigma(\{X_t\}_{t \in T})$  is the  $\sigma$ -field generated by the random variables  $\{X_t\}_{t \in T}$ . The random process is then called  $\phi$ -mixing when  $\lim_{s\to\infty} \phi(s) = 0$ . Additionally, we define the triangular coupling matrix  $\Gamma$  (Samson, 2000) as follows. For  $t, s \in [n]$ ,

$$\Gamma(t,s) \doteq \begin{cases} 1 & \text{when } t = s \\ 0 & \text{when } t > s \\ \sqrt{2\phi(\sigma(X_1,\dots,X_t),\sigma(X_s,\dots,X_n))} & \text{otherwise.} \end{cases}$$
(17)

**Theorem 17 (Variance-aware confidence interval with**  $\phi$ -mixing data) Let  $\delta \in (0, 1)$ , and  $X_1, \ldots, X_n$  be a stationary  $\phi$ -mixing sequence with marginal distribution  $\mathbb{P}$ . We let  $k: \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  be a reproducing kernel with  $\overline{k} = \sup_{x \in \mathcal{X}} k(x, x)$ . With probability at least  $1 - \delta$ , it holds that

$$\|\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n)-\mu_{\mathbb{P}}\|_{\mathcal{H}_k} \leq \mathcal{B}^{\phi}_{k,\delta}(\mathbb{P},n),$$

with

$$\mathcal{B}_{k,\delta}^{\phi}(\mathbb{P},n) = \mathcal{B}_{\delta}^{\phi} \doteq \sqrt{\frac{v + \Sigma_n}{n}} + 4\sqrt{\frac{2v \|\Gamma\|_2 \log(1/\delta)}{n}} + \frac{8\overline{k} \|\Gamma\|_2 \log(1/\delta)}{n}$$

where  $\Sigma_n$  is defined in Lemma 16,  $\Gamma$  is defined in Eq. (17) and  $\|\cdot\|_2$  is the spectral norm.

In particular, when the process is iid,  $\Gamma$  is the identity matrix, hence  $\|\Gamma\|_2 = 1$ .

**Example 3 (Uniformly ergodic Markov chains)** Suppose that there exists  $\phi_0, \phi_1 \in \mathbb{R}_+$  such that for all  $t \in \mathbb{R}$ ,  $\phi(t) \leq \phi_0 \exp(-\phi_1 t)$ . Then it holds (refer for example to Samson (2000)) that

$$\|\Gamma\|_2 \le \frac{\sqrt{2\phi_0}}{1 - e^{-\phi_1/2}}.$$

In particular, this is known to hold for a uniformly ergodic Markov chain (Roberts and Rosenthal, 2004) with transition operator P and stationary distribution  $\pi$ . In this case we can choose,  $\phi_0 = 4$  and  $\phi_1 = \log(2)/t_{\text{mix}}$ , where

$$t_{\mathsf{mix}} \doteq \min_{t \in \mathbb{N}} \left\{ \sup_{x \in \mathcal{X}} \left\| P^t(x, \cdot) - \pi(\cdot) \right\|_{\mathsf{TV}} < 1/4 \right\},\,$$

is called the mixing time of the chain with respect to total variation (Levin et al., 2009); here we obtain  $\|\Gamma\|_2 \leq 10t_{\text{mix}}$ .

#### 5.2 For $\beta$ -Mixing Processes

We next consider  $\beta$ -mixing, a common assumption in the machine learning literature (Mohri and Rostamizadeh, 2010). Recall that the  $\beta$ -mixing coefficient (Bradley, 2005; Doukhan, 2012) is defined for two  $\sigma$ -fields  $\mathcal{A}$  and  $\mathcal{B}$ ,

$$\beta(\mathcal{A}, \mathcal{B}) \doteq \sup \frac{1}{2} \sum_{i=1}^{I} \sum_{j=1}^{J} |\mathbb{P}(A_i \cap B_j) - \mathbb{P}(A_i)\mathbb{P}(B_j)|,$$

where the supremum is taken over all pairs of finite partitions  $\{A_1, \ldots, A_I\}$  and  $\{B_1, \ldots, B_J\}$ of  $\mathcal{X}$  such that for any  $1 \leq i \leq I$  and any  $1 \leq j \leq J$ ,  $A_i \in \mathcal{A}$  and  $B_j \in \mathcal{B}$  (Bradley, 2005, Equation 7). For  $s \in \mathbb{N}$ , we further define

$$\beta(s) \doteq \sup_{r \in \mathbb{N}} \beta\left(\sigma\left(\{X_t \colon t \le r\}\right), \sigma\left(\{X_t \colon t \ge r+s\}\right)\right),$$

where for  $T \subset \mathbb{N}$ ,  $\sigma(\{X_t\}_{t \in T})$  is the  $\sigma$ -field generated by the random variables  $\{X_t\}_{t \in T}$ . The random process is then called  $\beta$ -mixing<sup>6</sup> when  $\lim_{s\to\infty} \beta(s) = 0$ . In this case, for  $\xi \in (0, 1)$ , we define the  $\beta$ -mixing time as

$$t_{\min}^{\beta}(\xi) \doteq \operatorname*{arg\,min}_{t \in \mathbb{N}} \left\{ \beta(t) < \xi \right\}.$$

We observe that when the process is a stationary time-homogeneous Markov chain,  $t_{mix}^{\beta}$  corresponds to the definition of the average-mixing time (Münch and Salez, 2023; Wolfer

<sup>6.</sup> A  $\beta$ -mixing process is also called *absolutely regular* in the literature.

and Alquier, 2024). Note that  $\beta(\mathcal{A}, \mathcal{B}) \leq \phi(\mathcal{A}, \mathcal{B})$  (Bradley, 2005, Equation 1.11), and no reverse inequality holds for any universal constant, thus  $\beta$ -mixing is a strictly weaker notion than  $\phi$ -mixing.

**Theorem 18 (Variance-aware confidence interval with**  $\beta$ -mixing data) We let  $\delta \in (0,1)$ , and  $X_1, \ldots, X_n$  be a stationary  $\beta$ -mixing sequence. We let  $k: \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  be a reproducing kernel with  $\overline{k} = \sup_{x \in \mathcal{X}} k(x, x)$ . Finally, we suppose for simplicity<sup>7</sup> that n is a multiple of  $\tau_{n,\delta}^{\beta}$  where

$$\tau_{n,\delta}^{\beta} \doteq \operatorname*{arg\,min}_{s \in \mathbb{N}} \left\{ s \geq t_{\mathsf{mix}}^{\beta} \left( \frac{\delta}{6(n/(2s)-1)} \right) \right\}.$$

With probability at least  $1 - \delta$ , it holds that

$$\|\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n)-\mu_{\mathbb{P}}\|_{\mathcal{H}_k} \leq \mathcal{B}_{k,\delta}^{\beta}(\mathbb{P},n),$$

with

$$\mathcal{B}_{k,\delta}^{\beta}(\mathbb{P},n) = \mathcal{B}_{\delta}^{\beta} \doteq 2\sqrt{\left(v + \Sigma_{\tau_{n,\delta}^{\beta}}\right)\frac{\log(3/\delta)}{n} + \frac{8}{3}\tau_{n,\delta}^{\beta}\sqrt{\overline{k}}\frac{\log(3/\delta)}{n}},$$

where  $\Sigma_s$  is defined in Lemma 16.

Example 4 (Countable state time-homogeneous Markov chains) Irreducible, aperiodic and stationary countable Markov chains are always  $\beta$ -mixing (Bradley, 2005). Assume that there exist  $\beta_1 \in \mathbb{R}_+$  and  $b \in (1, \infty)$  such that for any  $t \in \mathbb{N}, \beta(t) \leq \beta_1/t^b$ , that is, the chain is only known to be algebraically mixing. Interestingly, such a chain may not be  $\phi$ -mixing, thus this example cannot be recovered from Theorem 17.

### 6. Applications

We put the apparatus developed in Section 2 and Section 3 to application in the context of hypothesis testing and robust parametric estimation. In this section, all considered kernels will be TI unless otherwise specified, and  $X_1, \ldots, X_n$  are i.i.d. from  $\mathbb{P} \in \mathcal{P}(\mathcal{X})$ . We introduce a statistical model  $\mathcal{M} = \{\mathbb{P}_{\theta} : \theta \in \Theta\}$  indexed by the parameter space  $\Theta$ . Examples of models studied in the literature on MMD include parametric models such as the Gaussian model  $\mathbb{P}_{\theta} = \mathcal{N}(\theta, \sigma^2 I_d)$  (with  $\sigma^2$  known) and mixture of Gaussians (Briol et al., 2019; Chérief-Abdellatif and Alquier, 2022), copulas Alquier et al. (2023) but also more complex models such as generative adversarial networks (Dziugaite et al., 2015; Li et al., 2015), stochastic volatility models and stochastic differential equations (Briol et al., 2019). Following Briol et al. (2019), our estimator will be the closest element<sup>8</sup> of the model

<sup>7.</sup> An adaptation of the proof removes this assumption and recovers similar bounds up to universal constants.

<sup>8.</sup> Note that it might be that the infimum is reached by multiple elements of  $\Theta$ , or is not reached. The first case does not lead to any difficulty, we must simply define a rule to break ties (for example, we can equip  $\Theta$  with a total order and chose the smallest minimizer according to this order). While Briol et al. (2019) provide sufficient conditions to ensure that the infimum is reached, the non-existence of a minimizer is also not a problem in practice: all the non-asymptotic results in Briol et al. (2019) and Chérief-Abdellatif and Alquier (2022) can easily be extended to any  $\epsilon$ -minimizer, for  $\epsilon$  small enough.

 $\mathcal{M}$  to the empirical measure obtained from X, where the distance is measured in the RKHS,

$$\left\|\mu_{\mathbb{P}_{\widehat{\theta}_n(X_1,\dots,X_n)}} - \widehat{\mu}_{\mathbb{P}}(X_1,\dots,X_n)\right\|_{\mathcal{H}_k} = \inf_{\theta\in\Theta} \left\|\mu_{\mathbb{P}_\theta} - \widehat{\mu}_{\mathbb{P}}(X_1,\dots,X_n)\right\|_{\mathcal{H}_k}.$$
 (18)

The computation of  $\hat{\theta}_n(X_1, \ldots, X_n)$  is usually done via stochastic gradient descent and variants, see Remark 13 above.

### 6.1 Hypothesis Testing

In this subsection we study hypothesis testing based on MMD. In the literature, two kinds of tests were proposed and studied: two-sample tests, and goodness of fit. In the twosample test problem, we are given  $X_1, \ldots, X_n$  iid from some  $\mathbb{P}_X$  and  $Y_1, \ldots, Y_m$  iid from some  $\mathbb{P}_Y$ , and we want to test  $\mathbf{H}_0 : \mathbb{P}_X = \mathbb{P}_Y$  against the alternative  $\mathbf{H}_1 : \mathbb{P}_X \neq \mathbb{P}_Y$ . In the goodness of fit problem, we are given  $X_1, \ldots, X_n$  iid from  $\mathbb{P}$  and we wish to test the hypothesis  $\mathbf{H}_0 : \mathbb{P} \in \mathcal{M} = {\mathbb{P}_{\theta} : \theta \in \Theta}$ , against the alternative  $\mathbf{H}_1 : \mathbb{P} \notin {\mathbb{P}_{\theta} : \theta \in \Theta}$ .

In the MMD literature, two-sample tests are more prevalent (Gretton et al., 2009, 2012). Recent work tackle goodness of fit testing (Chwialkowski et al., 2016; Jitkrittum et al., 2017), albeit with an asymptotic treatment. We will study both problems here. In both cases, we propose a non-asymptotic treatment. That is, the level of the test is smaller than  $\alpha$ for a finite sample size, and not only asymptotically. Moreover, in both cases, our datadependent bound (Theorem 12) allows to increase the power of the test, when compared to the procedure that would be based on the McDiarmid-based bound in (7).

### 6.1.1 Goodness-of-Fit Test

Recall that we define the significance level  $\alpha \in [0, 1]$  of a test as the probability of outputting  $\mathbf{H}_1$  when  $\mathbf{H}_0$  is true. Taking advantage of a non-asymptotic bound  $\mathcal{B}(X_1, \ldots, X_n, \alpha)$  (for example  $\mathcal{B}(X_1, \ldots, X_n, \alpha) = \hat{\mathcal{B}}_{\alpha}$  given in Theorem 12), we can design a test with prescribed level  $\alpha$  for any n as follows. We introduce the test statistic

$$T(X_1,\ldots,X_n) \doteq \inf_{\theta \in \Theta} \|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}_{\theta}}\|_{\mathcal{H}_k},$$

and reject  $\mathbf{H}_0$  on the "critical set"

$$\mathcal{C}(X_1,\ldots,X_n) \doteq \{T > \mathcal{B}(X_1,\ldots,X_n,\alpha)\}.$$

We show that the probability of rejection under the null hypothesis is at most  $\alpha$  and that the test is consistent.

**Theorem 19** Let  $n \in \mathbb{N}$ ,  $X_1, \ldots, X_n \sim \mathbb{P}$ , and k a translation invariant kernel. Let  $\mathcal{B}(\cdot, \alpha)$  be such that

(a) for any  $\alpha$ , with probability at least  $1 - \alpha$ ,  $\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \leq \mathcal{B}(X_1, \ldots, X_n, \alpha)$ .

Then

(i) 
$$\mathbb{P}_{\mathbf{H}_0} \left( \mathcal{C}(X_1, \ldots, X_n) \right) \leq \alpha.$$

If, moreover,

(b) for any fixed  $\alpha$ ,  $\mathcal{B}(X_1, \ldots, X_n, \alpha) \xrightarrow[n \to \infty]{a.s.} 0$ ,

(c) there is a sequence  $\alpha_n \to 0$  such that  $\mathcal{B}(X_1, \ldots, X_n, \alpha_n) \xrightarrow[n \to \infty]{a.s.} 0$ ,

then

(ii) when the model  $\mathcal{M}$  is closed with respect to the MMD metric <sup>9</sup>,

$$\lim_{n \to \infty} \mathbb{P}_{\mathbf{H}_1} \left( \mathcal{C}(X_1, \dots, X_n) \right) = 1,$$

more precisely,

$$1 - \mathbb{P}_{\mathbf{H}_1} \left( \mathcal{C}(X_1, \dots, X_n) \right) = \mathcal{O}(\alpha_n).$$

Obviously, the empirical bound  $\widehat{\mathcal{B}}_{k,\alpha}$  of Theorem 12 satisfies (a) and (b) (regardless of  $\mathbb{P}$ ), and (c) with  $\alpha_n = \exp(-n^{1-\epsilon})$  for any fixed  $\epsilon > 0$ . So does the McDiarmid-based bound in (7). On the other hand, we claimed that the empirical bound is smaller than bounds that does not take the variance into account. In other words, the power of the test  $\mathbb{P}_{\mathbf{H}_1}(\mathcal{C}(X_1,\ldots,X_n))$  for a finite *n* will be larger if we use the variance aware bound.

**Example 5** We first consider a single hypothesis test (that is a special case of goodnessof-fit when  $\Theta$  is a singleton). We consider data in  $\mathbb{R}^2$ : here  $\mathbb{P} = \mathcal{N}(0, \sigma^2 I_2)$  with  $\mathbb{P}_{\theta} =$  $\mathcal{N}(\theta, I_2)$ , and  $\Theta = \{(1, 1)\}$ . In other words,  $\mathbf{H}_0$  is true iff  $\sigma = 1$ . Note that in this case, if  $Y_1, \ldots, Y_n$  are iid from  $\mathbb{P}_{\theta} = \mathbb{P}_{\{(1,1)\}}$ , then  $q_{1-\alpha}$  defined as the  $(1-\alpha)$ -quantile of  $\|\mu_{\mathbb{P}_0} - \widehat{\mu}_{\mathbb{P}}(Y_1, \dots, Y_n)\|_{\mathcal{H}_k}$  allows to define a test with rejection zone given by  $\mathcal{C}'(X_1, \dots, X_n) = \mathcal{C}'(X_1, \dots, X_n)$  $\{T > q_{1-\alpha}\}$  that satisfies  $\mathbb{P}_{\mathbf{H}_0}(\mathcal{C}'(X_1,\ldots,X_n)) = \alpha$  by definition. It is a very natural procedure to estimate  $q_{1-\alpha}$  by Monte-Carlo, by sampling multiple times  $Y_1, \ldots, Y_n$  from  $\mathbb{P}_{\theta} = \mathbb{P}_{\{(1,1)\}}$ . This leads to a Monte-Carlo estimator  $\hat{q}_{1-\alpha}$  of  $q_{1-\alpha}$ . In our simulations, we sample  $X_1, \ldots, X_n$  from  $\mathbb{P}$ , perform the test based on  $\hat{q}_{1-\alpha}$ , the test based on the empirical bound and the test based on McDiarmid bound. This is repeated 100 times for each value  $\sigma \in \{0, 1/50, 2/50, \dots, 1\}$ . We report the frequency of rejections in Figure 1. The kernel used is a Gaussian kernel with  $\gamma = 1$ , and we consider sample sized  $n \in \{16, 40, 100, 250\}$ . We observe that, when compared to test based  $\hat{q}_{1-\alpha}$ , both tests based on bounds have a weak power (note that we did not try to optimize  $\gamma$  for now, this question will be tackled later). However, the test based on the empirical bound indeed rejects more often  $\mathbf{H}_0$  when  $\mathbf{H}_1$  is true. The improvement is clearer for small sample sizes.

**Example 6** We now consider a proper goodness-of-fit test in  $\mathbb{R}^2$ : we still consider  $\mathbb{P} = \mathcal{N}(0, \sigma^2 I_2)$  and  $\mathbb{P}_{\theta} = \mathcal{N}(\theta, I_2)$ , with  $\Theta = \mathbb{R}^2$ . It is important to observe that in this case, because  $\mathbf{H}_0$  is composite, we don't have a natural definition for  $q_{1-\alpha}$  as in the previous example. We sample  $X_1, \ldots, X_n$  from  $\mathbb{P}$ , perform the test based on the empirical bound and the test based on McDiarmid bound. This is repeated 100 times for each value  $\sigma \in \{0, 1/50, 2/50, \ldots, 1\}$ . We report the frequency of rejections in Figure 2. The kernel used is a Gaussian kernel with  $\gamma = 1$ , and we consider sample sized  $n \in \{16, 40, 100, 250\}$ . Similar comments to the previous case apply, but in this case, this makes the test based on the empirical Bernstein bound the best test available.

<sup>9.</sup> For any  $\mathbb{Q}$ , if there is a sequence  $(\theta_h)_{h\in\mathbb{N}}$  of elements of  $\Theta$  such that  $\left\|\mu_{\mathbb{Q}} - \mu_{\mathbb{P}_{\theta_h}}\right\|_{\mathcal{H}_k} \xrightarrow[h \to \infty]{} 0$ , then  $\mathbb{Q} = \mathbb{P}_{\theta}$  for some  $\theta \in \Theta$ .



Figure 1: Comparison of the test based on the Bernstein empirical (EmpBer) bound, versus the test based on McDiarmid bound (McDia), and the test based on the Monte-Carlo estimation of the quantile  $q_{1-\alpha}$ . Frequency of rejection of  $\mathbf{H}_0: \mathbb{P} \in \{\mathcal{N}((1,1), I_2)\}$  as a function of  $\sigma$  with  $\mathbb{P} = \mathcal{N}(0, \sigma^2 I_2)$ .



Figure 2: Comparison of the test based on the Bernstein empirical (EmpBer) bound, versus the test based on McDiarmid bound (McDia). Frequency of rejection of  $\mathbf{H}_0 : \mathbb{P} \in \{\mathcal{N}(\theta, I_2), \theta \in \mathbb{R}^2\}$  as a function of  $\sigma$  with  $\mathbb{P} = \mathcal{N}(0, \sigma^2 I_2)$ .

# 6.1.2 Two-Sample Test

Let  $(n,m) \in \mathbb{N}^2$ ,  $X_1, \ldots, X_n \sim \mathbb{P}_X$  and  $Y_1, \ldots, Y_m \sim \mathbb{P}_Y$ . Here, we wish to test  $\mathbf{H}_0 : \mathbb{P}_X = \mathbb{P}_Y$  against  $\mathbf{H}_1 : \mathbb{P}_X \neq \mathbb{P}_Y$ . This time, we use the test statistic

$$T_2(X_1,\ldots,X_n,Y_1,\ldots,Y_m) \doteq \|\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) - \widehat{\mu}_{\mathbb{P}}(Y_1,\ldots,Y_m)\|_{\mathcal{H}_k},$$

and reject  $\mathbf{H}_0$  on the "critical set"

$$\mathcal{C}_2(X_1,\ldots,X_n,Y_1,\ldots,Y_m) \doteq \{T_2 > \mathcal{B}(X_1,\ldots,X_n,\alpha/2) + \mathcal{B}(Y_1,\ldots,Y_m,\alpha/2)\}.$$

**Theorem 20** Let  $(n,m) \in \mathbb{N}^2$ ,  $X_1, \ldots, X_n \sim \mathbb{P}_X$  and  $Y_1, \ldots, Y_m \sim \mathbb{P}_Y$ , and k a translation invariant kernel. Let  $\mathcal{B}(\cdot, \alpha)$  be such that

(a) for any  $\alpha$ , each of these inequalities hold with probability at least  $1 - \alpha$  each:

$$\begin{aligned} \|\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) - \mu_{\mathbb{P}_X}\|_{\mathcal{H}_k} &\leq \mathcal{B}(X_1,\ldots,X_n,\alpha) \\ \|\widehat{\mu}_{\mathbb{P}}(Y_1,\ldots,Y_m) - \mu_{\mathbb{P}_Y}\|_{\mathcal{H}_k} &\leq \mathcal{B}(Y_1,\ldots,Y_m,\alpha). \end{aligned}$$

Then

(i) 
$$\mathbb{P}_{\mathbf{H}_0} \left( \mathcal{C}_2(X_1, \dots, X_n, Y_1, \dots, Y_m) \right) \leq \alpha.$$

Moreover, if

(b) for any fixed 
$$\alpha$$
,  $\mathcal{B}(X_1, \ldots, X_n, \alpha) \xrightarrow[n \to \infty]{a.s.} 0$  and  $\mathcal{B}(Y_1, \ldots, Y_n, \alpha) \xrightarrow[m \to \infty]{a.s.} 0$ ,

(c) there is a sequence  $\alpha_n \to 0$  such that

$$\mathcal{B}(X_1,\ldots,X_n,\alpha_n) \xrightarrow[n\to\infty]{a.s.} 0 \text{ and } \mathcal{B}(Y_1,\ldots,Y_n,\alpha_m) \xrightarrow[m\to\infty]{a.s.} 0$$

then

(*ii*) 
$$\lim_{n,m\to\infty} \mathbb{P}_{\mathbf{H}_1} \left( \mathcal{C}_2(X_1,\ldots,X_n,Y_1,\ldots,T_m) \right) = 1$$

Here again, the empirical bound  $\widehat{\mathcal{B}}_{k,\alpha}$  of Theorem 12 satisfies (a) and (b), and (c) with  $\alpha_n = \exp(-n^{1-\epsilon})$  for any  $\epsilon > 0$ .

#### 6.2 Robust Parametric Estimation under Huber Contamination

Chérief-Abdellatif and Alquier (2022, Proof of Theorem 3.1) show how to upper-bound (18) by the estimation error of the empirical measure in the RKHS and the approximation error of the model,

$$\left\|\mu_{\mathbb{P}_{\widehat{\theta}_n}} - \mu_{\mathbb{P}}\right\|_{\mathcal{H}_k} \le \inf_{\theta \in \Theta} \left\|\mu_{\mathbb{P}_{\theta}} - \mu_{\mathbb{P}}\right\|_{\mathcal{H}_k} + 2 \left\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\right\|_{\mathcal{H}_k}.$$
(19)

Recall that in the Huber contamination model (Huber, 2011), the observations  $X_1, \ldots, X_n$  are drawn independently from the mixture

$$\mathbb{P} = (1 - \xi)\mathbb{P}_{\theta_0} + \xi\mathbb{H},$$

where  $\xi \in [0, 1/2)$  is the contamination rate,  $\mathbb{H}$  is some unknown noise distribution and  $\theta_0 \in \Theta$ . In this setting, we can control the approximation error of the model: starting from (19), and adapting the argument of Chérief-Abdellatif and Alquier (2022, Corollary 3.3), we obtain

$$\left\|\mu_{\mathbb{P}_{\widehat{\theta}_n}} - \mu_{\mathbb{P}_{\theta_0}}\right\|_{\mathcal{H}_k} \le 4\xi \sqrt{\overline{\psi}} + 2 \left\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\right\|_{\mathcal{H}_k}.$$

We wish to apply Theorem 1 to the second term. However, when the sample is contaminated, the variance term  $v(\mathbb{P})$  we collect also depends on the properties of  $\mathbb{H}$ , which are unknown for all we know  $\mathbb{H}$  does not even have finite moments. Fortunately, we now see how to bound  $v(\mathbb{P})$  is terms of  $v(\mathbb{P}_{\theta_0})$  and  $\xi$ .

6.2.1 Improved Confidence Bounds in the Huber Contamination Model

**Lemma 21** In the Huber contamination model, that is,  $\mathbb{P} = (1 - \xi)\mathbb{P}_{\theta_0} + \xi\mathbb{H}$ , writing  $v = v(\mathbb{P})$  and  $v_0 = v(\mathbb{P}_{\theta_0})$ , it holds that

$$v \le v_0 + 2\xi(\Delta \psi - v_0) + \xi^2(v_0 - \Delta \psi).$$
 (20)

The bound can be streamlined; when  $\psi \geq 0$ ,

$$v \le (1 - 2\xi)v_0 + 2\xi\overline{\psi},$$

and otherwise

$$v \le (1 - 2\xi)v_0 + 2\xi\Delta\psi.$$

**Remark 22** In (20), observe that when  $\psi \geq 0$  and  $\xi \to 1$ , the first term vanishes, and  $v \to \overline{\psi}$ , recovering the distribution independent rates. Conversely, when  $\xi \to 0$ , we confirm that  $v \to v_0$ , which could be further bounded using properties of the model. In practice, we will focus on small values  $\xi \ll 1/2$ , and the simplified bounds will be sufficient.

An application of Theorem 1 yields the following.

**Corollary 23** In the Huber contamination model  $\mathbb{P} = (1 - \xi)\mathbb{P}_{\theta_0} + \xi\mathbb{H}$ , writing  $v = v(\mathbb{P})$ and  $v_0 = v(\mathbb{P}_{\theta_0})$ , with probability at least  $1 - \delta$ ,

$$\left\|\mu_{\mathbb{P}_{\widehat{\theta}_n}} - \mu_{\mathbb{P}_{\theta_0}}\right\|_{\mathcal{H}_k} \le 4\xi\sqrt{\overline{\psi}} + 2\sqrt{2[(1-2\xi)v_0 + 2\xi\overline{\psi}]}\frac{\log(2/\delta)}{n} + 2\frac{4\sqrt{\overline{\psi}}}{3}\frac{\log(2/\delta)}{n}.$$

When  $\overline{\psi} = 1, v_0 \ll 1, \xi \ll 1$ , the bound offers a significant improvement over Chérief-Abdellatif and Alquier (2022, Corollary 3.4).

#### 6.2.2 Improved Confidence Bounds in the Parameter Space

In this subsection we still employ the Gaussian kernel  $k_{\gamma}$  (9). It is instructive to analyze how robust estimation with respect to the MMD distance translates into what happens in the space of parameters. In fact, since we have freedom over the lengthscale parameter, we may understandably want to select  $\gamma$  such that the distance in the space of parameters is kept small as well. **Definition 24 (Link function** F) Let  $\mathcal{M} = \{\mathbb{P}_{\theta} : \theta \in \Theta\}$  indexed by the parameter space  $\Theta \subset \mathcal{P}(\mathbb{R}^d)$ . We say that a non-decreasing function  $F_k : \mathbb{R}_+ \to \mathbb{R}_+$  is a link function for the model  $\mathcal{M}$  and kernel k when for any  $\theta, \theta' \in \Theta$ ,

$$\left\|\theta - \theta'\right\|_{2} \leq F_{k}\left(\left\|\mu_{\mathbb{P}_{\theta}} - \mu_{\mathbb{P}_{\theta'}}\right\|_{\mathcal{H}_{k}}\right).$$

**Remark 25** Note that a link function  $F_k$  will only provide nontrivial information if  $F_k(0) = 0$  and  $F_k$  is continuous at 0. The existence of such a nontrivial link function implies that the model is identifiable: if  $\mu_{\mathbb{P}_{\theta}} = \mu_{\mathbb{P}_{\theta'}}$  then  $0 = F_k(\|\mu_{\mathbb{P}_{\theta}} - \mu_{\mathbb{P}_{\theta'}}\|_{\mathcal{H}_k}) \geq \|\theta - \theta'\|_2$ , hence  $\theta = \theta'$ . In this case, there is a unique function  $p : \mathcal{M} \to \Theta$  such that  $p(\mathbb{P}_{\theta}) = \theta$ , and  $F_k$  is a modulus of continuity of p.

A direct application of Corollary 23 gives the following.

**Corollary 26** In the setting of Corollary 23, assume that  $\mathcal{M}$  and k are such that a link function  $F_k$  exists. With probability at least  $1 - \delta$ ,

$$\left\|\widehat{\theta}_n - \theta_0\right\|_2 \le F_k \left( 4\xi \sqrt{\overline{\psi}} + 2\sqrt{2[(1-2\xi)v_0 + 2\xi\overline{\psi}]\frac{\log(2/\delta)}{n}} + 2\frac{4\sqrt{\overline{\psi}}}{3}\frac{\log(2/\delta)}{n} \right).$$

Gaussian location model, continued. We continue here Example 1. A direct computation (see for instance Chérief-Abdellatif and Alquier 2022) shows that

$$\left\|\mu_{\mathbb{P}_{\theta}}-\mu_{\mathbb{P}_{\theta'}}\right\|_{\mathcal{H}_{k_{\gamma}}}^{2}=2\left(\frac{\gamma^{2}}{2\sigma^{2}+\gamma^{2}}\right)^{d/2}\left[1-\exp\left(-\frac{\left\|\theta'-\theta\right\|_{2}^{2}}{4\sigma^{2}+2\gamma^{2}}\right)\right].$$

In other words, a link function for the Gaussian location model is explicitly given by

$$F_{k_{\gamma}}^{2}(h) = -2(2\sigma^{2} + \gamma^{2})\log\left(1 - \frac{h}{2}\left(1 + 2\frac{\sigma^{2}}{\gamma^{2}}\right)^{d/2}\right).$$

An application of Corollary 26, and setting  $\gamma = \lambda \sigma \sqrt{d}$ , for some  $\lambda > 0$ , proves that, with probability at least  $1 - \delta$ ,

$$\left\| \theta - \widehat{\theta}_n \right\|_2^2 \leq -2\sigma^2 (2 + d\lambda^2) \log \left\{ 1 - \left( 1 + \frac{2}{d\lambda^2} \right)^{d/2} \left( 4\xi + \sqrt{2 \left[ (1 - 2\xi) \left( 1 - e^{-\frac{1}{\lambda^2}} \right) + 2\xi \right] \frac{\log(2/\delta)}{n}} + \frac{4}{3} \frac{\log(2/\delta)}{n} \right)^2 \right\}$$
(21)  
$$\doteq G_{d,\sigma,n}(\lambda,\xi).$$

For comparison, we provide the bound of Chérief-Abdellatif and Alquier (2022, Proposition 4.1-Equation 2) that is obtained by plugging (7) into the link function:

$$\left\| \theta - \widehat{\theta}_n \right\|_2^2 \le -2\sigma^2 (2 + d\lambda^2) \log\left\{ 1 - \left(1 + \frac{2}{d\lambda^2}\right)^{d/2} \left(4\xi + \frac{1 + \sqrt{2\log(1/\delta)}}{\sqrt{n}}\right)^2 \right\}$$
(22)  
$$\doteq H_{d,\sigma,n}(\lambda,\xi).$$



Figure 3: Comparison of the bounds in (21) versus Chérief-Abdellatif and Alquier (2022, Proposition 4.1-Equation 2) as a function of  $\gamma$ .



Figure 4: Comparison of Chérief-Abdellatif and Alquier (2022) and (21) for the optimal hyper-parameter  $\gamma$  as a function of n.

In order to compare the tightness of the bounds in Eq. (22) and Eq. (21), we run experiments for fixed sample size (n = 10000), fixed confidence level  $(\delta = 0.1)$ , fixed variance parameter  $(\sigma = 3)$  and two different contamination levels  $(\xi = 0.01 \text{ and } \xi = 0.2)$ . We compute and plot the bound obtained when varying the scale parameter  $\gamma$  (see Figure 3). For a contamination level  $\xi = 0.1$ , we also compare the quantitative decay of the two bounds—optimized for  $\gamma$ —as the sample size n grows and report our results in 4. Based on our experiments, we make the following observations.

- (i) The new bound is always tighter than the one of Chérief-Abdellatif and Alquier (2022).
- (ii) When the contamination level increases, the two bounds are getting closer together. This is expected since we are losing the benefit of the variance factor.
- (*iii*) Especially under weak contamination, the new bound is much flatter in the sense where overshooting for the value of  $\gamma$  does not lead to a catastrophic degradation of the bound.
- (*iv*) The new bound performs exceptionally well in the small sample setting.

We conclude this section by showing heuristically that the bound in (21) can always be made smaller than the bound in (22), at least asymptotically in n, with an adequate choice of  $\lambda$ . For the sake of simplicity, we work in the non-contaminated setting  $\xi = 0$ . First,

$$H_{d,\sigma,n}(\lambda,0) = 4\sigma^2 (2+d\lambda^2) \left(1+\frac{2}{d\lambda^2}\right)^{d/2} \frac{\left(1+\sqrt{2\log(1/\delta)}\right)^2}{n} (1+o(1)),$$

and the first order term is exactly minimized for  $\lambda = 1$ , it leads to

$$H_{d,\sigma,n}(1,0) = 4\sigma^2(2+d)\left(1+\frac{2}{d}\right)^{d/2}\frac{\left(1+\sqrt{2\log(1/\delta)}\right)^2}{n}(1+o(1)).$$

An exact minimization of the bound in (21) is not feasible, but we will propose a choice of  $\lambda = \lambda_n$  that will lead to an improvement on  $H_{d,\sigma,n}(1,0)$ . First, let us assume that  $\lambda_1 > 0$ and that  $(\lambda_n)_{n \in \mathbb{N}}$  is a non-decreasing sequence. We obtain:

$$\begin{aligned} G_{d,\sigma,n}(\lambda_n, 0) &= 4\sigma^2 (2 + d\lambda_n^2) \left( 1 + \frac{2}{d\lambda_n^2} \right)^{d/2} \left( 1 - e^{-\frac{1}{\lambda_n^2}} \right) \frac{\log(2/\delta)}{n} (1 + o(1)) \\ &= 4\sigma^2 (2 + d\lambda_n^2) \left( 1 + \frac{2}{d\lambda_n^2} \right)^{d/2} \frac{\log(2/\delta)}{\lambda_n^2 n} (1 + o(1)) \\ &= 4d\sigma^2 \left( 1 + \frac{2}{d\lambda_n^2} \right)^{1+d/2} \frac{\log(2/\delta)}{n} (1 + o(1)). \end{aligned}$$

It is clear that choosing  $\lambda_n \to \infty$  when  $n \to \infty$  gives:

$$G_{d,\sigma,n}(\lambda_n, 0) = 4d\sigma^2 \frac{\log(2/\delta)}{n} (1 + o(1)).$$



Figure 5: Comparison of the bounds in (21) versus Chérief-Abdellatif and Alquier (2022, Proposition 4.1-Equation 2) as a function of  $\gamma$ , versus the empirical MSE on 50 experiments, n = 500, log-scale,  $\xi = 0$ ,  $\delta = 0.05$ . Right panel: we zoom on the MSE.

Finally, observe that

$$\frac{H_{d,\sigma,n}(1,0)}{G_{d,\sigma,n}(\lambda_n,0)} = 2\underbrace{\left(1+\frac{2}{d}\right)^{d/2+1}}_{\underset{d\to\infty}{\overset{\searrow}{\longrightarrow}e}} \underbrace{\frac{(1+\sqrt{2\log 1/\delta})^2}{(\sqrt{2\log(2/\delta)})^2}}_{\ge 1 \text{ for } \delta \in (0,9/10)} (1+o(1)) \ge 2e(1+o(1))$$

Unfortunately, this first order analysis is too crude to provide an accurate recommendation on the choice of  $\lambda_n$ . On the other hand, it shows that our variance-aware bound leads to a significant improvement over the bound of Chérief-Abdellatif and Alquier (2022) for Gaussian mean estimation. Moreover, the choice  $\lambda_n \to \infty$  is in accordance with the considerations on the asymptotic variance in Briol et al. (2019) in this model.

We illustrate this with an experiment in Figure 5. We observe that: indeed, the varianceaware bound decrease with  $\gamma \propto \lambda$ , while the McDiarmid bound has a minimum. Interestingly, the true MSE also has a minimum. Note however how the MSE is actually very flat as a function of  $\gamma$ : in other words, in this experiment, the choice of  $\gamma$  does not matter so much on the performance of the MMD estimator. The possibility to use the variance-aware bound to calibrate  $\gamma$  in practice should be investigated further in other models.

### 7. Proofs

In this section, we provide the detailed technical proofs that were deferred in earlier parts of the paper.

### 7.1 Proof of Theorem 1

In Tolstikhin et al. (2017), the authors mention but do not pursue the idea that Bernstein's inequality in separable Hilbert spaces yields the proper rate of  $\mathcal{O}_P(n^{-1/2})$ . We follow-up on

their idea in this proof. First recall that

$$\widehat{\mu}_{\mathbb{P}} \doteq \frac{1}{n} \sum_{t=1}^{n} k(X_t, \cdot), \qquad \mu_{\mathbb{P}} \doteq \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X, \cdot) \right],$$

and that we can write

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} = \frac{1}{n} \left\| \sum_{t=1}^n Z_t \right\|_{\mathcal{H}_k},$$

where for  $1 \leq t \leq n$ ,

$$Z_t \doteq k(X_t, \cdot) - \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X, \cdot) \right],$$

is centered. For  $\varepsilon > 0$ , we are thus interested in bounding the probability of the following deviation,

$$\mathbb{P}\left(\max_{1\leq s\leq n}\left\|\sum_{t=1}^{s} Z_{t}\right\|_{\mathcal{H}_{k}} > n\varepsilon\right).$$

We rely on a Bernstein-type inequality in Hilbert spaces, credited to Pinelis and Sakhanenko (1985) (see also Yurinsky 1995, Theorem 3.3.4), which is reported in Theorem 32. We first observe that the RKHS  $\mathcal{H}_k$  is separable by separability of the topological space  $\mathcal{X}$  and continuity of k. In order to invoke Theorem 32, we need to control from above the  $\sum_{t=1}^{n} \mathbb{E} ||Z_t||_{\mathcal{H}_k}^p$ , for any  $p \geq 3$ . For p = 2,

$$\sum_{t=1}^{n} \mathbb{E} \left\| k(X_t, \cdot) - \mu_{\mathbb{P}} \right\|_{\mathcal{H}_k}^2 = nv.$$

Moving on to higher moments, for  $p \ge 3$ , we have that

$$\begin{aligned} \|Z_t\|_{\mathcal{H}_k} &\leq \|k(X_t, \cdot)\|_{\mathcal{H}_k} + \|\mathbb{E}_{X \sim \mathbb{P}} \left[k(X, \cdot)\right]\|_{\mathcal{H}_k} \\ &\leq \sqrt{k(X_t, X_t)} + \sup_{x \in \mathcal{X}} \sqrt{k(x, x)} \\ &\leq 2\sqrt{\overline{k}}. \end{aligned}$$

As a result,

$$\sum_{t=1}^{n} \mathbb{E} \|Z_t\|_{\mathcal{H}_k}^p \leq \left(\sum_{t=1}^{n} \mathbb{E} \|Z_t\|_{\mathcal{H}_k}^2\right) \left(2\sqrt{\overline{k}}\right)^{p-2}$$
$$\leq nv(2\sqrt{\overline{k}})^{p-2} \leq \frac{p!}{2}nv\left(\frac{2\sqrt{\overline{k}}}{3}\right)^{p-2}$$

Invoking Theorem 32 with  $G^2 = nv$ ,  $H = 2\sqrt{\overline{k}}/3$ , we obtain that for any  $\varepsilon > 0$ ,

$$\mathbb{P}\left(\max_{1\leq s\leq n}\left\|\sum_{t=1}^{s} Z_{t}\right\|_{\mathcal{H}_{k}} > n\varepsilon\right) \leq 2\exp\left(-\frac{(\varepsilon n)^{2}}{2(nv+(\varepsilon n)2\sqrt{\overline{k}}/3)}\right)$$
$$= 2\exp\left(-\frac{n\varepsilon^{2}}{2\left(v+2\varepsilon\sqrt{\overline{k}}/3\right)}\right).$$

# 7.2 Proof of Lemma 10

We first verify the bounded-differences property.

$$\begin{split} \widehat{v}(x) &- \inf_{x_{t_0}' \in \mathcal{X}} \widehat{v} \left( x^{(t_0)} \right) \\ &= \overline{\psi} - \frac{1}{(n-1)n} \sum_{t \neq s} \psi(x_t - x_s) - \inf_{x_{t_0}' \in \mathcal{X}} \left[ \overline{\psi} - \frac{1}{(n-1)n} \sum_{t \neq s} \psi \left( x_t^{(t_0)} - x_s^{(t_0)} \right) \right] \\ &= \frac{1}{(n-1)n} \left( \sup_{x_{t_0}' \in \mathcal{X}} \left\{ \sum_{t \neq s} \psi \left( x_t^{(t_0)} - x_s^{(t_0)} \right) \right\} - \sum_{t \neq s} \psi(x_t - x_s) \right) \\ &= \frac{2}{(n-1)n} \left( \sup_{x_{t_0}' \in \mathcal{X}} \left\{ \sum_{t \in [n]} \psi \left( x_{t_0}' - x_t \right) \right\} - \sum_{t \in [n]} \psi(x_{t_0} - x_t) \right) \\ &\leq \frac{2}{(n-1)n} \left( (n-1)\overline{\psi} - \sum_{\substack{t \in [n] \\ t \neq t_0}} \psi(x_{t_0} - x_t) \right) \\ &\leq \frac{2}{n} \left( \overline{\psi} - \frac{1}{n-1} \sum_{\substack{t \in [n] \\ t \neq t_0}} \psi(x_{t_0} - x_t) \right) \end{split}$$

It follows from  $(\star)$  that,

$$\begin{split} \sum_{t_0=1}^n \left( \widehat{v}(x) - \inf_{x_{t_0}' \in \mathcal{X}} \widehat{v}\left(x^{(t_0)}\right) \right)^2 &\leq \sum_{t_0=1}^n \left( \frac{2}{n} \left( \overline{\psi} - \frac{1}{n-1} \sum_{\substack{t \in [n] \\ t \neq t_0}} \psi(x_{t_0} - x_t) \right) \right)^2 \\ &= \frac{4}{n^2} \sum_{t_0=1}^n \left( \frac{1}{n-1} \sum_{\substack{t \in [n] \\ t \neq t_0}} \left( \overline{\psi} - \psi(x_{t_0} - x_t) \right) \right)^2 \\ &\leq \frac{4}{n^2} \sum_{t_0=1}^n \frac{1}{n-1} \sum_{\substack{t \in [n] \\ t \neq t_0}} \left( \overline{\psi} - \psi(x_{t_0} - x_t) \right)^2 \\ &\leq \frac{4\Delta\psi}{n} \sum_{t_0=1}^n \frac{1}{(n-1)n} \sum_{\substack{t \in [n] \\ t \neq t_0}} \left( \overline{\psi} - \psi(x_{t_0} - x_t) \right) \\ &= \frac{4\Delta\psi}{n} \widehat{v} \end{split}$$

where the second inequality is Jensen's.

# 7.3 Proof of Lemma 11

From Lemma 10, the function

$$\mathcal{X}^n \to \mathbb{R}, x \mapsto \frac{n}{2\Delta\psi}\widehat{v}(x),$$

is weakly (2,0)-self-bounding and has the bounded differences property in the sense of Definition 7. From the second statement of Maurer (2006, Theorem 13),

$$\mathbb{P}\left(\mathbb{E}\left[\frac{n}{2\Delta\psi}\widehat{v}(X)\right] - \frac{n}{2\Delta\psi}\widehat{v}(X) > t\right) \le \exp\left(-\frac{t^2}{4\mathbb{E}\left[\frac{n}{2\Delta\psi}\widehat{v}(X)\right]}\right),$$

that can be rewritten,

$$\mathbb{P}\left(v - \hat{v}(X) > t\right) \le \exp\left(-\frac{nt^2}{8\Delta\psi v}\right).$$

In other words, with probability at least  $1 - \delta$ ,

$$v - 2\sqrt{v}\sqrt{\frac{2\Delta\psi\log(1/\delta)}{n}} \le \widehat{v}(X).$$

Completing the square, and by sub-additivity of the square root function, we successively have with probability  $1 - \delta$ ,

$$\left(\sqrt{v} - \sqrt{\frac{2\Delta\psi\log(1/\delta)}{n}}\right)^2 \le \widehat{v}(X) + \frac{2\Delta\psi\log(1/\delta)}{n}$$

$$\begin{split} \sqrt{v} - \sqrt{\frac{2\Delta\psi\log(1/\delta)}{n}} &\leq \sqrt{\widehat{v}(X) + \frac{2\Delta\psi\log(1/\delta)}{n}} \\ \sqrt{v} &\leq \sqrt{\widehat{v}(X)} + 2\sqrt{\frac{2\Delta\psi\log(1/\delta)}{n}}, \end{split}$$

which proves the lemma for b = -1. On the other hand, from the first inequality in Maurer (2006, Theorem 13),

$$\mathbb{P}\left(\frac{n}{2\Delta\psi}\widehat{v}(X) - \mathbb{E}\left[\frac{n}{2\Delta\psi}\widehat{v}(X)\right] > t\right) \le \exp\left(-\frac{t^2}{4\mathbb{E}\left[\frac{n}{2\Delta\psi}\widehat{v}(X)\right] + 2t}\right),$$

that can be rewritten as

$$\mathbb{P}\left(\widehat{v}(X) - v > t\right) \le \exp\left(-\frac{nt^2}{8\Delta\psi(v + t/2)}\right).$$

The positive solution  $t_+$  for

$$nt^2 = 8\Delta\psi(v+t/2)\log(1/\delta)$$

is readily given by

$$t_{+} = \frac{2\Delta\psi\log(1/\delta)}{n} + \frac{4}{n}\sqrt{\left(\frac{\Delta\psi}{2}\log(1/\delta)\right)\left(\frac{\Delta\psi}{2}\log(1/\delta) + nv\right)}.$$

Thus, it holds with probability at least  $1 - \delta$  that

 $\widehat{v}(X) \le v + t_+,$ 

and from sub-additivity of the square root, successively, with probability at least  $1 - \delta$ , we have

$$\begin{split} \widehat{v}(X) &\leq v + \frac{4\Delta\psi\log(1/\delta)}{n} + 2\sqrt{\frac{2v\Delta\psi\log(1/\delta)}{n}},\\ \widehat{v}(X) &\leq \left(\sqrt{v} + \sqrt{\frac{2\Delta\psi\log(1/\delta)}{n}}\right)^2 + \frac{2\Delta\psi\log(1/\delta)}{n},\\ \sqrt{\widehat{v}(X)} &\leq \sqrt{\left(\sqrt{v} + \sqrt{\frac{2\Delta\psi\log(1/\delta)}{n}}\right)^2 + \frac{2\Delta\psi\log(1/\delta)}{n},\\ \sqrt{\widehat{v}(X)} &\leq \sqrt{v} + 2\sqrt{\frac{2\Delta\psi\log(1/\delta)}{n}}, \end{split}$$

which finishes proving the lemma<sup>10</sup> for b = 1.

<sup>10.</sup> In this proof, Maurer (2006) is sufficient. We do not need the stronger results of Boucheron et al. (2009).

# 7.4 Proof of Theorem 12

From Theorem 1, with probability at least  $1 - \delta/2$ , it holds that

$$\|\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) - \mu_{\mathbb{P}}\|_{\mathcal{H}_{\psi}} \le \sqrt{2v \frac{\log(4/\delta)}{n}} + \frac{4}{3}\sqrt{\psi} \frac{\log(4/\delta)}{n}.$$
 (23)

Invoking Lemma 11 for b = -1, at confidence  $1 - \delta/2$  it holds that

$$\sqrt{v} \le \sqrt{\widehat{v}(X)} + 2\sqrt{\frac{2\Delta\psi\log(2/\delta)}{n}} \le \sqrt{\widehat{v}(X)} + 2\sqrt{\frac{2\Delta\psi\log(4/\delta)}{n}},\tag{24}$$

thus by combining (23) and (24) and a union bound yields that with probability at least  $1-\delta$ ,

$$\begin{aligned} \|\widehat{\mu}_{\mathbb{P}}(X_{1},\ldots,X_{n})-\mu_{\mathbb{P}}\|_{\mathcal{H}_{\psi}} \\ &\leq \left(\sqrt{\widehat{v}(X)}+2\sqrt{\frac{2\Delta\psi\log(4/\delta)}{n}}\right)\sqrt{\frac{2\log(4/\delta)}{n}}+\frac{4}{3}\sqrt{\overline{\psi}}\frac{\log(4/\delta)}{n} \\ &=\sqrt{2\widehat{v}(X)\frac{\log(4/\delta)}{n}}+4\left(\sqrt{\Delta\psi}+\frac{\sqrt{\overline{\psi}}}{3}\right)\frac{\log(4/\delta)}{n} \\ &\leq \sqrt{2\widehat{v}(X)\frac{\log(4/\delta)}{n}}+\frac{16}{3}\sqrt{\Delta\psi}\frac{\log(4/\delta)}{n}. \end{aligned}$$
(25)

# 7.5 Proof of Lemma 16

$$\mathbb{E} \|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{k}}^{2} = \mathbb{E} \langle \frac{1}{n} \sum_{t=1}^{n} k(X_{t}, \cdot) - \mu_{\mathbb{P}}, \frac{1}{n} \sum_{s=1}^{n} k(X_{s}, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}}$$

$$= \frac{1}{n^{2}} \sum_{s=1}^{n} \sum_{t=1}^{n} \mathbb{E} \langle k(X_{t}, \cdot) - \mu_{\mathbb{P}}, k(X_{s}, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}}$$

$$= \frac{2}{n^{2}} \sum_{t=1}^{n} \sum_{s

$$+ \frac{1}{n^{2}} \sum_{t=1}^{n} \mathbb{E} \langle k(X_{t}, \cdot) - \mu_{\mathbb{P}}, k(X_{t}, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}}$$

$$= \frac{2}{n^{2}} \sum_{t=2}^{n} (n-t+1) \mathbb{E} \langle k(X_{t}, \cdot) - \mu_{\mathbb{P}}, k(X_{1}, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}} + \frac{v}{n}.$$$$

### 7.6 Proof of Theorem 17

An application of Jensen's inequality and Lemma 16 immediately yield that

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \le \sqrt{\frac{v + \Sigma_n}{n}},$$

where  $\Sigma_n$  is a total measure of covariance in the RKHS introduced in Lemma 16. Additionally, recall that we can express the norm of the deviation of  $\hat{\mu}_{\mathbb{P}}$  from its mean  $\mu_{\mathbb{P}}$  in the RKHS as a supremum of empirical processes as follows,

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{k}} = \sup_{\substack{f \in \mathcal{H}_{k} \\ \|f\|_{\mathcal{H}_{k}} \leq 1}} \langle f, \widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}} = \frac{1}{n} \sup_{\substack{f \in \mathcal{H}_{k} \\ \|f\|_{\mathcal{H}_{k}} \leq 1}} \sum_{t=1}^{n} \left( f(X_{t}) - \mathbb{E}f \right) = \frac{1}{n} \sup_{g \in \mathcal{F}} \sum_{t=1}^{n} g(X_{t}),$$

where  $\mathcal{F}$  is the re-centered class,

$$\mathcal{F} \doteq \left\{ g \colon \mathcal{X} \to \mathbb{R}, g = f - \mathbb{E}f, f \in \mathcal{H}_k, \|f\|_{\mathcal{H}_k} \le 1 \right\}$$

In particular, observe that  $g \in \mathcal{F} \implies |g| \leq \overline{k}$ . Furthermore, the variance v in the RKHS can be expressed as

$$v = \mathbb{E} \left\| k(X, \cdot) - \mu_{\mathbb{P}} \right\|_{\mathcal{H}_{k}}^{2}$$
$$= \mathbb{E} \left( \sup_{\substack{f \in \mathcal{H}_{k} \\ \|f\|_{\mathcal{H}_{k}} \leq 1}} \langle f, k(X, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}} \right)^{2}$$
$$= \mathbb{E} \sup_{\substack{f \in \mathcal{H}_{k} \\ \|f\|_{\mathcal{H}_{k}} \leq 1}} \langle f, k(X, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}}^{2}$$
$$= \mathbb{E} \sup_{\substack{f \in \mathcal{H}_{k} \\ \|f\|_{\mathcal{H}_{k}} \leq 1}} (f(X) - \mathbb{E}f)^{2}.$$

We invoke Theorem 33 for  $\phi$ -dependent random sequences, with  $Z = n \|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k}$  and observing that

$$\mathbb{E}\left[V\right] = \sum_{t=1}^{n} \mathbb{E}\sup_{g \in \mathcal{F}} g(X_t)^2 = nv,$$

with probability at least  $1 - \delta$  it holds that

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \le \sqrt{\frac{v + \Sigma_n}{n}} + 4\sqrt{\frac{2v \|\Gamma\|_2 \log(1/\delta)}{n}} + \frac{8\overline{k} \|\Gamma\|_2 \log(1/\delta)}{n}.$$

# 7.7 Proof of Theorem 19

First, assume that  $\mathbf{H}_0$  is true, that is: there is a  $\theta_0 \in \Theta$  such that  $\mathbb{P} = \mathbb{P}_{\theta_0}$ . Then,

$$\mathbb{P}_{\mathbf{H}_{0}}(\mathcal{C}) = \mathbb{P}_{\mathbf{H}_{0}}(T > \mathcal{B}(X_{1}, \dots, X_{n}, \alpha))$$
$$= \mathbb{P}_{\mathbf{H}_{0}}\left(\inf_{\theta \in \Theta} \|\widehat{\mu_{\mathbb{P}}} - \mu_{\mathbb{P}_{\theta}}\|_{\mathcal{H}_{k}} > \mathcal{B}(X_{1}, \dots, X_{n}, \alpha)\right)$$
$$\leq \mathbb{P}_{\mathbf{H}_{0}}\left(\left\|\widehat{\mu_{\mathbb{P}}} - \mu_{\mathbb{P}_{\theta_{0}}}\right\|_{\mathcal{H}_{k}} > \mathcal{B}(X_{1}, \dots, X_{n}, \alpha)\right)$$
$$\leq \alpha$$

thanks to (a).

Now, assume that  $\mathbf{H}_0$  is not true, that is  $\mathbb{P} \notin \{\mathbb{P}_{\theta}, \theta \in \Theta\}$ . Note that under the assumption that the model is closed, we have

$$0 < \Delta := \inf_{\theta \in \Theta} \|\mu_{\mathbb{P}} - \mu_{\mathbb{P}_{\theta}}\|_{\mathcal{H}_{k}}.$$

By the triangle inequality, for any  $\theta$ ,

$$\|\mu_{\mathbb{P}_{\theta}} - \widehat{\mu}_{\mathbb{P}}\|_{\mathcal{H}_{k}} \geq \|\mu_{\mathbb{P}_{\theta}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{k}} - \|\mu_{\mathbb{P}} - \widehat{\mu}_{\mathbb{P}}\|_{\mathcal{H}_{k}}$$

and thus, taking the infimum w.r.t  $\theta$  on both sides,

$$T \ge \Delta - \|\mu_{\mathbb{P}} - \widehat{\mu}_{\mathbb{P}}\|_{\mathcal{H}_k}.$$

We apply (a) with confidence level  $\alpha_n$  and obtain

$$\mathbb{P}_{\mathbf{H}_1}\left(\|\mu_{\mathbb{P}} - \widehat{\mu}_{\mathbb{P}}\|_{\mathcal{H}_k} \ge \mathcal{B}(X_1, \dots, X_n, \alpha_n)\right) \le \alpha_n$$

and thus

$$\mathbb{P}_{\mathbf{H}_1} \left( T \leq \Delta - \mathcal{B}(X_1, \dots, X_n, \alpha_n) \right) \leq \alpha_n.$$

As  $\mathcal{B}(X_1, \ldots, X_n, \alpha_n) \to 0$  when  $n \to \infty$  thanks to (c), there is a N large enough such that, for any  $n \ge N$ ,  $\mathcal{B}(X_1, \ldots, X_n, \alpha_n) \le \Delta/2$ , and thus

$$\mathbb{P}_{\mathbf{H}_1}\left(T \le \frac{\Delta}{2}\right) \le \alpha_n.$$

Moreover, as we also have  $\mathcal{B}(X_1, \ldots, X_n, \alpha_n) \to 0$  when  $n \to \infty$  thanks to (b), there is a N' large enough such that, for any  $n \ge N'$ ,  $\mathcal{B}(n, \alpha) \le \Delta/2$  and thus

$$\mathbb{P}_{\mathbf{H}_1} \left( T \leq \mathcal{B}(X_1, \dots, X_n, \alpha_n) \right) \leq \alpha_n.$$

Finally,

$$\mathbb{P}_{\mathbf{H}_1}(\mathcal{C}) = 1 - \mathbb{P}_{\mathbf{H}_1}(T \le \mathcal{B}(X_1, \dots, X_n, \alpha_n)) \ge 1 - \alpha_n$$

as soon as  $n \ge \max(N, N')$ .

# 7.8 Proof of Theorem 20

First, under  $\mathbf{H}_0$ , we have  $\mathbb{P}_X = \mathbb{P}_Y$  and thus

$$\mathbb{P}_{\mathbf{H}_{0}}(\mathcal{C}_{2}) = \mathbb{P}_{\mathbf{H}_{0}}(T_{2} > \mathcal{B}(X_{1}, \dots, X_{n}, \alpha/2) + \mathcal{B}(Y_{1}, \dots, Y_{m}, \alpha/2))$$

$$= \mathbb{P}_{\mathbf{H}_{0}}\left(\|\widehat{\mu}_{\mathbb{P}}(X_{1}, \dots, X_{n}) - \widehat{\mu}_{\mathbb{P}}(Y_{1}, \dots, Y_{m})\|_{\mathcal{H}_{k}}\right)$$

$$> \mathcal{B}(X_{1}, \dots, X_{n}, \alpha/2) + \mathcal{B}(Y_{1}, \dots, Y_{m}, \alpha/2)$$

$$= \mathbb{P}_{\mathbf{H}_{0}}\left(\|\widehat{\mu}_{\mathbb{P}}(X_{1}, \dots, X_{n}) - \mu_{\mathbb{P}_{X}} + \mu_{\mathbb{P}_{X}} - \widehat{\mu}_{\mathbb{P}}(Y_{1}, \dots, Y_{m})\|_{\mathcal{H}_{k}}\right)$$

$$> \mathcal{B}(X_{1}, \dots, X_{n}, \alpha/2) + \mathcal{B}(Y_{1}, \dots, Y_{m}, \alpha/2)$$

$$\leq \mathbb{P}_{\mathbf{H}_{0}}\left(\|\widehat{\mu}_{\mathbb{P}}(X_{1}, \dots, X_{n}) - \mu_{\mathbb{P}_{X}}\|_{\mathcal{H}_{k}} > \mathcal{B}(X_{1}, \dots, X_{n}, \alpha/2)\right)$$

$$+ \mathbb{P}_{\mathbf{H}_{0}}\left(\|\widehat{\mu}_{\mathbb{P}}(Y_{1}, \dots, Y_{m}) - \mu_{\mathbb{P}_{X}}\|_{\mathcal{H}_{k}} > \mathcal{B}(Y_{1}, \dots, Y_{m}, \alpha/2)\right)$$

$$\leq \alpha/2 + \alpha/2 = \alpha.$$

thanks to (a).

Now, assume that  $\mathbf{H}_0$  is not true, that is  $\mathbb{P}_X \neq \mathbb{P}_Y$ , and put

$$\Delta := \|\mu_{\mathbb{P}_X} - \mu_{\mathbb{P}_Y}\|_{\mathcal{H}_k} > 0.$$

By the triangle inequality,

$$T_{2} = \|\widehat{\mu_{\mathbb{P}}}(X_{1},\ldots,X_{n}) - \widehat{\mu_{\mathbb{P}}}(Y_{1},\ldots,Y_{m})\|_{\mathcal{H}_{k}}$$
  
$$\geq \Delta - \|\widehat{\mu_{\mathbb{P}}}(X_{1},\ldots,X_{n}) - \mu_{\mathbb{P}_{X}}\|_{\mathcal{H}_{k}} - \|\widehat{\mu_{\mathbb{P}}}(Y_{1},\ldots,Y_{n}) - \mu_{\mathbb{P}_{Y}}\|_{\mathcal{H}_{k}}.$$

We apply (a) with respectice confidence levels  $\alpha_n$  and  $\alpha_m$  to get

$$\mathbb{P}_{\mathbf{H}_1}\left(\left\|\mu_{\mathbb{P}_X} - \widehat{\mu}_{\mathbb{P}}(X_1, \dots, X_n)\right\|_{\mathcal{H}_k} \ge \mathcal{B}(X_1, \dots, X_n, \alpha_n)\right) \le \alpha_n$$

and

$$\mathbb{P}_{\mathbf{H}_1}\left(\|\mu_{\mathbb{P}_Y} - \widehat{\mu}_{\mathbb{P}}(Y_1, \dots, Y_m)\|_{\mathcal{H}_k} \ge \mathcal{B}(Y_1, \dots, Y_m, \alpha_m)\right) \le \alpha_m.$$

Thus,

$$\mathbb{P}_{\mathbf{H}_1}\left(T_2 \leq \Delta - \mathcal{B}(X_1, \dots, X_n, \alpha_n) - \mathcal{B}(Y_1, \dots, Y_m, \alpha_m)\right) \leq \alpha_n + \alpha_m.$$

Thanks to (c), there is a N large enough such that, as soon as both  $n, m \geq N$ , both  $\mathcal{B}(X_1, \ldots, X_n, \alpha_n) \leq \Delta/4$  and  $\mathcal{B}(Y_1, \ldots, Y_m, \alpha_m) \leq \Delta/4$  hold, and thus

$$\mathbb{P}_{\mathbf{H}_1}\left(T \le \frac{\Delta}{2}\right) \le \alpha_n + \alpha_m.$$

Moreover, thanks to (b), there is N' large enough such that, when  $n, m \geq N'$ , both  $\mathcal{B}(X_1, \ldots, X_n, \alpha/2) \leq \Delta/4$  and  $\mathcal{B}(Y_1, \ldots, Y_m, \alpha/2) \leq \Delta/4$  hold, and thus

$$\mathbb{P}_{\mathbf{H}_1}\left(T_2 \leq \mathcal{B}(X_1, \dots, X_n, \alpha/2) + \mathcal{B}(Y_1, \dots, Y_m, \alpha/2)\right) \leq \alpha_n + \alpha_m.$$

Finally,

$$\mathbb{P}_{\mathbf{H}_1}(\mathcal{C}_2) = 1 - \mathbb{P}_{\mathbf{H}_1}(T_2 \leq \mathcal{B}(X_1, \dots, X_n, \alpha/2) + \mathcal{B}(Y_1, \dots, Y_m, \alpha/2)) \geq 1 - \alpha_n - \alpha_m$$

as soon as  $n \ge \max(N, N')$ , and thus

$$\mathbb{P}_{\mathbf{H}_{1}}\left(\mathcal{C}_{2}\right) \geq 1 - \alpha_{n} - \alpha_{m} \xrightarrow[n,m \to \infty]{} 1.$$

# 7.9 Proof of Lemma 21

$$\begin{aligned} v(\mathbb{P}) &= \mathbb{E}_{X \sim \mathbb{P}} \| k(X, \cdot) - \mu_{\mathbb{P}} \|_{\mathcal{H}_{k}}^{2} \\ &= \overline{\psi} - \int \int \psi(x' - x) d(\mathbb{P} \times \mathbb{P})(x, x') \\ &= \overline{\psi} - \int \int \psi(x' - x) d\mathbb{P}(x) d\mathbb{P}(x') \\ &= \overline{\psi} - \int \int \psi(x' - x) d((1 - \xi) \mathbb{P}_{\theta_{0}}(x) + \xi \mathbb{H}(x)) d((1 - \xi) \mathbb{P}_{\theta_{0}}(x') + \xi \mathbb{H}(x')) \\ &= \overline{\psi} - (1 - \xi)^{2} \int \int \psi(x' - x) d\mathbb{P}_{\theta_{0}}(x) d\mathbb{P}_{\theta_{0}}(x') \\ &\quad - 2\xi(1 - \xi) \int \int \psi(x' - x) d\mathbb{P}_{\theta_{0}}(x) d\mathbb{H}(x') \\ &\quad - \xi^{2} \int \int \psi(x' - x) d\mathbb{H}(x) d\mathbb{H}(x') \\ &\leq \overline{\psi} - (1 - \xi)^{2} (\overline{\psi} - \psi) - 2\xi(1 - \xi) \underline{\psi} - \xi^{2} \underline{\psi} \\ &= \overline{\psi} - (1 - 2\xi + \xi^{2}) (\overline{\psi} - \psi) - \xi(2 - \xi) \underline{\psi} \\ &= \overline{\psi} - (\overline{\psi} - \psi) + 2\xi(\overline{\psi} - \psi) - \xi^{2} (\overline{\psi} - \psi) - \xi(2 - \xi) \underline{\psi} \\ &\leq (1 - 2\xi) \psi + 2\xi \overline{\psi} - \xi(2 - \xi) \underline{\psi}, \\ &= (1 - 2\xi) \psi + 2\xi \Delta \psi + \xi^{2} \underline{\psi}, \end{aligned}$$

and the lemma follows by disjunction of cases.

## Acknowledgments

The authors declare that there is no conflict of interest. Part of this research was conducted when GW was supported by the Special Postdoctoral Researcher Program (SPDR) of RIKEN. Part of the research in this paper was done when PA was working at RIKEN AIP. We thank Bastien Dussap for pointing out a missing factor two in an earlier version of this manuscript. We thank the anonymous reviewers for their useful comments. In particular, the comparison to the approach based on Bernstein's inequality for U-statistics in Section 4 was suggested by Reviewer 1 and the analysis for time-dependent processes in Section 5 was suggested by Reviewer 2. All the remaining mistakes are ours.

### Appendix A. Extension to Non Translation Invariant Kernels

Let us define diag  $k: \mathcal{X} \to \mathbb{R}$  by diag k(x) = k(x, x), and use the shorthands diag k, diag k and  $\Delta$  diag k introduced in (2) accordingly. We show that the properties obtained in Section 3 can be extended to non TI kernel when  $\Delta$  diag k is controlled from above.

**Lemma 27** Let k be a characteristic reproducing kernel. The function

$$\mathcal{X}^n \to \mathbb{R}, x \mapsto \frac{n}{\Delta \operatorname{diag} k + 2\Delta k} \widehat{v}(x),$$

is weakly

$$\left(2,2n\frac{\Delta\operatorname{diag}k}{\Delta\operatorname{diag}k+2\Delta k}\right)$$
-self-bounding,

and has bounded differences in the sense of Definition 7.

**Lemma 28** For  $b \in \{-1, 1\}$ , with probability at least  $1 - \delta$ ,

$$b\left[\sqrt{\hat{v}(X) + \Delta \operatorname{diag} k} - \sqrt{v + \Delta \operatorname{diag} k}\right] \le 2\sqrt{\frac{(\Delta \operatorname{diag} k + 2\Delta k)\log(1/\delta)}{n}}.$$
 (26)

Attempting to recover concentration of  $\sqrt{\hat{v}}$  around  $\sqrt{v}$  with self-boundedness leads to an additional term in  $\mathcal{O}\left(n^{-1/4}\right)$  in (26) when  $\Delta \operatorname{diag} k \neq 0$ . We thus settle for concentration of  $\sqrt{\hat{v} + \Delta \operatorname{diag} k}$  around  $\sqrt{v + \Delta \operatorname{diag} k}$  instead.

**Theorem 29 (Confidence interval with empirical variance for general kernel)** Let  $n \in \mathbb{N}, X_1, \ldots, X_n \sim \mathbb{P}$ , let k be a characteristic reproducing kernel defined from a positive definite function  $\psi$  [see (4)]. Then with probability at least  $1 - \delta$ , it holds that

$$\|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \leq \mathcal{B}_{k,\delta}(X_1, \dots, X_n),$$

with

$$\widehat{\mathcal{B}}_{k,\delta}(X_1,\ldots,X_n) = \widehat{\mathcal{B}}_{\delta} \doteq \sqrt{2(\widehat{v}(X_1,\ldots,X_n) + \Delta \operatorname{diag} k) \frac{\log(4/\delta)}{n}} + \left(\frac{16}{3}\sqrt{\Delta k} + 2\sqrt{2}\sqrt{\Delta \operatorname{diag} k}\right) \frac{\log(4/\delta)}{n},$$

and where  $\hat{v}$  is the empirical variance proxy defined in (12).

**Remark 30** In particular, for a TI kernel  $\Delta \operatorname{diag} k = 0$ , recovering the results of the previous section. It is currently unclear whether the term in  $\Delta \operatorname{diag} k$  is necessary or if it is an artifact of our proof.

# Appendix B. Additional Proofs

In this section, we present supplementary technical proofs that were postponed in previous sections of the paper

# B.1 Proof of Lemma 8

For some  $t_0$ , we let  $x^{(t_0)} = (x_1, \ldots, x_{t_0-1}, x'_{t_0}, x_{t_0+1}, \ldots, x_n)$  where  $x_{t_0}$  was replaced with  $x'_{t_0}$ .

$$\frac{1}{d}\operatorname{Tr}\widehat{\Sigma}(x) - \frac{1}{d}\operatorname{Tr}\widehat{\Sigma}\left(x^{(t_0)}\right) \\
= \frac{1}{d}\sum_{i=1}^{d}\frac{1}{2n(n-1)}\left(\sum_{s=1}^{n}\sum_{t=1}^{n}\left(x_t^i - x_s^i\right)^2 - \sum_{s=1}^{n}\sum_{t=1}^{n}\left(x_t^{(t_0)i} - x_s^{(t_0)i}\right)^2\right) \\
\leq \frac{1}{d}\sum_{i=1}^{d}\left(\frac{1}{n-1}\sum_{s=1}^{n}\left(x_{t_0}^i - x_s^i\right)^2\right) \leq 1,$$
(27)

thus

$$\frac{1}{d}\operatorname{Tr}\widehat{\Sigma}(x) - \inf_{x_{t_0}' \in \mathcal{X}} \frac{1}{d}\operatorname{Tr}\widehat{\Sigma}\left(x^{(t_0)}\right) \leq 1.$$

Furthermore,

$$\sum_{t_0=1}^{n} \left( \frac{1}{d} \operatorname{Tr} \widehat{\Sigma}(x) - \frac{1}{d} \operatorname{Tr} \widehat{\Sigma}\left(x^{(t_0)}\right) \right)^2 \stackrel{(i)}{\leq} \sum_{t_0=1}^{n} \left( \frac{1}{d} \sum_{i=1}^{d} \left( \frac{1}{n-1} \sum_{s=1}^{n} \left( x_{t_0}^i - x_s^i \right)^2 \right) \right)^2$$
$$\stackrel{(ii)}{\leq} \frac{1}{d} \sum_{i=1}^{d} \sum_{t_0=1}^{n} \left( \frac{1}{n-1} \sum_{s=1}^{n} \left( x_{t_0}^i - x_s^i \right)^2 \right)^2$$
$$\stackrel{(iii)}{\leq} \frac{1}{d} \sum_{i=1}^{d} \frac{n}{n-1} \widehat{v}^i(X)$$
$$= \frac{n}{n-1} \frac{1}{d} \operatorname{Tr} \widehat{\Sigma}(X),$$

where (i) follows from (27), (ii) is Jensen's inequality, and (iii) stems from Maurer and Pontil (2009, Corollary 9).

# B.2 Proof of Lemma 9

We verify that  $\hat{v}$  is an unbiased estimator for v.

$$\begin{split} & \mathbb{E}_{X_1,\dots,X_n \sim \mathbb{P}} \left[ \widehat{v}(X_1,\dots,X_n) \right] \\ = & \mathbb{E}_{X_1,\dots,X_n \sim \mathbb{P}} \left[ \frac{1}{n-1} \sum_{t=1}^n \left( k(X_t,X_t) - \frac{1}{n} \sum_{s=1}^n k(X_t,X_s) \right) \right] \\ & = \frac{1}{n-1} \sum_{t=1}^n \left( \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X,X) \right] - \frac{1}{n} \sum_{s=1}^n \mathbb{E}_{X_t,X_s \sim \mathbb{P}} \left[ k(X_t,X_s) \right] \right) \\ & = \frac{n}{n-1} \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X,X) \right] - \frac{1}{n(n-1)} \sum_{t=1}^n \sum_{s=1}^n \mathbb{E}_{X_t,X_s \sim \mathbb{P}} \left[ k(X_t,X_s) \right] \\ & = \frac{n}{n-1} \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X,X) \right] - \frac{1}{(n-1)} \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X,X) \right] \\ & \quad - \frac{n^2 - n}{n(n-1)} \mathbb{E}_{X,X' \sim \mathbb{P}} \left[ k(X,X') \right] \\ & = \mathbb{E}_{X \sim \mathbb{P}} \left[ k(X,X) \right] - \mathbb{E}_{X,X' \sim \mathbb{P}} \left[ k(X,X') \right] \\ & = v. \end{split}$$

## B.3 Proof of Theorem 18

For  $t \in \mathbb{N}$ , we write  $Z_t \doteq k(X_t, \cdot) - \mu_{\mathbb{P}}$ , thus

$$\|\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n)-\mu_{\mathbb{P}}\|_{\mathcal{H}_k}=\frac{1}{n}\left\|\sum_{t=1}^n Z_t\right\|_{\mathcal{H}_k}.$$

We assume that n = 2Bs, where  $B, s \in \mathbb{N}$  will be determined later. For  $b \in [B]$ , we denote

$$Z^{[2b]} \doteq Z_1^{[2b]}, \dots, Z_s^{[2b]} \doteq Z_{(2b-1)s+1}, \dots, Z_{2bs},$$
$$Z^{[2b-1]} \doteq Z_1^{[2b-1]}, \dots, Z_s^{[2b-1]} \doteq Z_{(2b-2)s+1}, \dots, Z_{(2b-1)s},$$

and we decompose the random process into blocks

$$Z^{[1]}, Z^{[2]}, \dots, Z^{[2B]}.$$

From sub-additivity of the norm and the blocking method described in Yu (1994, Corollary 2.7), it holds that

$$\mathbb{P}\left(\left\|\sum_{t=1}^{n} Z_{t}\right\|_{\mathcal{H}_{k}} > n\varepsilon\right) \leq \sum_{\sigma \in \{0,1\}} \mathbb{P}\left(\left\|\sum_{b=1}^{B} \left(\sum_{t=1}^{s} Z_{t}^{[2b-\sigma]}\right)\right\|_{\mathcal{H}_{k}} > n\varepsilon/2\right)$$
$$\leq \sum_{\sigma \in \{0,1\}} \mathbb{P}\left(\left\|\sum_{b=1}^{B} \left(\sum_{t=1}^{s} \widetilde{Z}_{t}^{[2b-\sigma]}\right)\right\|_{\mathcal{H}_{k}} > n\varepsilon/2\right) + 2\left(B-1\right)\beta(s)$$

where for  $b \in [2]$ , the blocks  $\widetilde{Z}^{[2b]} = \widetilde{Z}_1^{[2b]}, \ldots, \widetilde{Z}_s^{[2b]}$  are mutually independent and a similar fact holds for the blocks  $\widetilde{Z}^{[2b-1]} = \widetilde{Z}_1^{[2b-1]}, \ldots, \widetilde{Z}_s^{[2b-1]}$ . As a result, for  $b \in [B]$ , the  $\sum_{t=1}^s \widetilde{Z}_t^{[2b-\sigma]}$  are independent, centered and valued in  $\mathcal{H}_k$ , thus the problem above has been reduced to controlling the norm of a sum of B iid random vectors in the Hilbert space  $\mathcal{H}_k$ . Let us fix  $\sigma \in \{0, 1\}$ . Similar to our proof in the iid setting, our goal is to invoke Theorem 32 (Yurinsky, 1995, Theorem 3.3.4). In order to do that, we need to control from above the  $\sum_{b=1}^{B} \mathbb{E} \left\| \sum_{t=1}^s \widetilde{Z}_t^{[2b-\sigma]} \right\|_{\mathcal{H}_k}^p$ , for any  $p \geq 3$ . For p = 2, by linearity of the expectation, we have

$$\sum_{b=1}^{B} \mathbb{E} \left\| \sum_{t=1}^{s} \widetilde{Z}_{t}^{[2b-\sigma]} \right\|_{\mathcal{H}_{k}}^{2} = \sum_{b=1}^{B} \sum_{t=1}^{s} \sum_{r=1}^{s} \mathbb{E} \langle \widetilde{Z}_{t}^{[2b-\sigma]}, \widetilde{Z}_{r}^{[2b-\sigma]} \rangle_{\mathcal{H}_{k}}.$$

By stationarity, we can write more simply

$$\begin{split} \sum_{b=1}^{B} \mathbb{E} \left\| \sum_{t=1}^{s} \widetilde{Z}_{t}^{[2b-\sigma]} \right\|_{\mathcal{H}_{k}}^{2} &= \sum_{b=1}^{B} \sum_{t=1}^{s} \sum_{r=1}^{s} \mathbb{E} \langle k(X_{(2b-\sigma-1)s+t}, \cdot) - \mu_{\mathbb{P}}, k(X_{(2b-\sigma-1)s+r}, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}} \\ &= B \sum_{t=1}^{s} \sum_{r=1}^{s} \mathbb{E} \langle k(X_{t-r+1}, \cdot) - \mu_{\mathbb{P}}, k(X_{1}, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}} \\ &= B \left( \sum_{t=1}^{s} \langle k(X_{1}, \cdot) - \mu_{\mathbb{P}}, k(X_{1}, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}} \\ &+ 2 \sum_{t=2}^{s} (s-t+1) \mathbb{E} \langle k(X_{t}, \cdot) - \mu_{\mathbb{P}}, k(X_{1}, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_{k}} \right). \end{split}$$

It is then natural to consider the following covariance coefficients in the RKHS introduced by Chérief-Abdellatif and Alquier (2022),

$$\rho_t \doteq \mathbb{E} \langle k(X_t, \cdot) - \mu_{\mathbb{P}}, k(X_1, \cdot) - \mu_{\mathbb{P}} \rangle_{\mathcal{H}_k},$$

and in particular, observe that we recover  $\rho_1 = v$ . It follows that

$$\sum_{b=1}^{B} \mathbb{E} \left\| \sum_{t=1}^{s} \widetilde{Z}_{t}^{[2b-\sigma]} \right\|_{\mathcal{H}_{k}}^{2} \leq \frac{n}{2} \left( v + \frac{2}{s} \sum_{t=2}^{s} (s-t+1)\rho_{t} \right) = \frac{n}{2} \left( v + \Sigma_{s} \right),$$

where the mixing coefficient  $\Sigma_s$  was introduced in Lemma 16. In particular,  $\Sigma_s$  vanishes for iid processes. For  $p \geq 3$ , by sub-additivity of the norm,

$$\left\|\sum_{t=1}^{s} \widetilde{Z}_{t}^{[2b-\sigma]}\right\|_{\mathcal{H}_{k}} \leq \sum_{t=1}^{s} \left\|\widetilde{Z}_{t}^{[2b-\sigma]}\right\|_{\mathcal{H}_{k}} \leq 2s\sqrt{\overline{k}}.$$

Therefore, invoking Theorem 32 with  $G^2 = \frac{n}{2}(v + \Sigma_s)$  and  $H = 2s\sqrt{k}/3$ , we obtain for  $\sigma \in \{0, 1\}$  and for any  $\varepsilon > 0$  that,

$$\mathbb{P}\left(\left\|\sum_{b=1}^{B}\left(\sum_{t=1}^{s}\widetilde{Z}_{t}^{[2b-\sigma]}\right)\right\|_{\mathcal{H}_{k}} > \frac{n\varepsilon}{2}\right) \leq 2\exp\left(-\frac{((\varepsilon/2)n)^{2}}{2(n(v+\Sigma_{s})/2 + ((\varepsilon/2)n)2s\sqrt{\overline{k}}/3)}\right)$$
$$= 2\exp\left(-\frac{n\varepsilon^{2}}{4\left(v+\Sigma_{s}+2\varepsilon s\sqrt{\overline{k}}/3\right)}\right).$$

It remains to control the term involving the  $\beta$ -mixing coefficient by suitable choosing the blocking size. For

$$\tau_{n,\delta}^{\beta} \doteq \operatorname*{arg\,min}_{s \in \mathbb{N}} \left\{ s \geq t_{\mathsf{mix}}^{\beta} \left( \frac{\delta}{6(n/(2s) - 1)} \right) \right\},$$

it holds that  $2(B-1)\beta(s) \leq \delta/3$ , which concludes the proof.

# B.4 Proof of Lemma 27

We will rely on the following property for general kernels.

**Lemma 31** Let  $x = (x_1, \ldots, x_{t_0}, \ldots, x_n) \in \mathcal{X}^n$  and  $x'_{t_0} \in \mathcal{X}$ . Writing

$$x^{(t)} = (x_1, \dots, x_{t_0-1}, x'_{t_0}, x_{t_0+1}, \dots, x_n),$$

it holds that

$$\widehat{v}(x) - \widehat{v}\left(x^{(t_0)}\right) = \frac{1}{n} \left(k\left(x_{t_0}, x_{t_0}\right) - k\left(x'_{t_0}, x'_{t_0}\right)\right) \\ + \frac{2}{(n-1)n} \sum_{t \in [n], t \neq t_0} \left(k\left(x'_{t_0}, x_t\right) - k\left(x_{t_0}, x_t\right)\right).$$

From Lemma 31, if follows that

$$\begin{aligned} \widehat{v}(x) &- \inf_{x_{t_0}' \in \mathcal{X}} \widehat{v}\left(x^{(t_0)}\right) \stackrel{(\dagger)}{\leq} \frac{1}{n} \left(k\left(x_{t_0}, x_{t_0}\right) - \underline{\operatorname{diag}}\,k\right) \\ &+ \frac{2}{(n-1)n} \sum_{t \in [n], t \neq t_0} \left(\overline{k} - k\left(x_{t_0}, x_t\right)\right) \\ \stackrel{(\dagger)}{\leq} \frac{\Delta \operatorname{diag}\,k + 2\Delta k}{n}, \end{aligned}$$

and the bounded differences property follows directly from (‡). Furthermore, as a result of (†),

$$\begin{split} &\sum_{t_0=1}^n \left( \widehat{v}(x) - \inf_{x_{t_0} \in \mathcal{X}} \widehat{v}\left(x^{(t_0)}\right) \right)^2 \\ &\leq \sum_{t_0=1}^n \left( \frac{1}{n} \left( k\left(x_{t_0}, x_{t_0}\right) - \underline{\operatorname{diag}} k \right) + \frac{2}{(n-1)n} \sum_{t \in [n], t \neq t_0} \left( \overline{k} - k\left(x_{t_0}, x_t\right) \right) \right)^2 \\ &= \frac{1}{n^2} \sum_{t_0=1}^n \left( \frac{1}{(n-1)} \sum_{t \in [n], t \neq t_0} k\left(x_{t_0}, x_{t_0}\right) - \underline{\operatorname{diag}} k + 2\overline{k} - 2k\left(x_{t_0}, x_t\right) \right)^2 \\ &\leq \frac{1}{n^2} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n], t \neq t_0} \left( k\left(x_{t_0}, x_{t_0}\right) - \underline{\operatorname{diag}} k + 2\overline{k} - 2k\left(x_{t_0}, x_t\right) \right)^2 \\ &\leq \frac{\Delta \operatorname{diag} k + 2\Delta k}{n^2} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n], t \neq t_0} \left( k\left(x_{t_0}, x_{t_0}\right) - \underline{\operatorname{diag}} k + 2\overline{k} - 2k\left(x_{t_0}, x_t\right) \right) \\ &\leq \frac{\Delta \operatorname{diag} k + 2\Delta k}{n^2} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n], t \neq t_0} \left( 2k\left(x_{t_0}, x_{t_0}\right) - 2\underline{\operatorname{diag}} k + 2\overline{k} - 2k\left(x_{t_0}, x_t\right) \right) \\ &= 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n^2} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n], t \neq t_0} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n^2} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n], t \neq t_0} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n} \left( \overline{k} - \underline{\operatorname{diag}} k \right) \\ &= 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n^2} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n]} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n^2} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n]} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n^2} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n]} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n]} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n]} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n]} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n]} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag} k + 2\Delta k}{n} \sum_{t_0=1}^n \frac{1}{(n-1)} \sum_{t \in [n]} \left( k\left(x_{t_0}, x_{t_0}\right) - k\left(x_{t_0}, x_t\right) \right) \\ &+ 2 \frac{\Delta \operatorname{diag$$

where for  $(\star)$  we relied on  $\overline{k} = \overline{\operatorname{diag} k}$  for a positive definite kernel.

#### 

# B.5 Proof of Lemma 28

Lemma 27 established that the function

$$\mathcal{X}^n \to \mathbb{R}, x \mapsto \frac{n}{\Delta \operatorname{diag} k + 2\Delta k} \widehat{v}(x),$$

is weakly self-bounding and has the bounded differences property. From an application of Boucheron et al. (2009, Theorem 1),

$$\begin{split} & \mathbb{P}\left(\widehat{v} - v > t\right) \leq \exp\left(-\frac{nt^2}{4(\Delta \operatorname{diag} k + 2\Delta k)(v + \Delta \operatorname{diag} k + t/2)}\right), \\ & \mathbb{P}\left(v - \widehat{v} > t\right) \leq \exp\left(-\frac{nt^2}{4(\Delta \operatorname{diag} k + 2\Delta k)(v + \Delta \operatorname{diag} k)}\right). \end{split}$$

The claim follows from rewriting

$$\widehat{v} - v = [\widehat{v} + \Delta \operatorname{diag} k] - [v + \Delta \operatorname{diag} k],$$

and solving quadratic inequalities as in Lemma 11.

## B.6 Proof of Theorem 29

From Theorem 1, and since  $\Delta \operatorname{diag} k > 0$ , with probability at least  $1 - \delta/2$ , it holds that

$$\|\widehat{\mu}_{\mathbb{P}}(X_1,\ldots,X_n) - \mu_{\mathbb{P}}\|_{\mathcal{H}_k} \le \sqrt{2(v + \Delta \operatorname{diag} k) \frac{\log(4/\delta)}{n}} + \frac{4}{3}\sqrt{k} \frac{\log(4/\delta)}{n}.$$
 (28)

By Lemma 28 for b = -1, at confidence  $1 - \delta/2$  it holds that

$$\sqrt{v + \Delta \operatorname{diag} k} \le \sqrt{\widehat{v} + \Delta \operatorname{diag} k} + 2\sqrt{\frac{(\Delta \operatorname{diag} k + 2\Delta k)\log(4/\delta)}{n}},\tag{29}$$

thus by combining (28) and (29) and a union bound yields that with probability at least  $1-\delta$ ,

$$\begin{aligned} \|\widehat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|_{\mathcal{H}_{k}} \\ &\leq \left(\sqrt{\widehat{v} + \Delta \operatorname{diag} k} + 2\sqrt{\frac{(\Delta \operatorname{diag} k + 2\Delta k)\log(4/\delta)}{n}}\right)\sqrt{\frac{2\log(4/\delta)}{n}} + \frac{4}{3}\sqrt{\overline{k}}\frac{\log(4/\delta)}{n} \\ &= \sqrt{2(\widehat{v} + \Delta \operatorname{diag} k)\frac{\log(4/\delta)}{n}} + \left(2\sqrt{2}\sqrt{\Delta \operatorname{diag} k + 2\Delta k} + \frac{4}{3}\sqrt{\overline{k}}\right)\frac{\log(4/\delta)}{n} \\ &\leq \sqrt{2(\widehat{v} + \Delta \operatorname{diag} k)\frac{\log(4/\delta)}{n}} + \left(\frac{16}{3}\sqrt{\Delta k} + 2\sqrt{2}\sqrt{\Delta \operatorname{diag} k}\right)\frac{\log(4/\delta)}{n}, \end{aligned}$$
(30)

where the last inequality is by sub-additivity of the square root.

# B.7 Proof of Lemma 31

$$\begin{aligned} \widehat{v}(x) &- \widehat{v}\left(x^{(t)}\right) \\ &= \frac{1}{n-1} \sum_{r=1}^{n} \left( k\left(x_{r}, x_{r}\right) - \frac{1}{n} \sum_{s=1}^{n} k\left(x_{r}, x_{s}\right) \right) \\ &- \frac{1}{n-1} \sum_{r=1}^{n} \left( k\left(x_{r}^{(t)}, x_{r}^{(t)}\right) - \frac{1}{n} \sum_{s=1}^{n} k\left(x_{r}^{(t)}, x_{s}^{(t)}\right) \right) \\ &= \frac{1}{n-1} \sum_{r=1}^{n} \left( k\left(x_{r}, x_{r}\right) - k\left(x_{r}^{(t)}, x_{r}^{(t)}\right) \right) \\ &- \frac{1}{(n-1)n} \sum_{(r,s) \in [n]^{2}} \left( k\left(x_{r}, x_{s}\right) - k\left(x_{r}^{(t)}, x_{s}^{(t)}\right) \right) \end{aligned}$$

thus,

$$\begin{split} \hat{v}(x) &- \hat{v}\left(x^{(t)}\right) \\ &= \frac{1}{n-1} \left(k\left(x_{t}, x_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) - \frac{1}{(n-1)n} \sum_{(r,s) \in [n]^{2}, r=s} \left(k\left(x_{r}, x_{s}\right) - k\left(x^{(t)}_{r}, x^{(t)}_{s}\right)\right) \right) \\ &- \frac{1}{(n-1)n} \sum_{(r,s) \in [n]^{2}, r\neq s} \left(k\left(x_{r}, x_{s}\right) - k\left(x^{(t)}_{r}, x^{(t)}_{s}\right)\right) \right) \\ &= \frac{1}{n-1} \left(k\left(x_{t}, x_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) - \frac{1}{(n-1)n} \sum_{r=1}^{n} \left(k\left(x_{r}, x_{r}\right) - k\left(x^{(t)}_{r}, x^{(t)}_{r}\right)\right) \right) \\ &- \frac{1}{(n-1)n} \sum_{(r,s) \in [n]^{2}, r\neq s, r=t} \left(k\left(x_{r}, x_{s}\right) - k\left(x^{(t)}_{r}, x^{(t)}_{s}\right)\right) \\ &- \frac{1}{(n-1)n} \sum_{(r,s) \in [n]^{2}, r\neq s, s=t} \left(k\left(x_{r}, x_{s}\right) - k\left(x^{(t)}_{r}, x^{(t)}_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x_{t}, x_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) \\ &- \frac{1}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x_{t}, x_{s}\right) - k\left(x'_{t}, x_{s}\right)\right) \\ &- \frac{1}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x_{t}, x_{s}\right) - k\left(x'_{t}, x_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x_{t}, x_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) + \frac{2}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x'_{t}, x_{s}\right) - k\left(x_{t}, x_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x_{t}, x_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) + \frac{2}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x'_{t}, x_{s}\right) - k\left(x_{t}, x_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x_{t}, x_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) + \frac{2}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x'_{t}, x_{s}\right) - k\left(x_{t}, x_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x'_{t}, x_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) + \frac{2}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x'_{t}, x_{s}\right) - k\left(x_{t}, x_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x'_{t}, x_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) + \frac{2}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x'_{t}, x_{s}\right) - k\left(x'_{t}, x_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x'_{t}, x'_{t}\right) - k\left(x'_{t}, x'_{t}\right)\right) + \frac{2}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x'_{t}, x'_{s}\right) - k\left(x'_{t}, x'_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x'_{t}, x'_{t}\right) - k\left(x''_{t}, x''_{t}\right)\right) + \frac{2}{(n-1)n} \sum_{s \in [n], s\neq t} \left(k\left(x'_{t}, x'_{s}\right) - k\left(x'_{t}, x'_{s}\right)\right) \\ &= \frac{1}{n} \left(k\left(x'_{t}, x'_{t}\right) - k\left(x''_{t}, x''_{t}\right)\right) + \frac{1}{n} \left(k\left(x'_{t}, x''_{t}\right) + \frac{1}{n} \left(k\left(x'_{t}, x''_{t}\right)\right) + \frac{1}{n} \left(k\left(x'_{t}, x''_{t}\right) + \frac{1}{n} \left(k\left(x'_{t}, x''_{t}\right)\right) + \frac{1}{n} \left(k\left(x'_{t}, x''_{t}\right)\right) + \frac{1}{n} \left(k\left(x$$

# Appendix C. Tools from the Literature

We report here tools from the literature.

**Theorem 32 (Yurinsky, 1995, Theorem 3.3.4)** Let  $Z_1, \ldots, Z_n$  be independent random variables in a separable Hilbert space  $\mathcal{H}$  that are centered, that is,

$$\forall t \in [n], \mathbb{E}\left[Z_t\right] = 0,$$

If there exists real numbers  $G, H \ge 0$  such that for any  $p \ge 2$ ,

$$\sum_{t=1}^{n} \mathbb{E} \| Z_t \|_{\mathcal{H}}^p \le \frac{1}{2} p! G^2 H^{p-2},$$

it holds that for any  $\varepsilon > 0$ ,

$$\mathbb{P}\left(\max_{1\leq s\leq n} \left\|\sum_{t=1}^{s} Z_t\right\|_{\mathcal{H}} > \varepsilon\right) \leq 2\exp\left(-\frac{\varepsilon^2/2}{G^2 + \varepsilon H}\right)$$

**Theorem 33 (Samson, 2000, Theorem 3)** Let  $X_1, \ldots, X_n$  be a stationary  $\phi$ -mixing sequence. For every  $\varepsilon > 0$ ,

$$\mathbb{P}\left(Z \geq \mathbb{E}Z + \varepsilon\right) \leq \exp\left(-\frac{1}{8\left\|\Gamma\right\|_{2}}\min\left\{\frac{\varepsilon}{C}, \frac{\varepsilon^{2}}{4\mathbb{E}\left[V\right]}\right\}\right)$$

where

$$Z \doteq \sup_{g \in \mathcal{F}} \left| \sum_{t=1}^{n} g(X_t) \right|,$$

with  $\mathcal{F}$  an arbitrary class of real bounded functions with  $|g| \leq C$ , and where the random variable V is defined as

$$V \doteq \sum_{t=1}^{n} \sup_{g \in \mathcal{F}} g(X_t)^2,$$

and  $\Gamma$  is the coupling matrix (refer to Equation 17).

### References

- P. Alquier, B.-E. Chérief-Abdellatif, A. Derumigny, and J.-D. Fermanian. Estimation of copulas via maximum mean discrepancy. *Journal of the American Statistical Association*, 118(543):1997–2012, 2023.
- M. A. Arcones. A Bernstein-type inequality for U-statistics and U-processes. Statistics & probability letters, 22(3):239–247, 1995.
- J.-Y. Audibert, R. Munos, and C. Szepesvári. Tuning bandit algorithms in stochastic environments. In *International conference on algorithmic learning theory*, pages 150–165. Springer, 2007.
- P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.

- V. Bentkus, N. Kalosha, and M. Van Zuijlen. On domination of tail probabilities of (super) martingales: explicit bounds. *Lithuanian Mathematical Journal*, 46(1):1–43, 2006.
- D. Berend and A. Kontorovich. A sharp estimate of the binomial mean absolute deviation with applications. *Statistics & Probability Letters*, 83(4):1254–1259, 2013.
- A. Berlinet and C. Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.
- S. Boucheron, G. Lugosi, and P. Massart. On concentration of self-bounding functions. *Electronic Journal of Probability*, 14:1884–1899, 2009.
- R. C. Bradley. Basic properties of strong mixing conditions. a survey and some open questions. *Probability Surveys*, 2:107–144, 2005.
- F.-X. Briol, A. Barp, A. B. Duncan, and M. Girolami. Statistical inference for generative models with maximum mean discrepancy. arXiv preprint arXiv:1906.05944, 2019.
- B.-E. Chérief-Abdellatif and P. Alquier. Finite sample properties of parametric MMD estimation: robustness to misspecification and dependence. *Bernoulli*, 28(1):181–213, 2022.
- K. Chwialkowski, H. Strathmann, and A. Gretton. A kernel test of goodness of fit. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2606–2615, New York, New York, USA, 20–22 Jun 2016. PMLR.
- D. Cohen, A. Kontorovich, and G. Wolfer. Learning discrete distributions with infinite support. In Advances in Neural Information Processing Systems, volume 33, pages 3942– 3951, 2020.
- M. Cuturi. Positive definite kernels in machine learning. arXiv preprint arXiv:0911.5367, 2009.
- J. Diestel and J. Uhl. Vector measures, vol. 15 of mathematical surveys. American Mathematical Society, Providence, RI, USA, 1977.
- P. Doukhan. Mixing: properties and examples, volume 85. Springer Science & Business Media, 2012.
- G. K. Dziugaite, D. M. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. In *Proceedings of the Thirty-First Conference* on Uncertainty in Artificial Intelligence, pages 258–267, 2015.
- M. G. Genton. Classes of kernels for machine learning: a statistics perspective. Journal of machine learning research, 2(Dec):299–312, 2001.
- A. Gretton. Introduction to RKHS, and some simple kernel algorithms. Adv. Top. Mach. Learn. Lecture Conducted from University College London, 16(5-3):2, 2013.

- A. Gretton, K. Fukumizu, Z. Harchaoui, and B. K. Sriperumbudur. A fast, consistent kernel two-sample test. *Advances in neural information processing systems*, 22, 2009.
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel twosample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- W. Hoeffding. Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association, 58(301):13–30, 1963.
- P. J. Huber. Robust statistics. In International encyclopedia of statistical science, pages 1248–1251. Springer, 2011.
- I. A. Ibragimov. Some limit theorems for stationary processes. Theory of Probability & Its Applications, 7(4):349–382, 1962.
- W. Jitkrittum, W. Xu, Z. Szabó, K. Fukumizu, and A. Gretton. A linear-time kernel goodness-of-fit test. Advances in Neural Information Processing Systems, 30, 2017.
- M. Kearns, Y. Mansour, D. Ron, R. Rubinfeld, R. E. Schapire, and L. Sellie. On the learnability of discrete distributions. In *Proceedings of the twenty-sixth annual ACM* symposium on Theory of computing, pages 273–282, 1994.
- A. K. Kuchibhotla and Q. Zheng. Near-optimal confidence sequences for bounded random variables. In *International Conference on Machine Learning*, pages 5827–5837. PMLR, 2021.
- D. A. Levin, Y. Peres, and E. L. Wilmer. Markov chains and mixing times, second edition. American Mathematical Soc., 2009.
- Y. Li, K. Swersky, and R. Zemel. Generative moment matching networks. In International conference on machine learning, pages 1718–1727. PMLR, 2015.
- D. Lopez-Paz, K. Muandet, B. Schölkopf, and I. Tolstikhin. Towards a learning theory of cause-effect inference. In *International Conference on Machine Learning*, pages 1452– 1461. PMLR, 2015.
- A. Maurer. Concentration inequalities for functions of independent variables. Random Structures & Algorithms, 29(2):121–138, 2006.
- A. Maurer and M. Pontil. Empirical Bernstein bounds and sample variance penalization. arXiv preprint arXiv:0907.3740, 2009.
- C. McDiarmid et al. On the method of bounded differences. *Surveys in combinatorics*, 141 (1):148–188, 1989.
- M. Mohri and A. Rostamizadeh. Stability bounds for stationary  $\varphi$ -mixing and  $\beta$ -mixing processes. Journal of Machine Learning Research, 11(2), 2010.
- M. Mohri, A. Rostamizadeh, and A. Talwalkar. Foundations of machine learning. MIT press, 2018.

- K. Muandet, K. Fukumizu, B. Sriperumbudur, A. Gretton, and B. Schölkopf. Kernel mean estimation and Stein effect. In *International Conference on Machine Learning*, pages 10–18. PMLR, 2014.
- K. Muandet, B. Sriperumbudur, K. Fukumizu, A. Gretton, and B. Schölkopf. Kernel mean shrinkage estimators. *Journal of Machine Learning Research*, 17, 2016.
- F. Münch and J. Salez. Mixing time and expansion of non-negatively curved Markov chains. Journal de l'École polytechnique — Mathématiques, 10:575–590, 2023. doi: 10.5802/jep. 226.
- T. Peel, S. Anthoine, and L. Ralaivola. Empirical Bernstein inequalities for U-statistics. Advances in Neural Information Processing Systems, 23, 2010.
- I. F. Pinelis and A. I. Sakhanenko. Remarks on inequalities for the probabilities of large deviations. *Teoriya Veroyatnostei i ee Primeneniya*, 30(1):127–131, 1985.
- G. O. Roberts and J. S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1(none):20 71, 2004. doi: 10.1214/154957804100000024.
- P.-M. Samson. Concentration of measure inequalities for Markov chains and Φ-mixing processes. The Annals of Probability, 28(1):416 461, 2000. doi: 10.1214/aop/1019160125.
- M. L. Stein. Interpolation of spatial data: some theory for kriging. Springer Science & Business Media, 1999.
- I. Tolstikhin, B. K. Sriperumbudur, and K. Muandet. Minimax estimation of kernel mean embeddings. *The Journal of Machine Learning Research*, 18(1):3002–3048, 2017.
- A. B. Tsybakov. Introduction to nonparametric estimation. Springer Science & Business Media, 2008.
- I. Waudby-Smith and A. Ramdas. Estimating means of bounded random variables by betting. Journal of the Royal Statistical Society Series B: Statistical Methodology, 86(1): 1-27, 2024.
- G. Wolfer and P. Alquier. Optimistic estimation of convergence in Markov chains with the average-mixing time. arXiv:2402.10506, 2024.
- B. Yu. Rates of convergence for empirical processes of stationary mixing sequences. The Annals of Probability, pages 94–116, 1994.
- V. Yurinsky. Sums and Gaussian Vectors. Lecture Notes in Mathematics. Springer Berlin, Heidelberg, 1995.